



THE HONG KONG
POLYTECHNIC UNIVERSITY

香港理工大學

Pao Yue-kong Library

包玉剛圖書館

Copyright Undertaking

This thesis is protected by copyright, with all rights reserved.

By reading and using the thesis, the reader understands and agrees to the following terms:

1. The reader will abide by the rules and legal ordinances governing copyright regarding the use of the thesis.
2. The reader will use the thesis for the purpose of research or private study only and not for distribution or further reproduction or any other purpose.
3. The reader agrees to indemnify and hold the University harmless from and against any loss, damage, cost, liability or expenses arising from copyright infringement or unauthorized usage.

IMPORTANT

If you have reasons to believe that any materials in this thesis are deemed not suitable to be distributed in this form, or a copyright owner having difficulty with the material being included in our database, please contact lbsys@polyu.edu.hk providing details. The Library will look into your claim and consider taking remedial action upon receipt of the written requests.

THREE ESSAYS IN OPERATIONS MANAGEMENT:
INVENTORY MANAGEMENT WITH DEMAND
LEARNING, QUALITY SIGNALING AND
BAYESIAN PERSUASION IN SERVICE

LUO ZHENWEI

PhD

The Hong Kong Polytechnic University

2021

The Hong Kong Polytechnic University

Department of Logistics and Maritime Studies

**Three Essays in Operations Management: Inventory
Management with Demand Learning, Quality Signaling
and Bayesian Persuasion in Service**

Luo Zhenwei

A thesis submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

June 2021

CERTIFICATE OF ORIGINALITY

I hereby declare that this thesis is my own work and that, to the best of my knowledge and belief, it reproduces no material previously published or written, nor material that has been accepted for the award of any other degree or diploma, except where due acknowledgement has been made in the text.

_____ (Signed)

_____ Luo Zhenwei (Name of student)

Abstract

The information uncertainty is quite common for decision makers in operations management. This thesis includes three different settings with information uncertainty, where the Bayesian updating framework is adopted as a parametric learning approach. In the first topic, we consider that an airline company offers early-bird-discount seats to customers and aims to maximize the expected profit via optimally allocating the seats for discount sales. The beliefs on demand parameters and buy-up substitution probability are updated using demand observations considering an “exploration-exploitation” tradeoff. Classic literature finds that unobservability of lost sales is a driving force for the “stock more” result, namely the Bayesian-optimal inventory level shall be kept higher than the myopic one to allow a better observability of demand. In contrast, we find that one can infer some information about lost sales from the substitution behavior of unsatisfied customers and hence may “stock less”. We also find that to better observe the primary demand for the regular-price seat, one shall “stock more” discounted seats to reduce the chance of tangling the substitution demand with the primary demand. And, to better observe the substitution probability, one shall “stock less” discounted seats to observe substitutions.

In the second topic, we consider a single-server queueing system whose service quality is either high or low. The server knows the actual quality level, and can signal it to customers via revealing or concealing his queue length. A signaling game is formed, and we adopt the sequential equilibrium concept to solve our game and apply the perfect sequential equilibrium as an equilibrium-refinement criterion. Under a general scenario in which the market is composed of both quality informed and uninformed customers, the unique equilibrium outcome is a pooling strategy when the market size is either below a lower threshold or above an upper threshold. And the separating equilibria may exist only when the market size falls between these two thresholds, under which uninformed customers can fully infer the server’s quality type based on his queue disclosure behavior.

In the third topic, we study a server's best queue-disclosure strategy in a single-server service system with uncertain quality level. We consider this problem as a Bayesian persuasion game. The server can commit to a strategy that states whether or not the queue length will be revealed to customers upon their arrival, given a realized quality level. We reformulate the server's decision problem as looking for the best Bayes-plausible distribution of customers' posteriors on service quality, which can be solved via a geometric approach. We also show that when the market size is sufficiently small (resp. large), the server always conceals (resp. reveals) the queue regardless of the realized service quality. In a medium-sized market, however, we numerically find that the server's optimal commitment strategy is often hybrid or mixed, that is, randomized over queue disclosure and concealment. We also extend our analysis to another scenario where the server is a social planner.

Publications Arising from the Thesis

Luo, Z., Guo, P. and Y. Wang. 2021. Manage inventories with learning on demands and buy-up substitution probability. Under 2nd-round major revision for *Manufacturing & Service Operations Management*.

Guo, P., Haviv, M., **Luo, Z.*** and Y. Wang. 2021. Signaling service quality via queue disclosure. Under 1st-round major revision for *Manufacturing & Service Operations Management*.

Guo, P., Haviv, M., **Luo, Z.*** and Y. Wang. 2021. Optimal queue length information disclosure when service quality is uncertain. Under 2nd-round review for *Production and Operations Management*.

Acknowledgments

The motto of PolyU says “To learn and to apply, for the benefit of mankind.” It perfectly summarizes my PhD study. I used to take everything existing for granted. But my PhD experience greatly changes my thinking way, and now I find that even the simplest ideas existing in our mind did not appear without cause, but were generated after huge efforts in thinking and creating. The simpler, the harder. It is really my great honour to receive PhD education and create some new knowledges for others!

I am very fortunate to get the guidances from my supervisors Prof. Guo Pengfei and Prof. Wang Yulan. They involved me in advanced research projects, provided me with new ideas when I were stagnant, and inspired me by their deep and creative thinking. Not only in my study and research, my supervisors also helped me a lot in my life and research career. Under their all-round care, I could fully devote myself to study and research, and gradually grow up to be an adequate researcher. The PhD journey is really long and hard. As my experience, only under loving care from supervisors can a PhD candidate develop well in body and mind. In this sense, I really appreciate my supervisors.

I gratefully acknowledge my university PolyU and my department LMS for offering me great material supports such as adequate financial funding, convenient working facilities and rich research resources. Additionally, I sincerely thank those people who helped me a lot in my PhD study: Prof. Zhang Hanqin and Prof. Wang Jinting, for your recommendations; Prof. Moshe Haviv, for your guidance and cooperation in my research papers; Prof. Song Miao, for your teaching and discussions with me; and Leung Lorraine HT, Lam Irene and Wong Anne LY, for your supports on my departmental affairs. Last but not least, my great thanks go to my parents, Mr. Luo Yongtao and Mrs. Jiang Yunxiang, who are always selflessly in favour of me and warmly tolerant of me.

Table of Contents

Title Page	i
Certificate of Originality	iii
Abstract	v
Publications Arising from the Thesis	vii
Acknowledgements	ix
Table of Contents	xi
1 Introduction	1
1.1 Information Uncertainty and Bayesian Updating	1
1.2 Layout of the Thesis	2
2 Manage Inventories with Learning on Demands and Buy-up Substitution	
Probability	5
2.1 Introduction	5
2.2 Literature Review	11
2.3 One-period Model	14
2.3.1 Model Description	14
2.3.2 Repeated Decision Making	16
2.4 Multi-period Bayesian Inventory Management	18
2.4.1 Four Information Scenarios and Bayesian Learning	18
2.4.2 Updating Only Demand Parameter θ	23
2.4.3 Updating Both Demand Parameter θ and Substitution Probability α	29

2.4.4	Comparison of Expected Profits under Four Information Scenarios . . .	30
2.5	Numeric Study	30
2.6	Conclusions and Suggestions for Future Research	35
3	Signaling Service Quality via Queue Disclosure	37
3.1	Introduction	37
3.2	Literature Review	41
3.3	Model Setup and Equilibrium Concepts	42
3.3.1	Timing of Signaling Game	43
3.3.2	Definitions of Sequential Equilibria and Perfect Sequential Equilibria .	44
3.4	Equilibrium Analysis	46
3.4.1	Customers' Equilibrium Queueing Strategies and Effective Arrival Rates	46
3.4.2	Sequential Equilibria Analysis	47
3.5	Signaling Game with Heterogeneous Customers	51
3.5.1	Customers' Equilibrium Queueing Strategies and Effective Arrival Rates	51
3.5.2	Sequential Equilibria Analysis	57
3.5.3	Effects of Using Queue Disclosure as Signal	62
3.6	Discussions	67
3.6.1	The Impact of Customer Type Composition	67
3.6.2	The Impact of Service Price	70
3.7	Conclusions and Suggestions for Future Research	72
4	Optimal Queue Length Information Disclosure When Service Quality Is Uncertain	75
4.1	Introduction	76
4.2	Literature Review	80
4.3	Model Description	83
4.4	Optimal Queue Disclosure Strategy	86
4.4.1	Geometric Approach	86
4.4.2	The Impact of Market Size	91
4.5	Social Planner	94
4.6	Conclusions and Suggestions for Future Research	98
	Appendix A Supplements and Proofs for Chapter 2	101

A.1	A General Two-Decision-Variable Problem	101
A.1.1	One-period Model	101
A.1.2	Multi-period Bayesian Inventory Management	102
A.1.3	Upper Bounds on Optimal Inventory Levels	105
A.2	Proofs for Chapter 2 and Appendix A.1	106
Appendix B Supplements and Proofs for Chapter 3		119
B.1	Sequential Equilibria Analysis of Hybrid and Mixed Strategies with Homogeneous Customers	119
B.2	Customers' Equilibrium Queueing Strategies in Unobservable Queues with $\delta^C = 0$ or 1	120
B.3	Review of Hole-avoiding Decision Process in Debo et al. (2012)	123
B.4	Sequential Equilibria Analysis of Hybrid and Mixed Strategies with Heterogeneous Customers	125
B.5	Proofs for Chapter 3 and Appendix B.2	126
Appendix C Proofs for Chapter 4		135
References		139

Chapter 1

Introduction

1.1 Information Uncertainty and Bayesian Updating

In the practice of operations management, the information uncertainty is a common phenomenon, where some players in the process lack necessary information in their decision makings. For example, an inventory manager, who is deciding how many to order, may not know the exact demand information; or a customer, who is deciding whether or not to join a queue to gain service, may not know the service quality of the server. They have to learn the needed information from related observations and signals, and such learnings can affect the whole operations process.

In this thesis, the information learning issues under three different settings in operations management are investigated. Under each setting, the decision maker knows that the concerned information follows some distribution, but is not sure of the exact parameter value of this distribution. In this sense, the Bayesian updating framework becomes a desirable parametric approach since we can incorporate previous experiences and intuitive knowledge as the prior belief on the unknown parameter and use new observations or information to update it. Another advantage of Bayesian updating is that it well depicts the nervous mechanism behind human's cognition and learning (Knill and Pouget, 2004; Doya et al., 2007; and Glimcher and Fehr, 2013).

For readers' interests, we provide a brief review on the Bayesian updating process as follows. The uncertain state of the world (e.g., the parameter of demand distribution, or the quality level of the server) is denoted by a parameter θ with $\theta \in \Theta$, where Θ is the *parameter space*. Before obtaining new observations, we hold a *prior belief* $p(\theta)$, which describes the

probability that the true state is θ . The observations make up a dataset x with $x \in X$, where X the *sample space*. Given $\theta \in \Theta$ and $x \in X$, the *likelihood function* $p(x|\theta)$ describes the probability that the observation would be x if θ is the true state. Once we observe x , we get the *posterior belief* $p(\theta|x)$ according to the *Bayes' rule* as

$$p(\theta|x) = \frac{p(x|\theta)p(\theta)}{\int_{\Theta} p(x|\theta')p(\theta')d\theta'}.$$

1.2 Layout of the Thesis

Now, let us see the three settings with information uncertainty and Bayesian updating in this thesis. We first focus on the information learning at the individual level, where a decision maker's learning and optimizing behaviors affect only his own profit. In Chapter 2, we consider an airline seat allocation problem. The airline company offers early-bird discount to customers. Those who purchase the tickets early before a deadline can get the discount. And if some demands for the discount are not satisfied, they may choose to buy the regular-price ticket as the substitute, which is referred to as the *buy-up substitution*. The manager needs to decide how many seats for discount among a fixed capacity of seats on a single flight to maximize the expected revenue. To make the allocation decision, the demand information and buy-up substitution probability are needed. The manager can adopt the *myopic optimization* policy, where he maximizes only the current-period profit using his beliefs and updates the beliefs using new sales information. Such a myopic policy neglects the impact of the current-period inventory decision on the demand realization and all the following periods, and thus fails to be optimal. An optimal policy should consider both the profit in the current period and the above impact, i.e., the “*exploration-exploitation*” tradeoff. The inventory management literature with such a tradeoff is named as *Bayesian inventory management*. In Chapter 2, we analyze the Bayesian inventory management under the airline setting, and investigate how the “exploration-exploitation” tradeoff can influence the optimal inventory level compared with the myopic optimal one. We consider four information scenarios based on whether or not lost sales can be observed and whether or not the substitution demands can be separated from the primary demands for the regular-price seats. The demand observations under different scenarios are different, which result in different analytical results. The classical structural result under Bayesian inventory management is that unobservable lost sales lead to the “stock more” result, i.e., the Bayesian optimal inventory level shall be set higher than the myopic

optimal one to yield more accurate demand observation (Lariviere and Porteus, 1999; Ding et al., 2002; and Chen and Plambeck, 2008). However, this result may not hold under our setting. We find the amount of the substitution demands also contains some information on the primary demand for the discounted seats. So, when the manager only gains very limited demand information directly from manipulating the inventory levels, he has the incentive to stock less discounted seats to induce more demand information from the substitutions. Other main findings include that the demand tangling causes a “stock more” driving force in learning the primary demand for the regular-price seats to yield better observation on it, and estimating the buy-up substitution probability requires “stock less” to let more substitution trials happen. Our analysis and results not only extend the Bayesian inventory management literature to a setting with two kinds of demands and a limited capacity of substitutes, but also provide managerial insights and heuristic bounds for the practical airline seat allocation.

We then turn the information learning model from the individual level to the game theoretical level, where the learning and optimizing behaviors of one player affect the profits of all players in the concerned operations process. Specifically, there are at least two players. Some players have private informations (named as *senders*), who try to signal such informations to other players (named as *receivers*). The payoff of each player is determined by his private information and the actions adopted by all players. In Chapters 3 and 4, we study such information uncertainty issues under two different but related queueing settings.

In Chapter 3, we consider that a server’s quality level, which can be high or low, is uncertain to the customers. Debo et al. (2012) show that the queue length can be a quality signal. Intuitively, a longer queue length indicates a higher quality level. We take a step further, and investigate whether the observability of the queue length itself can be a quality signal. To achieve this, we formulate the problem as a *signaling game* between a server and the customers. At first, nature decides whether the quality type is high or low according to the prior belief, which is a common knowledge to all players. After the server knows his quality type, he chooses to either reveal or conceal the queue length to maximize the customers’ effective arrival rate. The customers observe the server’s queue-disclosure action, update their beliefs based on the server’s action and the queue length in case of an observable queue, and finally decide whether to join or not. We use the *sequential equilibrium* concept (Kreps and Wilson, 1982) to analyze the equilibrium outcomes of the signaling game, and apply the *perfect sequential equilibrium* (Grossman and Perry, 1986) as the refinement criterion. The major takeaway is that the separating equilibria, where different quality types of the

server adopt totally different queue-disclosure actions, may appear in a medium-sized market when some customers are informed while others are not. Under such separating equilibria, the observability of the queue length becomes a quality signal, from which those uninformed customers can fully infer the quality type. Our results thus extend the finding in Debo et al. (2012). Compared with the research tradition where the queue disclosure action is not regarded as a quality signal, such a signaling effect increases (resp. decreases) the effective arrival rate to the high-quality (resp. low-quality) server, and increases the total utility of all customers from the low-quality server.

In the above signaling game, when all customers are uninformed, we prove that the queue-disclosure action signals no quality information. Notice that the server can only decide his action after the quality type is realized. What will happen if, before the quality type is realized, the server owns a commitment power whereby he can pre-determine his queue-disclosure action in case of different quality type? This is the research question of Chapter 4. To analyze it, we adopt the *Bayesian persuasion* framework in Kamenica and Gentzkow (2011). In their work, the persuasion signal is only restricted to some costless ones, such as words or announcements, which do not affect the payoffs directly. Indeed, the signals can be broadly in any forms as long as it can be used to persuade the receivers so as to improve the payoff of the sender. In our setting, the persuasion signal becomes the server's queue-disclosure action (i.e., revealing or concealing the queue length), which directly affects the payoffs of both the server and customers. We transform the problem into finding out the optimal Bayes-plausible distribution of customers' posteriors on service quality, and this can then be solved via a geometric approach. Specifically, the maximal expected effective arrival rate, as a function of the prior, can be graphed as the upper envelope of all convex combinations of points on the effective arrival rate function of the revealed queue and those of the concealed queue. The key finding is that in a medium-sized market, the server can design and commit to a randomized queue-disclosure policy to persuade more customers to join the queue compared with the traditional either-revealing-or-concealing paradigm. Our geometric approach has wide suitability, and we further apply it to another scenario where the server is a social planner.

Chapter 2

Manage Inventories with Learning on Demands and Buy-up Substitution Probability

In this chapter, we consider information uncertainty and learning with only one decision maker. An airline manager aims to optimize the airline seat allocation to maximize the expected profit from one flight with a fixed capacity of seats, but he does not know the exact demand information. The manager updates his belief using demand observations dynamically. Further, the “exploration-exploitation” tradeoff is involved in the model, and the manager needs to not only maximize the profit in the current period but also consider the impact of the current decision on all the following periods. Due to a severe curse of dimensionality inherent in the Bayesian dynamic programming problem, the exact optimal solution cannot be obtained theoretically, and our focus is to get some structural properties of the Bayesian optimal inventory level, which can provide some managerial insights and heuristic bounds for practical use.

2.1 Introduction

Airline companies often offer early-bird booking discounts for passengers. Those who satisfy the advance purchase requirement can get the best flight rate and those who purchase near the departure date may have to pay a much higher price. This practice is based on market segmentation – customers have different price sensitivity and time sensitivity – and offering the

early-bird booking discounts can help to stimulate more demands to fill some otherwise vacant seats. According to Shaw (1982), customers like the business travelers are time-sensitive but price-insensitive and they generally purchase late due to the tight time schedules and are willing to pay a high fare. For convenience, we call these customers *regular customers*. Others, like the leisure or vacation travelers, are not that time-sensitive, and would like to purchase early for a lower fare. We call them *early-bird customers*.

For airline companies, in order to maximize their revenues, an important decision is to determine how to allocate seats under multiple fares.^{2.1} For example, U.S. Department of Transportation (2019) points it out that “Some airlines set aside only a few seats on each flight at the lower rates.” Such decision is also called booking limit decision or seat protection level decision for different fare seats. A tradeoff exists with such an inventory decision: on one hand, reserving few seats for the low-fare tickets may result in the revenue loss by losing some price-sensitive passengers; on the other hand, reserving too many seats for the low-fare tickets can lead to fewer passengers purchasing the high-fare seats.

According to U.S. Department of Transportation (2019), the discounted seats can often be sold out very quickly. In practice, when some customers come for the discounted seats but find that they stock out, they may choose to buy high-fare tickets for substitution. This phenomenon is named as the *buy-up substitution*. According to Belobaba (1987), it is a very important concept in the airline seat allocation problem. And ignoring such buy-up substitution could result in a severe *spiral-down effect* as demonstrated by Cooper et al. (2006): setting a low protection level of the high-fare seats (or reserving too many low-fare tickets) results in a low estimation on the demand for high-fare seats, which, in turn, causes an even lower protection level of the high-fare seats in the following periods. Cooper et al. (2006) state that considering the buy-up substitution can effectively avoid such a spiral-down effect, and Cooper and Li (2012) further demonstrate the benefit of incorporating the buy-up substitution into the airline seat management problem.

We can formulate the seat allocation problem as a newsvendor-type model, and the optimal seat protection level for different fare seats can then be obtained by solving such a model; see Littlewood (1972) and Belobaba (1987). To make the optimal decision on the airline seat inventory control problem, managers shall have the information on the demand distributions

^{2.1}Another important approach in the revenue management is dynamic pricing. However, according to Belobaba (1987), the dynamic price changes may lead to an irrational price war with the competitors. Unlike the price adjustment, the seat inventory control aims to allocate the seats under multiple fares properly, which is easy to manipulate and hidden from the competitors, and thus it becomes a practically feasible strategy in revenue management.

for different fare tickets and the buy-up substitution probability. For the airline industry, the demand learning is very essential. Jack Bovey, the revenue optimization manager at British Airways, said that “For us, it was the first step towards what will hopefully become an important part of how we forecast, and hence price, flights.” (see The Alan Turing Institute, 2020) The information on the demand and the buy-up substitution probability has to be learned from historical sales data, which, in turn, are affected by the past inventory allocation decisions. How to dynamically allocate seats among different fare tickets through multiple time periods, with learning on both demand distribution and buy-up substitution probability? This is an important decision problem faced by airline managers. Here, we aim to investigate this problem and provide a solution.

Considering the tractability of analysis, we restrict to a representative simplified setting in this work. There exists a fixed amount of seats, which can be sold in two phases, an early-bird phase with price discount followed by a regular-price phase. Customers are segmented into two types, early-bird customers who prefer the discounted seats and regular customers who prefer the regular-price seats. If the stock-out occurs, unsatisfied early-bird customers may simply leave or choose to buy regular-price seats for substitution. We first present a baseline *single-period* inventory decision problem with buy-up substitution. Based on the optimal inventory allocation of the single-period model, we consider a myopic policy for the airline company in a multi-period setting: in each period, the company updates its belief on demand and substitution probability and then based on the updated belief, it adopts the one-period optimal inventory allocation decision. With a numerical example, we show that with and without considering buy-up substitution yield quite different outcomes, and ignoring buy-up substitution could lead to a 53.31% loss in revenue. The myopic policy, though easy to be implemented in practice, does not consider the impact on the future information gaining of the seat allocation decision. Hence, we then construct a dynamic programming model to investigate the optimal multi-period inventory management problem with learning on the demand and buy-up substitution probability. To generate insights and inspire heuristic algorithms for inventory managers, we mainly compare the Bayesian optimal inventory levels with the corresponding myopic inventory levels. In the following analysis, the “stock more (resp. less)” result means that the inventory level under the Bayesian inventory management is larger (resp. smaller) than the corresponding one under the myopic decision rule.

There exist two types of *censored data* in our setting. If lost sales are unobservable, the sales data provide censored demand information for the airline company. For example, if the

airline company sells out all early-bird tickets of a flight, say 50 seats, then the company knows that the early-bird demand is at least 50 but it does not have the exact number for it. The second type of censored data comes from early-bird customers' buy-up substitution behavior: the primary demand and the substitution demand are tangled. For example, suppose that the company observes that stockout happens during the early-bird-discount phase and in the following regular-price phase, the sales amount of regular-price seats is 100. Then the airline company does not know how many of these 100 sales come from the regular customers and how many from early-bird customers who encounter stockout of discounted seats. According to whether or not lost sales are observable and whether or not the substitution demand can be separated out from the primary demand for the regular-price seat, we have four information scenarios as listed in Table 2.1. The complete observation case is scenario \mathcal{OS} and the least observation case is \mathcal{UT} . There are two partial observation cases: scenario \mathcal{OT} and scenario \mathcal{US} .

Table 2.1: Four Information Scenarios with Acronyms

		Substitution Demand	
		<i>Separated</i>	<i>Tangled</i>
Lost Sales	<i>Observable</i>	\mathcal{OS}	\mathcal{OT}
	<i>Unobservable</i>	\mathcal{US}	\mathcal{UT}

For each information scenario, we derive the Bayesian updating formula and the Bellman equation of our dynamic programming model for finding the optimal inventory allocation decisions. Our focus is on investigating whether the classical “stock more” or “stock less” result still holds in our setting.

Table 2.2 summarizes our main comparison results under different information scenarios and pinpoints our contributions to the Bayesian inventory management literature. In the table, the symbol “ \gtrsim ” represents that the relationship can be either “ \geq ” or “ \leq ”.

We first consider the setting where the substitution probability is known and only the demand parameter needs to be estimated. We reach the following conclusions.

- In the complete-observation scenario \mathcal{OS} , the inventory manager does not need to increase the inventory level to gain more information on the demand, yielding the same decision as the myopic one.
- In the partial-observation scenario \mathcal{OT} , lost sales are observable but there exist tangling demands for the regular-price seat. To better estimate the primary demand for

the regular-price seat, one needs to stock more discounted seats so as to reduce the occurrence of the buy-up demand substitution.

- In the partial-observation scenario US , lost sales are unobservable. There is a famous “stock more” result in the Bayesian inventory management literature, that is, the optimal inventory level shall be set higher than the myopic inventory level to gain a better observation of demand; see, e.g., Lariviere and Porteus (1999). Chen and Plambeck (2008) further demonstrate that the unobservable lost sales still lead to the “stock more” result under the assumption that the substitute product is always available. Different from their setting, the limited seat capacity here does not guarantee the complete observations on the substitution demand. And we provide an example showing that the unobservability of lost sales may even be a “stock less” driving force. Intuitively, the information on the primary demand can be obtained not only directly through the sales amounts but also indirectly from the observation of buy-up substitutions. Since one can infer demand parameter from the substitutions, discounted seats can be stocked less to induce more substitution trials.

We then consider the setting where both the demand parameter and the substitution probability need to be estimated. To better learn the information on the substitution probability requires “stocking less” discounted seats so that substitution can happen more frequently. Such a “stock less” driving force is similar to the one in Chen and Plambeck (2008). Due to the interplay of the multiple driving forces, there are generally ambiguous results on “stock more” or “stock less” when at least one source of demand censoring exists.

The total inventory level is assumed to be a constant in our model. This fits well the airline industry where the size of the airplane is fixed. We also consider a more general setting with the total inventory level to be a decision variable. This might fit some other two-phase selling situations, such as food catering with the early-bird discount policy. We find that our main insights still hold in such a general setting. We refer the interested readers to the online Appendix [A.1](#) for the detailed analysis and discussion of this generalized model.

Our main contributions are as follows. First, our results can help the airline company to optimally determine the booking limit for the discounted seats using the effective Bayesian inventory management. Second, we identify conflicting driving forces behind the optimal inventory level decision and find that the classic “stock more” result may not hold anymore, which enriches the Bayesian inventory management literature.

The remainder of this chapter is organized as follows. We review related literature in

Table 2.2: Bayesian Optimal Inventory Levels versus the Corresponding Myopic Ones under Different Information Scenarios

	Information scenario	Relationship between Bayesian-optimal solutions and myopic-optimal ones	Contributions to the Bayesian inventory management literature
Updating demand parameter only (see Section 2.4.2)	\mathcal{OS}	$y^{\mathcal{OS}} = y^m$	Consistent with the literature.
	\mathcal{OT}	$y^{\mathcal{OT}} \geq y^m$	We identify a new driving force for “stocking more” the discounted product: reduce the tangling of the primary demand and the substitution demand.
	$US \& UT$	$y^{US(\setminus UT)} \geq y^m$	<ol style="list-style-type: none"> 1. Under a special scenario \widetilde{US}, we extend the traditional “stock more” result driven by the unobservable lost sales from one kind of demand to two kinds, the observations on which are affected by the decision variable at the same time. 2. Another striking finding is that the unobservable lost sales may become a “stock less” driving force. Chen and Plambeck (2008) get the “stock more” result by assuming that the substitute is always available, which may not hold anymore with a limited capacity of substitutes. When the sales amounts contain limited demand information, then one can further deduce such information from the substitutions, which drives the “stock less” result to generate more substitution trials.
Updating both demand parameter and substitution probability (see Section 2.4.3)	\mathcal{OS}	$y^{\mathcal{OS}} \leq y^m$	Better estimation of the substitution probability requires “stocking less” the discounted product, a result consistent with that in Chen and Plambeck (2008). Yet, it generalizes from a setting with unlimited capacity of substitutes and sole source of demands for the substitutes in Chen and Plambeck (2008) to a setting with limited capacity of substitutes and two sources of demands for the substitutes.
	$\mathcal{OT} \& US \& UT$	$y^{\mathcal{OT}(\setminus US \setminus UT)} \geq y^m$	Multiple driving forces co-exist. The numeric study shows that the final comparison result relies on factors such as Bayesian manager’s prior beliefs and seats’ prices.

Section 2.2. In Section 2.3, we present a baseline one-period model of optimizing the booking limit of the discounted seats under the airline setting. After that, we construct a multi-period model with learning about the demand parameter and the substitution probability under Bayesian inventory management in Section 2.4, where we mainly compare the Bayesian optimal inventory level with the corresponding myopic one under four information scenarios. In Section 2.5, we provide the numeric study. Section 2.6 concludes this chapter. We analyze the case in which both the total inventory level and the inventory level of the discounted product are decision variables in the online Appendix A.1. All the proofs are relegated to the online Appendix A.2.

2.2 Literature Review

Our study is closely related to the literature on Bayesian inventory management. In the early stage of this research stream, researchers mainly consider settings with observable lost sales. Scarf (1959) formulates a Bayesian inventory dynamic programming model with two state variables (inventory level and demand parameter). It is shown later in Scarf (1960) that the problem can be reduced into one state variable with a gamma demand distribution. Succeeding studies such as Azoury (1985) and Miller (1986) extend Scarf’s method to other demand distributions. Lovejoy (1990) provides myopic policies by reducing the single state to zero-dimensional state space, i.e., a static optimization problem. Later, researchers start to consider settings with lost sales being unobservable. Such demand censoring brings difficulty to the demand estimation. According to Braden and Freimer (1991), only special types of distribution, namely the newsvendor distributions defined by them, allow parsimonious information updating. By utilizing the newsvendor distribution and Scarf’s method on state-space reduction, Lariviere and Porteus (1999) successfully obtain analytical results on the optimal inventory decision in a multi-period newsvendor setting with unobservable lost sales. They demonstrate that the “stock more” result holds. Ding et al. (2002) further extend their model by considering a general demand distribution with perishable products and show that the “stock more” conclusion still holds. The proof of this conclusion is later rectified by Lu et al. (2005) and further simplified by Bensoussan et al. (2009). Chen (2010) develops bounds and heuristics for the optimal solutions by considering a single-product periodic-review inventory control problem. Here, we also develop an upper bound for our early-bird-discount model when both the total inventory level and the inventory level of the discounted product are

decision variables (see the online Appendix [A.1](#)). As the structure of our model is much more complicated, we adopt an enlarging technique to make the upper bound analysis feasible. Jain et al. (2015) and Bensoussan and Guo (2015) utilize the information on the stock-out time to estimate demand distribution with perishable and non-perishable products, respectively. Both papers demonstrate that when lost sales are unobservable, managers can utilize this little bit more information about stock-out times to improve the profit. Bensoussan et al. (2016) consider the incomplete inventory and demand information caused by the invisible demand such as spoilage, damage, pilferage and returns. They study the inventory management problem with only sales information and develop an iterative algorithm to solve the problem approximately.

Among all the studies on Bayesian inventory management, Chen and Plambeck (2008) is the first one considering the substitution issue when the stockout occurs. Our research also considers the substitution issue but it differs greatly from Chen and Plambeck (2008). First, Chen and Plambeck (2008) assume that the substitutable product is always available when customers' desired product is stocked-out. However, our work considers the buy-up substitution of using the regular-price product to substitute the early-bird-discount product, and thus the capacity of the substitutable product in our model is limited. Additionally, the demand for the substitutable product in Chen and Plambeck (2008) only comes from the stockout-based substitution, while in our model, such demand comes from two sources, the buy-up substitution and the primary demand for the regular-price product.

In contrast to the multi-period setting in the aforementioned studies, there exists a stream of literature studying profit maximization in a one-period setting with Bayesian learning. This type of studies ignore the estimation-and-optimization cycle but focus on issues such as shrinkage and pricing; see, e.g., Li and Ryan (2011), Harrison et al. (2012) and Li et al. (2019).

Our work is also related with studies of inventory management with demand estimation based on censored demand observation. Some studies utilize expectation-maximization (EM) algorithm to estimate the demand and substitution probability parameters, including Anupindi et al. (1998), Kök and Fisher (2007), Ulu et al. (2012), Vulcano et al. (2012) and Chen and Chao (2019). Some studies utilize non-parametric approaches in demand learning, including Huh and Rusmevichientong (2009), Feng and Shanthikumar (2017), Chen and Chao (2020) and Yuan et al. (2021), and some develop operational statistics to integrate demand estimation and inventory optimization together including Liyanage and Shanthikumar (2005)

and Chu et al.(2008).

Finally, our study is related with the seat allocation problem in the airline revenue management literature. According to McGill and van Ryzin (1999), the early bird discount selling strategy was first adopted by the airline companies such as BOAC (now British Airways) in the early 1970s. By doing so, airline companies could gain extra revenue from selling the seats that would be empty without offering discounts. Littlewood (1972) provides an optimal rule for this from the perspective of benefit maximization, which lays a foundation for many yield control models. This rule is extended by Belobaba (1987) to multiple fare classes by using the Expected Marginal Seat Revenue Model (EMSR) and Pfeifer (1989) obtains a similar result but in a different approach. Using the marginal analysis as in Belobaba (1987), Brumelle et al. (1990) formally prove that a variant of Littlewoods rule could be optimal under a general model of the seat allocation problem. van Ryzin and McGill (2000) provide a simple adaptive approach to optimize seat protection levels. Cooper et al. (2006) demonstrate that simply following the Littlewood’s rule without considering the buy-up substitution can cause a serious spiral-down effect, resulting in a big revenue loss. Cooper and Li (2012) further demonstrate the benefit of incorporating the buy-up substitution into the airline seat management problem. All of these studies do not combine the learning and inventory decision together. Ours is the first one combining the dynamic learning of the demand distribution and substitution probability and inventory decision together for the airline seat inventory management problem.

In our two-phase selling model, the first-phase price is lower than that of the second phase. A symmetric setting exists in business practice with the first-phase price higher than the second phase’s. Hu et al. (2015) study such a markdown inventory management problem. In their model, there are also two selling phases in each period, a clearance phase (modeled as the first phase) with a lower price and a regular-sales phase (modeled as the second phase) with a full price. Those customers who do not get the products in the clearance phase can choose to substitute in the following regular-sales phase. The main differences between their markdown model and our early bird discount model are that the inventory used in their clearance phase is part of those unsold leftover products from the previous period, that is, they are not newly produced, and the leftover products from the clearance phase cannot be sold in the following regular-sales phase. Thus, the selling periods in their markdown setting are inter-correlated, while the selling periods are independent in our early bird discount model. Another main difference is that Hu et al. (2015) only consider a static model and there is

no learning about the demand parameter and the substitution probability while we consider learning of both.

2.3 One-period Model

In this section, we first review a baseline one-period inventory management problem with two selling phases. The optimal inventory decision can be expressed in a similar way as a newsvendor-problem solution. Next, we present a myopic decision policy where the company repeatedly makes the one-period decision along the time line, with demand and substitution probability information updated according to Bayes' rule. We then construct a numerical example to illustrate the importance of incorporating the buy-up substitution probability into the model.

2.3.1 Model Description

Past studies on airline seat protection problems mainly focus on a single period setting, see, e.g., Brumelle et al. (1990). For completeness, we briefly review this model and state it with our notations. Consider a single selling period with two selling phases: an early-bird-discount phase and a regular-price phase. The corresponding selling prices are denoted as p_1 and p_2 , respectively, where $p_1 < p_2$. The primary demands for the discounted seat and the regular-price seat are D_1 and D_2 , respectively, which can be correlated. In this work, we consider them to be discrete random variables. Let $f_{12}(\cdot, \cdot | \theta)$ be their joint probability mass function, where θ is an unknown parameter with $\theta \in \Theta$. (Note that θ can be a vector of unknown parameters.) Denote the marginal probability mass functions of D_1 and D_2 as $f_1(\cdot | \theta)$ and $f_2(\cdot | \theta)$, respectively. Let M be the total number of available seats of the flight, which is a fixed number. The firm's objective is to determine the optimal amount of the discounted seats, denoted as y , to maximize its total expected profit over the two selling phases.

There exists a tradeoff associated with the inventory decision y . When the firm allocates too few seats for the early-bird-discount sales (i.e., y is very small), the primary demand for the early-bird-discount seat may not be fully satisfied and some of them may be lost, losing the opportunity to sell more. On the other hand, if the firm allocates too many seats for the early-bird-discount sales (i.e., y is very big), the firm may lose a chance to force some customers to buy-up their seats, because when stockout happens for the discounted seats, some customers who come for the discounted seats may choose to buy the regular-price seats.

As the primary demand for the discounted seat is D_1 , the realized sales of the discounted seat can be expressed as $D_1 \wedge y$, where $a \wedge b = \min(a, b)$. If there are leftover seats at the end of the early-bird-discount phase, they are sold in the regular-price phase as well. Thus, the amount of inventory available for the regular-price sales is $(M - y \wedge D_1)$. Note that the demand for the regular-price seat comes from two sources: those from unsatisfied customers in the early-bird-discount phase who choose the buy-up substitution, denoted by a random variable K , and the primary demand for the regular-price seat (i.e., D_2). We assume that each unsatisfied customer's substitution decision is a Bernoulli trial with probability α , called the *buy-up substitution probability*. The random variable K then follows a binomial distribution with parameters $((D_1 - y)^+, \alpha)$, where $x^+ = \max(0, x)$. Under given values of θ and α , the firm makes the inventory level decision to maximize its total expected profit $\pi(y|\theta, \alpha)$ as follows:

$$\begin{aligned} \max_y \pi(y|\theta, \alpha) &= p_1 E[D_1 \wedge y|\theta] + p_2 E[(K + D_2) \wedge (M - y \wedge D_1)|\theta, \alpha] \\ \text{s.t. } &0 < y \leq M. \end{aligned} \quad (2.1)$$

According to Brumelle et al. (1990), the optimal inventory level of the discounted seat y^* can be expressed as follow:

$$y^* = \max \left\{ 0 < y \leq M : Pr(K + D_2 > M - y | D_1 \geq y, \theta, \alpha) < \frac{p_1 - \alpha p_2}{(1 - \alpha)p_2} \right\}. \quad (2.2)$$

Specifically, if the substitution probability α is 0 (i.e., $K = 0$ with probability 1), the above optimal solution y^* reduces to the result of Littlewood's rule (Littlewood, 1972).

We have the following conclusion about the sensitivity of the optimal decision y^* with respect to the substitution probability α .

Proposition 2.1. *The profit function $\pi(y|\theta, \alpha)$ is submodular in (y, α) ; that is, $\partial[\pi(y + 1|\theta, \alpha) - \pi(y|\theta, \alpha)]/\partial\alpha < 0$. Therefore, the optimal level of the discounted inventories y^* decreases with the substitution probability α .*

Proposition 2.1 shows that, if the substitution probability is larger, the inventory manager should set a lower inventory level for the discounted seats. The behind reason is that those unsatisfied customers are more likely to buy the regular-price seats.

2.3.2 Repeated Decision Making

The one-period model assumes that the demand and substitution probability parameters are known. In practice, such information is often unknown to the inventory manager, who may not even realize the existence of buy-up substitution. When decisions are made repeatedly across multiple periods, unknown parameters can be learned from past sales. We construct a simple numerical example to illustrate repeated inventory decisions with learning on demand information through Bayesian updating. In particular, we consider scenarios with and without incorporating the buy-up substitution into demand estimation and compare the results.

Example 2.1. (*Repeated Decisions with and without Incorporating Buy-up Substitution*) Consider that an airline company offers an early bird sale of a flight with a medium-sized jet. The total number of seats is $M = 120$. The regular price is set at $p_2 = 2000$, and the early-bird-discount price $p_1 = 750$. Lost sales are unobservable, and the substitution demand and the primary demand for the regular-price seat cannot be separated. The primary demands for the discounted and normal-price seats are D_1 and D_2 following truncated Poisson distributions ($0 \leq D_1 \leq 160$, $0 \leq D_2 \leq 160$) with parameters λ_1 and λ_2 , respectively. The demand parameter θ may take value 1 or 2. When $\theta = 1$, $\lambda_1 = 36$ and $\lambda_2 = 20$; and when $\theta = 2$, $\lambda_1 = 80$ and $\lambda_2 = 0$. The substitution probability α may be low ($\alpha = 0.1$) or high ($\alpha = 0.8$).

When the inventory manager is aware of buy-up substitution, the repeated decisions can be made as follows. At the beginning of the first selling period, the inventory manager holds the prior beliefs $(Pr(\theta = 1), Pr(\theta = 2)) = (0.8, 0.2)$ and $(Pr(\alpha = 0.1), Pr(\alpha = 0.8)) = (0.5, 0.5)$. The optimal inventory allocation in period 1 can be calculated based on the one-period optimal inventory decision formula (2.2). In the following period, the inventory manager first updates the beliefs on demand and substitution probability based on the observed sales data, and then makes the optimal inventory allocation decision by adopting again the formula (2.2). This policy repeats for the remaining periods. In a similar vein, one can calculate the corresponding optimal inventory allocation decisions when the manager ignores the buy-up substitution. Denote the optimal solutions in period i ($i \geq 1$) with and without considering the buy-up substitution as y_i^{sub} and y_i^{no} , respectively.

Now suppose that the underlying true parameter values are $\theta = 2$ and $\alpha = 0.8$. That is, all customers belong to early-bird customers and they will buy regular-price seats with a probability of 0.8 if the discounted seats are stocked out.

When the inventory manager does not realize the buy-up substitution issue, we can calcu-

late his optimal inventory decision on the discounted seats in the first period as $y_1^{no} = 106$. His belief is then updated according to Bayes' rule. It can be shown that, with a probability of 0.99, the observed sales data is $(k, 0)$, where $59 \leq k < 106$. That is, under this situation, the company only observes the sales in the early-bird-discount phase. Hence, for $59 \leq k < 106$, it follows from Bayes' rule that the posterior belief that $\theta = 2$ is

$$\Pr(\theta = 2|(k, 0)) = \frac{0.2 \times \frac{80^k e^{-80}/k!}{\sum_{i=0}^{160} 80^i e^{-80}/i!}}{0.2 \times \frac{80^k e^{-80}/k!}{\sum_{i=0}^{160} 80^i e^{-80}/i!} + 0.8 \times \frac{36^k e^{-36}/k!}{\sum_{i=0}^{160} 36^i e^{-36}/i!} \times \frac{20^0 e^{-20}/0!}{\sum_{i=0}^{160} 20^i e^{-20}/i!}} \approx 1.00.$$

Therefore, after the first period, the optimal inventory decision for the second period becomes $y_2^{no} = 120$. As the protection level of the regular-price seats reduces to 0, the manager can never observe the sales of the regular-price seats. It then follows that the zero protection level shall be kept in all the following periods. By contrast, when the substitution issue is considered, the belief of the inventory manager can be updated very close to the underlying true value (i.e., $\Pr(\theta = 2, \alpha = 0.8) > 0.95$) on almost all the sample paths after five periods according to our simulation study, and the optimal inventory level correspondingly becomes $y_i^{sub} = 1$ ($i \geq 6$). The simulation result shows that from the sixth period, the ignorance of the substitution issue causes a 53.31% loss in revenue on average.

Example 2.1 shows that, when demand information is updated according to Bayes' rule, ignoring substitution can cause a severe loss of revenue. This finding is similar to the *spiral-down* effect illustrated in Cooper et al. (2006), who demonstrate that, if buy-up substitution is ignored by the manager, the protection level for regular-price seats will be low, and the regular-price sales will decrease, resulting in lower future estimates of demand for regular-price seats, which, in turn, leads to a lower protection level. In Cooper et al. (2006), demand information is updated in a non-Bayesian way. Here, we demonstrate that the similar spiral-down effect also exists if demand is updated according to Bayes' rule.

Repeated decision making based on the one-period model, although it is simple and easy to implement, has a big defect: it ignores the effect of the current inventory decision on learning demands and substitution probability in the following periods. Hence, we call it the *myopic decision* in a multi-period setting. In the following section, we will present a dynamic programming model, which considers not only the current-period revenue but also the efficiency of learning on demands and substitution probability in the following periods.

2.4 Multi-period Bayesian Inventory Management

We now consider a multi-period inventory management model with learning on the unknown demand parameter θ and the unknown buy-up substitution probability α . The firm's objective is to maximize the total discounted expected profit in N periods, where the discount factor is denoted by δ ($0 < \delta \leq 1$). The selling phases in each period are the same as those in the one-period setting.

We first describe the four information scenarios and derive the Bayesian learning formula for unknown parameters in each scenario. We then specifically consider two settings. In the first setting, the inventory manager cares most about the demand parameter by assuming that the value of the substitution probability parameter α is known from the prior knowledge and experience. Thus, only the demand parameter θ needs to be learned. In the second setting, both the demand parameter θ and the substitution probability α need to be learned. For each setting, we formulate the corresponding dynamic programming model for the optimal inventory decisions. Recall that the decision variable is the inventory level of the discounted seats y . We shall conduct the comparison between the Bayesian optimal inventory level and the corresponding myopic one, where the myopic one maximizes only that period's expected profit without considering the decision's impact on the future demand-information learning.

2.4.1 Four Information Scenarios and Bayesian Learning

In our study, there are two sources of demand censoring. One, lost sales may not be observable. Two, the demand for the regular-price seat is composed of both the substitution demand and the primary demand for the regular-price seat, which may not be separated. We then have four information scenarios based on whether or not lost sales are observable and the substitution demand can be separated out, as shown in Table 2.3.

Table 2.3: Four Information Scenarios with Available Information

		Substitution Demand	
		<i>Separated</i>	<i>Tangled</i>
Lost Sales	<i>Observable</i>	$\mathcal{OS}(x_1, x_{21}, x_{22})$	$\mathcal{OT}(x_1, x_2)$
	<i>Unobservable</i>	$\mathcal{US}(s_1, s_{21}, s_{22})$	$\mathcal{UT}(s_1, s_2)$

When lost sales are observable and the substitution demand can be separated out (denoted as the \mathcal{OS} scenario), we have the complete observations: the realized demands for the

discounted and regular-price seats, x_1 and x_2 are both observable. Moreover, in the composition of x_2 , the substitution demand x_{21} and the primary demand for the regular-price seat x_{22} are also known. This separation observation is feasible under some cases where according to experience, all primary demands for the discounted seats can be reasonably assumed to arrive only in the early-bird-discount phase and all primary demands for the regular-price seats arrive only in the regular-price phase. And thus, the substitutions happen only in the early-bird-discount phase, and the sales of the regular-price seats in the early-bird-discount (resp. regular-price) phase come only from the substitution demands (resp. primary demands for the regular-price seats).

Given the inventory level y , the demand parameter θ and the substitution probability α , the likelihood of observing demand realizations x_1 , x_{21} and x_{22} can be written as

$$f_{OS}^y(x_1, x_{21}, x_{22}|\theta, \alpha) = \begin{cases} f_{12}(x_1, x_{22}|\theta) \binom{x_1-y}{x_{21}} \alpha^{x_{21}} (1-\alpha)^{x_1-y-x_{21}}, & \text{if } x_1 > y; \\ f_{12}(x_1, x_{22}|\theta), & \text{if } x_1 \leq y. \end{cases}$$

When lost sales are unobservable but the substitution demand can be separated out (denoted as the \mathcal{US} scenario), sales for the discounted seats s_1 and that for the regular-price seats s_2 are both observable. Moreover, s_2 can be separated as sales from the substitution demand s_{21} and sales from the primary demand for the regular-price seat s_{22} . Note that sales quantities s_1 , s_{21} and s_{22} are censored data of demand realizations x_1 , x_{21} and x_{22} , respectively. The likelihood function takes different expressions depending on whether the stockout happens and if it indeed happens, which kind of seats is stocked out. For example, consider $s_1 = y$, $s_{21} < M - y$ and $s_{22} = M - y - s_{21}$. The condition $s_1 = y$ indicates that the discounted seat is sold out. Among those excess demands in D_1 , only s_{21} unsatisfied customers choose to buy the regular-price seats as substitutes, which is observable. Hence, the primary demand for the discounted seat shall be no less than $y + s_{21}$. When $s_{22} = M - y - s_{21}$, it indicates that the regular-price product is sold out in the regular-price phase, and hence the primary demand for the regular-price seat shall be no less than $M - y - s_{21}$. The likelihood of observing the aforementioned sales quantities in two phases can be written as

$$\sum_{i=y+s_{21}}^{+\infty} \sum_{j=M-y-s_{21}}^{+\infty} f_{12}(i, j|\theta) \binom{i-y}{s_{21}} \alpha^{s_{21}} (1-\alpha)^{i-y-s_{21}}.$$

Analogously, we can derive the likelihood functions of other observations. We can show that

under the \mathcal{US} scenario, the likelihood of observing sales quantities (s_1, s_{21}, s_{22}) is

$$f_{\mathcal{US}}^{y,M}(s_1, s_{21}, s_{22}|\theta, \alpha) = \begin{cases} f_{12}(s_1, s_{22}|\theta), & \text{if } s_1 < y, s_{21} = 0 \text{ and } s_{22} < M - s_1; \\ \sum_{j=M-s_1}^{+\infty} f_{12}(s_1, j|\theta), & \text{if } s_1 < y, s_{21} = 0 \text{ and } s_{22} = M - s_1; \\ \sum_{i=y+s_{21}}^{+\infty} f_{12}(i, s_{22}|\theta) \binom{i-y}{s_{21}} \alpha^{s_{21}} (1-\alpha)^{i-y-s_{21}}, & \text{if } s_1 = y, s_{21} < M - y \text{ and } s_{22} < M - y - s_{21}; \\ \sum_{i=y+s_{21}}^{+\infty} \sum_{j=M-y-s_{21}}^{+\infty} f_{12}(i, j|\theta) \binom{i-y}{s_{21}} \alpha^{s_{21}} (1-\alpha)^{i-y-s_{21}}, & \text{if } s_1 = y, s_{21} < M - y \text{ and } s_{22} = M - y - s_{21}; \\ \sum_{j=M-y}^{+\infty} \sum_{i=y+j}^{+\infty} f_1(i|\theta) \binom{i-y}{j} \alpha^j (1-\alpha)^{i-y-j}, & \text{if } s_1 = y, s_{21} = M - y \text{ and } s_{22} = 0. \end{cases}$$

In some other real settings, some customers may strictly prefer the regular-price seat, and purchase it once it is available. So, some primary demands for the regular-price seats may arrive in the early-bird-discount phase after the discounted seats sell out and the regular-price ones become available. On the other hand, some unsatisfied demands for the discount may choose to substitute in the regular-price phase considering that they may hesitate to pay a higher price. In this case, the airline manager cannot distinguish the substitution demands and the primary demands for the regular-price seats after the stockout of the discounted seats. The \mathcal{OT} information scenario considers observable lost sales and such tangled demands for the regular-price seat.

If the demand realization for the discounted seat x_1 satisfies $x_1 \leq y$, no substitution will happen. Then, the demand realization for the regular-price seat x_2 fully represents the realized primary demand D_2 . Hence, the likelihood of observing the demand realization (x_1, x_2) , where $x_1 \leq y$, is $f_{12}(x_1, x_2|\theta)$. When $x_1 > y$, there must exist some unsatisfied demand in the early-bird-discount phase, and hence x_2 can be the sum of the substitution demand from those unsatisfied customers in the early-bird-discount phase and the primary demand for the regular-price seat. Suppose that the substitution demand is i . Then, the likelihood of observing (x_1, x_2) can be written as

$$\binom{x_1 - y}{i} \alpha^i (1 - \alpha)^{x_1 - y - i} f_{12}(x_1, x_2 - i|\theta).$$

As the substitution demand i is latent, one needs to sum up above likelihoods over all the possible values of i , which fall within the range $[0, (x_1 - y) \wedge x_2]$. Hence, the likelihood of observing the demand realization (x_1, x_2) , where $x_1 > y$, is

$$\sum_{i=0}^{(x_1-y) \wedge x_2} \binom{x_1-y}{i} \alpha^i (1-\alpha)^{x_1-y-i} f_{12}(x_1, x_2 - i | \theta).$$

In summary, under the \mathcal{OT} scenario, the likelihood of observing the demand realization (x_1, x_2) can be written as

$$f_{\mathcal{OT}}^y(x_1, x_2 | \theta, \alpha) = \begin{cases} \sum_{i=0}^{(x_1-y) \wedge x_2} \binom{x_1-y}{i} \alpha^i (1-\alpha)^{x_1-y-i} f_{12}(x_1, x_2 - i | \theta), & \text{if } x_1 > y; \\ f_{12}(x_1, x_2 | \theta), & \text{if } x_1 \leq y. \end{cases}$$

The least information scenario is the one where lost sales are unobservable and the primary demand for the regular-price seat is tangled with the substitution demand from those unsatisfied customers in the early-bird-discount phase (denoted as the \mathcal{UT} scenario). Under the \mathcal{UT} scenario, one can only observe sales quantities s_1 and s_2 for the discounted seat and the regular-price seat, respectively. The exact expression of the likelihood function depends on whether the inventory stockout happens or not and, if it happens, which seat is stocked-out. For example, consider $s_1 = y$ and $s_2 = M - y$. $s_1 = y$ implies that the discounted seats are sold out, and $s_2 = M - y$ indicates that the regular-price seats are also sold out. Therefore, we can infer that the primary demand for the discounted seat is no less than y and the tangled demand (the sum of the substitution demand and the primary demand for the regular-price seat) is no less than $M - y$. In this case, to derive its likelihood, we shall first consider the complete observation of a realized primary demand i in the early-bird-discount phase, the substitution demand j , and the realized tangled demand k for the regular-price seat. The corresponding likelihood can be written as

$$\binom{i-y}{j} \alpha^j (1-\alpha)^{i-y-j} f_{12}(i, k - j | \theta).$$

We can then sum over all the possible values of i , j and k to obtain the likelihood of observing sales quantities (s_1, s_2) . Note that the substitution demand j should be no more than both $i - y$, the excess (unsatisfied) demand for the discounted seat and k , the tangled demand for the regular-price seat. Then, the likelihood of observing sales quantities $s_1 = y$ and $s_2 = M - y$

can be derived as

$$\sum_{i=y}^{+\infty} \sum_{k=M-y}^{+\infty} \left[\sum_{j=0}^{(i-y)\wedge k} \binom{i-y}{j} \alpha^j (1-\alpha)^{i-y-j} f_{12}(i, k-j|\theta) \right].$$

Similarly, we can derive the likelihoods of all the possible observations, which are summarized as follows:

$$f_{\mathcal{UT}}^y(s_1, s_2|\theta, \alpha) = \begin{cases} f_{12}(s_1, s_2|\theta), & \text{if } s_1 < y \text{ and } s_2 < M - s_1; \\ \sum_{k=M-s_1}^{+\infty} f_{12}(s_1, k|\theta), & \text{if } s_1 < y \text{ and } s_2 = M - s_1; \\ \sum_{i=y}^{+\infty} \left[\sum_{j=0}^{(i-y)\wedge s_2} \binom{i-y}{j} \alpha^j (1-\alpha)^{i-y-j} f_{12}(i, s_2-j|\theta) \right], & \text{if } s_1 = y \text{ and } s_2 < M - y; \\ \sum_{i=y}^{+\infty} \sum_{k=M-y}^{+\infty} \left[\sum_{j=0}^{(i-y)\wedge k} \binom{i-y}{j} \alpha^j (1-\alpha)^{i-y-j} f_{12}(i, k-j|\theta) \right], & \text{if } s_1 = y \text{ and } s_2 = M - y. \end{cases}$$

Let I_{scen}^y denote the information set that contains all the available information for a given inventory level y under the information scenario $scen$, where $scen \in \{\mathcal{OS}, \mathcal{OT}, \mathcal{US}, \mathcal{UT}\}$. For example, under the \mathcal{OS} scenario where lost sales are observable and the substitution demand can be separated out, $I_{\mathcal{OS}}^y = \{(x_1, x_{21}, x_{22}) : 0 \leq x_{21} \leq (x_1 - y)^+, x_1, x_{21}, x_{22} \in N_+\}$, where N_+ is a set of all nonnegative integers; that is, the information set contains all possibilities of both the realized primary demands in two phases and the substitution demand. Similarly, under the \mathcal{UT} scenario where lost sales are unobservable and the substitution demand is tangled with the primary demand, $I_{\mathcal{UT}}^y = \{(s_1, s_2) : 0 \leq s_1 \leq y, 0 \leq s_2 \leq M - s_1, s_1, s_2 \in N_+\}$; that is, the information set contains all possibilities of observed sales quantities.

Denote the joint prior distribution of θ and α in period i ($i = 1, 2, \dots, N$) as $\phi_i(\theta, \alpha)$. Given $\phi_i(\theta, \alpha)$ for period i , the posterior distribution $\phi_{i+1}(\theta, \alpha)$ derived based on the information observed in period i serves as the prior for the following time period $i + 1$. Under each information scenario $scen \in \{\mathcal{OS}, \mathcal{OT}, \mathcal{US}, \mathcal{UT}\}$, given the data observations for period i , $\xi \in I_{scen}^y$, the posterior distribution $\phi_{i+1}(\theta, \alpha)$ can be derived by using the corresponding likelihood function according to Bayes' rule as follow:

$$\phi_{i+1}(\theta, \alpha|\xi, y, \phi_i) = \frac{f_{scen}^y(\xi|\theta, \alpha)\phi_i(\theta, \alpha)}{\int_0^1 \int_{\Theta} f_{scen}^y(\xi|\theta', \alpha')\phi_i(\theta', \alpha')d\theta'd\alpha'}. \quad (2.3)$$

Let $v_i^{scen}(\phi_i)$ be the firm's maximum total discounted expected profit over periods i to N with a prior distribution ϕ_i for period i ($i = 1, 2, \dots, N$) under the information scenario $scen$, where $scen \in \{\mathcal{OS}, \mathcal{OT}, \mathcal{US}, \mathcal{UT}\}$. Then, we can write the Bayesian dynamic optimality equations as

$$v_i^{scen}(\phi_i) = \max_{0 < y \leq M} E_{\phi_i(\theta, \alpha)} \left\{ \pi(y|\theta, \alpha) + \delta \sum_{\xi \in I_{scen}^y} v_{i+1}^{scen}(\phi_{i+1}) f_{scen}^y(\xi|\theta, \alpha) \right\}, i = 1, \dots, N - 1,$$

and

$$v_N^{scen}(\phi_N) = \max_{0 < y \leq M} E_{\phi_N(\theta, \alpha)} \{ \pi(y|\theta, \alpha) \}.$$

For ease of exposition, we use $G_i^{scen}(y, \phi_i)$ to denote the corresponding objective function of $v_i^{scen}(\phi_i)$ ($i = 1, \dots, N$), which means that

$$G_i^{scen}(y, \phi_i) = E_{\phi_i(\theta, \alpha)} \left\{ \pi(y|\theta, \alpha) + \delta \sum_{\xi \in I_{scen}^y} v_{i+1}^{scen}(\phi_{i+1}) f_{scen}^y(\xi|\theta, \alpha) \right\}, i = 1, \dots, N - 1,$$

and

$$G_N^{scen}(y, \phi_N) = E_{\phi_N(\theta, \alpha)} \{ \pi(y|\theta, \alpha) \}.$$

The myopic inventory level in period i ($i = 1, \dots, N$) maximizes only that period's expected profit and is denoted as y_i^m . Hence, it is the optimal solution of the corresponding one-period model with prior belief $\phi_i(\theta, \alpha)$. For ease of exposition, we use $G_i^m(y, \phi_i)$ and $v_i^m(\phi_i)$ to denote the firm's objective function and the corresponding optimal value function in period i under the myopic setting, respectively.

2.4.2 Updating Only Demand Parameter θ

In this subsection, we consider the substitution probability α to be known and we only need to estimate the demand parameter θ . Doing this allows us to study the driving forces to have a better estimation of the demand parameter. We are particularly interested in examining whether the inventory level shall be kept higher than the myopic one in order to better learn demand information. Since α is given, $\phi_i(\theta, \alpha)$, the prior joint distribution for period

i ($i = 1, \dots, N$), reduces to a one-variable distribution. Let $\phi'_i(\theta)$ denote the prior marginal distribution of θ for period i , where $\phi'_i(\theta) = \frac{\phi_i(\theta, \alpha)}{\int_{\Theta} \phi_i(\theta', \alpha) d\theta'}$.

Below, we first consider the \mathcal{OS} scenario where we have complete observations. For period i ($i = 1, \dots, N$), the impact of increasing the inventory level of the discounted seats y by one unit satisfies

$$\begin{aligned} G_i^{\mathcal{OS}}(y+1, \phi'_i) - G_i^{\mathcal{OS}}(y, \phi'_i) &= E_{\phi'_i(\theta)} \{ \pi(y+1|\theta, \alpha) - \pi(y|\theta, \alpha) \} \\ &= G_i^m(y+1, \phi'_i) - G_i^m(y, \phi'_i). \end{aligned} \quad (2.4)$$

Equation (2.4) implies that the marginal impact of increasing the inventory level y on the objective function under the Bayesian inventory management scheme remains the same as that under the myopic decision scheme. It then follows that the Bayesian optimal inventory level shall be the same as the myopic one, which is formally stated in the following proposition.

Proposition 2.2. *When the substitution probability α is known, for any period i ($i = 1, \dots, N$), given the same prior distribution $\phi'_i(\theta)$, the Bayesian optimal inventory level under the \mathcal{OS} scenario is equal to the corresponding myopic one; that is, $y_i^{\mathcal{OS}} = y_i^m$.*

The underlying reason is that when we have complete observations, there is no need to manipulate the inventory level to observe more demand information. Hence, the decision maker only needs to maximize the current-period expected profit. Such an equality between $y_i^{\mathcal{OS}}$ and y_i^m serves as a benchmark for the following comparisons in other scenarios.

Next, we turn to the \mathcal{OT} scenario where the substitution demand is tangled with the primary demand for the regular-price seat. One may believe that the information is complete to estimate the demand parameter. However, there exists tangled demand that does not provide the complete observation of D_2 , the primary demand for the regular-price seat. Consider the marginal impact of increasing the inventory level y . For period i ($i = 1, \dots, N-1$), we can show that

$$\begin{aligned} G_i^{\mathcal{OT}}(y+1, \phi'_i) - G_i^{\mathcal{OT}}(y, \phi'_i) &= E_{\phi'_i(\theta)} \left\{ \pi(y+1|\theta, \alpha) - \pi(y|\theta, \alpha) \right. \\ &\quad \left. + \delta \left[\sum_{x_1} \sum_{x_2} v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^{y+1}(x_1, x_2|\theta, \alpha) - \sum_{x_1} \sum_{x_2} v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^y(x_1, x_2|\theta, \alpha) \right] \right\}. \end{aligned} \quad (2.5)$$

To investigate the relationship between the Bayesian optimal inventory level $y_i^{\mathcal{OT}}$ and the

myopic one y_i^m , we need to consider the following term stated in (2.5):

$$E_{\phi'_i(\theta)} \left\{ \sum_{x_1} \sum_{x_2} v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^{y+1}(x_1, x_2 | \theta, \alpha) - \sum_{x_1} \sum_{x_2} v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^y(x_1, x_2 | \theta, \alpha) \right\}.$$

Lemma 2.1. *When the substitution probability α is known, under the \mathcal{OT} scenario, given any prior distribution $\phi_i(\theta)$ for period i ($i = 1, \dots, N - 1$) with $0 < y < M$, we have*

$$E_{\phi_i(\theta)} \left\{ \sum_{x_1} \sum_{x_2} v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^{y+1}(x_1, x_2 | \theta, \alpha) \right\} \geq E_{\phi_i(\theta)} \left\{ \sum_{x_1} \sum_{x_2} v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^y(x_1, x_2 | \theta, \alpha) \right\}.$$

Lemma 2.1 implies that increasing the discounted seat's inventory level y in a period yields a larger total discounted expected profit for the following periods. From Lemma 2.1, we can further obtain that

$$G_i^{\mathcal{OT}}(y + 1, \phi'_i) - G_i^{\mathcal{OT}}(y, \phi'_i) \geq G_i^m(y + 1, \phi'_i) - G_i^m(y, \phi'_i). \quad (2.6)$$

The inequality (2.6) allows us to obtain the following “stock more” result.

Proposition 2.3. *When the substitution probability α is known, for any period i ($i = 1, \dots, N$), given the same prior distribution $\phi'_i(\theta)$, the Bayesian optimal inventory level under the \mathcal{OT} scenario is no less than the corresponding myopic one; that is, $y_i^{\mathcal{OT}} \geq y_i^m$.*

In inventory management literature, “stock more” is mainly driven by the lack of observability of lost sales (Lariviere and Porteus, 1999; Ding et al., 2002; and Chen and Plambeck, 2008). Here, Proposition 2.3 provides another driving force for the “stock more” result: *reducing demand substitution so as to better observe the primary demand*. With a higher inventory level of the discounted seat, the chance of buy-up substitution is less, which results in a better observation of the primary demand for the regular-price seat.

We now consider the \mathcal{US} scenario where lost sales are unobservable but the substitution demand can be separated out in the sales data. Is the unobservable lost sales still a driving force to “stock more” under our model setting? To answer this question, let us first consider a special scenario $\widetilde{\mathcal{US}}$ where we assume that the two primary demands D_1 and D_2 are independent and no unsatisfied demand chooses to substitute (i.e., $\alpha = 0$). In this scenario, $\phi'_i(\theta)$, the prior marginal distribution of θ for period i , can be further specified as $\phi'_{i,1}(\theta_1)\phi'_{i,2}(\theta_2)$ with $\theta_i \in \Theta_i$ ($i = 1, 2$), where $\phi'_{i,1}(\theta_1)$ and $\phi'_{i,2}(\theta_2)$ are the marginal distributions of D_1 and D_2 , respectively. Since the number of substitution demands s_{21} is always 0, the likelihood of

observing sales quantities (s_1, s_{22}) can be expressed as

$$f_{\widetilde{\mathcal{US}}}^y(s_1, s_{22}|\theta_1, \theta_2) = \begin{cases} f_1(s_1|\theta_1)f_2(s_{22}|\theta_2), & \text{if } s_1 < y \text{ and } s_{22} < M - s_1; \\ f_1(s_1|\theta_1) \sum_{j=M-s_1}^{+\infty} f_2(j|\theta_2), & \text{if } s_1 < y \text{ and } s_{22} = M - s_1; \\ \sum_{i=y}^{+\infty} f_1(i|\theta_1)f_2(s_{22}|\theta_2), & \text{if } s_1 = y \text{ and } s_{22} < M - y; \\ \sum_{i=y}^{+\infty} f_1(i|\theta_1) \sum_{j=M-y}^{+\infty} f_2(j|\theta_2), & \text{if } s_1 = y \text{ and } s_{22} = M - y. \end{cases}$$

It is easy to verify that the posterior distributions of D_1 and D_2 are still independent. For the $\widetilde{\mathcal{US}}$ scenario, we obtain the following result.

Proposition 2.4. *In the $\widetilde{\mathcal{US}}$ scenario, when the demand parameter θ_1 (resp. θ_2) is unknown but θ_2 (resp. θ_1) is known, for any period i ($i = 1, \dots, N$), given the same prior distribution $\phi'_{i,1}(\theta_1)$ (resp. $\phi'_{i,2}(\theta_2)$), the Bayesian optimal inventory level under the $\widetilde{\mathcal{US}}$ scenario is no less (resp. no larger) than the corresponding myopic one; that is, $y_i^{\widetilde{\mathcal{US}}} \geq y_i^m$ (resp. $y_i^{\widetilde{\mathcal{US}}} \leq y_i^m$).*

According to Proposition 2.4, the unobservable lost sale is still a “stock more” driving force considering either kind of the primary demands under the $\widetilde{\mathcal{US}}$ scenario. Intuitively, when only one kind of primary demands needs to be learned but another kind is already known, we should stock more seats for the unknown demand to better estimate it. As a change in the decision variable y affects the observations on two primary demands D_1 and D_2 , the “stock more” result here can be regarded as an extension of the traditional “stock more” result driven by the unobservable lost sales when there is only one kind of demand (Lariviere and Porteus, 1999; Ding et al., 2002; and Chen and Plambeck, 2008).

The result in Proposition 2.4 is obtained through assuming that the substitution probability is 0. Will the unobservable lost sale still be a “stock more” driving force when $\alpha > 0$ under our setting? Surprisingly, the answer is no; see the following “stock less” result under a special case.

Proposition 2.5. *Consider that the primary demand for the regular-price seat D_2 is always zero and the substitution probability are known with $\alpha > 0$. Only the demand parameter of D_1 , θ_1 , is unknown. The value of D_1 cannot be 1; i.e., $f_1(1|\theta_1) = 0$ for all $\theta_1 \in \Theta_1$. The total number of available seats is $M = 2$. Then, for any period i ($i = 1, \dots, N$), given the*

same prior distribution $\phi'_{i,1}(\theta_1)$, the Bayesian optimal inventory level is no larger than the corresponding myopic one; that is, $y_i^{\mu S} \leq y_i^m$.

In Proposition 2.5, no matter whether the inventory decision y is set to be 1 or 2, given the same realized early-bird demand, the observed early-bird sales amount conveys the same demand information. For example, if the realized demand is 0, the sales amount is 0 in both $y = 1$ and $y = 2$ cases. If the realized demand is 2 or larger, then the sales amount is $s_1 = 1$ in the case $y = 1$ and $s_1 = 2$ in the case $y = 2$. However, both $s_1 = 1$ and $s_1 = 2$ convey the same demand information $D_1 \geq 2$ because the demand cannot be 1. Hence, setting the different values of y here does not affect the demand information gaining in the first phase. However, the inventory manager can infer some demand information from the substitution demand when $y = 1$. In other words, “stocking less” discounted seats can induce more information about the early-bird demand, in sharp contrast to the literature result.

The next example shows that the strict less ‘<’ in Proposition 2.5 can be achieved. Moreover, such a “stock less” phenomenon can appear in more general settings with other common demand distributions and larger M .

Example 2.2. (“Stock Less” Driven by Unobservable Lost Sales) Under the US scenario, the following cases show that the “stock less” result can hold when only the demand parameter of D_1 is learned. In all three cases, the inventory manager aims to determine the optimal number of discounted seats to maximize the total expected profit in two periods (with discount factor $\delta = 1$). The optimal inventory levels of the discounted seat in the first period under the Bayesian inventory management and myopic optimization are $y_1^{\mu S}$ and y_1^m , respectively.

(a) **One-point Distribution with $M = 2$**

The discount and regular prices are set at $p_1 = 800$ and $p_2 = 1200$, respectively. The primary demand for the discounted seat D_1 follows a one-point distribution with an unknown parameter θ , and the primary demand for the regular-price seats D_2 is known as 0. The demand parameter θ takes value 1 or 2: when $\theta = 1$, $D_1 = 2$; otherwise, $D_1 = 3$. The buy-up substitution probability α is known as 0.5. At the beginning of the first selling period, when the inventory manager holds the prior beliefs $\Pr(\theta = 1) = 0.35$ and $\Pr(\theta = 2) = 0.65$, the optimal inventory levels for this period satisfy that $(y_1^{\mu S} = 1) < (y_1^m = 2)$.

(b) **Two-point Distribution with $M = 150$**

The discount and regular prices are set at $p_1 = 700$ and $p_2 = 1200$, respectively. The primary demand for the discounted seat D_1 follows a two-point distribution with an unknown parameter θ , and the primary demand for the regular-price seats is known as $D_2 = 0$. The demand parameter θ takes value 1 or 2: when $\theta = 1$, $Pr(D_1 = 150) = Pr(D_1 = 180) = \frac{1}{2}$; otherwise, $Pr(D_1 = 200) = Pr(D_1 = 230) = \frac{1}{2}$. The buy-up substitution probability α is known as 0.1. At the beginning of the first selling period, when the inventory manager holds the prior beliefs $Pr(\theta = 1) = 0.4$ and $Pr(\theta = 2) = 0.6$, the optimal inventory levels for this period satisfy that $(y_1^{\mathcal{US}} = 147) < (y_1^m = 148)$.

(c) **Truncated Poisson Distribution with $M = 100$**

The discount and regular prices are set as $p_1 = 700$ and $p_2 = 1200$, respectively. The primary demand for the discounted seat D_1 follows a truncated Poisson distribution ($0 \leq D_1 \leq 300$) with an unknown parameter λ_1 , and the primary demand for the regular-price seats D_2 also follows a truncated Poisson distribution ($0 \leq D_2 \leq 30$) whose parameter is known as 5. The unknown demand parameter λ_1 takes value 150 or 256. The buy-up substitution probability α is known as 0.2. At the beginning of the first selling period, when the inventory manager holds the prior beliefs $Pr(\lambda_1 = 150) = 0.8$ and $Pr(\lambda_1 = 256) = 0.2$, the optimal inventory levels for this period satisfy that $(y_1^{\mathcal{US}} = 77) < (y_1^m = 79)$.

Intuitively, besides the sales amount of the discounted seats s_1 , the observed substitution demand s_{21} can convey some information about the primary demand for the discounted seat D_1 , which yields a “stock less” driving force. This is different from the conclusion in Chen and Plambeck (2008), in which they assume that the substitute product is always available. Thus, in their paper, complete observations on the substitution demand can be achieved, and unobservable lost sales lead to “stock more”. Here, observations on the substitution demand are limited due to the total inventory constraint. “Stocking less” discounted inventory induces more observations on substitutions. Hence, the “stock more” result no longer holds.

The above analysis implies that the unobservable lost sale contains counter “stock more” and “stock less” driving forces, which makes the relationship between $y_i^{\mathcal{US}}$ and y_i^m generally uncertain. For the \mathcal{UT} scenario where both unobservable lost sales and tangled demands exist, there is an additional “stock more” driving force due to the tangled demands compared with the \mathcal{US} scenario, which brings more uncertainty to the final comparison result.

2.4.3 Updating Both Demand Parameter θ and Substitution Probability α

In this subsection, we consider that both the demand parameter θ and substitution probability α are unknown.

Let us still first investigate the simplest \mathcal{OS} scenario. This scenario actually provides complete observations to estimate the demand parameter: lost sales are observable and the substitution demand can be separated from the primary demand. However, due to the need of estimating the substitution probability, inventory levels shall be manipulated so that we can obtain more observations on the customer's substitution behavior. For this scenario, we can obtain the following comparison result regarding the marginal impact of increasing the inventory level y between the Bayesian-inventory-management setting and the myopic setting, where the proof of inequality (2.7) can be found in the online Appendix A.2:

$$G_i^{\mathcal{OS}}(y+1, \phi_i) - G_i^{\mathcal{OS}}(y, \phi_i) \leq G_i^m(y+1, \phi_i) - G_i^m(y, \phi_i), \quad (2.7)$$

Based on (2.7), we can obtain the following result.

Proposition 2.6. *For any period i ($i = 1, \dots, N$), given the same prior $\phi_i(\theta, \alpha)$, the Bayesian optimal inventory level under the \mathcal{OS} scenario is no larger than the corresponding myopic one; that is, $y_i^{\mathcal{OS}} \leq y_i^m$.*

In comparison with the setting where only the demand parameter θ needs estimation, Proposition 2.6 shows that introducing a new parameter, the substitution probability, generates a driving force for “stock less”. Note that under the \mathcal{OS} scenario, although the inventory level of the discounted seats y is reduced to allow more substitution demands, the observations on the primary demands are unaffected. Such a “stock less” result is similar to that obtained in Chen and Plambeck (2008). Yet, our result generalizes from their setting with unlimited capacity of substitutes and sole source of demands for the substitutes to a setting with limited capacity of substitutes and two sources of demands for the substitutes.

For the other three information scenarios, those driving forces identified in Section 2.4.2 still exist besides the “stock less” result driven by learning the substitution probability. This brings more uncertainty to the aggregate effect on the optimal inventory level. The final comparison result generally depends on factors such as the Bayesian inventory manager's prior beliefs and seats' prices (see Section 2.5).

2.4.4 Comparison of Expected Profits under Four Information Scenarios

We now compare the system performance under the four information scenarios. The following proposition lists the comparison results, which help us to understand the value of information.

Proposition 2.7. *For period i ($i = 1, \dots, N$), given the same prior distribution ϕ_i ,*

- (a) *the objective functions under the four information scenarios satisfy the following relationship: $\sum_{n=0}^{N-i} \delta^n \cdot \pi(y, \phi_i) \leq G_i^{\mathcal{UT}}(y, \phi_i) \leq G_i^{\mathcal{OT}}(y, \phi_i) (G_i^{\mathcal{US}}(y, \phi_i), \text{ resp.}) \leq G_i^{\mathcal{OS}}(y, \phi_i)$;*
- (b) *the optimality value functions under the four information scenarios satisfy the following relationship: $\sum_{n=0}^{N-i} \delta^n \cdot \max_{0 < y \leq M} \pi(y, \phi_i) \leq v_i^{\mathcal{UT}}(\phi_i) \leq v_i^{\mathcal{OT}}(\phi_i) (v_i^{\mathcal{US}}(\phi_i), \text{ resp.}) \leq v_i^{\mathcal{OS}}(\phi_i)$.*

In Proposition 2.7, the term $\sum_{n=0}^{N-i} \delta^n \cdot \pi(y, \phi_i)$ represents the discounted profit without any learning. This lemma shows that an inventory system with Bayesian learning can always achieve a profit at least as high as the one without any learning. As shown in Proposition 2.7, among the four information scenarios, the most informative \mathcal{OS} scenario generates the highest profit while the least informative \mathcal{UT} scenario generates the lowest profit. The performances of the \mathcal{OT} and \mathcal{US} scenarios lie in between and cannot be simply compared: the former one has censored observations on the primary demand due to demand tangling while the latter one has unobservable lost sales.

2.5 Numeric Study

In this section, we conduct the numerical experiments to analyze our Bayesian inventory management in the airline setting. We examine the system parameters' impact on the “stock more” and “stock less” results by varying the parameter values including the Bayesian inventory manager's prior beliefs and seats' prices under a two-period setting. In particular, we consider the \mathcal{US} scenario under a two-period setting, under which we can easily solve the dynamic programming model. Such a setting is also enough for us to identify the key “exploration-exploitation” tradeoff.

An airline company offers an early bird sale of a flight with a medium-sized jet. The total number of seats is $M = 120$. The regular price is set at $p_2 = 1200$. The inventory manager aims to determine the optimal number of discounted seats to maximize the total expected

profit in two periods (with discount factor $\delta = 1$) under the \mathcal{US} scenario. The primary demands for the discounted and regular-price seats are D_1 and D_2 following truncated Poisson distributions ($0 \leq D_1 \leq 300$, $0 \leq D_2 \leq 100$) with parameters λ_1 and λ_2 , respectively. The two parameters λ_1 and λ_2 are correlated, and their relationship is indicated by a parameter θ , which takes value 1, 2 or 3. When $\theta = 1$, $\lambda_1 = 160$ and $\lambda_2 = 5$; when $\theta = 2$, $\lambda_1 = 160$ and $\lambda_2 = 20$; and when $\theta = 3$, $\lambda_1 = 270$ and $\lambda_2 = 5$. The buy-up substitution probability α may be low ($\alpha = 0.2$) or high ($\alpha = 0.7$). At the beginning of the first selling period, the inventory manager holds the prior beliefs $\tilde{u} = (Pr(\theta = 1), Pr(\theta = 2), Pr(\theta = 3))$ and $\tilde{w} = (Pr(\alpha = 0.2), Pr(\alpha = 0.7))$. Let the optimal inventory levels of the discounted seat in the first period under the Bayesian inventory management and myopic optimization be $y_1^{\mathcal{US}}$ and y_1^m , respectively.

We first fix the early-bird-discount price at $p_1 = 700$, and then investigate how different prior beliefs affect the “stock more” and “stock less” results.

Case (i). *Unknown D_1 with $\tilde{u} = (u_1, 0, 1 - u_1)$ ($0 \leq u_1 \leq 1$) and $\tilde{w} = (1, 0)$.*

In this case, the demand parameter of D_2 is known as $\lambda_2 = 5$ and the buy-up substitution probability $\alpha = 0.2$. Only the primary demand for the discounted seat D_1 is unknown. Note that when we vary the prior belief \tilde{u} , a larger u_1 leads to a lower expectation of D_1 . Table 2.4(i) summarizes the changes of the optimal inventory levels and the corresponding expected profits under the Bayesian inventory management and the myopic optimization as u_1 increases. We can see that under the myopic optimization, the optimal inventory level y_1^m increases, but the corresponding total expected profit decreases as u_1 becomes larger. According to Table 2.4(i), when $u_1 = 0$, the inventory manager is sure that the demand parameter of D_1 is $\lambda_1 = 270$, and the two optimal inventory levels $y_1^{\mathcal{US}}$ and y_1^m are equal. As u_1 becomes 0.1, the demand information on D_1 becomes uncertain, and the manager raises the inventory level of the discounted seats to better observe D_1 . However, when u_1 takes larger values (i.e., 0.2 to 0.5), the myopic optimal inventory level y_1^m also becomes larger such that the benefit brought by “stock more” is less than the loss caused by it, and thus the manager should not manipulate the inventory level. After u_1 further increases to 0.6, y_1^m becomes large enough. In this case, reducing discounted seats induces more substitution trials, from which the manager can infer more demand information, leading to the “stock less” result. And under larger u_1 (i.e., 0.7 to 0.9), such a “stock less” result still holds. Especially, the gap between $y_1^{\mathcal{US}}$ and y_1^m reaches the maximum at $u_1 = 0.8$, where y_1^m is very large and the uncertainty on D_1 is still relatively high. But as u_1 increases to 0.9, the uncertainty on D_1 weakens, and so does the incentive to

Table 2.4: The Impacts of System Parameters on Optimal Inventory Levels ($y_1^{\mathcal{MS}}$ and y_1^m) and Corresponding Expected Profits ($v_1^{\mathcal{MS}}(\phi_1)$ and $v_1^m(\phi_1) + v_2^m(\phi_2)$) under Bayesian Inventory Management and Myopic Optimization

(i) Unknown D_1 : $\tilde{u} = (u_1, 0, 1 - u_1)$ and $\tilde{w} = (1, 0)$					
u_1	$y_1^{\mathcal{MS}}$	y_1^m	$y_1^{\mathcal{MS}} - y_1^m$	$v_1^{\mathcal{MS}}$ ($\times 10^5$)	$v_1^m + v_2^m$ ($\times 10^5$)
0	77	77	0	2.0505	2.0505
0.1	78	77	+1	2.0141	2.0140
0.2	79	79	0	1.9790	1.9790
0.3	80	80	0	1.9450	1.9450
0.4	84	84	0	1.9128	1.9128
0.5	90	90	0	1.8840	1.8840
0.6	97	98	-1	1.8622	1.8618
0.7	99	100	-1	1.8435	1.8428
0.8	99	102	-3	1.8254	1.8241
0.9	101	103	-2	1.8088	1.8066
1	104	104	0	1.7995	1.7995

(ii) Unknown D_2 : $\tilde{u} = (u_2, 1 - u_2, 0)$ and $\tilde{w} = (1, 0)$					
u_2	$y_1^{\mathcal{MS}}$	y_1^m	$y_1^{\mathcal{MS}} - y_1^m$	$v_1^{\mathcal{MS}}$ ($\times 10^5$)	$v_1^m + v_2^m$ ($\times 10^5$)
0	84	84	0	1.9710	1.9710
0.1	86	86	0	1.9470	1.9470
0.2	87	87	0	1.9242	1.9242
0.3	88	89	-1	1.9019	1.9018
0.4	90	91	-1	1.8809	1.8808
0.5	93	95	-2	1.8612	1.8606
0.6	95	98	-3	1.8435	1.8423
0.7	98	100	-2	1.8278	1.8268
0.8	100	102	-2	1.8151	1.8144
0.9	102	103	-1	1.8057	1.8054
1	104	104	0	1.7995	1.7995

(iii) Unknown α : $\tilde{u} = (1, 0, 0)$ and $\tilde{w} = (w, 1 - w)$					
w	$y_1^{\mathcal{MS}}$	y_1^m	$y_1^{\mathcal{MS}} - y_1^m$	$v_1^{\mathcal{MS}}$ ($\times 10^5$)	$v_1^m + v_2^m$ ($\times 10^5$)
0	1	1	0	2.7305	2.7305
0.1	1	1	0	2.5923	2.5923
0.2	1	1	0	2.4541	2.4541
0.3	1	1	0	2.3159	2.3159
0.4	8	8	0	2.1785	2.1785
0.5	52	52	0	2.0611	2.0611
0.6	95	98	-3	1.9960	1.9928
0.7	96	100	-4	1.9412	1.9344
0.8	97	102	-5	1.8876	1.8700
0.9	97	103	-6	1.8333	1.8097
1	104	104	0	1.7995	1.7995

(iv) Varying p_1 : $\tilde{u} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ and $\tilde{w} = (\frac{1}{2}, \frac{1}{2})$					
p_1	$y_1^{\mathcal{MS}}$	y_1^m	$y_1^{\mathcal{MS}} - y_1^m$	$v_1^{\mathcal{MS}}$ ($\times 10^5$)	$v_1^m + v_2^m$ ($\times 10^5$)
600	1	1	0	2.1445	2.1445
650	1	1	0	2.1670	2.1670
700	25	29	-4	2.1933	2.1932
750	60	69	-9	2.2390	2.2344
800	65	75	-10	2.2920	2.2868
850	71	79	-8	2.3501	2.3383
900	72	83	-11	2.4092	2.3853
950	72	86	-14	2.4689	2.4358
1000	73	91	-18	2.5290	2.5016
1050	74	97	-23	2.5900	2.5874
1100	98	103	-5	2.6788	2.6762
1150	102	107	-5	2.7737	2.7714
1200	120	120	0	2.8800	2.8800

stock less, shortening the gap between $y_1^{\mathcal{US}}$ and y_1^m . Finally, such a gap disappears when the manager is sure that $u_1 = 1$.

Case (ii). *Unknown D_2 with $\tilde{u} = (u_2, 1 - u_2, 0)$ ($0 \leq u_2 \leq 1$) and $\tilde{w} = (1, 0)$.*

In this case, the demand parameter of D_1 is known as $\lambda_1 = 160$ and the buy-up substitution probability $\alpha = 0.2$. Only the demand parameter of D_2 is unknown. As we vary the prior belief \tilde{u} by increasing u_2 , the expectation of D_2 decreases, making the optimal inventory level y_1^m under the myopic policy increase and the corresponding total expected profit decrease; see Table 2.4(ii). On the whole range of u_2 , learning D_2 drives the “stock less” result whereby the Bayesian inventory manager can better observe it. Further, according to Table 2.4(ii), when the demand uncertainty regarding D_2 is relatively low (i.e., $u_2 \leq 0.2$), $y_1^{\mathcal{US}}$ is equal to y_1^m . As the uncertainty becomes higher (i.e., $u_2 > 0.2$), the extra demand information gained from lowering the inventory level of discounted seats brings more profit to the manager, driving him to stock less. Specifically, as u_2 increases from 0.3 to 0.6 and then to 1, the gap between $y_1^{\mathcal{US}}$ and y_1^m first increases, then reaches the maximum at $u_2 = 0.6$, and finally decreases to 0.

Case (iii). *Unknown α with $\tilde{u} = (1, 0, 0)$ and $\tilde{w} = (w, 1 - w)$ ($0 \leq w \leq 1$).*

In this case, the demand parameters of both D_1 and D_2 are known as $\lambda_1 = 160$ and $\lambda_2 = 5$ but the buy-up substitution probability is unknown. We then vary the prior belief on the substitution probability α by changing the parameter value of w to investigate its impact on the “stock more” and “stock less” results. In accord with Proposition 2.1, as α becomes smaller (i.e., with a larger w), the optimal inventory level y_1^m under the myopic policy increases and the corresponding total expected profit decreases as shown in Table 2.4(iii). In Section 2.4.3, we show that under the \mathcal{OS} scenario, learning α yields a “stock less” driving force to generate more substitution trials. In this example, under the \mathcal{US} scenario, “stock less” is still the dominant driving force. In Table 2.4(iii), when $w \leq 0.5$, y_1^m is very small, and thus it is not good to stock less for very limited extra substitution observations at the cost of a big revenue loss. Things become different when w is larger than 0.5. In these cases, y_1^m becomes relatively large, and now “stocking less” for more substitution observations dominates the final comparison result, yielding $y_1^{\mathcal{US}} \leq y_1^m$. Especially, as w increases from 0.6 to 0.9, the gap between $y_1^{\mathcal{US}}$ and y_1^m becomes larger. This implies that the biggest difference between $y_1^{\mathcal{US}}$ and y_1^m does not necessarily appear when the uncertainty is highest. Indeed, besides the uncertainty in belief, such a difference is affected by the magnitudes of expected profits and inventory levels. In this example, at $w = 0.9$, although the uncertainty on the substitution probability is not that large, “stocking less” still brings much benefit and causes little revenue

loss since the myopic optimal inventory level y_1^m is very close to the total amount of seats M .

Decomposing the prior belief into the above three cases enables us to have a better picture of what driving forces the learning on each kind of demands and the buy-up substitution probability can generate. Besides the prior beliefs, the difference between the two prices is another important factor that affects the comparison result between the Bayesian optimal inventory level and the corresponding myopic one. Below, we investigate the effect of the price difference between p_1 and p_2 .

Case (iv). *Varying p_1 ($600 \leq p_1 \leq 1200$) with $\tilde{u} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ and $\tilde{w} = (\frac{1}{2}, \frac{1}{2})$.*

In this case, we fix both the regular price $p_2 = 1200$ and the prior beliefs about the demands and the buy-up substitution probability \tilde{u} and \tilde{w} . We then vary the early-bird-discount price p_1 to investigate its impact on the optimal inventory level decisions. As p_1 increases, both the optimal inventory level y_1^m under the myopic policy and the corresponding total expected profit increase. From Table 2.4(iv), we can see that the change of p_1 complicates the comparison result. When the early-bird-discount price p_1 is very low (i.e., $p_1 \leq 650$), y_1^m takes the smallest value 1. Under this situation, the value of gaining more demand information is less than the potential revenue loss caused by stocking more discounted seats than the myopic ones due to the extremely low discounted price. Hence, $y_1^{\mu S} = y_1^m$. As p_1 becomes larger (i.e., $700 \leq p_1 \leq 1050$), the inventory level of the discounted seats under the myopic policy becomes relatively high, “stocking less” can help to gain some information on demands and the buy-up substitution probability, which yields more value than the potential revenue loss induced by it due to the moderate early-bird-discount price p_1 . Thus, “stock less” occurs. Note that as p_1 increases, the difference between $y_1^{\mu S}$ and y_1^m does not necessarily show monotonicity because it depends jointly on the uncertainty in belief and the magnitudes of expected profits and inventory levels. As p_1 further increases (i.e., $p_1 \geq 1100$), the price difference between the early-bird-discount price p_1 and the regular price p_2 reduces. Now, “stocking less” is still the dominant driving force, but the manager should reduce the gap between $y_1^{\mu S}$ and y_1^m to avoid much revenue loss due to the high p_1 .

The above four cases systematically show how various system parameters, i.e., the prior beliefs and prices, separately affect the results regarding “stock more” or “stock less”. For a general problem, one should take all possible driving forces into consideration, and try to identify those dominant ones to design heuristic algorithms for practical use.

2.6 Conclusions and Suggestions for Future Research

In this study, we investigate an airline company’s optimal seat allocation decision when it provides the early-bird booking discount under a multi-period setting. The number of total seats is fixed and within each period, the company offers the price discount in the early-bird-discount phase and charges the full price in the regular-price phase. Setting a proper inventory level for the early-bird-discount phase is critical for the company’s revenue management. Too few seats reserved for the early bird discount, on one hand, could lead to the loss of some potential sales, but on the other hand, could achieve the goal of forcing some unsatisfied customers in the early-bird-discount phase to purchase the regular-price seats as substitutes.

An optimal decision on the inventory level requires the decision maker to have knowledge of the primary demands for the seats under the two prices and the buy-up substitution probability of those unsatisfied early-bird customers purchasing the regular-price seats. In this chapter, we consider a dynamic inventory management model with Bayesian learning on both the demand parameter and the buy-up substitution probability. We adopt the Bayesian inventory management to dynamically learn and optimize. Specifically, we examine four information scenarios based on whether or not lost sales are observable and the substitution demand from those unsatisfied early-bird customers can be separated from the primary demand for the regular-price seat. Under each scenario, we compare the Bayesian optimal inventory level of the discounted seat with the corresponding myopic one. Such comparison results can provide insights to the heuristic algorithms considering the “exploration-exploitation” tradeoff.

When the buy-up substitution probability is known and only the demand parameter needs to be learned, we identify a new driving force for the classical “stock more” result: demand tangling between the substitution demand and the primary demand. When lost sales are observable but the substitution demand cannot be separated from the primary demand, we should increase the inventory level of the discounted seat in order to reduce the likelihood of substitution so that we can better learn the primary demand. In the literature on Bayesian inventory management, such as Lariviere and Porteus (1999) and Ding et al. (2002), the unobservable lost sale is regarded as a driving force of “stock more”. This is also identified by Chen and Plambeck (2008) in which they consider the stock-out-based substitution and the amount of available substitutes is infinite. Here, we show that the unobservable lost sale still is a driving force of “stock more” to better learn each kind of demands when there is no buy-up substitution. However, if the substitution probability is non-zero, the unobservable lost sale may now become a driving force of “stock less”. The behind reason is that in our setting, the

capacity of substitutes is limited; when very limited demand information is contained in the sales amounts, one can further infer some information from the observed substitutions, which induces the manager to “stock less” discounted seats to yield more substitution trials. When we need to update both the demand parameter and the buy-up substitution probability, we find that there exists a driving force for “stock less” as in Chen and Plambeck (2008): reducing the inventory level of the discounted seat can induce more substitution events to occur, benefiting the learning on the substitution probability. The interplay of the aforementioned multiple driving forces determines the final inventory outcome regarding whether to “stock more” or “stock less”. The numeric study shows that the final result relies on factors such as Bayesian manager’s prior beliefs and seats’ prices, and is determined by the strongest driving force.

To gain deeper understanding of the two-phase early bird policy, we further extend the model with a fixed total inventory level to a setting where both the total inventory level in two phases and the inventory level of the discounted product are decision variables, and analyze it in the online Appendix A.1. We find that the above driving forces still exist. Also, when lost sales are unobservable, the “stock more” result still may not hold. Instead, we may need to stock less one kind of the product as a result of the tradeoff between the revenue that can be improved by utilizing the demand information and the corresponding overstocking risk. We then derive the upper bound on the Bayesian optimal inventory levels for this two-decision-variable problem, which enriches the Bayesian inventory management literature.

Overall, our two-fare model has identified the main “exploration-exploitation” tradeoffs under the early bird policy. A direct extension of this study is to develop efficient heuristic algorithms for multi-fare and multi-flight airline problems based on those driving forces. It is also worthwhile to apply Bayesian inventory management to the markdown inventory management problem (see, e.g., Hu et al., 2015), where new tradeoffs may appear. Although the markdown model seems quite similar to our early-bird model, the related analysis will be much more complex and challenging due to the quasiconvex single-period profit function and inter-correlated demands over the two adjacent periods. We would like to leave them to future research.

Chapter 3

Signaling Service Quality via Queue Disclosure

In Chapter 2, there is only one decision maker who tries to learn the unknown information for decision making. From this chapter, we extend the setting to multiple players, where the information senders own private information on their types, and try to signal such information to the receivers. Specifically, we consider the queueing setting. A server knows his quality type, but the customers do not. The server chooses the optimal queue disclosure strategy to attract customers to join the queue as many as possible. From the server's action, the customers infer some information on the quality type, and then decide the joining strategy.

In this chapter, we consider that the server decides the queue disclosure strategy after he realizes the quality type. This leads to a signaling game. And we mainly investigate whether the server's queue-disclosure action (i.e., revealing or concealing) itself can be a quality signal, and if so, what impacts it can exert. In the next chapter, we will consider that the server can design and commit to a queue disclosure strategy before the quality type is realized. This different timing yields a totally different Bayesian persuasion game. And we will mainly investigate the underlying mechanism and find out the optimal queue disclosure strategy.

3.1 Introduction

In many service systems, the service quality of a server (he) is unknown to some customers (she). Take the restaurant as an example. The food quality depends on the chefs' skills and the sources of the ingredients, and the nearby residents are likely to know the food quality

while outsiders might not. As another example, the service quality of online consulting or call centers relies on the professional knowledge and seniority of the consultant or operator. Those regular customers can be regarded as informed of the service quality, while others are not. Upon customers' arrivals, a server may decide to reveal the queue length to the customers or not before they make the joining-or-balking decisions. For example, when a customer enters a restaurant, she should order the meal or reserve a seat first, and then waits for the service. Before the customer places her order, the server can display the number of waiting customers on the display screens or queueing machines (take KFC and McDonald's as examples). If so, the customer is informed of the queue length, based on which she decides whether to join the queue (i.e., place an order or make a reservation) or not. Otherwise, the queue is unobservable to her. Similarly, for the online queueing setting, the server can achieve revealing (resp. concealing) the queue length through informing (resp. not informing) an incoming customer of the waiting line through the softwares, telephones, or mobile apps, etc.

The values of the services mentioned above cannot be easily communicated. Debo et al. (2012) point out that uninformed customers can infer the server's quality by inspecting the queue length and then follow a 'hole-avoiding' strategy: they behave almost the same as the customers informed of high quality, except that at a certain queue length, called the hole, they do not join. Based on the pioneering work of Debo et al. (2012), one can further ask the following questions: Why is a server willing to reveal the queue length to customers? Isn't the visibility of the queue length itself a signal indicating the service quality? In this chapter, we aim to investigate these questions by considering the visibility of queue length as a signal of service quality. Correspondingly, in our setting, when the queue is observable, an uninformed customer can infer some quality information by inspecting not only the queue length but also the server's queue-disclosure action itself.

Consider the following signaling game setting between a server and customers. Customers arrive to a single-server queueing system according to a Poisson process. Their service times follow an exponential distribution. The server's service quality is either high or low, which is determined by nature via a Bernoulli trial. A fraction of customers are informed ones who know exactly the server's quality level. They can be either *positively informed* if the server is of high quality or *negatively informed* if the server is of low quality. Uninformed customers, however, hold some prior belief on the service quality. Customers are otherwise identical, i.e., they receive the same service reward and bear the same per-unit-time delay cost. The server

knows his own quality and can fine-tune the queue disclosure action (revealing or concealing his queue length) to signal the quality type to maximize the system demand, that is, the effective arrival rate. Upon observing the server's queue disclosure action and the actual queue length of an observable queue, customers update their beliefs (if uninformed), and then decide whether to join the queue or balk.

To solve this signaling game, we investigate the *sequential equilibrium* (Kreps and Wilson, 1982), and use the *perfect sequential equilibrium* (Grossman and Perry, 1986) as the refinement criterion. Three types of sequential equilibria are considered: the *pure strategy* in which both high- and low-quality types of the server choose to either always reveal or always conceal the queue, the *mixed strategy* in which both types of the server randomize between revealing and concealing the queue, and the *hybrid strategy* in which one type of the server chooses to either always reveal or always conceal the queue and the other type randomizes between these two actions. Furthermore, for the pure strategy, when both types of the server adopt the same action, we call it a *pooling strategy*; but when their actions are different, we call it a *separating strategy*.

We first consider a basic scenario where all customers are uninformed of the service quality. Hence, the effective arrival rate to an observable or unobservable queue is independent of the quality type as all customers hold the same belief. This property leads to the result that except at some discrete values of the market size (i.e., the potential arrival rate), the unique pure-strategy perfect sequential equilibrium is that both types choose to conceal (resp. reveal) the queue when the market size is below (resp. above) a threshold. We further prove that the pooling perfect sequential equilibria generate larger effective arrival rates to two types of the server compared with the hybrid or mixed sequential equilibria. Under pooling equilibria, the queue disclosure action conveys no quality information, and the effective arrival rates to the high-quality and low-quality servers remain the same as the ones under the *non-signaling case* where the queue-disclosure action is not regarded as a quality signal.

The situation becomes quite different when the customers become heterogeneous, i.e., some of them are informed but others are uninformed. We first identify the heterogeneous customers' equilibrium queueing strategies in both the unobservable and observable queues, and then analyze the sequential equilibria of the whole signaling game. Similarly to the homogeneous-customer scenario, we show that there exist two thresholds with respect to the market size, below the lower one both types of the server choose to always conceal the queue in equilibrium, and above the upper one both choose to always reveal the queue. That is, the

unique sequential equilibrium is still a pooling strategy when the market size is small or large enough. We further show that when the market size falls between these two thresholds, the separating sequential equilibria can be sustained as the equilibrium outcome, and their exact existence ranges are depicted. Intuitively, as the positively informed customers are always more likely to join the queue than the negatively informed ones, the effective arrival rate to the high-quality server in an observable or unobservable queue is larger than the corresponding one to the low-quality server under a fixed belief of uninformed customers. And this may generate different incentives for two server types to reveal or conceal the queue length in a medium-sized market, which makes it possible for separating sequential equilibria to appear.

Under the separating sequential equilibria, the uninformed customers can infer full information on the quality type from the queue-disclosure action, and thus behave in the same way as the informed customers. When the separating sequential equilibria exist, the maximal effective arrival rate to the high-quality (resp. low-quality) server considering all pure-strategy perfect sequential equilibria is larger (resp. smaller) than the maximal one in the non-signaling case. Uninformed customers cause an overcrowded queueing system of the low-quality server, which may hurt the total utility of all customers and make it negative. But after they completely infer the quality type under the separating sequential equilibria, their incentives to join become weaker. This makes the system back to ‘normal’, and thus improves customers’ total utility to be nonnegative. However, we cannot guarantee a similar definite result regarding the high-quality server. On one hand, a larger effective arrival rate to the high-quality server under the separating sequential equilibrium may either reduce or improve the total utility. On the other hand, the high-quality server may adopt different queue-disclosure actions in the signaling and non-signaling cases. Finally, we conduct some sensitivity analyses to examine how the customer type composition (i.e., the proportion of informed customers in the market) and the service price affect the separating sequential equilibria. And we find that the relationships are non-monotone.

Besides the aforementioned equilibrium results of the signaling game, our study also contributes to a better understanding of the customer queueing strategies with unknown service quality in a twofold way. First, we fully characterize customers’ equilibrium queueing strategies with unknown service quality when the queue is *unobservable*. Our result complements the queueing game literatures on unknown service quality, which all consider observable queues (see, e.g., Debo et al., 2012). Second, when the queue disclosure action itself is treated as a quality signaling device, uninformed customers’ ‘hole-avoiding’ queueing strategy obtained

in Debo et al. (2012) shall be rectified for a medium-sized market, under which a separating equilibrium could exist and uninformed customers behave exactly the same as informed customers.

The remainder of this chapter is organized as follows. We review the related literature in Section 3.2. In Section 3.3, we present the signaling game and introduce the definitions of sequential equilibrium and perfect sequential equilibrium. The sequential equilibria of a basic model with all customers uninformed are analyzed in Section 3.4. And the equilibrium analyses under the general scenario with heterogeneous customers are presented in Section 3.5. Furthermore, we conduct the sensitivity analyses to investigate the impacts of customer type composition and service price on the separating sequential equilibria in Section 3.6. Concluding remarks are provided in Section 3.7. All supplementary contents for readers' interest are relegated to Appendixes B.1-B.4, and all the proofs can be found in Appendix B.5.

3.2 Literature Review

Our research is closely related to the literature on signaling games in queueing systems. Allon et al. (2011) considers a *cheap talk* game in which a server sends a queue-length-dependent signal to the customers. Yu et al. (2018) further study a cheap talk game with heterogeneous customers, and they show that the server can infer customer types through customers' reaction toward the server's delay announcement. Veeraraghavan and Debo (2009, 2011) consider two parallel queues, where uninformed customers can infer some quality information from observable queue lengths. Debo et al. (2012) consider an observable queue with both informed and uninformed customers. They show that uninformed customers' pure equilibrium joining strategy is a hole-avoiding one. Many recent studies consider other quality signals, such as service or waiting time (Debo and Veeraraghavan, 2014; Kremer and Debo, 2016), price and wait lines (Debo et al., 2020) and information generated by customers (Yu et al., 2016; Wang and Hu, 2020). Different from the aforementioned work, here we consider a signaling game where the server's queue disclosure behavior is a signal of his quality level.

Under the signaling game, the sender takes a signaling action after the realization of the state of the world. We note that some recent studies about the queueing system consider a different timing sequence of the game by adopting Bayesian persuasion (Kamenica and Gentzkow, 2011), under which the sender pre-commits to a strategy before the state of the

world is realized. Lingenbrink and Iyer (2019) apply Bayesian persuasion to a queueing setting. They consider that the server pre-commits to a queue-length-dependent signaling strategy. Guo et al. (2020) consider the uncertain service quality and the server can ex ante commit to a quality-dependent queue-disclosure strategy to persuade more customers to join the system. In their paper, the queue-disclosure strategy commitment is made before the service quality is realized. By contrast, our research considers that the server has no commitment power and makes his queue disclosure decision after the quality type is realized.

Our research is also related to the stream of research on information provision and purchase in queues. Hassin and Haviv (1994) examine a parallel queueing system in which customers could buy information on the queue lengths to join the shorter queue. Hassin (2007) studies a scenario in which the server knows his service quality and other system parameters and decides whether or not to disclose such information. In Hassin and Roet-Green (2017), customers may balk, join directly, or buy the queue-length information first and then make their joining-or-balking decisions. Hassin and Roet-Green (2018) further consider a setting with parallel servers in which an uninformed customer becomes informed after paying for inspecting the queues. In our research, we do not consider information purchase. Instead, uninformed customers can infer the server's quality based on his queue disclosure action.

The research considering strategic customers in queueing systems originates from Naor (1969). In this research stream, our work is related to those studies on delay announcements. Hassin (1986) investigates a server's incentive to disclose the queue length information and find that the server prefers concealing (resp., revealing) the queue in a small-sized (resp., large-sized) market. Other studies that investigate the impact of delay announcements include Whitt (1999), Armony and Maglaras (2004a, 2004b), Burnetas and Economou (2007), Guo and Zipkin (2007), Armony et al. (2009), Guo and Hassin (2011), Yu et al. (2016), Ibrahim et al. (2017), Yu et al. (2017), Hu et al. (2018), and Yu et al. (2021), etc. We refer the interested readers to the two survey books, Hassin and Haviv (2003) and Hassin (2016), and the review papers of Aksin et al. (2007) and Ibrahim (2018) and references therein for the works in this research stream.

3.3 Model Setup and Equilibrium Concepts

In this section, we first describe our signaling game, and then present the equilibrium concepts used in the game analysis.

3.3.1 Timing of Signaling Game

Consider a single-server queueing system. Nature moves first and determines the server's quality type t according to a Bernoulli distribution: with probability δ , the server (he) is of the high-quality type (labelled H) and with probability $1 - \delta$, he is of the low-quality type (labelled L), where $0 < \delta < 1$. After observing his own quality type t , $t \in \mathbb{T} := \{H, L\}$, the server can use the queue disclosure action, i.e., *revealing* the queue denoted by R or *concealing* the queue denoted by C as a signal to convey his quality information to customers. Let $\mathbb{S} = \{R, C\}$ denote the server's signal set. Customers arrive at the server according to a Poisson process with rate λ , and they are all uninformed of the service quality. Service times are independent and identically distributed exponential random variables with rate parameter μ . Let $\rho := \lambda/\mu$. Customers who join the system receive the same quality of service and incur the same waiting cost of θ per unit time in the system (waiting time plus service time). When a customer (she) is served by a high-quality (resp. low-quality) server, she receives a monetary reward V_H (resp. V_L). The inequality $V_H > V_L > \frac{\theta}{\mu}$ is required to ensure that at least one customer joins the system. Upon observing the server's queue disclosure behavior (revealing or concealing), uninformed customers update their beliefs about the server being the high-quality type accordingly, denoted by δ^R and δ^C , respectively. Customers then decide whether or not to join the system. When the server conceals the queue, the queue is unobservable (labelled U), and each (identical) customer has two pure strategies: to join the queue or not to join. Then, a pure or mixed strategy can be expressed as the joining probability. When the server reveals the queue, the queue becomes observable (labelled O). In steady state, the probability of the queue length being i ($i = 0, 1, \dots$) in case of the high-quality (resp. low-quality) server is denoted as $\pi_{i,H}(\delta^R)$ (resp. $\pi_{i,L}(\delta^R)$). And after observing the queue length i ($i = 0, 1, \dots$), the uninformed customers further update their belief (denoted as the probability of high quality) as $Pr(H|i, \delta^R)$. Then, the customers make their joining-or-balking decisions at each queue length. We normalize the server's reward from serving a customer to be one, and hence the server's payoff is equal to the customers' effective arrival rate. In summary, the timing of our signaling game is as follows:

- (1) Nature chooses the server's service quality type $t \in \mathbb{T} := \{H, L\}$ from a Bernoulli distribution.
- (2) The server learns his quality type and then chooses a queue disclosure action from the set $\mathbb{S} = \{R, C\}$ with some probability.

- (3) Uninformed customers update their beliefs about the server's service quality based on his queue disclosure action and the queue length in a revealed queue.
- (4) Customers make their respective joining-or-balking decisions.

And we can present the signaling game using the extensive form in Figure 3.1.

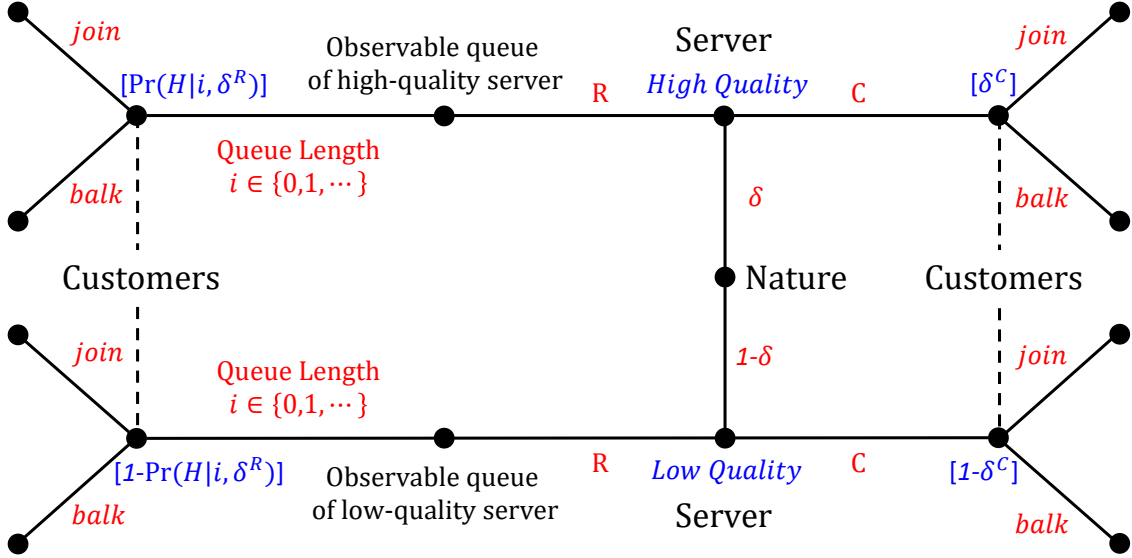


Figure 3.1: The extensive-form representation of the signaling game (with payoffs ignored for simplicity)

3.3.2 Definitions of Sequential Equilibria and Perfect Sequential Equilibria

In this work, we apply the *sequential equilibrium* concept (Kreps and Wilson, 1982) to solve our signaling game. The sequential equilibrium of our signaling game is defined as below.

Definition 3.1. (*Sequential Equilibrium*) A sequential equilibrium of the signaling game is a behavior-belief profile consisting of the server's signaling rules $f(s|t)$, where $f(s|t)$ specifies the probability that the type- t ($t \in \mathbb{T}$) server chooses signal s ($s \in \mathbb{S}$), customers' joining rules, and customers' beliefs δ^C , δ^R and $Pr(H|i, \delta^R)$ ($i = 0, 1, \dots$), which shall satisfy the following two conditions:

- (i) (**Sequential Rationality**) No player deviates from the equilibrium strategy on his or her each information set under the belief specified on it.

(ii) (**Consistency**) When the server sends signal $s \in \mathbb{S}$ with positive probability, customers update their beliefs using signal s according to the Bayes' rule; that is, if $\delta f(s|H) + (1 - \delta)f(s|L) > 0$, then $\delta^s = \frac{\delta f(s|H)}{\delta f(s|H) + (1 - \delta)f(s|L)}$. And after observing queue length i ($i = 0, 1, \dots$) in a revealed queue, the uninformed customers further update their belief as $Pr(H|i, \delta^R) = \frac{\delta^R \pi_{i,H}(\delta^R)}{\delta^R \pi_{i,H}(\delta^R) + (1 - \delta^R) \pi_{i,L}(\delta^R)}$.

In an equilibrium, if the server sends a signal with a positive probability, we say that this signal is *on the equilibrium path*; otherwise, it is *off the equilibrium path*. The condition (ii) in the above definition does not put any restriction on the customers' off-equilibrium-path posterior beliefs after seeing R or C . This may lead to multiple equilibria, and some of them may be unreasonable. We then adopt the *perfect sequential equilibrium* (Grossman and Perry, 1986) as a further refinement criterion to impose restrictions on the off-equilibrium-path beliefs. Besides the above-described two conditions, the perfect sequential equilibrium essentially requires the following *credibility (of the updating rule)* for our signaling game.

Definition 3.2. (Credible Updating Rule) For a signal s off the equilibrium path, given the customers' equilibrium queueing strategies under the signal s and new belief (satisfying the credible updating rule), denote the set of types of the server that can be strictly better off by deviating from the equilibrium path to s by \mathbb{T}' , and the set of types of the server that are indifferent between deviating to s and staying at the equilibrium path by \mathbb{T}'' . Let $h(t)$ be the probability of type- t ($t \in \mathbb{T}$) server deviating from the equilibrium path to s , which shall satisfy $h(t) = 1$ if $t \in \mathbb{T}'$, $h(t) \in [0, 1]$ if $t \in \mathbb{T}''$, and $h(t) = 0$ if $t \in \mathbb{T}/(\mathbb{T}' \cup \mathbb{T}'')$. If there exists a nonempty set $\mathbb{T}' \cup \mathbb{T}''$, then

(a) the customers' posterior belief on type t upon observing the signal s is $\frac{w(t)h(t)}{\sum_{t' \in \mathbb{T}' \cup \mathbb{T}''} w(t')h(t')}$ (we require $\sum_{t \in \mathbb{T}' \cup \mathbb{T}''} h(t) > 0$), where we use $w(\cdot)$ to denote the prior belief (i.e., $w(H) = \delta$ and $w(L) = 1 - \delta$),^{3.1}

(b) sets \mathbb{T}' and \mathbb{T}'' remain unchanged under the above posterior beliefs;

otherwise, there is no restriction on customers' posterior belief after seeing the off-equilibrium-path signal s .

Under Definition 3.2, when the uninformed customers hold a posterior belief γ satisfying the credible updating rule after seeing an off-equilibrium-path signal, their corresponding

^{3.1}There may exist many posterior beliefs satisfying the credible updating rule. The *perfect sequential equilibrium* concept, however, does not specify the selection criteria.

equilibrium queueing strategy will make the server of the quality type in the set \mathbb{T}' strictly better off by deviating with probability 1 and the server of the type in the set \mathbb{T}'' indifferent between staying and deviating with some probability of deviating such that the customers' posterior belief is γ . The process of identifying a belief satisfying the credible updating rule is essentially a fixed-point argument.

3.4 Equilibrium Analysis

Below, we first analyze customers' equilibrium queueing strategies in both the revealed and concealed queues under any given beliefs $\delta^R \in [0, 1]$ and $\delta^C \in [0, 1]$, and then conduct the sequential equilibrium analysis for the whole signaling game.

3.4.1 Customers' Equilibrium Queueing Strategies and Effective Arrival Rates

Since all customers are regarded as indifferent, we only consider symmetric strategies in the customer games in a revealed or concealed queue. If the server reveals his queue, uninformed customers update their beliefs as $Pr(H|i, \delta^R) = \frac{\delta^R \pi_{i,H}(\delta^R)}{\delta^R \pi_{i,H}(\delta^R) + (1-\delta^R) \pi_{i,L}(\delta^R)}$ after observing queue length i ($i = 0, 1, \dots$). As all customers are uninformed, the probabilities for the queue length being i are the same for both the high-quality and low-quality servers (i.e., $\pi_{i,H}(\delta^R) = \pi_{i,L}(\delta^R)$). Hence, we can get that $Pr(H|i, \delta^R) = \delta^R$ for all queue length i . Based on Naor (1969), we know that the uninformed customers *all join* if and only if the queue length (including the one in service) upon arrival does not exceed the threshold $n(\delta^R) := \lfloor [\delta^R V_H + (1 - \delta^R) V_L] \mu / \theta \rfloor - 1$, where $\lfloor \cdot \rfloor$ is the floor function. In the following analysis, we simply use $n(\delta^R)$ to denote the customers' equilibrium queueing strategy in a revealed queue. In steady state, the system under a revealed queue is an $M/M/1/(n(\delta^R) + 1)$ queue with a capacity constraint of $n(\delta^R) + 1$. Let $p_{n(\delta^R)+1}$ be the probability that the queue length is $n(\delta^R) + 1$. Then, $p_{n(\delta^R)+1} = \rho^{n(\delta^R)+1} / \sum_{i=0}^{n(\delta^R)+1} \rho^i$. Denote the customers' effective arrival rate to this observable queue as $\lambda^O(\delta^R)$. We then have

$$\lambda^O(\delta^R) = \lambda(1 - p_{n(\delta^R)+1}) = \frac{\lambda \sum_{k=0}^{n(\delta^R)} \rho^k}{\sum_{i=0}^{n(\delta^R)+1} \rho^i}. \quad (3.1)$$

Notice that we can rewrite $\lambda^O(\delta^R) = \mu - \frac{\mu}{\sum_{i=0}^{n(\delta^R)+1} \rho^i}$, from which we can easily see that $\lambda^O(\delta^R)$ is strictly increasing with the potential arrival rate λ .

Next, consider the case where the server conceals the queue with uninformed customers' belief as δ^C . According to Edelson and Hildebrand (1975), we know that if the potential arrival rate λ is small enough ($\lambda < \mu - \theta/[\delta^C V_H + (1 - \delta^C)V_L]$), all customers join the system; otherwise, customers in equilibrium adopt a mixed strategy, joining the system with probability $\frac{\mu - \theta/[\delta^C V_H + (1 - \delta^C)V_L]}{\lambda}$. Denote the customers' equilibrium joining probability by $p(\delta^C)$. Then, we have

$$p(\delta^C) = \begin{cases} 1, & \text{if } \lambda < \mu - \theta/[\delta^C V_H + (1 - \delta^C)V_L]; \\ \frac{\mu - \theta/[\delta^C V_H + (1 - \delta^C)V_L]}{\lambda}, & \text{otherwise.} \end{cases} \quad (3.2)$$

Let $\lambda^U(\delta^C)$ be the customers' effective arrival rate to this unobservable queue, which then can be derived as

$$\lambda^U(\delta^C) = \begin{cases} \lambda, & \text{if } \lambda < \mu - \theta/[\delta^C V_H + (1 - \delta^C)V_L]; \\ \mu - \theta/[\delta^C V_H + (1 - \delta^C)V_L], & \text{otherwise.} \end{cases} \quad (3.3)$$

3.4.2 Sequential Equilibria Analysis

We are now ready to derive the server's equilibrium signaling strategy. When all customers are uninformed, the sequential rationality condition in Definition 3.1 is indeed equivalent to the following requirements: the customers' joining rule is $(n(\delta^R), p(\delta^C))$, and the server's signaling rule maximizes his expected payoff such that $\forall t \in \mathbb{T}$, $f(R|t) > 0$ (resp. $f(C|t) > 0$) only if $\lambda^O(\delta^R) \geq \lambda^U(\delta^C)$ (resp. $\lambda^U(\delta^C) \geq \lambda^O(\delta^R)$). And we can express a sequential equilibrium as

$$[(f(R|H), f(R|L)), (n(\delta^R), p(\delta^C)), \delta^R, \delta^C].$$

In particular, there exist three kinds of signaling strategies: a *pure strategy* in which both high- and low-quality types of the server choose to either always reveal or always conceal the queue (i.e., both $f(R|H)$ and $f(R|L)$ are either 0 or 1), a *mixed strategy* in which both types of the server randomize between revealing and concealing the queue (i.e., both $f(R|H)$ and $f(R|L)$ are strictly between 0 and 1), and a *hybrid strategy* in which one type of the server chooses to either always reveal or always conceal the queue but the other type randomizes between these two actions (i.e., exactly one of the two probabilities $f(R|H)$ and $f(R|L)$ is

either 0 or 1 but the other one is strictly between 0 and 1). The pure strategies can be further classified into two types: the *pooling* strategy in which both types of the server send the same signal, and the *separating* strategy in which two types of the server send different signals. In our signaling game with two types of the server and two possible signals, there are four possible pure-strategy sequential equilibria. In the following analysis, we first analyze the pure strategies and then the hybrid and mixed ones.

Pure strategy analysis

We now investigate the four pure strategies one by one. For simplicity, let (s', s'') with $s', s'' \in \{R, C\}$ denote the pure strategy played by two types of the server under which the high-quality type server always chooses signal s' while the low-quality one always chooses s'' , i.e., $f(s'|H) = 1$ and $f(s''|L) = 1$.

(1) (R, R) , i.e., *Pooling on R*. Then, R is on the equilibrium path, and by Bayes' rule, customers' updated belief after observing R is still $\delta^R = \delta$. Hence, the payoffs to both types of the server are the same, $\lambda^O(\delta)$. To check whether both types of the server are willing to stay on R , we need to specify the off-equilibrium-path belief δ^C . The sequential equilibrium concept does not put any restriction on δ^C . As long as the off-equilibrium-path belief δ^C leads to $\lambda^U(\delta^C) \leq \lambda^O(\delta)$, both types of the server have no incentive to deviate to C . So, a pooling sequential equilibrium $[(R, R), (n(\delta^R), p(\delta^C)), \delta^R = \delta, \delta^C]$ with δ^C satisfying $\lambda^U(\delta^C) \leq \lambda^O(\delta)$ is an equilibrium outcome.

(2) (C, C) , i.e., *Pooling on C*. Similarly, a pooling sequential equilibrium $[(C, C), (n(\delta^R), p(\delta^C)), \delta^R, \delta^C = \delta]$ with the off-equilibrium-path belief δ^R satisfying $\lambda^O(\delta^R) \leq \lambda^U(\delta)$ is an equilibrium outcome.

(3) (R, C) , i.e., *Separation with the H-type sending R and L-type sending C*. If the server adopts this separating strategy, then both R and C are on the equilibrium path, and by Bayes' rule, customers' beliefs upon observing the signal are updated as $\delta^R = 1$ and $\delta^C = 0$. We now check when this separating strategy can be sustained. Note that if $\lambda^O(1) > \lambda^U(0)$, the low-quality server becomes better off by deviating to R , while if $\lambda^O(1) < \lambda^U(0)$, the high-quality server benefits by deviating to C . That being so, only when $\lambda^O(1) = \lambda^U(0)$ can the separating sequential equilibrium $[(R, C), (n(\delta^R), p(\delta^C)), \delta^R = 1, \delta^C = 0]$ be sustained as an equilibrium outcome.

(4) (C, R) , i.e., *Separation with H-type sending C and L-type sending R*. Similarly, the separating sequential equilibrium $[(C, R), (n(\delta^R), p(\delta^C)), \delta^R = 0, \delta^C = 1]$ can be sustained as

an equilibrium outcome only when $\lambda^O(0) = \lambda^U(1)$.

Denote the unique crossing point of $\lambda^O(\delta)$ and $\lambda^U(\delta)$ as $\hat{\lambda}$. From above analyses, we can obtain the following result.

Proposition 3.1. *Consider that all customers are uninformed of service quality. Then, if $\lambda < \hat{\lambda}$, the unique pure-strategy perfect sequential equilibrium is pooling on C (i.e., (C, C)); otherwise, it is pooling on R (i.e., (R, R)), except at those potential arrival rates under which $\lambda^O(1) = \lambda^U(0)$ and $\lambda^O(0) = \lambda^U(1)$.*

Proposition 3.1 shows that the separating equilibria exist only when the potential arrival rate λ takes some specific values; otherwise, the pure-strategy perfect sequential equilibria must be pooling. We can see that when all customers are uninformed, there exists a potential arrival rate threshold, below which both types of the server prefer to conceal the queue and above which they both prefer to reveal the queue. Similar results have been obtained in Hassin (1986) and Chen and Frank (2004). The only difference is that in both papers, there is no signaling issue and the server simply compares his payoffs under observable and unobservable queues to decide whether or not to reveal his queue. Here, we consider a signaling game under which the server's queue disclosure behavior is a signal of service quality and obtain similar results.

Hybrid and mixed strategies

The hybrid and mixed strategies can be similarly analyzed. To simplify the presentation, we relegate the detailed analysis to Appendix B.1. We next compare the server's payoffs (i.e., effective arrival rates) under different types of sequential equilibria and obtain the following proposition.

Proposition 3.2. *The effective arrival rates of both types of the server under the pure-strategy perfect sequential equilibria are larger than the respective ones under the mixed- and hybrid-strategy sequential equilibria.*

We now provide a numerical example illustrating it; see Figure 3.2.

Example 3.1. *Suppose $V_H = 3$, $V_L = 2$, $\mu = 1$, $\theta = 0.5$, and $\delta = 0.5$. From Figure 3.2, we can see that the pure-strategy perfect sequential equilibria exist on the whole range of λ , while the hybrid- or mixed-strategy sequential equilibria can be sustained only on several bounded ranges of λ . Note that no matter what the equilibria are, the effective arrival rates to*

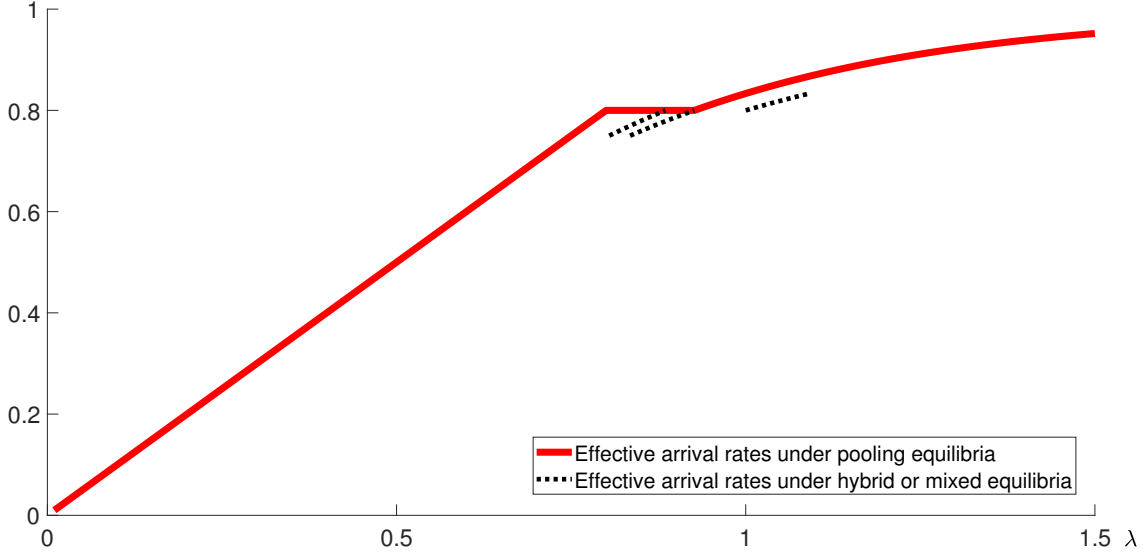


Figure 3.2: Comparison of effective arrival rates under various sequential equilibria: $V_H = 3$, $V_L = 2$, $\mu = 1$, $\theta = 0.5$ and $\delta = 0.5$

the high- and low-quality servers are always the same when all customers are homogeneously uninformed. Hence, in Figure 3.2, we do not distinguish the server type for each effective arrival rate. Figure 3.2 confirms that the effective arrival rates under the hybrid or mixed sequential equilibria indeed cannot exceed the ones under the pure-strategy perfect sequential equilibria.

Impact of the signaling effect

We refer to our setting with the signaling effect of the queue disclosure action as the *signaling case* and the setting without as the *non-signaling case*. Under the non-signaling case, it can be easily verified that the server prefers to conceal the queue if $\lambda^U(\delta) > \lambda^O(\delta)$ and to reveal it otherwise. The non-signaling case then performs the same as the signaling case when we consider the pooling perfect sequential equilibria. That being so, by Propositions 3.1 and 3.2, we have the following result.

Corollary 3.1. *When all customers are uninformed of service quality, the effective arrival rate under the pooling perfect sequential equilibria is the same as the maximal one under the non-signaling case.*

Corollary 3.1 implies that due to pooling being the equilibrium strategy of the server, using the queue disclosure action as a signal of service quality has no effect on the server's effective arrival rate when all the customers are uninformed.

3.5 Signaling Game with Heterogeneous Customers

Until now, we have considered homogeneous customers who are all uninformed of the service quality. In this section, we extend this assumption by allowing some of the customers to be informed of the service quality. The informed customers can either be *positively informed customers* if the server is of high quality or *negatively informed customers* if the server is of low quality. We use the variable q ($0 < q < 1$) to represent the fraction of the informed customers. The signaling game with such heterogeneous customers becomes more complicated. In the following analysis, we first investigate the equilibrium queueing strategies of both the informed and uninformed customers in unobservable and observable queues, and then analyze the sequential equilibria.

3.5.1 Customers' Equilibrium Queueing Strategies and Effective Arrival Rates

We now analyze the customers' equilibrium queueing strategies and derive the effective arrival rates given that the queue is concealed from or revealed to customers.

Concealed Queue

Consider the case where the server conceals the queue. Assume that upon observing the server's queue concealment behavior, uninformed customers hold a belief that the server's service quality is high with probability δ^C ($0 < \delta^C < 1$).^{3.2} Uninformed customers then join the system with probability $p_{un}(\delta^C)$. As to the informed customers who know the service quality, they join the system with probability $p_H(\delta^C)$ (resp. $p_L(\delta^C)$) when the server's service quality is high (resp. low). In such a static game with incomplete information, we denote the customer queueing strategy by the triplet (p_L, p_{un}, p_H) and the equilibrium strategy profile by (p_L^U, p_{un}^U, p_H^U) with δ^C omitted for notational convenience in the following analysis.

Given customers' queueing strategy (p_L, p_{un}, p_H) , the expected utility of a positively informed customer is $u_H(p_{un}, p_H) := V_H - \frac{\theta}{\mu - \lambda(qp_H + (1-q)p_{un})}$, the one of a negatively informed customer is $u_L(p_L, p_{un}) := V_L - \frac{\theta}{\mu - \lambda(qp_L + (1-q)p_{un})}$, and the one of an uninformed customer is $u_{un}(p_L, p_{un}, p_H) := \delta^C \left[V_H - \frac{\theta}{\mu - \lambda(qp_H + (1-q)p_{un})} \right] + (1 - \delta^C) \left[V_L - \frac{\theta}{\mu - \lambda(qp_L + (1-q)p_{un})} \right]$. The equilibrium queueing behaviors of all types of customers are totally determined by the above

^{3.2}The special cases where δ^C is 0 or 1 are analyzed in Appendix B.2.

Table 3.1: Equilibrium joining strategy (p_L^U, p_{un}^U, p_H^U) in unobservable queues

Range of λ	$(0, \mu - \frac{\theta}{V_L}]$	$(\mu - \frac{\theta}{V_L}, \min(\lambda_1, \lambda_2)]$	$(\min(\lambda_1, \lambda_2), \bar{\lambda}]$	$(\bar{\lambda}, \frac{\theta(V_H - V_L)}{qV_H V_L}]$	$(\frac{\theta(V_H - V_L)}{qV_H V_L}, +\infty)$
Case 1: $\lambda_1 < \lambda_2$	(1, 1, 1)	$(p_L^U, 1, 1)$	(0, 1, 1)	$(0, p_{un}^U, 1)$	(p_L^U, p_{un}^U, p_H^U)
Case 2: $\lambda_1 \geq \lambda_2$	(1, 1, 1)	$(p_L^U, 1, 1)$	(p_L^U, p_{un}^U, p_H^U)		

three utilities. The following proposition shows customers' equilibrium queueing strategies (p_L^U, p_{un}^U, p_H^U) under various cases. When any one element in the triplet (p_L^U, p_{un}^U, p_H^U) equals zero or one, we simply write it as 0 or 1. For example, the triplet $(0, p_{un}^U, 1)$ represents that $p_L^U = 0$, $p_H^U = 1$ and $0 \leq p_{un}^U \leq 1$.

Proposition 3.3. *When the queue is unobservable, customers' equilibrium queueing strategies (p_L^U, p_{un}^U, p_H^U) under various cases are summarized in Table 3.1 with the parameter values specified as follows: $\lambda_1 = \frac{\mu - \theta/V_L}{1 - q}$, $\lambda_2 = \mu - \frac{\theta}{V_H}$, $\bar{\lambda}$ is the unique solution for λ satisfying $0 < \lambda < \mu$ in the equation $u_{un}(0, 1, 1) = 0$, the value of p_L^U in $(p_L^U, 1, 1)$ is $\frac{\mu - \theta/V_L}{\lambda q} - \frac{1 - q}{q}$, p_{un}^U in $(0, p_{un}^U, 1)$ is the unique solution for p_{un} satisfying $0 < p_{un} < \min\left\{\frac{\mu - q\lambda}{(1 - q)\lambda}, 1\right\}$ in the equation $u_{un}(0, p_{un}, 1) = 0$, and (p_L^U, p_{un}^U, p_H^U) represents a continuum of equilibria with any $p_{un}^U \in \left[\max\left\{0, \frac{\mu - \theta/V_H}{\lambda(1 - q)} - \frac{q}{1 - q}\right\}, \min\left\{1, \frac{\mu - \theta/V_L}{\lambda(1 - q)}\right\}\right]$ and the corresponding $p_H^U = \frac{\mu - \theta/V_H}{\lambda q} - \frac{(1 - q)p_{un}^U}{q}$ and $p_L^U = \frac{\mu - \theta/V_L}{\lambda q} - \frac{(1 - q)p_{un}^U}{q}$.*

As illustrated in Table 3.1, under Case 1, as the potential arrival rate λ increases across $\frac{\theta(V_H - V_L)}{qV_H V_L}$, the equilibrium outcome evolves from a unique equilibrium to multiple equilibria. Let us investigate this interesting phenomenon in detail. For $\lambda < \frac{\theta(V_H - V_L)}{qV_H V_L}$, the expected utility of at least one type of customers is strictly positive, and at least one of the three probabilities (p_L^U, p_{un}^U, p_H^U) is specified as 0 or 1 with others uniquely identified by making the corresponding expected utility as 0, which makes the final equilibrium triplet unique. However, when $\lambda > \frac{\theta(V_H - V_L)}{qV_H V_L}$, the expected utilities of all types of customers are 0 in equilibrium, and this yields three equations $u_H(p_{un}^U, p_H^U) = 0$, $u_{un}(p_L^U, p_{un}^U, p_H^U) = 0$ and $u_L(p_L^U, p_{un}^U) = 0$, any one of which is redundant given the other two. Then, two equations with three variables yield multiple equilibria with $p_{un}^U \in \left[\max\left\{0, \frac{\mu - \theta/V_H}{\lambda(1 - q)} - \frac{q}{1 - q}\right\}, \min\left\{1, \frac{\mu - \theta/V_L}{\lambda(1 - q)}\right\}\right]$. When $\lambda = \frac{\theta(V_H - V_L)}{qV_H V_L}$, we get that $0 < \frac{\mu - \theta/V_H}{\lambda(1 - q)} - \frac{q}{1 - q} = \frac{\mu - \theta/V_L}{\lambda(1 - q)} < 1$, and thus the value of p_{un}^U is uniquely as $\frac{qV_H V_L(\mu - \theta/V_L)}{(1 - q)\theta(V_H - V_L)}$ with the uniquely corresponding $p_H^U = 1$ and $p_L^U = 0$. As λ becomes a little bit larger than $\frac{\theta(V_H - V_L)}{qV_H V_L}$, we get that $0 < \frac{\mu - \theta/V_H}{\lambda(1 - q)} - \frac{q}{1 - q} < \frac{\mu - \theta/V_L}{\lambda(1 - q)} < 1$, and thus the length of the feasible range of p_{un}^U becomes strictly positive, yielding multiple equilibria immediately. Similarly to the above analysis, the evolution from a unique equilibrium to multiple equilibria

Table 3.2: Effective joining rates $\lambda_H^U(\delta^C)$ and $\lambda_L^U(\delta^C)$ in unobservable queues

Range of λ		$(0, \mu - \frac{\theta}{V_L}]$	$(\mu - \frac{\theta}{V_L}, \min(\lambda_1, \lambda_2)]$	$(\min(\lambda_1, \lambda_2), \bar{\lambda}]$	$(\bar{\lambda}, \frac{\theta(V_H - V_L)}{qV_H V_L}]$	$(\frac{\theta(V_H - V_L)}{qV_H V_L}, +\infty)$
Case 1: $\lambda_1 < \lambda_2$	$\lambda_H^U(\delta^C)$	λ			$x(\lambda)$	$\mu - \frac{\theta}{V_H}$
	$\lambda_L^U(\delta^C)$	λ	$\mu - \frac{\theta}{V_L}$	$(1 - q)\lambda$	$x(\lambda) - q\lambda$	$\mu - \frac{\theta}{V_L}$
Case 2: $\lambda_1 \geq \lambda_2$	$\lambda_H^U(\delta^C)$	λ			$\mu - \frac{\theta}{V_H}$	
	$\lambda_L^U(\delta^C)$	λ	$\mu - \frac{\theta}{V_L}$			

in Case 2 can be illustrated.

An interesting observation is that, in Case 1, the negatively informed customers surely join in a small-sized market and join with some probability in a large-sized market. However, they never join when the market size is moderate (i.e., $\lambda \in (\lambda_1, \frac{\theta(V_H - V_L)}{qV_H V_L}]$). Intuitively, the expected quality level of uninformed customers is higher than the one of the negatively informed customers. In that moderate market range, the uninformed customers join the queue at a positive probability. But this makes the expected utility of a negatively informed customer strictly negative (i.e., $u_L(0, p_{un}^U) < 0$), preventing them from joining.

Based on Proposition 3.3, we can further derive the effective arrival rates under any given market size λ for both high- and low-quality servers. Denote $\lambda_H^U(\delta^C)$ and $\lambda_L^U(\delta^C)$ as the effective arrival rates of the high- and low-quality servers when the uninformed customers hold the belief that the server is of high quality with probability δ^C , respectively. Then, we can get the following result.

Proposition 3.4. *When the queue is unobservable, $\lambda_H^U(\delta^C)$ and $\lambda_L^U(\delta^C)$, the effective arrival rates of the high- and low-quality servers, are summarized in Table 3.2, where $x(\lambda)$ is the unique solution for x satisfying $q\lambda < x < \mu$ in the equation $\delta^C V_H + (1 - \delta^C)V_L = \delta^C \frac{\theta}{\mu - x} + (1 - \delta^C) \frac{\theta}{\mu - (x - q\lambda)}$, and all other parameter values are specified in Proposition 3.3.*

In Case 1, $\lambda_H^U(\delta^C)$ is non-decreasing with the potential arrival rate λ , and $\lambda_L^U(\delta^C)$ is decreasing with λ when $\lambda \in (\bar{\lambda}, \frac{\theta(V_H - V_L)}{qV_H V_L}]$ and non-decreasing otherwise. While in Case 2, both $\lambda_H^U(\delta^C)$ and $\lambda_L^U(\delta^C)$ are non-decreasing with λ .

Although this queueing game may have multiple equilibria on some ranges of the potential arrival rate, Proposition 3.4 shows that the effective arrival rates for both types of the server are in fact unique.^{3.3} This is because multiple equilibria occur only when all types of customers

^{3.3}Different from the general case $0 < \delta^C < 1$, when $\delta^C = 0$ or 1 , such uniqueness does not hold any more. For the convenience and consistency of the following sequential equilibrium analysis, we only consider $\lambda_H^U(0) = \lim_{\delta^C \rightarrow 0^+} \lambda_H^U(\delta^C)$ and $\lambda_L^U(1) = \lim_{\delta^C \rightarrow 1^-} \lambda_L^U(\delta^C)$ based on the continuities of $\lambda_H^U(\delta^C)$ and $\lambda_L^U(\delta^C)$ in δ^C ($0 < \delta^C < 1$). This part of analysis can be found in Appendix B.2.

obtain an expected utility of zero (see Proposition 3.3 and its proof). Under such a scenario, different equilibria only affect the composition of the effective arrival rate, i.e., the proportion of those joining customers who are informed or uninformed.

In the first case of Proposition 3.4, when the market size (reflected by the potential arrival rate) falls into the range $\lambda \in \left(\bar{\lambda}, \frac{\theta(V_H - V_L)}{qV_H V_L}\right]$, increasing the market size actually reduces the demand (reflected by the effective arrival rate) for the low-quality server. This is quite counter-intuitive. The explanation is as follows. When the server is of low quality and the potential arrival rate λ falls into this range, informed customer will not join the system. Only part of the uninformed customers join with an expected utility $u_{un}(0, p_{un}^U, 1) = 0$. However, for the uninformed customers, as λ increases, p_{un}^U need be decreased to keep $u_{un}(0, p_{un}^U, 1) = 0$, which means that uninformed customers have less incentives to join, leading to $\lambda_L^U(\delta^C)$ decreasing with λ . Figure 3.3 illustrates the changes of effective arrival rates to both high- and low-quality servers under Case 1 of Proposition 3.4.




Figure 3.3: Illustration of effective arrival rates to the high- and low-quality servers under a concealed queue: $V_H = 7$, $V_L = 2$, $\mu = 1$, $\theta = 1$, $\delta^C = 0.5$ and $q = 0.3$

Revealed Queue

When the server reveals the queue, all customers, both informed and uninformed, inspect the queue length upon arrival and then decide whether or not to join. Upon observing the server's queue revelation behavior, uninformed customers hold a belief that the server is of high quality with probability δ^R . Next, we first focus on the general case $0 < \delta^R < 1$, and then

analyze two special cases where δ^R is 0 or 1. For simplicity and notational convenience in the following analysis, we omit the term δ^R . For an informed customer, the equilibrium strategy is simple and definite: when the server is of high (resp. low) quality, they join the queue unless it is longer than a threshold $n(1) := \lfloor V_H \mu / \theta \rfloor - 1$ (resp. $n(0) := \lfloor V_L \mu / \theta \rfloor - 1$). In other words, a positively informed customer joins the queue with probability $p_H^O(i) = 1$ at queue length i when $i = 0, 1, \dots, n(1)$, and with probability $p_H^O(i) = 0$ otherwise; a negatively informed customer joins the queue with probability $p_L^O(i) = 1$ at queue length i when $i = 0, 1, \dots, n(0)$, and with probability $p_L^O(i) = 0$ otherwise. For uninformed customers with $0 < \delta^R < 1$, they can infer quality information from the queue length (see Debo et al., 2012), and thus the decision making of an uninformed customer is much more complicated. Clearly, if the queue is shorter than or equal to $n(0)$, the uninformed customer joins the system, i.e., her joining probability $p_{un}^O(i) = 1$ for $i = 0, 1, \dots, n(0)$. Likewise, if it is longer than $n(1)$, she balks, i.e., her joining probability $p_{un}^O(i) = 0$ for $i = n(1) + 1, \dots$. We now need to derive the uninformed customers' joining probability at queue length i for $i = n(0) + 1, \dots, n(1)$. We assume that the joining decision of uninformed customers is made only based on the current queue length when they arrive (see Debo et al., 2012). In this dynamic game with incomplete information, the customers' equilibrium queueing strategy profile is denoted by a set of the triplet $\{(p_L^O(i), p_{un}^O(i), p_H^O(i))\}_{i=0}^{+\infty}$.

Here, our game, where the service rate is the same for two types of the server, is a special case of the “consumer game” in Debo et al. (2012), where different types of the server can adopt different service rates. According to Debo et al. (2012), an equilibrium pure strategy for an uninformed customer is a *hole-avoiding* strategy. Specifically, an uninformed customer behaves as a positively informed customer except at queue length denoted by n_{hole} (namely, the hole) where she plans not to join. The queue-length joining set is hence $\{0, \dots, n_{hole} - 1, n_{hole} + 1, \dots, n(1)\}$. The underlying reason behind such hole-avoiding strategy is as follows. Given that all uninformed customers behave in this way, the fact that an uninformed customer observes a queue length longer than n_{hole} upon arrival implies that sometime in the past, an informed customer had inspected a queue length of n_{hole} and joined the system. The service quality hence must be high as otherwise the informed customer would have balked. For the sake of reading convenience, we illustrate the hole-avoiding decision process of uninformed customers in Debo et al. (2012) using our terms. To keep brevity, we relegate the related review to Appendix B.3. Note that we let $\lambda_{i,H}$ (resp. $\lambda_{i,L}$) be the effective arrival rate at queue length i and $\pi_{i,H}$ (resp. $\pi_{i,L}$) the limiting probability that the number of customers in

the system equals i when the server is of high (resp. low) quality, where $i = 0, 1, \dots, n(1) + 1$.

Denote $\lambda_H^O(\delta^R)$ and $\lambda_L^O(\delta^R)$ as the effective arrival rates of the high- and low-quality servers in equilibrium when uninformed customers hold the belief that the server is of high quality with probability δ^R , respectively. Then, according to the time reversibility of the above ergodic BD processes in steady state, we can easily get that

$$\lambda_H^O(\delta^R) = \mu(1 - \pi_{0,H}) = \mu \left(1 - \frac{1}{\sum_{i=0}^{n_{hole}} \rho^i + q \sum_{i=n_{hole}+1}^{n(1)+1} \rho^i} \right),$$

and

$$\lambda_L^O(\delta^R) = \mu(1 - \pi_{0,L}) = \mu \left(1 - \frac{1}{\sum_{i=0}^{n(0)+1} \rho^i + \sum_{i=n(0)+2}^{n_{hole}} (1-q)^{i-n(0)-1} \rho^i} \right).$$

For two special cases where δ^R is 0 or 1, the analyses are as follows.

- (i) When uninformed customers believe that the server's service quality is low (i.e., $\delta^R = 0$), if the server is indeed of low quality, then customers, both informed and uninformed, hold the same belief $\delta^R = 0$. Hence, the queue-length joining set of all customers is $\{0, \dots, n(0)\}$, and the effective arrival rate is $\lambda_L^O(0) = \lambda^O(0)$. While if the server is of high quality, then uninformed customers and positively informed customers hold totally opposite beliefs, and the queue-length joining sets of the uninformed customers and positively informed customers are $\{0, \dots, n(0)\}$ ^{3.4} and $\{0, \dots, n(1)\}$, respectively. So, the effective arrival rate becomes $\lambda_H^O(0) = \mu \left(1 - \frac{1}{\sum_{i=0}^{n(0)+1} \rho^i + \sum_{i=n(0)+2}^{n(1)+1} q^{i-n(0)-1} \rho^i} \right)$.
- (ii) When uninformed customers believe that the server's service quality is high (i.e., $\delta^R = 1$), if the server is indeed of high quality, then all customers hold the same belief $\delta^R = 1$. Hence, the queue-length joining set of all customers is $\{0, \dots, n(1)\}$, and the effective arrival rate is $\lambda_H^O(1) = \lambda^O(1)$. While if the server is of low quality, then uninformed customers and negatively informed customers hold totally opposite beliefs, and their queue-length joining sets are $\{0, \dots, n(1)\}$ and $\{0, \dots, n(0)\}$, respectively. So, the effective arrival rate becomes $\lambda_L^O(1) = \mu \left(1 - \frac{1}{\sum_{i=0}^{n(0)+1} \rho^i + \sum_{i=n(0)+2}^{n(1)+1} (1-q)^{i-n(0)-1} \rho^i} \right)$.

^{3.4}The belief $\delta^R = 0$ means that the uninformed customers make sure that the service quality is low after seeing a revealed queue, and they do not further infer quality information from the queue length. In other words, a queue longer than $n(0) + 1$ does not convey any quality-related information to all convinced uninformed customers. In the following sequential equilibrium analysis, the term $\lambda_H^O(0)$ only serves for judging an equilibrium, but cannot be the final equilibrium effective arrival rate.

3.5.2 Sequential Equilibria Analysis

The sequential equilibrium of the signaling game with heterogeneous customers can be defined after Definition 3.1 in section 3.3.2. Now, the *Sequential Rationality* requires that the customers' joining rules are (p_L^U, p_{un}^U, p_H^U) and $\{(p_L^O(i), p_{un}^O(i), p_H^O(i))\}_{i=0}^{+\infty}$, and the server's signaling rule maximizes his expected payoff such that $\forall t \in \mathbb{T}$, $f(R|t) > 0$ (resp. $f(C|t) > 0$) only if $\lambda_t^O(\delta^R) \geq \lambda_t^U(\delta^C)$ (resp. $\lambda_t^U(\delta^C) \geq \lambda_t^O(\delta^R)$). Then, we can express a sequential equilibrium in this setting as

$$[(f(R|H), f(R|L)), \{(p_L^U, p_{un}^U, p_H^U), \{(p_L^O(i), p_{un}^O(i), p_H^O(i))\}_{i=0}^{+\infty}\}, \delta^R, \delta^C].$$

We are now ready to analyze the sequential equilibria of the signaling game. In the following analysis, we theoretically investigate the pure strategies and examine whether and when they can be sustained as an equilibrium outcome of our signaling game. The hybrid and mixed strategies can only be analyzed numerically, and the related sequential equilibrium analysis can be found in Appendix B.4. There are still four pure strategies in this signaling game with heterogeneous customers, which are specified one by one as follows.

(1) (R, R) , i.e., *Pooling on R*, under which both types of the server choose to always reveal their queues. Then, R is on the equilibrium path, and by Bayes' rule, uninformed customers' updated belief after observing R is still $\delta^R = \delta$. Let the effective arrival rates to the high- and low-quality servers be $\lambda_H^O(\delta)$ and $\lambda_L^O(\delta)$, respectively. To check whether both types of the server are willing to stay on R , we need to specify the off-equilibrium-path belief δ^C . As long as the off-equilibrium-path belief δ^C leads to $\lambda_H^U(\delta^C) \leq \lambda_H^O(\delta)$ and $\lambda_L^U(\delta^C) \leq \lambda_L^O(\delta)$, both types of the server have no incentive to deviate to C . Hence, with such off-equilibrium-path beliefs, this pooling strategy can be sustained as a sequential equilibrium outcome.

(2) (C, C) , i.e., *Pooling on C*, under which both types of the server choose to always conceal their queues. Then, C is on the equilibrium path, and by Bayes' rule, uninformed customers' updated belief after observing C is still $\delta^C = \delta$. Similarly, as long as the off-equilibrium-path belief δ^R leads to $\lambda_H^O(\delta^R) \leq \lambda_H^U(\delta)$ and $\lambda_L^O(\delta^R) \leq \lambda_L^U(\delta)$, both types of the server have no incentive to deviate to R . Then, with such off-equilibrium-path beliefs, this pooling sequential equilibrium can be sustained as an equilibrium outcome.

The following proposition summarizes the sufficient conditions under which the sequential equilibrium is uniquely a pooling one.

Proposition 3.5. *When the market consists of both informed and uninformed customers,*

there exist two potential arrival rate thresholds, $\hat{\lambda}_C$ and $\hat{\lambda}_R$ satisfying $\hat{\lambda}_R > \hat{\lambda}_C$,^{3.5} such that when $\lambda < \hat{\lambda}_C$, (C, C) , pooling on C is the unique sequential equilibrium with the off-equilibrium-path belief $\delta^R \in [0, 1]$, while when $\lambda > \hat{\lambda}_R$, (R, R) , pooling on R is with the off-equilibrium-path belief $\delta^C \in [0, 1]$.

Proposition 3.5 shows that for the signaling game with heterogeneous customers, concealing (resp. revealing) is still the unique dominant strategy for two types of the server when the market size is small (resp. large) enough. For a medium-sized market, the separating sequential equilibria may exist. We then depict the ranges of the market size λ on which the separating strategies can be sustained as sequential equilibria.

(3) (R, C) , i.e., *Separation with the H-type sending R and L-type sending C*, under which the high-quality server always reveals the queue while the low-quality server always conceals the queue. If the server adopts this separating strategy, then both R and C are on the equilibrium path, and by Bayes' rule, uninformed customers' beliefs upon observing the signal are updated as $\delta^R = 1$ and $\delta^C = 0$. We now check whether this separating strategy can be sustained. Note that if $\lambda_L^O(1) > \lambda_L^U(0)$, the low-quality server becomes strictly better off by deviating to R , and if $\lambda_H^U(0) > \lambda_H^O(1)$, the high-quality server benefits by deviating to C . That being so, only when $\lambda_H^O(1) \geq \lambda_H^U(0)$ and $\lambda_L^U(0) \geq \lambda_L^O(1)$ can this separating sequential equilibrium be sustained.

(4) (C, R) , i.e., *Separation with H-type sending C and L-type sending R*, under which the high-quality server always conceals the queue while the low-quality server always reveals the queue. If the server adopts this separating strategy, then both R and C are on the equilibrium path, and by Bayes' rule, uninformed customers' beliefs upon observing the signal are updated as $\delta^C = 1$ and $\delta^R = 0$. Similarly, this separating sequential equilibrium can be sustained only if $\lambda_H^U(1) \geq \lambda_H^O(0)$ and $\lambda_L^O(0) \geq \lambda_L^U(1)$.

Intuitively, since some of the customers are informed of the true quality type, the high- and low-quality servers may have different incentives to reveal or conceal the queue in a medium-sized market. At some market sizes, the high-quality (resp. low-quality) server prefers revealing (resp. concealing). In this case, the uninformed customers can infer the true quality type from the queue disclosure action. If the high-quality (resp. low-quality) server deviates to concealing (resp. revealing), he knows that the uninformed customers must believe that he is of the low-quality (resp. high-quality) type after observing a concealed (resp. revealed) queue, and he finds that such a deviation will make him worse off, which

^{3.5}The detailed expressions of $\hat{\lambda}_C$ and $\hat{\lambda}_R$ can be found in the proof of Proposition 3.5.

makes the separating sequential equilibrium (R, C) sustained. Similarly, another separating sequential equilibrium (C, R) can be intuitively understood.

Based on the above analysis, the separating sequential equilibrium (R, C) can be sustained only if $\lambda_H^O(1) \geq \lambda_H^U(0)$ and $\lambda_L^U(0) \geq \lambda_L^O(1)$, and another one (C, R) can be sustained only if $\lambda_H^U(1) \geq \lambda_H^O(0)$ and $\lambda_L^O(0) \geq \lambda_L^U(1)$. Let $\Lambda_{O \geq U}^{(R,C)} := \{\lambda | \lambda_H^O(1) \geq \lambda_H^U(0)\}$, $\Lambda_{U \geq O}^{(R,C)} := \{\lambda | \lambda_L^U(0) \geq \lambda_L^O(1)\}$, $\Lambda_{O \geq U}^{(C,R)} := \{\lambda | \lambda_L^O(0) \geq \lambda_L^U(1)\}$, and $\Lambda_{U \geq O}^{(C,R)} := \{\lambda | \lambda_H^U(1) \geq \lambda_H^O(0)\}$. Then, we can get the following results on the separating sequential equilibria.

Proposition 3.6. *For a medium-sized market $\lambda \in (\hat{\lambda}_C, \hat{\lambda}_R)$ (except at several threshold points), at most one of the two separating strategies $((R, C)$ and $(C, R))$ can be sustained as an equilibrium outcome. Specifically, the ranges of the market size λ where the separating sequential equilibria (R, C) and (C, R) exist can be expressed as $\Lambda^{(R,C)} := \Lambda_{O \geq U}^{(R,C)} \cap \Lambda_{U \geq O}^{(R,C)}$ and $\Lambda^{(C,R)} := \Lambda_{O \geq U}^{(C,R)} \cap \Lambda_{U \geq O}^{(C,R)}$, respectively.*

In a separating sequential equilibrium, the server's queue disclosure behavior, revealing or concealing the queue, exactly signals his service quality. Note that we cannot rule out the possibility that pooling and separating sequential equilibria coexist in a medium-sized market. Proposition 3.6 depicts the exact ranges of the market size λ where separating sequential equilibria exist. For $t \in \{H, L\}$ and $d \in \{0, 1\}$, the expressions of $\lambda_t^U(d)$ and $\lambda_t^O(d)$ can be explicitly got. And thus, given the values of all system parameters, we can definitely identify the key ranges $\Lambda^{(R,C)}$ and $\Lambda^{(C,R)}$ in Proposition 3.6. For example, when $(\mu - \theta/V_L)/(1 - q) \geq \mu - \theta/V_H$, according to Proposition 3.4 and Lemma B.1 (see Appendix B.5), we can get the unique crossing point of $\lambda_H^O(1)$ and $\lambda_H^U(0)$ as $\hat{\lambda}_{H1}$, of $\lambda_L^O(1)$ and $\lambda_L^U(0)$ as $\hat{\lambda}_{L1}$, of $\lambda_H^O(0)$ and $\lambda_H^U(1)$ as $\hat{\lambda}_{H0}$, and of $\lambda_L^O(0)$ and $\lambda_L^U(1)$ as $\hat{\lambda}_{L0}$. Then, if $\hat{\lambda}_{H1} \leq \hat{\lambda}_{L1}$, (R, C) is the unique separating sequential equilibrium for $\lambda \in [\hat{\lambda}_{H1}, \hat{\lambda}_{L1}]$; and if $\hat{\lambda}_{H0} \geq \hat{\lambda}_{L0}$, (C, R) is the unique separating sequential equilibrium for $\lambda \in [\hat{\lambda}_{L0}, \hat{\lambda}_{H0}]$. By contrast, when $(\mu - \theta/V_L)/(1 - q) < \mu - \theta/V_H$, the corresponding crossing points may not be unique, and this may make each concerned range (i.e., $\Lambda^{(R,C)}$ or $\Lambda^{(C,R)}$) composed of several separate ranges of λ .

In section 3.4.2 where all customers are uninformed, the effective arrival rates to both the high-quality and low-quality servers are always the same under any sequential equilibrium. However, such a result does not hold any more under the heterogeneous-customers scenario. The comparison result is shown in the following corollary.

Corollary 3.2. *When the market consists of both informed and uninformed customers, the*

effective arrival rate of the high-quality server is weakly larger than that of the low-quality server under any (pure, mixed or hybrid) sequential equilibrium.

Clearly, a positively informed customer is always more likely to join a queue than a negatively informed one. An uninformed customer, however, cannot make the exact inference about the service quality, and thus her joining decision is the same regardless of the server's type. A combination of the above observations then leads to the result stated in Corollary 3.2.

Below, we provide a simple example to illustrate the pure-strategy sequential equilibria. We also numerically show the hybrid- and mixed-strategy equilibria as well. For the equilibrium queueing strategy of uninformed customers in an observable queue, we give priority to the pure strategy with the smallest hole value.

Example 3.2. Consider the parameter values to be $V_H = 2.5$, $V_L = 2$, $\mu = 1$, $\theta = 0.5$, $\delta = 0.5$, and $q = 0.5$. Under this setting, we have $(\mu - \theta/V_L)/(1 - q) > \mu - \theta/V_H$, and thus the above-mentioned crossing points are unique as $\hat{\lambda}_C = \hat{\lambda}_{L1} (= 0.8580) < \hat{\lambda}_{L0} (= 0.8882) < \hat{\lambda}_{H1} (= 0.9265) < \hat{\lambda}_R = \hat{\lambda}_{H0} (= 0.9579)$.

By Proposition 3.5, we know that the unique pure-strategy sequential equilibrium is (C, C) with the off-equilibrium-path belief $\delta^R \in [0, 1]$, namely pooling on C , when the potential arrival rate $\lambda < \hat{\lambda}_C$, and it is (R, R) with the off-equilibrium-path belief $\delta^C \in [0, 1]$, namely pooling on R , when $\lambda > \hat{\lambda}_R$. When $\hat{\lambda}_{L0} \leq \lambda \leq \hat{\lambda}_{H0}$, according to Proposition 3.6, the strategy (C, R) in which the high-quality server conceals the queue and the low-quality server reveals the queue, is a pure-strategy separating sequential equilibrium. Also, for $\hat{\lambda}_{L0} < \lambda < \hat{\lambda}_{H0}$, no pooling sequential equilibrium can be sustained. Another separating strategy (R, C) can never be sustained as a sequential equilibrium. For the remaining range $\hat{\lambda}_C \leq \lambda < \hat{\lambda}_{L0}$, we can show that $\lambda_H^O(\delta^R) < \lambda_H^U(\delta)$ always holds for any belief $\delta^R \in [0, 1]$. Then, as long as the belief δ^R satisfies $\lambda_L^O(\delta^R) \leq \lambda_L^U(\delta)$ (i.e., $0 \leq \delta^R < 1$ in this example), the pure strategy (C, C) can be sustained as a sequential equilibrium. Also, note that in this situation, the set $\mathbb{T}' \cup \mathbb{T}''$ is empty. Therefore, the credible updating rule does not put any restriction on δ^R . As such, (C, C) is a perfect sequential equilibrium for $\lambda \in [\hat{\lambda}_C, \hat{\lambda}_{L0})$ with the off-equilibrium-path belief $\delta^R \in [0, 1]$. See Figure 3.4 for the illustration of our pure-strategy sequential equilibrium outcome and the corresponding effective arrival rates to both types of the server. Figure 3.4 also reconfirms Corollary 3.2 that in equilibrium, the effective arrival rate to the high-quality server is always no less than the one to the low-quality server.

Next, we investigate the mixed and hybrid strategies following the analysis in Appendix

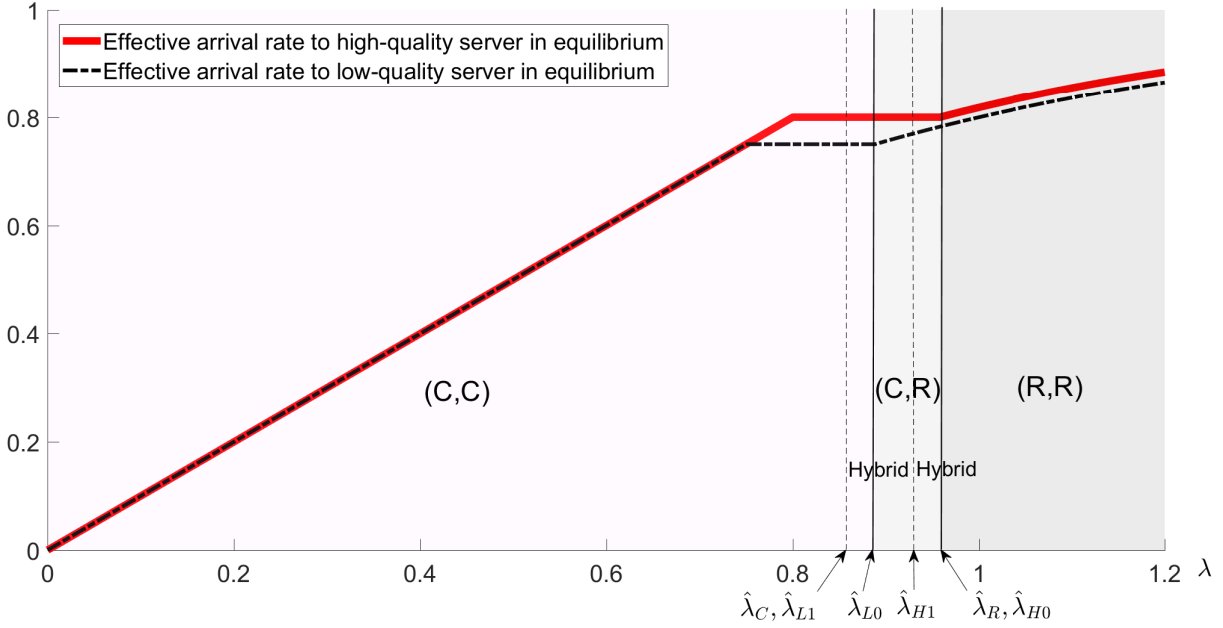


Figure 3.4: Sequential equilibrium outcome and effective arrival rates in equilibrium: $V_H = 2.5$, $V_L = 2$, $\mu = 1$, $\theta = 0.5$, $\delta = 0.5$ and $q = 0.5$

B.4. First, consider the hybrid strategy $f(R|H) = 1$ and $0 < f(R|L) < 1$. By Bayes' rule, the posterior beliefs of uninformed customers are $\delta^C = 0$ and $\delta^R = \frac{\delta}{\delta + (1-\delta)f(R|L)} \in (\delta, 1)$. When $\lambda_L^O(\delta^R) = \lambda_L^U(0)$, we find that the high-quality server strictly prefers to deviate from revealing (R) to concealing (C), improving his effective arrival rate from $\lambda_H^O(\delta^R)$ to $\lambda_H^U(0)$. So, this hybrid strategy cannot be sustained as an equilibrium. Similarly, the hybrid strategy $0 < f(R|H) < 1$ and $f(R|L) = 0$ cannot be sustained as an equilibrium. Then, consider the hybrid strategy $f(R|H) = 0$ and $0 < f(R|L) < 1$. Under this strategy, the posterior beliefs of uninformed customers are $\delta^R = 0$ and $\delta^C = \frac{\delta}{\delta + (1-\delta)(1-f(C|L))} \in (\delta, 1)$. Only at the unique crossing point of $\lambda_L^O(0)$ and $\lambda_L^U(\delta^C)$ (i.e., $\lambda = \hat{\lambda}_{L0}$) is the low-quality server indifferent between R and C. It can be verified that the high-quality server has no incentive to deviate at this crossing point, and thus this hybrid strategy can be sustained at $\lambda = \hat{\lambda}_{L0}$. Similarly, we can show that the hybrid strategy $0 < f(R|H) < 1$ and $f(R|L) = 1$ can be sustained as an equilibrium only at $\lambda = \hat{\lambda}_{H0}$; see Figure 3.4 for the illustration of the existence of the hybrid-strategy sequential equilibrium. Finally, the mixed strategy with $0 < f(R|H) < 1$ and $0 < f(R|L) < 1$ requires that $\lambda_t^O(\delta^R) = \lambda_t^U(\delta^C)$, $t = H, L$. It can be verified that the mixed strategy can never be sustained as a sequential equilibrium in this example.

3.5.3 Effects of Using Queue Disclosure as Signal

Here, we conduct a comparison of the system performances with and without using the queue-disclosure action as a signaling device. Since the hybrid and mixed sequential equilibria can only be numerically identified and their existence ranges of λ are rather limited (see Example 3.2 as an illustration), we only consider the pure-strategy sequential equilibria in this subsection. In the non-signaling case, uninformed customers make their joining decisions based on their prior beliefs when the queue is concealed; while when the queue is revealed, they adopt the ‘hole-avoiding’ strategy discussed in Debo et al. (2012) by utilizing the queue-length information. Anticipating customers’ joining decisions, the server then makes his queue disclosure decision. Specifically, the t -type ($t = H, L$) server conceals the queue if $\lambda_t^U(\delta) \geq \lambda_t^O(\delta)$ and reveals it if $\lambda_t^U(\delta) < \lambda_t^O(\delta)$. By comparing the equilibrium outcomes in the signaling and non-signaling cases, we can obtain the following results regarding the effective arrival rates.

Proposition 3.7. *When the market is composed of both informed and uninformed customers,*

- (i) *the equilibrium effective arrival rates of both types of the server under the signaling case equal the corresponding maximal ones under the non-signaling case if the potential arrival rate λ is either smaller than $\hat{\lambda}_C$ or larger than $\hat{\lambda}_R$.*
- (ii) *for $\lambda \in [\hat{\lambda}_C, \hat{\lambda}_R]$, when a separating sequential equilibrium can be sustained, the maximal effective arrival rate to the high-quality (resp. low-quality) server under all pure-strategy perfect sequential equilibria is no less (resp. no greater) than the maximal one in the non-signaling case.*

Proposition 3.7 implies that in both small and large-sized markets, the queue revelation and concealment convey no further quality information to customers. Only in a medium-sized market may the signaling mechanism work. This is caused by the existence of separating equilibria in which uninformed customers can fully infer the server’s type and behave the same as informed customers. Such signaling effects lead to a non-larger effective arrival rate for the low-quality server and a non-smaller effective arrival rate for the high-quality server compared to those in the non-signaling case.

We then turn to customers’ total utility and investigate the impact of separating equilibria on it. In a revealed queue with the belief of uninformed customers as δ^R , the total utility of

all customers from the type- t ($t \in \mathbb{T}$) server can be derived as

$$u_t^O(\delta^R) = \sum_{i=0}^{n(1)} \lambda_{i,t} \pi_{i,t} \left(V_t - \frac{(i+1)\theta}{\mu} \right). \quad (3.4)$$

Similarly, in a revealed queue with the belief of uninformed customers as δ^C , the total utility of all customers from the type- t ($t \in \mathbb{T}$) server can be written as

$$u_t^U(\delta^C) = \lambda_t^U(\delta^C) \left(V_t - \frac{\theta}{\mu - \lambda_t^U(\delta^C)} \right). \quad (3.5)$$

The following proposition shows that the separating equilibria benefit customers' total utility from the low-quality server. Since multiple pure-strategy sequential equilibria may be sustained at the same time, when we mention customers' total utility from the high-quality (resp. low-quality) server in the signaling case, we consider by default the pure-strategy perfect sequential equilibrium where the high-quality (resp. low-quality) server obtains the maximal effective arrival rate, which is consistent with Proposition 3.7(ii).

Proposition 3.8. *When a separating sequential equilibrium can be sustained, the total utility of all customers from the low-quality server in the signaling case is no less than the one in the non-signaling case.*

Intuitively, under the separating sequential equilibria, all customers become negatively informed when the server is of low quality. As the expected quality level of uninformed customers is higher than the low quality level V_L , the effective joining rate of the low-quality server with heterogeneous customers is larger than the one with only negatively informed customers. Such an overcrowded queueing system makes the actual utility of some uninformed customers negative, which can be improved to be nonnegative after uninformed customers infer the true quality type according to the separating sequential equilibria.

Then, consider the high-quality server. Since the expected quality level of uninformed customers is lower than the high quality level V_H , the effective joining rate of the high-quality server with heterogeneous customers is smaller than the one with only positively informed customers. However, a larger effective joining rate in the signaling case cannot definitely yield a benefit or loss for customers from the high-quality server. On one hand, it increases the workload of the queueing system, which is a driving force of reducing utility. On the other hand, the joining probability of uninformed customers may be properly increased in a concealed queue or at a not-that-long queue length in a revealed queue, which leads to a

counter driving force of increasing utility. Additionally, the uncertain relationship between the total utilities from the high-quality server may be caused by different queue disclosure actions in two cases.

Next, we use two numerical examples to illustrate all above comparison results on the effective arrival rates and customers' total utility.

Example 3.3. Consider the parameter values to be $V_H = 4$, $V_L = 1$, $\mu = 2$, $\theta = 1$, $\delta = 0.25$, and $q = 0.3$. The values of the key points in Figure 3.5 are $\hat{\lambda}_C = \hat{\lambda}_{L1} (= 1.0749) < \hat{\lambda}_{L\delta} (= 1.1224) < \hat{\lambda}_{L0} (= 1.2361) < \hat{\lambda}_{H\delta} (= 2.3208) < \hat{\lambda}'_{H\delta} (= 2.3429) < \hat{\lambda}_R = \hat{\lambda}_{H0} (= 3.2997)$.




Figure 3.5: Comparisons of the maximal effective arrival rates to the high-quality (resp. low-quality) server, λ_H^{non} and λ_H^{sig} (resp. λ_L^{non} and λ_L^{sig}), and the corresponding customers' total utilities from the high-quality (resp. low-quality) server, u_H^{non} and u_H^{sig} (resp. u_L^{non} and u_L^{sig}) in the non-signaling and signaling cases: $V_H = 4$, $V_L = 1$, $\mu = 2$, $\theta = 1$, $\delta = 0.25$ and $q = 0.3$

In the non-signaling case, the high-quality server conceals (resp. reveals) the queue when $\lambda \leq \hat{\lambda}_{H\delta}$ (resp. $\lambda > \hat{\lambda}_{H\delta}$), and the low-quality server conceals (resp. reveals) the queue when $\lambda \leq \hat{\lambda}_{L\delta}$ (resp. $\lambda > \hat{\lambda}_{L\delta}$). In the signaling case, by Propositions 3.5 and 3.7(i), we know that only a pooling strategy can be sustained as a sequential equilibrium when $\lambda < \hat{\lambda}_C$ and $\lambda > \hat{\lambda}_R$, and the equilibrium effective arrival rates to both types of the server remain unchanged regardless of whether the queue disclosure action is used as a signaling device or not; see Figure 3.5. When $\hat{\lambda}_{L0} \leq \lambda \leq \hat{\lambda}_{H0}$, the separating sequential equilibrium (C, R) can be sustained as an equilibrium outcome. In the subrange $\hat{\lambda}_{L0} < \lambda < \hat{\lambda}'_{H\delta}$, (C, R) is the unique pure-strategy perfect sequential equilibrium, and we have $\lambda_H^U(1) > \lambda_H^O(\delta)$ and $\lambda_H^U(1) \geq \lambda_H^U(\delta)$;

while in the subrange $\hat{\lambda}'_{H\delta} < \lambda \leq \hat{\lambda}_{H0}$, we get that $\lambda_H^O(\delta) > \lambda_H^U(1)$, and the pooling strategy (R, R) can also be sustained as a perfect sequential equilibrium with the off-equilibrium-path belief $\delta^C \in [0, 1]$. Therefore, for $\hat{\lambda}_{L0} \leq \lambda \leq \hat{\lambda}_{H0}$, signaling through the queue disclosure can make the high-quality server better off and the low-quality server worse off considering all pure-strategy perfect sequential equilibria, which is consistent with Proposition 3.7(ii). For the remaining range $\hat{\lambda}_C \leq \lambda < \hat{\lambda}_{L0}$, only the pooling sequential equilibrium (C, C) can be sustained with the off-equilibrium-path belief δ^R satisfying $\lambda_L^O(\delta^R) \leq \lambda_L^U(\delta)$ (e.g., $\delta^R = 0$), which is also a perfect sequential equilibrium because the credible updating rule puts no restriction on the off-equilibrium-path belief δ^R . Now, the high-quality server conceals the queue under both the signaling and non-signaling cases, and thus the optimal effective arrival rates to the high-quality server under two cases keep the same. This observation holds for the low quality server when $\hat{\lambda}_C \leq \lambda \leq \hat{\lambda}_{L\delta}$. However, for $\hat{\lambda}_{L\delta} < \lambda < \hat{\lambda}_{L0}$, although the signaling effect does not change the belief of uninformed customers, the effective arrival rate to the low-quality server in the signaling case becomes strictly smaller than the one under the non-signaling case due to different queue-disclosure actions: without considering the signaling effect, the low-quality server reveals the queue, while at this moment, (C, C) is the unique pure-strategy perfect sequential equilibrium in the signaling case.

Regarding the customers' total utility from the low-quality server, for $\hat{\lambda}_{L0} \leq \lambda \leq \hat{\lambda}_{H0}$ where the separating sequential equilibrium (C, R) exists, we can see that it becomes weakly larger in the signaling case than the one in the non-signaling case, which echoes Proposition 3.8. But the total utility from the high-quality server in the signaling case becomes weakly smaller than the one in the non-signaling case. Intuitively, for $\hat{\lambda}_{L0} \leq \lambda \leq \hat{\lambda}_{H\delta}$, concealing is the choice of the high-quality server in both the signaling and non-signaling cases, but the larger effective arrival rate under the signaling case decreases the total expected utility from a concealed queue; and in the subrange $\hat{\lambda}_{H\delta} < \lambda \leq \hat{\lambda}'_{H\delta}$, the high-quality server conceals (resp. reveals) the queue in the signaling (resp. non-signaling) case, which generates a zero (resp. positive) total utility. Note that a relatively larger effective arrival rate in a concealed queue does not always decrease the total utility from the high-quality server.^{3.6} Another interesting observation is that for $\hat{\lambda}_{L\delta} < \lambda < \hat{\lambda}_{L0}$, the total utility from the low-quality server in the signaling case becomes strictly less than the one under the non-signaling case, which is caused

^{3.6}For example, when the parameter values are $V_H = 10$, $V_L = 1$, $\mu = 2$, $\theta = 1$, $\delta = 0.01$ and $q = 0.3$, at the market size $\lambda = 1.5450$, the high-quality server conceals the queue in both the signaling and non-signaling case with the effective arrival rates as $\lambda_H^U(1) = 1.5451$ and $\lambda_H^U(\delta) = 1.5370$, respectively. And the corresponding total utilities are $u_H^U(1) = 12.0544$ and $u_H^U(\delta) = 12.0503$, respectively. In this case, $u_H^U(1) > u_H^U(\delta)$.

by different queue-disclosure actions in two cases as mentioned before.

In Example 3.3, only one kind of the separating sequential equilibria, (C, R) , appears. Next, let us turn to another example where another separating sequential equilibrium (R, C) exists, and mainly see what impacts (R, C) can cause to the system performances.

Example 3.4. Consider the parameter values as $V_H = 1.01$, $V_L = 0.91$, $\mu = 2$, $\theta = 1$, $\delta = 0.3$, and $q = 0.9$. Figure 3.6 shows the comparison results, where the values of the key points are $\hat{\lambda}_C = \hat{\lambda}_{H1} (= 1.2539) < \hat{\lambda}_{H0} (= 1.2905) < \hat{\lambda}_{L1} (= 1.5239) < \hat{\lambda}_R = \hat{\lambda}_{L0} (= 1.6400)$. In the non-signaling case, the high-quality server conceals (resp. reveals) the queue when $\lambda \leq \hat{\lambda}_{H0}$ (resp. $\lambda > \hat{\lambda}_{H0}$), and the low-quality server conceals (resp. reveals) the queue when $\lambda \leq \hat{\lambda}_{L0}$ (resp. $\lambda > \hat{\lambda}_{L0}$). In the signaling case, only a pooling strategy can be sustained as a sequential equilibrium when $\lambda < \hat{\lambda}_C$ and $\lambda > \hat{\lambda}_R$; for $\hat{\lambda}_{H1} < \lambda < \hat{\lambda}_{L1}$, (R, C) is the unique pure-strategy perfect sequential equilibrium; and for $\hat{\lambda}_{L1} < \lambda < \hat{\lambda}_{L0}$, no (pure, hybrid or mixed) sequential equilibrium exists even though all possibilities of customers' equilibrium queueing strategies are considered. The classic results show that the sequential equilibria exist for every finite extensive game (see Selten, 1975; and Kreps and Wilson, 1982). However, since the players in our game involve infinite customers, our signaling game is not a finite one, and thus the existence of sequential equilibria cannot be guaranteed.^{3.7} Even so, we have provided some key definite results on the existence of the sequential equilibria (see Propositions 3.1, 3.5 and 3.6).

Here, we only focus on the effect caused by the separating sequential equilibrium (R, C) for $\hat{\lambda}_{H1} \leq \lambda \leq \hat{\lambda}_{L1}$. In both the signaling and non-signaling cases, the low-quality server conceals the queue. Since only a few of customers are uninformed (i.e., $1 - q = 0.1$), the overall performance of the queueing system is not affected by the different beliefs in two cases, and thus the effective arrival rates to the low-quality server keep the same. Also, the total utilities of all customers from the low-quality server in two cases are equal. Then, consider the high-quality server. As the uninformed customers become fully informed in the signaling case, the effective arrival rate to it is strictly larger than the one in the non-signaling case. For $\hat{\lambda}_{H0} \leq \lambda \leq \hat{\lambda}_{L1}$, the high-quality server reveals the queue in both cases. And the larger effective arrival rate in the signaling case decreases the total utility. Then, for $\hat{\lambda}_{H1} \leq \lambda < \hat{\lambda}_{H0}$, the high-quality server conceals the queue in the non-signaling case, which makes the total utility as 0, but the total utility in the signaling case is strictly larger than 0 in the signaling case because the high-quality server now reveals the queue.

^{3.7}The Nash equilibrium, which is weaker than the sequential equilibrium, must exist in our signaling game. For example, the Nash equilibria where the server adopts the pooling strategy (C, C) or (R, R) can always be sustained for all $\lambda \in (0, +\infty)$ (see the proof in Appendix B.5).




Figure 3.6: Comparisons of the maximal effective arrival rates to the high-quality (resp. low-quality) server, λ_H^{non} and λ_H^{sig} (resp. λ_L^{non} and λ_L^{sig}), and the corresponding customers' total utilities from the high-quality (resp. low-quality) server, u_H^{non} and u_H^{sig} (resp. u_L^{non} and u_L^{sig}) in the non-signaling and signaling cases: $V_H = 1.01$, $V_L = 0.91$, $\mu = 2$, $\theta = 1$, $\delta = 0.3$ and $q = 0.9$

3.6 Discussions

Section 3.5.3 shows that the signaling effect of the queue disclosure action influences the system performances only under separating sequential equilibria in a medium-sized market. Through helping the uninformed customers to identify the true quality type, it generates a larger effective arrival rate to the high-quality server and improves the customers' total utility in case of low quality. In this section, we focus on the existence of the separating sequential equilibrium and investigate how the customer type composition and service price affect it.

3.6.1 The Impact of Customer Type Composition

A close look at the results stated in §3.4.2 and §3.5.2 reveals that the existence of a separating equilibrium in our signaling game requires the coexistence of informed and uninformed customers in the market. This makes us wonder *whether increasing q , the proportion of informed customers in the market, can induce the separating sequential equilibria to occur more likely*. We examine this question in this subsection.

Proposition 3.6 depicts the exact existence ranges of the separating sequential equilibria, and basically they are determined by eight effective-arrival-rate functions (i.e., $\lambda_t^U(d)$ and

$\lambda_t^O(d)$ with $t \in \{H, L\}$ and $d \in \{0, 1\}$). Then, the impact of q on the separating sequential equilibria can be equivalently converted to its impact on these eight functions. Next, we investigate its impact under two cases.

Case 1. *The fraction of informed customers is large enough (i.e., $q \geq \hat{q} := 1 - \frac{\mu - \theta/V_L}{\mu - \theta/V_H}$).*

In this case, a small number of uninformed customers do not affect the resulting effective arrival rates in concealed queues (see Appendix B.2), and thus $\lambda_t^U(d)$ ($t \in \{H, L\}$ and $d \in \{0, 1\}$) is independent of q , which implies that the impact of q is exerted through changing the effective arrival rates to revealed queues.

Then, let us first consider the separating sequential equilibrium (C, R) . Under (C, R) , the uninformed customers hold a posterior belief $\delta^C = 1$ when seeing a concealed queue and $\delta^R = 0$ when seeing a revealed queue. Then, all customers become negatively informed faced with a revealed queue, and thus the resulting effective arrival rate $\lambda_L^O(0)$ does not change as q varies, which means that a change in q does not affect the low-quality server's incentive to stay at R . However, as more customers become informed, the effective arrival rate to a revealed queue in case of high quality server (i.e., $\lambda_H^O(0)$) becomes larger, and this increases the high-quality server's incentive to deviate from C to R , making the equilibrium (C, R) less likely to be sustained. Therefore, the range for which the separating sequential equilibrium (C, R) can be sustained, if exists, becomes smaller as q increases.

Similarly, we can analyze the separating sequential equilibrium (R, C) . The uninformed customers update the belief as $\delta^C = 0$ when seeing a concealed queue and $\delta^R = 1$ when seeing a revealed queue. Then, all customers become positively informed faced with a revealed queue, and thus the resulting effective arrival rate $\lambda_H^O(1)$ is irrelevant to q , implying that changing q does not affect the high-quality server's incentive to stay at R . By contrast, as more customers become informed, the effective arrival rate to a revealed queue in case of low quality server (i.e., $\lambda_L^O(1)$) becomes smaller. Hence, a larger value of q reduces the low-quality server's incentive of mimicking the high-quality server's behavior, making the separating equilibrium (R, C) more likely to be sustained. Therefore, the range for which the separating sequential equilibrium (R, C) can be sustained, if exists, becomes larger as q increases.

Case 2. *The fraction of informed customers is small enough (i.e., $q < \hat{q}$).*

In this case, a substantial amount of uninformed customers play a crucial role in the resulting effective arrival rates to concealed queues, making $\lambda_t^U(d)$ ($t \in \{H, L\}$ and $d \in \{0, 1\}$) dependent on q . This leads to an uncertain relationship between the existence of the separating sequential equilibria and the fraction of informed customers q . For example, let us consider the separating sequential equilibrium (C, R) . As q increases, for the high-quality server, the increasing number of positively informed customers improves the effective arrival rate to a revealed queue (i.e., $\lambda_H^O(0)$), and thus the high-quality server has a stronger incentive to deviate from C to R , causing a driving force of narrowing the existence ranges. On the other hand, according to Appendix B.2, as more customers are negatively informed, the effective arrival rate to the concealed queue in case of low-quality server becomes weakly smaller (i.e., $\lambda_L^U(1)$ is nonincreasing in q). This reduces the low-quality server's incentive of mimicking the high-quality server's behavior, leading to a driving force of expanding the existence ranges. Taken altogether, increasing q changes both high- and low-quality servers' queue-disclosure incentives, and causes a pair of counter driving forces that determine the existence ranges of (C, R) . So, there is no monotonic relationship between q and the existence ranges of the separating sequential equilibrium (C, R) . Similarly, it can be verified that such an uncertain relationship applies to the separating sequential equilibrium (R, C) .

We now use the following numeric example to illustrate the impact of q on the occurrence of separating sequential equilibria.




Figure 3.7: Impact of q on the ranges of λ where separating sequential equilibria exist: (a) $V_H = 4$, $V_L = 1$, $\mu = 2$, $\theta = 1$; (b) $V_H = 1.01$, $V_L = 0.91$, $\mu = 2$, $\theta = 1$

Example 3.5. We use two different sets of parameter values to illustrate the impact of q on the existence ranges of (C, R) and (R, C) separately. First, consider the same set of parameter values used in Example 3.3 except that now we vary the value of q from 0 to 1. In this case, if $q < \hat{q} := 0.4286$, $(\mu - \theta/V_L)/(1 - q) < \mu - \theta/V_H$; otherwise, $(\mu - \theta/V_L)/(1 - q) \geq \mu - \theta/V_H$. Figure 3.7(a) depicts the change of the range(s) of the potential arrival rate λ where the separating sequential equilibrium (C, R) exists as the fraction of informed customers q increases. From Figure 3.7(a), we can see that for $q \in [0, \hat{q})$, q has a non-monotonic impact on the occurrence of the separating sequential equilibrium (C, R) : when $q \leq 0.2899$, the existence ranges expand as q increases, where the fishtail shape is caused by the bulge shape of $\lambda_L^U(1)$ (see Figure 3.3 for illustration); and when $0.2899 < q < \hat{q}$, the existence ranges narrow as q increases. Once q surpasses the threshold \hat{q} , further increasing q surely makes (C, R) less likely to appear.

Then, consider parameter values used in Example 3.4 except that we vary the value of q from 0 to 1. Under this scenario, the key threshold of q becomes $\hat{q} := 0.1077$. Figure 3.7(b) shows that the separating sequential equilibrium (R, C) exists only when $q > \hat{q}$. And as q increases, the existence range of the separating sequential equilibrium (R, C) monotonically expands.

3.6.2 The Impact of Service Price

Until now, the service price is normalized to one, and customers' monetary rewards gained from service (i.e., V_H and V_L) are indeed the service value minus the service price. In this subsection, we investigate the impact of the service price p , which is exogenously given. The service value of the high-quality (resp. low-quality) server is denoted as \mathcal{V}_H (resp. \mathcal{V}_L), and then a customer receives a monetary reward $V_H = \mathcal{V}_H - p$ (resp. $V_L = \mathcal{V}_L - p$) if she is served by the high-quality (resp. low-quality) server. Next, we study the relationship between p and the existence ranges of the separating sequential equilibria. Note that since we require that $V_L > \frac{\theta}{\mu}$, the service price should satisfy that $0 \leq p < \mathcal{V}_L - \frac{\theta}{\mu}$.

Similarly to section 3.6.1, the impact of p on the separating sequential equilibria can be equivalently converted to its impact on eight effective arrival rates (i.e., $\lambda_t^U(d)$ and $\lambda_t^O(d)$ with $t \in \{H, L\}$ and $d \in \{0, 1\}$). It can be easily verified that these effective arrival rates are all decreasing in p . Intuitively, the higher the service price p is, the lower the monetary rewards that customers gain from service become, and thus the less motivated the customers are to join the queue. There is a difference in the decreasing patterns of eight effective arrival rates:

the decreasing of $\lambda_t^U(d)$ ($t \in \{H, L\}$ and $d \in \{0, 1\}$) happens in a continuous way, while $\lambda_t^O(d)$ ($t \in \{H, L\}$ and $d \in \{0, 1\}$) keeps piecewise constant due to the floor function in $n(0)$ and $n(1)$ and only decreases (or down jumps) at several threshold values of p at which $\frac{V_t\mu}{\theta}$ ($t \in \{H, L\}$) takes integer values. Then, the impact of p on the equilibrium outcome can be analyzed in the following two cases.

Case 1. *For the ranges of the service price p on which both $n(0)$ and $n(1)$ keep unchanged, $\lambda_t^O(d)$ ($t \in \{H, L\}$ and $d \in \{0, 1\}$) is constant in p , and thus the changes of the existence ranges of the separating sequential equilibria are induced by the decreasing of $\lambda_t^U(d)$ in p ($t \in \{H, L\}$ and $d \in \{0, 1\}$). Such decreases result in the existence ranges of two separating sequential equilibria shifting towards smaller λ . Intuitively, for the separating sequential equilibrium (C, R) (resp. (R, C)), the increasing p attracts fewer and fewer customers to join the concealed queue. This makes the low-quality (resp. high-quality) server more likely to reveal the queue at a smaller market size λ . By contrast, the high-quality (resp. low-quality) server has less incentive to conceal the queue at a larger market size λ . These two factors together lead to the down shifting of the existence ranges of the separating sequential equilibrium (C, R) (resp. (R, C)).*

Case 2. *At those threshold values of p where $\frac{V_H\mu}{\theta}$ or $\frac{V_L\mu}{\theta}$ takes integer values, $\lambda_t^U(d)$ ($t \in \{H, L\}$ and $d \in \{0, 1\}$) can be regarded as constant, but $\lambda_t^O(d)$ ($t \in \{H, L\}$ and $d \in \{0, 1\}$) jumps down as p increases across these thresholds. This makes each existence range of separating sequential equilibrium, if still exists at these values of p , shifts towards larger λ .^{3.8} The behind mechanism can be intuitively understood following the analysis in Case 1.*

We now use the following numeric example to illustrate the impact of p on the occurrence of separating sequential equilibria.

Example 3.6. *In this example, we use two different sets of parameter values to illustrate the impact of the service price p on the existence ranges of (C, R) and (R, C) separately. First, consider $\mathcal{V}_H = 6$, $\mathcal{V}_L = 3$, $\mu = 2$, $\theta = 1$ and $q = 0.5$. We vary the value of p from 0 to 2.5^- . Figure 3.8(a) shows that the separating sequential equilibrium (C, R) always exists for all $p \in [0, 2.5)$ and how its existence ranges change as the service price p increases. We*

^{3.8}The bulge shape of $\lambda_L^U(1)$ (see Figure 3.3 for illustration) may narrow some subrange of λ where (C, R) can be sustained as a separating sequential equilibrium. But the direct effect of these threshold values of p is still to raise up the existence ranges. So, we do not mention the special change caused by the bulge shape when analyzing the impact of p .




Figure 3.8: Impact of p on the ranges of λ where separating sequential equilibria exist: (a) $\mathcal{V}_H = 6$, $\mathcal{V}_L = 3$, $\mu = 2$, $\theta = 1$, $q = 0.5$; (b) $\mathcal{V}_H = 5$, $\mathcal{V}_L = 4.9$, $\mu = 2$, $\theta = 1$, $q = 0.9$

can see that the changing patterns mentioned in Cases 1 and 2 happen. For example, for $p \in (1, 1.5]$, we have that $n(0) = 3$ and $n(1) = 9$, and the shadowed range shifts downwards as p increases. When p increases a little bit to be larger than 1.5, $n(0)$ and $n(1)$ decrease to 2 and 8, respectively, and the shadowed range jumps up at $p = 1.5^+$.

Then, consider $\mathcal{V}_H = 5$, $\mathcal{V}_L = 4.9$, $\mu = 2$, $\theta = 1$ and $q = 0.9$. We vary the value of p from 0 to 4.4. Figure 3.8(b) shows that the separating sequential equilibrium (R, C) only appears when the service price p falls into several separate intervals, and on each of these continuous intervals of p , the shadowed existence range shifts downwards as p increases. For example, for $p \in (p_1, 2]$ with $p_1 = 1.9$, we have that $n(0) = 5$ and $n(1) = 6$, and the shadowed range shifts downwards as p increases. When p is slightly larger than 2, the separating sequential equilibrium (R, C) does not hold until p increases to be slightly larger than $p_2 := 2.4$ where $n(0) = 4$ and $n(1) = 5$.

3.7 Conclusions and Suggestions for Future Research

In many service systems, the service quality is unknown to some incoming customers. Uninformed customers often gather quality information through multiple sources, among which a simple way is to inspect the queue length. The queue length can convey some quality information because it contains some information about those informed customers' behavior. Analysis of customers' equilibrium queueing strategy in observable queues with unknown service quality has been well done in Debo et al. (2012). One can further think about the

following question: why does a server allow his queue to be observable/unobservable? To examine this question, we study a signaling game for the server with the queue disclosure action as a signaling device. In our model, customers in an observable queue obtain the service quality information not only through inspecting the queue length but also by considering the server's incentives on the queue disclosure actions. For this signaling game, we investigate its sequential equilibria and adopt the perfect sequential equilibrium concept as further refinement of the equilibrium concept whenever needed.

Specifically, we consider two scenarios, a basic one with only uninformed customers and a general one with both informed and uninformed customers existing in the system. Our major takeaway is that *a separating equilibrium exists only when the market size is moderate and the system has both informed and uninformed customers*. This has multiple implications. One, in a circumstance where all customers are uninformed, the pooling equilibrium dominates other equilibria, and thus the queue disclosure action itself conveys no valuable quality information. Two, considering heterogeneous customers, when the market size is very small, both types of the server tend to conceal their queues. Hence, customers cannot infer service quality from their queue concealment behavior. Similarly, if the market size is very large, both types of the server tend to reveal their queues and hence the quality cannot be inferred there either. Three, in a circumstance where a separating equilibrium prevails, the queue disclosure action fully conveys the service quality information to uninformed customers. Consequently, uninformed customers behave exactly the same as informed ones in determining their queueing strategies. Furthermore, the effective arrival rate of the high-quality (resp. low-quality) server is weakly larger (resp. smaller) than that without considering the queue disclosure action as a signal of service quality, and the total utility of all customers from the low-quality server can be improved considering such a signaling effect.

In our study, the server uses his queue disclosure action as a signaling device. In reality, a server can also signal his quality information through other devices such as price; see Debo et al. (2020) for the study on this. It would be interesting to consider a signaling game in which the server employs price and queue disclosure action jointly to signal his quality. We leave it for future research.

Chapter 4

Optimal Queue Length Information Disclosure When Service Quality Is Uncertain

Let us still follow the setting in Chapter 3 and consider that all customers are uninformed. From the results in the signaling game, we know that with all customers uninformed, the server's queue disclosure strategy signals no quality information, and the customers make the joining decision still based on the prior belief.

In reality, customers can tolerate a higher level of congestion when facing a higher level of service quality. By taking this into account, when no other quality signal is available, the server can tailor-make its queue-disclosure strategy, according to the realized service quality, to attract more customers to join. Now, consider that before the quality type is realized, the server owns a commitment power that enables him to design and commit to an ex-ante queue-disclosure policy that states whether or not the queue length will be revealed to customers upon their arrival, given a realized quality level. Then, some quality information can be inferred from in the server's queue disclosure action. Can such a commitment strategy persuade more customers to join compared with the traditional revealing-or-concealing paradigm? We formulate this problem as a Bayesian persuasion model, and use a graphical geometric approach to solve it following the one in Kamenica and Gentzkow (2011). Such a geometric approach can be applied to various scenarios such as profit maximizer and social planner.

4.1 Introduction

Whether queue length information should be provided to customers, who in turn decide whether or not to join the queue, is a classic research topic. It is well documented in the queueing game literature (e.g., Hassin and Haviv, 2003, p. 51) that there exists a threshold on the arrival rate below which the server conceals the queue length and above which he reveals it. The underlying reason is that when the potential arrival rate is low, all customers join a concealed queue due to their positive expected utility, while a disclosed queue always comes with a chance, albeit small, for blocking away some incoming customers who face a long queue upon their arrival. By contrast, when the potential arrival rate is sufficiently high, a concealed queue's effective arrival rate is fixed to the one that comes with an expected utility of zero in case of joining, while that of a revealed queue always increases with the potential arrival rate as some customers, albeit a small fraction, may nevertheless face a short queue upon their arrival.

The foregoing queue-disclosure strategy is based on a setting with known service quality. In real practice, service quality provided by service providers may be uncertain. For example, the food quality of restaurants may be uncertain as it is affected by factors such as the ingredients' freshness and the chefs' skill. The service quality of online consulting, like online healthcare diagnosis and telephone hotline, heavily relies on the skills and expertise levels of consultants/agents, which are uncertain to customers particularly when consultants/agents are taking rotations in their schedules. Our research question is: in such service systems with uncertain service quality, how shall the server conduct his queue-disclosure strategy?

A simple way is that the server commits to fully revealing or concealing the queue regardless of the realized service quality. Under such a commitment, customers cannot infer any quality information from the server's queue-disclosure action and they have to make their joining-or-balking decisions based on their prior beliefs. Consequently, the server's ex-ante commitment problem degenerates to the one without commitment, as illustrated in the following example.

Example 4.1. (The Decoupling of Queue-Disclosure Strategy from Service Quality) *A server (he) provides some service with uncertain quality level. Nature decides whether the service quality is high with a value of 2 or low with a value of 1, according to a Bernoulli trial with probabilities 0.33 and 0.67, respectively. Customers' service times follow an exponential distribution with rate 1.1. Potential customers (she) arrive according to a Poisson*

process with rate 0.5. The unit time waiting cost is 1. Both the server and the customers hold the same prior belief regarding the service quality.

When the queue length is always revealed to customers, customers' belief on the service quality is the expected one 1.33. Under this scenario, customers join the observable queue if and only if the queue is empty, and the effective arrival rate can be calculated to be 0.3438. Similarly, when the queue is always concealed from customers, customers' belief on the quality is still the expected one 1.33; the effective arrival rate can be derived by setting the customer joining utility to be zero (since the potential arrival rate is high enough), which yields an equilibrium arrival rate of 0.3481. Thus, the optimal strategy for the server is to conceal the queue length.

In the foregoing example, the queue-disclosure strategy is decoupled from the realized quality level. As such, customers have to make their queueing decisions based on the expected service quality. Perhaps, the server can persuade more customers to join the system by linking the queue-disclosure strategy with the realized service quality. Speaking mathematically, the server can design a queue disclosure strategy characterized by two conditional queue-disclosure probabilities, $\pi(\cdot|high)$ and $\pi(\cdot|low)$, which correspond to the realized service quality being high or low, respectively. The server then commits to it before the realization of service quality. Clearly, if the server commits to a strategy with $\pi(concealing|high) = 1$ and $\pi(revealing|low) = 1$, incoming customers can exactly infer the service quality by the visibility of the queue and thus do not rely on the quality expectation to make their queueing decision. In fact, one can show that the above quality-dependent queue-disclosure strategy can improve the expected effective arrival rate to 0.3953, a 13.56% improvement over the one when the server always conceals the queue.

As we will show in Section 4.4.1, the optimal disclosure strategy for the server under the setting given in Example 4.1 shall be

$$\pi(revealing|high) = 0, \quad \pi(revealing|low) = 0.7537;$$

$$\pi(concealing|high) = 1, \quad \pi(concealing|low) = 0.2463.$$

That is, when the service quality turns out to be low, instead of fully disclosing the queue length to the customers (as discussed above), the server randomizes queue length disclosure and concealment with the probability of revealing being 0.7537. Then, with probability 0.5050 (resp. 0.4950), customers see a revealed (resp. concealed) queue. Thus, the posterior

probability for the service quality being high becomes 0 (resp. 0.6667), leading to an effective arrival rate of 0.3438 (resp. 0.5). As a result, the effective arrival rate of a concealed queue is still 0.5 but the probability of its occurrence increases from 0.33 to 0.4950. The expected effective arrival rate is now 0.4211, a further 6.53% improvement over the above one under the full quality revelation strategy. This example provides two important insights for service providers. One, for a server with uncertain service quality, it is better to link the queue-disclosure strategy with the realized quality level. Two, a randomization strategy on queue disclosure may make the server better off as customers cannot fully infer the realized quality.

The above example indicates that the quality-linked queue-disclosure strategy can be beneficial to the server. However, it requires the server to pre-commit to his queue-disclosure strategy and, once the quality level is realized, the corresponding queue-disclosure action has to be performed without manipulations. One way to interpret such pre-commitment towards the randomized queue-disclosure strategy is to consider a “long-run” server who aims to maximize his long-run average profit when facing “short-run” customers (see, e.g., Rayo and Segal, 2010). Customers can then infer the server’s randomization strategy from their long-term experiences/observations of the server’s queue-disclosure actions. Note that with the advancement in information technology, it becomes relatively easy for service providers to change the visibility status of their queue. For example, through turning on or off the display screens or through controlling the provision of the real-time queue information on online platforms or mobile apps such as Dianping.com and Yelp.com, the queue length information can be revealed to or concealed from customers.

In this study, our main target is to illustrate the underlying mechanism why a randomized quality-linked queue-disclosure strategy can yield a larger effective arrival rate for the server. We also provide an approach to find such a strategy. The quality-linked queue-disclosure strategy considered in our study can help service providers to persuade more customers to join their system. The well-constructed quality-linked queue-disclosure strategy can also help a social planner to better regulate customer arrivals.

Specifically, we consider a stylized single-server service system. Customers arrive according to a Poisson process and service times are exponentially distributed. The server’s service quality, however, is random and takes the value of either high type (labeled as h) or low type (labeled as l). The server can observe the realized quality but customers cannot. The probability for the quality being high is common knowledge and hence is customers’ prior belief about service quality. Customers are homogeneous: they have the same prior knowl-

edge, receive the same service reward, and incur the same unit-time delay cost. Before service quality is realized, the server announces his queue-disclosure strategy, characterized by conditional probabilities of disclosing the queue length given each type of realized service quality, and commits to it. Once the service quality is realized, the corresponding queue-disclosure action is performed. Based on the visibility of the queue, customers update their beliefs about the service quality according to Bayes' rule and then make their joining-or-balking decisions accordingly to maximize their utilities.

By Bayes' rule, the expected posterior probability regarding the server's service quality equals its prior. In reverse, if a distribution of posteriors satisfies this property, it is called *Bayes plausible* (see Kamenica and Gentzkow, 2011). We show that any Bayes-plausible distribution of customers' posteriors corresponds to a unique queue disclosure strategy that can induce it. Thus, searching for the optimal disclosure strategy is equivalent to searching for the optimal Bayes-plausible distribution of customers' posteriors. Such a reformulation of the problem provides a useful geometric approach to our aim of deriving the optimal disclosure strategy. First, we can plot the effective arrival rates of the revealed and concealed queues as two functions of the probability for the service quality being high. Next, we demonstrate that any convex combination of the two points from these two functions can be generated through a properly-designed queue-disclosure strategy. We further show that any point on the upper envelope of all the convex combinations represents the maximal effective arrival rate under the corresponding prior. As such, we can graphically determine whether the server can benefit from the randomized queue-disclosure strategy by simply checking whether the upper envelope is strictly above the two effective arrival rate functions.

After the derivation of the optimal disclosure strategy, we turn to examining the impact of market size (i.e., the potential total arrival rate) on the optimal disclosure strategy. We show that when the market size is sufficiently small, the server always conceals the queue length information no matter whether the realized service quality is high or low; however, when the market size is very large, the server always reveals the queue. These two results are consistent with those stated in the literature when the service quality is known (see, e.g., Hassin and Haviv, 2003, p. 51). However, when the market size is medium, we numerically find that it is often optimal for the server to adopt a quality-dependent queue-disclosure strategy, which can help to increase the server's effective arrival rate. Moreover, such a strategy is often hybrid or mixed, that is, randomizing queue disclosure and concealment actions.

We then extend our analysis to a setting where the server acts as a social planner and

aims to maximize the social welfare. We show that we can still apply our geometric approach to find out the best queue-disclosure strategy for the social planner. In contrast to the classic literature result that revealing the queue length is always socially optimal (see, e.g., Hassin and Haviv, 2003; and Hassin and Roet-Green, 2017), we find that when service quality is uncertain, a randomized queue-disclosure strategy can make the social planner better off.

The rest of this chapter is organized as follows. Section 4.2 reviews the related literature. The formal model is presented in Section 4.3. We investigate the optimal queue disclosure strategy in Section 4.4. Section 4.5 examines a situation in which the server is a social planner. Concluding remarks are provided in Section 4.6. All the proofs are relegated to Appendix C.

4.2 Literature Review

Our work is closely related to the studies on quality disclosure. In economics, Grossman (1981) investigates product quality disclosure problems through ex post verifiable disclosure and warranties. Grossman shows that the seller would voluntarily disclose the private information in equilibrium if the disclosure is costless and information is verifiable. Milgrom (1981) characterizes the favorableness of news and introduces the novel persuasion game. Milgrom shows that, in a sales encounter model, the salesman always reports the most favorable data about his product. In operations management, there are some studies on customers' queueing strategy with unknown service quality. Veeraraghavan and Debo (2009, 2011) consider the quality issue in a two-parallel-observable-queue setting. They show that in equilibrium it might be optimal for customers to join a longer queue. Debo et al. (2012) examine customers' queueing strategy when queues are observable and service quality is unknown to certain customers. They conclude that uninformed customers adopt a hole-avoiding strategy; i.e., they do not join when the queue is at certain length (called hole) but otherwise behave in the same way as the positively-informed customers. In these works, customers are heterogeneous, with some customers being informed with the quality information, and thus the queue length can provide some information for the uninformed customers. Different from the above works, in our model, customers are all uninformed and thus the queue length itself cannot convey the quality information. It is the server's disclosure action –concealing or revealing the queue –that provides some information on service quality.

Our work is related to studies on product and service quality revelation by signaling games, including Debo and Veeraraghavan (2014), Kremer and Debo (2016), Yu et al. (2016), Wang

and Özkan-Seely (2018), Debo et al. (2020), Wang and Hu (2020), Guo et al. (2020), etc. The main difference between a quality-signaling game and our queue-disclosure game is the timing difference. In a signaling game, the server signals his type *after* his quality type is realized. However, in our persuasion game, the server commits to a queue-disclosure strategy *before* his quality type is realized. Also, in our game, the server must commit to his pre-determined queue-disclosure strategy once the quality is realized, but there is no such requirement in a signaling game.

The solution technique we are using is closely related to the one used in *Bayesian persuasion* games as introduced in Kamenica and Gentzkow (2011). They study how a sender designs a signaling system and commits to it in order to induce preferred actions from an information receiver. Kamenica and Gentzkow (2011) demonstrate that the concavification of the value function identifies whether the sender benefits from persuasion, but the structure of the optimal signal can be very hard to derive when the state space is large. Gentzkow and Kamenica (2016) show the optimal signal structure of a particular class of Bayesian-persuasion games where the receiver’s optimal action depends only on the expectation of the unknown state and sender’s payoff is independent of the state. Lingenbrink and Iyer (2019) pioneer in introducing the Bayesian persuasion game into a queueing setting. There, the unknown state of the world is only the queue length and they prove that the optimal signaling mechanism is a binary threshold signal that is queue-length-dependent. Different from Lingenbrink and Iyer (2019) where service quality is given, we consider uncertain service quality. Despite this key difference, both Lingenbrink and Iyer (2019) and our work demonstrate that by pre-committing to a queue-disclosure strategy, the server can persuade more customers to join.

Our work is also related to studies on information provision and purchase in queues. Hassin and Haviv (1994) consider a case in which customers arriving at two parallel queues can choose to buy information on queue length at a price so as to join the shorter queue. Hassin (2007) examines a scenario where service quality and some other system parameters are known to the server but not to the customers. The server can choose whether or not to disclose his private information to customers. Hassin and Roet-Green (2017) study information purchase in a one-server queue setting. In their study, incoming customers can buy information on the queue length. Hassin and Roet-Green (2018) consider a setting where customers coming to parallel servers try to deduce from the queue length of one server whether to join this queue or to inspect another queue. Those who have inspected other queues play a role of informed

customers. The fraction of informed customers is not predetermined but rather an artifact of the strategy used by customers. Yang et al. (2019) further study the consumers' search among queues when both the quality and queue length are uncertain.

The study on the impact of delay announcements on queues is also related. Allon et al. (2011) consider a *cheap talk* game between the server and the customers, where the server knows the state of the system and then sends a signal and the customers use the signal to update their belief on the expected waiting time. The difference between the cheap talk game and our ex-ante commitment approach is that there the sender is not committed to a signaling rule. Another key difference is that Allon et al. (2011) do not consider quality issues while we do. Yu et al. (2018) further study such a cheap talk game in a setting with heterogeneous customers and show that customers' response to a delay announcement can be used to elicit information on customer types. Other related works in this stream include Hassin (1986), Whitt (1999), Armony and Maglaras (2004a, 2004b), Burnetas and Economou (2007), Guo and Zipkin (2007), Armony et al. (2009), Guo and Hassin (2011), Yu et al. (2016), Yu et al. (2017), Hu et al. (2018), and Yu et al. (2021), etc. We refer the interested readers to two survey books, Hassin and Haviv (2003) and Hassin (2016), and the survey papers by Aksin et al. (2007) and Ibrahim (2018) for more works in this research stream. Recently, Li et al. (2020) study the optimal queue disclosure strategy but consider that service quality is known. They demonstrate that it is socially optimal to disclose the queue length only if the queue is either very short or very long. Different from them, we consider the server's optimal queue-disclosure strategy when service quality is uncertain and our queue disclosure strategy is service-quality-dependent.

Interestingly, our conclusion on the optimal queue disclosure strategy in different sized markets is similar to the findings of Hassin and Roet-Green (2017) and Hu et al. (2018). Their studies and ours all find that to maximize the effective arrival rate, queue concealment shall be adopted in a small-size market, queue revealing shall be adopted in a large-size market, and partial queue disclosure is optimal in a moderate-size market. However, the settings and underlying driving forces are quite different. In Hassin and Roet-Green (2017), the partial information disclosure is achieved through imposing an inspection cost while in Hu et al. (2018), it is achieved through informing part of customers. In contrast, in our setting, partial queue disclosure is scenario-based: based on the realized service quality, incoming customers are either all informed or all uninformed of the queue length according to the server's pre-determined probability. Customer-based information disclosure in Hassin and Roet-Green

(2017) and Hu et al. (2018) helps the server to extract customers' surplus in certain conditions. Scenario-based information disclosure in our work helps the server to manipulate customers' posterior belief about the uncertain state (quality level) so as to attract arrivals.

4.3 Model Description

Consider a single-server queueing system. Potential customers arrive according to a Poisson process with rate λ . Their service times follow an exponential distribution with mean $1/\mu$. Let $\rho := \lambda/\mu$. Service quality can be of high value V_h with probability δ_0 or of low value V_l with probability $1 - \delta_0$. All customers (she), in case of joining the queue, receive the same quality of service and incur a waiting cost of θ per unit time. We require $V_h > V_l > \frac{\theta}{\mu}$ to ensure that at least one customer joins the system. Customers make their joining-or-balking decisions to maximize their own utility. Nature decides the value of service quality and the lottery is done once. All the above information is common knowledge, known to both the server and customers. Before the realization of service quality, the server decides his queue length disclosure strategy by selecting two conditional probabilities, f_h and f_l , that represent the probabilities that the queue length information is revealed to all incoming customers when the realized service quality is high and low, respectively. Then, $1 - f_h$ (resp. $1 - f_l$) is the corresponding probability of concealing the queue length information from customers when the realized service quality is high (resp. low). The server then commits to this strategy and announces it to all customers. The goal of the server is to maximize the expected effective arrival rate of his service system.

After the service quality is realized, the corresponding queue-disclosure action is conducted, following the pre-announced strategy. Upon observing the server's queue length disclosure action, customers update their beliefs about the service quality according to Bayes' rule. Specifically, when the queue length is revealed, customers assess the service quality to be high with probability $P_{H|R}(f_h, f_l) = \frac{\delta_0 f_h}{\delta_0 f_h + (1 - \delta_0) f_l}$ and to be low with complementary probability $1 - P_{H|R}(f_h, f_l)$. They then decide whether or not to join the queue under the assumption that the expected service value is $V_R(f_h, f_l) = P_{H|R}(f_h, f_l) V_h + (1 - P_{H|R}(f_h, f_l)) V_l$. According to Naor (1969), customers adopt a threshold policy for joining: they join the queue if and only if the queue length is smaller than some threshold $n_e(f_h, f_l) := \lfloor V_R(f_h, f_l) \mu / \theta \rfloor$. Thus, the queue in equilibrium becomes an $M/M/1/n_e(f_h, f_l)$ system and the corresponding

effective arrival rate, denoted by $\lambda_e^R(f_h, f_l)$, can be calculated as

$$\lambda_e^R(f_h, f_l) = \lambda \left(1 - \frac{\rho^{n_e(f_h, f_l)}}{\sum_{j=0}^{n_e(f_h, f_l)} \rho^j} \right).$$

In a similar way, define $P_{H|C}(f_h, f_l)$ and $V_C(f_h, f_l)$ for the case where the server conceals his queue length. When the queue is concealed, the customers' equilibrium queueing strategy can be represented by their joining probability (see Edelson and Hildebrand, 1975), denoted as $p_e(f_h, f_l)$, which equals 1 if $\lambda < \mu - \theta/V_C(f_h, f_l)$ and equals $\frac{\mu - \theta/V_C(f_h, f_l)}{\lambda}$ otherwise. The effective arrival rates, denoted as $\lambda_e^C(f_h, f_l)$, are then λ and $\mu - \theta/V_C(f_h, f_l)$, respectively.

Given the pre-determined queue-disclosure strategy, the queue is revealed (resp. concealed) with probability $\delta_0 f_h + (1 - \delta_0) f_l$ (resp. $\delta_0(1 - f_h) + (1 - \delta_0)(1 - f_l)$), and the posterior probability of high-quality service is $P_{H|R}(f_h, f_l)$ (resp. $P_{H|C}(f_h, f_l)$). The distribution of posteriors is *Bayes plausible* because the expected posterior is equal to the prior, i.e.,

$$\delta_0 = [\delta_0 f_h + (1 - \delta_0) f_l] P_{H|R}(f_h, f_l) + [\delta_0(1 - f_h) + (1 - \delta_0)(1 - f_l)] P_{H|C}(f_h, f_l).$$

The summary of the sequence of events is as follows. First, the server chooses a queue-disclosure strategy profile (f_h, f_l) and commits to it. After that, nature determines the service quality and the server makes his queue-disclosure decision based on (f_h, f_l) . Upon observing the server's disclosure action, customers update their beliefs about the service quality being high $P_{H|R}(f_h, f_l)$ (if the queue length is revealed) or $P_{H|C}(f_h, f_l)$ (if the queue length is concealed). Customers then make their corresponding joining-or-balking decisions. See Figure 4.1 for an illustration. Backward induction is adopted to derive the game outcome.




Figure 4.1: The sequence of events

First, given the queue-disclosure strategy profile (f_h, f_l) and the server's action, we can derive customers' queueing strategy $(n_e(f_h, f_l), p_e(f_h, f_l))$. We then solve the optimization

problem for the server who aims to maximize his expected effective arrival rate

$$\lambda_e(f_h, f_l) = [\delta_0 f_h + (1 - \delta_0) f_l] \lambda_e^R(f_h, f_l) + [\delta_0(1 - f_h) + (1 - \delta_0)(1 - f_l)] \lambda_e^C(f_h, f_l).$$

Denote the optimal queue-disclosure strategy for the server as (f_h^e, f_l^e) .

Given the server's queue-disclosure strategy profile (f_h, f_l) , the total utility of all customers under a revealed queue can be derived as

$$u_e^R(f_h, f_l) = \lambda \sum_{j=0}^{n_e(f_h, f_l)-1} p_j^{n_e(f_h, f_l)} \left(V_R(f_h, f_l) - \frac{(j+1)\theta}{\mu} \right),$$

where $p_j^m = \frac{\rho^j}{\sum_{k=0}^m \rho^k}$ ($0 \leq j \leq m$). Similarly, the total utility of all customers under a concealed queue can be written as

$$u_e^C(f_h, f_l) = \lambda p_e(f_h, f_l) \left(V_C(f_h, f_l) - \frac{\theta}{\mu - \lambda p_e(f_h, f_l)} \right).$$

Then, the expected total utility across customers can be expressed as

$$u_e(f_h, f_l) = [\delta_0 f_h + (1 - \delta_0) f_l] u_e^R(f_h, f_l) + [\delta_0(1 - f_h) + (1 - \delta_0)(1 - f_l)] u_e^C(f_h, f_l).$$

The optimal disclosure strategy may turn out to be a *pure* strategy, in which the server commits to fully disclosing or concealing his queue length at both quality levels (i.e., f_h and f_l can only be 0 or 1), or a *mixed* strategy, in which the server randomizes queue revealing and queue concealing at both quality levels (i.e., f_h and f_l are both larger than 0 and less than 1), or a *hybrid* strategy, in which the server randomizes revealing and concealing at one quality level and fully reveals/conceals the queue length at the other level (i.e., one of f_h and f_l is either 0 or 1 and the other is strictly between 0 and 1). We further call a queue disclosure strategy *quality-independent* if the server reveals the queue length with the same probability at both high and low quality levels and *quality-dependent* if these two probabilities are different. Clearly, the quality-independent disclosure strategy conveys no information on service quality and hence the posterior equals the prior. Only a quality-dependent disclosure strategy conveys some quality information to customers.

4.4 Optimal Queue Disclosure Strategy

In this section, we analyze the server's optimal queue disclosure strategy. First, we reformulate the server's decision problem into a nonlinear programming, the optimal solution of which can be derived in a geometric way through convex combination. After that, we investigate the impact of market size on system performances.

4.4.1 Geometric Approach

One can directly maximize the server's expected effective arrival rate by considering the disclosure probabilities (f_h, f_l) as decision variables. This approach does not yield closed-form solutions and thus cannot provide useful insights. Below, we consider the problem from another angle. We first demonstrate that there exists a one-to-one correspondence between the server's queue disclosure strategy and the Bayes-plausible posterior probabilities. We then transform the server's optimization problem into a new problem of finding the best Bayes-plausible posterior distribution. Based on that, we provide a geometric approach for deriving the optimal disclosure strategy.

In Section 4.3, we have shown that the server's queue disclosure strategy yields a unique Bayes-plausible posterior distribution. Conversely, any Bayes-plausible posterior distribution corresponds to a unique queue disclosure strategy. The details are as follows. Suppose that customers observe a revealed queue with probability p^R and a concealed queue with probability $p^C = 1 - p^R$. The posterior belief on the service quality being high conditional on a revealed queue is $p_{H|R}$ and the effective arrival rate is a function of this posterior belief, denoted by $\lambda_e^R(p_{H|R})$. Similarly, denote the posterior belief on the service quality being high conditional on a concealed queue by $p_{H|C}$ and the corresponding effective arrival rate as a function of this belief by $\lambda_e^C(p_{H|C})$. We have the following proposition.

Proposition 4.1. *Consider a prior δ_0 and two posteriors, $p_{H|R}$ with probability p^R when the queue length is revealed and $p_{H|C}$ with probability p^C when the queue length is concealed. If such a distribution of posteriors is Bayes-plausible (i.e., $\delta_0 = p^R p_{H|R} + p^C p_{H|C}$), it can be induced by a queue disclosure strategy with $f_h = p^R p_{H|R} / \delta_0$ and $f_l = p^R (1 - p_{H|R}) / (1 - \delta_0)$.*

Based on Proposition 4.1, we can transform the problem of searching for the optimal disclosure strategy into a problem of searching for the best Bayes-plausible distribution of posteriors. Mathematically, we can rewrite the server's effective arrival rate maximization

problem as follows:

$$\begin{aligned}
\lambda_e(f_h^e, f_l^e) &= \max_{p^R, p^C, p_{H|R}, p_{H|C}} p^R \lambda_e^R(p_{H|R}) + p^C \lambda_e^C(p_{H|C}) \\
s.t. \quad & p^R + p^C = 1 \\
& \delta_0 = p^R p_{H|R} + p^C p_{H|C} \\
& 0 \leq p^R, p^C, p_{H|R}, p_{H|C} \leq 1.
\end{aligned} \tag{4.1}$$

This optimization problem can be solved through a geometric approach, which we now describe in detail. Let δ represent the parameter of the posterior belief (i.e., the probability for the service quality being high). The effective arrival rate conditional on a revealed queue or a concealed queue, is a function of the expected service quality, which is determined by customers' posterior belief δ . Therefore, we can express the effective arrival rates as functions of the posterior belief δ . Now, consider the two effective arrival rate functions $\lambda_e^R(\delta)$ and $\lambda_e^C(\delta)$ in the domain $\delta \in [0, 1]$. Recall that $p^R + p^C = 1$. When p^R changes from 0 to 1, the value of $p^R \lambda_e^R(p_{H|R}) + p^C \lambda_e^C(p_{H|C})$ lies on the line segment connecting the two points $(p_{H|R}, \lambda_e^R(p_{H|R}))$ and $(p_{H|C}, \lambda_e^C(p_{H|C}))$. The crossing point of this line segment with the vertical line $\delta = \delta_0$ satisfies the Bayes plausibility requirement (4.1). Therefore, to find the optimal solution, we only need to consider all the segments connecting a point on the function curve of $\lambda_e^R(\delta)$ and a point on the function curve of $\lambda_e^C(\delta)$. The highest crossing point of all the possible line segments with the vertical line $\delta = \delta_0$ represents the maximal effective arrival rate that can be achieved through the server's queue disclosure strategy.

To facilitate the derivation of structural properties of this reformulated optimization problem, we first provide the following lemma on the shapes of the two effective arrival rate functions.

Lemma 4.1. *The two effective arrival rate functions, $\lambda_e^R(\delta)$ and $\lambda_e^C(\delta)$, exhibit the following properties:*

- (i) $\lambda_e^R(\delta)$ is a piecewise constant function with some up jumps as δ increases from 0 to 1;
- (ii) $\lambda_e^C(\delta)$ is concave and nondecreasing in δ .

The shapes of two effective arrival rate functions can be used to derive the optimal queue disclosure strategy. For the sake of analysis, we further define the point set

$$co(\lambda_e^R(\cdot), \lambda_e^C(\cdot)) = \{\alpha(\delta_1, \lambda_e^R(\delta_1)) + (1 - \alpha)(\delta_2, \lambda_e^C(\delta_2)) | 0 \leq \alpha, \delta_1, \delta_2 \leq 1\},$$

which is the convex combination of one point $(\delta_1, \lambda_e^R(\delta_1))$ on the function $\lambda_e^R(\cdot)$ and another point $(\delta_2, \lambda_e^C(\delta_2))$ on the function $\lambda_e^C(\cdot)$, where $0 \leq \delta_1, \delta_2 \leq 1$. The significance of constructing $co(\lambda_e^R(\cdot), \lambda_e^C(\cdot))$ is demonstrated in the following proposition.

Proposition 4.2. *Given a prior belief δ_0 , there exists a queue disclosure strategy (f_h, f_l) that results in an expected effective arrival rate $\lambda_e(f_h, f_l)$ if and only if $(\delta_0, \lambda_e(f_h, f_l)) \in co(\lambda_e^R(\cdot), \lambda_e^C(\cdot))$.*

Proposition 4.2 ensures that the maximal effective arrival rate needs to be searched only in the set $co(\lambda_e^R(\cdot), \lambda_e^C(\cdot))$. Define

$$\Lambda_e(\delta) := \max\{\Lambda \mid (\delta, \Lambda) \in co(\lambda_e^R(\cdot), \lambda_e^C(\cdot))\}. \quad (4.2)$$

Then, function $\Lambda_e(\delta)$, $\delta \in [0, 1]$, is the upper envelope of the set $co(\lambda_e^R(\cdot), \lambda_e^C(\cdot))$.

Based on Proposition 4.2, we have the following conclusion on the optimal queue disclosure strategy.

Proposition 4.3. *Given the prior δ_0 , the server's maximal payoff under the optimal queue disclosure strategy is $\Lambda_e(\delta_0)$.*

Proposition 4.3 indicates that a pre-committed queue-disclosure strategy helps at the given prior δ_0 only if $\Lambda_e(\delta_0) > \max\{\lambda_e^R(\delta_0), \lambda_e^C(\delta_0)\}$. A similar upper envelope is provided in Kamenica and Gentzkow (2011). However, in Kamenica and Gentzkow (2011), different signals correspond to the same value function of the sender, and thus the upper envelope is formed through the concavification of that value function. Differently, in our work, signals are the queue-disclosure actions—revealing and concealing the queue length. These two signals correspond to two different value functions. Under the Bayes plausibility condition, the upper envelope is formed through the convex combination of these two value functions. We now illustrate the aforementioned geometric approach in the following example.

Example 4.2. (Illustration of the Upper Envelope) *Consider the parameter values to be $V_h = 2$, $V_l = 1$, $\mu = 1.1$, $\theta = 1$ and $\lambda = 0.6$. The dashed curve in Figure 4.2 represents the effective arrival rate function $\lambda_e^C(\delta)$ and the dotted piecewise flat line represents the effective arrival rate function $\lambda_e^R(\delta)$. Clearly, the upper envelope formed by all the segments connecting two arbitrary points on these two effective arrival rate functions is the solid line connecting the two points $(0, \lambda_e^R(0))$ and $(1, \lambda_e^C(1))$. The first point $(0, \lambda_e^R(0))$ represents the effective arrival rate of a revealed queue with a posterior belief $p_{H|R} = 0$, and the second point $(1, \lambda_e^C(1))$*




Figure 4.2: The upper envelope Λ_e : $V_h = 2$, $V_l = 1$, $\mu = 1.1$, $\theta = 1$ and $\lambda = 0.6$

represents the effective arrival rate of a concealed queue with a posterior belief $p_{H|C} = 1$. Given any prior belief δ_0 , say $\delta_0 = 0.3$, we can recover the probability p^R by solving the Bayes plausibility condition $p^R * 0 + (1 - p^R) * 1 = 0.3$, which yields $p^R = 0.7$. Then, according to Proposition 4.1, we can recover the optimal queue disclosure strategy as follows: $f_h = p^R p_{H|R} / \delta_0 = 0$ and $f_l = p^R (1 - p_{H|R}) / (1 - \delta_0) = 1$. One can easily check that for any prior belief $\delta_0 \in (0, 1)$, the optimal queue disclosure strategy is to always conceal the queue length when the realized service quality is high but to always reveal it when the realized service quality is low, i.e., $(f_h^e, f_l^e) = (0, 1)$. That is, the server's optimal queue-disclosure strategy is pure and quality-dependent, which fully conveys the quality information to customers.

Example 4.2 shows that a pure queue-disclosure strategy can be the server's optimal strategy. Below, we will demonstrate that a hybrid disclosure strategy works best for the server under the setting given in Example 4.1 (stated in the Introduction).

Example 4.3. (Illustration of Example 4.1 via Geometric Approach) The motivating Example 4.1 is illustrated in Figure 4.3. In this example, the effective arrival rate function of the concealed queue reaches a flat line at $d_{12} = 0.6667$. The upper envelope formed by all the segments connecting two arbitrary points of the two effective arrival rate functions is represented by the solid line. Given the prior belief $\delta_0 = 0.33$, the maximal effective arrival rate is on the segment connecting the two points $(0, \lambda_e^R(0))$ and $(d_{12}, \lambda_e^C(d_{12}))$. The first point $(0, \lambda_e^R(0))$ represents the effective arrival rate in a revealed queue with a posterior belief




Figure 4.3: The upper envelope Λ_e : $V_h = 2$, $V_l = 1$, $\mu = 1.1$, $\theta = 1$ and $\lambda = 0.5$

$p_{H|R} = 0$, and the second point $(d_{12}, \lambda_e^C(d_{12}))$ represents the effective arrival rate in a concealed queue with a posterior belief $p_{H|C} = 0.6667$. Given $\delta_0 = 0.33$, we can recover the probability p^R by solving the Bayes plausibility condition $p^R * 0 + (1 - p^R) * 0.6667 = 0.33$, which yields $p^R = 0.5050$. Then, according to Proposition 4.1, we can recover the optimal queue disclosure strategy as follows: $f_h = p^R p_{H|R} / \delta_0 = 0$ and $f_l = p^R (1 - p_{H|R}) / (1 - \delta_0) = 0.7537$. This is a hybrid strategy and it conveys partial information about service quality to customers. By checking the graph of the upper envelope, we can see that for any prior belief $\delta_0 \in (0, d_{12})$, $\Lambda_e(\delta_0)$ is located above the two effective arrival rate functions and the corresponding hybrid queue-disclosure strategy is beneficial to the server.

It is worth mentioning that, although the maximal effective arrival rate is unique for a given prior belief δ_0 , the corresponding optimal queue-disclosure strategy is not necessarily unique. A point on the upper envelope may correspond to multiple pairs of posteriors whose distribution is Bayes-plausible. Let us revisit Example 2 and the upper envelope plotted in Figure 2. We still keep the parameter values $V_h = 2$, $V_l = 1$, $\mu = 1.1$ and $\theta = 1$, but change the value of λ from 0.6 to 0.7160. In such a setting, the up-jumping point of $\lambda_e^R(\delta)$, $(0.8182, 0.5698)$, happens to locate on the upper envelope $\Lambda_e(\delta)$ that is a segment connecting the two points $(0, \lambda_e^R(0))$ and $(1, \lambda_e^C(1))$, where $\lambda_e^R(0) = 0.4337$ and $\lambda_e^C(1) = 0.6000$. When the prior is $\delta_0 = 0.8182$, we can obtain the following two optimal queue-disclosure strategies: $(f_h^e, f_l^e) = (0, 1)$ or $(1, 1)$.

Based on the above geometric approach, we can further obtain the following lemma.

Lemma 4.2. *The optimal queue disclosure strategy is pure and quality-independent in the following two situations:*

(i) $(f_h^e, f_l^e) = (0, 0)$ if $\lambda_e^C(\delta) \geq \lambda_e^R(\delta)$ for all $\delta \in [0, 1]$.

(ii) $(f_h^e, f_l^e) = (1, 1)$ if $\lambda_e^R(\delta)$ is a constant function (i.e., a horizontal line) and $\lambda_e^R(\delta) \geq \lambda_e^C(\delta)$ for all $\delta \in [0, 1]$.

The first statement of Lemma 4.2 requires that the effective arrival rate function of a concealed queue is located above that of a revealed queue. By considering its concavity property (see Lemma 4.1), we can conclude that the upper envelope function coincides with the effective arrival rate function of a concealed queue. Hence, the optimal queue disclosure strategy is to always conceal the queue regardless of the realized service quality. The second statement of Lemma 4.2 provides a sufficient condition for the optimal queue disclosure strategy to be always revealing, regardless of the realized service quality. Note that this condition not only requires the effective arrival rate function of a revealed queue to be located above that of a concealed queue but also requires the former to be a constant function, that is, no jumps occur for this function in the whole domain $\delta \in [0, 1]$.

4.4.2 The Impact of Market Size

In this section, we fix the prior belief δ_0 and explore the impact of market size (i.e., the potential arrival rate) λ on the server's optimal queue disclosure strategy.

When the service quality is certain, the impact of market size on the delay announcement strategy has been well studied in the literature. According to Hassin (1986) and Chen and Frank (2004), when the market size λ is below a threshold value, concealing the queue makes the server better off; otherwise, revealing the queue is preferred. Moreover, when λ is very small, customers 'all join' in the unobservable queue setting while some customers balk in the observable queue setting. Hence, concealing the queue-length information is the better option for servers with very small λ . As λ becomes large enough, the effective arrival rate becomes a constant in unobservable queues because customers' joining utility is now zero and no more customers want to join. However, in an observable queue setting, the queue is stochastically longer as λ increases and hence the effective arrival rate is strictly increasing in λ : there always exists a chance for the increased amount of customers to observe a short

queue and join. Therefore, revealing the queue is preferred by the server when λ is very large. In our work, the aforementioned results and insights can still hold for the sufficiently small and sufficiently large markets under certain conditions, as implied by Lemma 4.2. We now formally show that these results also hold for our optimal queue disclosure strategy.

Proposition 4.4. *The optimal queue disclosure strategy (f_h^e, f_l^e) satisfies the following two properties:*

- (i) *If the potential arrival rate $\lambda < \mu - \frac{\theta}{\delta V_h + (1-\delta)V_l}$, then the server's optimal strategy is to always conceal the queue; that is, $(f_h^e, f_l^e) = (0, 0)$.*
- (ii) *There exists a threshold denoted by $\bar{\lambda}^e$ (which is greater than $\mu - \frac{\theta}{\delta V_h + (1-\delta)V_l}$) such that if $\lambda > \bar{\lambda}^e$,^{4.1} the server's optimal strategy is to always reveal the queue; that is, $(f_h^e, f_l^e) = (1, 1)$.*

Indeed, when the market size is very small, all customers join the concealed queue regardless of the value of service quality and thus concealing the queue is the server's optimal strategy. Similarly, when the market size is very large, revealing the queue is the optimal strategy. However, when the market size λ is intermediate, things become tricky and the optimal disclosure strategy depends on the tradeoff between the value of informing customers of the queue length and the value of providing partial quality information. We will use the following numerical example to illustrate.

Example 4.4. (*Sensitivity Analysis: The Impact of Market Size λ on the Server's Optimal Queue Disclosure Strategy and System Performance*) Consider the parameter values $V_h = 18$, $V_l = 3$, $\mu = 3$, $\theta = 8$ and $\delta_0 = 0.1$. There are three key market size thresholds as shown in Figure 4.4: $\lambda_{A3} = 1.2222$, $\lambda_{B3} = 2.0621$ and $\lambda_{C3} = 17.2492$. There, the bottom subplot depicts the server's optimal queue disclosure strategy (f_h^e, f_l^e) as a function of λ , the middle subplot shows the customers' total utility, while the upper subplot depicts the maximal effective arrival rate that can be achieved by adopting the pre-committed optimal queue disclosure strategy (f_h^e, f_l^e) .

Figure 4.4 shows that when the market size is small ($\lambda < \lambda_{A3}$), fully concealing the queue (i.e., $(f_h^e, f_l^e) = (0, 0)$) is the dominant strategy because in such a situation, all customers join the concealed queue. When the market size reaches the threshold λ_{A3} , balking is possible as customers' expected joining utility is now reduced to zero. As the market size further increases

^{4.1}The definition of $\bar{\lambda}^e$ can be found in the proof of Proposition 4.4.




Figure 4.4: The impact of market size on the optimal queue disclosure strategy, maximal effective arrival rate and customers' total utility: $V_h = 18$, $V_l = 3$, $\mu = 3$, $\theta = 8$ and $\delta_0 = 0.1$

and becomes larger than λ_{A3} , the server still conceals the queue when the realized service quality is high but starts to randomize concealing and revealing when the realized service quality is low, with the probability of revealing the queue increasing in the market size λ . Performing such a randomization, on the one hand, can strengthen the customers' belief on the high quality when they observe that the queue length is concealed but, on the other hand, provides a chance for the customers to see a revealed queue, from which they then infer that the service quality is low. Overall, the increase of the customers' effective arrival rate in the former case surpasses the reduction of the customers' effective arrival rate in the latter case, thereby benefiting the server. When the market size reaches λ_{C3} and keeps further increasing, it is no longer beneficial to conceal the queue if the realized service quality is high as now revealing the queue induces more customers to join. Then, fully revealing the queue is the dominant strategy; that is, $(f_h^e, f_l^e) = (1, 1)$.

Consequently, as shown in the upper subplot of Figure 4.4, in a small-sized market ($\lambda \in (0, \lambda_{A3})$), the maximal effective arrival rate coincides with that of a fully concealed queue, while in a large-sized market ($\lambda \in (\lambda_{C3}, +\infty)$), it coincides with that of a fully revealed queue. However, for a medium-sized market ($\lambda \in (\lambda_{A3}, \lambda_{C3})$), the maximal effective arrival rate is strictly larger than that of either a fully revealed or fully concealed queue. The difference between them can be used as a measure of the value of providing the quality-dependent disclosure

strategy.

Regarding the customers' total utility, to better understand the impact of the quality-dependent queue disclosure strategy, we also derive the quality-independent optimal queue disclosure strategy, denoted by $(\hat{f}_h^e, \hat{f}_l^e)$, and the corresponding customers' total utility. The middle subplot of Figure 4.4 shows that compared to a quality-independent queue disclosure strategy, our quality-dependent queue disclosure strategy can improve customers' total utility only when the market size λ falls into a relatively small range $(\lambda_{A3}, \lambda_{B3})$. However, in a relatively large market size range $\lambda \in (\lambda_{B3}, \lambda_{C3})$, customers' total utility is smaller under our quality-dependent optimal queue disclosure strategy than that under the quality-independent optimal queue disclosure strategy. Therefore, although the pre-committed queue disclosure strategy can be used to attract more customers to join the service system, it does not necessarily benefit them.

Note that in Example 4.4, when the market size falls into the range $\lambda \in (\lambda_{A3}, \lambda_{C3})$, the equilibrium is hybrid: the server randomizes concealing and revealing the queue only when the realized service quality is low. We also conduct other numerical examples and find that it is also possible that the equilibrium is fully mixed; that is, the server randomizes the queue disclosure and concealment at both the high- and low-quality states.

4.5 Social Planner

In the previous section, we consider a profit-maximizing server and study his optimal queue disclosure strategy. In reality, however, servers can be social planners whose aim is to maximize the overall social welfare (i.e., customers' total utility in our work). We now extend our commitment game into this setting. We will show that our geometric approach is robust and provides some new insights on the optimal queue-disclosure strategy.

The social welfare in our setting is the sum of customers' utilities, defined as

$$u_e(f_h, f_l) = [\delta_0 f_h + (1 - \delta_0) f_l] u_e^R(f_h, f_l) + [\delta_0(1 - f_h) + (1 - \delta_0)(1 - f_l)] u_e^C(f_h, f_l),$$

and the details on $u_e^R(f_h, f_l)$ and $u_e^C(f_h, f_l)$ are provided in section 4.3. To analyze the social planner's optimal queue disclosure strategy, we first investigate the geometry behind the functions involved, and obtain the following results regarding the shapes of two utility functions, $u_e^R(\cdot)$ and $u_e^C(\cdot)$.

Lemma 4.3. *The two utility functions, $u_e^R(\delta)$ and $u_e^C(\delta)$, exhibit the following properties:*

- (i) $u_e^R(\delta)$ is a piecewise linear and increasing function, with some down jumps, as δ increases from 0 to 1;
- (ii) if $\lambda \geq \mu - \frac{\theta}{V_h}$, $u_e^C(\delta)$ equals 0 for all $\delta \in [0, 1]$; if $\lambda \leq \mu - \frac{\theta}{V_l}$, $u_e^C(\delta)$ is linear and increasing in δ for $\delta \in [0, 1]$; otherwise, $u_e^C(\delta)$ equals 0 for $\delta \in \left[0, \frac{\theta/(\mu-\lambda)-V_l}{V_h-V_l}\right]$ and linear and increasing in δ for $\delta \in \left(\frac{\theta/(\mu-\lambda)-V_l}{V_h-V_l}, 1\right]$;
- (iii) $u_e^R(\delta) > u_e^C(\delta)$ for all $\delta \in [0, 1]$.

The third statement of Lemma 4.3 shows that from the viewpoint of welfare maximization, revealing the queue is always better than concealing it (see also Hassin and Haviv, 2003; and Hassin and Roet-Green, 2017). This is intuitive as revealing the queue can help customers to make an informed decision. However, considering the optimal queue disclosure strategy, this classical result no longer holds. We will show that concealing the queue length with a strictly positive probability can be socially desired under some situations.

Using the geometric approach, we can construct the convex-combination point set

$$co(u_e^R(\cdot), u_e^C(\cdot)) = \{\alpha(\delta_1, u_e^R(\delta_1)) + (1 - \alpha)(\delta_2, u_e^C(\delta_2)) \mid 0 \leq \alpha, \delta_1, \delta_2 \leq 1\},$$

and the upper envelope of it

$$U_e(\delta) := \max\{U \mid (\delta, U) \in co(u_e^R(\cdot), u_e^C(\cdot))\},$$

from which the social planner's optimal queue disclosure strategy $(\tilde{f}_h^e, \tilde{f}_l^e)$ can be derived.

Example 4.5. (Illustration of the Social Planner's Optimal Queue Disclosure Strategy) Consider a setting with $V_h = 2$, $V_l = 1$, $\mu = 1.1$, $\theta = 1$ and $\lambda = 10$. The dotted piecewise increasing line in Figure 4.5 represents the utility function $u_e^R(\delta)$, the flat line represents the utility function $u_e^C(\delta)$, and the upper envelope formed by all the segments connecting two arbitrary points on these two utility functions is the solid line. Given the prior belief $\delta_0 = 0.83$, the maximal effective arrival rate is on the segment connecting the two points $(\hat{\delta}^-, u_e^R(\hat{\delta}^-))$ and $(1, u_e^C(1))$, where $\hat{\delta}$ is the down-jumping point of the utility function $u_e^R(\delta)$, $\hat{\delta}^- = \lim_{\delta \rightarrow \hat{\delta}, \delta < \hat{\delta}} \delta$ and $u_e^R(\hat{\delta}^-)$ is the left-hand limit of $u_e^R(\cdot)$ at the point $\hat{\delta}$. The first point $(\hat{\delta}^-, u_e^R(\hat{\delta}^-))$ represents the effective arrival rate in a revealed queue with a posterior

belief $p_{H|R} = 0.8181$, and the second point $(1, u_e^C(1))$ represents the effective arrival rate in a concealed queue with a posterior belief $p_{H|C} = 1$. Given $\delta_0 = 0.83$, we can recover the probability p^R by solving the Bayes plausibility condition $p^R * 0.8181 + (1 - p^R) * 1 = 0.83$, which yields $p^R = 0.9346$. Then, according to Proposition 4.1, we can obtain the optimal queue disclosure strategy: $\tilde{f}_h^e = p^R p_{H|R} / \delta_0 = 0.9212$ and $\tilde{f}_l^e = p^R (1 - p_{H|R}) / (1 - \delta_0) = 1.0000$. Clearly, this optimal strategy is a hybrid one. Figure 4.5 indicates that under the prior $\delta_0 = 0.83$, the expected total utility under the optimal queue disclosure strategy (i.e., $U_e(\delta_0) = 0.8419$) achieves a 659% improvement over the one under the ‘always revealing’ strategy (i.e., $u_e^R(\delta_0) = 0.1109$).




Figure 4.5: The upper envelope U_e : $V_h = 2$, $V_l = 1$, $\mu = 1.1$, $\theta = 1$ and $\lambda = 10$

Let $\hat{\Delta}$ denote the set of values of δ where $u_e^R(\delta)$ jumps down (or specifically, the term $[\delta V_h + (1 - \delta)V_l]\mu/\theta$ takes integer values). Then, consider any prior $\delta_0 \in [\hat{\delta}, \hat{\delta} + \epsilon)$ where $\hat{\delta} \in \hat{\Delta}$, ϵ is sufficiently small and $\hat{\delta} + \epsilon < 1$. For the point (δ_0, \bar{u}) on the segment connecting two points $(\hat{\delta}^-, u_e^R(\hat{\delta}^-))$ and $(1, u_e^C(1))$, \bar{u} is strictly larger than $u_e^R(\delta_0)$, which implies that the optimal queue disclosure strategy must achieve a larger expected utility than the ‘always revealing’ strategy does. Also, we note that $\hat{\delta}$ is independent of the market size λ . To summarize, we have the following conclusion.

Proposition 4.5. *Given any prior $\delta_0 \in [\hat{\delta}, \hat{\delta} + \epsilon)$ where $\hat{\delta} \in \hat{\Delta}$, ϵ is sufficiently small and $\hat{\delta} + \epsilon < 1$, the optimal queue disclosure strategy achieves a larger expected customer utility than the ‘always revealing’ strategy does; that is, $U_e(\delta_0) > u_e^R(\delta_0)$ for all $\lambda \in (0, +\infty)$.*

Proposition 4.5 implies that full disclosure is not necessarily socially desired, and it may be in the best interest of the social planner to conceal the queue length information with positive probability. This conclusion echoes the ones in Cui and Veeraraghavan (2016), Hu et al. (2018) and Li et al. (2020). According to Naor (1969), tolls/taxes can be levied in queueing systems to control arrivals in order to improve welfare. The lack of information, according to Cui and Veeraraghavan (2016), acts as an information tax that deters admission, leading to improved welfare. Similar rationale holds here.

Different from the profit-maximizing case where the server’s optimal queue disclosure strategy is ‘always concealing’ (resp. ‘always revealing’) when the market size λ is small (resp. large) enough, the optimal queue disclosure strategy of a social planner can be quality-dependent over the whole range of λ , as illustrated in the following example.




Figure 4.6: The impact of market size on the social planner’s optimal queue disclosure strategy and customers’ total utility: $V_h = 2$, $V_l = 1$, $\mu = 1.1$, $\theta = 1$ and $\delta_0 = 0.83$

Example 4.6. (*Sensitivity Analysis: The Impact of Market Size λ on the Social Planner’s Optimal Queue Disclosure Strategy and System Performance*) Consider the setting in Example 4.5, where the prior belief δ_0 is very close to the down-jumping point $\hat{\delta}$. Figure 4.6 indicates that on the whole range of the market size ($\lambda \in (0, +\infty)$), the optimal queue disclosure strategy $(\tilde{f}_h^e, \tilde{f}_l^e)$ achieves a strictly larger expected utility than the ‘always revealing’ strategy $(f_h^e, f_l^e) = (1, 1)$.

In the profit-maximizing case, more arrivals are always preferred. But in the welfare maximization case, it is sometimes socially desired to persuade fewer customers to join as an overly crowded system can reduce the overall utility for customers. To discourage some customers from joining, the social planner should convince customers that the service quality can be low. In Figure 4.6, for $0 < \lambda < \lambda_D (= 0.1000)$, the optimal queue disclosure strategy is a pure one with $(\tilde{f}_h^e, \tilde{f}_l^e) = (0, 1)$, which indirectly provides full information on the quality type. In this case, as the market size λ is very small, the expected total utility from joining a concealed queue is strictly positive when the posterior belief $p_{H|C}$ is 1, which is even larger than the one from joining a revealed queue with prior belief $\delta_0 = 0.83$. This provides an incentive for the social planner to conceal the queue in case of high quality. And to achieve a higher overall utility, the server should reveal the queue in case of low quality. For $\lambda_D \leq \lambda < +\infty$, the utility from a concealed queue becomes relatively small, and the one from a revealed queue with the prior $\delta_0 = 0.83$ is also not that large. In this case, the social planner should randomize revealing and concealing to reduce customers' belief about the service quality being high (i.e., $p_{H|R}$). Just like the special case with $\lambda = 10$ in Example 4.5, the optimal queue disclosure strategy is $(\tilde{f}_h^e, \tilde{f}_l^e) = (0.9212, 1.0000)$. Under such a strategy, the customers' belief about the service quality being high after seeing a revealed queue becomes $0.8181 (< \delta_0)$, which decreases the maximal queue length from 2 to 1 and achieves a higher expected utility than the 'always revealing' strategy does.

4.6 Conclusions and Suggestions for Future Research

In some service systems, the service quality is generally uncertain. In this work, we examine a situation in which before the realization of service quality, the server can design for his benefit a queue-disclosure strategy that links the queue concealment and revelation with the realized service quality and ex-ante commits to it. We demonstrate that the commitment to such a disclosure strategy helps the server to attract more customers to join the service system than otherwise.

We transfer the problem of searching for the optimal queue disclosure strategy to an equivalent problem of searching for the optimal Bayes-plausible posterior distribution. Based on the reformulated optimization problem, we then provide a geometric approach to obtain the optimal strategy. We show that as long as the upper envelope of all the convex combinations of one point from the effective arrival rate function of a concealed queue and another point

from that of a revealed queue is located above these two functions, a properly designed queue disclosure strategy (which might involve randomization) can be utilized to help attracting more customers to join the service system. We also investigate the impact of the market size on the server's optimal queue disclosure strategy. We show that it is always in the server's best interest to conceal the queue in a very small-sized market but to reveal it in a very large-sized market. In a medium-sized market, through the numeric study, we find that it is often optimal for the server to randomize queue concealment and revelation to either fully or partially convey the service quality information to customers. We then extend our analysis to a setting where the server is a welfare-maximizing social planner. We show that the geometric approach can be easily applied to this situation. We find that it may be beneficial for the social planner to randomize revealing and concealing the queue over the whole range of the market size. This result is in sharp contrast to the one stated in the classical literature (see, e.g., Hassin and Haviv, 2003; and Hassin and Roet-Green, 2017) that it is always socially optimal to reveal the queue.

Our work demonstrates that the quality-linked queue-disclosure strategy can be used to persuade more customers to join a queueing system. We further present an intuitive geometric approach on how to find such an optimal strategy. Admittedly, our model has limitations. First, we restrict the signal to be a binary one, namely concealing or revealing the queue. Under this assumption, the effective arrival rates can be easily calculated. It would be an interesting topic to extend our approach to other types of delay announcements such as informing customers about the exact waiting time (Guo and Zipkin, 2007) or announcing the waiting time of the last customer to enter service (Ibrahim et al., 2017). Second, in our information disclosure scheme, only the quality type is regarded as the uncertain state of the world. It would be an interesting research to find out the optimal persuasion mechanism by taking both the quality type and queue length as the joint uncertain states of the world. Despite such limitations, we hope that our work can serve as a stepping stone for further studies on the combination of information disclosure and Bayesian persuasion in queueing systems.

Appendix A

Supplements and Proofs for Chapter 2

A.1 A General Two-Decision-Variable Problem

In this extension part, we consider a general setting where both M and y are decision variables. Such an extended model corresponds to some real settings, especially in the catering industry. For example, in restaurants, the early-bird discount on menu items has been an effective way to shift customers' dining time to off-peak hours (see Susskind et al., 2004).^{A.1} There, the product becomes the meal on the menu, and the limited seat capacity generally does not affect the sales of meals because those customers who find no available seats can choose to wait or do the order take-out. Our two-decision-variable analysis extends the existing Bayesian inventory management literature in which only one decision variable is considered. We also provide the upper bounds of the myopic- and Bayesian-optimal inventory levels.

A.1.1 One-period Model

To be consistent with the main content, hereafter we call the inventories reserved for the early bird discount 'the discounted product' and the inventories reserved for the regular-price sales 'the regular-price product'. The baseline one-period model and all notations are all the same as those in §2.3 except an extra cost term cM that is proportional to the total inventory level M with a per-unit cost parameter $c > 0$. Given the demand parameter θ and the buy-up substitution probability α , the inventory manager determines the optimal two-dimensional

^{A.1}Take Cafe De Coral, a representative restaurant in Hong Kong, as an example. It provides an early-bird-discount price on hot-pot meals to those customers arriving during one hour before the regular dinner time from Monday to Friday (except public holidays); for more details, see <https://www.jetsoclub.com/2019/09/cafe-de-coral-hot-pot-early-bird-offer-0930.html>.

inventory levels (y, M) ^{A.2} to maximize the total expected profit $\pi(y, M|\theta, \alpha)$ over the two selling phases as follows:

$$\begin{aligned} \max_{y, M} \pi(y, M|\theta, \alpha) &= p_1 E[D_1 \wedge y|\theta] + p_2 E[(K + D_2) \wedge (M - y \wedge D_1)|\theta, \alpha] - cM \text{ (A.1)} \\ \text{s.t. } &0 < y \leq M, \end{aligned}$$

We get the following important *supermodularity* properties (Topkis, 2011) that will be used in the analysis of Bayesian inventory management and upper bounds.

Proposition A.1. *The profit function $\pi(y, M)$ is supermodular in (y, M) ; that is, $\pi(y + 1, M + 1|\theta, \alpha) + \pi(y, M|\theta, \alpha) \geq \pi(y + 1, M|\theta, \alpha) + \pi(y, M + 1|\theta, \alpha)$. It is also supermodular in y and p_1 , i.e., $\partial[\pi(y + 1, M|\theta, \alpha) - \pi(y, M|\theta, \alpha)]/\partial p_1 \geq 0$.*

Proposition A.1 indicates that the optimal value of the inventory level of the discounted product y is increasing in both the total number of the product M and the early-bird-discount price p_1 , and the optimal value of the total inventory level M is increasing in the inventory level of the discounted product y .

A.1.2 Multi-period Bayesian Inventory Management

The multi-period setting is quite similar to the one stated in §2.4 except that we now have two decision variables y and M . The likelihood functions under the \mathcal{OS} and \mathcal{OT} scenarios have nothing to do with M . Hence, the relationships between the Bayesian optimal inventory level of the discounted product and the corresponding myopic one remain intact under these two scenarios, and the comparison over the total inventory level is determined only by the supermodularity between y and M as stated in Proposition A.1. In summary, we can directly obtain the following results.

Corollary A.1. *For any period i ($i = 1, \dots, N$), when the buy-up substitution probability α is known, given the same prior distribution $\phi'_i(\theta)$,*

- (a) *the Bayesian optimal inventory levels under the \mathcal{OS} scenario are equal to the corresponding myopic ones; that is, $y_i^{\mathcal{OS}} = y_i^m$ and $M_i^{\mathcal{OS}} = M_i^m$; and*

^{A.2}To keep consistency, in all the following one-period and multi-period analysis, when there exist multiple optimal inventory-level pairs, we first pick out those optima with the smallest total inventory level M , from which we then choose the one with the smallest inventory level of the discounted product y as the final optimal solution.

(b) the Bayesian optimal inventory levels under the \mathcal{OT} scenario are no less than the corresponding myopic ones; that is, $y_i^{\mathcal{OT}} \geq y_i^m$ and $M_i^{\mathcal{OT}} \geq M_i^m$.

When both the demand parameter θ and the buy-up substitution probability α are unknown, given the same prior $\phi_i(\theta, \alpha)$, the Bayesian optimal inventory levels under the \mathcal{OS} scenario are no larger than the corresponding myopic ones; that is, $y_i^{\mathcal{OS}} \leq y_i^m$ and $M_i^{\mathcal{OS}} \leq M_i^m$.

Next, we turn to the \mathcal{US} and \mathcal{UT} scenarios, and mainly investigate whether the unobservable lost sale only yields the “stock more” result when α is known. Following the similar analysis stated in the proof of Lemma 2.1, we can obtain the following results.

Lemma A.1. *When the buy-up substitution probability α is known, under the scen scenario ($scen \in \{\mathcal{US}, \mathcal{UT}\}$), for any period i ($i = 1, \dots, N - 1$) with $0 < y \leq M$, given the prior distribution $\phi_i(\theta)$, we have*

$$E_{\phi_i(\theta)} \left\{ \sum_{\xi \in I_{scen}^{y+1, M+1}} v_{i+1}^{scen}(\phi'_{i+1}) f_{scen}^{y+1, M+1}(\xi | \theta, \alpha) \right\} \geq E_{\phi_i(\theta)} \left\{ \sum_{\xi \in I_{scen}^{y, M}} v_{i+1}^{scen}(\phi'_{i+1}) f_{scen}^{y, M}(\xi | \theta, \alpha) \right\},$$

and

$$E_{\phi_i(\theta)} \left\{ \sum_{\xi \in I_{scen}^{y, M+1}} v_{i+1}^{scen}(\phi'_{i+1}) f_{scen}^{y, M+1}(\xi | \theta, \alpha) \right\} \geq E_{\phi_i(\theta)} \left\{ \sum_{\xi \in I_{scen}^{y, M}} v_{i+1}^{scen}(\phi'_{i+1}) f_{scen}^{y, M}(\xi | \theta, \alpha) \right\}.$$

Lemma A.1 implies that increasing the inventory level of either the discounted or regular-price product can increase the total discounted expected profit for the following periods, which seems to verify that the “stock more” result applies here. Is this conjecture true? The answer is *no*. To illustrate this, let us consider the following example under the \mathcal{UT} scenario.

Example A.1. *Consider a two-period Bayesian inventory management problem under the \mathcal{UT} scenario. The related parameter values are set as follows: the discounted price $p_1 = 100$, the regular price $p_2 = 110$, the per-unit cost $c = 97$, the buy-up substitution probability $\alpha = 0.5$ and the discount factor $\delta = 1$. The demand parameter θ can only be 1 or 2. When $\theta = 1$, the primary demand for the discounted product is $D_1 = 4$ and that for the regular-price product is $D_2 = 2$; when $\theta = 2$, $D_1 = 4$ and $D_2 = 4$. At the beginning of the first selling period, if the inventory manager holds a prior belief that $Pr(\theta = 1) = 0.3$ and $Pr(\theta = 2) = 0.7$, the optimal inventory levels for this period satisfy that $(y_1^{\mathcal{UT}} = 1) < (y_1^m = 4)$ and $(M_1^{\mathcal{UT}} =$*

5) $< (M_1^m = 6)$; if the prior belief is that $Pr(\theta = 1) = 0.2$ and $Pr(\theta = 2) = 0.8$, we have $(y_1^{UT} = 4) > (y_1^m = 1)$ and $(M_1^{UT} = 7) > (M_1^m = 5)$.

As the UT scenario contains the stronger “stock more” driving forces compared with the US scenario, Example A.1 implies that the unobservable lost sale cannot guarantee the occurrence of “stock more”, which is again in sharp contrast to the existing “stock more” finding (Lariviere and Porteus, 1999; Ding et al., 2002; and Chen and Plambeck, 2008). Intuitively, the marginal benefit gained from increasing one unit of inventory differs between the discounted and regular-price products due to factors such as the prices and the degree of demand uncertainty. When the marginal benefit gained from increasing one type of the product dominates that of the other type, the Bayesian inventory manager has the incentive to hold more its inventory. However, doing that might increase the overstocking risk. One way to mitigate such risk is to reduce the inventory level of the other type. Furthermore, we can show in the following proposition that under the $scen$ scenario ($scen \in \{US, UT\}$), once the inventory level of one type of the product is reduced, the inventory level of the other one must be increased.

Proposition A.2. *Under the $scen$ scenario ($scen \in \{US, UT\}$), for any period i ($i = 1, \dots, N$), we have*

- (a) *if the Bayesian optimal inventory level of the discounted product is lower than its myopic one ($y_i^{scen} < y_i^m$), then the Bayesian optimal inventory level of the regular-price product shall be higher than its myopic one, i.e., $M_i^{scen} - y_i^{scen} > M_i^m - y_i^m$;*
- (b) *if the Bayesian optimal inventory level of the regular-price product is lower than its myopic one ($M_i^{scen} - y_i^{scen} < M_i^m - y_i^m$), then the Bayesian optimal inventory level of the discounted product shall be higher than its myopic one, i.e., $y_i^{scen} > y_i^m$.*

Under the US and UT scenarios, stocking more to gain more information about demands is still a trend. It may be necessary for us to enhance the inventory levels of both discounted and regular-price products, and it is also possible to lower the inventory level of one type of the product to mitigate the overstocking risk. However, it is never true that we lower the inventory levels of both types of products simultaneously. Proposition A.2 tells us that the aim of such inventory level reduction is to allocate more inventory for the other type as the marginal benefit of increasing that type’s inventory level is larger.

A.1.3 Upper Bounds on Optimal Inventory Levels

We now derive the upper bounds for the optimal total inventory level y and the inventory level of the discounted product M . First, let us consider the one-period setting. To obtain the globally optimal values of y and M , one can numerically conduct the one-dimensional search through the variable space such as M . Given each M , one can then compute the optimal $y^*(M)$ based on equation (2.2) stated in §2.3. Note that one can notably reduce the search range and shorten the computational time with the knowledge about the upper bounds of the optimal y and M . Using the supermodularity properties in Proposition A.1, we can obtain an upper bound for the two optimal inventory levels by solving the optimization problem with $p_1 = p_2$. When the early-bird-discount price p_1 is increased to p_2 , there is no discount. Hence, the inventory decision problem is reduced to be a newsvendor problem. We then only need to consider the total demand for the regular-price product, which can be written as $(D_1 \wedge M + K_M + D_2)$, where K_M follows a binomial distribution with parameters $((D_1 - M)^+, \alpha)$. Next, we consider a stochastically larger demand $D_1 + D_2$ to replace $(D_1 \wedge M + K_M + D_2)$, with which we can derive an upper bound for the optimal inventory levels.

Similarly, using the above enlarging technique, we can develop the upper bounds for the Bayesian optimal inventory levels in a multi-period setting. As the first step, we derive an enlarged concave objective function by increasing p_1 to p_2 . Then, we have the regular-price product only. Again, we can simply consider a stochastically larger demand $D_1 + D_2$ as the demand for the regular-price product. Here, we use a superscript p_2 to represent this case. Then, the objective function becomes $\pi^{p_2}(M, M|\theta)$, which relies only on the demand parameter θ and is a concave function of M . In Bayesian inventory management, the inventory decision for the current period will affect all the following periods. To simplify the upper bound analysis, we need to toss out such delayed effect. Let $\pi^{p_2, sup} := \sup_{\theta \in \Theta} \{\max_{M>0} \pi^{p_2}(M, M|\theta)\}$, denoting a constant upper bound on $\pi^{p_2}(M, M|\theta)$. Then, based on the results in the one-period model, we can show that for any information scenario $scen \in \{\mathcal{OS}, \mathcal{OT}, \mathcal{US}, \mathcal{UT}\}$, its objective function satisfies

$$G_i^{scen}(y, M, \phi_i) \leq G_i^{p_2, sup}(M, \phi'_i) := E_{\phi'_i(\theta)} \{\pi^{p_2}(M, M|\theta)\} + \sum_{n=1}^{N-i} \delta^n \cdot \pi^{p_2, sup}, i = 1, \dots, N,$$

where $\phi'_i(\theta) = \int_0^1 \phi_i(\theta, \alpha) d\alpha$, which is the marginal prior distribution of θ for period i . In this way, we enlarge the objective functions under all the information scenarios into a newsvendor-

type objective function $G_i^{p_2, sup}(M, \phi'_i)$, which is concave in the inventory stocking decision M . Denote the optimal solution of $\max_{M>0} G_i^{p_2, sup}(M, \phi'_i)$ as $M_i^{p_2, sup}(\phi'_i)$. Then, we can derive the upper bound by following a similar procedure stated in the proof of Proposition 3 of Chen (2010). For $i = 1, \dots, N$, let

$$M_i(\phi'_i) = \min \left\{ M \in N_+ : M \geq M_i^{p_2, sup}(\phi'_i), G_i^{p_2, sup}(M, \phi'_i) \leq \sum_{n=0}^{N-i} \delta^n \cdot \max_{0 < y \leq M} \pi(y, M, \phi_i) \right\}. \quad (\text{A.2})$$

Then, as $G_i^{p_2, sup}(M, \phi'_i)$ is concave, $G_i^{p_2, sup}(M, \phi'_i) \leq \sum_{n=0}^{N-i} \delta^n \cdot \max_{0 < y \leq M} \pi(y, M, \phi_i)$ for any $M \geq M_i(\phi'_i)$. It can be easily verified that the comparison results stated in Proposition 2.7 still hold in this two-decision-variable problem. And thus, we can get that for any $0 < y \leq M$ ($M \geq M_i(\phi'_i)$),

$$G_i^{scen}(y, M, \phi_i) \leq G_i^{p_2, sup}(M, \phi'_i) \leq \sum_{n=0}^{N-i} \delta^n \cdot \max_{0 < y \leq M} \pi(y, M, \phi_i) \leq v_i^{scen}(\phi_i).$$

This implies that $M_i(\phi'_i)$ must be an upper bound on the Bayesian optimal inventory levels in period i ($i = 1, \dots, N$) under all the information scenarios, which is formally stated in the following proposition and can be easily computed.

Proposition A.3. *For period i ($i = 1, \dots, N$), given the prior distribution $\phi_i(\theta, \alpha)$, the Bayesian optimal inventory levels y_i^{scen} and M_i^{scen} ($scen \in \{\mathcal{OS}, \mathcal{OT}, \mathcal{US}, \mathcal{UT}\}$) are bounded above by $M_i(\phi'_i)$ stated in (A.2).*

A.2 Proofs for Chapter 2 and Appendix A.1

Proof of Proposition 2.1. With some derivation effort, we can get that

$$\begin{aligned} & \pi(y+1|\theta, \alpha) - \pi(y|\theta, \alpha) \\ &= (p_1 - p_2) \sum_{i=y+1}^{+\infty} f_1(i|\theta) \\ & \quad + (1 - \alpha)p_2 \sum_{i=y+1}^{+\infty} \left\{ \sum_{j=0}^{M-y-1} \binom{i-y-1}{j} \alpha^j (1 - \alpha)^{i-y-j-1} \left[\sum_{k=0}^{M-y-j-1} f_{12}(i, k|\theta) \right] \right\} \end{aligned}$$

$$= (p_1 - p_2) \sum_{i=y+1}^{+\infty} f_1(i|\theta) + (1 - \alpha)p_2 \sum_{i=y+1}^{+\infty} f_1(i)Pr(K + D_2 \leq M - y - 1|D_1 = i, \theta, \alpha),$$

where $Pr(\cdot)$ denotes probability. Consider two substitution probabilities α^1 and α^2 ($\alpha^1 < \alpha^2$), and the corresponding variables of K are K^1 and K^2 , respectively. It is obvious that the variable $K^2 + D_2$ is stochastically larger than $K^1 + D_2$. And thus, we get that

$$Pr(K^1 + D_2 \leq M - y - 1|D_1 = i, \theta, \alpha^1) \geq Pr(K^2 + D_2 \leq M - y - 1|D_1 = i, \theta, \alpha^2).$$

Then, it is easy to verify that $\frac{\partial[\pi(y+1|\theta, \alpha) - \pi(y|\theta, \alpha)]}{\partial \alpha} < 0$, based on which the result in Proposition 2.1 can be got.

Lemma A.2. For any given $\theta \in \Theta$, $\alpha \in [0, 1]$, and $0 < y < M$, the likelihood functions $f_{\mathcal{OT}}^y(x_1, x_2|\theta, \alpha)$ and $f_{\mathcal{OT}}^{y+1}(x_1, x_2|\theta, \alpha)$ satisfy the following relationship:

$$f_{\mathcal{OT}}^y(x_1, x_2|\theta, \alpha) = \begin{cases} f_{\mathcal{OT}}^{y+1}(x_1, x_2|\theta, \alpha), & \text{if } x_1 \leq y; \\ (1 - \alpha)f_{\mathcal{OT}}^{y+1}(x_1, x_2|\theta, \alpha), & \text{if } x_1 > y \text{ and } x_2 = 0; \\ (1 - \alpha)f_{\mathcal{OT}}^{y+1}(x_1, x_2|\theta, \alpha) + \alpha f_{\mathcal{OT}}^{y+1}(x_1, x_2 - 1|\theta, \alpha), & \text{if } x_1 > y \text{ and } x_2 > 0. \end{cases}$$

Proof of Lemma A.2. Here, we only investigate the case where $x_1 > y$ and $x_2 > 0$, since the results under other cases can be easily obtained. When $x_1 > y$ and $x_2 > 0$, we can show that

$$\begin{aligned} f_{\mathcal{OT}}^y(x_1, x_2|\theta, \alpha) &= \sum_{i=0}^{(x_1-y)\wedge x_2} \binom{x_1-y}{i} \alpha^i (1 - \alpha)^{x_1-y-i} f_{12}(x_1, x_2 - i|\theta) \\ &= \sum_{i=0}^{(x_1-y)\wedge x_2} \binom{x_1-y-1}{i} \alpha^i (1 - \alpha)^{x_1-y-i} f_{12}(x_1, x_2 - i|\theta) \\ &\quad + \sum_{i=0}^{(x_1-y)\wedge x_2} \binom{x_1-y-1}{i-1} \alpha^i (1 - \alpha)^{x_1-y-i} f_{12}(x_1, x_2 - i|\theta) \\ &= \sum_{i=0}^{(x_1-y-1)\wedge x_2} \binom{x_1-y-1}{i} \alpha^i (1 - \alpha)^{x_1-y-i} f_{12}(x_1, x_2 - i|\theta) \\ &\quad + \sum_{i=0}^{(x_1-y-1)\wedge (x_2-1)} \binom{x_1-y-1}{i} \alpha^{i+1} (1 - \alpha)^{x_1-y-i-1} f_{12}(x_1, x_2 - i - 1|\theta) \end{aligned}$$

$$\begin{aligned}
&= (1 - \alpha) \sum_{i=0}^{(x_1-y-1) \wedge x_2} \binom{x_1-y-1}{i} \alpha^i (1 - \alpha)^{x_1-y-i-1} f_{12}(x_1, x_2 - i | \theta) \\
&\quad + \alpha \sum_{i=0}^{(x_1-y-1) \wedge (x_2-1)} \binom{x_1-y-1}{i} \alpha^i (1 - \alpha)^{x_1-y-i-1} f_{12}(x_1, x_2 - i - 1 | \theta) \\
&= (1 - \alpha) f_{\mathcal{OT}}^{y+1}(x_1, x_2 | \theta, \alpha) + \alpha f_{\mathcal{OT}}^{y+1}(x_1, x_2 - 1 | \theta, \alpha).
\end{aligned}$$

Proof of Lemma 2.1. We prove the proposition for the case $\alpha > 0$. We can easily get the result for the case $\alpha = 0$ by following the same steps. Since α is known, we write $f_{\mathcal{OT}}^y(x_1, x_2 | \theta)$ as shorthand for $f_{\mathcal{OT}}^y(x_1, x_2 | \theta, \alpha)$. By utilizing Lemma A.2, we first show that the followings hold: for $i = 1, \dots, N - 1$, $0 < y < M$, and any $\phi_i(\theta)$,

$$\begin{aligned}
&E_{\phi'_i(\theta)} \{v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^y(x_1, x_2 | \theta)\} \\
&\leq \begin{cases} E_{\phi'_i(\theta)} \{v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^{y+1}(x_1, x_2 | \theta)\}, & \text{if } x_1 \leq y; \\ (1 - \alpha) E_{\phi'_i(\theta)} \{v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^{y+1}(x_1, x_2 | \theta)\}, & \text{if } x_1 > y \text{ and } x_2 = 0; \\ (1 - \alpha) E_{\phi'_i(\theta)} \{v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^{y+1}(x_1, x_2 | \theta)\} + \alpha E_{\phi'_i(\theta)} \{v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^{y+1}(x_1, x_2 - 1 | \theta)\}, & \text{if } x_1 > y \text{ and } x_2 > 0. \end{cases}
\end{aligned}$$

For above inequalities, below we prove the one under the case that $x_1 > y$ and $x_2 > 0$ for illustration purpose only. Similarly, we can easily show that other ones also hold.

According to the backward induction, when $i = N - 1$,

$$\begin{aligned}
&E_{\phi_{N-1}(\theta)} \{v_N^{\mathcal{OT}}(\phi'_N) f_{\mathcal{OT}}^y(x_1, x_2 | \theta)\} \\
&= \int_{\Theta} \max_{0 < y' \leq M} \left\{ \int_{\Theta} \pi(y' | \theta') \phi'_N(\theta' | x_1, x_2, y, \phi_{N-1}) d\theta' \right\} f_{\mathcal{OT}}^y(x_1, x_2 | \theta) \phi_{N-1}(\theta) d\theta \\
&= \max_{0 < y' \leq M} \left\{ \int_{\Theta} \pi(y' | \theta') \frac{f_{\mathcal{OT}}^y(x_1, x_2 | \theta') \phi_{N-1}(\theta')}{\int_{\Theta} f_{\mathcal{OT}}^y(x_1, x_2 | \theta) \phi_{N-1}(\theta) d\theta} d\theta' \right\} \cdot \int_{\Theta} f_{\mathcal{OT}}^y(x_1, x_2 | \theta) \phi_{N-1}(\theta) d\theta \\
&= \max_{0 < y' \leq M} \int_{\Theta} \pi(y' | \theta) f_{\mathcal{OT}}^y(x_1, x_2 | \theta) \phi_{N-1}(\theta) d\theta \\
&= \max_{0 < y' \leq M} \int_{\Theta} \pi(y' | \theta) [(1 - \alpha) f_{\mathcal{OT}}^{y+1}(x_1, x_2 | \theta) + \alpha f_{\mathcal{OT}}^{y+1}(x_1, x_2 - 1 | \theta)] \phi_{N-1}(\theta) d\theta \\
&\leq (1 - \alpha) \max_{0 < y' \leq M} \int_{\Theta} \pi(y' | \theta) f_{\mathcal{OT}}^{y+1}(x_1, x_2 | \theta) \phi_{N-1}(\theta) d\theta \\
&\quad + \alpha \max_{0 < y' \leq M} \int_{\Theta} \pi(y' | \theta) f_{\mathcal{OT}}^{y+1}(x_1, x_2 - 1 | \theta) \phi_{N-1}(\theta) d\theta
\end{aligned}$$

$$=(1 - \alpha)E_{\phi_{N-1}(\theta)} \{v_N^{\mathcal{OT}}(\phi'_N)f_{\mathcal{OT}}^{y+1}(x_1, x_2|\theta)\} + \alpha E_{\phi_{N-1}(\theta)} \{v_N^{\mathcal{OT}}(\phi'_N)f_{\mathcal{OT}}^{y+1}(x_1, x_2 - 1|\theta)\}.$$

Assume the result holds for period $i + 1$ ($i = 1, \dots, N - 2$). Next, we check that for period i as follows.

$$\begin{aligned} & E_{\phi_i(\theta)} \{v_{i+1}^{\mathcal{OT}}(\phi'_{i+1})f_{\mathcal{OT}}^y(x_1, x_2|\theta)\} \\ &= \max_{0 < y' \leq M} \left\{ \int_{\Theta} \pi(y'|\theta)f_{\mathcal{OT}}^y(x_1, x_2|\theta)\phi_i(\theta)d\theta \right. \\ & \quad \left. + \delta \sum_{x'_1} \sum_{x'_2} \int_{\Theta} v_{i+2}^{\mathcal{OT}}(\phi'_{i+2}(\theta'|x_1, x_2, y|x_1, x_2, y|\phi_i))f_{\mathcal{OT}}^y(x'_1, x'_2|\theta)f_{\mathcal{OT}}^y(x_1, x_2|\theta)\phi_i(\theta)d\theta \right\} \\ &= \max_{0 < y' \leq M} \left\{ \int_{\Theta} \pi(y'|\theta)f_{\mathcal{OT}}^y(x_1, x_2|\theta)\phi_i(\theta)d\theta \right. \\ & \quad \left. + \delta \sum_{x'_1} \sum_{x'_2} \int_{\Theta} v_{i+2}^{\mathcal{OT}}(\phi'_{i+2}(\theta'|x_1, x_2, y|x_1, x_2, y|\phi_i))f_{\mathcal{OT}}^y(x_1, x_2|\theta)f_{\mathcal{OT}}^y(x'_1, x'_2|\theta)\phi_i(\theta)d\theta \right\} \\ &= \max_{0 < y' \leq M} \left\{ \int_{\Theta} \pi(y'|\theta)f_{\mathcal{OT}}^y(x_1, x_2|\theta)\phi_i(\theta)d\theta \right. \\ & \quad \left. + \delta \sum_{x'_1} \sum_{x'_2} E_{\phi'_{i+1}(\theta|x_1, x_2, y, \phi_i)} [v_{i+2}^{\mathcal{OT}}(\phi'_{i+2}(\theta'|x_1, x_2, y, \phi'_{i+1}))f_{\mathcal{OT}}^y(x_1, x_2|\theta)] \right. \\ & \quad \left. \cdot \int_{\Theta} f_{\mathcal{OT}}^y(x'_1, x'_2|\theta)\phi_i(\theta)d\theta \right\} \\ &\leq \max_{0 < y' \leq M} \left\{ \int_{\Theta} \pi(y'|\theta)f_{\mathcal{OT}}^y(x_1, x_2|\theta)\phi_i(\theta)d\theta \right. \\ & \quad \left. + \delta \sum_{x'_1} \sum_{x'_2} E_{\phi'_{i+1}(\theta|x_1, x_2, y, \phi_i)} [(1 - \alpha)v_{i+2}^{\mathcal{OT}}(\phi'_{i+2}(\theta'|x_1, x_2, y + 1, \phi'_{i+1}))f_{\mathcal{OT}}^{y+1}(x_1, x_2|\theta) \right. \\ & \quad \left. + \alpha v_{i+2}^{\mathcal{OT}}(\phi'_{i+2}(\theta'|x_1, x_2 - 1, y + 1, \phi'_{i+1}))f_{\mathcal{OT}}^{y+1}(x_1, x_2 - 1|\theta)] \int_{\Theta} f_{\mathcal{OT}}^y(x'_1, x'_2|\theta)\phi_i(\theta)d\theta \right\} \\ &= \max_{0 < y' \leq M} \left\{ \int_{\Theta} \pi(y'|\theta)[(1 - \alpha)f_{\mathcal{OT}}^{y+1}(x_1, x_2|\theta) + \alpha f_{\mathcal{OT}}^{y+1}(x_1, x_2 - 1|\theta)]\phi_i(\theta)d\theta \right. \\ & \quad \left. + \delta \sum_{x'_1} \sum_{x'_2} (1 - \alpha) \int_{\Theta} v_{i+2}^{\mathcal{OT}}(\phi'_{i+2}(\theta'|x_1, x_2, y|x_1, x_2, y + 1|\phi_i)) \right. \\ & \quad \left. \cdot f_{\mathcal{OT}}^y(x'_1, x'_2|\theta)f_{\mathcal{OT}}^{y+1}(x_1, x_2|\theta)\phi_i(\theta)d\theta \right. \\ & \quad \left. + \delta \sum_{x'_1} \sum_{x'_2} \alpha \int_{\Theta} v_{i+2}^{\mathcal{OT}}(\phi'_{i+2}(\theta'|x_1, x_2, y|x_1, x_2 - 1, y + 1|\phi_i)) \right. \\ & \quad \left. \cdot f_{\mathcal{OT}}^y(x'_1, x'_2|\theta)f_{\mathcal{OT}}^{y+1}(x_1, x_2 - 1|\theta)\phi_i(\theta)d\theta \right\} \\ &\leq (1 - \alpha) \max_{0 < y' \leq M} \left\{ \int_{\Theta} \pi(y'|\theta)f_{\mathcal{OT}}^{y+1}(x_1, x_2|\theta)\phi_i(\theta)d\theta \right. \end{aligned}$$

$$\begin{aligned}
& + \delta \sum_{x'_1} \sum_{x'_2} \int_{\Theta} v_{i+2}^{\mathcal{OT}}(\phi'_{i+2}(\theta'|x_1, x_2, y|x_1, x_2, y+1|\phi_i)) f_{\mathcal{OT}}^y(x'_1, x'_2|\theta) f_{\mathcal{OT}}^{y+1}(x_1, x_2|\theta) \phi_i(\theta) d\theta \Big\} \\
& + \alpha \max_{0 < y' \leq M} \left\{ \int_{\Theta} \pi(y'|\theta) f_{\mathcal{OT}}^{y+1}(x_1, x_2 - 1|\theta) d\theta \right. \\
& + \delta \sum_{x'_1} \sum_{x'_2} \int_{\Theta} v_{i+2}^{\mathcal{OT}}(\phi'_{i+2}(\theta'|x_1, x_2, y|x_1, x_2 - 1, y+1|\phi_i)) \\
& \quad \left. \cdot f_{\mathcal{OT}}^y(x'_1, x'_2|\theta) f_{\mathcal{OT}}^{y+1}(x_1, x_2 - 1|\theta) \phi_i(\theta) d\theta \right\} \\
& = (1 - \alpha) E_{\phi_i(\theta)} \{ v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^{y+1}(x_1, x_2|\theta) \} + \alpha E_{\phi_i(\theta)} \{ v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^{y+1}(x_1, x_2 - 1|\theta) \}.
\end{aligned}$$

Now, we can show that

$$\begin{aligned}
& E_{\phi_i(\theta)} \left\{ \sum_{x_1} \sum_{x_2} v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^y(x_1, x_2|\theta) \right\} \\
& = \sum_{x_1=0}^y \sum_{x_2=0}^{+\infty} E_{\phi_i(\theta)} \{ v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^y(x_1, x_2|\theta) \} + \sum_{x_1=y+1}^{+\infty} E_{\phi_i(\theta)} \{ v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^y(x_1, 0|\theta) \} \\
& \quad + \sum_{x_1=y+1}^{+\infty} \sum_{x_2=1}^{+\infty} E_{\phi_i(\theta)} \{ v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^y(x_1, x_2|\theta) \} \\
& \leq (1 - \alpha) \sum_{x_1=y+1}^{+\infty} \sum_{x_2=1}^{+\infty} E_{\phi_i(\theta)} \{ v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^{y+1}(x_1, x_2|\theta) \} \\
& \quad + \alpha \sum_{x_1=y+1}^{+\infty} \sum_{x_2=1}^{+\infty} E_{\phi_i(\theta)} \{ v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^{y+1}(x_1, x_2 - 1|\theta) \} \\
& \quad + \sum_{x_1=0}^y \sum_{x_2=0}^{+\infty} E_{\phi_i(\theta)} \{ v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^{y+1}(x_1, x_2|\theta) \} \\
& \quad + (1 - \alpha) \sum_{x_1=y+1}^{+\infty} E_{\phi_i(\theta)} \{ v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^{y+1}(x_1, 0|\theta) \} \\
& = \sum_{x_1=0}^{+\infty} \sum_{x_2=0}^{+\infty} E_{\phi_i(\theta)} \{ v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^{y+1}(x_1, x_2|\theta) \} \\
& = E_{\phi_i(\theta)} \left\{ \sum_{x_1} \sum_{x_2} v_{i+1}^{\mathcal{OT}}(\phi'_{i+1}) f_{\mathcal{OT}}^{y+1}(x_1, x_2|\theta) \right\}.
\end{aligned}$$

Proof of Proposition 2.4. Under the $\widetilde{\mathcal{US}}$ scenario, we first assume that the demand parameter θ_2 is known. Then, we can write $f_{\widetilde{\mathcal{US}}}^y(s_1, s_{22}|\theta_1)$ as shorthand for $f_{\widetilde{\mathcal{US}}}^y(s_1, s_{22}|\theta_1, \theta_2, \alpha)$. We first show that the followings hold: for $i = 1, \dots, N - 1$, $0 < y < M$, and any $\phi_{i,1}(\theta)$,

(a) when $s_1 < y$ and $s_{22} < M - s_1$,

$$E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{U}}\mathcal{S}}(\phi'_{i+1,1}) f_{\widetilde{\mathcal{U}}\mathcal{S}}^y(s_1, s_{22}|\theta_1) \right\} = f_2(s_{22}|\theta_2) \cdot E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{U}}\mathcal{S}}(\phi'_{i+1,1}) f_1(s_1|\theta_1) \right\};$$

(b) when $s_1 < y$ and $s_{22} = M - s_1$,

$$E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{U}}\mathcal{S}}(\phi'_{i+1,1}) f_{\widetilde{\mathcal{U}}\mathcal{S}}^y(s_1, M - s_1|\theta_1) \right\} = \left[\sum_{j=M-s_1}^{+\infty} f_2(j|\theta_2) \right] \cdot E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{U}}\mathcal{S}}(\phi'_{i+1,1}) f_1(s_1|\theta_1) \right\};$$

(c) when $s_1 = y$ and $s_{22} < M - y$,

$$\begin{aligned} E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{U}}\mathcal{S}}(\phi'_{i+1,1}) f_{\widetilde{\mathcal{U}}\mathcal{S}}^y(y, s_{22}|\theta_1) \right\} &= f_2(s_{22}|\theta_2) \cdot E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{U}}\mathcal{S}}(\phi'_{i+1,1}) \sum_{i=y}^{+\infty} f_1(i|\theta_1) \right\} \\ &\leq f_2(s_{22}|\theta_2) E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{U}}\mathcal{S}}(\phi'_{i+1,1}) f_1(y|\theta_1) \right\} + f_2(s_{22}|\theta_2) E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{U}}\mathcal{S}}(\phi'_{i+1,1}) \sum_{i=y+1}^{+\infty} f_1(i|\theta_1) \right\}; \end{aligned}$$

(d) when $s_1 = y$ and $s_{22} = M - y$,

$$\begin{aligned} &E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{U}}\mathcal{S}}(\phi'_{i+1,1}) f_{\widetilde{\mathcal{U}}\mathcal{S}}^y(y, M - y|\theta_1) \right\} \\ &= \left[\sum_{j=M-y}^{+\infty} f_2(j|\theta_2) \right] \cdot E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{U}}\mathcal{S}}(\phi'_{i+1,1}) \sum_{i=y}^{+\infty} f_1(i|\theta_1) \right\} \\ &\leq \left[\sum_{j=M-y}^{+\infty} f_2(j|\theta_2) \right] E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{U}}\mathcal{S}}(\phi'_{i+1,1}) f_1(y|\theta_1) \right\} \\ &\quad + \left[\sum_{j=M-y}^{+\infty} f_2(j|\theta_2) \right] E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{U}}\mathcal{S}}(\phi'_{i+1,1}) \sum_{i=y+1}^{+\infty} f_1(i|\theta_1) \right\}. \end{aligned}$$

The above four relationships can be proved following the procedure in Lemma 2.1 through splitting off the likelihood function $f_{\widetilde{\mathcal{U}}\mathcal{S}}^y(s_1, s_{22}|\theta_1)$. Here, we omit the details. Then, with some derivation effort, we can get that

$$\begin{aligned} &E_{\phi_{i,1}(\theta_1)} \left\{ \sum_{s_1} \sum_{s_{22}} v_{i+1}^{\widetilde{\mathcal{U}}\mathcal{S}}(\phi'_{i+1,1}) f_{\widetilde{\mathcal{U}}\mathcal{S}}^y(s_1, s_{22}|\theta_1) \right\} \\ &\leq \left[\sum_{j=M-y}^{+\infty} f_2(j|\theta_2) \right] E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{U}}\mathcal{S}}(\phi'_{i+1,1}) f_1(y|\theta_1) \right\} \end{aligned}$$

$$\begin{aligned}
& + \left[\sum_{j=M-y}^{+\infty} f_2(j|\theta_2) \right] E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{US}}}(\phi'_{i+1,1}) \sum_{i=y+1}^{+\infty} f_1(i|\theta_1) \right\} \\
& + \sum_{s_{22}=0}^{M-y-1} f_2(s_{22}|\theta_2) E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{US}}}(\phi'_{i+1,1}) f_1(y|\theta_1) \right\} \\
& + \sum_{s_{22}=0}^{M-y-1} f_2(s_{22}|\theta_2) E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{US}}}(\phi'_{i+1,1}) \sum_{i=y+1}^{+\infty} f_1(i|\theta_1) \right\} \\
& + \sum_{s_1=0}^{y-1} \sum_{s_{22}=0}^{M-s_1-1} f_2(s_{22}|\theta_2) E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{US}}}(\phi'_{i+1,1}) f_1(s_1|\theta_1) \right\} \\
& + \sum_{s_1=0}^{y-1} \left[\sum_{j=M-s_1}^{+\infty} f_2(j|\theta_2) \right] E_{\phi'_{i,1}(\theta_1)} \left\{ v_{i+1}^{\widetilde{\mathcal{US}}}(\phi'_{i+1,1}) f_1(s_1|\theta_1) \right\} \\
& = E_{\phi_{i,1}(\theta_1)} \left\{ \sum_{s_1} \sum_{s_{22}} v_{i+1}^{\widetilde{\mathcal{US}}}(\phi'_{i+1,1}) f_{\widetilde{\mathcal{US}}}^{y+1}(s_1, s_{22}|\theta_1) \right\},
\end{aligned}$$

which implies that $y_i^{\widetilde{\mathcal{US}}} \geq y_i^m$.

Following the same procedure, we can prove that under the $\widetilde{\mathcal{US}}$ scenario, when the demand parameter θ_1 is known, for any period i ($i = 1, \dots, N$), and given the same prior distribution $\phi'_{i,2}(\theta_2)$, learning the demand parameter θ_2 requires $y_i^{\widetilde{\mathcal{US}}} \leq y_i^m$. In this case, we write $f_{\widetilde{\mathcal{US}}}^y(s_1, s_{22}|\theta_2)$ as shorthand for $f_{\widetilde{\mathcal{US}}}^y(s_1, s_{22}|\theta_1, \theta_2, \alpha)$. Now, the following four relationships are needed: for $i = 1, \dots, N-1$, $1 < y \leq M$, and any $\phi_{i,2}(\theta)$,

(a') when $s_1 < y$ and $s_{22} < M - s_1$,

$$E_{\phi'_{i,2}(\theta_2)} \left\{ v_{i+1}^{\widetilde{\mathcal{US}}}(\phi'_{i+1,2}) f_{\widetilde{\mathcal{US}}}^y(s_1, s_{22}|\theta_2) \right\} = f_1(s_1|\theta_1) \cdot E_{\phi'_{i,2}(\theta_2)} \left\{ v_{i+1}^{\widetilde{\mathcal{US}}}(\phi'_{i+1,2}) f_2(s_{22}|\theta_2) \right\};$$

(b') when $s_1 < y$ and $s_{22} = M - s_1$,

$$E_{\phi'_{i,2}(\theta_2)} \left\{ v_{i+1}^{\widetilde{\mathcal{US}}}(\phi'_{i+1,2}) f_{\widetilde{\mathcal{US}}}^y(s_1, M - s_1|\theta_2) \right\} = f_1(s_1|\theta_1) \cdot E_{\phi'_{i,2}(\theta_2)} \left\{ v_{i+1}^{\widetilde{\mathcal{US}}}(\phi'_{i+1,2}) \sum_{j=M-s_1}^{+\infty} f_2(j|\theta_2) \right\};$$

(c') when $s_1 = y$ and $s_{22} < M - y$,

$$E_{\phi'_{i,2}(\theta_2)} \left\{ v_{i+1}^{\widetilde{\mathcal{US}}}(\phi'_{i+1,2}) f_{\widetilde{\mathcal{US}}}^y(y, s_{22}|\theta_2) \right\} = \left[\sum_{i=y}^{+\infty} f_1(i|\theta_1) \right] \cdot E_{\phi'_{i,2}(\theta_2)} \left\{ v_{i+1}^{\widetilde{\mathcal{US}}}(\phi'_{i+1,2}) f_2(s_{22}|\theta_2) \right\};$$

(d') when $s_1 = y$ and $s_{22} = M - y$,

$$\begin{aligned}
& E_{\phi'_{i,2}(\theta_2)} \left\{ v_{i+1}^{\widetilde{US}}(\phi'_{i+1,2}) f_{\widetilde{US}}^y(y, M - y | \theta_2) \right\} \\
&= \left[\sum_{i=y}^{+\infty} f_1(i | \theta_1) \right] \cdot E_{\phi'_{i,2}(\theta_2)} \left\{ v_{i+1}^{\widetilde{US}}(\phi'_{i+1,2}) \sum_{j=M-y}^{+\infty} f_2(j | \theta_2) \right\} \\
&\leq \left[\sum_{i=y}^{+\infty} f_1(i | \theta_1) \right] \cdot E_{\phi'_{i,2}(\theta_2)} \left\{ v_{i+1}^{\widetilde{US}}(\phi'_{i+1,2}) f_2(M - y | \theta_2) \right\} \\
&\quad + \left[\sum_{i=y}^{+\infty} f_1(i | \theta_1) \right] \cdot E_{\phi'_{i,2}(\theta_2)} \left\{ v_{i+1}^{\widetilde{US}}(\phi'_{i+1,2}) \sum_{j=M-y+1}^{+\infty} f_2(j | \theta_2) \right\}.
\end{aligned}$$

Proof of Proposition 2.5. Under the setting of Proposition 2.5, since D_2 is always 0 and α is known, we can write $f_{US}^y(s_1, s_{21} | \theta_1)$ as shorthand for $f_{US}^y(s_1, s_{21}, s_{22} | \theta, \alpha)$. As $M = 2$, the value of y can only be 1 or 2. We can get the following relationship between two likelihood functions $f_{US}^1(s_1, s_{21} | \theta_1)$ and $f_{US}^2(s_1, s_{21} | \theta_1)$ based on the assumption that $f_1(1 | \theta_1) = 0$ for all $\theta_1 \in \Theta_1$:

$$\begin{aligned}
& f_{US}^2(s_1, s_{21} | \theta_1) \\
&= \begin{cases} f_1(0 | \theta_1), & \text{if } s_1 = 0 \text{ and } s_{21} = 0, \\ \sum_{i=2}^{+\infty} f_1(i | \theta_1), & \text{if } s_1 = 2 \text{ and } s_{21} = 0, \end{cases} \\
&= \begin{cases} f_1(0 | \theta_1), & \text{if } s_1 = 0 \text{ and } s_{21} = 0, \\ \sum_{i=2}^{+\infty} f_1(i | \theta_1) (1 - \alpha)^{i-1} + \sum_{i=2}^{+\infty} \sum_{k=1}^{i-1} f_1(i | \theta_1) \binom{i-1}{k} \alpha^k (1 - \alpha)^{i-k-1}, & \text{if } s_1 = 2 \text{ and } s_{21} = 0, \end{cases} \\
&= \begin{cases} f_{US}^1(0, 0 | \theta_1), & \text{if } s_1 = 0 \text{ and } s_{21} = 0, \\ f_{US}^1(1, 0 | \theta_1) + f_{US}^1(1, 1 | \theta_1), & \text{if } s_1 = 2 \text{ and } s_{21} = 0, \end{cases}
\end{aligned}$$

Following the same logic in Lemma 2.1 and Proposition 2.3, we can get the ‘‘stock less’’ result.

Proof of Inequality (2.7). For $i = 1, \dots, N - 1$, we have that

$$\begin{aligned}
& G_i^{\mathcal{OS}}(y + 1, \phi_i) - G_i^{\mathcal{OS}}(y, \phi_i) \\
&= E_{\phi_i(\theta, \alpha)} \left\{ \pi(y + 1 | \theta, \alpha) - \pi(y | \theta, \alpha) \right\}
\end{aligned}$$

$$\begin{aligned}
& + \delta \sum_{x_1=y+1}^{+\infty} \sum_{x_{22}=0}^{+\infty} \left[\sum_{x_{21}=0}^{x_1-y-1} v_{i+1}^{\mathcal{OS}}(\phi_{i+1}) \binom{x_1-y-1}{x_{21}} \alpha^{x_{21}} (1-\alpha)^{x_1-y-1-x_{21}} \right. \\
& \quad \left. - \sum_{x_{21}=0}^{x_1-y} v_{i+1}^{\mathcal{OS}}(\phi_{i+1}) \binom{x_1-y}{x_{21}} \alpha^{x_{21}} (1-\alpha)^{x_1-y-x_{21}} \right] f_{12}(x_1, x_{22}|\theta) \Big\}.
\end{aligned}$$

Similarly to the proof of Proposition 4 in Chen and Plambeck (2008), we can get that for any $i = 1, \dots, N-1$ and prior $\phi_i(\theta, \alpha)$,

$$\begin{aligned}
& E_{\phi_i(\theta, \alpha)} \left\{ \sum_{x_1=y+1}^{+\infty} \sum_{x_{22}=0}^{+\infty} \sum_{x_{21}=0}^{x_1-y-1} v_{i+1}^{\mathcal{OS}}(\phi_{i+1}) \binom{x_1-y-1}{x_{21}} \alpha^{x_{21}} (1-\alpha)^{x_1-y-1-x_{21}} f_{12}(x_1, x_{22}|\theta) \right\} \\
& \leq E_{\phi_i(\theta, \alpha)} \left\{ \sum_{x_1=y+1}^{+\infty} \sum_{x_{22}=0}^{+\infty} \sum_{x_{21}=0}^{x_1-y} v_{i+1}^{\mathcal{OS}}(\phi_{i+1}) \binom{x_1-y}{x_{21}} \alpha^{x_{21}} (1-\alpha)^{x_1-y-x_{21}} f_{12}(x_1, x_{22}|\theta) \right\}.
\end{aligned}$$

Thus, the inequality (2.7) holds.

Proof of Proposition 2.7. We can see that (a) implies (b) by taking the maximum over the inequalities in (a). So, we only need to prove (a) here. Moreover, below we focus on proving

$$\sum_{n=0}^{N-i} \delta^n \cdot \pi(y, \phi_i) \leq G_i^{\mathcal{UT}}(y, \phi_i) \leq G_i^{\mathcal{OT}}(y, \phi_i) \leq G_i^{\mathcal{OS}}(y, \phi_i).$$

Similarly, we can show that $G_i^{\mathcal{UT}}(y, \phi_i) \leq G_i^{\mathcal{US}}(y, \phi_i) \leq G_i^{\mathcal{OS}}(y, \phi_i)$.

First, we use the backward induction to show that $\sum_{n=0}^{N-i} \delta^n \cdot \pi(y, \phi_i) \leq G_i^{\mathcal{UT}}(y, \phi_i)$ for $i = 1, \dots, N$. When $i = N$, it holds for sure. Assume the result holds for period $i+1$ ($i = 1, \dots, N-1$), which means that $\sum_{n=0}^{N-i-1} \delta^n \cdot \pi(y, \phi_{i+1}) \leq G_{i+1}^{\mathcal{UT}}(y, \phi_{i+1})$, and thus $\sum_{n=0}^{N-i-1} \delta^n \cdot \max_{0 < y \leq M} \pi(y, \phi_{i+1}) \leq v_{i+1}^{\mathcal{UT}}(\phi_{i+1})$. Now, for period i , we have

$$\begin{aligned}
& G_i^{\mathcal{UT}}(y, \phi_i) \\
& = E_{\phi_i(\theta, \alpha)} \left\{ \pi(y|\theta, \alpha) \right. \\
& \quad + \delta \sum_{s_1=0}^{y-1} \sum_{s_2=0}^{M-s_1-1} v_{i+1}^{\mathcal{UT}}(\phi_{i+1}) f_{\mathcal{UT}}^y(s_1, s_2|\theta, \alpha) + \delta \sum_{s_2=0}^{M-y-1} v_{i+1}^{\mathcal{UT}}(\phi_{i+1}) f_{\mathcal{UT}}^y(y, s_2|\theta, \alpha) \\
& \quad \left. + \delta \sum_{s_1=0}^{y-1} v_{i+1}^{\mathcal{UT}}(\phi_{i+1}) f_{\mathcal{UT}}^y(s_1, M-s_1|\theta, \alpha) + \delta v_{i+1}^{\mathcal{UT}}(\phi_{i+1}) f_{\mathcal{UT}}^y(y, M-y|\theta, \alpha) \right\}
\end{aligned}$$

$$\begin{aligned}
&\geq E_{\phi_i(\theta, \alpha)} \left\{ \pi(y|\theta, \alpha) + \delta \sum_{s_1=0}^{y-1} \sum_{s_2=0}^{M-s_1-1} \left[\sum_{n=0}^{N-i-1} \delta^n \cdot \max_{0 < y \leq M} \pi(y, \phi_{i+1}) \right] f_{\mathcal{UT}}^y(s_1, s_2|\theta, \alpha) \right. \\
&\quad + \delta \sum_{s_1=0}^{y-1} \left[\sum_{n=0}^{N-i-1} \delta^n \cdot \max_{0 < y \leq M} \pi(y, \phi_{i+1}) \right] f_{\mathcal{UT}}^y(s_1, M - s_1|\theta, \alpha) \\
&\quad + \delta \sum_{s_2=0}^{M-y-1} \left[\sum_{n=0}^{N-i-1} \delta^n \cdot \max_{0 < y \leq M} \pi(y, \phi_{i+1}) \right] f_{\mathcal{UT}}^y(y, s_2|\theta, \alpha) \\
&\quad \left. + \delta \left[\sum_{n=0}^{N-i-1} \delta^n \cdot \max_{0 < y \leq M} \pi(y, \phi_{i+1}) \right] f_{\mathcal{UT}}^y(y, M - y|\theta, \alpha) \right\} \\
&\geq E_{\phi_i(\theta, \alpha)} \{ \pi(y|\theta, \alpha) \} \\
&\quad + \sum_{n=1}^{N-i} \delta^k \sum_{s_1=0}^{y-1} \sum_{s_2=0}^{M-s_1-1} \pi(y, \phi_{i+1}(\theta', \alpha' | s_1, s_2, y, \phi_i)) E_{\phi_i(\theta, \alpha)} \{ f_{\mathcal{UT}}^y(s_1, s_2|\theta, \alpha) \} \\
&\quad + \sum_{n=1}^{N-i} \delta^k \sum_{s_1=0}^{y-1} \pi(y, \phi_{i+1}(\theta', \alpha' | s_1, M - s_1, y, \phi_i)) E_{\phi_i(\theta, \alpha)} \{ f_{\mathcal{UT}}^y(s_1, M - s_1|\theta, \alpha) \} \\
&\quad + \sum_{n=1}^{N-i} \delta^k \sum_{s_2=0}^{M-y-1} \pi(y, \phi_{i+1}(\theta', \alpha' | y, s_2, y, \phi_i)) E_{\phi_i(\theta, \alpha)} \{ f_{\mathcal{UT}}^y(y, s_2|\theta, \alpha) \} \\
&\quad + \sum_{n=1}^{N-i} \delta^k \pi(y, \phi_{i+1}(\theta', \alpha' | y, M - y, y, \phi_i)) E_{\phi_i(\theta, \alpha)} \{ f_{\mathcal{UT}}^y(y, M - y|\theta, \alpha) \} \\
&= \sum_{n=0}^{N-i} \delta^n \cdot \pi(y, \phi_i),
\end{aligned}$$

where the last equality is based on the law of total expectation, whose formal proof is similar to that of Lemma 1(a) in Chen (2010).

Next, we use the backward induction to show $G_i^{\mathcal{UT}}(y, \phi_i) \leq G_i^{\mathcal{OT}}(y, \phi_i)$ ($i = 1, \dots, N$). When $i = N$, it holds for sure. Assume that the result holds for period $i+1$ ($i = 1, \dots, N-1$), which means that $G_{i+1}^{\mathcal{UT}}(y, \phi_{i+1}) \leq G_{i+1}^{\mathcal{OT}}(y, \phi_{i+1})$, and thus $v_{i+1}^{\mathcal{UT}}(\phi_{i+1}) \leq v_{i+1}^{\mathcal{OT}}(\phi_{i+1})$. Then, for period i , we have

$$\begin{aligned}
&G_i^{\mathcal{UT}}(y, \phi_i) \\
&= E_{\phi_i(\theta, \alpha)} \left\{ \pi(y|\theta, \alpha) + \delta \sum_{s_1=0}^{y-1} \sum_{s_2=0}^{M-s_1-1} v_{i+1}^{\mathcal{UT}}(\phi_{i+1}) f_{\mathcal{UT}}^y(s_1, s_2|\theta, \alpha) + \delta v_{i+1}^{\mathcal{UT}}(\phi_{i+1}) f_{\mathcal{UT}}^y(y, M - y|\theta, \alpha) \right. \\
&\quad \left. + \delta \sum_{s_1=0}^{y-1} v_{i+1}^{\mathcal{UT}}(\phi_{i+1}) f_{\mathcal{UT}}^y(s_1, M - s_1|\theta, \alpha) + \delta \sum_{s_2=0}^{M-y-1} v_{i+1}^{\mathcal{UT}}(\phi_{i+1}) f_{\mathcal{UT}}^y(y, s_2|\theta, \alpha) \right\}
\end{aligned}$$

$$\begin{aligned}
&= E_{\phi_i(\theta, \alpha)} \{ \pi(y|\theta, \alpha) \} \\
&\quad + \delta \sum_{s_1=0}^{y-1} \sum_{s_2=0}^{M-s_1-1} v_{i+1}^{\mathcal{UT}}(\phi_{i+1}(\theta', \alpha' | x_1 = s_1, x_2 = s_2, y, \phi_i)) E_{\phi_i(\theta, \alpha)} \{ f_{\mathcal{UT}}^y(s_1, s_2 | \theta, \alpha) \} \\
&\quad + \delta \sum_{s_1=0}^{y-1} v_{i+1}^{\mathcal{UT}}(\phi_{i+1}(\theta', \alpha' | x_1 = s_1, x_2 \geq M - x_1, y, \phi_i)) E_{\phi_i(\theta, \alpha)} \{ f_{\mathcal{UT}}^y(s_1, M - s_1 | \theta, \alpha) \} \\
&\quad + \delta \sum_{s_2=0}^{M-y-1} v_{i+1}^{\mathcal{UT}}(\phi_{i+1}(\theta', \alpha' | x_1 \geq y, x_2 = s_2, y, \phi_i)) E_{\phi_i(\theta, \alpha)} \{ f_{\mathcal{UT}}^y(y, s_2 | \theta, \alpha) \} \\
&\quad + \delta v_{i+1}^{\mathcal{UT}}(\phi_{i+1}(\theta', \alpha' | x_1 \geq y, x_2 \geq M - y, y, \phi_i)) E_{\phi_i(\theta, \alpha)} \{ f_{\mathcal{UT}}^y(y, M - y | \theta, \alpha) \} \\
&\leq E_{\phi_i(\theta, \alpha)} \{ \pi(y|\theta, \alpha) \} \\
&\quad + \delta \sum_{s_1=0}^{y-1} \sum_{s_2=0}^{M-s_1-1} v_{i+1}^{\mathcal{OT}}(\phi_{i+1}(\theta', \alpha' | x_1 = s_1, x_2 = s_2, y, \phi_i)) E_{\phi_i(\theta, \alpha)} \{ f_{\mathcal{OT}}^y(s_1, s_2 | \theta, \alpha) \} \\
&\quad + \delta \sum_{s_1=0}^{y-1} v_{i+1}^{\mathcal{OT}}(\phi_{i+1}(\theta', \alpha' | x_1 = s_1, x_2 \geq M - x_1, y, \phi_i)) E_{\phi_i(\theta, \alpha)} \{ f_{\mathcal{OT}}^y(s_1, M - s_1 | \theta, \alpha) \} \\
&\quad + \delta \sum_{s_2=0}^{M-y-1} v_{i+1}^{\mathcal{OT}}(\phi_{i+1}(\theta', \alpha' | x_1 \geq y, x_2 = s_2, y, \phi_i)) E_{\phi_i(\theta, \alpha)} \{ f_{\mathcal{OT}}^y(y, s_2 | \theta, \alpha) \} \\
&\quad + \delta v_{i+1}^{\mathcal{OT}}(\phi_{i+1}(\theta', \alpha' | x_1 \geq y, x_2 \geq M - y, y, \phi_i)) E_{\phi_i(\theta, \alpha)} \{ f_{\mathcal{OT}}^y(y, M - y | \theta, \alpha) \} \\
&\leq E_{\phi_i(\theta, \alpha)} \{ \pi(y|\theta, \alpha) \} + \delta \sum_{x_1=0}^{y-1} \sum_{x_2=0}^{M-x_1-1} v_{i+1}^{\mathcal{OT}}(\phi_{i+1}(\theta', \alpha' | x_1, x_2, y, \phi_i)) E_{\phi_i(\theta, \alpha)} \{ f_{\mathcal{OT}}^y(x_1, x_2 | \theta, \alpha) \} \\
&\quad + \delta \sum_{x_1=0}^{y-1} \sum_{x_2=M-x_1}^{+\infty} v_{i+1}^{\mathcal{OT}}(\phi_{i+1}(\theta', \alpha' | x_1, x_2, y, \phi_i)) E_{\phi_i(\theta, \alpha)} \{ f_{\mathcal{OT}}^y(x_1, x_2 | \theta, \alpha) \} \\
&\quad + \delta \sum_{x_1=y}^{+\infty} \sum_{x_2=0}^{M-y-1} v_{i+1}^{\mathcal{OT}}(\phi_{i+1}(\theta', \alpha' | x_1, x_2, y, \phi_i)) E_{\phi_i(\theta, \alpha)} \{ f_{\mathcal{OT}}^y(x_1, x_2 | \theta, \alpha) \} \\
&\quad + \delta \sum_{x_1=y}^{+\infty} \sum_{x_2=M-y}^{+\infty} v_{i+1}^{\mathcal{OT}}(\phi_{i+1}(\theta', \alpha' | x_1, x_2, y, \phi_i)) E_{\phi_i(\theta, \alpha)} \{ f_{\mathcal{OT}}^y(x_1, x_2 | \theta, \alpha) \} \\
&= G_i^{\mathcal{OT}}(y, \phi_i),
\end{aligned}$$

where the last inequality can be formally proved by following the procedure stated in the proof of Lemma 1(b) in Chen (2010). Here, we omit the details. Similarly, we can prove that $G_i^{\mathcal{OT}}(y, \phi_i) \leq G_i^{\mathcal{OS}}(y, \phi_i)$ ($i = 1, \dots, N$).

Proof of Proposition A.1. We can easily obtain the results by using the expression of $\pi(y+1, M|\theta, \alpha) - \pi(y, M|\theta, \alpha)$ in the proof of Proposition 2.1. We thus omit the detail.

Proof of Proposition A.2. We prove the results by contradiction under *scen* scenario ($scen \in \{\mathcal{US}, \mathcal{UT}\}$). For (a), assume that if $y_i^{scen} < y_i^m$, $M_i^{scen} - y_i^{scen} \leq M_i^m - y_i^m$ ($i = 1, \dots, N$), which implies $M_i^{scen} < M_i^m$. First, we can increase y_i^{scen} to y_i^m and keep the difference between M_i^{scen} and y_i^{scen} unchanged through increasing both y_i^{scen} and M_i^{scen} by the same amount. Then, with $y_i^{scen} = y_i^m$, we increase M_i^{scen} to M_i^m . In this way, we increase y_i^{scen} to y_i^m and increase M_i^{scen} to M_i^m simultaneously. Thus, according to Lemma A.1, we have

$$\begin{aligned} E_{\phi'_i(\theta)} & \left\{ \sum_{s_1} \sum_{s_2} v_{i+1}^{scen}(\phi'_{i+1}) f_{scen}^{y_i^m, M_i^m}(s_1, s_2 | \theta, \alpha) \right\} \\ & \geq E_{\phi'_i(\theta)} \left\{ \sum_{s_1} \sum_{s_2} v_{i+1}^{scen}(\phi'_{i+1}) f_{scen}^{y_i^{scen}, M_i^{scen}}(s_1, s_2 | \theta, \alpha) \right\}. \end{aligned}$$

In addition, based on the results in the baseline one-period model, we can obtain that $E_{\phi'_i(\theta)} \{\pi(y_i^m, M_i^m | \theta, \alpha)\} > E_{\phi'_i(\theta)} \{\pi(y_i^{scen}, M_i^{scen} | \theta, \alpha)\}$ according to the optimality of y_i^m and M_i^m . Combining the above results, we have $G_i^{scen}(y_i^m, M_i^m, \phi'_i) > G_i^{scen}(y_i^{scen}, M_i^{scen}, \phi'_i)$, which contradicts the Bayesian optimality of y_i^{scen} and M_i^{scen} .

Similarly, we can prove that if $M_i^{scen} - y_i^{scen} < M_i^m - y_i^m$, $y_i^{scen} > y_i^m$ ($i = 1, \dots, N$).

Appendix B

Supplements and Proofs for Chapter 3

B.1 Sequential Equilibria Analysis of Hybrid and Mixed Strategies with Homogeneous Customers

In both the hybrid and mixed strategies, two signals R and C are on the equilibrium path. In the basic model where all customers are uninformed, the effective arrival rate to an observable or unobservable queue remains the same to both types of the server under a given belief. When one type of the server randomizes, he must be indifferent between R and C . This means that the condition $\lambda^O(\delta^R) = \lambda^U(\delta^C)$ must be satisfied. Below, we consider various hybrid and mixed strategies one by one.

(1) $f(R|H) = 1$ and $0 < f(R|L) < 1$. Only the low-quality server sends the signal C with positive probability, and thus $\delta^C = 0$. When the customers see the signal R , they update their belief as $\delta^R = \frac{\delta}{\delta + (1-\delta)f(R|L)}$ by the Bayes' rule. In this case, we have $\delta < \delta^R < 1$. When $\lambda^O(\delta^R) = \lambda^U(0)$, the high-quality server has no incentive to deviate to C . Also, the low-quality server is indifferent between the two signals, and he randomizes between them. So, this hybrid strategy $[(f(R|H) = 1, 0 < f(R|L) < 1), (n(\delta^R), p(\delta^C)), \delta^R, \delta^C = 0]$ can be sustained as an equilibrium if $\lambda^O(\delta^R) = \mu - \theta/V_L$ with the belief $\delta < \delta^R < 1$.

Considering the floor function in $n(\delta^R)$, denote $\mathbb{N}^O := \{i \in \mathbb{N}_+ | i = n(\delta^R), \delta < \delta^R < 1\}$ as the set of all possible largest queue lengths for customers to join, where \mathbb{N}_+ is the set of all nonnegative integers. It is clear that $\lambda^O(\cdot)$ is uniquely determined by $n(\delta^R)$, and thus each element in \mathbb{N}^O corresponds a unique increasing line of $\lambda^O(\cdot)(\lambda)$ in λ . Each above increasing line has a unique crossing point with the function $\lambda^U(0)$. The condition $\lambda^O(\delta^R) = \mu - \theta/V_L$

holds only at these crossing points. And the total number of these crossing points is the cardinality of the set $|\mathbb{N}^O|$. In other words, only when the potential arrival rate λ takes several discrete values may this hybrid equilibrium appear.

(2) $0 < f(R|H) < 1$ and $f(R|L) = 1$. In this case, we have $\delta^C = 1$ and $0 < \delta^R = \frac{\delta f(R|H)}{\delta f(R|H)+1-\delta} < \delta$. The hybrid strategy $[(0 < f(R|H) < 1, f(R|L) = 1), (n(\delta^R), p(\delta^C)), \delta^R, \delta^C = 1]$ can be sustained as an equilibrium if $\lambda^O(\delta^R) = \mu - \theta/V_H$ with the belief $0 < \delta^R < \delta$, which holds at several discrete values of the potential arrival rate λ due to the floor function in $\lambda^O(\delta^R)$.

(3) $f(R|H) = 0$ and $0 < f(R|L) < 1$. Now, we have $\delta^R = 0$ and $\delta < \delta^C = \frac{\delta}{\delta+(1-\delta)(1-f(C|L))} < 1$. The hybrid strategy $[(f(R|H) = 0, 0 < f(R|L) < 1), (n(\delta^R), p(\delta^C)), \delta^R = 0, \delta^C]$ can be sustained as an equilibrium if $\lambda^O(0) = \lambda^U(\delta^C)$ with the belief $\delta < \delta^C < 1$. Denote the unique crossing point of $\lambda^O(0)$ and $\lambda^U(\delta)$ as $\lambda_{0\delta}$ and that of $\lambda^O(0)$ and $\lambda^U(1)$ as λ_{01} . Then, it can be easily verified that the condition $\lambda^O(0) = \lambda^U(\delta^C)$ ($\delta < \delta^C < 1$) holds only for $\lambda \in (\lambda_{0\delta}, \lambda_{01})$.

(4) $0 < f(R|H) < 1$ and $f(R|L) = 0$. Similarly, we have $\delta^R = 1$ and $0 < \delta^C = \frac{\delta(1-f(R|H))}{\delta(1-f(R|H))+1-\delta} < \delta$. The hybrid strategy $[(0 < f(R|H) < 1, f(R|L) = 0), (n(\delta^R), p(\delta^C)), \delta^R = 1, \delta^C]$ can be sustained as an equilibrium if $\lambda^O(1) = \lambda^U(\delta^C)$ ($0 < \delta^C < \delta$). Denote the unique crossing point of $\lambda^O(1)$ and $\lambda^U(0)$ as λ_{10} and that of $\lambda^O(1)$ and $\lambda^U(\delta)$ as $\lambda_{1\delta}$. Then, only for $\lambda \in (\lambda_{10}, \lambda_{1\delta})$ can the condition $\lambda^O(1) = \lambda^U(\delta^C)$ ($0 < \delta^C < \delta$) hold with a corresponding value of $f(R|H) \in (0, 1)$.

(5) $0 < f(R|H) < 1$ and $0 < f(R|L) < 1$. Under this mixed strategy, the posterior beliefs satisfy that $0 < \delta^R = \frac{\delta f(R|H)}{\delta f(R|H)+(1-\delta)f(R|L)} < 1$ and $0 < \delta^C = \frac{\delta f(C|H)}{\delta f(C|H)+(1-\delta)f(C|L)} < 1$. Then, the mixed strategy $[(0 < f(R|H) < 1, 0 < f(R|L) < 1), (n(\delta^R), p(\delta^C)), \delta^R, \delta^C]$ ($0 < \delta^R, \delta^C < 1$) can be sustained if $\lambda^O(\delta^R) = \lambda^U(\delta^C)$.

B.2 Customers' Equilibrium Queueing Strategies in Unobservable Queues with $\delta^C = 0$ or 1

When $\delta^C = 0$, the uninformed customers believe that the quality level must be low. Then, given customers' queueing strategy (p_L, p_{un}, p_H) , both the expected utilities of a negatively informed customer and an uninformed one are $u_L(p_L, p_{un}) := V_L - \frac{\theta}{\mu - \lambda(qp_L + (1-q)p_{un})}$. Then, customers' equilibrium queueing strategies (p_L^U, p_{un}^U, p_H^U) evolve as follows.

First, when $\lambda \leq \mu - \theta/V_L$, it is still true that joining is a dominant strategy for all customers, and thus $(1, 1, 1)$ is the unique equilibrium profile. After λ becomes larger than

Range of λ	$(0, \mu - \frac{\theta}{V_L}]$	$(\mu - \frac{\theta}{V_L}, \min(\lambda_1, \lambda_2)]$	$(\min(\lambda_1, \lambda_2), \frac{\theta(V_H - V_L)}{qV_H V_L}]$	$(\frac{\theta(V_H - V_L)}{qV_H V_L}, +\infty)$
Case 1: $\lambda_1 < \lambda_2$	$(1, 1, 1)$	$(p_L^U, 1, 1)$	$(0, p_{un}^U, 1)$	$(0, p_{un}^U, p_H^U)$
Case 2: $\lambda_1 \geq \lambda_2$	$(1, 1, 1)$	$(p_L^U, 1, 1)$	$(0, p_{un}^U, p_H^U)$	

Table B.1: Equilibrium joining strategy (p_L^U, p_{un}^U, p_H^U) in unobservable queues with $\delta^C = 0$

$\mu - \theta/V_L$, it is still dominant for the positively informed customers to join. However, the negatively informed and uninformed customers now adopt mixed strategies p_L and p_{un} such that

$$u_L(p_L, p_{un}) = V_L - \frac{\theta}{\mu - (qp_L\lambda + (1-q)p_{un}\lambda)} = 0. \quad (\text{B.1})$$

Thus, the equilibrium is $(p_L^U, p_{un}^U, 1)$ with p_L^U and p_{un}^U solve $u_L(p_L^U, p_{un}^U) = 0$. When λ further increases to $\frac{\mu - \theta/V_H}{q + (1-q)p_{un}^U}$, we get that $u_H(p_{un}^U, 1) = 0$. Then, the equilibrium joining strategy is (p_L^U, p_{un}^U, p_H^U) , where p_L^U and p_{un}^U solve $u_L(p_L^U, p_{un}^U) = 0$ and p_H^U solves $u_H(p_{un}^U, p_H^U) = 0$ given p_{un}^U .

Then, based on $\lambda_L^U(0) = \lambda(qp_L^U + (1-q)p_{un}^U)$ and $\lambda_H^U(0) = \lambda(qp_H^U + (1-q)p_{un}^U)$, we can get the effective arrival rates of the low- and high-quality servers. The effective arrival rate of the low quality server equals λ when $\lambda \leq \mu - \theta/V_L$ and $\mu - \theta/V_L$ otherwise. However, different from the general case $0 < \delta^C < 1$ in Proposition 3.4 where the effective arrival rates for both types of the server are finally unique, as now p_L^U and p_{un}^U sometimes are only determined by $u_L(p_L^U, p_{un}^U) = 0$, the resulting effective arrival rate of the high quality server $\lambda_H^U(0)$ may not be unique depending on the composition of p_L^U and p_{un}^U . In our research, we specify the value of $\lambda_H^U(0)$ based on the continuity of $\lambda_H^U(\delta^C)$ in δ^C as stated in the following corollary.

Corollary B.1. *In unobservable queues, both the effective arrival rates $\lambda_H^U(\delta^C)$ and $\lambda_L^U(\delta^C)$ are nondecreasing as the belief of all uninformed customers $\delta^C \in (0, 1)$ increases.*

According to Corollary B.1, $\lambda_H^U(\delta^C)$ is nondecreasing in δ^C . To simplify and unify the sequential equilibrium analysis, we focus on the case where $\lambda_H^U(0) = \lim_{\delta^C \rightarrow 0^+} \lambda_H^U(\delta^C)$. This corresponds to the equilibrium queueing strategies in Table B.1 and the effective arrival rates in Table B.2.

Following the above analysis, we can analyze the case where $\delta^C = 1$. Now, the uninformed customers believe that the quality level must be high. Then, given customers' queueing strategy (p_L, p_{un}, p_H) , both the expected utilities of a positively informed customer and an uninformed one are $u_H(p_{un}, p_H) := V_H - \frac{\theta}{\mu - \lambda(qp_H + (1-q)p_{un})}$. Then, customers' equilibrium queueing strategies (p_L^U, p_{un}^U, p_H^U) evolve as follows.

Range of λ		$(0, \mu - \frac{\theta}{V_L}]$	$(\mu - \frac{\theta}{V_L}, \min(\lambda_1, \lambda_2)]$	$(\min(\lambda_1, \lambda_2), \frac{\theta(V_H - V_L)}{qV_H V_L}]$	$(\frac{\theta(V_H - V_L)}{qV_H V_L}, +\infty)$
Case 1: $\lambda_1 < \lambda_2$	$\lambda_H^U(0)$	λ		$\mu - \frac{\theta}{V_L} + q\lambda$	$\mu - \frac{\theta}{V_H}$
	$\lambda_L^U(0)$	λ	$\mu - \frac{\theta}{V_L}$		
Case 2: $\lambda_1 \geq \lambda_2$	$\lambda_H^U(0)$	λ		$\mu - \frac{\theta}{V_H}$	
	$\lambda_L^U(0)$	λ	$\mu - \frac{\theta}{V_L}$		

Table B.2: Effective arrival rates $\lambda_H^U(0)$ and $\lambda_L^U(0)$ corresponding to the equilibrium joining strategies in Table B.1

Range of λ	$(0, \mu - \frac{\theta}{V_L}]$	$(\mu - \frac{\theta}{V_L}, \min(\lambda_1, \lambda_2)]$	$(\min(\lambda_1, \lambda_2), \lambda_2]$	$(\lambda_2, \frac{\theta(V_H - V_L)}{qV_H V_L}]$	$(\frac{\theta(V_H - V_L)}{qV_H V_L}, +\infty)$
Case 1: $\lambda_1 < \lambda_2$	$(1, 1, 1)$	$(p_L^U, 1, 1)$	$(0, 1, 1)$	$(0, p_{un}^U, 1)$	(p_L^U, p_{un}^U, p_H^U)
Case 2: $\lambda_1 \geq \lambda_2$	$(1, 1, 1)$	$(p_L^U, 1, 1)$	(p_L^U, p_{un}^U, p_H^U)		

Table B.3: Equilibrium joining strategy (p_L^U, p_{un}^U, p_H^U) in unobservable queues with $\delta^C = 1$

First, when $\lambda \leq \mu - \theta/V_L$, we still get the unique equilibrium profile $(1, 1, 1)$. As λ becomes larger than $\mu - \theta/V_L$, it is still dominant for the uninformed and positively informed customers to join. However, the negatively informed now adopts a mixed strategy p_L such that

$$u_L(p_L, 1) = V_L - \frac{\theta}{\mu - (qp_L\lambda + (1 - q)\lambda)} = 0. \quad (\text{B.2})$$

Hence, the equilibrium is $(p_L^U, 1, 1)$, where $p_L^U = \max\{0, p_L\}$ with p_L solving (B.2). When λ further increases to $\mu - \theta/V_H$, we get that $u_H(1, 1) = 0$. Then, the equilibrium joining strategy is (p_L^U, p_{un}^U, p_H^U) , where p_{un}^U and p_H^U solve $u_H(p_{un}^U, p_H^U) = 0$ and $p_L^U = \max\{0, p_L\}$ with p_L solving $u_L(p_L, p_{un}^U) = 0$ given p_{un}^U .

Then, based on above analysis, we can get the effective arrival rates. The effective arrival rate of the high quality server equals λ when $\lambda \leq \mu - \theta/V_H$ and $\mu - \theta/V_H$ otherwise. Different compositions of p_{un}^U and p_H^U may still yield different values of $\lambda_L^U(1)$. Using the nondecreasing property of $\lambda_L^U(\delta^C)$ in δ^C (see Corollary B.1), we only consider $\lambda_L^U(1) = \lim_{\delta^C \rightarrow 1^-} \lambda_L^U(\delta^C)$. This corresponds to the equilibrium queueing strategies in Table B.3 and the effective arrival rates in Table B.4.

Range of λ		$(0, \mu - \frac{\theta}{V_L}]$	$(\mu - \frac{\theta}{V_L}, \min(\lambda_1, \lambda_2)]$	$(\min(\lambda_1, \lambda_2), \lambda_2]$	$(\lambda_2, \frac{\theta(V_H - V_L)}{qV_H V_L}]$	$(\frac{\theta(V_H - V_L)}{qV_H V_L}, +\infty)$
Case 1: $\lambda_1 < \lambda_2$	$\lambda_H^U(1)$	λ			$\mu - \frac{\theta}{V_H}$	
	$\lambda_L^U(1)$	λ	$\mu - \frac{\theta}{V_L}$	$(1-q)\lambda$	$\mu - \frac{\theta}{V_H} - q\lambda$	$\mu - \frac{\theta}{V_L}$
Case 2: $\lambda_1 \geq \lambda_2$	$\lambda_H^U(1)$	λ			$\mu - \frac{\theta}{V_H}$	
	$\lambda_L^U(1)$	λ	$\mu - \frac{\theta}{V_L}$			

Table B.4: Effective arrival rates $\lambda_H^U(1)$ and $\lambda_L^U(1)$ corresponding to the equilibrium joining strategies in Table B.3

B.3 Review of Hole-avoiding Decision Process in Debo et al. (2012)

Since some customers are informed and others are not, and the join-or-balk actions of the informed customers vary with the information they possess, the progression of the queue length is quality-dependent. Both high-quality and low-quality queues can be modeled as birth-and-death (BD) processes. Recall that $\lambda_{i,H}$ (resp. $\lambda_{i,L}$) is the effective arrival rate at queue length i in case the service quality is high (resp. low), where $i = 0, 1, \dots, n(1) + 1$. Then, we have $\lambda_{i,H} = \lambda[qp_H^O(i) + (1-q)p_{un}^O(i)]$ and $\lambda_{i,L} = \lambda[qp_L^O(i) + (1-q)p_{un}^O(i)]$. And recall that $\pi_{i,H}$ (resp. $\pi_{i,L}$) is the limiting probability that the number of customers in the system equals i when the server is of high (resp. low) quality, $i = 0, 1, \dots, n(1) + 1$. Clearly,

$$\pi_{i,H} = \pi_{0,H} \prod_{j=0}^{i-1} \lambda_{j,H} / \mu^i, \quad (\text{B.3})$$

where

$$\pi_{0,H} = \left(1 + \sum_{i=1}^{n(1)+1} \prod_{j=0}^{i-1} \lambda_{j,H} / \mu^i \right)^{-1}. \quad (\text{B.4})$$

And similar expressions can be derived for $\pi_{i,L}$, $i \geq 0$. Suppose that the queue length is i . Then, the posterior probability that the service is of high quality equals

$$\Pr(H|i) \equiv \frac{\delta^R \pi_{i,H}}{\delta^R \pi_{i,H} + (1 - \delta^R) \pi_{i,L}} = \frac{\delta^R}{\delta^R + (1 - \delta^R) \frac{\pi_{0,L}}{\pi_{0,H}} \prod_{j=0}^{i-1} \frac{\lambda_{j,L}}{\lambda_{j,H}}}. \quad (\text{B.5})$$

If the decision problem for a customer is whether to join or not, then, on seeing a queue length i , one should join iff

$$\Pr(H|i) \geq \frac{\theta(i+1)/\mu - V_L}{V_H - V_L}. \quad (\text{B.6})$$

Note that the right hand side (rhs) of (B.6) is a linear increasing function of i .

After finding the best response of both informed and uninformed customers, we then search for the equilibrium, i.e., the determination of n_{hole} . The value of n_{hole} is uniquely determined by the likelihood ratio ϕ_0 , where $\phi_0 \equiv \pi_{0,L}/\pi_{0,H}$, in the following way. Given ϕ_0 , one can calculate $\Pr(H|1)$ according to (B.5) and check whether the inequality (B.6) is violated or not to determine the decision of uninformed customers at state 1; if the decision is to join, one can further determine $\Pr(H|2)$ and check the decision of uninformed customers at state 2. Doing this recursively, one would find the first state violating inequality (B.6), namely n_{hole} . Therefore, instead of searching for an equilibrium n_{hole} , one can alternatively search for the equilibrium value of ϕ_0 . Note that ϕ_0 is a function of n_{hole} , denoted as $\phi_0(n_{hole})$. Following the standard analysis of BD processes for both high-quality and low-quality queues, we obtain that

$$\phi_0(n_{hole}) = \frac{1 + \sum_{i=1}^{n(1)+1} \prod_{j=0}^{i-1} \lambda_{j,h} / \mu^i}{1 + \sum_{i=1}^{n_{hole}} \prod_{j=0}^{i-1} \lambda_{j,l} / \mu^i} = \frac{\sum_{i=0}^{n_{hole}} \rho^i + q \sum_{i=n_{hole}+1}^{n(1)+1} \rho^i}{\sum_{i=0}^{n(0)+1} \rho^i + \sum_{n(0)+2}^{n_{hole}} (1-q)^{i-n(0)-1} \rho^i}. \quad (\text{B.7})$$

In order to find the equilibrium value of n_{hole} , the following two algorithms are designed.

Algorithm 1: Consider all the integers in the set $\{n_l + 1, \dots, n_h + 1\}$ as potential integers for the position of the hole. Suppose all other uninformed customers are adopting a hole-avoiding strategy with a hole positioned at m_{hole} . Let $n_{hole}(m_{hole})$ be the best response strategy for the tagged uninformed customer. Given m_{hole} , the corresponding $\phi_0(m_{hole})$ can be calculated according to (B.7). Then one can derive $n_{hole}(m_{hole})$, the smallest queue length violating the condition (B.6). The crossing point, if exists, of this best response function with the 45-degree line in the coordinate plane would be the equilibrium hole.

Algorithm 2: Directly search over ϕ_0 . First, for each given ϕ_0 , one can derive the corresponding $n_{hole}(\phi_0)$, the smallest queue length violating condition (B.6). This is a non-increasing step function. Second, for each given $n_{hole}(\phi_0)$, one can calculate the corresponding ϕ_0 according to (B.7) and obtain $\phi_0(n_{hole}(\phi_0))$. This, too, is a step function. The equilibrium value of ϕ_0 is then the crossing point, if exists, of this step function with the 45-degree line in the coordinate plane, from which the equilibrium n_{hole} can be calculated.

B.4 Sequential Equilibria Analysis of Hybrid and Mixed Strategies with Heterogeneous Customers

For the general scenario where both informed and uninformed customers exist, we investigate various hybrid and mixed strategies one by one as follows. Note that both signals R and C are on the equilibrium path here.

(1) $f(R|H) = 1$ and $0 < f(R|L) < 1$. Only the low-quality server sends the signal C with positive probability, and thus $\delta^C = 0$. When uninformed customers see the signal R , they update their belief as $\delta^R = \frac{\delta}{\delta + (1-\delta)f(R|L)}$ by the Bayes' rule. In this case, $\delta < \delta^R < 1$. We need the first condition that $\lambda_L^O(\delta^R) = \lambda_L^U(0)$, under which the low-quality server is indifferent between the two signals and thus randomizes between them. The second condition is that $\lambda_H^O(\delta^R) \geq \lambda_H^U(0)$, and then the high-quality server has no incentive to deviate to C . So, this hybrid-strategy sequential equilibrium

$$[(1, f(R|L)), \{(p_L^U, p_{un}^U, p_H^U), \{(p_L^O(i), p_{un}^O(i), p_H^O(i))\}_{i=0}^{+\infty}\}, \delta^R, 0]$$

can be sustained if $\lambda_H^O(\delta^R) \geq \lambda_H^U(0)$ and $\lambda_L^O(\delta^R) = \lambda_L^U(0)$, where δ^R satisfies $\delta < \delta^R < 1$.

(2) $0 < f(R|H) < 1$ and $f(R|L) = 1$. In this case, we have $\delta^C = 1$ and $0 < \delta^R = \frac{\delta f(R|H)}{\delta f(R|H) + 1 - \delta} < \delta$. Similarly to (1), we can get that the hybrid-strategy sequential equilibrium

$$[(f(R|H), 1), \{(p_L^U, p_{un}^U, p_H^U), \{(p_L^O(i), p_{un}^O(i), p_H^O(i))\}_{i=0}^{+\infty}\}, \delta^R, 1]$$

can be sustained if $\lambda_H^O(\delta^R) = \lambda_H^U(1)$ and $\lambda_L^O(\delta^R) \geq \lambda_L^U(1)$ with δ^R satisfying $0 < \delta^R < \delta$.

(3) $f(R|H) = 0$ and $0 < f(R|L) < 1$. Now, we have $\delta^R = 0$ and $\delta < \delta^C = \frac{\delta}{\delta + (1-\delta)(1-f(C|L))} < 1$. We can get that the hybrid-strategy sequential equilibrium

$$[(0, f(R|L)), \{(p_L^U, p_{un}^U, p_H^U), \{(p_L^O(i), p_{un}^O(i), p_H^O(i))\}_{i=0}^{+\infty}\}, 0, \delta^C]$$

can be sustained if $\lambda_H^U(\delta^C) \geq \lambda_H^O(0)$ and $\lambda_L^O(0) = \lambda_L^U(\delta^C)$ with δ^C satisfying $\delta < \delta^C < 1$.

(4) $0 < f(R|H) < 1$ and $f(R|L) = 0$. We have $\delta^R = 1$ and $0 < \delta^C = \frac{\delta(1-f(R|H))}{\delta(1-f(R|H)) + 1 - \delta} < \delta$. And the hybrid-strategy sequential equilibrium

$$[(f(R|H), 0), \{(p_L^U, p_{un}^U, p_H^U), \{(p_L^O(i), p_{un}^O(i), p_H^O(i))\}_{i=0}^{+\infty}\}, 1, \delta^C]$$

can be sustained if $\lambda_H^O(1) = \lambda_H^U(\delta^C)$ and $\lambda_L^U(\delta^C) \geq \lambda_L^O(1)$ with δ^C satisfying $0 < \delta^C < \delta$.

(5) $0 < f(R|H) < 1$ and $0 < f(R|L) < 1$. Under this mixed strategy, the posterior beliefs satisfy that $0 < \delta^R = \frac{\delta f(R|H)}{\delta f(R|H) + (1-\delta)f(R|L)} < 1$ and $0 < \delta^C = \frac{\delta f(C|H)}{\delta f(C|H) + (1-\delta)f(C|L)} < 1$. Then, the mixed-strategy sequential equilibrium

$$[(f(R|H), f(R|L)), \{(p_L^U, p_{un}^U, p_H^U), \{(p_L^O(i), p_{un}^O(i), p_H^O(i))\}_{i=0}^{+\infty}\}, \delta^R, \delta^C]$$

can be sustained if $\lambda_H^O(\delta^R) = \lambda_H^U(\delta^C)$ and $\lambda_L^O(\delta^R) = \lambda_L^U(\delta^C)$ with δ^R and δ^C satisfying $0 < \delta^R, \delta^C < 1$.

B.5 Proofs for Chapter 3 and Appendix B.2

Proof of Proposition 3.1. First consider pooling on C . For $\lambda \in (\hat{\lambda}, +\infty)$, the set \mathbb{T}' in Definition 3.2 is equal to $\mathbb{T} = \{H, L\}$, and the credible updating rule must set $\delta^R = \delta$. The behind reason is that if one type of the server benefits by deviating, then so does the other type because they share the same payoff functions. This means that the two types of the server always have the same incentive to deviate, leading to the off-equilibrium-path beliefs being equal to the prior based on the credible updating rule. Then, under the belief $\delta^R = \delta$, both types of the server deviate to R . So, pooling on C cannot be sustained as a perfect sequential equilibrium on $\lambda \in (\hat{\lambda}, +\infty)$. By contrast, for $\lambda \in (0, \hat{\lambda})$, it can be verified that the set $\mathbb{T}' \cup \mathbb{T}''$ in Definition 3.2 is empty, and thus the credible updating rule puts no restriction on customers' posterior belief after observing the off-equilibrium-path signal R . So, the pooling strategy $[(C, C), (n(\delta^R), p(\delta^C)), \delta^R, \delta^C = \delta]$ with the off-equilibrium-path belief δ^R satisfying $\lambda^O(\delta^R) \leq \lambda^U(\delta)$ can be sustained as a perfect sequential equilibrium for $\lambda \in (0, \hat{\lambda})$.

The same argument applies to pooling on R . It can be easily verified that the credible updating rule requires that $\delta^C = \delta$ and thus filters out pooling on R as a perfect sequential equilibrium for $\lambda \in (0, \hat{\lambda})$. While for $\lambda \in (\hat{\lambda}, +\infty)$, the credible updating rule exerts no restriction on δ^C , and the pooling strategy $[(R, R), (n(\delta^R), p(\delta^C)), \delta^R = \delta, \delta^C]$ with the off-equilibrium-path belief δ^C satisfying $\lambda^U(\delta^C) \leq \lambda^O(\delta)$ can be sustained as a perfect sequential equilibrium.

Proof of Proposition 3.2. Considering the relationship between $f(R|H)$ and $f(R|L)$ in the hybrid or mixed equilibria, we prove the result based on the following three cases.

1. $f(R|H) > f(R|L)$. We have that $0 \leq \delta^C < \delta < \delta^R \leq 1$ in the equilibria, and thus

the corresponding effective arrival rates to two types of the server must be smaller than those under the pooling-on- C equilibrium.

2. $f(R|H) < f(R|L)$. In this case, we have $0 \leq \delta^R < \delta < \delta^C \leq 1$ in the equilibria, and thus the corresponding effective arrival rates to two types of the server must be no larger than those under the pooling-on- R equilibrium.
3. $f(R|H) = f(R|L)$. In this case, we have $0 < f(R|H) = f(R|L) < 1$ and $\delta^R = \delta^C = \delta$ in the equilibria. Such a mixed-strategy sequential equilibrium can only be sustained at $\hat{\lambda}$ with the resulting effective arrival rates equal to $\lambda^O(\delta)(= \lambda^U(\delta))$, which are just the ones under the pooling perfect sequential equilibria.

Proof of Proposition 3.3. When $\lambda \leq \mu - \theta/V_L$, we have $V_L - \frac{\theta}{\mu - \lambda} \geq 0$. This means a negatively informed customer should join even when all other customers (positively informed as well as uninformed) join. This makes joining a dominant strategy for all and therefore $(1, 1, 1)$ is the unique equilibrium profile.

When λ increases and becomes larger than $\mu - \theta/V_L$, it is still strictly dominant for the positively informed and uninformed customers to join. However, the negatively informed customers now adopt a mixed strategy p_L such that

$$u_L(p_L, 1) = V_L - \frac{\theta}{\mu - (qp_L\lambda + (1 - q)\lambda)} = 0. \quad (\text{B.8})$$

Thus, the equilibrium is $(p_L^U, 1, 1)$, where p_L^U is the solution of (B.8), i.e.,

$$p_L^U = \frac{\mu - \theta/V_L}{\lambda q} - \frac{1 - q}{q}. \quad (\text{B.9})$$

When λ further increases, exactly one of the following two cases occurs: (1) Case 1: p_L^U reaches the value of zero while all others still strictly prefer to join (i.e., $u_H(1, 1) > 0$ and $u_{un}(0, 1, 1) > 0$); or (2) Case 2: before or when p_L^U decreases to 0, both $u_H(1, 1)$ and $u_{un}(p_L^U, 1, 1)$ become 0 (recall that we already have $u_L(p_L^U, 1) = 0$).

We now consider Case 1 and its condition. According to (B.9), we see that p_L^U reaches the value of zero when $\lambda = \lambda_1 := \frac{\mu - \theta/V_L}{1 - q}$. Note that $u_L(0, 1) = 0$ for $\lambda = \lambda_1$. If the condition $u_H(1, 1) > 0$ holds for this value of λ , then $u_{un}(0, 1, 1) > 0$, and thus $(0, 1, 1)$ is the unique equilibrium. Plugging $\lambda = \lambda_1$ into this condition, we have $\lambda_1 < \lambda_2 := \mu - \frac{\theta}{V_H}$, which coincides with the condition stated for Case 1 in the proposition.

Further increasing λ makes $u_L(0, 1) < 0$, $u_{un}(0, 1, 1) > 0$ and $u_H(1, 1) > 0$. And thus the equilibrium keeps as $(0, 1, 1)$ until λ reaches a level where uninformed customers begin to randomize between joining and balking. This threshold value of λ solves the equation $u_{un}(0, 1, 1) = 0$, which means that uninformed customers get a utility zero. Denote its unique solution satisfying $0 < \lambda < \mu$ by $\bar{\lambda}$.

Increasing λ further leads to the equilibrium $(0, p_{un}^U, 1)$. And as λ increases, p_{un}^U decreases. The equilibrium $(0, p_{un}^U, 1)$ stays for a while until λ reaches the point where the utility of a negatively informed customer becomes zero (i.e., $u_L(0, p_{un}^U) = 0$). In this case, recall that we already have $u_{un}(0, p_{un}^U, 1) = 0$, implying $u_H(p_{un}^U, 1) = 0$. Then, solving any two of the three equations $u_H(p_{un}^U, 1) = 0$, $u_{un}(0, p_{un}^U, 1) = 0$ and $u_L(0, p_{un}^U) = 0$ yields the critical value of λ as $\frac{\theta(V_H - V_L)}{qV_HV_L}$.

When λ further increases, the expected utilities of all customers become zero in equilibrium, which leads to the three equations $u_H(p_{un}^U, p_H^U) = 0$, $u_{un}(p_L^U, p_{un}^U, p_H^U) = 0$ and $u_L(p_L^U, p_{un}^U) = 0$. Among these three equations, any one of them is redundant given the other two. So, the equilibrium queueing strategy is identified by a system of two nonlinear equations with three variables. As a result, multiple equilibria typically exist. Specifically, any choice of p_{un}^U leads to the corresponding $p_H^U = \frac{\mu - \theta/V_H}{\lambda q} - \frac{(1-q)p_{un}^U}{q}$ and $p_L^U = \frac{\mu - \theta/V_L}{\lambda q} - \frac{(1-q)p_{un}^U}{q}$, which are uniquely identified by $u_H(p_{un}^U, p_H^U) = 0$ and $u_L(p_L^U, p_{un}^U) = 0$, respectively. It is clear that $p_L^U < p_H^U$ but what remains to be needed is that $0 \leq p_L^U, p_{un}^U, p_H^U \leq 1$. Since p_L^U and p_H^U are both decreasing with p_{un}^U , the smallest possible value for p_{un}^U should be no less than the one where the corresponding p_H^U is 1. Solving $\frac{\mu - \theta/V_H}{\lambda q} - \frac{(1-q)p_{un}^U}{q} = 1$ leads to $p_{un}^U = \frac{\mu - \theta/V_H}{\lambda(1-q)} - \frac{q}{1-q}$. And thus, $p_{un}^U \geq \max\left\{0, \frac{\mu - \theta/V_H}{\lambda(1-q)} - \frac{q}{1-q}\right\}$. As p_{un}^U starts to increase from this value, to compensate for that, both p_H^U and p_L^U have to be reduced. Then, the largest possible value for p_{un}^U should be no larger than the one where the corresponding p_L^U is 0. Solving $\frac{\mu - \theta/V_L}{\lambda q} - \frac{(1-q)p_{un}^U}{q} = 0$ leads to $p_{un}^U = \frac{\mu - \theta/V_L}{\lambda(1-q)}$. Since now $\lambda > \frac{\mu - \theta/V_L}{(1-q)}$, we get that $\frac{\mu - \theta/V_L}{\lambda(1-q)} < 1$. Therefore, the feasible range for p_{un}^U is $\left[\max\left\{0, \frac{\mu - \theta/V_H}{\lambda(1-q)} - \frac{q}{1-q}\right\}, \frac{\mu - \theta/V_L}{\lambda(1-q)}\right]$.

Next, consider Case 2 where $\lambda_1 \geq \lambda_2$. When $\lambda > \lambda_2$, it can be verified that the expected utilities of all customers become zero in equilibrium. Similarly to the arguments in Case 1, we can show that the next pattern is (p_L^U, p_{un}^U, p_H^U) with $\max\left\{0, \frac{\mu - \theta/V_H}{\lambda(1-q)} - \frac{q}{1-q}\right\} \leq p_{un}^U \leq \min\left\{1, \frac{\mu - \theta/V_L}{\lambda(1-q)}\right\}$.

Proof of Proposition 3.4. Below we provide the detailed proof for the high-quality server under Case 1. As the proofs of other parts follow a similar logic, we omit the details.

First, note that $\lambda_H^U(\delta^C) = \lambda(qp_H^U + (1-q)p_{un}^U)$. According to Proposition 3.3, $qp_H^U + (1 -$

$q)p_{un}^U = 1$ for $\lambda \leq \bar{\lambda}$. Second, for $\bar{\lambda} < \lambda \leq \frac{\theta(V_H - V_L)}{qV_H V_L}$, p_{un}^U is between 0 and 1, $p_H^U = 1$ and $p_L^U = 0$, which means that the arrivals into a high-quality server consist of all informed customers (a rate of $q\lambda$) and some of uninformed customers, while the arrivals into the low-quality server consist of only that part of uninformed customers. Hence, denote the effective arrival rate of the high-quality server as x and thus, the effective arrival rate of the low-quality server is $x - q\lambda$. Since uninformed customers adopt a mixed strategy, their expected utility is 0, and thus x is the unique solution satisfying $q\lambda < x < \mu$ in the following equation

$$\delta^C V_H + (1 - \delta^C) V_L = \delta^C \frac{\theta}{\mu - x} + (1 - \delta^C) \frac{\theta}{\mu - (x - q\lambda)}. \quad (\text{B.10})$$

Note that the right hand side (rhs) of (B.10) is increasing in x and decreasing in λ considering $q\lambda < x < \mu$, resulting in the increasing property of $x(\lambda)$ in λ . Similarly, to show that $y := x(\lambda) - q\lambda$ is decreasing in λ , one only need to rewrite (B.10) as follow:

$$\delta^C V_H + (1 - \delta^C) V_L = \delta^C \frac{\theta}{\mu - (y + q\lambda)} + (1 - \delta^C) \frac{\theta}{\mu - y}. \quad (\text{B.11})$$

Finally, for $\lambda \geq \frac{\theta(V_H - V_L)}{qV_H V_L}$, the expected utility is 0 in a high-quality-server queue, and thus the effective arrival rate of the high-quality server solves $V_H - \frac{\theta}{\mu - \lambda_H^U(\delta^C)} = 0$ or $\lambda_H^U(\delta^C) = \mu - \frac{\theta}{V_H}$.

Lemma B.1. *When the queue is observable, the effective arrival rate of the type- t server ($t = H, L$) $\lambda_t^O(\delta^R)$ satisfies $\lambda_t^O(0) \leq \lambda_t^O(\delta^R) \leq \lambda_t^O(1)$, where both $\lambda_t^O(0)$ and $\lambda_t^O(1)$ are strictly increasing in λ .*

Proof of Lemma B.1. If $\delta^R = 0$ (resp. $\delta^R = 1$), i.e., when all uninformed customers believe that the service quality must be low (resp. high), they will behave as negatively (resp. positively) informed customers and join the queue if the queue length does not exceed $n(0)$ (resp. $n(1)$). For $0 < \delta^R < 1$, according to Debo et al. (2012), uninformed customers join with a probability at queue length i ($i \in \{n(0) + 1, \dots, n(1)\}$). Both the queues of two types of the server can be modeled as birth-and-death processes. And the idle probability of the queueing system with $0 < \delta^R < 1$ is no less than the one with $\delta^R = 1$ and no larger than the one with $\delta^R = 0$. Therefore, we can use the time reversibility of an ergodic BD process in steady state to obtain that $\lambda_t^O(0) \leq \lambda_t^O(\delta^R) \leq \lambda_t^O(1)$ ($t = H, L$). Additionally, from the analysis in section 3.5.1, we can directly see that $\lambda_t^O(0)$ and $\lambda_t^O(1)$ ($t = H, L$) are all strictly increasing functions in λ .

Proof of Proposition 3.5. Using the result in Lemma B.1, we have that $\lambda_t^O(\delta) \geq \lambda_t^O(0)$ ($t = H, L$). Furthermore, denote the unique crossing point of $\lambda_H^U(1)$ and $\lambda_H^O(0)$ as $\hat{\lambda}_{R1}$, and the maximal crossing point of $\lambda_L^U(1)$ and $\lambda_L^O(0)$ as $\hat{\lambda}_{R2}$. Let $\hat{\lambda}_R = \max\{\hat{\lambda}_{R1}, \hat{\lambda}_{R2}\}$. Then, for $\lambda > \hat{\lambda}_R$, no matter what value δ^C takes, we must have that $\lambda_t^O(\delta) \geq \lambda_t^O(0) > \lambda_t^U(1) \geq \lambda_t^U(\delta^C)$ ($t = H, L$), and thus R is the unique strictly dominant strategy for two types of the server. So, we get a unique sequential equilibrium of pooling on R with the off-equilibrium-path belief $\delta^C \in [0, 1]$.

Similarly, for pooling on C , denote the minimal crossing point of $\lambda_H^U(0)$ and $\lambda_H^O(1)$ as $\hat{\lambda}_{C1}$, and the unique crossing point of $\lambda_L^U(0)$ and $\lambda_L^O(1)$ as $\hat{\lambda}_{C2}$. Let $\hat{\lambda}_C = \min\{\hat{\lambda}_{C1}, \hat{\lambda}_{C2}\}$. It can be easily verified that $\hat{\lambda}_C < \hat{\lambda}_R$. From Lemma B.1, we have that $\lambda_t^O(\delta) \leq \lambda_t^O(1)$ ($t = H, L$). So, for $0 < \lambda < \hat{\lambda}_C$, no matter what value δ^R takes, we must have that $\lambda_t^O(\delta^R) \leq \lambda_t^O(1) < \lambda_t^U(0) \leq \lambda_t^U(\delta)$ ($t = H, L$), and thus C is the unique strictly dominant strategy for two types of the server. So, we get a unique sequential equilibrium of pooling on C with the off-equilibrium-path belief $\delta^R \in [0, 1]$.

Proof of Proposition 3.6. To get rid of the effect of some threshold values of λ , we consider the open intervals of λ where a separating sequential equilibrium exists. In any open interval where (R, C) can be sustained as a sequential equilibrium, we must have that $\lambda_H^O(1) > \lambda_H^U(0)$ and $\lambda_L^U(0) > \lambda_L^O(1)$. So, another separating sequential equilibrium (C, R) cannot be sustained on this interval because $\lambda_L^U(1) \geq \lambda_L^U(0) > \lambda_L^O(1) \geq \lambda_L^O(0)$, which means that the low-quality server will deviate from R to C . Similarly, it is easy to prove that when the separating sequential equilibrium (C, R) exists on some open intervals, then (R, C) cannot be sustained as a sequential equilibrium at the same time.

Proof of Corollary 3.2. Let us first consider the pure-strategy sequential equilibria. Given the belief of uninformed customers, the effective arrival rate to an unobservable (resp. observable) queue is larger when the service quality is high than that when it is low. So, the result holds under any pooling equilibrium. For the separating equilibria, we prove the result by contradiction. Assume that the result does not hold under some separating equilibrium; i.e., the effective arrival rate to the low-quality server is larger than that to the high-quality server. If the high-quality server deviates and mimics the low-quality server, then he can obtain an effective arrival rate that is no less than that of the low-quality server, which is larger than the one he obtains by staying on the equilibrium path. This implies that such an equilibrium cannot be sustained, leading to a contradiction.

Above argument applies to the hybrid or mixed sequential equilibria. Note that in the

hybrid or mixed sequential equilibria, when one type of the server randomizes between queue revelation and concealment, the effective arrival rates to his revealed and concealed queues must be the same.

Proof of Proposition 3.7.

(i) For $0 < \lambda < \hat{\lambda}_C$, the unique sequential equilibrium is pooling on C . According to the definition of $\hat{\lambda}_C$ in Proposition 3.5, we have $\lambda_t^U(\delta) > \lambda_t^O(\delta)$ ($t = H, L$). So, both the high- and low-quality servers choose to conceal the queue length in the non-signaling case with the corresponding effective arrival rates $\lambda_H^U(\delta)$ and $\lambda_L^U(\delta)$, which are equal to the ones under the pooling sequential equilibrium. Similarly, we can easily prove the result for $\lambda > \hat{\lambda}_R$.

(ii) For $\hat{\lambda}_C \leq \lambda \leq \hat{\lambda}_R$, let us first suppose that the separating strategy (R, C) can be sustained as a sequential equilibrium on some intervals of λ . We must have that $\lambda_H^O(1) \geq \lambda_H^U(0)$ and $\lambda_L^U(0) \geq \lambda_L^O(1)$. Then, we get that $\lambda_L^U(\delta) \geq \lambda_L^U(0) \geq \lambda_L^O(1) \geq \lambda_L^O(\delta)$. So, in the non-signaling case, concealing the queue length yields the maximal effective arrival rate of the low-quality server as $\lambda_L^U(\delta) (\geq \max\{\lambda_L^U(0), \lambda_L^U(\delta)\})$. This means that the low-quality server under the signaling case becomes worse off. Next, we show that the high-quality server becomes better off by considering two cases. Note that we always have $\lambda_H^O(1) \geq \lambda_H^O(\delta)$. In the first case where $\lambda_H^O(1) \geq \lambda_H^U(\delta)$, the result is obvious. In the second case where $\lambda_H^O(1) < \lambda_H^U(\delta)$, we can easily obtain that pooling on C can also be sustained as a perfect sequential equilibrium: when $\lambda_L^U(\delta) > \lambda_L^O(0)$, the off-equilibrium-path belief satisfies $\delta^R \in [0, 1]$; and when $\lambda_L^U(\delta) = \lambda_L^O(0)$, the off-equilibrium-path belief is $\delta^R = 0$ according to the credible updating rule. Then, the result can be proved.

Following the above procedure, the result for the separating strategy (C, R) can be proved. Here, we omit the details.

Proof of Proposition 3.8. First, assume that the separating strategy (R, C) can be sustained as a sequential equilibrium. According to the proof of Proposition 3.7, the low-quality server obtains the maximal effective arrival rate under either (R, C) or (C, C) in the signaling case, and must conceal the queue length in the non-signaling case. Then, the total utilities of all customers from the low-quality server in the signaling and non-signaling cases are $u_L^U(0)$ (or $u_L^U(\delta)$) and $u_L^U(\delta)$, respectively. Based on Proposition 3.4 and Appendix B.2, the difference between $u_L^U(0)$ and $u_L^U(\delta)$ exists for $\lambda \in \left(\lambda_1, \frac{\theta(V_H - V_L)}{qV_HV_L}\right)$ when $\lambda_1 < \lambda_2$, where $\lambda_L^U(\delta) > \mu - \frac{\theta}{V_L}$, leading to $u_L^U(\delta) < 0 = u_L^U(0)$. And under other cases, we have $u_L^U(\delta) = u_L^U(0)$. Therefore, we can conclude that $u_L^U(0) \geq u_L^U(\delta)$, implying that the customers' total utility becomes weakly larger in the signaling case.

Similarly, let us investigate the separating sequential equilibrium (C, R) . The low-quality server obtains the maximal effective arrival rate under either (C, R) or (R, R) in the signaling case, and must reveal the queue in the non-signaling case. Then, the total utilities of all customers from the low-quality server in the signaling and non-signaling cases are $u_L^O(0)$ (or $u_L^O(\delta)$) and $u_L^O(\delta)$, respectively. Note that in the utility function (3.4) with the low-quality server, the effective arrival rate at queue length i , $\lambda_{i,L}$, and the limiting probability of queue length i , $\pi_{i,L}$, are dependent on δ^R , which can be further specified as $\lambda_{i,L}(\delta^R)$ and $\pi_{i,L}(\delta^R)$. According to the definition of $n(0)$, it can be easily verified that $\lambda_{i,L}(0) = \lambda_{i,L}(\delta) = \lambda$ for $i = 0, 1, \dots, n(0)$, $\pi_{i,L}(\delta) \leq \pi_{i,L}(0)$ for $i = 0, 1, \dots, n(0)$, and $V_L - \frac{(i+1)\theta}{\mu} < 0$ for $i = n(0) + 1, \dots, n(1)$. So, we can get that

$$u_L^O(\delta) \leq \sum_{i=0}^{n(0)} \lambda_{i,L}(\delta) \pi_{i,L}(\delta) \left(V_L - \frac{(i+1)\theta}{\mu} \right) \leq \sum_{i=0}^{n(0)} \lambda_{i,L}(0) \pi_{i,L}(0) \left(V_L - \frac{(i+1)\theta}{\mu} \right) = u_L^O(0),$$

which indicates that the customers' total utility becomes weakly larger in the signaling case.

Proof of the statement “the Nash equilibria where the server adopts the pooling strategy (C, C) or (R, R) can always be sustained for all $\lambda \in (0, +\infty)$ ”: When both the high-quality and low-quality servers choose to conceal the queue, then the customers' equilibrium queueing strategy profile in a concealed queue is indeed equivalent to the one specified in section 3.5.1 with the belief of the uninformed customers equal to the prior δ . Since no path with positive probability arrives at the terminal nodes following each possible queue length in a revealed queue, no matter what the customers' queueing strategy in a revealed queue is, the expected utilities of all customers are equal to the ones under a concealed queue. In this case, we can assume that no (informed and uninformed) customer joins a revealed queue, which leads to zero effective arrival rates in revealed queues of both the high-quality and low-quality servers. And such an assumed strategy makes no player in the game deviate, which sustains the strategies of all players mentioned above as a Nash equilibrium of the overall signaling game.

Similarly, we can show that when both types of the server reveal the queue, there exist some customers' equilibrium queueing strategies that make the corresponding Nash equilibrium sustained. Here, we omit the details.

Proof of Corollary B.1. We have two cases on $\lambda_H^U(\delta^C)$ and $\lambda_L^U(\delta^C)$ as stated in Proposition 3.4. Case 2 is trivial as both effective arrival rates are independent of δ^C . And the result

under Case 1 holds as a result of the following three facts.

Fact 1. $\frac{d\bar{\lambda}}{d\delta^C} > 0$. Recall that in Proposition 3.3, $\bar{\lambda}$ is the unique solution for λ satisfying $0 < \lambda < \mu$ in the equation:

$$\delta^C V_H + (1 - \delta^C) V_L = \delta^C \frac{\theta}{\mu - \lambda} + (1 - \delta^C) \frac{\theta}{\mu - \lambda(1 - q)}.$$

Let $F(\delta^C, \bar{\lambda}) := \delta^C [V_H - \theta/(\mu - \bar{\lambda})] + (1 - \delta^C) \{V_L - \theta/[\mu - \bar{\lambda}(1 - q)]\}$. Then, we have that

$$\frac{d\bar{\lambda}}{d\delta^C} = -\frac{\partial F/\partial \delta^C}{\partial F/\partial \bar{\lambda}} = \frac{[V_H - \theta/(\mu - \bar{\lambda})] - \{V_L - \theta/[\mu - \bar{\lambda}(1 - q)]\}}{\delta^C \theta/(\mu - \bar{\lambda})^2 + (1 - \delta^C)(1 - q)\theta/[\mu - \bar{\lambda}(1 - q)]^2}.$$

From the proof of Proposition 3.3, we can see that $V_H > \theta/(\mu - \bar{\lambda})$ and $V_L < \theta/[\mu - \bar{\lambda}(1 - q)]$.

Therefore, we can get that $\frac{d\bar{\lambda}}{d\delta^C} \geq 0$.

Fact 2. $\frac{dx(\lambda)}{d\lambda} > 0$ for $\lambda \in \left(\bar{\lambda}, \frac{\theta(V_H - V_L)}{qV_H V_L}\right)$. The behind logic is similar to the one of Fact 1, and thus we omit the details here.

Fact 3. $x(\lambda) < \lambda$ for $\lambda \in \left(\bar{\lambda}, \frac{\theta(V_H - V_L)}{qV_H V_L}\right)$. This inequality holds because $x(\lambda) = \lambda[q p_h^U + (1 - q)p_{un}^U]$ where $p_h^U = 1$ and $0 < p_{un}^U < 1$ for $\lambda \in \left(\bar{\lambda}, \frac{\theta(V_H - V_L)}{qV_H V_L}\right)$.

Appendix C

Proofs for Chapter 4

Proof of Proposition 4.1. If a queue disclosure strategy (f_h, f_l) can induce the Bayes-plausible distribution of posteriors as presented in Proposition 4.1, it should satisfy that $p^R = \delta_0 f_h + (1 - \delta_0) f_l$ and $p_{H|R} = P_{H|R}(f_h, f_l) = \frac{\delta_0 f_h}{\delta_0 f_h + (1 - \delta_0) f_l}$. From these two equations, we obtain that $f_h = p^R p_{H|R} / \delta_0$ and $f_l = p^R (1 - p_{H|R}) / (1 - \delta_0)$.

Proof of Lemma 4.1. First, $\lambda_e^R(\delta) = \lambda \left(1 - \frac{\rho^{n(\delta)}}{\sum_{i=0}^{n(\delta)} \rho^i} \right) = \mu - \frac{\mu}{\sum_{k=0}^{n(\delta)} \rho^k}$, where $n(\delta) = \lfloor [\delta V_h + (1 - \delta) V_l] \mu / \theta \rfloor$. So, $\lambda_e^R(\delta)$ increases in $n(\delta)$. Second, as δ increases from 0 to 1, $n(\delta)$ repeats the following pattern: it first remains unchanged for a while, then increases by 1, and then remains unchanged, etc. The change of $\lambda_e^R(\delta)$ in δ is a consequence of this change pattern of $n(\delta)$ in δ .

Note that when $\lambda \geq \mu - \theta / V_h$, $\lambda_e^C(\delta) = \mu - \frac{\theta}{\delta V_h + (1 - \delta) V_l}$ for all $0 \leq \delta \leq 1$, which is concave and increasing in δ . When $\lambda < \mu - \theta / V_h$, $\lambda_e^C(\delta)$ consists of two pieces, first an increasing and concave function $\mu - \frac{\theta}{\delta V_h + (1 - \delta) V_l}$ on the domain $\delta \in \left[0, \frac{\theta - (\mu - \lambda) V_l}{(\mu - \lambda)(V_h - V_l)} \right]$ and then a constant λ on the domain $\delta \in \left[\frac{\theta - (\mu - \lambda) V_l}{(\mu - \lambda)(V_h - V_l)}, 1 \right]$. The overall function is still concave.

Proof of Proposition 4.2. As shown in Section 4.3, any queue disclosure strategy (f_h, f_l) yields a Bayes-plausible distribution of posteriors (i.e., $\delta_0 = [\delta_0 f_h + (1 - \delta_0) f_l] P_{H|R}(f_h, f_l) + [\delta_0 (1 - f_h) + (1 - \delta_0)(1 - f_l)] P_{H|C}(f_h, f_l)$), and an objective value $\lambda_e(f_h, f_l) = [\delta_0 f_h + (1 - \delta_0) f_l] \lambda_e^R(f_h, f_l) + [\delta_0 (1 - f_h) + (1 - \delta_0)(1 - f_l)] \lambda_e^C(f_h, f_l)$. It is straightforward to show that the point $(\delta_0, \lambda_e(f_h, f_l))$ can be regarded as the convex combination of two points $(P_{H|R}(f_h, f_l), \lambda_e^R(P_{H|R}(f_h, f_l)))$ and $(P_{H|C}(f_h, f_l), \lambda_e^C(P_{H|C}(f_h, f_l)))$, and thus $(\delta_0, \lambda_e(f_h, f_l)) \in \text{co}(\lambda_e^R, \lambda_e^C)$. On the other hand, given $(\delta_0, \Lambda) \in \text{co}(\lambda_e^R, \lambda_e^C)$, there exist δ_1, δ_2 and $\hat{\alpha}$ such that $\hat{\alpha} \delta_1 + (1 - \hat{\alpha}) \delta_2 = \delta_0$ and $\hat{\alpha} \lambda_e^R(\delta_1) + (1 - \hat{\alpha}) \lambda_e^C(\delta_2) = \Lambda$ ($0 \leq \delta_1, \delta_2, \hat{\alpha} \leq 1$). This indicates that

when the prior probability for the service quality to be high is δ_0 , we can always identify a queue disclosure strategy that yields the corresponding payoff value Λ for the server according to Proposition 4.1.

Proof of Proposition 4.3. According to Proposition 4.2, under the given prior δ_0 , all effective arrival rates that can be induced by feasible queue disclosure strategies constitute the set $\{\Lambda | (\delta_0, \Lambda) \in \text{co}(\lambda_e^R(\cdot), \lambda_e^C(\cdot))\}$. The conclusion then follows by the definition of $\Lambda_e(\delta_0)$.

Proof of Lemma 4.2. Since $\lambda_e^C(\delta) \geq \lambda_e^R(\delta)$ for all $\delta \in [0, 1]$ and $\lambda_e^C(\delta)$ is concave in δ , all convex combinations between a point on $\lambda_e^R(\cdot)$ and a point on $\lambda_e^C(\cdot)$ fall on or below the function curve $\lambda_e^C(\cdot)$. Therefore, $\Lambda_e(\delta) = \lambda_e^C(\delta)$ and $(f_h^e, f_l^e) = (0, 0)$. Similarly, when $\lambda_e^R(\delta)$ is a horizontal line and $\lambda_e^R(\delta) \geq \lambda_e^C(\delta)$ for all $\delta \in [0, 1]$, these convex combinations fall on or below the flat line $\lambda_e^R(\cdot)$. Therefore, $\Lambda_e(\delta) = \lambda_e^R(\delta)$, and $(f_h^e, f_l^e) = (1, 1)$.

Proof of Proposition 4.4. Part (i) clearly holds because, in this case, the effective arrival rate equals the potential arrival rate under the ‘always concealing’ strategy.

For part (ii), based on the relationship between f_h and f_l , we consider two cases: $f_h < f_l$ and $f_h \geq f_l$. In the first case, $P_{H|R} < \delta_0$ (and hence, $P_{H|C} > \delta_0$). It then follows that $\lambda_e^R(0, 1) \leq \lambda_e^R(f_h, f_l) \leq \lambda_e^R(1, 1)$ and $\lambda_e^C(0, 0) \leq \lambda_e^C(f_h, f_l) \leq \lambda_e^C(0, 1)$. We know that the function $\lambda_e^C(0, 1)$ becomes flat when λ increases to a certain value while $\lambda_e^R(1, 1)$ always strictly increases with λ . Hence, $\lambda_e^R(1, 1)$ crosses $\lambda_e^C(0, 1)$ exactly once and from below as λ increases. Denote this crossing point by $\bar{\lambda}_1^e$. It then follows that as long as $\lambda > \bar{\lambda}_1^e$, $\lambda_e^R(1, 1) > \lambda_e^C(f_h, f_l)$. Together with $\lambda_e^R(1, 1) \geq \lambda_e^R(f_h, f_l)$, we can conclude that ‘always revealing’ is the optimal choice for the service provider.

We now show the case where $f_h \geq f_l$. In this case, $P_{H|R} \geq \delta_0$ (and hence, $P_{H|C} \leq \delta_0$), and thus, $\lambda_e^R(1, 1) \leq \lambda_e^R(f_h, f_l) \leq \lambda_e^R(1, 0)$ and $\lambda_e^C(1, 0) \leq \lambda_e^C(f_h, f_l) \leq \lambda_e^C(0, 0)$. To show that ‘always revealing’ is the best strategy, we first introduce an arrival rate function which is always no less than $\lambda_e(f_h, f_l)$, and then show that the effective arrival rate under ‘always revealing’ can still outperform this arrival rate. Define

$$\bar{\lambda}_e(f_h, f_l) := [\delta_0 f_h + (1 - \delta_0) f_l] \lambda_e^R(f_h, f_l) + [\delta_0(1 - f_h) + (1 - \delta_0)(1 - f_l)] \lambda_e^C(0, 0).$$

This new function replaces the term $\lambda_e^C(f_h, f_l)$ in the expression of $\lambda_e(f_h, f_l)$ with a larger value term $\lambda_e^C(0, 0)$ and thus $\lambda_e(f_h, f_l) \leq \bar{\lambda}_e(f_h, f_l)$. Denote the optimal solution of maximizing $\bar{\lambda}_e(f_h, f_l)$ by $(\bar{f}_h^e, \bar{f}_l^e)$. Then, $\lambda_e(f_h^e, f_l^e) \leq \bar{\lambda}_e(\bar{f}_h^e, \bar{f}_l^e)$.

We now show that there exists a threshold $\bar{\lambda}_2^e$ such that when $\lambda > \max\{\bar{\lambda}_1^e, \bar{\lambda}_2^e\}$, ‘always

revealing' yields an effective arrival rate no less than $\bar{\lambda}_e(f_h^e, f_l^e)$. Recall that when the server reveals the queue length, customers join if and only if the queue length upon arrival (including themselves) is no greater than $n_e(f_h, f_l)$, where $n_e(f_h, f_l) = \lfloor V_R(f_h, f_l)\mu/\theta \rfloor$. Define a set of integers $S := \{n \in N_+ : n_e(1, 1) < n \leq n_e(1, 0)\}$, where N_+ is the set of all nonnegative integers. Clearly, when queue is observable, incoming customers' threshold $n_e(f_h, f_l)$ always falls into the set $S \cup \{n_e(1, 1)\}$, because the strategy $(1, 1)$ yields the lowest expected service value and the strategy $(1, 0)$ yields the largest expected service value for incoming customers. Therefore, we must have $n_e(\bar{f}_h^e, \bar{f}_l^e) \in S \cup \{n_e(1, 1)\}$. If the set S is empty, let $\bar{\lambda}_2^e = 0$; otherwise, we define $\bar{\lambda}_2^e$ through the following procedure. Let $n_1, \dots, n_{|S|}$ be all the elements in set S , where $|S|$ is the cardinality of S . We now fix the joining threshold $n_e(f_h, f_l) = n_i$ ($i = 1, \dots, |S|$) and consider the range $\lambda > \bar{\lambda}_1^e$. Consider the constrained maximization problem as follows:

$$(\bar{f}_h^i, \bar{f}_l^i) = \arg \max_{(f_h, f_l)} \{\bar{\lambda}_e(f_h, f_l) | n_e(f_h, f_l) = n_i, \lambda > \bar{\lambda}_1^e\}.$$

According to the definition of $\bar{\lambda}_1^e$, we have that when $\lambda > \bar{\lambda}_1^e$, $\lambda_e^R(1, 1) > \lambda_e^C(0, 1) > \lambda_e^C(0, 0)$. Also, as $f_h \geq f_l$, we have $\lambda_e^R(f_h, f_l) \geq \lambda_e^R(1, 1)$. Considering these two inequalities together, we get $\lambda_e^R(f_h, f_l) > \lambda_e^C(0, 0)$. With this inequality, we can then check the expression of $\bar{\lambda}_e(f_h, f_l)$. Now, the value of the term $\lambda_e^R(f_h, f_l)$ is fixed due to a fixed joining threshold n_i and the term $\lambda_e^C(0, 0)$ reaching a fixed value when $\lambda > \bar{\lambda}_1^e$. Maximizing $\bar{\lambda}_e(f_h, f_l)$ then requires to maximize the term $\delta_0 f_h + (1 - \delta_0) f_l$, which yields $\bar{f}_h^i = 1$ and $0 \leq \bar{f}_l^i < 1$ (note that \bar{f}_l^i cannot equal 1 under the constraint $n_e(f_h, f_l) = n_i$). Therefore, within the range $\lambda > \bar{\lambda}_1^e$, $(\bar{f}_h^e, \bar{f}_l^e)$ must be $(1, 1)$ or one of $(1, \bar{f}_l^i)$ ($i = 1, \dots, |S|$). Furthermore, we have that

$$\lim_{\lambda \rightarrow +\infty} [\lambda_e(1, 1) - \bar{\lambda}_e(1, \bar{f}_l^i)] = \mu - \left\{ [\delta_0 + (1 - \delta_0)\bar{f}_l^i]\mu + (1 - \delta_0)(1 - \bar{f}_l^i) \left(\mu - \frac{\theta}{V_l} \right) \right\} > 0.$$

Then, for $\bar{\lambda}_e(\bar{f}_h^i, \bar{f}_l^i)$, we can find a threshold for the potential arrival rate, $\bar{\lambda}_2^i$ ($\bar{\lambda}_2^i \geq 0$), such that when $\lambda > \bar{\lambda}_2^i$, $\lambda_e(1, 1) > \bar{\lambda}_e(\bar{f}_h^i, \bar{f}_l^i)$. Let $\bar{\lambda}_2^e$ be $\max\{\bar{\lambda}_2^1, \dots, \bar{\lambda}_2^{|S|}\}$ when the set S is nonempty. It follows that when $\lambda > \max\{\bar{\lambda}_1^e, \bar{\lambda}_2^e\}$, $\lambda_e(1, 1) \geq \bar{\lambda}_e(f_h, f_l) \geq \lambda_e(f_h, f_l)$ for $f_h \geq f_l$.

Finally, let $\bar{\lambda}^e := \max\{\bar{\lambda}_1^e, \bar{\lambda}_2^e\}$. We can then conclude that $(1, 1) = \arg \max_{(f_h, f_l)} \lambda_e(f_h, f_l)$ for $\lambda > \bar{\lambda}^e$.

Proof of Lemma 4.3.

(i) Recall that $u_e^R(\delta) = \lambda \sum_{j=0}^{n_e(\delta)-1} p_j^{n_e(\delta)} \left[\delta(V_h - V_l) + V_l - \frac{(j+1)\theta}{\mu} \right]$, with $n_e(\delta) = \lfloor [\delta V_h + (1 -$

$\delta)V_l)]\mu/\theta]$. As δ increases from 0 to 1, $n_e(\delta)$ repeats the following pattern: it first remains unchanged for a while, then increases by 1, and then remains unchanged, etc. When $n_e(\delta)$ remains unchanged, $u_e^R(\delta)$ is a linear function in δ with the slope being $\lambda(V_h - V_l) \left(1 - p_{n_e(\delta)}^{n_e(\delta)}\right)$. And when $n_e(\delta)$ increases by 1 at some $\delta = \hat{\delta}$, we have that $\hat{\delta}(V_h - V_l) + V_l - \frac{n_e(\hat{\delta})\theta}{\mu} = 0$. Notice that $p_j^{n_e(\hat{\delta})-1} < p_j^{n_e(\hat{\delta})}$ for $j = 0, \dots, n_e(\hat{\delta}) - 2$. Then, we can get that

$$\begin{aligned} \lim_{\delta \rightarrow \hat{\delta}^-} u_e^R(\delta) &= \lambda \sum_{j=0}^{n_e(\hat{\delta})-2} p_j^{n_e(\hat{\delta})-1} \left[\delta(V_h - V_l) + V_l - \frac{(j+1)\theta}{\mu} \right] \\ &> \lambda \sum_{j=0}^{n_e(\hat{\delta})-2} p_j^{n_e(\hat{\delta})} \left[\delta(V_h - V_l) + V_l - \frac{(j+1)\theta}{\mu} \right] = \lambda \sum_{j=0}^{n_e(\hat{\delta})-1} p_j^{n_e(\hat{\delta})} \left[\delta(V_h - V_l) + V_l - \frac{(j+1)\theta}{\mu} \right] \\ &= u_e^R(\hat{\delta}), \end{aligned}$$

which means that $u_e^R(\delta)$ jumps down at $\delta = \hat{\delta}$.

(ii) Recall that $u_e^C(\delta) = \lambda p_e(\delta) \left[\delta(V_h - V_l) + V_l - \frac{\theta}{\mu - \lambda p_e(\delta)} \right]$ with $p_e(\delta) = 1$ if $\lambda < \mu - \frac{\theta}{\delta(V_h - V_l) + V_l}$ and $p_e(\delta) = \frac{\mu - \theta / [\delta(V_h - V_l) + V_l]}{\lambda}$ otherwise. If $\lambda \geq \mu - \frac{\theta}{V_h}$, then $p_e(\delta) = \frac{\mu - \theta / [\delta(V_h - V_l) + V_l]}{\lambda}$ for all $\delta \in [0, 1]$, which makes $u_e^C(\delta)$ constant as 0; if $\lambda \leq \mu - \frac{\theta}{V_l}$, $p_e(\delta) = 1$ for all $\delta \in [0, 1]$, and thus $u_e^C(\delta)$ is linear increasing in δ with the slope $\lambda(V_h - V_l)$; otherwise, $p_e(\delta) = \frac{\mu - \theta / [\delta(V_h - V_l) + V_l]}{\lambda}$ for $\delta \in \left[0, \frac{\theta / (\mu - \lambda) - V_l}{V_h - V_l}\right]$, which makes $u_e^C(\delta)$ equal to 0, and $p_e(\delta) = 1$ for $\delta \in \left(\frac{\theta / (\mu - \lambda) - V_l}{V_h - V_l}, 1\right]$, which makes $u_e^C(\delta)$ linear increasing in δ .

(iii) Now, let us compare $u_e^R(\delta)$ and $u_e^C(\delta)$ under a given δ ($0 \leq \delta \leq 1$). First, when $\lambda \geq \mu - \frac{\theta}{\delta V_h + (1-\delta)V_l}$, we have $u_e^R(\delta) > 0$ but $u_e^C(\delta) = 0$, which directly yield $u_e^R(\delta) > u_e^C(\delta)$. Then, consider $0 < \lambda < \mu - \frac{\theta}{\delta V_h + (1-\delta)V_l}$. Under this case, we have $0 < \rho < 1$ and $p_e(\delta) = 1$. Note that in an $M/M/1/n_e(\delta)$ queue, the expected queue length is $E[L] := \sum_{j=0}^{n_e(\delta)} j p_j^{n_e(\delta)} = \frac{\rho}{1-\rho} - \frac{[n_e(\delta)+1]\rho^{n_e(\delta)+1}}{1-\rho^{n_e(\delta)+1}}$. Then, we can express $u_e^R(\delta)$ as

$$u_e^R(\delta) = \lambda[\delta V_h + (1-\delta)V_l] - \frac{(E[L] + 1)\lambda\theta}{\mu} - \lambda p_{n_e(\delta)}^{n_e(\delta)} \left[\delta V_h + (1-\delta)V_l - \frac{(n_e(\delta) + 1)\theta}{\mu} \right].$$

According to the definition of $n_e(\delta)$, we have $\delta V_h + (1-\delta)V_l - \frac{(n_e(\delta)+1)\theta}{\mu} < 0$. To sum up, we can get that

$$\begin{aligned} u_e^R(\delta) - u_e^C(\delta) &> \lambda[\delta V_h + (1-\delta)V_l] - \frac{(E[L] + 1)\lambda\theta}{\mu} - \lambda \left[\delta V_h + (1-\delta)V_l - \frac{\theta}{\mu - \lambda} \right] \\ &= \frac{\lambda\theta}{\mu - \lambda} - \frac{(E[L] + 1)\lambda\theta}{\mu} = \frac{[n_e(\delta) + 1]\rho^{n_e(\delta)+1}\lambda\theta}{(1 - \rho^{n_e(\delta)+1})\mu} > 0. \end{aligned}$$

References

- Aksin, Z., Armony, M. and V. Mehrotra. 2007. The modern call-center: A multi-disciplinary perspective on operations management research. *Production and Operations Management* 16(6): 665-688.
- Allon, G., Bassamboo, A. and I. Gurvich. 2011. “We will be right with you”: Managing customer expectations with vague promises and cheap talk. *Operations Research* 59(6): 1382-1394.
- Anupindi, R., Dada, M. and S. Gupta. 1998. Estimation of consumer demand with stock-out based substitution: An application to vending machine products. *Marketing Science* 17(4): 406-423.
- Armony, M. and C. Maglaras. 2004a. Contact centers with a call-back option and real-time delay information. *Operations Research* 52(4): 527-545.
- Armony, M. and C. Maglaras. 2004b. On customer contact centers with a call-back option: Customer decisions, routing rules and system design. *Operations Research* 52(2): 271-292.
- Armony, M., Shimkin, N. and W. Whitt. 2009. The Impact of delay announcements in many-server queues with abandonment. *Operations Research* 57(1): 66-81.
- Azoury, K. S. 1985. Bayes solution to dynamic inventory models under unknown demand distribution. *Management Science* 31(9): 1150-1160.
- Belobaba, P. 1987. Air travel demand and airline seat inventory management. Massachusetts Institute of Technology.
- Bensoussan, A., Çakanyldrm, M., Li, M., et al. 2016. Managing inventory with cash register information: Sales recorded but not demands. *Production and Operations Management* 25(1): 9-21.
- Bensoussan, A., Çakanyldrm, M. and S. P. Sethi. 2009. A note on The censored newsvendor and the optimal acquisition of information. *Operations Research* 57(3): 791-794.

- Bensoussan, A. and P. Guo. 2015. Technical note Managing nonperishable inventories with learning about demand arrival rate through stockout times. *Operations Research* 63(3): 602-609.
- Braden, D. J. and M. Freimer. 1991. Informational dynamics of censored observations. *Management Science* 37(11): 1390-1404.
- Brumelle, S. L., McGill, J. I., Oum, T. H., et al. 1990. Allocation of airline seats between stochastically dependent demands. *Transportation Science* 24(3): 183-192.
- Burnetas, A. and A. Economou. 2007. Equilibrium customer strategies in a single server Markovian queue with setup times. *Queueing Systems: Theory and Applications* 56(3-4): 213-228.
- Chen, B. and X. Chao. 2019. Parametric demand learning with limited price explorations in a backlog stochastic inventory system. *IIE Transactions* 51(6): 605-613.
- Chen, B. and X. Chao. 2020. Dynamic inventory control with stockout substitution and demand learning. *Management Science* Articles in Advance: 1-20.
- Chen, H. and M. Frank. 2004. Monopoly pricing when customers queue. *IIE Transactions* 36(6): 1-13.
- Chen, L. 2010. Bounds and heuristics for optimal Bayesian inventory control with unobserved lost sales. *Operations Research* 58(2): 396-413.
- Chen, L. and E. L. Plambeck. 2008. Dynamic inventory management with learning about the demand distribution and substitution probability. *Manufacturing & Service Operations Management* 10(2): 236-256.
- Chu, L. Y., Shanthikumar, J. G. and Z. J. M. Shen. 2008. Solving operational statistics via a Bayesian analysis. *Operations Research Letters* 36(1): 110-116.
- Cooper, W. L. and L. Li. 2012. On the use of buy up as a model of customer choice in revenue management. *Production and Operations Management* 21(5): 833-850.
- Cooper, W. L., Homem-de-Mello, T. and A. J. Kleywegt. 2006. Models of the spiral-down effect in revenue management. *Operations Research* 54(5): 968-987.
- Cui, S. and S. Veeraraghavan. 2016. Blind queues: The impact of consumer beliefs on revenues and congestion. *Management Science* 62(12): 3656-3672.
- Debo, L., Parlur, C. and U. Rajan. 2012. Signalling quality via queues. *Management Science* 58(5): 876-891.

- Debo, L., Rajan, U. and S. Veeraraghavan. 2020. Signaling quality via long lines and uninformative prices. *Manufacturing & Service Operations Management* 22(3): 513-527.
- Debo, L. and S. Veeraraghavan. 2014. Equilibrium in queues under unknown service times and service value. *Operations Research* 62(1): 38-57.
- Ding, X., Puterman, M. L. and A. Bisi. 2002. The censored newsvendor and the optimal acquisition of information. *Operations Research* 50(3): 517-527.
- Doya, K., Ishii, S., Pouget, A. and R. P. Rao. 2007. Bayesian brain: Probabilistic approaches to neural coding. MIT Press, Cambridge.
- Edelson, N. and D. Hildebrand. 1975. Congestion tolls for Poisson queueing processes. *Econometrica* 43(1): 81-92.
- Feng, Q. and J. G. Shanthikumar. 2017. Supply and demand functions in inventory models. *Operations Research* 66(1): 77-91.
- Gentzkow, M. and E. Kamenica. 2016. A Rothschild-Stiglitz approach to Bayesian persuasion. *American Economic Review* 106(5): 597-601.
- Glimcher, P. W. and E. Fehr. 2013. Neuroeconomics: Decision making and the brain. Academic Press.
- Grossman, S. J. 1981. The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics* 24(3): 461-483.
- Grossman, S. J. and M. Perry. 1986. Perfect sequential equilibrium. *Journal of Economic Theory* 39(1): 97-119.
- Guo, P., Haviv, M., Luo, Z. and Y. Wang. 2020. Optimal queue length information disclosure when service quality is uncertain: A Bayesian persuasion approach. Available at SSRN: <https://ssrn.com/abstract=3599693>
- Guo, P. and P. Zipkin. 2007. Analysis and comparison of queues with different levels of delay information. *Management Science* 53(6): 962-970.
- Guo, P. and R. Hassin. 2011. Strategic behavior and social optimization in Markovian vacation queues. *Operations Research* 59(4): 986-997.
- Harrison, J. M., Keskin, N. B. and A. Zeevi. 2012. Bayesian dynamic pricing policies: Learning and earning under a binary prior distribution. *Management Science* 58(3): 570-586.
- Hassin, R. 1986. Consumer information in markets with random product quality: The case

- of queues and balking. *Econometrica* 54(5): 1185-1195.
- Hassin, R. 2007. Information and uncertainty in a queueing system. *Probability in the Engineering and Informational Sciences* 21(3): 361-380.
- Hassin, R. 2016. *Rational Queueing*. Chapman and Hall/CRC. London.
- Hassin, R. and M. Haviv. 1994. Equilibrium strategies and the value of information in a two line queueing system with threshold jockeying. *Stochastic Models* 10(2): 415-435.
- Hassin, R. and M. Haviv. 2003. *To queue or not to queue: Equilibrium behavior in queueing systems*. Kluwer Academic Publishers. London.
- Hassin, R. and R. Roet-Green. 2017. The impact of inspection cost on equilibrium, revenue, and social welfare in a single-server queue. *Operations Research* 65(3): 804-820.
- Hassin, R. and R. Roet-Green. 2018. Cascade equilibrium strategies in a two-server queueing system with inspection cost. *European Journal of Operational Research* 267(3): 1014-1026.
- Hu, M., Li, Y. and J. Wang. 2018. Efficient ignorance: Information heterogeneity in a queue. *Management Science* 64(6): 2650-2671.
- Hu, P., Stephen, S. and Y. Man. 2015. Joint inventory and markdown management for perishable goods with strategic consumer behavior. *Operations Research* 64(1): 118-134.
- Huh, W. T. and P. Rusmevichientong. 2009. A nonparametric asymptotic analysis of inventory planning with censored demand. *Mathematics of Operations Research* 34(1): 103-123.
- Ibrahim, R. 2018. Sharing delay information in service systems: A literature survey. *Queueing Systems* 89(1-2): 49-79.
- Ibrahim, R., Armony, M. and A. Bassamboo. 2017. Does the past predict the future? The case of delay announcements in service systems. *Management Science* 63(6): 1762-1780.
- Jain, A., Rudi, N. and T. Wang. 2015. Demand estimation and ordering under censoring: Stock-out timing is (almost) all you need. *Operations Research* 63(1): 134-150.
- Kamenica, E. and M. Gentzkow. 2011. Bayesian persuasion. *American Economic Review* 101(6): 2590-2615.
- Knill, D. C. and A. Pouget. 2004. The Bayesian brain: The role of uncertainty in neural coding and computation. *TRENDS in Neurosciences* 27(12): 712-719.
- Kök, A. G. and M. L. Fisher. 2007. Demand estimation and assortment optimization under substitution: Methodology and application. *Operations Research* 55(6): 1001-1021.

- Koutsoupias, E. and C. Papadimitriou. 1999. Worst-case equilibria. *Annual Symposium on Theoretical Aspects of Computer Science*. Springer. Berlin.
- Kremer, M. and L. Debo. 2016. Inferring quality from wait time. *Management Science* 62(10): 3023-3038.
- Kreps, D. M. and R. Wilson. 1982. Sequential equilibria. *Econometrica* 50(4): 863-894.
- Lariviere, M. A. and E. L. Porteus. 1999. Stalking information: Bayesian inventory management with unobserved lost sales. *Management Science* 45(3): 346-363.
- Li, K., Cui S. and J. Wang. 2020. On the optimal disclosure of queue length information. *Naval Research Logistics* forthcoming.
- Li, R. and J. K. Ryan. 2011. A Bayesian inventory model using real-time condition monitoring information. *Production and Operations Management* 20(5): 754-771.
- Li, R., Song, J. S., Sun, S., et al. 2019. Fight inventory shrinkage: Simultaneous learning of inventory level and shrinkage rate. Working paper, Syracuse University.
- Lingenbrink, D. and K. Iyer. 2019. Optimal signaling mechanisms in unobservable queues. *Operations Research* 67(5): 1397-1416.
- Littlewood, K. 1972. Forecasting and control of passenger bookings. *Airline Group International Federation of Operational Research Societies Proceedings* 12: 95-117.
- Liyanage, L. H. and J. G. Shanthikumar. 2005. A practical inventory control policy using operational statistics. *Operations Research Letters* 33(4): 341-348.
- Lovejoy, W. S. 1990. Myopic policies for some inventory models with uncertain demand distributions. *Management Science* 36(6): 724-738.
- Lu, X., Song, J. S. and K. Zhu. 2005. On The censored newsvendor and the optimal acquisition of information. *Operations Research* 53(6): 1024-1026.
- McGill, J. I. and G. van Ryzin. 1999. Revenue management: Research overview and prospects. *Transportation Science* 33(2): 233-256.
- Milgrom, P. R. 1981. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*: 380-391.
- Miller, B. L. 1986. Scarf's state reduction method, flexibility, and a dependent demand inventory model. *Operations Research* 34(1): 83-90.
- Naor, P. 1969. The regulation of queue size by levying tolls. *Econometrica* 37(1): 15-24.

- Pfeifer, P. E. 1989. The airline discount fare allocation problem. *Decision Sciences* 20(1): 149-157.
- Rayo, L. and I. Segal. 2010. Optimal information disclosure. *Journal of Political Economy* 118(5): 949-987.
- Scarf, H. 1959. Bayes solutions of the statistical inventory problem. *The Annals of Mathematical Statistics* 30(2): 490-508.
- Scarf, H. 1960. Some remarks on Bayes solutions to the inventory problem. *Naval Research Logistics* 7(4): 591-596.
- Selten, R. 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* 4: 25-55.
- Shaw, S. 1982. Air transport: A marketing perspective. Pitman Books, London
- Susskind, A. M., Reynolds, D. and E.Tsuchiya. 2004. An evaluation of guests preferred incentives to shift time-variable demand in restaurants. *Cornell Hotel and Restaurant Administration Quarterly* 45(1): 68-84.
- The Alan Turing Institute. 2020. Dynamic forecasting with British Airways. Retrieved Jul 2021, from <https://www.turing.ac.uk/research/impact-stories/dynamic-forecasting-british-airways>.
- Topkis, D. M. 2011. Supermodularity and complementarity. Princeton University Press.
- U.S. Department of Transportation. 2019. Fly rights: A consumer guide to air travel. Retrieved Jul 2021, from <https://www.transportation.gov/airconsumer/fly-rights>.
- Ulu, C., Honhon, D. and A. Alptekinolu. 2012. Learning consumer tastes through dynamic assortments. *Operations Research* 60(4): 833-849.
- van Ryzin, G. and J. McGill. 2000. Revenue management without forecasting or optimization: An adaptive algorithm for determining airline seat protection levels. *Management Science* 46(6): 760-775.
- Veeraraghavan, S. and L. Debo. 2009. Joining longer queues: Information externalities in queue choice. *Manufacturing & Service Operations Management* 11(4): 543-562.
- Veeraraghavan, S. and L. Debo. 2011. Herding in queues with waiting costs: Rationality and regret. *Manufacturing & Service Operations Management* 13(3): 329-346.
- Vulcano, G., van Ryzin, G. and R. Ratliff. 2012. Estimating primary demand for substitutable products from sales transaction data. *Operations Research* 60(2): 313-334.

- Wang, J. and M. Hu. 2020. Efficient inaccuracy: User-generated information sharing in a queue. *Management Science* articles in advance.
- Wang, S. and G. F. Özkan-Seely. 2018. Signaling product quality through a trial period. *Operations Research* 66(2): 301-312.
- Whitt, W. 1999. Improving service by informing customers about anticipated delays. *Management Science* 45(2): 192-207.
- Yang, L., Debo, L. and V. Gupta. 2019. Search among queues under quality differentiation. *Management Science* 65(8): 3605-3623.
- Yu, M., Debo, L. and R. Kapuscinski. 2016. Strategic waiting for consumer-generated quality information: Dynamic pricing of new experience goods. *Management Science* 62(2): 410-435.
- Yu, Q., Allon, G. and A. Bassamboo. 2017. How do delay announcements shape customer behavior? An empirical study. *Management Science* 63(1): 1-20.
- Yu, Q., Allon, G. and A. Bassamboo. 2021. The reference effect of delay announcements: A field experiment. *Management Science* Articles in Advance: 1-19.
- Yu, Q., Allon, G., Bassamboo, A. and S. Iravani. 2018. Managing customer expectations and priorities in service systems. *Management Science* 64(8): 3942-3970.
- Yuan, H., Luo, Q. and C. Shi. 2021. Marrying stochastic gradient descent with bandits: Learning algorithms for inventory systems with fixed costs. *Management Science* Articles in Advance: 1-27.

