

2016

Advance reservations and information sharing in queues with strategic customers

<https://hdl.handle.net/2144/19499>

"Downloaded from OpenBU. Boston University's institutional repository."

BOSTON UNIVERSITY
COLLEGE OF ENGINEERING

Dissertation

**ADVANCE RESERVATIONS AND INFORMATION
SHARING IN QUEUES WITH STRATEGIC CUSTOMERS**

by

ERAN SIMHON

B.S., Technion - Israel Institute of Technology, 2006
M.S., Technion - Israel Institute of Technology, 2012

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

2016

© 2016 by
ERAN SIMHON
All rights reserved

Approved by

First Reader

David Starobinski, PhD
Professor of Electrical and Computer Engineering
Professor of Systems Engineering

Second Reader

Christos G. Cassandras, PhD
Professor of Electrical and Computer Engineering
Professor of Systems Engineering

Third Reader

Ari Trachtenberg, PhD
Professor of Electrical and Computer Engineering
Professor of Systems Engineering

Fourth Reader

Yezekael Havel, PhD
Professor of Computer Science and Engineering
University of Avignon

Acknowledgments

First and foremost, I would like to express my deep appreciation and gratitude to my supervisor, Prof. David Starobinski, for his endless support, advice and guidance through our research.

Great thanks to my collaborators Prof. Yezekael Hayel, Prof. Quanyan Zhu, Carrie Cramer, Zachary Lister, Eddy Luo, Yoni Klausner and Christopher Liao. Special thanks are due to Prof. Ari Trachtenberg for providing me valuable feedback throughout my PhD research.

I would also like to thank my lab mates and friends at BU: Liangxiao Xin, Feng Nan, Sepideh Pourazarm, Theodora Brisimi, Maxim Timchenko, Emir Kavurmacioglu, Morteza Hashemi and Wei Si, for their good advice and for making my days joyful.

I would also like to thank my committee members Prof. Christos Cassandras and Prof. Iohannis Paschalis for their valuable comments and suggestions. I wish to express my gratitude to the Division of Systems Engineering and the Center for Information and Systems Engineering at Boston University for the financial and institutional support. Special thanks to Elizabeth Flagg, Denise Joseph and Christina Polyzos, for their endless patience and willingness to help.

Finally, I would like to thank my family. To my parents for their love and support. To my spouse Michael, whose endless love and encouragement gave me the strength to accomplish this dissertation.

ADVANCE RESERVATIONS AND INFORMATION SHARING IN QUEUES WITH STRATEGIC CUSTOMERS

ERAN SIMHON

Boston University, College of Engineering, 2016

Major Professor: David Starobinski, PhD
Professor of Electrical and Computer Engineering
Professor of Systems Engineering

ABSTRACT

In many branches of the economy, including transportation, lodging, and more recently cloud computing, users can reserve resources in advance. Although advance reservations are gaining popularity, little is known about the strategic behavior of customers facing the decision whether to reserve a resource in advance or not.

Making an advance reservation can reduce the waiting time or the probability of not getting service, but it is usually associated with an additional cost. To evaluate this trade-off, we develop a game-theoretic framework, called *advance reservation games*, that helps in reasoning about the strategic behavior of customers in systems that allow advance reservations. Using this framework, we analyze several advance reservation models, in the context of slotted loss queues and waiting queues. The analysis of the economic equilibria, from the provider perspective, yields several key insights, including: *(i)* If customers have no a-priori information about the availability of servers, then only customers granted service should be charged a reservation fee; *(ii)* Informing customers about the exact number of available servers is less profitable than only informing them that servers are available; *(iii)* In many cases, the reservation fee

that leads to the equilibrium with maximum possible profit leads to other equilibria, including one resulting with no profit; *(iv)* If the game repeats many times and customers update their strategy after observing actions of other customers at previous stage, then the system converges to an equilibrium where no one makes an advance reservation, if such an equilibrium exists. Else, the system cycles and yields positive profit to the provider

Finally, we study the impact of information sharing in $M/M/1$ queues with strategic customers. We analyze the intuitive policy of sharing the queue length with customers when it is small and hiding it when it is large. We prove that, from the provider perspective, such a policy is never optimal. That is, either always sharing the queue length or always hiding it maximizes the average number of customers joining the queue.

Contents

1	Introduction and Motivation	1
2	Background and Related Work	12
2.1	Game Theory	12
2.2	Queueing Theory	14
2.3	Queueing Games	15
2.4	Advance Reservations	17
2.5	Learning Models	18
3	Advance Reservation Games	20
3.1	Model Description	20
3.2	Equilibria Analysis	23
3.2.1	Classification of the equilibria	23
3.2.2	Equilibria structure	29
3.3	Revenue Maximization	34
3.4	Price of Conservatism	38
3.4.1	Single-server Case	40
3.4.2	Many-server Case	43
3.5	Summary	47
4	Advance Reservation Games with Information Sharing	49
4.1	Binary-information Game	50
4.2	Full-information Game	52
4.2.1	Preliminaries	52

4.2.2	Equilibria Analysis	56
4.2.3	Impact of the AR Cost	62
4.2.4	Poisson Distributed Demand	71
4.3	Partial-information Model	74
4.4	Comparison of Information Sharing Policies	79
4.4.1	Simulation Results	80
4.4.2	Numerical Example	83
4.5	Summary	84
5	Waiting Queues with Advance Reservations	86
5.1	Game Description	86
5.2	Equilibrium Analysis	88
5.3	Revenue Maximization	94
5.3.1	Price of Conservatism	96
5.4	Summary	99
6	Dynamic AR Games	100
6.1	Preliminaries and Learning Process	100
6.2	Learning Analysis	103
6.3	Stability Analysis	107
6.4	The Dynamics of a Game with Unique Equilibrium	110
6.4.1	Strategy-learning	110
6.4.2	Action-learning	111
6.5	The Dynamics of a Game with Multiple Equilibria	115
6.5.1	Strategy-learning	115
6.5.2	Action-learning	117
6.6	Learning in Loss systems	118
6.7	Summary	119

7	Information Sharing in Waiting Queues	121
7.1	Model Description	122
7.2	Equilibrium Analysis of Threshold Policies	124
7.3	Comparison of Policies	128
7.4	Summary	134
8	Concluding Remarks and Future Directions	135
	References	139
	Curriculum Vitae	144

List of Tables

3.1	Payoff summary.	22
6.1	Stability summary.	110

List of Figures

3·1	Slotted system with two servers. The demand for each slot is independent. The demand for slots 1, 2 and 3, is respectively 3, 2 and 4.	22
3·2	Example of a realization of the demand in a given slot. The service starts at arrival time 1. As customers have smaller (normalized) arrival times, they have the opportunity to reserve a server earlier.	24
3·3	An example with 10 servers, Poisson distributed demand with mean 15 and AR cost $C = 0.45$. The line C and the function $C(\tau_e)$ intersect twice. Hence, there are two <i>some-make-AR</i> equilibria. Since $C(1) < C$, there is also a <i>none-make-AR</i> equilibrium.	32
3·4	The cost function in the two <i>no-information</i> models with $N = 10$. The demand is a Poisson random variable with mean equals 10.	34
3·5	An example with $N = 15$ servers and Poisson distributed demand with parameter $\lambda = 20$. To fee that maximizes the revenue belongs to a range with multiple equilibria, including <i>none-make-AR</i>	39
3·6	$\lambda = 1.2$, $N = 1$. The reservation fee $C(\tau_e)$ is a monotonically increasing function of the threshold τ_e . The revenue function $R(\tau_e)$ is concave with maximum value at $\tau_e^* = 0.523$	42

4.1	An example of a system with 6 servers and Poisson distributed demand with mean $\lambda = 10$. The equilibrium is determined by the expected payoff of a customer that arrives at the 6-th threshold and observes 6 available servers. If that customer is better off not making AR (a), then there is a none-make-AR equilibrium. If there is a threshold such that the customer is indifferent between the two actions (b), then there is a some-make-AR equilibrium.	71
4.2	Illustration of Example 5. For each number of available servers $i \in 1, 2, \dots, 6$, the probability of service π_i as a function of the threshold τ (which represents the expected payoff of a threshold customer that is not making AR) intersects once with the line $1 - C$ (which represents the payoff of any customer making AR), thus showing the uniqueness of the equilibrium. As the number of available servers i decreases, the time threshold at which customers stop making AR increases (i.e., customers are more likely to make AR when observing fewer available servers).	74
4.3	In example 6, the probability of service of the uninformed threshold customer as a function of the threshold $\pi_u(\tau_u)$ intersects twice with the line 0.85, which is the payoff if making AR.	77
4.4	The expected payoff of the uninformed threshold customer as a function of the threshold. The number of times the line $1 - C$ intersects this function determines the set of equilibria.	79
4.5	A comparison between the number of customers making AR in a system with $N = 10$ servers. As less information is shared, more customers are likely to make AR.	82

4·6	The probability of service of the threshold customers as a function of the thresholds. The fact that τ^e is much closer to τ_2^e than to τ_1^e explains why the number of reservations is larger with the <i>binary-information</i> policy than with the <i>full-information</i> policy.	84
5·1	An illustration of the model with three customers. Customer 1 makes a reservation at time t_1 , and he is served upon arrival at s_1 . Customer 2 also makes AR and is served upon arrival, but his service is preempted by customer 1 which made reservation earlier. Customer 3 is served only when the service of customer 2 is completed.	88
5·2	$C(\tau)$ when $\rho > 0.5$. The number and types of equilibria is determined by the value of C	94
5·3	When the utilization $\rho < 2/3$, the optimal fee C^* leads to a unique equilibrium (a). When $\rho > 2/3$, C^* leads to multiple equilibria (b).	97
5·4	The maximum possible revenue and the maximum guaranteed revenue in a system with parameters $\lambda = 45$ and $\mu = 60$	98
6·1	An example with $\lambda = 45$ and $\mu = 60$. The <i>some-make-AR</i> equilibria split the belief range to intervals. At each interval, one expected cost is smaller than the other.	105
6·2	Simulation results. In a game with unique equilibrium, more customers make AR under <i>action-learning</i> than under <i>strategy-learning</i>	114
6·3	Convergence to equilibrium under <i>strategy-learning</i>	116

7.1	A Markov chain representation of the system evolution. D is the threshold above which customers are not informed about the queue length and q is the probability that uninformed customers join the queue. In this example, we assume $D < L - 1$, that is, informed customers always join the queue.	126
7.2	A queue with $\lambda = 9.8$, $\mu = 10$ and $U = 2$. The optimal policy is u_- . . .	133
7.3	A queue with $\lambda = 11$, $\mu = 10$ and $U = 2$. The optimal policy is u_+ . . .	134

List of Abbreviations

AR	Advance reservation
CCDF	Complimentary cumulative distribution function
CDF	Cumulative distribution function
FCFS	First-come-first-served
M/D/1	A single server queue with memoryless arrivals and deterministic service	
M/M/1	A single server queue with memoryless arrivals and service
PMF	Probability mass function
PoC	Price of Conservatism

Chapter 1

Introduction and Motivation

Advance reservation (AR) services form a pillar of the economy. They are widely deployed in the industries of transportation (e.g., for reserving airplane and train tickets), lodging (e.g., for booking hotel rooms), and health care (e.g., for scheduling medical appointments). AR is also gaining popularity in communication networks (Guok et al., 2006; Charbonneau and Vokkarane, 2012) and cloud computing (Sotomayor et al., 2009; Xie et al., 2012; Niu et al., 2012). The rise of the shared economy expands the applications of AR to new fields such as shared cars (Bardhi and Eckhardt, 2012) and shared parking spaces (Geng and Cassandras, 2012).

Supporting AR benefits the service provider since knowledge about future demand can improve resource management and quality-of-service (Charbonneau and Vokkarane, 2012). Customers are motivated to reserve in advance, since such a reservation decreases their expected waiting time or their probability of not getting service.

A large portion of the existing research of advance reservations focuses on algorithmic aspects, such as scheduling and routing (Gu erin and Orda, 2000; Cohen et al., 2009; Fazlollahi and Starobinski, 2015). Yet, in services supporting AR, it is often up to the customers to decide whether to make a reservation in advance. Hence, understanding the strategic behavior of customers in systems that support AR is a fundamental problem.

When designing a system that allows advance reservations, the service provider

controls several parameters such as:

- The AR fee (price);
- The charging scheme;
- The information shared with customers.

In practice, services that allow advance reservations often use different parameters. Yet, the impact of these parameters on the decisions of customers and on the performance of the system is not well understood.

In this dissertation, we introduce a game-theoretic framework, called *Advance Reservation (AR) games*, that helps in reasoning about the strategic behavior of customers in systems supporting advance reservations. Throughout the dissertation, we use this framework to evaluate how different policies and prices, established by a service provider, affect the strategic behavior of customers and, in turn, the economic outcomes of different queueing systems.

Queueing systems are typically divided into two main types. The first type is *waiting queues*. In such queues, customers wait for service. Examples include emergency rooms, banks, or cloud resources. Reserving resources in advance in such systems decreases the expected waiting time. The second type is *loss queues*. In such queues, customers that encounter a busy system (i.e., all servers are busy) leave the queue without being served (e.g., it is assumed that alternative services exist). For example, consider hotel rooms or car rentals. Reserving resources in advance in such systems decreases the probability of not getting service (also known as *blocking probability*). In this dissertation, we use the AR games framework to study both types of queues. Note that throughout the dissertation, we refer to a service provider as a “she” and to a customer as a “he”.

In Chapters 3 and 4, we study the behavior of customers in a slotted loss queue with N servers that supports AR. The demand for service follows a general discrete distribution. We assume that customers differ by how much time in advance they realize that service is required (we refer to the realization point as the *arrival time*). One of our goals is to capture the impact of the customers' arrival times on their strategic behavior. As shown throughout the dissertation, at equilibrium, customers that arrive early tend to reserve in advance in contrast to customers that arrive late.

Upon realizing that service is required, customers can either make advance reservation, or defer their decisions and request service on the spot. Making AR decreases the probability of being blocked, but bears an additional *reservation cost*. The AR cost may be a fee set by the provider (common in buying tickets for sport and entertainment events), the time or resources required for making the reservation (consider the case where making reservation requires making a phone call or physical presence), or the cost of financing advance payment (common when buying air tickets).

In Chapter 3, we consider a setting in which information about availability of servers is not shared with the customers. Under this setting, customers may make AR request but not get service. We study two versions of this game. In the first version, customers that try to reserve a server in advance but fail do not bear the reservation costs. For example, assume that the reservation cost corresponds to the financial cost of an advance payment. If a reservation is not made, this cost is spared. In the second version, we assume that the cost of reservation is applied to all customers attempting making AR. For example, consider a customer that waits in line to buy a ticket for a show only to find out at the end that the show is sold out. In this case, the reservation cost corresponds to the time spent waiting, and this cost is applied whether service is provided or not.

We determine the equilibrium structures of those two games. We show that,

at equilibrium, the strategy followed by all customers has a threshold form, where only customers with arrival time smaller than some threshold make AR. We refer to an equilibrium with threshold zero as a *none-make-AR* equilibrium and to an equilibrium with threshold between 0 and 1 as a *some-make-AR* equilibrium. In the former equilibrium, it is guaranteed that none of the customers will make AR, while in the latter equilibrium, due to the stochastic nature of the problem, the number of reservations is a random variable. We show that the equilibrium is not necessarily unique and in some situations a game can have both *some-make-AR* and *none-make-AR* equilibria.

We next assume that the AR cost is a fee charged by the provider, who can decide whether to charge the fee from all customers attempting AR or only from those granted service. By comparing the equilibrium outcomes of those two policies, we show that charging a fee from all customers attempting to make AR, including those not granted service, is never optimal.

After showing that charging AR fee only from customers granted service is preferable, another question arises: What is the AR fee that maximizes the providers expected revenue from AR fees? The answer to this question turns out to be more complicated. We show that the fee which leads to the equilibrium with the maximum revenue may also lead to an equilibrium with no revenue. Therefore, in order to properly set the AR fee, the provider should consider both the fee yielding the maximum possible revenue and the fee yielding the maximum guaranteed revenue. For this purpose, we introduce the concept of price of conservatism (PoC), which corresponds to the ratio of the maximum possible revenue to the maximum guaranteed revenue. We assume that the demand follows a Poisson distribution and derive the price of conservatism in different settings. First, we analyze the case of a single server, where we prove that $\text{PoC} = 1$ (i.e., no loss). Next, we conduct the analysis of a

many-server system and prove that the price of conservatism can be arbitrarily high. This situation occurs when the system is slightly overloaded. We note that since AR games are zero-sum games, the social welfare (i.e., the total payoff of all players in the game, including the provider) is not affected by the decisions of the provider and customers. Therefore, the price of anarchy (Koutsoupias and Papadimitriou, 1999) in such games always equals one. In contrast, the price of conservatism measures the loss of profit from the viewpoint of the provider.

In Chapter 4, we use the framework of AR games to explore the impact of sharing information about server availability with customers. In recent years, several works used game-theoretic tools to study the impact of information on the decision whether to join a queue or to balk. For examples, see (Guo and Zipkin, 2007; Hassin, 2007). Those papers show that providing additional information (besides the queue parameters) is beneficial in some cases, while in other cases it can hurt either the provider revenue or the social welfare. In this dissertation, we aim to determine how different information sharing policies impact the decision of customers whether to make AR or not.

A provider, aiming to maximize the number of advance reservations, can adopt several information sharing policies. For example, she may share with customers exact information about the number of available servers. This practice is common in entertainment services, where customers are allowed to choose their seats and observe the exact number of available seats before making the reservation (e.g., using Ticketmaster reservation services¹). Another option is to hide information about the number of available servers and only to inform customers if servers are available or not. This policy is common with airlines such as Delta Airlines which allows customers to

¹See www.ticketmaster.com, accessed on 06/09/2016.

choose their seats but only after buying their tickets². The provider may also choose a middle ground solution where information about the number of available servers is shared, but only when a few servers remain available. This policy is implemented in lodging websites, such as Booking.com, that alert potential customers when five rooms or less are left³.

To evaluate the impact of information on the decisions of customers, we consider the following set-up: customers first make a reservation inquiry and receive information about the number of available servers. If no servers are available, they leave the system with no gain or loss. Otherwise, they decide whether to make AR or not. Making AR guarantees service but is associated with an additional cost.

First, we study a *binary-information-game*. In this version of the game, when customers make an inquiry, they are only informed whether servers are available or not. We show that this game has the same outcome as the first version of games with *no-information* sharing.

Next, we study a *full-information game*. In this game, customers observe the exact number of available servers. Therefore, their decisions depend on both their arrival times and the number of available servers. We show that, at equilibrium, customers follow a threshold strategy as in the previous games. However, a threshold strategy, in this game, consists of N thresholds instead of one. We also show that this game has a unique equilibrium. High AR costs lead to a *none-make-AR* equilibrium, while low AR costs lead to a *some-make-AR* equilibrium.

We then consider a *partial-information game*. In this game, the provider informs customers about the number of available servers only if this number falls below a certain value. We assume that customers that are not informed realize that the

²See www.delta.com, accessed on 06/09/2016.

³See www.booking.com, accessed on 06/09/2016

number of available servers is greater than that threshold and take this fact under consideration upon making their decisions. We show that, in this game, the same types of equilibria exist (i.e., *none-make-AR* and *some-make-AR*). However, the game may have multiple equilibria.

Informing customers about the (exact) number of available servers leads to a trade-off. On the one hand, customers that observe that all or almost all servers are available are less likely to make a reservation (compared to a system that does not share availability information). On the other hand, customers that observe that only a few servers are available are more likely to make a reservation. To evaluate this trade-off we use simulations. We show that, on average, the number of reservations decreases as more information is provided to the customers. More specifically, the *full-information* policy yields the lowest number of reservations. In the *partial-information* policies, the number of customers making advance reservation increases as the threshold is lowered, and the greatest number of reservations is achieved when no information at all is provided.

In Chapter 5, we switch our focus from loss queues to waiting queues. Many services, such as health care, banking and cloud computing, combine both first-come-first-served (FCFS) scheduling policy and advance reservations. The motivation to make reservations in advance in such queues is no longer to decrease or eliminate the probability of not getting service, but to decrease the waiting time.

We assume that the time axis is divided into two time periods: *reservations period* and *service period*. This restriction simplifies the analysis and is common in the literature on advance reservations (Virtamo, 1992; Yessad et al., 2007; Syed et al., 2008). During the reservation period, each customer realizes that he will need service at a specific future time point. Upon realizing that, they decide whether to make a reservation or not. Since customers are assumed to be strategic, they will make

reservations only if it reduces their expected cost which consists of the reservation cost (if making a reservation) and the cost of waiting.

We derive the equilibrium structure of the game and show that, as in the loss queue, two possible types of equilibria prevail. In one equilibrium, none of the customers makes AR. In the second equilibrium customers that realize early enough that they will need future service make AR.

We show that if the utilization of the queue (i.e., the ratio between the arrival rate and the service rate) is smaller than 0.5, the game has a unique equilibrium. Low AR costs lead to a *some-make-AR* equilibrium, while high AR costs lead to a *none-make-AR* equilibrium. If the utilization is greater than 0.5, a middle range of AR costs also exists. An AR cost belonging to that range leads to three equilibria, including *none-make-AR* and two *some-make-AR*. For this model, we derive a closed form terms for the critical values of AR cost that determines which equilibria prevail.

Assuming that the AR cost is a fee charged by the service provider, we next analyze the game from the prospective of a provider aiming to maximize her revenue from AR fees. We show that if the utilization of the queue is greater than $2/3$, then the maximizing fee leads to multiple equilibria. Thus, charging that fee may lead to the highest possible revenue but also to no revenue.

The AR games analyzed so far assume that all customers follow an equilibrium strategy. To relax this assumption, we study, in Chapter 6, a dynamic version of the game, in which customers initially follow an arbitrary threshold strategy. Our goal is to find whether the system converges to an equilibrium or cycles. If it converges and multiple equilibria exist, we aim to find to which equilibrium the system converges.

We use *best response dynamics* and distinguish between *strategy-learning* and *action-learning*. In *strategy-learning*, customers obtain information about strategies adopted at previous steps, while in *action-learning*, customers *estimate* the previous

strategies by obtaining information about the actions taken at previous steps. Our analysis shows that starting with any initial belief about customers behavior (i) Under *strategy-learning*, the system always converges to an equilibrium; (ii) Under *action-learning*, convergence is not guaranteed and if convergence occurs, it can only be to a *none-make-AR* equilibrium; (iii) If the equilibrium is unique, more customers, on average, make reservations under *action-learning* than under *strategy-learning*. Those results are valid for both waiting queues and loss queues with no information sharing.

In Chapter 7, we study the impact of information sharing in $M/M/1$ queues with strategic customers that need to decide whether to join the queue or to balk. The literature on the strategic behavior of customers in $M/M/1$ queues is traditionally divided into the classical observable and unobservable queues. In the former case (Naor, 1969), customers are informed about the current queue length before deciding whether to join or balk. In the latter case (Edelson and Hilderbrand, 1975), customers make their decisions based on statistical information (e.g., the queue parameters). In both cases, customers behave strategically and join the queue only if their expected waiting cost is smaller than the reward obtained upon being served.

The goal of our work is to find out whether there are situations where the service provider can increase her revenue by combining those two frameworks. In particular, we are interested in studying policies where the provider shares information with some customers and hides it from others, depending on the actual queue length. We assume that the service provider has a fixed income from each customer that joins the queue. Thus, in order to maximize the revenue, the provider should maximize the effective arrival rate, which is the rate of customers that join the queue, or equivalently, minimize the idle period of the system.

We compare between the outcomes of the following three policies: (i) always inform customers about the queue length; (ii) never inform customers about the

queue length; (iii) inform customers based on a threshold policy in which queue length information is provided when the queue length is below the specified threshold and is hidden otherwise. After finding the equilibrium structure of the third policy, we prove that although that policy seems intuitive, either sharing information with all customers or hiding information from all customers always yields greater expected revenue.

In summary, in this dissertation we:

1. Develop a stochastic game-theoretic framework that helps in reasoning about the strategic behavior of customers in loss queues and waiting queues that support advance reservations.
2. Analyze the impact of pricing on the decisions of customers and how those decisions impact the revenue of the service provider. The analysis includes both equilibrium outcomes and the outcomes of dynamic games that repeat many times.
3. Evaluate the impact of several information sharing policies on the decisions of strategic customers whether to make advance reservation or not in a loss queue and whether to join or not a waiting queue.

The rest of the dissertation is organized as follows. In Chapter 2, we provide relevant background on game theory and queueing theory and survey related work. In Chapter 3, we formally define AR games in loss queues and study two versions of the game with no information sharing. In Chapter 4, we study three versions of the game, each one with a different information sharing policy. In Chapter 5, we define and study AR games in waiting queues. In Chapter 6, we study dynamic versions of AR games. In Chapter 7, we study the impact of information on $M/M/1$ queues.

Finally, in Chapter 8, we conclude the dissertation and provide directions for future work.

Chapter 2

Background and Related Work

In this chapter, we first review concepts of game theory and queueing theory related to this dissertation. We then review related work on applying game theory to queues, and relevant literature in the areas of advance reservations and learning.

2.1 Game Theory

Game theory is the mathematical study of the interaction among independent, strategic and rational players (customers in this dissertation). It is rooted in the seminal book “Theory of games and economic behavior” (Von Neumann and Morgenstern, 1947) and has been applied to a wide range of fields including economics, political science, biology, psychology and engineering.

A game consists of players, the action space of these players, the outcomes of these actions and the payoffs which result from these outcomes. Players may be of different types and also have different information about the state of the system. Due to those differences, players may have different preferences for the same outcome of a game. A payoff is a numerical value that represents the utility of a player. Given player i , and two outcomes A and B , if the payoff function $U_i(\cdot)$ satisfies $U_i(A) < U_i(B)$, then player i prefers B over A .

The action space is the set of actions a player can take. Once each player chooses

an action (each player may have a different action set to choose from), the game has an outcome. Given an outcome, each player can compute its payoff, which reflects the utility of each player for that outcome.

A *pure strategy* is a mapping of any situation a player can face into an action from his action space. A *mixed strategy* is a mapping that assigns a probability to each action. A Nash equilibrium is a balanced state, in which none of the users has any incentive to deviate from his chosen strategy after observing the strategies chosen by the other users. Naturally, equilibrium strategies are those of interest. A Nash equilibrium may not be unique. Although other types of equilibria exist in the literature, throughout this dissertation, any equilibrium refers to a Nash equilibrium. An equilibrium is said to be a *symmetric equilibrium* if, under that equilibrium, all players follow the same strategy. Such a strategy maps the space of all possible situations players can face into an action (pure strategy) or a set of probabilities (mixed strategy).

Typically, the number of players in a game is known by all players. However, in many applications, uncertainty about the number of participants is a crucial element. For example, when a customer decides whether to make an advance reservation, he usually does not know how many other customers seek service at the same time. To model such situations, one can treat the number of players as a random variable. In such games, it is usually assumed that the players know the distribution of that random variable. A seminal work on games with a random number of players is performed in (Myerson, 1998) which introduces *Poisson games*. In that paper, Myerson shows that if and only if the number of players is Poisson distributed with parameter λ , then the number of other players, as seen by a player who does not count himself, is also Poisson distributed with the same parameter λ .

In this dissertation, we focus on *mechanism design*, which is an engineering ap-

proach to game theory. The goal is to design mechanisms and policies that optimize desired objectives, under the assumption that all players follow an equilibrium strategy. Mechanism design approach was first introduced in (Hurwicz, 1973) and is widely used in several fields such as auctions (Cramton et al., 2006), scheduling (Heydenreich et al., 2007) and communication networks (Srivastava et al., 2005).

2.2 Queueing Theory

Queueing theory is rooted in Erlang's seminal work (Erlang, 1909), which developed mathematical models to describe the Copenhagen telephone exchange. Queueing theory has applications in telecommunication, public transportation, call centers and hospitals. In this section, we briefly describe queueing models that are relevant to this dissertation.

A queueing system has several properties which can be described by Kendall's notation (Kendall, 1953) which is a common notation to classify queueing systems. With this notation, a queue is described by the four factors $A/S/c/K$, where A denotes the inter-arrival distribution, S denotes the service distribution, c denotes the number of servers and K is the capacity of the queue.

A queue with capacity equal to the number of servers is called a *loss queue*. In such a queue, customers that cannot be served at their desired service time leave the system. Throughout the dissertation, we assume that those customers do not make another trial (a common assumption made in the literature of loss queues for the sake of tractability (Ross, 1995)).

A queue whose size is greater than the number of servers is called a *waiting queue*. In such a queue, customers that cannot be served upon arrival wait for service (unless the queue is full). The capacity of the queue can be finite or infinite. In this

dissertation, we only consider waiting queues of infinite capacity.

We consider two waiting queues: $M/M/1$ and $M/D/1$. In the former queue, both the inter-arrival time and the service time follow independent Exponential distributions. The letter M stand for Markovian or memoryless, while the number 1 stands for one server. In the $M/D/1$ queue, the service rate is a constant (D stands for deterministic). We follow the standard convention of not specifying the queue capacity when it is infinite.

The dynamic evolution of many queueing systems can be modeled using *Markov processes*, which are stochastic processes with *Markovian property*. We say that a stochastic process is Markovian if the conditional probability distribution of future states of the process depends only on the present state, and not on past states or events.

2.3 Queueing Games

Queueing games is the study of a queueing system under the assumption that customers are strategic. That is, the decisions of customers are not only influenced by prices and policies set by providers but also by their beliefs about the decisions of other customers. Typically, in a queueing game, each customer needs to make one specific decision such as to join the queue or to balk, to buy priority or not, and so on. The main objective when studying a queueing game is to find how strategic customers respond to different system's parameters and, in turn, how their decisions impact the throughput of the system and/or the social welfare.

The application of game theory to analyze the strategic behavior of customers in queues was pioneered by (Naor, 1969). In that paper, the author considers an $M/M/1$ queue where customers observe the queue length and then decide whether to join or

balk. Follow-up works analyze the behavior of customers in different $M/M/1$ queueing systems. For example, (Edelson and Hilderbrand, 1975) analyze an unobservable $M/M/1$ queue, where customers decide whether to join or balk without knowing the queue length. (Balachandran, 1972) analyzes an observable $M/M/1$ queue with priorities, where customers decide on a payment and accordingly priorities are assigned.

Over the years, research on queueing games branched to other types of queueing systems. For example, (Haviv et al., 2010) analyze an unobservable $M/M/N/N$ (i.e., a loss queue with N servers) that is initially empty and customers decide whether to join or balk based on their arrival time. (Altman and Shimkin, 1998) analyze an observable processor sharing system, where customers decide whether or not to join after observing the number of customers in the system. (Jain et al., 2011) introduce concert queueing games, where N customers, interested in early service with minimal wait, choose their arrival time into a system with a specific opening time. (Haviv and Roughgarden, 2007) analyze an unobservable system with non-identical servers, where customers aim to minimize their waiting time and select a server accordingly.

In recent years, research on the impact of information in queueing games has emerged. (Hassin and Roet-Green, 2011) give examples of queueing systems providing on-line queue length information, including occupancy of emergency rooms in hospitals, voting locations, security gates at international airports, and amusement parks. (Shone et al., 2013) compare between the performance of observable and unobservable $M/M/1$ queues and find the conditions under which the joining rates of the observable queue and unobservable queue are equal. (Hassin, 2007) also studies the information impact on an $M/M/1$ queue but considers different kinds of information such as the current service rate or the current quality of service (it is assumed that both parameters may change overtime). The authors show that in some cases it is preferable to share information and in other cases it is preferable to hide it.

For a comprehensive review of queueing game, see (Hassin and Haviv, 2003) and (Hassin, 2016).

2.4 Advance Reservations

Queueing systems and communication networks that support advance reservations have extensively been researched. Most of the research focuses on performance evaluation and algorithmic aspects of AR systems. For example, (Smith et al., 2000) suggest a scheduling model that supports AR and evaluate several performance metrics. (Kaushik et al., 2006) suggest an AR model with flexible time window and show that this model has a lower blocking probability and a higher utilization than a model without window. (Guérin and Orda, 2000) analyze the effect of AR on the complexity of path selection. (Virtamo, 1992) evaluates the impact of advance reservations on server utilization. (Buyya et al., 2009) report a simulation-based comparison between different payment mechanisms. (Cohen et al., 2009) propose algorithms for network routing that support advance channel reservations. For a survey of the field, see (Charbonneau and Vokkarane, 2012).

Advance reservation models have been also studied in the field of revenue management. One of the first works in this field was conducted by (Lieberman and Yechiali, 1978) which analyze a hotel reservation system where overbooking is allowed and the goal is to find the optimal overbooking level. (Reiman and Wang, 2008) propose an admission control strategy for a reservation system with different classes of customers. (Bertsimas and Shioda, 2003) propose a policy for accepting/rejecting restaurant reservations. (Nasiry and Popescu, 2012) assume that customers have uncertain valuations and they need to decide whether to purchase in advance or not before their valuation is realized. The provider, aiming to maximize the revenue, can

set different prices for advance sell and spot sell. None of this prior work considers the strategic behavior of customers, namely, that decisions of customers are not only influenced by prices and policies set by providers but also by their beliefs about the decisions of other customers.

2.5 Learning Models

The concept of learning an equilibrium is rooted in Cournot's duopoly model (Cournot, 1897) and has been extensively researched since. In Cournot's model there are N firms and each one simultaneously sets a quantity to produce. The market price is set at a level such that demand equals the total quantity produced by all firms. The process repeats many times. At each step, all firms observe the quantities produced by other firms, and then each one sets a new quantity, naively assuming that all other firms will not change their quantity. Cournot shows that the game converges to an equilibrium. Furthermore, he shows that as the number of firms increases the price decreases and eventually converges to the cost of production (i.e., the profit converges to zero).

Traditionally, learning models are used for fixed-player games (the same players participate at each iteration) with a static environment, see, for instance, (Lakshmi-varahan, 1981; Fudenberg, 1998; Milgrom and Roberts, 1991). In recent years, more papers have focused on learning under stochastic settings. For example, in (Liu and van Ryzin, 2011) customers choose between buying a product at full price or waiting for a discount period. Decisions are made based on observing past capacities. (Altman and Shimkin, 1998) analyze a processor sharing model. In this model, customers choose between joining or balking after observing the performance history. (Zohar et al., 2002) present a model of abandonment from unobservable queues. The decision is based on the expected waiting time which is formed through accumulated experi-

ence. (Fu and van der Schaar, 2009) assume that the same set of players participate in a bid for wireless resources at each stage. However, the number of packets that need to be transmitted at each iteration is a random variable.

Different learning models differ by their learning rule. A learning rule defines what kind of information players gain and how they use it. In this dissertation, we focus on *best response dynamics*. According to this rule, which is rooted in Cournot's work, players observe the most recent actions and assume that those actions will be repeated in the next step. Another popular learning rule is *fictitious play* which assumes that, at each iteration, players observe the actions made in all previous steps and they best-respond to the empirical frequency of observed actions. This rule was suggested by (Brown, 1951). In contrast, (Littman, 1994) and (Tan, 1993) assume that players only observe their own payoffs and learn by trail-and-error. This learning rule is known as *reinforcement learning*.

We conclude that both advance reservations and queueing games have been studied extensively in the past twenty years. However, to our knowledge, applying game-theoretic tools to queueing systems that support advance reservations, was not done before, except for (Oh and Su, 2012). In that paper, the authors suggest policies to mitigate the problem of no-show in restaurants: to punish no-shows by charging fees and to encourage show-ups by giving discounts.. In this dissertation, we focus on determining optimal policies that maximize the number of customers making reservations and the revenue from reservation fees in a game-theoretic setting.

Chapter 3

Advance Reservation Games

In this chapter, we consider a slotted loss queue that allows customers to reserve resources in advance. We assume that the provider does not share any information about availability of servers. We first provide a detailed description of the model under consideration. Next, we find the equilibrium structure under the assumption that AR cost is applied only to served customers and when the AR cost is applied to all customers that make AR. Then, we compare between the outcomes of those two versions, assuming that the AR cost is a fee charged by the provider. Finally, we aim to determine which fee maximizes the revenue from AR fees.

3.1 Model Description

First, we describe the assumptions that are common to both models:

1. The system consists of N servers.
2. The service time axis is slotted. That is, in each slot, customers are served from the beginning till the end of the slot.
3. The *demand*, which represents the number of customers that request service in a specific slot (each customer requests one server) is a random variable, denoted by D . Customers that do not get service in a given slot do not make

another trial (a common assumption in the literature of loss queues (Ross, 1995)). Thus, the demand in each slot is independent of the history and follows a *general* probability distribution P_D , supported in $[a, b]$, where $0 \leq a < N$ and $N < b \leq \infty$ (i.e, the demand has a positive probability of being smaller, equal, or larger than the number of servers). The distribution P_D is public information.

4. The customers of each slot differ by the time elapsing between considering making AR (i.e, realizing that service will be needed in that future slot) and the slot starting time. We refer to this time interval as the *arrival time*. The arrival times of all customers are independent and identically distributed random variables, supported in \mathbb{R}^+ , with cumulative distribution function (CDF) denoted $F(\cdot)$.
5. Each customer chooses one of two actions: make AR or not make AR, denoted AR and AR' respectively.
6. If the demand for a slot is larger than N , the servers are allocated to the first N customers that made AR. If fewer than N customers made AR, the remaining servers are arbitrarily allocated between the customers that did not make AR.
7. The customers and the provider know the number of servers N and statistical information on the system (i.e., the distribution of the demand and the arrival times).
8. Making reservation is associated with a fixed reservation cost C . All the customers have the same utility U from the service. Without loss of generality, we set $U = 1$.

Figure 3.1 illustrates the model




Figure 3.1: Slotted system with two servers. The demand for each slot is independent. The demand for slots 1, 2 and 3, is respectively 3, 2 and 4.

-	Make AR		Not make AR	
Model	Served	Not served	Served	Not served
1	$1 - C$	0	1	0
2	$1 - C$	$-C$	1	0

Table 3.1: Payoff summary.

The two models analyzed in this chapter differ in their reservation mechanisms as follows:

1. **No-information model 1:** customers have no information regarding the availability of servers at the time of reservation. If a customer makes an AR request, he is then informed whether a server will be allocated at the requested slot or not. In the first case, the reservation cost is applied. In the second case, the customer leaves the system with no gain or cost.
2. **No-information model 2:** as in the first model, customers have no information regarding the availability of servers at the time of reservation. In this model, however, the reservation cost is applied on each customer that makes an AR request.

The possible payoffs of the two models are summarized in Table 3.1.

3.2 Equilibria Analysis

3.2.1 Classification of the equilibria

We analyze the two models as non-cooperative games where each player (customer) aims to maximize his payoff. Since the demand for the different slots are identical and independent (i.i.d) random variables, the analysis of a single slot is sufficient for analyzing the game. Any cost greater or equal to one has a trivial result where none of the customers makes AR. Zero cost or a negative cost have the trivial result of all customers making AR. Hence, in our analysis, we consider only costs between zero and one (i.e., $0 < C < 1$).

First, we note that the demand seen by a customer may be different from that seen by an external observer (the provider, for instance). Indeed, the fact that a customer seeks service in a given time slot affects his estimation of the number of other customers seeking service in that time slot. On the one hand, a customer is more likely to fall in a slot with large demand than in one with small demand. On the other hand, he must exclude himself. This phenomenon is known as the discrete case of the *waiting time paradox* (or *residual life paradox*). We define \tilde{D} as the number of customers seen by a customer beside himself. The probability distribution function (PDF) of \tilde{D} is known to be (Avineri, 2004):

$$\mathbb{P}(\tilde{D} = j) = \mathbb{P}(D = j + 1) \frac{(j + 1)}{\mathbb{E}[D]}. \quad (3.1)$$

The following lemma states that each customer makes a decision at his arrival time.

Lemma 1. *For all customers, making AR at the arrival time yields at least the same payoff as making AR later on.*




Figure 3.2: Example of a realization of the demand in a given slot. The service starts at arrival time 1. As customers have smaller (normalized) arrival times, they have the opportunity to reserve a server earlier.

Proof. In both models, if a customer makes AR when all servers are already reserved, his payoff will be the same as if making AR later on (0 in the first model and $-C$ in the second model). If a customer makes AR when there is at least one free server, his payoff will be $1 - C$. If he makes a reservation at a later point, his payoff will be the same if there is still at least one free server. If there is no more server available, his payoff will be zero in the first model and negative in the second model. \square

Given an arrival time γ , we set $t \triangleq F(\gamma)$ and refer to it as the *normalized arrival time*. Due to the probability integral transformation theorem (Dodge, 2006, p. 320), we know that the normalized arrival time is a random variable, uniformly distributed in $[0, 1]$, where 1 is the service start time. Note that $F(\gamma)$ is also the average fraction of customers with arrival time smaller than γ . Since we are only interested in the normalized arrival time, for brevity, we henceforth refer to it as arrival time. Fig. 3.2 illustrates the new notion of arrival time.

For each game, we define a strategy function $\sigma: t \rightarrow [0, 1]$, which represents the probability that a tagged customer with arrival time $t \in [0, 1]$ makes AR. Since we consider only symmetric equilibria, all customers follow the same strategy function. Through conditioning, given that there are k other customers with arrival times t_1, \dots, t_k that follow strategy σ , the tagged customer can find his probability of getting

service (we denote that event by S) for each action he chooses. His probability to get service, when choosing action $\mathcal{A} \in \{AR, AR'\}$ is

$$\mathbb{P}(S|t, \mathcal{A}, \sigma) = \mathbb{P}(\tilde{D} < N) + \sum_{k=N}^{\infty} \mathbb{P}(\tilde{D} = k) \int_{t_1=0}^1 \cdots \int_{t_k=0}^1 \mathbb{P}(S|t, \mathcal{A}, k, t_1, \dots, t_k, \sigma) dt_1 \cdots dt_k. \quad (3.2)$$

The first term in (3.2) is the probability that the number of customers (beside the tagged customer) is smaller than N . In this case, all customers get service, regardless of their decisions. The second term is the weighted sum of the probabilities of getting service when the number of customers (beside the tagged customer) is at least N . In this case, the probability that the tagged customer gets service depends on his action and on the strategy followed by the other customers and their arrival times (note that the PDF of the random variable t_j equals 1, for each $j \in \{1, \dots, k\}$). As shown in the sequel, deriving an explicit expression for $\mathbb{P}(S|t, \mathcal{A}, k, t_1, \dots, t_k, \sigma)$ is not required for the equilibria analysis.

Given the model and strategy function followed by all other customers, one can express the expected payoff, denoted $U_\sigma(t, \mathcal{A})$, for each action \mathcal{A} by multiplying $\mathbb{P}(S|t, \mathcal{A}, \sigma)$ and $1 - \mathbb{P}(S|t, \mathcal{A}, \sigma)$ with the relevant payoffs, as summarized in Table 3.1. In the first model, the expected payoffs are:

$$U_\sigma(t, AR) = \mathbb{P}(S|t, AR, \sigma) \cdot (1 - C) + (1 - \mathbb{P}(S|t, AR, \sigma)) \cdot 0 \quad (3.3)$$

and

$$U_\sigma(t, AR') = \mathbb{P}(S|t, AR', \sigma) \cdot 1 + (1 - \mathbb{P}(S|t, AR', \sigma)) \cdot 0. \quad (3.4)$$

In the second model, the expected payoffs are:

$$U_\sigma(t, AR) = \mathbb{P}(S|t, AR, \sigma) (1 - C) + (1 - \mathbb{P}(S|t, AR, \sigma)) (-C) \quad (3.5)$$

and

$$U_\sigma(t, AR') = \mathbb{P}(S|t, AR', \sigma) \cdot 1 + (1 - \mathbb{P}(S|t, AR', \sigma)) \cdot 0. \quad (3.6)$$

At equilibrium, each customer chooses an action that maximizes his expected payoff. Thus, we define an equilibrium strategy (i.e., a strategy that leads to equilibrium) as follows.

Definition 1. *Strategy σ is an equilibrium strategy if the following holds for any arrival time $t \in [0, 1]$:*

1. *If $\sigma(t) = 0$ then $U_\sigma(t, AR) \leq U_\sigma(t, AR')$.*
2. *If $0 < \sigma(t) < 1$ then $U_\sigma(t, AR) = U_\sigma(t, AR')$.*
3. *If $\sigma(t) = 1$ then $U_\sigma(t, AR) \geq U_\sigma(t, AR')$.*

That is, at equilibrium, a customer chooses the action AR' , only if he is (weakly) better off not making AR; he randomizes his action, only if he is indifferent between the two outcomes; and he chooses the action AR , only if he is (weakly) better off making AR.

Next, we show that at equilibrium all customers follow the same *threshold strategy*, defined below.

Definition 2. *A threshold strategy has the following structure:*

$$\sigma(t) = \begin{cases} 0 & \text{if } t > \tau \\ 1 & \text{if } t \leq \tau. \end{cases}$$

where τ is a threshold value in the interval $[0, 1]$.

Theoretically, the game could have infinite number of equilibria. For example, suppose that $\sigma(t) = 0$ and $(1 - C)\mathbb{P}(S|t, AR, \sigma) = \mathbb{P}(S|t, AR', \sigma)$, for all $0 \leq t \leq 1$

(i.e., all customers choose AR' and they are indifferent between the two actions). In this case, σ is an equilibrium strategy. But

$$\sigma(t) = \begin{cases} 1 & \text{if } t = 0.5, \\ 0 & \text{otherwise,} \end{cases}$$

is also an equilibrium strategy, since the probability that a customer will arrive exactly at 0.5 is zero. In practice, in both equilibria all customers choose AR' . In our analysis, we wish to ignore such degenerate cases. For this purpose, we assume the following:

Assumption 1. *At equilibrium, given an arrival time $t = t_1$, if $\sigma(t_1) = x$, then there is a non-zero measure interval T_1 such that $t_1 \in T_1$ and $\sigma(t) = x$ for any $t \in T_1$.*

Lemma 2. *In the no-information models, at equilibrium, all customers follow a threshold strategy.*

Proof. Consider a tagged customer with arrival time t and assume that the rest of the customers follow a strategy function σ . Since information is not shared with the customers, all customers that do not make AR have the same estimate of their probability of service. More formally, let \tilde{D}_{AR} be the number of reservations. Then, the probability of service if not making AR is

$$\mathbb{P}(S|\tilde{D}=\tilde{d}, \tilde{D}_{AR}=\tilde{d}_{AR}) = \begin{cases} 1 & \text{if } \tilde{d} < N, \\ \frac{N-\tilde{d}_{AR}}{\tilde{d}+1-\tilde{d}_{AR}} & \text{if } \tilde{d} \geq N \text{ and } \tilde{d}_{AR} < N, \\ 0 & \text{otherwise.} \end{cases} \quad (3.7)$$

Thus, $U_\sigma(t, AR')$ does not depend on t . For brevity, we denote the expected payoff of not making AR by β . From Lemma 1, we know that the expected payoff when making AR $U_\sigma(t, AR)$ is a non-increasing function of t . Hence, the two expected payoffs can intersect at most once.

If $U_\sigma(t, AR) < \beta$ for all $t \in [0, 1]$, then σ is an equilibrium strategy only if none of

the customers makes AR (i.e., it is a threshold strategy with $\tau = 0$). If $U_\sigma(t, AR) > \beta$ for all $t \in [0, 1]$, then σ is an equilibrium strategy only if all customers make AR (i.e., it is a threshold strategy with $\tau = 1$). However, if all customers make AR, a customer with arrival time 1^- has the same probability to get service with and without AR. Thus, he is better off not making AR. Therefore, an equilibrium where all customers make AR cannot exist.

Finally, if the two expected payoff functions intersect, they can either intersect at a single point t_0 or along an interval $[t_1, t_2]$. In the first case, $U_\sigma(t, AR) > \beta$ for all $t < t_0$ and $U_\sigma(t, AR) < \beta$ for all $t > t_0$. Thus, in this case, σ is an equilibrium strategy only if it is a threshold strategy with $\tau = t_0$.

In the second case, $U_\sigma(t, AR)$ has the same value for all $t \in [t_1, t_2]$, which can only happen if $\sigma(t) = 0, \forall t \in [t_1, t_2]$ (by Assumption 1, we ignore the case of $\sigma(t) \neq 0$ over a measure zero subset of $[t_1, t_2]$). Since none of the customers make AR in the interval $[t_1, t_2]$, and since $U_\sigma(t, AR) > \beta$ for all $t < t_1$ and $U_\sigma(t, AR) < \beta$ for all $t > t_2$, we conclude that σ is an equilibrium strategy only if it is a threshold strategy with $\tau = t_1$. \square

After showing that a threshold strategy is the only possible equilibrium strategy, we distinguish between two types of equilibria.

Definition 3. *None-make-AR is an equilibrium in which all customers follow a threshold strategy with threshold $\tau_e = 0$.*

Definition 4. *Some-make-AR is an equilibrium in which all customers follow a threshold strategy with threshold $\tau_e \in (0, 1)$.*

Using the results obtained so far, we find next the equilibria structure for each model separately.

3.2.2 Equilibria structure

In this section, we show that different ranges of costs lead to different equilibria. The following theorem summarizes the main results.

Theorem 1. *For each model $i = 1, 2$, there exist quantities \underline{C} and $\overline{C}_i \geq \underline{C}$, such that:*

- *If $0 < C < \underline{C}$, there is at least one some-make-AR equilibrium.*
- *If $\underline{C} < C < \overline{C}_i$, there is a none-make-AR equilibrium and at least two some-make-AR equilibria.*
- *If $C > \overline{C}_i$, none-make-AR is the unique equilibrium.*

For simplicity, we do not consider the boundary cases $C = \underline{C}$ and $C = \overline{C}_i$ in our discussion.

No-information model 1

We consider the first *no-information* model. For each type of equilibria, we determine the range of costs in which they may occur.

Some-make-AR equilibria. If all customers follow a strategy with threshold τ_e , that strategy is an equilibrium strategy if and only if a customer with arrival time τ_e (referred to as a *threshold customer*) is indifferent between the actions AR and AR' . We denote $\pi_{AR}(\tau_e)$ the probability that a threshold customer gets service upon chosen action AR , and $\pi_{AR'}(\tau_e)$ the probability that a threshold customer gets service upon chosen action AR' . Hence, a strategy with threshold τ_e is an equilibrium if and only if

$$(1 - C) \pi_{AR}(\tau_e) = \pi_{AR'}(\tau_e), \quad (3.8)$$

where the left hand side of Eq. (3.8) is the expected payoff of AR and the right hand side is the expected payoff of AR' . Using Eq. (3.8), we express the cost as a function

of the threshold

$$C_1(\tau_e) \triangleq 1 - \frac{\pi_{AR'}(\tau_e)}{\pi_{AR}(\tau_e)}. \quad (3.9)$$

Next, we develop the expressions $\pi_{AR}(\tau_e)$ and $\pi_{AR'}(\tau_e)$. The former expression corresponds to the probability that either the demand is at most N or the demand exceeds N but fewer than N customers make AR. The number of customers making AR, given $\tilde{D} = j$ with $j \geq N$, is a random variable that follows a binomial distribution. The number of trials is j and the success probability is $1 - \tau_e$. The probability that the threshold customer gets service is equal to the probability that the number of successes is at most $N - 1$. By summing this probability over all possible values of j , we get:

$$\pi_{AR}(\tau_e) = \mathbb{P}(\tilde{D} < N) + \sum_{j=N}^{\infty} \sum_{i=0}^{N-1} \mathbb{P}(\tilde{D} = j) \tau_e^i (1 - \tau_e)^{j-i} \binom{j}{i}. \quad (3.10)$$

Likewise, we have

$$\pi_{AR'}(\tau_e) = \mathbb{P}(\tilde{D} < N) + \sum_{j=N}^{\infty} \sum_{i=0}^{N-1} \mathbb{P}(\tilde{D} = j) \tau_e^i (1 - \tau_e)^{j-i} \binom{j}{i} \frac{N - i}{j + 1 - i}. \quad (3.11)$$

In that case, if the demand exceeds N but fewer than N customers make AR, service is not guaranteed. Given a demand j and a number of reservations i , the probability to get service without AR is the ratio of the number of unreserved servers $N - i$ to the number of customers that did not make AR $j + 1 - i$.

Next, we prove that these two functions are continuous.

Lemma 3. *The functions $\pi_{AR'}$ and π_{AR} are continuous functions of τ_e in the range $[0, 1]$.*

Proof. Starting with Eq. (3.10) and ignoring the first term of the function which does

not depend on τ_e , we need to show that the second term is continuous. The inner sum of the second term is continuous, since it is a finite sum of polynomial functions. To prove that the outer sum is continuous, we use Cauchy's uniform convergence criterion (Trench, 2003). We shall show that for any $\epsilon > 0$ there exists an integer M such that

$$\sup_{0 \leq \tau_e \leq 1} \sum_{j=n}^m \sum_{i=0}^{N-1} \mathbb{P}(\tilde{D} = j) \tau_e^i (1 - \tau_e)^{j-i} \binom{j}{i} < \epsilon \quad \forall n, m \geq M. \quad (3.12)$$

The above expression is upper bounded by $\mathbb{P}(n \leq \tilde{D} \leq m)$ which, in turn, is upper bounded by $\mathbb{P}(\tilde{D} \geq n)$. For any discrete distribution and $\epsilon > 0$ there exists M such that $\mathbb{P}(\tilde{D} \geq n) < \epsilon$ for any $n > M$. Thus, we have shown that Eq. (3.12) holds true. Since, for any $\tau_e \in [0, 1]$, $\pi_{AR}(\tau_e) \geq \pi_{AR'}(\tau_e)$, the proof is also valid for $\pi_{AR'}$. \square

Since both $\pi_{AR'}$ and π_{AR} are continuous and positive in the range $\tau_e \in [0, 1]$, we deduce that $C_1(\tau_e)$ is a continuous function in this range. Next, we observe that if all customers make AR, then the probability of service of a customer with arrival time one (i.e., the last arriving customer) does not depend on his decision. Hence, $C_1(1) = 0$. In any other case, the probability to get service is greater when making AR. Hence, $C_1(\tau_e) > 0$ for any $0 \leq \tau_e < 1$. We denote the supremum value of $C_1(\tau_e)$ as

$$\bar{C}_1 \triangleq \sup_{0 < \tau_e < 1} C_1(\tau_e). \quad (3.13)$$

Since the equation $C_1(\tau_e) = C$ has a solution if and only if C is smaller than the supremum value of $C_1(\tau_e)$, we conclude that a *some-make-AR* equilibrium exists if $C < \bar{C}_1$ and does not exist if $C > \bar{C}_1$.

None-make-AR equilibrium. If none of the customers makes AR, all have the same expected payoff $\pi_{AR'}(0)$. A customer that deviates gets service with probability




Figure 3.3: An example with 10 servers, Poisson distributed demand with mean 15 and AR cost $C = 0.45$. The line C and the function $C(\tau_e)$ intersect twice. Hence, there are *two-some-make-AR* equilibria. Since $C(1) < C$, there is also a *none-make-AR* equilibrium.

$\pi_{AR}(0) = 1$ and his payoff is $1 - C$. Thus, if $1 - C < \pi_{AR'}(0)$, then none of the customers will have an incentive to deviate. On the other hand, if $1 - C > \pi_{AR'}(0)$, then all the customers will have an incentive to deviate. By defining $\underline{C} \triangleq C_1(0)$, we conclude that if $C > \underline{C}$, then a *none-make-AR* equilibrium exists. If $C < \underline{C}$, then a *none-make-AR* equilibrium does not exist.

By definition $\overline{C}_1 \geq \underline{C}$. Therefore, we have shown that for any value of $0 < C < 1$, at least one equilibrium exists. Furthermore, if the interval $I = (\underline{C}, \overline{C}_1)$ is not empty (i.e., the supremum point is not reached at $\tau_e = 1$), then for any $C \in I$, the equation $C = C_1(\tau_e)$ must have at least two solutions due to the continuity of the function. Therefore, any cost $C \in I$ has at least two different *some-make-AR* equilibria (the exact number of *some-make-AR* equilibria depends on the number of maximal points of the function $C(\tau_e)$). See Figure 3.3 for an illustration.

No-information model 2

In this section, we show that the second game has the same equilibria structure as the first one, but with different ranges.

Some-make-AR equilibria. If all the customers follow a strategy with threshold τ_e , the probability to get service with or without making AR is calculated in the same way as in the previous model. Thus, the functions π_{AR} and $\pi_{AR'}$ can be also used in the analysis of this model. As in the first game, at a *some-make-AR* equilibrium, the threshold customer is indifferent between the two actions AR and AR' . Thus,

$$\pi_{AR}(\tau_e) - C = \pi_{AR'}(\tau_e), \quad (3.14)$$

where the left hand side of the equation is the expected payoff of AR , while the right hand side is the expected payoff AR' . In this model, the cost as a function of the threshold is:

$$C_2(\tau_e) \triangleq \pi_{AR}(\tau_e) - \pi_{AR'}(\tau_e). \quad (3.15)$$

We define

$$\bar{C}_2 \triangleq \sup_{0 < \tau_e < 1} C_2(\tau_e). \quad (3.16)$$

As is the first model, $C_2(1) = 0$ and $C_2(\tau_e) > 0$ for any $\tau_e < 1$. Thus, a *some-make-AR* equilibrium exists if $C < \bar{C}_2$ and does not exist if $C > \bar{C}_2$.

None-make-AR equilibrium. If none of the customers makes AR, then the expected payoffs of not making AR and the expected payoff of deviating are the same as in the first model. Therefore, the range of costs that have a *none-make-AR* equilibrium is the same as in the first model.

In conclusion, the difference between the analyses of the two games is that \bar{C}_1




Figure 3.4: The cost function in the two *no-information* models with $N = 10$. The demand is a Poisson random variable with mean equals 10.

may be different from \bar{C}_2 . From Eq. (3.9) and (3.15), we obtain

$$\frac{C_1(\tau_e)}{C_2(\tau_e)} = \frac{1}{\pi_{AR}(\tau_e)} \quad \forall \tau_e \in (0, 1). \quad (3.17)$$

In any *some-make-AR* equilibrium, the probability to get service is smaller than one. Hence, for any $\tau_e \in (0, 1)$, $\bar{C}_1 > \bar{C}_2$ (as illustrated in Figure 3.4).

3.3 Revenue Maximization

In this section, we assume that the AR cost is a fee charged by the provider. She can either charge the reservation fee only from customer granted service (this policy is aligned with the first model discussed in the previous section) or to charge the AR fee from any customer attempting to make AR (this policy is aligned with the second model).

Intuitively, there is a trade-off between those two policies. Given an equilibrium with a certain expected fraction of customers making AR, the total expected number of customer paying the fee is greater when applying the second policy. However, Eq. (3.17) indicates that if the provider is aiming to achieve a certain fraction of customers making AR, she will have to advertise a lower fee if using the second policy. Our objective, in this section, is to determine which policy maximizes the provider revenue from AR fees.

We define the maximum possible revenue for model $i = 1, 2$ as

$$R_i^* = \sup_{0 < \tau_e < 1} R_i(\tau_e). \quad (3.18)$$

Under *some-make-AR* equilibrium with threshold τ_e , the number of reservations is $D_{AR}(\tau_e)$ or simply D_{AR} from now and on. In the first model, the expected revenue per server is the expected number of reserved servers, multiplied by the fee and normalized by the number of servers N :

$$R_1(\tau_e) = \frac{\mathbb{E}[\min(D_{AR}, N)]C_1(\tau_e)}{N}. \quad (3.19)$$

In the second model, it is the expected number of reservations, multiplied by the fee and normalized by N :

$$R_2(\tau_e) = \frac{\mathbb{E}[D]\tau_e C_2(\tau_e)}{N}. \quad (3.20)$$

By comparing these two functions, we state the following result:

Theorem 2. *In AR games, the maximum possible revenue, at equilibrium, is greater when charging the AR fee only from customers that get service and not from all customers that make AR requests.*

Proof. We prove the theorem by showing that for any given threshold, the first model

yields greater revenue than the second. Namely, we show that for any value of $\tau_e \in (0, 1)$ the following holds:

$$R_1(\tau_e) > R_2(\tau_e). \quad (3.21)$$

From Eqs. (3.19), (3.20) and (3.17), we obtain that showing that $R_1(\tau_e) > R_2(\tau_e)$ is equivalent to showing that

$$\mathbb{E}[\min(D_{AR}, N)] - \mathbb{E}[D] \cdot \tau_e \cdot \pi_{AR}(\tau_e) > 0. \quad (3.22)$$

First, we expand the first term of (3.22):

$$\mathbb{E}[\min(D_{AR}, N)] = \sum_{i=0}^N \mathbb{P}(D_{AR} = i) i + \sum_{i=N+1}^{\infty} \mathbb{P}(D_{AR} = i) N. \quad (3.23)$$

The PDF of D_{AR} is

$$\mathbb{P}(D_{AR} = i) = \sum_{j=i}^{\infty} \mathbb{P}(D = j) \tau_e^i (1 - \tau_e)^{j-i} \binom{j}{i}. \quad (3.24)$$

By combining Eq. (3.23) and (3.24), we get

$$\begin{aligned} \mathbb{E}[\min(D_{AR}, N)] &= \sum_{i=1}^N \sum_{j=i}^{\infty} \mathbb{P}(D = j) \tau_e^i (1 - \tau_e)^{j-i} \binom{j}{i} i \\ &\quad + \sum_{i=N+1}^{\infty} \sum_{j=i}^{\infty} \mathbb{P}(D = j) \tau_e^i (1 - \tau_e)^{j-i} \binom{j}{i} N. \end{aligned} \quad (3.25)$$

Next, we expand $\pi_{AR}(\tau_e)$:

$$\begin{aligned}
\pi_{AR}(\tau_e) &= \mathbb{P}(\tilde{D} < N) + \sum_{i=0}^{N-1} \sum_{j=N}^{\infty} \mathbb{P}(\tilde{D} = j) \tau_e^i (1-\tau_e)^{j-i} \binom{j}{i} \\
&= \sum_{i=0}^{N-1} \sum_{j=i}^{\infty} \mathbb{P}(D = j+1) \frac{(j+1)}{\mathbb{E}[D]} \tau_e^i (1-\tau_e)^{j-i} \binom{j}{i} \\
&= \sum_{i=0}^{N-1} \sum_{j=i+1}^{\infty} \mathbb{P}(D = j) \frac{j}{\mathbb{E}[D]} \tau_e^i (1-\tau_e)^{j-i-1} \binom{j-1}{i} \\
&= \sum_{i=1}^N \sum_{j=i}^{\infty} \mathbb{P}(D = j) \frac{j}{\mathbb{E}[D]} \tau_e^{i-1} (1-\tau_e)^{j-i} \binom{j-1}{i-1} \\
&= \sum_{i=1}^N \sum_{j=i}^{\infty} \mathbb{P}(D = j) \frac{j}{\mathbb{E}[D]} \tau_e^{i-1} (1-\tau_e)^{j-i} \binom{j}{i} \frac{i}{j}. \tag{3.26}
\end{aligned}$$

The explanation for Eq. (3.26) is as follows. We start from Eq. (3.10). We merge the two terms in Eq. (3.10) and substitute $P(\tilde{D} = j)$ by the right hand side of Eq. (3.1). Next, we replace j by $j-1$ and start the sum at $j=i$ instead of $j=i+1$. Next, we do a similar change with the variable i . Finally, we multiple and divide the expression by $\binom{j}{i}$.

In the next step, we multiply both sides of Eq. (3.26) by $\mathbb{E}[D] \cdot \tau_e$:

$$\mathbb{E}[D] \cdot \tau_e \cdot \pi_{AR}(\tau_e) = \sum_{i=1}^N \sum_{j=i}^{\infty} \mathbb{P}(D = j) \tau_e^i (1-\tau_e)^{j-i} \binom{j}{i} i. \tag{3.27}$$

Finally, we substitute the first term of the left hand side of Eq. (3.22) with the right hand side of Eq. (3.25) and the second term of the left hand side of Eq. (3.22) with the right hand side of Eq. (3.27). We then get

$$\mathbb{E}[\min(D_{AR}, N)] - \mathbb{E}[D] \cdot \tau_e \cdot \pi_{AR}(\tau_e) = \sum_{i=N+1}^{\infty} \sum_{j=i}^{\infty} \mathbb{P}(D = j) \tau_e^i (1-\tau_e)^{j-i} \binom{j}{i} N > 0, \tag{3.28}$$

which completes the proof. \square

The result of Theorem 2 is reassuring, since a mechanism that charges reservation fees only from customers getting service (first model) appears as more fair than a mechanism that charges reservation fees from all customers making AR requests (second model). Further, one may argue that the demand for service would decrease under the latter charging scheme. The theorem proves that the latter scheme is detrimental for the provider even if the demand for service were not reduced.

3.4 Price of Conservatism

In the previous section, we showed that in order to maximize the revenue, the provider should choose the first policy, which yields higher revenue for any given fee. In this section, we assume that the first policy is chosen and investigate different fees and their impact on the revenue.

By means of example, we next show that the fee that maximizes the revenue may yield more than one equilibrium, where one of them yields no revenue.

Example 1. *Consider a system with 15 servers and a Poisson distributed demand with parameter (mean) $\lambda = 20$. In this case, the maximum revenue per resource is $R_1^* = 0.41$ and it is achieved with fee $C_1^* = 0.47$. Since $\underline{C} = 0.26$, if charging C_1^* , then none-make-AR is also an equilibrium. Hence, charging C_1^* may yield the maximum possible revenue but may also yield no revenue. The revenue and fee functions are illustrated in Figure 3.5.*

If the fee that yields the maximum possible revenue is not unique, the provider may prefer a fee with smaller but guaranteed revenue. In order to weigh the different options, we propose the metric of *price of conservatism* (PoC). In the rest of this section, we formally define the term PoC and derive it for different settings. Since we only deal with the first model, the model index is removed in this section.




Figure 3.5: An example with $N = 15$ servers and Poisson distributed demand with parameter $\lambda = 20$. To see that the fee that maximizes the revenue belongs to a range with multiple equilibria, including *none-make-AR*.

In order to have a positive guaranteed revenue, the provider must choose a fee smaller than \underline{C} . Furthermore, if that fee has more than one equilibrium, then the guaranteed revenue is defined as the minimum between the revenues of the different *some-make-AR* equilibria. We define Z_C as the set of *some-make-AR* equilibria of the fee C , namely, $Z_C = \{\tau_e : C(\tau_e) = C, 0 < \tau_e < 1\}$. The maximum expected guaranteed revenue is defined as follows:

$$R_g^* = \sup_{0 < C < \underline{C}} \left(\inf_{\tau_e \in Z_C} R(\tau_e) \right). \quad (3.29)$$

The following definition captures the potential revenue loss resulting from a conservative pricing decision.

Definition 5. *The price of conservatism (PoC) is the ratio of the expected maximum possible revenue R^* to the expected maximum guaranteed revenue R_g^* .*

Next, we evaluate the provider's revenue and PoC under the assumption that the demand D is a Poisson random variable with parameter λ . We denote the number of

customers not making AR by $D_{AR'}$. Due to the properties of Poisson games (Myerson, 1998), D_{AR} and $D_{AR'}$ are independent Poisson random variables with parameter $\lambda\tau_e$ and $\lambda(1 - \tau_e)$, respectively. Furthermore, the total number of customers and the number of customers making each action, as seen by a customer if not counting himself, has the same distributions as D , D_{AR} and $D_{AR'}$ respectively.

3.4.1 Single-server Case

We start with the special case $N = 1$. If all customers follow a strategy with threshold τ_e , the probability that the threshold customer will get service is:

1. If making AR:

$$\pi_{AR}(\tau_e) = e^{-\lambda\tau_e}, \quad (3.30)$$

which is the probability that beside the customer that arrives at the threshold, no one makes AR (i.e., all the customers arrive after the threshold point).

2. If not making AR:

$$\pi_{AR'}(\tau_e) = e^{-\lambda\tau_e} \sum_{i=0}^{\infty} \frac{e^{-\lambda(1-\tau_e)} (\lambda(1-\tau_e))^i}{i!} \frac{1}{i+1} = \frac{e^{-\lambda} (-1 + e^{\lambda(1-\tau_e)})}{\lambda(1-\tau_e)}, \quad (3.31)$$

which is the probability that none of the customers makes AR, multiplied by the probability to get service given that none of the customers makes AR.

By substituting Eqs. (3.30) and (3.31) in Eq. (3.9), we get

$$C(\tau_e) = \frac{e^{-\lambda(1-\tau_e)} + \lambda(1-\tau_e) - 1}{\lambda(1-\tau_e)}. \quad (3.32)$$

Lemma 4. *For the case $N = 1$, $C(\tau_e)$ is a monotonically decreasing function.*

Proof. The derivative of $C(\tau_e)$ is:

$$\frac{dC}{d\tau_e} = \frac{e^{-\lambda(1-\tau_e)} (-e^{\lambda(1-\tau_e)} + 1 + \lambda(1-\tau_e))}{\lambda(1-\tau_e)^2}. \quad (3.33)$$

Since $\lambda(1-\tau_e)^2 \geq 0$, $e^{-\lambda(1-\tau_e)} \geq 0$ and $-e^{\lambda(1-\tau_e)} + 1 + \lambda(1-\tau_e) < 0$ for any $\lambda > 0$ and $0 < \tau_e < 1$, we conclude that $\frac{dC}{d\tau_e} < 0$. \square

From the lemma, we infer that $\bar{C} = C(0)$. Thus, by definition,

$$\bar{C} = \underline{C} = \frac{e^{-\lambda} + \lambda - 1}{\lambda}. \quad (3.34)$$

Therefore, for any fee smaller than \bar{C} there is no *none-make-AR* equilibrium. Furthermore, for any value of C between zero and \bar{C} , the equation $C = C(\tau_e)$ has a single solution and therefore the *some-make-AR* equilibrium is unique. The result is stated in the following theorem.

Theorem 3. *In a single server system, the equilibrium is unique and its type is:*

- *Some-make-AR equilibrium if $0 < C < \frac{e^{-\lambda} + \lambda - 1}{\lambda}$.*
- *None-make-AR equilibrium if $C > \frac{e^{-\lambda} + \lambda - 1}{\lambda}$.*

The expected revenue $R(\tau_e)$ in the case $N = 1$ is equal to the probability that at least one customer makes AR multiplied by the fee:

$$R(\tau_e) = (1 - e^{-\lambda\tau_e}) \left(\frac{e^{-\lambda(1-\tau_e)} + \lambda(1-\tau_e) - 1}{\lambda(1-\tau_e)} \right). \quad (3.35)$$

Since the equilibrium is unique, the provider will maximize her expected revenue by




Figure 3.6: $\lambda = 1.2$, $N = 1$. The reservation fee $C(\tau_e)$ is a monotonically increasing function of the threshold τ_e . The revenue function $R(\tau_e)$ is concave with maximum value at $\tau_e^* = 0.523$.

choosing $C^* = C(\tau_e^*)$ where:

$$\tau_e^* = \arg \max_{0 < \tau_e < 1} R(\tau_e). \quad (3.36)$$

Due to the uniqueness of the equilibrium, $R^* = R_g^* = R(\tau_e^*)$. Hence:

Corollary 1. *In a single server system, the price of conservatism is 1.*

Example 2. *We consider a system with $N = 1$ server and average demand $\lambda = 1.2$. For this system, the maximum fee that leads to a some-make-AR equilibrium is $\bar{C} = 0.417$. The optimal fee is obtained when $\tau_e^* = 0.477$ (i.e., when on average 47.7% of the customers make AR). This threshold is achieved when the provider sets a fee $C^* = 0.257$. The provider's expected revenue in this case is $R(\tau_e^*) = 0.112$. Figure 3.6 shows the fee and revenue as functions of the threshold τ_e .*

3.4.2 Many-server Case

In this section, we study the behavior of the system when the number of servers goes to infinity. We distinguish between overloaded and underloaded systems.

Overloaded system

We start with an overloaded system and we show that the PoC is a function of the ratio between the average demand and the number of servers.

Theorem 4. *In an overloaded many-server system, where $\lambda = \alpha N$ and $\alpha > 1$, the following holds:*

$$\lim_{N \rightarrow \infty} R^* = 1. \quad (3.37)$$

$$\lim_{N \rightarrow \infty} R_g^* = 1 - \frac{1}{\alpha}. \quad (3.38)$$

Hence,

$$PoC = \frac{\alpha}{\alpha - 1}. \quad (3.39)$$

Proof. In order to prove Eq. (3.37), we show that if the fee approaches one from below, there is a *some-make-AR* equilibrium where almost all servers are reserved.

Let $\tau_e = 1 - 1/\alpha$, hence D_{AR} is Poisson distributed with parameter N . The probability that the threshold customer gets service is equivalent to the probability that D_{AR} will be smaller than N , which, in turn, is equal to

$$\lim_{N \rightarrow \infty} \mathbb{P}(D_{AR} < N) = \frac{1}{2}. \quad (3.40)$$

Next, we show that when $\tau_e = 1 - 1/\alpha$, the probability to get service if not making AR tends to zero as $N \rightarrow \infty$. First, recall Chebyshev's inequality which states that

for any random variable X and real positive number Q

$$\mathbb{P}(|X - \mathbb{E}X| \geq Q) \leq \frac{\text{Var}X}{Q^2}. \quad (3.41)$$

Setting $Q = \delta\sqrt{N}$ where δ is a positive real number, we get from Eq. (3.41)

$$\mathbb{P}\left(|D_{AR} - N| \geq \delta\sqrt{N}\right) \leq \frac{1}{\delta^2}. \quad (3.42)$$

In the same way, setting $Q = \epsilon\sqrt{(\alpha - 1)N}$ where ϵ is a positive real number, we get

$$\mathbb{P}\left(|D_{AR'} - (\alpha - 1)N| \geq \epsilon\sqrt{(\alpha - 1)N}\right) \leq \frac{1}{\epsilon^2}. \quad (3.43)$$

Hence,

$$\mathbb{P}\left(D_{AR'} \leq (\alpha - 1)N - \epsilon\sqrt{(\alpha - 1)N}\right) \leq \frac{1}{\epsilon^2}. \quad (3.44)$$

From Eqs. (3.42) and (3.44), we deduce that with probability one the number of free servers $D_{AR} - N$ is $O(\sqrt{N})$ while $D_{AR'}$ is $(\alpha - 1)N + O(\sqrt{N})$. Hence, for any $\alpha > 1$, as $N \rightarrow \infty$, a customer that does not make AR will get service with probability zero.

We showed that when $N \rightarrow \infty$ and τ_e is such that on average N customers make AR, the expected payoff of the threshold customer tends to $0.5(1 - C)$ if making AR and to zero if not making AR. Thus, a strategy with threshold τ_e is an equilibrium only if C tends to one. Therefore, we conclude that there exists a value of τ_e such that on average N customers make AR while the fee is almost one. Therefore, in an overloaded system:

$$\lim_{N \rightarrow \infty} R^* = 1. \quad (3.45)$$

Next, we show that Eq. (3.38) holds. If none of the customers makes AR, they all

have the same probability to get service. Again, due to Chebyshev's inequality, for any $\delta > 1$, the following holds:

$$\mathbb{P}\left(\alpha N - \delta\sqrt{\alpha N} \leq D \leq \alpha N + \delta\sqrt{\alpha N}\right) \leq \frac{1}{\delta^2}. \quad (3.46)$$

In other words, with probability one the demand is $\alpha N + O(\sqrt{N})$. In this case, as $N \rightarrow \infty$ the fraction of customers getting service converges to $1/\alpha$. Hence,

$$\lim_{N \rightarrow \infty} \pi_{AR}(0) = \frac{1}{\alpha}. \quad (3.47)$$

On the other hand, when deviating from the *none-make-AR* strategy, the probability to get service is $\pi_{AR}(0) = 1$. Thus,

$$\lim_{N \rightarrow \infty} \underline{C} = 1 - \frac{1}{\alpha}. \quad (3.48)$$

Next, we show that in the overloaded system, if $C < \underline{C}$, then in any *some-make-AR* equilibrium almost all servers are reserved. By contradiction, we assume that there exists a *some-make-AR* equilibrium with threshold τ_e such that, $C(\tau_e) < 1 - 1/\alpha$ and $\mathbb{E}[D_{AR}] = \delta N$ where $0 < \delta < 1$. In this case, the probability of the threshold customer to get service if making AR converges to one as $N \rightarrow \infty$. Thus, his expected payoff is greater than $1/\alpha$. If not making AR, his probability to get service is smaller than $1/\alpha$ (which is the probability to get service if none makes AR). Therefore, the expected payoff of the threshold customer is greater if making AR than if not making AR, which contradicts the definition of a *some-make-AR* equilibrium. Hence, we have shown that the assumption cannot hold true. Thus, with probability one, the number of reservation will be at least $N + o(N)$. Therefore, with probability one, the ratio between the number of free servers and the number of servers is zero. The provider

will maximize her guaranteed revenue by advertising a fee just below \underline{C} and we finally obtain

$$\lim_{N \rightarrow \infty} R_g^* = 1 - \frac{1}{\alpha}. \quad (3.49)$$

□

The results indicate that if α is almost one and none of the customers makes AR then, in order to persuade customers to deviate, the provider will have to advertise a fee close to zero. Such fee will yield almost revenue per resource. In other words, although there is an equilibrium that yields a revenue per resource of almost one, if initially none of the customers makes AR, any fee the provider will advertise will not significantly increase her revenue.

Underloaded system

In an underloaded many-server system we show that any fee leads to an asymptotically zero revenue.

Theorem 5. *In an underloaded many-server system, where $\lambda = \alpha N$ and $\alpha < 1$, the following holds:*

$$\lim_{N \rightarrow \infty} R^* = R_g^* = 0. \quad (3.50)$$

Proof. Given any $\alpha < 1$ and any $\epsilon > 0$, we can find a large enough N such that

$$\mathbb{P}(D > N) \leq \epsilon. \quad (3.51)$$

In other words, for large enough N , the probability that the demand will exceed the number of servers tends to zero. In this case, the dominant strategy of all customers,

regardless of their arrival time, is not to make AR. Hence,

$$\lim_{N \rightarrow \infty} \bar{C} = 0. \quad (3.52)$$

□

3.5 Summary

In this chapter, we introduce AR games. We assume that customers have no information about the availability of servers and we analyzed two versions of the game. In one version, the AR cost is only applied to customers who get service, while in the second version, the AR cost is applied to all customers that attempt AR.

First, we show that, at equilibrium, either all customers that arrive earlier than some threshold point make AR or none of them makes AR. Next, we prove the existence of at least one Nash equilibrium and find the range of costs that determine each equilibrium. Furthermore, we show that in some situations there are multiple equilibria. By means of an example, we show that a given AR cost can lead to an equilibrium with zero reservations and also to an equilibrium where almost all servers are reserved in advance. We then assume that the AR cost is a fee charged by the provider. We show that charging a fee from all customers attempting to reserve a server can only reduce the provider's revenue.

In order for a provider to decide on a proper AR fee, we propose the concept of *Price of Conservatism (PoC)* which corresponds to the ratio of the maximum possible expected revenue to the maximum guaranteed expected revenue. A greater PoC indicates greater potential revenue loss if the provider opts to be conservative. We focus on the model where charges are collected only from the customers getting

service and assume that the demand is Poisson distributed. First, we show that in a single server system the equilibrium is unique. Thus, $PoC = 1$ and the provider experiences no loss. Next, we show that in an overloaded many-server system where the average demand is $\lambda = \alpha N$ with $\alpha > 1$, the maximum possible expected revenue tends to one, while, the maximum guaranteed expected revenue tends to $1 - 1/\alpha$ as $N \rightarrow \infty$. Hence $PoC = \alpha/(\alpha - 1)$, which increases in an unbounded fashion as α approaches 1 from above. Finally, we show that in an underloaded many-server system, the provider cannot make revenue from AR fees.

Chapter 4

Advance Reservation Games with Information Sharing

In the previous chapter, we assumed that the customers has no information about servers availability. In contrast, in this chapter, we assume that the provider shares some information with the customers. We analyze three versions of the game. In the first version (*binary-information game*), customers only know whether a server is available or not; in the second version (*full-information game*), customers know the exact number of available servers; finally, in the third version (*partial-information game*), the provider informs the customers about the exact number of available servers when this number is small, and hides this information otherwise.

After analyzing those three games separately, we use simulations and a numeric example to evaluate which policy maximizes the number of reservations.

The games in this chapter share the same assumptions as the first game described in Section 3.1 (i.e., the first *no-information model*), except for using different information sharing policies. Note that a version in which AR costs is applied on customers that do not get service (the *No-information model 2*) is not relevant for these three games since service is now granted to all customers that make AR.

4.1 Binary-information Game

In this model, customers make decisions based on statistical information and also based on the knowledge that a server is currently available at the desired slot. Next, we show that this additional information has no effect on the decisions of customers and each cost leads to the same set of equilibria as in the first *no-information* model.

Lemma 5. *In the binary-information-game, at equilibrium, all customers follow a threshold strategy.*

Proof. Suppose a customer is informed that a server is available. For this case, we show that the expected payoff of not making AR is a non-increasing function of the arrival time while the payoff of making AR is fixed to $1 - C$. Using similar arguments as in the *no-information* case (see Lemma 2), one can then show that the only possible equilibrium is a threshold strategy.

We define $\tilde{D}_{AR}(t)$ to be the number of reservations made before time t . We need to show that for any $t_1 > t_2$, regardless of the strategy σ followed by the rest of the customers, the following holds:

$$\mathbb{P}\left(S|t_1, \tilde{D}_{AR}(t_1) < N, AR'\right) \leq \mathbb{P}\left(S|t_2, \tilde{D}_{AR}(t_2) < N, AR'\right), \quad (4.1)$$

where the left (right) hand side is the probability of a customer with arrival time t_1 (t_2) to get service, given that the number of reservations made before his request is smaller than N and the chosen action is AR' . Using conditional probability, Eq. (4.1) can be rewritten as

$$\frac{\mathbb{P}\left(S, \tilde{D}_{AR}(t_1) < N|t_1, AR'\right)}{\mathbb{P}\left(\tilde{D}_{AR}(t_1) < N|t_1, AR'\right)} \leq \frac{\mathbb{P}\left(S, \tilde{D}_{AR}(t_2) < N|t_2, AR'\right)}{\mathbb{P}\left(\tilde{D}_{AR}(t_2) < N|t_2, AR'\right)}. \quad (4.2)$$

The event $\{S\}$ is contained the event $\{\tilde{D}_{AR}(\cdot) < N\}$. Moreover, under action AR'

the probability to get service does not depend on the arrival time. We deduce that the numerators on both sides of the equation above are equal.

Since $\tilde{D}_{AR}(t_2)$ is stochastically larger or equal to $\tilde{D}_{AR}(t_1)$ when $t_2 < t_1$, we deduce that the denominator of the right hand side of Eq. (4.2) is smaller or equal to the denominator of the left hand side of Eq. (4.2). Thus, we have shown that Eq. (4.1) holds. \square

Some-make-AR equilibrium. Consider first the *no-information* model and a *some-make-AR* equilibrium with threshold τ_e , but assume that the threshold customer is being informed that a server is available, namely $\tilde{D}_{AR}(\tau_e) < N$. If he makes AR, his payoff is $1 - C$. The expected payoff of not making AR is $\mathbb{P}(S|\tau_e, \tilde{D}_{AR}(\tau_e) < N, AR')$, which is the probability that the threshold customer gets service, given that all customers follow a strategy with threshold τ_e , that there is at least one free server and that the decision is AR' . Next, we show

$$\mathbb{P}(S|\tau_e, \tilde{D}_{AR}(\tau_e) < N, AR') = \frac{\pi_{AR'}(\tau_e)}{\pi_{AR}(\tau_e)}. \quad (4.3)$$

By conditioning on the event $\{\tilde{D}_{AR}(\tau_e) < N\}$, we get

$$\mathbb{P}(S|\tau_e, \tilde{D}_{AR}(\tau_e) < N, AR') = \frac{\mathbb{P}(S, \tilde{D}_{AR}(\tau_e) < N|\tau_e, AR')}{\mathbb{P}(\tilde{D}_{AR}(\tau_e) < N|\tau_e, AR')}. \quad (4.4)$$

Since a customer cannot get service when observing no free servers, the numerator $\mathbb{P}(S, \tilde{D}_{AR}(\tau_e) < N|\tau_e, AR')$ is equal to $\mathbb{P}(S|\tau_e, AR')$ which is equal by definition to $\pi_{AR'}(\tau_e)$.

The denominator $\mathbb{P}(\tilde{D}_{AR}(\tau_e) < N|\tau_e, AR')$ is the probability that the threshold customer will see the event $\{\tilde{D}_{AR}(\tau_e) < N\}$ (the fact that he does not make AR is irrelevant). This, in turn, can be rephrased as the probability to get service when

making AR exactly at the threshold point without knowing if there are free servers, which is the definition of $\pi_{AR}(\tau_e)$. Thus, $\mathbb{P}\left(\tilde{D}_{AR}(\tau_e) < N | \tau_e, AR'\right) = \pi_{AR}(\tau_e)$.

We have shown that Eq. (4.3) holds true for any τ_e . Using Eq. (3.8), we deduce that the threshold customer stays indifferent between the two actions after being informed that a server is available. Hence, we conclude that if a threshold strategy is an equilibrium strategy in the first model, it is also an equilibrium strategy in the third model.

None-make-AR equilibrium. If none of the customers makes AR, the expected payoffs of not making AR and the expected payoff of deviating are the same as in the first and second models. Therefore, the range of fees that have a *none-make-AR* equilibrium is the same as in the other two models.

Theorem 6. *In AR games, if the AR cost is applied only on served customers, then informing customers that servers are available or hiding this information lead to the same equilibria.*

4.2 Full-information Game

In this section, we assume that upon making an inquiry, customers observe the exact number of available servers $n \in \{0, N\}$. This change in the model greatly affects the analysis and nature of equilibria, since the decision of a customer affects information provided to other customers.

4.2.1 Preliminaries

First, we provide definitions and technical results that which we will later use to prove the main results of this section.

Probability of Service. We define $\mathbb{P}(S|\sigma, \tilde{D}_{AR}(t) = N - n)$ to be the *probability of service* of a customer that arrives at time t , observes $n \geq 1$ available servers and chooses AR' , while the rest of the customers follows the strategy σ . The probability of service applies only to customers that choose AR' , since service is granted to customers that choose AR upon observing $n \geq 1$ available servers.

Non-degenerate Equilibria. As in the previous chapter, we wish to ignore degenerate equilibria. Thus, we assume:

Assumption 2. *Let $n \geq 1$, $0 \leq t \leq 1$, and $0 \leq x \leq 1$. At equilibrium, if $\sigma(n, t) = x$, then there must exist a non-zero measure interval I such that $t \in I$ and $\sigma(n, t') = x$ for all $t' \in I$.*

Lemma 6. *Let X_t be a non-negative random variable with parameter t . Let $g(x)$ be a non-negative function of x whose derivative is positive with respect to x for $x \in [k, l]$ and non-negative elsewhere. If the complementary cumulative distribution function (CCDF) $\bar{F}_{X_t}(x)$ strictly increases with t for any $x \in [k, l]$ and increases with t elsewhere, then $\mathbb{E}[g(X_t)]$ strictly increases with t .*

Proof. Let $x \in [k, l]$, $\gamma = g(x)$ and $t_2 > t_1$. Then,

$$\mathbb{P}(g(X_{t_2}) > \gamma) = \mathbb{P}(X_{t_2} > x) > \mathbb{P}(X_{t_1} > x) = \mathbb{P}(g(X_{t_1}) > \gamma), \quad \forall x \in [k, l]. \quad (4.5)$$

Using the well-known formula for the expectation of a non-negative random variable (Ross et al., 1996, Chapter 9)

$$\mathbb{E}[X] = \int_0^{\infty} \mathbb{P}(X > \gamma) d\gamma, \quad (4.6)$$

we obtain

$$\mathbb{E}[g(X_{t_1})] = \int_0^{g(k)} \mathbb{P}(g(X_{t_1}) > \gamma) d\gamma + \int_{g(k)}^{g(l)} \mathbb{P}(g(X_{t_1}) > \gamma) d\gamma + \int_{g(l)}^{\infty} \mathbb{P}(g(X_{t_1}) > \gamma) d\gamma \quad (4.7)$$

and

$$\mathbb{E}[g(X_{t_2})] = \int_0^{g(k)} \mathbb{P}(g(X_{t_2}) > \gamma) d\gamma + \int_{g(k)}^{g(l)} \mathbb{P}(g(X_{t_2}) > \gamma) d\gamma + \int_{g(l)}^{\infty} \mathbb{P}(g(X_{t_2}) > \gamma) d\gamma. \quad (4.8)$$

From Eq. (4.5), we know that the middle term in the RHS of Eq. (4.7) is strictly smaller than that in Eq. (4.8), while the two other terms in Eq. (4.7) are no larger than the corresponding terms in Eq. (4.8). Thus, we deduce that $\mathbb{E}[g(X_{t_2})] > \mathbb{E}[g(X_{t_1})]$. \square

Definition 6. Let $\{a_k, \dots, a_l\}$ be a set of positive real numbers, t be a real number and $h(t)$ be a general positive function of t whose derivative is strictly negative with respect to t . A discrete non-negative random variable X_t supported in $[k, l]$ is said to belong to the distributions family \mathcal{F} if it has the following CDF:

$$F_{X_t}(x) \triangleq \mathbb{P}(X_t \leq x) = \frac{\sum_{j=k}^x a_j (h(t))^j}{\sum_{j=k}^l a_j (h(t))^j}, \quad \forall x \in [k, l]. \quad (4.9)$$

Lemma 7. Suppose X_t is a discrete random variable, supported in $[k, l]$ with a CDF F_{X_t} . If $F_{X_t} \in \mathcal{F}$, then for any $x \in [k, l]$, $F_{X_t}(x)$ strictly increases with t .

Proof. We compute the derivative of F_{X_t} (as defined in Eq. (4.9)) with respect to t

and show that it is positive:

$$\begin{aligned}
\frac{dF_{X_t}}{dt} &= \frac{-\sum_{i=k}^x a_i i |h'(t)| (h(t))^{i-1} \sum_{j=k}^l a_j (h(t))^j + \sum_{i=k}^x a_i (h(t))^i \sum_{j=k}^l a_j j |h'(t)| (h(t))^{j-1}}{\left(\sum_{j=k}^l a_j (h(t))^j\right)^2} \\
&= \frac{\sum_{i=k}^x \left(a_i (h(t))^{i-1} \left(-i |h'(t)| \sum_{j=k}^l a_j (h(t))^j + \sum_{j=k}^l a_j j |h'(t)| (h(t))^j \right) \right)}{\left(\sum_{j=k}^l a_j (h(t))^j\right)^2} \\
&= \frac{\sum_{i=k}^x a_i (h(t))^{i-1} \sum_{j=k}^l a_j (h(t))^j |h'(t)| (-i+j)}{\left(\sum_{j=k}^l a_j (h(t))^j\right)^2} \\
&= \frac{\sum_{i=k}^x a_i (h(t))^{i-1} \sum_{j=k}^x a_j (h(t))^j |h'(t)| (-i+j)}{\left(\sum_{j=k}^l a_j (h(t))^j\right)^2} + \frac{\sum_{i=k}^x a_i (h(t))^{i-1} \sum_{j=x+1}^l a_j (h(t))^j |h'(t)| (-i+j)}{\left(\sum_{j=k}^l a_j (h(t))^j\right)^2}.
\end{aligned} \tag{4.10}$$

In the numerator of Eq. (4.10), the first term is canceled out due to symmetry and the second term is positive for any $x \in [k, l)$ \square

Next, we show that for any strategy σ , the number of reservations (stochastically) increases with the demand.

Lemma 8. *If all customers follow an arbitrary strategy σ , then, $\mathbb{P}(\tilde{D}_{AR} > \tilde{d}_{AR} | \sigma, \tilde{D} = \tilde{d})$ increases with \tilde{d} .*

Proof. The proof is based on a coupling argument (Ross et al., 1996, Chapter 9). Consider a realization R_1 with demand \tilde{d} and set of arrival times $T = \{t_1, t_2, \dots, t_d\}$. Consider a second realization R_2 which is identical to R_1 but with an additional customer with arrival time t' that observes n' available servers. We respectively denote

$\tilde{D}_{AR}^1(t)$ and $\tilde{D}_{AR}^2(t)$ as the number of reservations made by time t in realizations R_1 and R_2 .

If $\sigma(t', n') = 0$ (i.e., the additional customer does not make AR), then R_1 and R_2 are identical in terms of the number of reservations. Otherwise, $\tilde{D}_{AR}^2(t') > \tilde{D}_{AR}^1(t')$. Let $T' = \{t \mid t \in T, t > t'\}$. If there is an arrival point $t'' \in T'$ such that $\tilde{D}_{AR}^2(t'') = \tilde{D}_{AR}^1(t'')$, then R_1 and R_2 merge (i.e., the decisions of all customers that arrive after t'' are identical in both realizations). In this case, the number of reservations in both realizations are equal. If there is no such t'' , then the total number of reservations in R_2 is larger than in R_1 . We conclude that, in any case, the number of reservations \tilde{D}_{AR} cannot decrease with the demand. \square

4.2.2 Equilibria Analysis

First, we show that even if customers are indifferent between the actions AR and AR' , a mixed strategy cannot be an equilibrium.

Lemma 9. *At equilibrium, none of the customers uses a mixed strategy.*

Proof. The proof goes by contradiction. Consider an interval of arrival times I and a number of available servers $1 \leq n \leq N$, such that all customers with arrival times in I that observe n available servers use a mixed strategy. In this case, there is a strictly positive probability that one or more reservations will be made during this time interval. Thus, the probability of service, when choosing AR' , depends on the arrival time. In contrast, the payoff of choosing AR is a constant. Thus, a non-zero measure interval in which customers are indifferent between the actions AR and AR' does not exist, and hence a mixed strategy cannot be an equilibrium strategy in that interval. Note that the case of a mixed strategy over a measure zero interval can be ignored by Assumption 2. \square

Next, we state that there exist N thresholds, such that, at equilibrium, a customer that observes n available servers will make a reservation only if he arrives before the n -th threshold.

Definition 7. Let $\tau^e = \{\tau_1^e, \tau_2^e, \dots, \tau_N^e\}$, where $0 \leq \tau_N^e \leq \tau_{N-1}^e \leq \dots \leq \tau_1^e \leq 1$. A strategy function $\sigma(t, n)$ is a threshold strategy if it has the form:

$$\sigma(t, n) = \begin{cases} 1 & \text{if } t < \tau_n^e, \\ 0 & \text{if } t \geq \tau_n^e. \end{cases} \quad (4.11)$$

Theorem 7. At equilibrium, all customers follow a threshold strategy.

Proof. We prove the theorem by induction. First, we show that, at equilibrium, a threshold strategy is followed by customers that observe N available servers. Then, we show that if all customers that observe $n \in [k, k + 1, \dots, N]$ available servers follow a threshold strategy, then a threshold strategy is also followed by customers that observe $k - 1$ available servers.

Base case. From Lemma 9, we know that, at equilibrium, the strategy followed by all customers that observe a certain number of available servers is a set of intervals where in each interval either all customers choose AR or all choose AR' . Consider the intervals $I_1 = (t_1, t_2)$ and $I_2 = (t_2, t_3)$, where $t_3 > t_2 > t_1$. Let assume by contradiction that there is an equilibrium strategy σ such that $\sigma(t, N) = 0$ for all $t \in I_1$ and $\sigma(t, N) = 1$ for all $t \in I_2$. In this case, the following must hold:

$$\mathbb{P}(S|\sigma, \tilde{D}_{AR}(t)=0) \geq 1 - C, \quad \forall t \in I_1 \quad (4.12)$$

and

$$\mathbb{P}(S|\sigma, \tilde{D}_{AR}(t)=0) \leq 1 - C, \quad \forall t \in I_2. \quad (4.13)$$

Next, we show that $\mathbb{P}(S|\sigma, \tilde{D}_{AR}(t) = 0)$ strictly increases with t within I_2 , and hence the payoff when choosing AR along the interval I_2 must be strictly larger than $1 - C$, which leads to a contradiction.

Using the law of total probability, conditioned on the demand, the probability of service can be written

$$\mathbb{P}(S|\sigma, \tilde{D}_{AR}(t) = 0) = \sum_{i=a}^b \mathbb{P}(S|\sigma, \tilde{D} = i, \tilde{D}_{AR}(t) = 0) \mathbb{P}(\tilde{D} = i|\sigma, \tilde{D}_{AR}(t) = 0). \quad (4.14)$$

Using Lemma 6, we next show that $\mathbb{P}(S|\sigma, \tilde{D}_{AR}(t) = 0)$ strictly increases with t . To do so, we show that

- (i) For any $\tilde{d} \in [a, b)$, $\mathbb{P}(\tilde{D} > \tilde{d}|\sigma, \tilde{D}_{AR}(t) = 0)$ strictly decreases with t (the conditional random variable $\tilde{D}|\{\sigma, \tilde{D}_{AR}(t) = 0\}$ corresponds to X_t in Lemma 6).
- (ii) Within the range $\tilde{d} \in [N, b]$, $\mathbb{P}(S|\sigma, \tilde{D} = \tilde{d}, \tilde{D}_{AR}(t) = 0)$ (which corresponds to $g(x)$ in Lemma 6) strictly decreases with \tilde{d} (which corresponds to x in Lemma 6).

Starting with (i), we define $T_{AR'}(t)$ to be the sum of length of intervals within customers that observe N available servers choose AR' prior to t under strategy σ . From Bayes' Theorem, the distribution of $\tilde{D}|\{\sigma, \tilde{D}_{AR}(t) = 0\}$ is

$$\mathbb{P}(\tilde{D} = i|\sigma, \tilde{D}_{AR}(t) = 0) = \frac{\mathbb{P}(\tilde{D} = i, \tilde{D}_{AR}(t) = 0|\sigma)}{\mathbb{P}(\tilde{D}_{AR}(t) = 0|\sigma)}, \quad \forall i \in [a, b]. \quad (4.15)$$

The term $\mathbb{P}(\tilde{D} = i, \tilde{D}_{AR}(t) = 0|\sigma)$ is the probability that i customers arrive, and each customer arrives either within intervals where customers do not make AR or after t . Since the arrival time of each customer is independent of others, the probability that a customer arrives either within intervals where customers do not make AR or after

t is $T_{AR'}(t) + (1 - t)$. Thus, for any $i \in [a, b]$,

$$\begin{aligned} \frac{\mathbb{P}(\tilde{D} \leq i, \tilde{D}_{AR}(t) = 0 | \sigma)}{\mathbb{P}(\tilde{D}_{AR}(t) = 0 | \sigma)} &= \frac{\sum_{j=a}^i \mathbb{P}(\tilde{D} = j) \mathbb{P}(\tilde{D}_{AR}(t) = 0 | \sigma, \tilde{D} = j)}{\sum_{j=a}^b \mathbb{P}(\tilde{D} = j) \mathbb{P}(\tilde{D}_{AR}(t) = 0 | \sigma, \tilde{D} = j)} \\ &= \frac{\sum_{j=a}^i \mathbb{P}(\tilde{D} = j) (T_{AR'}(t) + 1 - t)^j}{\sum_{j=a}^b \mathbb{P}(\tilde{D} = j) (T_{AR'}(t) + 1 - t)^j}. \end{aligned} \quad (4.16)$$

Since $T_{AR'}(t)$ is a constant within the interval I_2 , the term $T_{AR'}(t) + 1 - t$ decreases with t . By Definition 6, the conditional random variable $\tilde{D} | \{\sigma, \tilde{D}_{AR}(t) = 0\}$ belongs to the family of distributions \mathcal{F} , where $\mathbb{P}(\tilde{D} = j)$ corresponds to a_j and $T_{AR'}(t) + 1 - t$ corresponds to $h(t)$. Hence, by Lemma 7, for any $\tilde{d} \in [a, b)$, $\mathbb{P}(\tilde{D} > \tilde{d} | \sigma, \tilde{D}_{AR}(t) = 0)$ strictly decreases with t .

Next we prove (ii). Using the law of total probability, conditioned on the number of reservations, the probability of service can be written

$$\mathbb{P}(S | \sigma, \tilde{D} = \tilde{d}, \tilde{D}_{AR}(t) = 0) = \sum_{i=a}^b \mathbb{P}(S | \tilde{D} = \tilde{d}, \tilde{D}_{AR} = i) \mathbb{P}(\tilde{D}_{AR} = i | \sigma, \tilde{D} = \tilde{d}, \tilde{D}_{AR}(t) = 0). \quad (4.17)$$

From the definition of the model,

$$\mathbb{P}(S | \tilde{D} = \tilde{d}, \tilde{D}_{AR} = \tilde{d}_{AR}) = \begin{cases} 1 & \text{if } \tilde{d} < N, \\ \frac{N - \tilde{d}_{AR}}{\tilde{d} + 1 - \tilde{d}_{AR}} & \text{if } \tilde{d} \geq N \text{ and } \tilde{d}_{AR} < N, \\ 0 & \text{otherwise.} \end{cases} \quad (4.18)$$

From Eq. (4.18), one can see that $\mathbb{P}(S | \tilde{D} = \tilde{d}, \tilde{D}_{AR} = \tilde{d}_{AR})$ decreases with both \tilde{d} and \tilde{d}_{AR} . Furthermore, Lemma 8 showed that the number of reservations cannot decrease as the demand increases. Thus, for any demand realizations $\tilde{d}_1 > \tilde{d}_2$ the following

holds:

$$\mathbb{P}(S|\sigma, \tilde{D}=\tilde{d}_1, \tilde{D}_{AR}(t)=0) = \sum_{i=a}^b \mathbb{P}(S|\tilde{D}=\tilde{d}_1, \tilde{D}_{AR}=i)\mathbb{P}(\tilde{D}_{AR}=i|\sigma, \tilde{D}=\tilde{d}_1, \tilde{D}_{AR}(t)=0) \quad (4.19)$$

$$\leq \sum_{i=a}^b \mathbb{P}(S|\tilde{D}=\tilde{d}_2, \tilde{D}_{AR}=i)\mathbb{P}(\tilde{D}_{AR}=i|\sigma, \tilde{D}=\tilde{d}_1, \tilde{D}_{AR}(t)=0) \quad (4.20)$$

$$\leq \sum_{i=a}^b \mathbb{P}(S|\tilde{D}=\tilde{d}_2, \tilde{D}_{AR}=i)\mathbb{P}(\tilde{D}_{AR}=i|\sigma, \tilde{D}=\tilde{d}_2, \tilde{D}_{AR}(t)=0) \quad (4.21)$$

$$= \mathbb{P}(S|\sigma, \tilde{D}=\tilde{d}_2, \tilde{D}_{AR}(t)=0). \quad (4.22)$$

Moreover, from Eq. (4.18), it follows that if $\tilde{d}_1 \geq N$, then Eq. (4.19) is strictly smaller than Eq. (4.20).

We showed that both items 1 and 2 hold. Hence, $\mathbb{P}(S|\sigma, \tilde{D}_{AR}(t) = 0)$ strictly increases with t . Therefore, at equilibrium, an interval within which customers choose AR' cannot be followed by an interval within which customers choose AR . Hence, at equilibrium, the strategy followed by customers that observe N available servers must be a threshold strategy.

Inductive step. We assume that a threshold strategy with thresholds $\tau_k \geq \tau_{k+1}, \geq \dots \geq \tau_N$ is followed by all customers that observe $n \in [k, k+1, \dots, N]$ available servers for $1 < k < N$. We show that, at equilibrium, a threshold strategy with threshold $\tau_{k-1} \geq \tau_k$ must be followed by customers that observe $k-1$ available servers. We split the proof into two parts. In the first part, we show that a customer with arrival time $t' < \tau_k$ that observes $k-1$ available servers is better off choosing AR . Thus, there is a threshold $\tau_{k-1} \geq \tau_k$ such that all customers that arrive before

τ_{k-1} and observe $k-1$ available servers choose AR . In the second part, we show that all customers that arrive after τ_{k-1} and observe $k-1$ available servers choose AR' .

Consider a scenario where a customer arrives at time $t < \tau_k$ and observes k available servers and a second scenario where a customer arrives at the same time and observes $k-1$ available servers. We need to show that

$$\mathbb{P}(S|\sigma, \tilde{D}_{AR}(t) = N-k) \geq \mathbb{P}(S|\sigma, \tilde{D}_{AR}(t) = N-k+1). \quad (4.23)$$

From Eq. (4.19) - Eq. (4.22), it follows that the probability of service increases with the demand, regardless of the number of available servers observed upon arrival (i.e., item 2 in the base case holds also for $\tilde{D}_{AR}(t) \neq 0$). Thus, we only need to show that

$$\mathbb{P}(\tilde{D} > \tilde{d} | \sigma, \tilde{D}_{AR}(t) = N-k+1) \geq \mathbb{P}(\tilde{D} > \tilde{d} | \sigma, \tilde{D}_{AR}(t) = N-k). \quad (4.24)$$

From the induction's assumption, a customer that observes k available servers knows that exactly $N-k$ customers arrived earlier. Likewise, a customer that observes $k-1$ available servers knows that at least $N-k+1$ customers arrived earlier. Since the system is sampled at the same time and the same strategy is followed in both cases, and since arrivals are i.i.d, observing more reservations makes it more likely that the demand will be larger. Thus, Eq. (4.24) holds.

Next, we prove the second part. We do so by showing that both items 1 and 2 in the base case also hold when observing $k-1$ available servers. Starting with item 1, we show that the conditional random variable $\tilde{D}|\{\sigma, \tilde{D}_{AR}(t) = N-k+1\}$ belongs to \mathcal{F} . We denote by AR_{N-k+1} the event that the first $N-k+1$ customers choose AR (i.e., the i -th customer arrives before τ_i , for $i = 1, 2, \dots, N-k+1$). Let $T_{AR'}(t)$ be the sum of length of intervals within which customers that observe $k-1$ available

servers choose AR' under strategy σ . The CDF of $\tilde{D}|\{\sigma, \tilde{D}_{AR}(t) = N - k + 1\}$ is

$$\begin{aligned}
\mathbb{P}(\tilde{D} \leq j | \sigma, \tilde{D}_{AR}(t) = N - k + 1) &= \frac{\sum_{j=N-k+1}^i \mathbb{P}(\tilde{D} = j, \tilde{D}_{AR}(t) = N - k + 1 | \sigma)}{\mathbb{P}(\tilde{D}_{AR}(t) = N - k + 1)} \\
&= \frac{\sum_{j=N-k+1}^i \mathbb{P}(\tilde{D} = j) \binom{j}{N-k+1} \mathbb{P}(AR_{N-k+1})(T_{AR'}(t) + 1 - t)^{j-(N-k+1)}}{\sum_{j=N-k+1}^b \mathbb{P}(\tilde{D} = j) \binom{j}{N-k+1} \mathbb{P}(AR_{N-k+1})(T_{AR'}(t) + 1 - t)^{j-(N-k+1)}} \\
&= \frac{\sum_{j=N-k+1}^i \mathbb{P}(\tilde{D} = j) \binom{j}{N-k+1} (T_{AR'}(t) + 1 - t)^j}{\sum_{j=N-k+1}^b \mathbb{P}(\tilde{D} = j) \binom{j}{N-k+1} (T_{AR'}(t) + 1 - t)^j}, \quad \forall i \geq N - k + 1. \tag{4.25}
\end{aligned}$$

By Definition 6, this distribution belongs to \mathcal{F} , where $\mathbb{P}(\tilde{D} = j) \binom{j}{N-k+1}$ corresponds to a_j and $T_{AR'}(t) + 1 - t$ corresponds to $h(t)$. Thus, item 1 holds. As explained earlier in the proof, item 2 also holds for $D_{AR}(t) = N - k + 1$. We conclude that for $n = k - 1$ it is also true that an interval within which customers choose AR cannot follow an interval within which customers choose AR' . Hence, a threshold strategy is followed by all customers that observe $k - 1$ available servers. \square

4.2.3 Impact of the AR Cost

In this chapter, we show that any AR cost leads to a unique equilibrium. Furthermore, we show that there is a critical value \underline{C} such that costs higher than \underline{C} lead to an equilibrium with no reservations.

First, we denote by σ_{τ^e} an equilibrium strategy constructed by a set of N thresholds $\tau^e = \{\tau_1^e, \tau_2^e, \dots, \tau_N^e\}$. For such a strategy to be an equilibrium, the following must hold for any $n \in \{1, \dots, N\}$:

(i) If $\tau_n^e > 0$, then

$$\mathbb{P}(S|\sigma_{\tau^e}, \tilde{D}_{AR}(\tau_n^e) = N - n) = 1 - C. \quad (4.26)$$

(ii) If $\tau_n^e = 0$, then

$$\mathbb{P}(S|\sigma_{\tau^e}, \tilde{D}_{AR}(\tau_n^e) = N - n) \geq 1 - C. \quad (4.27)$$

To understand this property, consider a virtual customer that arrives at time τ_n^e and observes $1 \leq n \leq N$ available servers. We refer to such a customer as the *n-th threshold customer*. If $\tau_n^e > 0$, then customers that arrive before τ_n^e choose AR , while customers that arrive after τ_n^e choose AR' . Hence, a customer that arrives exactly at the threshold must be indifferent between the two options. If $\tau_n^e = 0$, then all customers choose AR' . Thus, the probability of service of all customers (including the *n*-th threshold customer) should be no smaller than $1 - C$.

To prove the uniqueness of the equilibrium, we will show that for each $1 \leq n \leq N$ there is exactly one threshold τ_n^e for which Eq. (4.26) or Eq. (4.27) hold. For this purpose, we develop an expression of the probability of service of the *n*-th threshold customer as a function of the set of thresholds $\boldsymbol{\tau} = \{\tau_1, \tau_2, \dots, \tau_N\}$. We denote this quantity by $\pi_n(\boldsymbol{\tau})$. If the *n*-th threshold customer chooses AR' , then all customers that arrive after his also choose AR' . In this case $\tilde{D}_{AR} = N - n$. The probability of service of the *n*-th threshold customer is

$$\pi_n(\boldsymbol{\tau}) = \sum_{i=N-n}^{N-1} \mathbb{P}(\tilde{D} = i | \tilde{D}_{AR} = N - n) \cdot 1 + \sum_{i=N}^b \mathbb{P}(\tilde{D} = i | \tilde{D}_{AR} = N - n) \cdot \frac{n}{i+1 - (N-n)}. \quad (4.28)$$

The first term is the probability that the total demand is smaller than N (recall that \tilde{D} is the number of customers as observed by a tagged customer excluding himself). In this case, service is guaranteed to all customers. The second term is the probability to get service when the demand exceeds N . In this case, the n unreserved servers will be arbitrarily allocated among customers that chose AR' . Next, we develop an expression for $\pi_n(\boldsymbol{\tau})$ that directly relates to the system parameters (i.e., an expression that does not include \tilde{D}_{AR}).

Lemma 10. *Let $1 \leq n \leq N$. The probability of service of the n -th threshold customer is*

$$\pi_n(\boldsymbol{\tau}) = \frac{\sum_{i=N-n}^{N-1} \mathbb{P}(\tilde{D}=i) \binom{i}{N-n} (1-\tau_n)^i + \sum_{i=N}^b \mathbb{P}(\tilde{D}=i) \binom{i}{N-n} (1-\tau_n)^i \frac{n}{i+1-(N-n)}}{\sum_{i=N-n}^b \mathbb{P}(\tilde{D}=i) \binom{i}{N-n} (1-\tau_n)^i}. \quad (4.29)$$

Proof. First, we find an expression for $\mathbb{P}(\tilde{D} = i | \tilde{D}_{AR} = N - n)$. For brevity, let $m \triangleq N - n$. Using Bayes' theorem, we obtain

$$\begin{aligned} \mathbb{P}(\tilde{D} = i | \tilde{D}_{AR} = m) &= \frac{\mathbb{P}(\tilde{D} = i) \mathbb{P}(\tilde{D}_{AR} = m | \tilde{D} = i)}{\mathbb{P}(\tilde{D}_{AR} = m)} \\ &= \frac{\mathbb{P}(\tilde{D} = i) \mathbb{P}(\tilde{D}_{AR} = m | \tilde{D} = i)}{\sum_{k=m}^b \mathbb{P}(\tilde{D} = k) \mathbb{P}(\tilde{D}_{AR} = m | \tilde{D} = k)}, \quad \forall i \in [m, b]. \end{aligned} \quad (4.30)$$

Next, we develop an expression for $\mathbb{P}(\tilde{D}_{AR} = m | \tilde{D} = i)$. Given that i customers request service, the probability that exactly m customers choose AR is the probability that the first m customers among them choose AR (i.e., the arrival time of the j -th customer is smaller than τ_{N-j} for all $j \in \{1, \dots, m\}$), while the last $i - m$ customers choose AR' (i.e., the arrival times of the $m + 1$ customer and on are larger than τ_n).

We denote by AR_m the event that a given subset of m customers choose AR in a system that it is initially empty. The probability that the remaining $i - m$ customers choose AR' in a system with n available servers is $(1 - \tau_n)^{i-m}$. Thus,

$$\mathbb{P}(\tilde{D}_{AR}=m|\tilde{D}=i) = \binom{j}{m} \mathbb{P}(AR_m)(1 - \tau_n)^{i-m}. \quad (4.31)$$

The term $\mathbb{P}(AR_m)$ will be canceled out, and therefore we do not develop an expression for it. Substituting Eq. (4.30) and Eq. (4.31) into Eq. (4.28), we get

$$\pi_n(\boldsymbol{\tau}) = \frac{\sum_{i=m}^{N-1} \mathbb{P}(\tilde{D}=i) \binom{i}{m} \mathbb{P}(AR_m)(1 - \tau_n)^{(i-m)} + \sum_{i=N}^b \mathbb{P}(\tilde{D}=i) \binom{i}{m} \mathbb{P}(AR_m)(1 - \tau_n)^{i-m} \frac{n}{i+1-m}}{\sum_{i=m}^b \mathbb{P}(\tilde{D}=i) \binom{i}{m} \mathbb{P}(AR_m)(1 - \tau_n)^{(i-m)}}. \quad (4.32)$$

Note that $\mathbb{P}(AR_m)$ and $(1 - \tau_n)^{-m}$ do not depend on i . Since those two terms appear in each term of the numerator and the denominator, we can cancel them out and conclude that Eq. (4.29) holds. \square

From Lemma 10, we deduce:

Corollary 2. $\pi_n(\boldsymbol{\tau})$ only depends on $\tau_n \in \boldsymbol{\tau}$.

We thus redefine the function $\pi_n(\cdot)$ such that the threshold τ_n is its only input.

Example 3. Assume that the demand is uniformly distributed between 1, 2 and 3.

We compute the distribution of \tilde{D} using Eq. (3.1):

$$\mathbb{P}(\tilde{D}=i) = \begin{cases} 1/6 & \text{if } i = 0, \\ 1/3 & \text{if } i = 1, \\ 1/2 & \text{if } i = 2. \end{cases}$$

Applying this demand distribution on Eq. (4.29), we get

$$\pi_2(\tau_2) = \frac{1 + 2(1 - \tau_2) + 2(1 - \tau_2)^2}{1 + 2(1 - \tau_2) + 3(1 - \tau_2)^2}, \quad (4.33)$$

and

$$\pi_1(\tau_1) = \frac{2 + 3(1 - \tau_1)}{2 + 6(1 - \tau_1)}. \quad (4.34)$$

We will later return to this example. Next, we show that $\pi_n(\tau_n)$ increases with both τ_n and n . Those two properties are required for proving the uniqueness and existence of the equilibrium.

Lemma 11. *For any $n \in \{1, 2, \dots, N\}$, the following holds:*

- (i) $\pi_n(\tau_n)$ is continuous and strictly increasing in the range $[0, 1)$.
- (ii) $\pi_n(\tau) > \pi_{n-1}(\tau)$, for any $\tau \in [0, 1)$.

Proof. Set $m = N - n$ and, for $j \in [m, b]$, define

$$\alpha_j \triangleq \mathbb{P}(\tilde{D} = j) \binom{j}{m}, \quad (4.35)$$

and

$$\beta_j \triangleq \begin{cases} 1 & \text{if } j < N, \\ \frac{n}{j+1-m} & \text{if } j \geq N. \end{cases} \quad (4.36)$$

We rewrite Eq. (4.29)

$$\pi_n(\tau_n) = \frac{\sum_{j=m}^b \alpha_j \beta_j (1 - \tau_n)^j}{\sum_{i=m}^b \alpha_i (1 - \tau_n)^i}. \quad (4.37)$$

For proving (i), we compute the derivative of $\pi_n(\tau_n)$ with respect to τ_n and show that it is positive for any $\tau_n \in [0, 1)$. We have

$$\begin{aligned} \frac{d\pi_n}{d\tau_n} &= \frac{-\sum_{j=m}^b \alpha_j \beta_j j (1-\tau_n)^{j-1} \sum_{i=m}^b \alpha_i (1-\tau_n)^i + \sum_{j=m}^b \alpha_j \beta_j (1-\tau_n)^j \sum_{i=m}^b \alpha_i i (1-\tau_n)^{i-1}}{\left(\sum_{i=m}^b \alpha_i (1-\tau_n)^i\right)^2} \\ &= \frac{\sum_{j=m}^b \sum_{i=m}^b \alpha_j \alpha_i \beta_j (1-\tau_n)^{i+j-1} (i-j)}{\left(\sum_{i=m}^b \alpha_i (1-\tau_n)^i\right)^2}. \end{aligned} \quad (4.38)$$

We need to show that the numerator of Eq. (4.38) is positive. Let consider any element of the numerator $\{i = k, j = l\}$ and its conjugate element $\{i = l, j = k\}$. By summing these two elements we get

$$\alpha_l \alpha_k (1-\tau_n)^{l+k-1} (k-l) (\beta_l - \beta_k). \quad (4.39)$$

Since $\beta_i = \beta_{i-1}$ for all $i < N$ and $\beta_i > \beta_{i-1}$ for all $i \geq N$, we deduce that $\beta_l - \beta_k = 0$ if both k and l are smaller than N . Otherwise, $\beta_l - \beta_k$ has the same sign as $k - l$. Thus, for any value of k and l , Eq. (4.39) is either zero or positive. We conclude that the sum of the numerator of Eq. (4.38) is positive, and hence the derivative of $\pi_n(\tau_n)$ is positive.

Next, we show that $\pi_n(\tau_n)$ is a continuous function in $[0, 1)$. In Eq. (4.32), both the numerator and denominator are probabilities. Hence, although they sum to infinity, they have a finite limit for any value of $\tau_n \in [0, 1)$. Since they are both polynomial expressions of τ_n with a finite limit, they are continuous (see Cauchy's uniform convergence criterion (Trench, 2003, p.246)). The denominator is equal to zero only at $\tau_n = 1$. Thus, we conclude that $\pi_n(\tau_n)$ is continuous in $[0, 1)$.

For proving (ii), we show that

$$\frac{\sum_{i=m}^b \alpha_i \beta_i (1-\tau)^i}{\sum_{i=m}^b \alpha_i (1-\tau)^i} - \frac{\sum_{j=m+1}^b \alpha_j \beta_j (1-\tau)^j}{\sum_{j=m+1}^b \alpha_j (1-\tau)^j} > 0. \quad (4.40)$$

This is equivalent to showing that

$$\begin{aligned} & \sum_{i=m}^b \alpha_i \beta_i (1-\tau)^i \sum_{j=m+1}^b \alpha_j (1-\tau)^j - \sum_{i=m}^b \alpha_i (1-\tau)^i \sum_{j=m+1}^b \alpha_j \beta_j (1-\tau)^j \\ &= \sum_{i=m}^b \alpha_i (1-\tau)^i \left(\beta_i \sum_{j=m+1}^b \alpha_j (1-\tau)^j - \sum_{j=m+1}^b \alpha_j \beta_j (1-\tau)^j \right) \\ &= \sum_{i=m}^b \alpha_i (1-\tau)^i \left(\sum_{j=m+1}^b \alpha_j (1-\tau)^j (\beta_i - \beta_j) \right) \\ &= \sum_{i=m}^b \alpha_i (1-\tau)^i \left(\sum_{j=m}^b \alpha_j (1-\tau)^j (\beta_i - \beta_j) \right) + \sum_{i=m}^b \alpha_i \alpha_m (1-\tau)^{i+m} (\beta_m - \beta_i) > 0. \end{aligned} \quad (4.41)$$

In Eq. (4.41), the first term is 0, while the second term is positive (note that, by Eq. (4.36), $\beta_m = 1$). \square

We define two types of threshold equilibria:

Definition 8. In a *none-make-AR* equilibrium, none of the customers, regardless of their arrival times, choose AR.

Definition 9. In a *some-make-AR* equilibrium, all customers follow a threshold strategy with a set of thresholds $0 < \tau_N^e < \tau_{N-1}^e < \dots < \tau_1^e < 1$.

We define the critical cost

$$\underline{C} \triangleq 1 - \pi_N(0). \quad (4.42)$$

We are now ready to state the main result of this section:

Theorem 8. *Given any number of servers N and any demand distribution D :*

- *If $C < \underline{C}$, there exists a unique some-make-AR equilibrium.*
- *If $C \geq \underline{C}$, there exists a unique none-make-AR equilibrium.*

Proof. From Theorem 7, we know that, at equilibrium, all customers follow a threshold strategy. First, we show that, at equilibrium, all thresholds are smaller than 1 (i.e., an equilibrium where all customers choose AR regardless of their arrival time does not exist). Let assume by contradiction that, for some $n \in \{1, N\}$, $\sigma(t, n) = 1$ for all $t \in [0, 1]$. In this case, the n -th threshold customer (which arrives at time 1) knows that all customers that arrives earlier chose AR and since no customer will arrive after him, his service is granted. That is,

$$\pi_n(1) = 1, \quad \forall n \in \{1, \dots, N\}. \quad (4.43)$$

Hence, the n -th threshold customer is better off choosing AR' which leads to a contradiction.

Next, assume that $C < \underline{C}$. From Lemma 11, we deduce that, in this case, there is exactly one set of thresholds $0 < \tau_N^e < \tau_{N-1}^e < \dots < \tau_1^e < 1$ such that, for any $n \in \{1, N\}$, $1 - C = \pi_n(\tau_n^e)$. We will now show this set represents a *some-make-AR* equilibrium. Assume that all customers follow this set of thresholds. Consider a customer with arrival time $t < \tau_n^e$ that observes n available servers. His probability of service is smaller than $\pi_n(t)$ since there is a positive probability that reservations will be made between t and τ_n^e , while $\pi_n(t)$ is the probability of service given that no reservations will be made after time t . From the first part of Lemma 11, we know

that $\pi_n(\tau_n^e) > \pi_n(t)$. Thus, his probability of service is smaller than the probability of service of the n -th threshold customer and he is better off choosing AR .

Now, consider a customer with arrival time $t > \tau_n^e$ that observes n available servers. Since all customers that arrive later than τ_n^e and observe n available servers choose AR' , observing n available servers at time t is equivalent to observing n available servers at time τ_n^e . Thus, this customer (just like the n -th threshold customer) is indifferent between the two actions and has no motivation to deviate. We conclude that this set of thresholds represents a *some-make-AR* equilibrium. A *none-make-AR* equilibrium does not exist, since if all customers choose AR' , their expected payoff is $1 - \underline{C}$ which is smaller than $1 - C$.

Finally, assume that $C \geq \underline{C}$. If all customers that observe N available servers choose AR' , then they all have the same probability of service $1 - \underline{C}$ which is greater than $1 - C$. Thus, none will deviate and *none-make-AR* is an equilibrium. A *some-make-AR* equilibrium does not exist, since there is no value of τ_N for which $1 - C = \pi_N(\tau_N)$. \square

The two types of equilibrium are illustrated in Figure 4.1

Example 4. *Continuing Example 3, we substitute Eq. (4.33) into Eq. (4.42) and get $\underline{C} = 1 - 5/6$. Thus, any C larger than $1/6$ leads to a none-make-AR equilibrium. Next, assume $C = 0.1$. By solving the equalities $\pi_2(\tau_2) = 1 - 0.1$ and $\pi_1(\tau_1) = 1 - 0.1$, using Eq. (4.33) and Eq. (4.34), we obtain that the thresholds at equilibrium are $\tau_1^e = 0.917$ and $\tau_2^e = 0.453$. That is, if the first customer making an inquiry arrives before $t = 0.453$, then he makes a reservation. If the demand is larger than 1 and the second customer making an inquiry arrives before 0.917, then the second server will also be reserved.*

An example with Poisson distributed demand with mean 10 and 6 servers. The




Figure 4.1: An example of a system with 6 servers and Poisson distributed demand with mean $\lambda = 10$. The equilibrium is determined by the expected payoff of a customer that arrives at the 6-th threshold and observes 6 available servers. If that customer is better off not making AR (a), then there is a none-make-AR equilibrium. If there is a threshold such that the customer is indifferent between the two actions (b), then there is a some-make-AR equilibrium.

figures show the expected payoff of a customer that observes an empty system and arrives exactly at the 6-th threshold. If, for any threshold, the customer is better off not making AR (a), then there is a none-make-AR equilibrium. If there is a threshold such that the customer is indifferent between the two actions (b), then there is a some-make-AR equilibrium.

4.2.4 Poisson Distributed Demand

In the previous section, we studied a game with demand that follows a general distribution. In this section, we apply the result on a system with Poisson distributed demand with mean λ . As was showed in (Myerson, 1998), in Poisson games (i.e., games in which the number of players is Poisson distributed) the distributions of D and \tilde{D} are identical. Furthermore, if players are randomly ascribed to different types with fixed probabilities, the number of players of each type is independent of the number of players of other types and is also Poisson distributed. Hence, from the perspective of the n -th threshold customer, the number of customers that arrive

after him is a Poisson random variable with parameter $\lambda(1 - \tau_n)$. If he chooses AR' , then all those customers also choose AR' . Using this property, we can express the probability of service in a simpler way than in Eq. (4.29), that is

$$\pi_n(\tau_n) = \sum_{i=0}^{N-1} \mathbb{P}(\tilde{D}_{AR'}=i) + \sum_{i=N}^{\infty} \mathbb{P}(\tilde{D}_{AR'}=i) \frac{n}{i+1}. \quad (4.44)$$

The first term is the probability that the number of customers choosing AR' is smaller than N . In this case, service is guaranteed to all customers. The second term is the probability to get service when the demand exceeds N . In this case, the n unreserved servers will be arbitrarily allocated to customers that chose AR' . We substitute

$$\mathbb{P}(\tilde{D}_{AR'}=i) = e^{-\lambda(1-\tau_n)} \frac{(\lambda(1-\tau_n))^i}{i!} \quad (4.45)$$

into Eq. (4.44) and obtain a closed form expression for $\pi_n(\tau_n)$, namely

$$\begin{aligned} \pi_n(\tau_n) &= e^{-\lambda(1-\tau_n)} \sum_{i=0}^{n-1} \frac{(\lambda(1-\tau_n))^i}{i!} + e^{-\lambda(1-\tau_n)} \sum_{i=n}^{\infty} \frac{(\lambda(1-\tau_n))^i}{i!} \frac{n}{i+1} \\ &= e^{-\lambda(1-\tau_n)} \sum_{i=0}^{n-1} \frac{(\lambda(1-\tau_n))^i}{i!} + \frac{n}{\lambda(1-\tau_n)} e^{-\lambda(1-\tau_n)} \sum_{i=n+1}^{\infty} \frac{(\lambda(1-\tau_n))^i}{i!} \\ &= e^{-\lambda(1-\tau_n)} \sum_{i=0}^{n-1} \frac{(\lambda(1-\tau_n))^i}{i!} + \frac{n}{\lambda(1-\tau_n)} \left(1 - e^{-\lambda(1-\tau_n)} \sum_{i=0}^n \frac{(\lambda(1-\tau_n))^i}{i!} \right). \end{aligned} \quad (4.46)$$

Using the upper incomplete Gamma function

$$\Gamma[s, x] \triangleq (s-1)! e^{-x} \sum_{k=0}^{s-1} \frac{x^k}{k!}, \quad (4.47)$$

we obtain

$$\pi_n(\tau_n) = \frac{n}{\lambda(1 - \tau_n)} + \frac{\Gamma[n, \lambda(1 - \tau_n)]}{(n - 1)!} - \frac{\Gamma[1 + n, \lambda(1 - \tau_n)]}{\lambda(1 - \tau_n)(n - 1)!}. \quad (4.48)$$

Example 5. Consider a game with $N = 6$ servers, average demand $\lambda = 6$ and reservation cost $C = 0.15$. To find the equilibrium strategy, we first check which type of equilibrium prevails. Since $1 - \pi_6(0) = 0.16 > 0.15$, we deduce that the game has a some-make-AR equilibrium (see Theorem 8). We then solve the following set of equations:

$$1 - 0.15 = \frac{n}{6\tau_n} + \frac{\Gamma[n, 6(1 - \tau_n)]}{(n - 1)!} - \frac{\Gamma[1 + n, 6(1 - \tau_n)]}{6(1 - \tau_n)(n - 1)!}, \quad \forall n \in \{1, 2, 3, 4, 5, 6\}. \quad (4.49)$$

The solution is the set of thresholds $\tau^e = \{0.944, 0.787, 0.605, 0.416, 0.223, 0.027\}$. That is, at equilibrium, a customer that observes an empty system (i.e., $n = N$) will choose AR only if he arrives before $t = 0.027$, while a customer that observes one available server will choose AR if he arrives before $t = 0.944$. Figure 4.2 shows the intersection of $1 - C$ with the functions $\pi_n(\tau_n)$ for $n = 1, 2, \dots, 6$.

In this example, the probability that all customers arrive after τ_6^e , and hence all customers choose AR' is

$$\sum_{i=0}^{\infty} \mathbb{P}(D=i)(1 - \tau_6^e)^i = e^{-6} \sum_{i=0}^{\infty} \frac{6^i}{i!} (0.973)^i = 0.850. \quad (4.50)$$

If the provider does not share information about the number of available servers, then under the same parameters there is a unique equilibrium and the probability that all customers choose AR' is smaller than 0.01.

The example illustrates the drawback of the *full-information* policy: even if a




Figure 4.2: Illustration of Example 5. For each number of available servers $i \in 1, 2, \dots, 6$, the probability of service π_i as a function of the threshold τ (which represents the expected payoff of a threshold customer that is not making AR) intersects once with the line $1 - C$ (which represents the payoff of any customer making AR), thus showing the uniqueness of the equilibrium. As the number of available servers i decreases, the time threshold at which customers stop making AR increases (i.e., customers are more likely to make AR when observing fewer available servers).

some-make-AR equilibrium exists, there is a high chance that no customer will make a reservation.

4.3 Partial-information Model

We next consider a *partial-information* policy. Under this policy, customers are informed about the number of available servers only if this number is below or equal to some threshold denoted M , where $M \in \{1, 2, \dots, N-1\}$. The drawback of this policy is that an uninformed customer can deduce that the number of available servers is larger than M . In this section, we analyze the equilibrium structure of the *partial-information* policy, assuming that customers know that this policy is used. Note that if $M = N - 1$, then an uninformed customer knows that all servers are avail-

able, and therefore a *partial-information* policy with $M = N - 1$ is equivalent to the *full-information* policy.

A result similar to Theorem 7 also holds for the *partial-information* game. Thus, this game has the same two types of equilibria (i.e., *none-make-AR* and *some-make-AR*), as the *full-information* game. In this game, all uninformed customers use the same threshold, denoted by τ_u , which replaces the thresholds $\{\tau_{M+1}, \dots, \tau_N\}$. Thus, the set of thresholds now contains $M + 1$ thresholds. Eq. (4.29), which describes the probability of service of the threshold customers in the *full-information* model, is still valid for the informed customers. In order to find the threshold followed by uninformed customers, we consider an *uninformed threshold customer*. As earlier, this represents a virtual uninformed customer that arrives exactly at the threshold followed by uninformed customers.

We denote the probability of service of the uninformed threshold customer by $\pi_u(\tau_u)$. We express it, using the law of total probability, by conditioning on the number of available servers $i \in \{M + 1, M + 2, \dots, N\}$, given that at least $N - M$ servers are available. As we showed in the previous section, the probability of service of the i -th threshold customer is independent of all other thresholds. Thus, we can construct $\pi_u(\tau_u)$ using $\pi_i(\cdot)$ as defined in Eq. (4.29), namely

$$\pi_u(\tau_u) = \sum_{i=M+1}^N \pi_i(\tau_u) \mathbb{P}(\tilde{D}_{AR} = N - i | \tilde{D}_{AR} < N - M). \quad (4.51)$$

The distribution of \tilde{D}_{AR} can be found using the law of total probability conditioned

on \tilde{D} . Given \tilde{D} , \tilde{D}_{AR} is a binomial random variable with success probability τ_u . Thus,

$$\begin{aligned}\mathbb{P}(\tilde{D}_{AR}=i) &= \sum_{j=i}^{\infty} \mathbb{P}(\tilde{D}=j)\mathbb{P}(\tilde{D}_{AR}=i|\tilde{D}=j) \\ &= \sum_{j=i}^{\infty} \mathbb{P}(\tilde{D}=j) \binom{j}{i} \tau_u^i (1-\tau_u)^{j-i}.\end{aligned}\quad (4.52)$$

Next, we show that, unlike the probability of service of an informed threshold customer, $\pi_u(\tau_u)$ is not necessarily an increasing function of the threshold τ_u . In this case, the equality $\pi_u(\tau_u) = 1 - C$ may hold for more than one value of τ_u and the equilibrium will not be unique.

We prove our claim by the means of an example. Assume that the demand is Poisson distributed with parameter λ . As in the previous section, we use the properties of Poisson games to simplify the expression of $\pi_u(\tau_u)$. If all uninformed customers follow a threshold-strategy with threshold τ_u , then the number of customers choosing AR (i.e., \tilde{D}_{AR}) and the number of customers choosing AR' (i.e., $\tilde{D}_{AR'}$) are independent Poisson random variables with parameters $\lambda\tau$ and $\lambda(1-\tau_u)$, respectively. The probability of service of the uninformed threshold customer is

$$\pi_u(\tau_u) = \sum_{i=0}^{N-M-1} \mathbb{P}(\tilde{D}_{AR}=i|\tilde{D}_{AR}<N-M) \left(\sum_{j=0}^{N-i-1} \mathbb{P}(\tilde{D}_{AR'}=j) + \sum_{j=N-i}^{\infty} \mathbb{P}(\tilde{D}_{AR'}=j) \frac{N-i}{j+1} \right), \quad (4.53)$$

which is the sum of the probabilities that i customers choose AR , each multiplied by the probability of service given i . The inner probabilities are

$$\mathbb{P}(\tilde{D}_{AR}=i|\tilde{D}_{AR}<N-M) = \frac{\frac{(\lambda\tau_u)^i}{i!}}{\sum_{k=0}^{N-M-1} \frac{(\lambda\tau_u)^k}{k!}}, \quad (4.54)$$

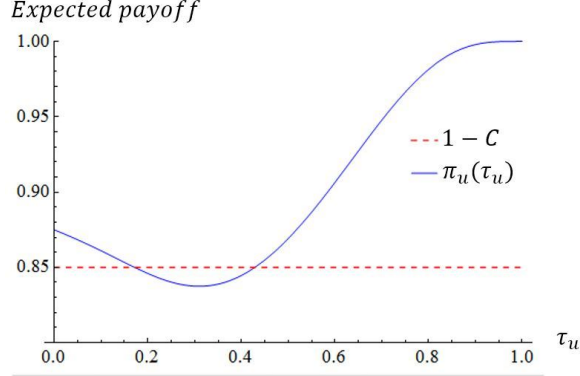


Figure 4.3: In example 6, the probability of service of the uninformed threshold customer as a function of the threshold $\pi_u(\tau_u)$ intersects twice with the line 0.85, which is the payoff if making AR.

and

$$\mathbb{P}(\tilde{D}_{AR'} = j) = \frac{e^{-\lambda(1-\tau_u)}(\lambda(1-\tau_u))^j}{j!}. \quad (4.55)$$

The following example shows that a game with Poisson distributed demand can have several equilibria.

Example 6. We consider a system with the following parameters $\lambda = 10$, $N = 10$, $C = 0.15$ and $M = 3$. From Eq. (4.53) - Eq. (4.55), we get (as illustrated in Figure 4.3) $\pi_u(0.429) = \pi_u(0.172) = 1 - 0.15$. Using Eq. (4.48), we find the n -th threshold for $n = 1, 2, 3$ and conclude that both $\{\tau_1 = 0.966, \tau_2 = 0.872, \tau_3 = 0.763, \tau_u = 0.172\}$ and $\{\tau_1 = 0.966, \tau_2 = 0.872, \tau_3 = 0.763, \tau_u = 0.429\}$ represent equilibria strategies.

After showing that multiple equilibria may exist, we continue studying the equilibrium structure of the game. First, we observe that if all customers choose AR' , then the system is always empty and the information sharing policy has no effect on the game's outcome. Thus, when the reservation cost is larger than \underline{C} (which is de-

defined in Eq. (4.42)), the *partial-information* game has a *none-make-AR* equilibrium, similar to the *full-information* game. However, in the *partial-information* game, the equilibrium is not necessarily unique.

In order to have a unique *none-make-AR* equilibrium, we must have

$$1 - C < \pi_u(\tau_u), \quad \forall \tau_u \in [0, 1]. \quad (4.56)$$

We define

$$\bar{C} = 1 - \inf_{0 \leq \tau_u \leq 1} \pi_u(\tau_u), \quad (4.57)$$

and conclude that any cost above \bar{C} leads to a unique *none-make-AR* equilibrium.

Next, we observe that if all customers choose *AR*, then the probability of service of an uninformed threshold customer (i.e., a customer that arrives at time $t = 1$) is 1. Thus, for any cost $C < \bar{C}$, there is at least one value of τ_u such that

$$1 - C = \pi_u(\tau_u). \quad (4.58)$$

The number of *some-make-AR* equilibria is determined by the number of values of τ_u for which Eq. (4.58) holds. If \bar{C} is strictly smaller than \underline{C} , then $\pi_u(\tau_u)$ is not a monotonic function. Hence, any C in the range (\underline{C}, \bar{C}) has at least two *some-make-AR* equilibria. The following theorem summarizes the equilibria structure of the *partial-information* game:

Theorem 9. *In the partial-information game, there exist quantities \underline{C} and $\bar{C} \geq \underline{C}$, such that*

- *If $0 < C < \underline{C}$, there is at least one some-make-AR equilibrium.*




Figure 4-4: The expected payoff of the uninformed threshold customer as a function of the threshold. The number of times the line $1 - C$ intersects this function determines the set of equilibria.

- If $\underline{C} < C < \bar{C}$, there is a none-make-AR equilibrium and at least two some-make-AR equilibria.
- If $C > \bar{C}$, none-make-AR is the unique equilibrium.

To simplify the presentation, the boundary cases $C = \underline{C}$ and $C = \bar{C}$ are ignored. In Figure 4-4, we illustrate how the AR cost defines the set of equilibria. We use the same system as in Example 6.

4.4 Comparison of Information Sharing Policies

In this section, we assume that the provider is interested to persuade as many customers as possible to make reservations. We resort to simulations to find out which policy maximizes the average number of reservations. We consider the *full-information*, *partial-information* policies and the *binary-information* policy (which has the same outcome as the *no-information* policy). Procedure 1 details the simu-

Procedure 1 Information-sharing Simulation(N, τ^e, P_D)

```

 $j \leftarrow 1$  {index}
 $D_{AR} \leftarrow 0$  {number of reservations}
 $D \leftarrow$  random variable from  $P_D$  {the demand}
 $\mathbf{t} \in \mathbb{R}^D \leftarrow$  vector with random values from  $U[0, 1]$  {arrival times}
Sort  $\mathbf{t}$  in ascending order {sorting from the first customer that makes an inquiry
to the last}
while  $j \leq \min\{D, N\}$  do {iterate over all customers or until all servers are re-
served}
    if  $t_j < \tau_{N-j+1}^e$  then {if the arrival time is smaller than the appropriate thresh-
old}
         $\tilde{D}_{AR} \leftarrow \tilde{D}_{AR} + 1$  {choosing  $AR$ }
         $j \leftarrow j + 1$  {increase the index by 1}
    else
        break {if a customer is better off choosing  $AR'$ , then all customers that arrive
after him are also better off choosing  $AR'$ }
    end if
end while
return  $D_{AR}$ 

```

lation steps at each iteration.

4.4.1 Simulation Results

We first examine an overloaded system where the demand is Poisson distributed with mean $\lambda = 15$, while the number of servers is $N = 10$. Using Eq. (4.48), we get that $\underline{C} = 1 - \pi_N(0) = 0.342$. We evaluate the performance of the difference policies for seven different reservation costs between 0 and \underline{C} .

For all costs in this range, all policies have a unique *some-make-AR* equilibrium. Each point is computed by averaging out results over 10,000 iterations. As shown in Figure 4-5(a), for each AR cost, the policy that maximizes the number of reservations is the *binary-information* policy. The *full-information* policy performs the worst, and the performance of the *partial-information* policy performs better as M (the threshold at which information starts to be shared) decreases.

The simulation results also indicate that the gap between the outcomes of the different policies increases as the AR cost increases. When the cost is low, the motivation to make reservations is high and almost all servers are reserved, regardless of the policy. As the cost increases, the different information sharing policies have larger impact on customers' decisions. For example, when the cost is $C = 0.34$, the average number of reservation is 20 times higher with the *binary-information* policy than with the *full-information* policy.

Next, we consider an underloaded system with mean demand $\lambda = 8$. This time, we use six different prices between 0 and $\underline{C} = 0.053$. Figure 4.5(b) shows similar results as Figure 4.5(a).



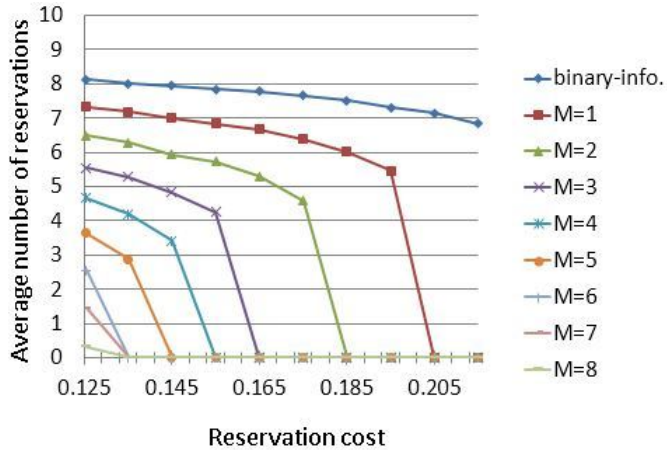
(a) Average demand $\lambda = 15$.(b) Average demand $\lambda = 8$.(c) Average demand of $\lambda = 10$.

Figure 4-5: A comparison between the number of customers making AR in a system with $N = 10$ servers. As less information is shared, more customers are likely to make AR.

Finally, we set the mean demand to $\lambda = 10$ and explore the case of multiple equilibria (i.e., AR costs larger than \underline{C}). In this case, the *full-information* policy yields zero reservations (see Theorem 8). We run simulation to compare between the performance of the *binary-information* and the *partial-information* policies. We choose different AR costs in the range $[0.215, 0.228]$. In order to compare performance when multiple equilibria exist, we always choose the one that leading to the largest expected number of reservations. As shown in Fig. 4.5(c), the results follow the same pattern as in the previous simulations, namely as less information is shared, the number of reservations increases.

4.4.2 Numerical Example

To better understand the gap between the outcomes of the *full-information* and *binary-information* policies, we return to Example 3.

Example 7. Consider the demand distribution of Example 3. By Eq. (4.3), the probability of service of the threshold customer in the binary-information model is

$$\pi_{no-info}(\tau) = \frac{3 + 3(1 - \tau) - (1 - \tau)^2}{3 + 6(1 - \tau) - 3(1 - \tau)^2}. \quad (4.59)$$

In Figure 4.6, we plot $\pi_2(\tau)$, $\pi_1(\tau)$ and $\pi_{no-info}(\tau)$. We can see in the graph that for any AR cost in $(0, 0.167)$, at equilibrium, τ^e (which is the fraction of customers making AR in a game with binary-information) is much closer to τ_2^e than to τ_1^e . That means that when sharing information, we dramatically decrease the probability that the first customer making an inquiry will choose AR, while we only slightly increase the probability that the second customer will choose AR (assuming the first one chose AR). Thus, we can expect that the average number of reservations will be higher when binary-information (or no-information) is shared.




Figure 4-6: The probability of service of the threshold customers as a function of the thresholds. The fact that τ^e is much closer to τ_2^e than to τ_1^e explains why the number of reservations is larger with the *binary-information* policy than with the *full-information* policy.

Based on the simulation results and this example, we make the following conjecture:

Conjecture 10. *In AR games, sharing no information about the number of available servers maximizes the average number of reservations.*

4.5 Summary

In this section, we focus on the impact of information on the strategic behavior of customers in a slotted loss queue that allows advance reservations. We first show that if the AR cost is applied only to customers getting service, then informing customers that servers are available or hiding this information lead to the same equilibria. Then, we find the equilibria structure under a *full-information* policy is being used. We show that customers are less likely to make AR as they observe more available servers. In order to lower the probability that the system will stay empty, we suggest the *partial-information* policy and find its equilibrium structure. Using simulation, we show that indeed, the number of customers making AR increases when using the

partial-information policy competing to *full-information* policy. However, the policy that maximizes the fraction of customers making AR is the *binary-information* policy. From the simulation results, we also infer that the customers' behavior is more sensitive to changes in the information disclosure policy as the reservation cost increases.

The results indicate that in order for a full-information policy to be beneficial to the service provider, the provider must encourage customers to make reservations also when all or almost all servers are available. For example, in entertainment events, customers have such motivation since they can choose better seats as they buy the ticket earlier.

Chapter 5

Waiting Queues with Advance Reservations

In this chapter, we switch our focus from loss queues to waiting queues. As in the previous chapters, our objective is twofold. First, we study the behavior of strategic customers and find the equilibrium structure of the game. Then, we analyze the problem from the perspective of a service provider aiming to maximize her revenue from AR fees.

5.1 Game Description

We consider a preemptive-resume $M/D/1$ queue that supports advance reservations. In our model, there is a reservation period which covers $[-T, 0]$ and is open for reservations from time $-T$ until time 0. Each customer $k = 1, 2, \dots$ is associated with a request time $-T \leq t_k \leq 0$ and a desired service starting time (shortly noted as arrival time) $s_k > 0$. That is, if $t_1 < t_2$, then customer 1 has the opportunity to reserve the server before customer 2. If $s_1 < s_2$, then customer 1 wants to be served earlier than customer 2. The service period starts only after the reservation period ends. The request time can be interpreted as how much time in advance a customer realizes that he will need service at a future time point.

The request times are derived from a general continuous distribution with cumulative distribution function $F(\cdot)$. The arrivals follow a Poisson process with rate λ . The service time is $1/\mu$ and we assume that $\mu > \lambda$.

Each customer, at his request time, decides whether to make a reservation or not. We denote those two actions by AR and AR' , respectively. If a customer makes a reservation but his desired service time is already reserved, the nearest future available time will be reserved for him. A customer that does not make a reservation is served on a first-come-first-served basis along periods of times over which the server is not reserved.

Customers do not know in advance what will be their waiting time if making or if not making AR . Their decisions are based on statistical information which is the values of λ , μ and F .

The cost of each customer consists of the reservation cost C (if making AR) and the cost of waiting which is a linear function of the waiting time. Without loss of generality, we assume that the cost of waiting is equal to the waiting time. Note that the waiting time when making AR is smaller than when not making AR . However, it may be greater than zero, since it is possible that the server is already reserved at the desired service time. For simplification, we assume that the service period is long enough such that we can ignore the transient phase before the queue reaches its steady state.

In a preemptive-resume queue, if a job is interrupted, then it later resumes and is not restarted. Due to this property, if the server is idle and a customer is waiting for service, he will be served even if service cannot be completed due to an existing reservation (in this case his service will be preempted and later resumed). Hence, supporting advance reservations in a preemptive-resume queue does not impact the utilization of the server which is $\rho = \lambda/\mu$. Figure 5.1 illustrates the model.




Figure 5.1: An illustration of the model with three customers. Customer 1 makes a reservation at time t_1 , and he is served upon arrival at s_1 . Customer 2 also makes AR and is served upon arrival, but his service is preempted by customer 1 which made reservation earlier. Customer 3 is served only when the service of customer 2 is completed.

5.2 Equilibrium Analysis

We can analyze this system as a priority queue where a priority between 0 (lowest priority) and 1 (highest priority) is assigned to each customer. A customer with request time t has priority 0 if not making AR and priority $p = 1 - F(t)$ if making AR. Customers that share the same priority are served on a first-come-first-served basis. We refer to p as the *potential priority*. Due to the probability integral transformation theorem (Dodge, 2006, p. 320), we know that p is a random variable, uniformly distributed in $[0, 1]$.

Since customers are statistically identical, we consider only symmetrical behavior. Thus, a decision of a tagged customer is a mapping of his potential priority p to the probability of making AR. We denote this strategy function by $\sigma(p)$. Consider a tagged customer with potential priority p . We define $W(\cdot)$ to be a mapping of the strategy followed by the rest of the customers and the priority of the tagged customer to his expected waiting time. Thus, the expected waiting time of the tagged customer is $W(\sigma, p)$ if he makes AR and $W(\sigma, 0)$ otherwise. Since customers are strategic, a customer with potential priority p will make AR only if

$$W(\sigma, p) + C \leq W(\sigma, 0). \quad (5.1)$$

Next, we define a threshold strategy and show that this is the only strategy that can lead to equilibria.

Definition 10. Let $\tau \in (0, 1]$. A strategy function $\sigma(p)$ is said to be a threshold strategy if it satisfies

$$\sigma(p) = \begin{cases} 1 & \text{if } p > \tau, \\ 0 & \text{if } p \leq \tau. \end{cases}$$

Lemma 12. At equilibrium, all customers follow a threshold strategy.

Proof. Consider any strategy function σ . Since the expected waiting time is non-increasing with the priority, there is either a single potential priority, or an interval of potential priorities, or no potential priority such that

$$W(\sigma, p) + C = W(\sigma, 0). \quad (5.2)$$

Note that the left hand side of Eq. (5.2) is the expected cost if making AR, while the right hand side is the expected cost if not making AR. If Eq. (5.2) holds for a single value p' , then a customer with potential priority greater (respectively, smaller) than p' is better off making (respectively, not making) AR. Therefore, σ is an equilibrium strategy only if it is a threshold strategy with threshold $\tau = p'$.

If Eq. (5.2) holds for an interval of values $[p', p'']$, then all customers with potential priority $p \in [p', p'']$ do not make AR (otherwise, $W(\sigma, p)$ would not be a constant over that interval). Therefore, σ is an equilibrium strategy only if it is a threshold strategy, with threshold $\tau = p'$.

Finally, suppose that Eq. (5.2) does not hold for any $p \in [0, 1]$. If $W(\sigma, p) + C < W(\sigma, 0)$ for all $p \in [0, 1]$, then all customers are better off not making AR. Therefore, σ is an equilibrium strategy only if it is a threshold strategy, with threshold $\tau = 1$.

Note that a situation where $W(\sigma, p) + C > W(\sigma, 0)$ for all $p \in [0, 1]$ does not exist,

since a user with potential priority zero has the same expected waiting time if making or avoiding AR. \square

Next, we define two types of equilibria.

Definition 11. *An equilibrium strategy with threshold τ is called a some-make-AR equilibrium if $\tau < 1$.*

Definition 12. *An equilibrium strategy with threshold τ is called a none-make-AR equilibrium if $\tau = 1$.*

Since the structure of the equilibrium depends on the reservation cost, we aim to determine the equilibrium to which a given cost leads. Given that all customers follow a threshold strategy, we define a *threshold customer* to be a customer with potential priority equals exactly to the threshold followed by all other customers.

Given a strategy with threshold τ , a threshold customer that makes AR observes three priority classes:

1. A lower priority class which contains all customers with priority smaller than his (none of them makes AR). The arrival rate of customers belonging to this class is $\lambda\tau$.
2. His own priority class which contains only himself (since the potential priority is a continuous random variable, the probability that two customers will have the same potential priority is zero). Thus, the arrival rate of customers belonging to this class is 0.
3. A higher priority class which contains all customers with greater priority (they all made AR before he did). The arrival rate of customers belonging to this class is $\lambda(1 - \tau)$.

If the threshold customer does not make AR, he only observes two classes:

1. His own priority class which also contains all customers with priority smaller than his. The arrival rate of customers belonging to this class is $\lambda\tau$.
2. A higher priority class which contains all customers with greater priority. The arrival rate of customers belonging to this class is $\lambda(1 - \tau)$.

By applying the known formula of the waiting time in an $M/G/1$ queue with preemptive-resume priorities (Conway et al., 2012, p.175), we obtain the following:

1. If the threshold customer makes AR, then his expected waiting time is

$$W_{AR}(\tau) = \frac{\mu - \frac{\lambda}{2}(1 - \tau)}{(\mu - \lambda(1 - \tau))^2} - \frac{1}{\mu}. \quad (5.3)$$

2. If the threshold customer does not make AR, then his expected waiting time is

$$W_{AR'}(\tau) = \frac{\mu - \frac{\lambda}{2}}{(\mu - \lambda(1 - \tau))(\mu - \lambda)} - \frac{1}{\mu}. \quad (5.4)$$

The condition for threshold $\tau < 1$ to be a *some-make-AR* equilibrium is

$$C + W_{AR}(1) = W_{AR'}(1). \quad (5.5)$$

That is, a customer with potential priority equals to the threshold is indifferent between the two actions. The condition for threshold $\tau = 1$ to be a *none-make-AR* equilibrium is

$$C + W_{AR}(\tau) \geq W_{AR'}(\tau). \quad (5.6)$$

That is, a customer with potential priority 1 (and hence, all customers) are better off not making AR.

By isolating C in Eq. (5.5), we define $C(\tau)$ to be a function that maps a threshold to the AR cost that leads to that threshold

$$C(\tau) \triangleq \frac{\lambda \cdot \mu \cdot \tau}{2(\mu - \lambda) \cdot (\mu - \lambda(1 - \tau))^2}. \quad (5.7)$$

We conclude that given cost C , the threshold $\tau_e \in (0, 1)$ represents a *some-make-AR* equilibrium if and only if $C = C(\tau_e)$. The threshold $\tau_e = 1$ represents a *none-make-AR* equilibrium if and only if $C \geq C(1)$. In order to find the equilibria structure, we next find the properties of $C(\tau)$.

Lemma 13. *If $\rho \leq 1/2$, then $C(\tau)$ is a monotonically increasing function. If $\rho > 1/2$, then $C(\tau)$ is a unimodal function with a global maximum.*

Proof. First, we compute the derivative of $C(\tau)$:

$$\frac{dC}{d\tau} = \frac{\lambda\mu(\lambda(1 + \tau) - \mu)}{2(\lambda - \mu)(\mu - \lambda(1 - \tau))^3}. \quad (5.8)$$

Since the denominator is negative for any τ , the sign of the derivative is determined by the sign of $\lambda(1 + \tau) - \mu$. If $\rho \leq 1/2$, then this expression is negative for any $\tau \in (0, 1)$ and the derivative of $C(\tau)$ is positive for any $\tau \in (0, 1)$. If $\rho > 1/2$, then the derivative of $C(\tau)$ is positive for any $\tau < (\mu - \lambda)/\lambda$; is equal to zero at $\tau = (\mu - \lambda)/\lambda$; and negative otherwise. Thus, for any value of $\rho > 1/2$, $C(\tau)$ is unimodal with a global maximum. □

Next, we define:

$$\underline{C} \triangleq C(1) = \frac{\lambda}{2\mu(\mu - \lambda)}, \quad (5.9)$$

and

$$\overline{C} \triangleq \frac{\mu}{8(\lambda - \mu)^2}. \quad (5.10)$$

Note that if $\rho \leq 0.5$, then \underline{C} is the maximum value of $C(\tau)$ and if $\rho > 0.5$, then \overline{C} is the maximum value of $C(\tau)$. We can now state the main result of this chapter:

Theorem 11. *The game has the following equilibrium structure.*

When $\rho < 1/2$:

- If $C < \underline{C}$, then there is a unique *some-make-AR* equilibrium.
- If $C > \underline{C}$, then there is a unique *none-make-AR* equilibrium.

When $\rho > 1/2$:

- If $C < \underline{C}$, then there is a unique *some-make-AR* equilibrium.
- If $\underline{C} < C < \overline{C}$, then there are two *some-make-AR* equilibria and a *none-make-AR* equilibrium.
- If $C > \overline{C}$, then there is a unique *none-make-AR* equilibrium.

Proof. We begin with $\rho \leq 0.5$. If $C < \underline{C}$, then there is a single value of τ such that $C = C(\tau)$ has a solution. Hence, there is one *some-make-AR* equilibrium. A *none-make-AR* equilibrium does not exist since $C(1) > C$. If $C > \underline{C}$, then there is no value of τ such that $C = C(\tau)$ has a solution. Hence, a *some-make-AR* equilibrium does not exist. On the other hand, a *none-make-AR* equilibrium exists since $C > C(1)$.

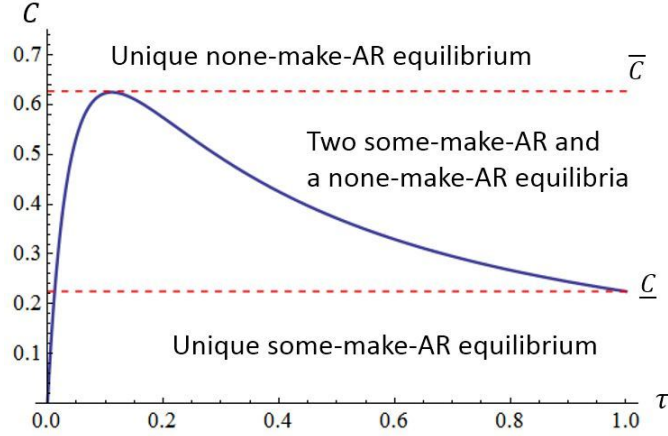


Figure 5.2: $C(\tau)$ when $\rho > 0.5$. The number and types of equilibria is determined by the value of C

Next, consider $\rho > 0.5$. In the range $[0, \underline{C}]$, the function $C(\tau)$ is monotonically increasing. Thus, if $C \in [0, \underline{C}]$, then there is a single value of τ such that $C = C(\tau)$ has a solution, and hence there is one *some-make-AR* equilibrium. In the range $[\underline{C}, \bar{C}]$, the function $C(\tau)$ is unimodal. Thus, if $C \in [\underline{C}, \bar{C}]$, then there exist two values of τ that solve $C = C(\tau)$, and hence there are two *some-make-AR* equilibria. The condition for the existence of a *none-make-AR* equilibrium is the same as in the case of $\rho \leq 0.5$. \square

In Figure 5.2, we plot $C(\tau)$ as defined in Eq. (5.7) for a queue with parameters $\lambda = 18$ and $\mu = 20$. Given the AR cost C , the equilibrium structure is determined by the number of times C intersects the function $C(\tau)$.

5.3 Revenue Maximization

In this section, we assume that the AR cost is a fee determined by the service provider. Our goal is to find which fee maximizes the provider revenue from AR fees. The revenue per time unit, at equilibrium with threshold τ_e , is the number of customers

making AR multiplied by the AR fee that leads to that equilibrium. The expected revenue is

$$\begin{aligned} R(\tau_e) &= \lambda(1 - \tau_e)C(\tau_e) \\ &= \frac{\lambda^2(1 - \tau_e)\tau_e\mu}{2(\mu - \lambda)(\lambda(\tau_e - 1) + \mu)^2}. \end{aligned} \quad (5.11)$$

With some manipulation, we get that the revenue function does not depend on the values of λ and μ but only on the utilization ρ :

$$R(\tau_e) = \frac{\rho(1 - \tau_e)\tau_e}{2(1 - \rho)(1 + \rho(\tau_e - 1))^2}. \quad (5.12)$$

At first glance, this result seems surprising since it implies that the revenue does not increase when scaling the system (i.e., increasing both arrival and service rates). However, in an $M/D/1$ queue, the waiting time decreases as the system gets larger, and hence customers are less motivated to make AR. Therefore, scaling the system has a trade off. For a given threshold, as we scale the system, more customers will make AR but they will pay a smaller fee.

By solving the equation $dR/d\tau_e = 0$, we find that the optimal threshold is $\tau_e^* = (1 - \rho)/(2 - \rho)$. By substituting τ_e^* into Eq. (5.7), we get that the optimal fee is

$$C^* = \frac{\lambda(2\mu - \lambda)}{8\mu(\mu - \lambda)^2}. \quad (5.13)$$

Similarly, by substituting τ_e^* into Eq. (5.12), we get that the maximum possible

revenue is

$$R^* = \frac{\rho^2}{8(1-\rho)^2}. \quad (5.14)$$

If $\rho \leq 0.5$, then setting the AR fee to $C = C^*$ maximizes the provider revenue. However, if $\rho > 0.5$, this fee may lead to multiple equilibria, including one with no reservations. From Theorem 11, a *none-make-AR* equilibrium exists if $C^* > \underline{C}$. From Eq. (5.9) and Eq. (5.13), we deduce that if

$$\frac{\lambda(2\mu - \lambda)}{8\mu(\mu - \lambda)^2} > \frac{\lambda}{2\mu(\mu - \lambda)}, \quad (5.15)$$

then C^* leads to multiple equilibria (including a *none-make-AR* equilibrium). One can show that the inequality above holds only if $\rho > 2/3$. We conclude with the following corollary, which is illustrated in Figure 5.3.

Corollary 3. *The revenue maximizing fee C^* leads to a unique some-make-AR equilibrium if $\rho < 2/3$ and to multiple equilibria, including a none-make-AR equilibrium, otherwise.*

5.3.1 Price of Conservatism

Assuming that $\rho > 2/3$, the provider can either be risk-averse and charge a fee that leads to a unique equilibrium with guaranteed revenue, or it can be risk-taking and charge a higher fee that may lead to greater revenue but also to zero revenue. To compare between the two options, we use the Price of Conservatism (PoC) metric, which was introduced in Chapter 3. PoC is the ratio between the maximum possible

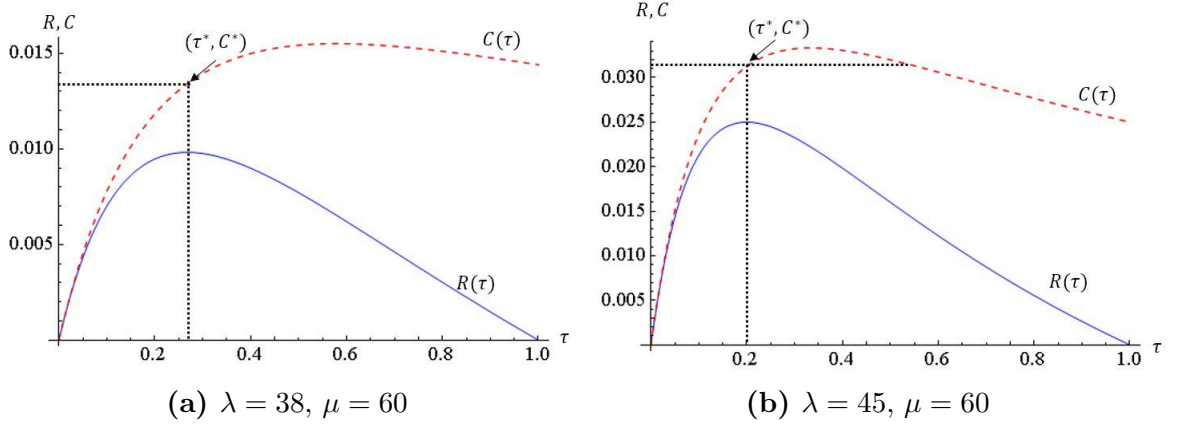


Figure 5.3: When the utilization $\rho < 2/3$, the optimal fee C^* leads to a unique equilibrium (a). When $\rho > 2/3$, C^* leads to multiple equilibria (b).

revenue R^* and the maximum guaranteed revenue R_g^* , which is defined as follows.

$$R_g^* = \max_{0 < \tau_e < 1} R(\tau_e). \quad (5.16)$$

s.t. $C(\tau_e) < \underline{C}$.

Since $R(\tau_e)$ has exactly one extreme point (which is τ_e^*), it is increasing in the range $[0, \tau_e^*)$. Therefore, the maximum guaranteed revenue is achieved when choosing the largest τ_e for which $C(\tau_e) < \underline{C}$. In other words, C should be slightly smaller than \underline{C} . By solving $C(\tau) = \underline{C}$, we get two solutions: $\tau_e^1 = 1$ and

$$\tau_e^2 = \left(\frac{1 - \rho}{\rho} \right)^2. \quad (5.17)$$

By substituting τ_e^2 into Eq. (5.12), we get

$$R_g^* = \frac{2\rho - 1}{2(1 - \rho)}, \quad (5.18)$$




Figure 5.4: The maximum possible revenue and the maximum guaranteed revenue in a system with parameters $\lambda = 45$ and $\mu = 60$.

and by dividing R^* by R_g^* , we get

$$PoC = \frac{\rho^2}{-8\rho^2 + 12\rho - 4}. \quad (5.19)$$

We conclude with the following theorem.

Theorem 12. *If $\rho < 2/3$, then $PoC = 1$. Else, $PoC = \frac{\rho^2}{-8\rho^2 + 12\rho - 4}$.*

Figure 5.4 shows the maximum possible revenue and the maximum guaranteed revenue in a system with parameters $\lambda = 45$ and $\mu = 60$. It can be seen in the figure that for those parameters, the difference between the two values is quite small. Thus, being risk averse seems preferable in this case.

By computing the derivative of PoC with respect to ρ , we get that for any $\rho > 2/3$

$$\frac{dPoC}{d\rho} = \frac{\rho(3\rho - 2)}{4(2\rho^2 - 3\rho + 1)^2} > 0 \quad (5.20)$$

Thus, we obtain the following corollary:

Corollary 4. *The price of conservatism increases with the utilization.*

That is, as the utilization increases, the ratio of the provider's revenue when she is risk-averse to the potential revenue when she is risk-taking increases as well.

5.4 Summary

In this chapter, we analyze a preemptive-resume $M/D/1$ queue that supports advance reservations. First, we find the equilibrium structure and show that if the utilization of the queue is smaller than 0.5, then the equilibrium is unique. Otherwise, the number of equilibria depends on the AR cost. We then study the revenue maximization problem, assuming that the AR cost is a fee charged by the provider. We show that if the utilization is smaller than $2/3$, then the revenue-maximizing fee leads to a unique equilibrium. Otherwise, the fee that leads to the maximum possible revenue also leads to an equilibrium with no revenue. Finally, we show that the ratio between the maximum possible revenue and maximum guaranteed revenue (PoC) increases unboundedly when the utilization of the queue approaches 1.

Chapter 6

Dynamic AR Games

In this chapter, we study a dynamic version of AR games. We consider two AR models: a loss queue with no information sharing (which was analyzed in Chapter 3) and a waiting queue (which was analyzed in the previous chapter). Since the equilibrium structure of the waiting queue model is more explicit than the equilibrium structure of the loss queue model, we focus on the former. Towards the end of this chapter, we explain how the results obtained for the waiting queue model can be extended to the loss queue model.

6.1 Preliminaries and Learning Process

In dynamic games (also known as learning models, since players learn over time the behavior of other players), it is assumed that the game repeats many times and that initially customers do not necessarily follow an equilibrium strategy. The goal is to find the long-term behavior of the customers. In our analysis, we use a *best response dynamic* model which is rooted in Cournot study of duopoly (Cournot, 1897). In this section, we describe the learning model.

At each game (step) a new set of customers participate. The learning process begins with an initial strategy function σ . All customers believe that this strategy will be followed by all other customers. We refer to this initial strategy as the *initial*

belief.

Next, we make the following assumption:

Assumption 3. *Customers that are indifferent between actions AR and AR' choose action AR' .*

Based on this assumption, and using the proof of Lemma 12, one can show that the best response of all customers to any initial belief is a threshold strategy.

We are interested to study the long-term outcomes of a dynamic game. In order to simplify the analysis and since a threshold strategy is followed at all steps, we assume the following.

Assumption 4. *The initial belief is a threshold strategy.*

We denote the threshold of the strategy followed at step $i \geq 1$ by $\tau_i \in [0, 1]$. Customers observe the proportion of customers that choose AR at the previous step and use it to estimate the strategy that was followed at that step. We denote the estimation of τ_i by $\hat{\tau}_i$. If the demand and the number of reservations at step i are D_i and D_{AR}^i respectively, then

$$\hat{\tau}_i = 1 - \frac{D_{AR}^i}{D_i}. \quad (6.1)$$

We distinguish between two types of learning:

1. **Strategy learning.** In this type of learning, it is assumed that at each step i , $\hat{\tau}_i = \tau_i$. That is, customers observe past strategies rather than actions. This can occur when the sample size (i.e., the number of customers at each step) is large.

2. **Action learning.** In this type of learning, if the threshold followed by all customers at step i is not 0 or 1, then there is a strictly positive probability that $\hat{\tau}_i \neq \tau_i$.

We use a generalized version of best response dynamics, presented in (Thorlund-Petersen, 1990). In this version, at each step, customers form a *belief* about the strategy that will be used by other customers at the current step. The belief at step i is constructed based on the previous belief and the last observation. Hence, the belief at step i is

$$\begin{aligned}\beta_i &= \beta_{i-1}(1 - \delta) + \hat{\tau}_{i-1}\delta, \\ \hat{\tau}_1 &= \beta_1,\end{aligned}\tag{6.2}$$

where $\delta \in (0, 1]$ is a parameter that defines the *weight* of the new observation compared to the previous belief, and β_1 represents the initial belief.

Since the best response of all customers to any belief is a threshold strategy, we can define a joint best response function $BR : [0, 1] \rightarrow [0, 1]$. The input is a belief about the threshold strategy that will be followed by all customers. The output is the best response threshold to that belief. Thus, we can describe the best response dynamics of the game as the following process:

$$\tau_i = BR(\beta_{i-1}(1 - \delta) + \hat{\tau}_{i-1}\delta).\tag{6.3}$$

Note that under *strategy-learning* this process is deterministic, while under *action-learning* this process is a stochastic Markov process (Gardiner et al., 1985, Chapter 3).

6.2 Learning Analysis

In this section, we focus on the behavior of customers at a given step. Thus, we remove the subscript i . We begin the analysis with the following observations:

- (i) Given a belief β (i.e., assuming that all other customers follow the threshold β), if a tagged customer with potential priority $p > \beta$ chooses AR , then all customers with potential priority greater than his will have higher priority and all customers with potential priority smaller than his will have lower priority. Therefore, his (believed) expected waiting time is equal to the expected waiting time of a threshold customer that chooses AR in a system where all customers follow the threshold p . Hence,

$$W(\beta, p) = W_{AR}(p) \quad \text{if } p \geq \beta, \quad (6.4)$$

where $W_{AR}(\cdot)$ is defined in Eq. (5.3).

- (ii) Given a belief β , if a tagged customer with potential priority $p < \beta$ chooses AR , then his (believed) expected waiting time is the same as the expected waiting time of the threshold customer (recall that each customer assumes that he is the only one deviating). Hence,

$$W(\beta, p) = W_{AR}(\beta) \quad \text{if } p < \beta. \quad (6.5)$$

- (iii) The expected waiting time of all customers that choose AR' are equal. Hence,

$$W(\beta, 0) = W_{AR'}(\beta) \quad \forall p \in [0, 1], \quad (6.6)$$

where $W_{AR'}(\cdot)$ is defined in Eq. (5.4).

Those properties will be used later to prove our main results.

Next, we show that the best response to a given belief β is determined by the sign of $W_{AR}(\beta) + C - W_{AR'}(\beta)$. That is, the system behave differently if a customer that has a potential priority equals to the belief β is better off choosing AR or AR' .

We split the belief range $[0, 1]$ to a set of intervals where in each interval either $W_{AR}(\beta) + C > W_{AR'}(\beta)$ or $W_{AR}(\beta) + C < W_{AR'}(\beta)$. Thus, the range $[0, 1]$ is split by the set of β values for which

$$W_{AR}(\beta) + C = W_{AR'}(\beta). \quad (6.7)$$

We use the following notations to distinguish between the two types of intervals.

Definition 13. *An interval I is called **AR-interval**, if for any belief $\beta \in I$*

$$W_{AR}(\beta) + C - W_{AR'}(\beta) < 0. \quad (6.8)$$

Definition 14. *An interval I is called **AR'-interval**, if for any belief $\beta \in I$*

$$W_{AR}(\beta) + C - W_{AR'}(\beta) > 0. \quad (6.9)$$

As was explained in the previous chapter, $W_{AR}(0) + C > W_{AR'}(0)$. Hence, the first interval (i.e., the interval with endpoint $\beta = 0$) is an *AR'-interval*. For simplification, we ignore the boundary case where the two expected cost functions are tangent to each other. Thus, any two adjacent intervals are of different types. In Figure 6-1, we illustrate the different intervals in a game with a unique equilibrium and in a game with multiple equilibria.

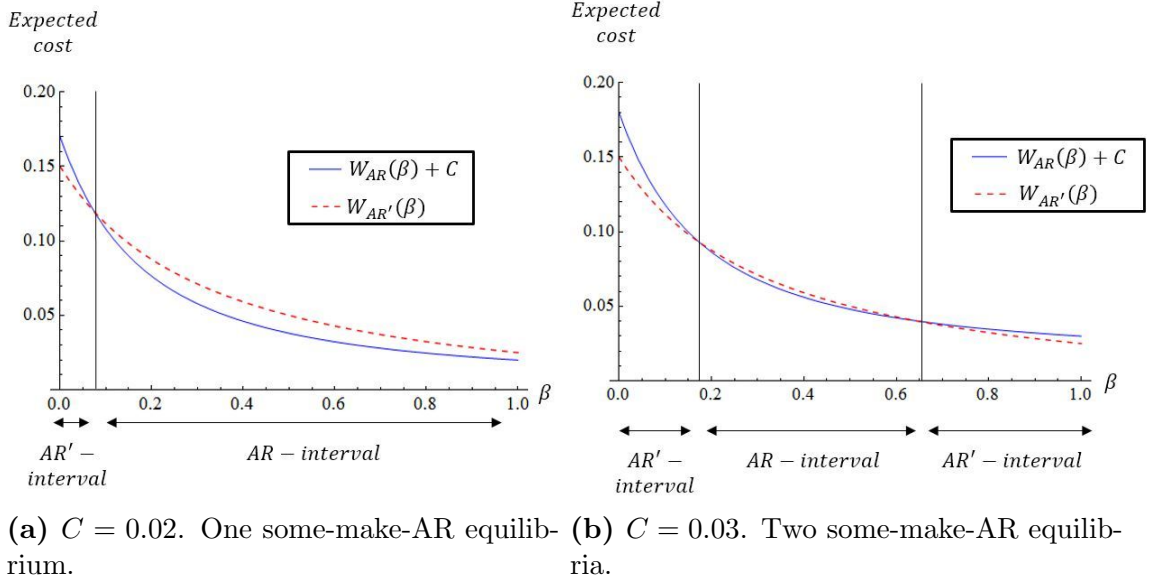


Figure 6.1: An example with $\lambda = 45$ and $\mu = 60$. The *some-make-AR* equilibria split the belief range to intervals. At each interval, one expected cost is smaller than the other.

In the following two lemmas, we separately characterize the best response of customers to a belief belonging to an *AR'-interval* and to that belonging to an *AR-interval*. We show that in the former case, the best response strategy is greater than the belief but still belongs to the same interval. In the latter case, the best response is 0, that is, all customers make AR.

By computing the derivative of $W_{AR}(\beta)$ and $W_{AR'}(\beta)$, one can verify that both functions are decreasing with β . This property will be used in the proofs of both lemmas.

Lemma 14. *If a belief β belongs to AR' -interval I , then*

(i) $BR(\beta) > \beta$;

(ii) $BR(\beta) \in I$.

Proof. First, we consider the case

$$W_{AR}(1) + C > W_{AR'}(\beta). \quad (6.10)$$

Since $W_{AR}(\cdot)$ and $W_{AR'}(\cdot)$ are decreasing functions and by Eq. (6.10), we deduce that the two expected cost functions do not intersect at any point greater than β . Therefore, β belongs to the last interval (i.e., an interval with endpoint $\beta = 1$). Given Eq. (6.4) and Eq. (6.5), and since $W_{AR}(\cdot)$ is decreasing, we deduce that the expected cost of all customers that make AR is at least $W_{AR}(1) + C$. Given that Eq. (6.10) holds, we deduce that no customer will make AR and the best response to β is $\tau = 1$. Thus, for this case, part (i) and (ii) hold.

If Eq. (6.10) does not hold, then there exists a unique value $\eta \in [\beta, 1)$ such that

$$W_{AR}(\eta) + C = W_{AR'}(\beta). \quad (6.11)$$

From Eq. (6.4) and Eq. (6.5), one can see that the expected costs of a customer with potential priority smaller than η is greater than $W_{AR}(\eta) + C$ if he chooses AR. Hence, he is better off choosing AR'. The expected cost of a customer with potential priority greater than η , if choosing AR, is at most $W_{AR}(\eta) + C$, and hence he is better off choosing AR. Thus, the best response of all customers is $\tau = \eta \geq \beta$.

We have shown that $\tau \geq \beta$. Next, we prove part (ii) of the lemma. Let $I = (I^-, I^+)$. We need to show that $\tau < I^+$. Assume by contradiction that $\tau > I^+$. Since both $W_{AR}(\cdot)$ and $W_{AR'}(\cdot)$ are decreasing functions and based on Eq. (6.11), we deduce that the following must hold:

$$W_{AR'}(I^+) < W_{AR}(\tau) + C = W_{AR'}(\beta) < W_{AR}(I^+) + C. \quad (6.12)$$

Sine I^+ is an equilibrium point, Eq. (6.12) contradicts Eq. (5.5). \square

Lemma 15. *If the belief β belongs to an AR-interval, then $BR(\beta) = 0$*

Proof. Based on Eqs. (6.4) and (6.5) and since $W_{AR}(\cdot)$ is a decreasing function, we deduce that

$$W(\beta, p) \leq W_{AR}(\beta), \quad \forall p \in [0, 1]. \quad (6.13)$$

From Eq. (6.8) and Eq. (6.13), we deduce that

$$W(\beta, p) + C < W_{AR'}(\beta), \quad \forall p \in [0, 1]. \quad (6.14)$$

That is, all customers are better off choosing AR and the best response to β is $\tau = 0$. \square

6.3 Stability Analysis

Using Lemma 14 and Lemma 15, we analyze the different equilibria in terms of stability. The stability analysis helps in establishing the long-term behavior of the dynamical system. We start by adapting the definitions of stability presented in (Fudenberg, 1998) to our model.

Definition 15. *A strategy with threshold τ is a steady-state strategy if and only if $\mathbb{P}(\tau_i = \tau | \beta_1 = \tau) = 1, \forall i > 0$.*

That is, only if the initial belief is being followed at all future steps, then it is a steady state. The definition implies that in order for a strategy to be a steady-state strategy it is necessary that it will be an equilibrium strategy but it is not a sufficient condition.

Definition 16. *A strategy with threshold τ is a stable steady-state strategy if it is a steady-state strategy and in addition, for every neighborhood θ of τ (which is any interval in $[0, 1]$ that contains τ) there exists a neighborhood $\theta_1 \in \theta$ of τ such that if $\beta_1 \in \theta_1$, then $\tau_i \in \theta, \forall i > 0$.*

That is, if the initial belief is close enough to the steady-state, it remains nearby. A steady-state strategy that is not stable is referred to as an *unstable* steady-state strategy.

Definition 17. *A strategy with threshold τ is an asymptotically stable steady-state strategy if it is stable and in addition, there exists a neighborhood θ of τ such that if $\beta_1 \in \theta$, then $\lim_{i \rightarrow \infty} \tau_i = \tau$.*

That is, if the initial belief is close enough to the steady-state, it will eventually converge to it.

First, we note that due to the stochastic nature of *action-learning*, a *some-make-AR* equilibrium is not a steady-state under *action-learning*. That is, even if at some point all customers follow an equilibrium strategy, with some probability, the fraction of customers choosing each action will defer from that strategy. In this case, it is not guaranteed that the equilibrium strategy will be followed in the next stage. Thus, under *action-learning*, the game can only converge to a *none-make-AR* equilibrium (if it exists). Under *strategy-learning*, all equilibria are steady-states. Thus, all equilibria are candidates for convergence. However, in the next section, we will show that if three equilibria exist, the game can only converge to two of them.

In the remaining of this section, we classify each equilibrium according to Definitions 15-17. Using Lemmas 14 and 15, we obtain the following results.

Theorem 13. *under strategy-learning, a some-make-AR is an unstable steady-state equilibrium.*

Proof. Consider a *some-make-AR* equilibrium with threshold τ^e . Let $\theta = [\tau^e - \gamma, \tau^e + \gamma]$, where $0 < \gamma < \tau^e$ (i.e., $0 \notin \theta$). Given any neighborhood $\theta_1 \in \theta$ of τ^e and a belief $\beta \in \theta_1$, it follows by Lemma 15 that either the best response to any $\beta < \tau_e$ is 0 or the best response to any $\beta > \tau_e$ is 0. Hence, θ does not contain a sub neighborhood of τ^e , such that the best response to any belief in this sub-neighborhood is in θ . This contradicts Definition 16. We conclude that a *some-make-AR* equilibrium is an unstable steady-state. □

Theorem 14. *Under both strategy-learning and action-learning, if a none-make-AR equilibrium exists, then it is an asymptotically stable steady-state equilibrium.*

Proof. With both *strategy-learning* and *action-learning*, if at some step i , $\beta_i = 1$ and *none-make-AR* is an equilibrium, then at all future steps all customers will keep not making AR. Thus, it is a steady-state.

Next, we show that under *action-learning*, a *none-make-AR* equilibrium is an asymptotically stable steady-state. *Strategy-learning* can be seen as a special case of *action-learning* where the estimators are always exact. Thus, it is sufficient to prove the result for *action-learning* only.

Let θ be a neighborhood of 1. Under *action-learning*, a steady-state is stable only if there is a neighborhood of 1 $\theta_1 \in \theta$, such that the best response to any threshold strategy in θ_1 is a *none-make-AR* strategy. If a *none-make-AR* equilibrium exists, then $C + W_{AR}(1) > W_{AR'}(1)$. Since $W_{AR'}(\cdot)$ is a continuous function, we deduce that there exists θ_1 such that, for any $\nu \in \theta_1$, $C + W_{AR}(1) > W_{AR'}(\nu)$. In other words, the best response strategy to a belief $\beta_i \in \theta_1$ is $\tau_i = 1$. Thus, $\beta_{i+1} > \beta_i$, and hence also $\tau_{i+1} = 1$. By induction, we conclude that the belief converges to 1. Hence, we have shown that if a *none-make-AR* is an equilibrium, then any neighborhood of 1 has a

sub-neighborhood of 1 θ_1 such that if $\beta_1 \in \theta_1$, then $\lim_{i \rightarrow \infty} \beta_i = 1$, and therefore it is an asymptotically stable steady-state.

□

In Table 6.1, we summarize the results of this section. In the next two sections, we use the results we obtained so far to separately analyze the dynamics of a system with a unique equilibrium and a system with multiple equilibria.

	<i>Strategy-learning</i>	<i>Action-learning</i>
<i>Some-make-AR equilibrium</i>	An unstable steady-state	Not a steady-state
<i>None-make-AR equilibrium</i>	An asymptotically stable steady-state	An asymptotically stable steady-state

Table 6.1: Stability summary.

6.4 The Dynamics of a Game with Unique Equilibrium

In this section, we assume that $C < \underline{C}$. Thus, the game has a unique *some-make-AR* equilibrium with threshold τ^e . Hence, there are two intervals $I_{AR} \triangleq [0, \tau^e)$ and $I_{AR'} \triangleq (\tau^e, 1]$.

6.4.1 Strategy-learning

Using the analysis of a single step which we carried at the previous sections, we next show that regardless of the initial belief, the customers' strategy always converges to τ^e .

Theorem 15. *Under strategy-learning, convergence to the unique some-make-AR equilibrium is guaranteed.*

Proof. We split the proof into two cases. In the first case, we assume that $\beta_1 \in I_{AR}$. From Lemma 14, we deduce that $\tau_1 \in (\beta_1, \tau^e]$. Since for any weight δ , $\beta_2 \in (\beta, \tau_1]$, we deduce that if $\beta_2 \neq \tau^e$, then $\beta_2 \in I_{AR}$. By induction, we deduce that, for any $j > 0$,

$$\tau_j \geq \beta_j \geq \beta_{j-1}, \quad (6.15)$$

$$\tau_j \leq \tau^e. \quad (6.16)$$

The set $\{\tau_i, i = 1, 2, \dots\}$ is a monotonically increasing sequence bounded by τ^e . Thus, it has a limit, denoted by L . From Eq. (6.2), we deduce that $\lim_{i \rightarrow \infty} \beta_i \rightarrow L$, and hence $\lim_{i \rightarrow \infty} BR(\beta_i) = L$. We conclude that the limit L is a fixed point of BR , and hence it must be the equilibrium point τ^e .

In the second case, we assume that β_1 belongs to the second interval. Therefore, based on Lemma 15, $\tau_1 = 0$ and $\beta_2 \in [0, \beta_1)$. If $\beta_2 \in I_{AR'}$, then the strategy will converge to τ^e as we showed earlier in the proof. If not, then $\tau_2 = 0$. In this case, $\beta_3 \in [0, \beta_2)$. By induction, we deduce that if at step $i > 1$ the belief still belongs to I_{AR} , then

$$\beta_i = \beta_1(1 - \delta)^{i-1}. \quad (6.17)$$

This guarantees that for some $i \geq 2$, $\beta_i \in I_{AR'}$ and the strategy must converge to τ^e . □

6.4.2 Action-learning

Since *some-make-AR* equilibrium is not a steady-state under *action-learning*, the customers' strategy does not converge to an equilibrium. Next, we show that, at each step, all customers follow a strategy with threshold between 0 and τ^e . Thus, the

expected number of customers making AR, at each step, is greater than in a system where all customers follow the equilibrium strategy.

Theorem 16. *In a dynamic game with a unique some-make-AR equilibrium, the average number of customers making AR under action-learning is greater than under strategy-learning.*

Proof. Consider an arbitrary step i and assume that *action-learning* is applied. If at step i the belief β_i belongs to an *AR-interval*, then all customers will choose *AR*. If β_i belongs to an *AR'-interval*, then $\tau_i \in (\beta_i, \tau^e)$. Thus, in any realization, the strategy followed by all customers in all steps is a random variable that takes values between 0 and τ^e . In *strategy-learning*, the strategy followed by all customers converges to τ^e . Thus, the average fraction of customers not making *AR* converges to a value between 0 and τ^e under *action-learning* and to τ^e under *strategy-learning*. \square

Next, we present a simulated example that compares between the revenue in the *action-learning* model and the *strategy-learning* model. The pseudo-code of the simulation is given in Procedure 2.

Procedure 2 Learning Simulation $(\lambda, \beta_1, C, \delta, l)$

```

for  $i \leftarrow 1$  to  $l$  {iterating over all steps} do
   $D_{AR} \leftarrow 0$  {the number of reservations}
   $D \leftarrow$  generate Poisson random variable {the number of customers}
  for  $j \leftarrow 1$  to  $D$  {iterating over all customers} do
     $p \leftarrow$  generate random variable from  $U(0,1)$  {the potential priority}
    if  $p > \beta_i$  {check if the potential priority is greater than the current belief}
      then
        if  $W_{AR}(p) + C < W_{AR'}(\beta_i)$  {check if the customer is better off making AR}
          then
             $D_{AR} \leftarrow D_{AR} + 1$  {increase the number of reservations by one}
          end if
        else
          if  $W_{AR}(\beta_i) + C < W_{AR'}(\beta_i)$  {check if the customer is better off making AR}
            then
               $D_{AR} \leftarrow D_{AR} + 1$  {increase the number of reservations by one}
            end if
          end if
        end for
      if strategy-learning then
         $\hat{\tau}_i \leftarrow \begin{cases} 0 & \text{if } \beta_i > \tau^e, \\ \tau : W_{AR}(\tau) + C = W_{AR'}(\beta_i) & \text{if } \beta_i \leq \tau^e, \end{cases}$ 
        {compute the current strategy}
      end if
      if action-learning then
         $\hat{\tau}_i \leftarrow 1 - \frac{D_{AR}}{D}$  {estimate the current strategy}
      end if
       $\beta_{i+1} \leftarrow \beta_i(1 - \delta) + \hat{\tau}_i\delta$  {compute the new belief}
    end for
  end for

```

Example 8. We consider a queue with parameters $\lambda = 45$ and $\mu = 60$. We set the AR cost to $C = 0.024$. The unique equilibrium is $\tau^e = 0.1026$ (i.e., on a static game, on average, 89.74% of the customers make AR). We run a simulation of 10,000 steps. Each step lasts for one time unit (i.e., the average demand at each step is 45). We set $\beta_1 = \tau^e$ and $\delta = 1$. The average number of reservations per time unit is 40.3 (i.e., on




Figure 6-2: Simulation results. In a game with unique equilibrium, more customers make AR under *action-learning* than under *strategy-learning*.

average, 89.5% make AR) under *strategy-learning* and 42.7 (i.e., on average, 94.9% make AR) under *action-learning*. We conclude that, as Theorem 15 states, when customers base their decisions on historic actions and not strategies, more customers make AR. Statistical analysis (one-tailed *t*-test) shows that the difference between the mean number of reservations under *action-learning* and under *strategy-learning* is statistically significant, with confidence level of 99%. In Figure 6-2, we plot the number of reservations, under *action-learning* and under *strategy-learning*. We use the same realization of customers in both cases and we can see that at each iteration, the number of reservations is greater (or equal) when applying *action-learning*.

We conclude that if the provider interest is that as many customers as possible will make reservations, then she is better off if customers gain information about previous actions rather than strategies.

6.5 The Dynamics of a Game with Multiple Equilibria

In this section, we assume that $\rho > 0.5$ and $\underline{C} < C < \overline{C}$. Thus, the game has two *some-make-AR* equilibria and a *none-make-AR* equilibrium. We denote the smaller and greater thresholds followed at the *some-make-AR* equilibria by τ_1^e and τ_2^e , respectively. The three intervals are $[0, \tau_1^e)$, (τ_1^e, τ_2^e) and $(\tau_2^e, 1)$. Note that the first and third intervals are *AR'-intervals*, while the second interval is an *AR-interval*.

6.5.1 Strategy-learning

Based on Lemma 14, Lemma 15 and the proof of Theorem 15, one can show that when the initial belief belongs $[0, \tau_2^e]$ the game converges to τ_1^e and when the initial belief belonging to $(\tau_2^e, 1]$ the game converges to 1.

Corollary 5. *Under strategy-learning, a game with multiple equilibria converges to τ_1^e , if $\beta_1 < \tau_2^e$, and to 1 otherwise.*

Example 9. *We consider the same system as described in Example 8. However, this time, we set the AR cost to $C = 0.032$. For this cost, the set of equilibria is $\{\tau_1^e = 0.222, \tau_2^e = 0.5, \tau_3^e = 1\}$. We set three different initial thresholds, each one belongs to a different interval: $\beta_1 = 0.2, \beta_2 = 0.4$ and $\beta_3 = 0.6$. We apply strategy-learning with $\delta = 0.8$. As Figure 6-3 shows, within a few steps, the system converges to an equilibrium.*

From Corollary 5, we deduce that if the initial belief is a random variable, then with probability $\mathbb{P}(\beta_1 < \tau_2^e)$ the strategy converges to τ_1^e and with probability $1 - \mathbb{P}(\beta_1 < \tau_2^e)$ it converges to one. Assuming that C is a fee charged by the provider, we aim to determine which fee maximizes the average revenue from AR fees in a dynamic game that repeats many times. By Corollary 5, the expected revenue depends on both




Figure 6.3: Convergence to equilibrium under *strategy-learning*

τ_1^e and τ_2^e . Thus, the first step is to find the relation between those two thresholds. By manipulating the equation $C(\tau_1^e) = C(\tau_2^e)$, we get the following:

$$\tau_2^e = \left(\frac{1-\rho}{\rho} \right)^2 \frac{1}{\tau_1^e}. \quad (6.18)$$

We denote the CDF of β_1 by F_β and the long-term expected revenue of the dynamic game (i.e., the expected revenue after convergence) by $R_D(\tau_1^e)$. Using Eq. (6.18), we obtain

$$R_D(\tau_1^e) = R(\tau_1^e) F_\beta \left(\left(\frac{1-\rho}{\rho} \right)^2 \frac{1}{\tau_1^e} \right) + 0 \cdot \left(1 - F_\beta \left(\left(\frac{1-\rho}{\rho} \right)^2 \frac{1}{\tau_1^e} \right) \right). \quad (6.19)$$

Given F_β , one can find, using Eq. (6.19) and Eq. (5.12), the value of τ_1^e that maximizes $R_D(\tau_1^e)$ and, in turn, the optimal fee. Let assume that F_β is uniformly distributed in $[0, 1]$. In this case, $R_D(\tau)$ becomes

$$R_D(\tau_1^e) = \frac{(1-\tau_1^e)(1-\rho)}{2\rho(1-\rho(\tau_1^e-1))^2}. \quad (6.20)$$

By computing the derivative of $R_D(\tau_1^e)$ with respect to τ_1^e , one can show that it decreases with τ_1^e . Thus, we deduce that when considering a threshold which is not unique (i.e., a threshold between $((1-\rho)/\rho)^2$ and 1), the optimal threshold is

$((1 - \rho)/\rho)^2$. Combining this result with Corollary 3 leads to the following theorem:

Theorem 17. *Under strategy-learning, if the initial belief is uniformly distributed between 0 and 1, then the optimal fee is \underline{C} when $\rho > 2/3$ and C^* when $\rho < 2/3$.*

6.5.2 Action-learning

Under *action-learning*, a system formed by a risk-taking provider has one steady-state equilibrium which is a *none-make-AR* equilibrium. In this section, we show that in some cases the system converges to this steady-state, while in other cases it cycles. The long-term behavior is determined by the initial belief β_1 and the weight δ .

Since *none-make-AR* equilibrium is an asymptotically stable steady-state there exists $\nu > 0$ such that if the belief at some step is in the range $[\nu, 1]$, then the customers' strategy converges to the *none-make-AR* equilibrium. Thus, we can conclude that regardless of the value of δ , if the initial belief is in the range $[\nu, 1]$, then convergence to the *none-make-AR* equilibrium is guaranteed. The more interesting question is how the system behaves when the initial belief is in the range $[0, \nu)$. Next, we show that if $\delta = 1$ (i.e., the belief is based only on the actions at the most recent step), then convergence to equilibrium is guaranteed, regardless of the initial belief.

Theorem 18. *Under action-learning, if $\delta = 1$ and multiple equilibria exist, then the system converges to a none-make-AR equilibrium with probability 1.*

Proof. At each step in which $\tau_i > 0$, there is a strictly positive probability that $\hat{\tau}_i$ will be in the range $[\nu, 1]$. At the boundary case $\tau_i = 0$, the probability that all customers will choose *AR'* is 0. However, $\tau_i = 0$ is not an equilibrium, and hence in the next step τ_{i+1} will be greater than zero. Since the game repeats infinitely many times, with probability 1, there exists a step j such that $\hat{\tau}_j$ will be in the range $[\nu, 1]$. Since

$\beta_{j+1} = \hat{\tau}_j$ when $\delta = 1$, we have that $\tau_{j+1} = 1$. Thus, convergence to a *none-make-AR* equilibrium is guaranteed. \square

Next, we show that if $\delta < 1$, the strategy does not necessarily converge to a *none-make-AR* equilibrium. We do it by the means of an example.

Example 10. *Consider the same game as in Example 9 but with AR cost $C = 0.028$ and $\delta = 0.5$. In this system, the set of equilibria is $\{0.142, 0.777, 1\}$. Starting with an initial belief $\beta_1 < 0.777$, the system can converge to 1 only if, at some point, the belief will belong to the last interval $(0.777, 1]$.*

By Lemma 15, the best response to any belief belonging to $(0.142, 0.777)$ is 0. Thus, the belief needs to “jump” from the interval $[0, 0.142]$ to the interval $[0.777, 1]$. This cannot happen since the belief, at any step, cannot be greater than $0.142 \cdot 0.5 + 1 \cdot 0.5 = 0.571$.

6.6 Learning in Loss systems

In this section, we consider a dynamic version of the loss queue game with *no-information* sharing. As it turns out, the results of this variant are almost identical to those of the dynamic waiting queue. To avoid repetition, we do not prove the results for the loss queue. For formal proofs, see (Simhon et al., 2015)

First, we observe that the stability analysis is the same for both models. Thus, Theorems 13 and 14 are also valid for the loss queue. Next, we separately analyze a game with a unique *some-make-AR* equilibrium and a game with multiple equilibria including a *none-make-AR* equilibrium. We begin by proving that there exists an AR cost that leads to a unique *some-make-AR* equilibrium.

Lemma 16. *For any number of servers and demand distribution, there exists a fee for which the game has a unique some-make-AR equilibrium.*

Proof. Since for any threshold $\tau > 0$, $\pi_{AR}(\tau) > \pi_{AR'}(\tau)$ and since $\pi_{AR}(0) = \pi_{AR'}(0)$ (see Section 3.2.2), we deduce that there exists $x \in [0, 1]$ such that $f(t) \triangleq \pi_{AR}(\tau) - \pi_{AR'}(\tau)$ is monotonically increasing in the range $[0, x]$. We define $\tau_{min} = \min_{x \leq \tau \leq 1} f(\tau)$. In case that the fee C is set such that $(1-C)\pi_{AR}(\tau_{min}) > \pi_{AR'}(\tau_{min})$, then the two cost functions intersect once. In this case, the following holds: $(1-C)\pi_{AR}(1) > \pi_{AR'}(1)$, and hence *none-make-AR equilibrium* does not exist and we conclude that there is a unique *some-make-AR equilibrium*. \square

The analysis of a loss queue with a unique *some-make-AR* equilibrium is the same as that of the waiting queue. Thus, Theorem 15 and 16 are valid also for the loss queue.

Next, we consider the case of having multiple equilibria including both *some-make-AR* and *none-make-AR*. The main difference between the waiting queue the loss queue is that in the former one there are exactly two *some-make-AR* equilibrium, while in the latter one the number of *some-make-AR* equilibria depends on the distribution of the demand and can be greater than two.

Under *strategy-learning*, if there are more than two *some-make-AR* equilibria, we cannot know to which equilibrium the game will converge. Thus, Corollary 5 becomes:

Corollary 6. *Under strategy-learning with multiple equilibria, the customers' strategy convergences to an equilibrium.*

Finally, under *action-learning*, the game can only converge to a *none-make-AR* equilibrium. Thus, Theorem 17 is also valid for the loss queue model.

6.7 Summary

In this chapter, we study a dynamic version of the game where it repeats many times and customers observe either past actions or strategies. First, we show that under

strategy-learning, the game converges to an equilibrium and under *action-learning* the game can only converge to a *none-make-AR* equilibrium, if exists.

Next, we analyze the system from the perspective of a provider that charges a reservation fee and is aiming to maximize her revenue. The analysis shows that if the provider is risk-averse and chooses a low fee that leads to a unique equilibrium, sharing information about the actions yields on average greater revenue than sharing information about the strategies. On the other hand, if the provider is risk-taking and charges a fee that leads to multiple equilibria, then sharing previous actions may eventually cause all customers not to make AR. Furthermore, we observe that the way the belief is constructed impacts the long term outcomes. In some cases, convergence is guaranteed, while in others, cycling is guaranteed. Thus, we conclude that the optimal pricing policy depends on the value of δ and whether the customers observes historical actions or strategies.

Chapter 7

Information Sharing in Waiting Queues

In this chapter, we explore different information sharing policies in an $M/M/1$ queue where customers need to decide whether to join the queue or to balk. The literature on the strategic behavior of customers in queues is traditionally divided into the classical observable and *no-information* queues. In the former case, introduced by Naor (Naor, 1969), customers are informed about the current queue length before deciding whether to join or balk. In the latter case, introduced by Edelson and Hilderbrand (Edelson and Hilderbrand, 1975), customers make their decisions based on statistical information (e.g., the queue parameters). In both cases, customers behave strategically and join the queue only if their expected waiting cost is smaller than the reward obtained upon being served.

The goal of this work is to find out whether there are situations where the service provider can increase her revenue by combining those two frameworks. In particular, we are interested in studying policies where the provider shares information with some customers and hides it from others, depending on the actual queue length. We assume that the service provider has a fixed income from each customer that joins the queue. Thus, in order to maximize her revenue, the provider should maximize the effective arrival rate, which is the rate of customers that join the queue, or equivalently, minimize the idle period of the system.

We compare between the outcomes of the following three policies: (i) always

inform customers about the queue length; (ii) never inform customers about the queue length; (iii) inform customers based on a threshold policy in which queue length information is provided when the queue length is below a specified threshold and is hidden otherwise. Although the third policy seems intuitive, we formally prove that, in any setting, either sharing information with all customers or hiding information from all customers always yields greater expected revenue.

7.1 Model Description

We consider a standard $M/M/1$ first-come-first-served queueing system. The arrival rate is a Poisson process with mean λ . The service rate is exponential with mean $1/\mu$. We denote $\rho = \lambda/\mu$ to be the *maximum* load of the queue (recall that customers do not always join the queue). The cost of each time unit spent at the queue (waiting and being served) is C . Without loss of generality, we set $C = 1$. All customers have the same reward U from service, where $U > 1/\mu$ (otherwise, no customer ever joins the queue).

The standard set-up in all the problems related to strategic behavior in queueing systems (Hassin and Haviv, 2003) is that a new customer decides to join or to balk depending on his expected sojourn time in the queue and the reward obtained for service completion. The behavior of a new customer is highly dependent on his knowledge/information about the current queue length (the queue length at the instant of his arrival). Naturally, a customer will join the queue if and only if the expected sojourn time (which depends on the information about the queue length provided to him) is smaller than the reward.

We consider a state-dependent information disclosure policy, such that the control parameter u is a mapping from the queue length i into the interval $[0, 1]$. Then, the

control $u(i)$ for all $i = 0, 1, \dots$ is the probability that the provider gives the information to an arrival customer when there are i customers in the queue. We only consider stationary information disclosure policies, i.e. control policies that do not depend on time. If $u(i) = 1$, for all $i \geq 0$, then it is exactly the (fully) observable model presented by Naor (Naor, 1969) and if $u(i) = 0$, for all $i \geq 0$, then it is the *no-information* model presented by Edelson and Hilderbrand (Edelson and Hilderbrand, 1975). For the rest of the paper, the former information disclosure policy is denoted u_+ and the latter policy is denoted u_- .

The provider's objective is to maximize her revenue, generated through the service completion. Thus, it aims to maximize the utilization of the queue which is equivalent to minimizing the idle stationary probability of the queue, denoted by π_0 . The optimization problem for the provider is the following:

$$\min_{u \in \mathcal{U}} \pi_0^u, \tag{7.1}$$

where \mathcal{U} is the set of information disclosure policies/mapping from \mathbb{N} to the interval $[0, 1]$. Based on Naor's model, we know that an informed customer will join if the queue length is strictly lower than the threshold $L = \lfloor U\mu \rfloor$. An uninformed customer makes his decision based on his expected sojourn time, denoted by W^{UI} . We assume that the uninformed customers are aware of the policy used by the provider (this information can be obtained by trials or via exogenous sources). Thus, an uninformed customer can evaluate his sojourn time given the system parameters and the provider policy.

Getting results for all type of information disclosure policies can be complicated. Therefore, we restrict our analysis to deterministic threshold information disclosure policies, which make sense in applicative contexts. The principle of revealing the

queue length when the system is empty and not revealing it when the system is full seems intuitive. Revealing information when the queue is small should increase the incoming rate of customers. At the opposite, when the system is overloaded, the provider may aim not to scare incoming customers and give them the information about the queue length.

7.2 Equilibrium Analysis of Threshold Policies

In this section, we consider threshold policies, denoted $u_D(\cdot)$, for which the provider informs all customers about the queue length if the actual queue length is below or equal some threshold D and does not inform them otherwise, that is,

$$u_D(i) = \begin{cases} 1 & \text{if } i \leq D, \\ 0 & \text{if } i > D. \end{cases} \quad (7.2)$$

Since an uninformed customer knows that the queue length is greater than D , his expected sojourn time, denoted $W^{UI}(D, q)$, depends on the threshold D and on the decision of the other uninformed customers, denoted by q , where $q \in [0, 1]$ represents the probability that an uninformed customer joins the queue (all customers are identical, hence we only consider symmetric strategies). We denote an equilibrium solution by q^* .

If $D \geq L - 1$, uninformed customers will never join the queue as their expected sojourn times, conditioned on the fact that the queue length is at least L , is necessarily higher than the reward U (recall that $L = \lfloor U\mu \rfloor$). Thus, the unique equilibrium is $q^* = 0$ and a threshold policy with a threshold $D \geq L - 1$ is equivalent to the observable model. Henceforth, we only consider threshold policies with $D \in \{0, 1, \dots, L - 2\}$ (which also means that we only consider $U\mu \geq 2$). Thus, at equilibrium, all informed

customers join, while all uninformed customers join with probability q .

The evolution of the queue length process forms a Markov chain with transition rate λ from state i to state $i + 1$ if $i \leq D$ and transition rate λq otherwise. The transition rate from i to $i - 1$ is μ for all $i > 0$. The Markov chain is illustrated in Figure 7.1. This Markov chain is a birth-death process and the stationary distribution is given by

$$\pi_i = \begin{cases} \pi_0 \left(\frac{\lambda}{\mu}\right)^i & \text{if } i \leq D + 1, \\ \pi_0 \frac{\lambda^i q^{i-D-1}}{\mu^i} & \text{if } i > D + 1, \end{cases} \quad (7.3)$$

where π_i is the stationary probability that the queue length is equal to i . The idle stationary probability (i.e., the probability that the queue is empty) as a function of the threshold D and the joining probability q is given by:

$$\begin{aligned} \pi_0(D, q) &= \left[\sum_{i=0}^{D+1} \left(\frac{\lambda}{\mu}\right)^i + \sum_{i=D+2}^{\infty} \frac{\lambda^i q^{i-D-1}}{\mu^i} \right]^{-1} \\ &= \frac{(\lambda - \mu)(\lambda q - \mu)}{\lambda^2(-1 + q)(\lambda/\mu)^D - \lambda q \mu + \mu^2}. \end{aligned} \quad (7.4)$$

We define $\pi(i|i > D)$ to be the probability that the queue length is i , given that the queue length is higher than D . Using conditional probability rules, we obtain

$$\pi(i|i > D) = \frac{\pi_i}{\sum_{j=D+1}^{\infty} \pi_j}. \quad (7.5)$$

The expected sojourn time for an uninformed new customer is

$$W^U(D, q) = \sum_{i=D+1}^{\infty} \frac{i+1}{\mu} \pi(i|i > D). \quad (7.6)$$




Figure 7.1: A Markov chain representation of the system evolution. D is the threshold above which customers are not informed about the queue length and q is the probability that uninformed customers join the queue. In this example, we assume $D < L - 1$, that is, informed customers always join the queue.

Using Eq. (7.3) and Eq. (7.5), we can express the expected sojourn time:

$$W^{UI}(D, q) = \frac{\sum_{i=D+1}^{\infty} \pi_i \frac{i+1}{\mu}}{\sum_{i=D+1}^{\infty} \pi_i} = \frac{1+D}{\mu} + \frac{1}{\mu - \lambda q}. \quad (7.7)$$

We next aim to determine the equilibrium strategy q^* which determines if an uninformed customer joins the queue or not. First, we prove the existence and uniqueness of the equilibrium.

Lemma 17. *A game with a threshold information policy has a unique equilibrium.*

Proof. Joining when the queue length is smaller than L and balking otherwise is a dominant strategy for the informed customers, hence unique. Given the behavior of the informed customers, we next analyze the behavior of the uninformed customers.

From Eq. (7.7), we deduce that the expected payoff of an uninformed customer that joins (i.e., $U - W^{UI}(D, q)$) is decreasing with q , while the expected payoff of balking (which is zero) does not depend on q . Hence, the two payoff functions either intersect once or do not intersect. In the latter case, we will have a unique pure equilibrium where all uninformed customers join ($q^* = 1$) or all of them balk ($q^* = 0$); depending on which expected payoff is greater. In the former case, the intersection

point forms a mixed equilibrium. The fact that the expected payoff of joining is decreasing with q , while the expected payoff of balking does not depend on q means that the game has the "avoid-the-crowd" property which guarantees that the equilibrium is unique (see Chapter 1 in (Hassin and Haviv, 2003)). \square

We start the equilibrium analysis with the simple case $\rho < 1$ and $U > 1/(\mu - \lambda)$. In this case, if policy u_- is used, then all customers join the queue because the expected sojourn time when all the customers join is $1/(\mu - \lambda)$. This outcome is obviously optimal for the provider and no other policy can outperform it. In any other case, an equilibrium where all customers join does not exist, namely at equilibrium $q < 1$. Henceforth, we focus on that scenario which implies that either the system is overloaded, that is $\rho > 1$, or it is underloaded with bounded reward, that is $\rho < 1$ and $U < 1/(\mu - \lambda)$.

The unique equilibrium might be a pure equilibrium with $q^* = 0$ or a mixed equilibrium with $q^* \in (0, 1)$. In the case of a pure equilibrium, an uninformed customer knows that the queue length is exactly $D + 1$ (since no uninformed customer joins the queue), and hence his expected sojourn time is $(D + 2)/\mu$. In order for $q^* = 0$ to be an equilibrium, an uninformed customer should not be better off by joining, which implies that

$$D + 2 \geq U\mu \geq \lfloor U\mu \rfloor = L. \quad (7.8)$$

Since we only consider threshold policies with $D \leq L - 2$, we deduce that a pure equilibrium exists only if $U\mu = D + 2$.

Next, we study the case of having a mixed equilibrium. We derive the fraction of uninformed customers that join the queue, by using the property that at a mixed equilibrium, each player must be indifferent between the actions of joining and not

joining the queue. Hence, at equilibrium, $W^{UI}(D, q^*) = U$. By isolating q in Eq. (7.7), we get

$$q^* = \frac{\mu(U\mu - D - 2)}{\lambda(U\mu - D - 1)}. \quad (7.9)$$

Since a pure equilibrium with $q^* = 0$ only exists when $U\mu = D + 2$, the equation above captures both the case of a mixed equilibrium and the case of a pure equilibrium.

From Eq. (7.9), we obtain the customers' equilibrium strategy under an information disclosure policy u_D adopted by the provider. This result is summarized with the following theorem.

Theorem 19. *If a provider uses the information disclosure policy $u_D(\cdot)$, all informed customers join the queue if the queue length is strictly smaller than L . Uninformed customers join the queue with probability $q^*(D)$, where*

$$q^*(D) = \begin{cases} 0, & \text{if } D \geq L - 1, \\ \frac{\mu(U\mu - D - 2)}{\lambda(U\mu - D - 1)}, & \text{otherwise.} \end{cases}$$

7.3 Comparison of Policies

Using Theorem 19, we next determine the optimal information disclosure policy that the provider should adopt in order to optimize her revenue. Toward this end, we derive and compare the idle stationary probabilities, at equilibrium, for the three types of policies: u_D, u_- and u_+ . We ignore the case of $D \geq L - 1$ since it is equivalent to the policy u_+ .

For deriving the idle stationary probability when a threshold policy u_D is used, we substitute Eq. (7.9) into Eq. (7.4) (i.e., replacing q with q^*). We obtain that at

equilibrium

$$\pi_0^*(D) = \frac{\mu - \lambda}{\mu - \lambda \rho^D [\rho(U\mu - 1 - D) - (U\mu - 2 - D)]}. \quad (7.10)$$

When sharing queue length information with all customers, we have, at equilibrium, an $M/M/1/L$ queue (i.e., a queue with finite capacity L). This type of queue was studied in (Takács, 1962) and the idle stationary probability is

$$\pi_0^{u+} = \frac{\mu - \lambda}{\mu - \lambda \rho^{\lfloor U\mu \rfloor}}. \quad (7.11)$$

Finally, when no information is shared with customers there is a unique mixed equilibrium (recall that $U > 1/\mu$ and either $\rho > 1$ or $U < 1/(\mu - \lambda)$ and, hence, there is no pure equilibrium in this case). The effective arrival rate is λq^* . Since all customers are indifferent between joining and balking, the following holds:

$$U = \frac{1}{\mu - \lambda q^*}. \quad (7.12)$$

By isolating q^* and substituting it in the equation $\pi_0^{u-} = 1 - \lambda q^*/\mu$, we get that

$$\pi_0^{u-} = \frac{1}{U\mu}. \quad (7.13)$$

To compare between the performances of the different policies, we distinguish between the case of an overloaded queue (i.e., $\rho > 1$) and that of an underloaded queue (i.e., $\rho < 1$). Note that the strategic behavior of the customers guarantees that the effective load is always smaller than one.

Next, we show that in an overloaded queue, for any value of t and $D \in \{0, 1, \dots, L - 2\}$, we always have $\pi_0^*(D) > \pi_0^{u+}$. Similarly, in an underloaded queue, $\pi_0^*(D) > \pi_0^{u-}$ always holds. Hence:

Theorem 20. *Considering the set of deterministic threshold based information disclosure policies. To maximize her revenue, a service provider must either use the full information policy u_+ or the empty information policy u_- .*

Proof. We split the proof into two cases. In the first case, we assume that $\rho > 1$. For this case, we will show that $\pi_0(D) > \pi_0^{u+}$. Then, we study the case of $\rho < 1$, for which we show that $\pi_0(D) > \pi_0^{u-}$.

From Eq. (7.10) and Eq. (7.11), we deduce that we can prove the first case by showing that:

$$\rho^{\lfloor U\mu \rfloor} > \rho^D [\rho(U\mu - 1 - D) - (U\mu - 2 - D)], \quad (7.14)$$

which is equivalent to showing that

$$\rho^{\lfloor U\mu \rfloor - D} > (\rho - 1)(U\mu - 1 - D) + 1. \quad (7.15)$$

We define $M = U\mu - D - 1$ and observe that

$$\rho^{\lfloor U\mu \rfloor - D} > \rho^{U\mu - D - 1} \quad (7.16)$$

(recall that $\rho > 1$). Thus, we can prove our claim by showing that for any $M \in [U\mu - \lfloor U\mu \rfloor + 1, U\mu - 1]$ and $\rho > 1$,

$$f(M) \triangleq \rho^M - (\rho - 1)M - 1 \geq 0. \quad (7.17)$$

We note that the smallest possible value of M is 1. Since $f(1) = 0$, showing that the derivative of $f(M)$ is non-negative will be the final step for proving the first case. The derivative of $f(M)$ with respect to M is given by

$$f' = \rho^M \ln(\rho) - \rho + 1. \quad (7.18)$$

Thus, we need to show that

$$\ln(\rho) \geq \frac{\rho - 1}{\rho^M}. \quad (7.19)$$

Using the First Mean Value Theorem for Integrals (Halmos, 1953), one can show that $\ln(x) \geq (x - 1)/x$, for any $x > 0$. Thus, Eq. (7.19) holds true and we conclude that the derivative of $f(M)$ is non-negative.

Next, we prove the other case where $\rho < 1$. We will show that

$$\frac{\mu - \lambda}{\mu - \lambda(\rho)^D[\rho(U\mu - 1 - D) - (U\mu - 2 - D)]} > \frac{1}{U\mu}, \quad (7.20)$$

which is equivalent to showing that

$$U(\mu - \lambda) > 1 - \rho^{D+1}[\rho(U\mu - 1 - D) - (U\mu - 2 - D)]. \quad (7.21)$$

We define $N = D + 1$ and with some algebra we get that Eq. (7.21) is equivalent to

$$U(\mu - \lambda) > 1 + \rho^N[(1 - \rho)(U\mu - N) - 1]. \quad (7.22)$$

Next, we show that the equation above holds true for $N = 1$, which is the smallest

possible value of N . We set $N = 1$ and we get:

$$U\mu(1 - \rho) > 1 + \rho[(1 - \rho)(U\mu - 1) - 1]. \quad (7.23)$$

With some algebra we get that the equation above is equivalent to:

$$U\mu - 1 > \rho(U\mu - 1). \quad (7.24)$$

The equation above holds true for any $U\mu \geq 2$ and for any $\rho < 1$. Hence, we showed that Eq. (7.20) holds true for the case $N = 1$. Next, we will show that the right hand side of Eq. (7.22) is non-increasing with N , which will conclude our proof. We define

$$g(N) = 1 + \rho^N[(1 - \rho)(U\mu - N) - 1]. \quad (7.25)$$

The derivative of $g(N)$ is

$$g' = -\rho^N \ln(\rho) + \rho^N \ln(\rho)(1 - \rho)(U\mu - N) - \rho^N(1 - \rho). \quad (7.26)$$

To show that the derivative is negative, we divide it by ρ^N and we need to show that:

$$-\ln(\rho) + \ln(\rho)(1 - \rho)(U\mu - N) - 1 + \rho \leq 0. \quad (7.27)$$

Since $\ln(\rho)(1 - \rho) < 0$ and $(U\mu - N) \geq 1$, the following holds:

$$\ln(\rho)(1 - \rho)(U\mu - N) \leq \ln(\rho)(1 - \rho). \quad (7.28)$$




Figure 7.2: A queue with $\lambda = 9.8$, $\mu = 10$ and $U = 2$. The optimal policy is u_- .

Thus, we can show that Eq. (7.27) holds true by showing that

$$-\ln(\rho) + \ln(\rho)(1 - \rho) - 1 + \rho \leq 0. \quad (7.29)$$

This is equivalent to showing that

$$\ln(\rho) \geq \frac{\rho - 1}{\rho}, \quad (7.30)$$

which holds for any $\rho > 0$, as we explained before. \square

Figure 7.2 shows the stationary idle probability when different threshold values are used in an underloaded queue with $\lambda = 9.8$, $\mu = 10$ and $U = 2$. In this case $L = 20$. Thus, we consider threshold policies from 0 to 18. We can see from the figure that the idle stationary probability of any threshold policy is always greater than u_- . For some threshold values, it is also greater than u_+ .

In Figure 7.3, we consider an overloaded queue with $\lambda = 11$, $\mu = 10$ and $U = 2$. In this case, the idle stationary probability of any threshold policy from 0 to 18 is bounded between u_- and u_+ . This time, the policy that minimizes the idle probability




Figure 7-3: A queue with $\lambda = 11$, $\mu = 10$ and $U = 2$. The optimal policy is u_+ .

is to always inform customers about the queue length, i.e. u_+ .

7.4 Summary

In this chapter, we evaluate different information sharing policies in $M/M/1$ queue. Our main goal is to find out how customers will behave under a policy which shares information when the queue length is small and hides it when the queue length is large. First, we find the equilibria structure and we show that the equilibrium is unique. Then, we prove that, in order to maximize the throughput, the service provider is always better off either sharing the information or hiding it.

Chapter 8

Concluding Remarks and Future Directions

In this dissertation, we introduce advance reservation games: games where customers must decide whether to reserve a future resource in advance. We study how different policies and prices impact the behavior of strategic customers and, in turn, the economic outcome of the system. We use a mechanism design approach to find policies and prices that optimize, at equilibrium, the objective of a service provider.

In Chapter 3, we consider a slotted loss system. We first show that, at equilibrium, either all customers that arrive before some threshold point make AR or none of them makes AR. Next, we prove the existence of one or more Nash equilibria and find the range of costs corresponding to each equilibrium. We further show that certain costs may yield more than one equilibrium. Next, we show that charging a fee from all customers attempting to reserve a server (including those not granted service) can only reduce the provider's revenue from AR fees.

In order for a provider to decide on a proper AR fee, we introduce the concept of *Price of Conservatism (PoC)* which corresponds to the ratio of the maximum possible expected profit to the maximum guaranteed expected profit. A greater PoC indicates greater potential profit loss if the provider opts to be conservative. We consider the model where charges are collected only from the customers getting service and assume

that the demand is Poisson distributed. We show that in a single server system the equilibrium is unique. Thus, $PoC = 1$ and the provider experiences no loss. In an overloaded many-server system where the average demand is $\lambda = \alpha N$ with $\alpha > 1$, the maximum possible expected profit tends to 1, while the maximum guaranteed expected profit tends to $1 - 1/\alpha$ as $N \rightarrow \infty$. Hence, $PoC = \alpha/(\alpha - 1)$. Thus, the price of conservatism increases in an unbounded fashion as α approaches 1 from above. Finally, we show that in an underloaded many-server system, the provider cannot make profit.

Extensions of the analysis of loss systems to more complex settings (e.g., with customers differing in their utilities or customers booking more than one server or slot) and analysis of PoC in other systems represent interesting directions for future work.

In Chapter 4, we study the impact of different information sharing policies on customers' behavior in loss systems. We consider three policies: *binary-information* sharing, *full-information* sharing and *partial-information* sharing. We find the equilibrium structure of each game and resort to simulations to find out which policy maximizes the number of advance reservations. The simulation results indicate that sharing more information decreases the average number of reservations. Furthermore, we observe that as the reservation cost increases, the impact of the information sharing policy on the behavior of customers increases.

Several open questions remain about the impact of information on customers' behavior in systems with advance reservations. One possible extension to this work is a game with dynamic pricing. For example, the provider can offer a reduced AR fee or a reduced service fee when the system is almost empty. Another possible extension is studying the impact of information on systems that allow cancellations. In this case, the number of available servers can both decrease or increase overtime.

In Chapter 5, we explore strategic behavior in a waiting queue that let customers reserve the server in advance. We find the equilibrium structure of this game which is similar to that in loss systems with *no-information* sharing. We then focus on the provider revenue. We show that, at equilibrium, the optimal revenue from AR fees is determined by the utilization of the queue. If the utilization is smaller than $2/3$, then the optimal fee leads to a unique equilibrium, while a higher fee leads to three equilibria, including one with no reservation. Once again, we use the PoC metric to evaluate the ratio between the maximum revenue with risk-averse pricing and with risk-taking pricing. We show that this ratio increases with the utilization. That is, as the system gets busier, the provider loses more by being risk-averse.

Studying advance reservations in other types of waiting queues represents an interesting area for future research. Studying the impact of information sharing in waiting queues could be another interesting area of research. For example, one can study a version of the game where before deciding whether to make AR or not, the customers know what will be their waiting times if making AR.

In Chapter 6, we study a dynamic version of AR games. Our goal is to shed light on the long-term behavior of customers. We explore two versions of best response dynamics. In one version, customers observe previous strategies. In this version, convergence to equilibrium is guaranteed. If multiple equilibria exist, then the game converges to either a *none-make-AR* equilibrium or to a *some-make-AR* equilibrium. The initial belief determines to which equilibrium the system will converge. In the second version, customers observe past actions rather than strategies. In this case, the game can only converge to a *none-make-AR* equilibrium (if it exists). Thus, if the provider's interest is that customers make reservations, then she should design the AR mechanism such that a *none-make-AR* equilibrium does not exist.

Future work includes studying different learning rules, such as fictitious play or

reinforcement learning. Another direction is to apply best response dynamic on AR games with full-information sharing.

In Chapter 7, we study the impact of information sharing in $M/M/1$ queues, where customers decide whether to join or balk. We seek to find out whether informing customers about the queue length when it is small and hiding this information when it is large can be beneficial for a provider aiming to maximize her revenue. We prove that such policy is never optimal. That is, a provider is better off by either always sharing information or always hiding it. Future work could investigate more complex information sharing policies. For example, sharing or hiding the information in a probabilistic fashion.

In conclusion, in this dissertation, we develop a framework that helps us to better understand the strategic behavior of customers in a system that allows advance reservations. We use this framework to analyze the impact of pricing, charging, and information sharing policies on the economic equilibria of the system and on its dynamic behavior. The analysis yields several insights that should prove useful to a service provider designing a system that allows advance reservations.

References

- Altman, E. and Shimkin, N. (1998). Individual equilibrium and learning in processor sharing systems. *Operations Research*, 46(6):776–784.
- Avineri, E. (2004). A cumulative prospect theory approach to passengers behavior modeling: waiting time paradox revisited. In *Intelligent Transportation Systems*, volume 8, pages 195–204. Taylor & Francis.
- Balachandran, K. (1972). Purchasing priorities in queues. *Management Science*, 18(5-Part-1):319–326.
- Bardhi, F. and Eckhardt, G. M. (2012). Access-based consumption: The case of car sharing. *Journal of Consumer Research*, 39(4):881–898.
- Bertsimas, D. and Shioda, R. (2003). Restaurant revenue management. *Operations Research*, 51(3):472–486.
- Brown, G. W. (1951). Iterative solution of games by fictitious play. *Activity analysis of production and allocation*, 13(1):374–376.
- Buyya, R., Yeo, C. S., Venugopal, S., Broberg, J., and Brandic, I. (2009). Cloud computing and emerging it platforms: Vision, hype, and reality for delivering computing as the 5th utility. *Future Generation computer systems*, 25(6):599–616.
- Charbonneau, N. and Vokkarane, V. M. (2012). A survey of advance reservation routing and wavelength assignment in wavelength-routed wdm networks. *IEEE Communications Surveys & Tutorials*, 14(4):1037–1064.
- Cohen, R., Fazlollahi, N., and Starobinski, D. (2009). Path switching and grading algorithms for advance channel reservation architectures. *IEEE/ACM Transactions on Networking*, 17(5):1684–1695.
- Conway, R. W., Maxwell, W. L., and Miller, L. W. (2012). *Theory of scheduling*. Courier Corporation.
- Cournot, A. A. (1897). Recherches sur les principes mathématiques de la théorie des richesses, paris 1838. *English transl. by NT Bacon under the title Researches into the Mathematical Principles of the Theory of Wealth*, New York.

- Cramton, P. C., Shoham, Y., Steinberg, R., et al. (2006). *Combinatorial auctions*. MIT press Cambridge.
- Dodge, Y. (2006). *The Oxford dictionary of statistical terms*. Oxford University Press on Demand.
- Edelson, N. M. and Hilderbrand, D. K. (1975). Congestion tolls for poisson queuing processes. *Econometrica*, 43(1):81–92.
- Erlang, A. K. (1909). The theory of probabilities and telephone conversations. *Nyt Tidsskrift for Matematik B*, 20(33-39):16.
- Fazlollahi, N. and Starobinski, D. (2015). Distance vector-based advance reservation with delay performance guarantees. *Theory of Computing Systems*, pages 1–28.
- Fu, F. and van der Schaar, M. (2009). Learning to compete for resources in wireless stochastic games. *IEEE Transactions on Vehicular Technology*, 58(4):1904–1919.
- Fudenberg, D. (1998). *The theory of learning in games*. MIT press.
- Gardiner, C. W. et al. (1985). *Handbook of stochastic methods*, volume 3. Springer Berlin.
- Geng, Y. and Cassandras, C. G. (2012). A new smart parking system infrastructure and implementation. *Procedia-Social and Behavioral Sciences*, 54:1278–1287.
- Guérin, R. A. and Orda, A. (2000). Networks with advance reservations: The routing perspective. In *IEEE INFOCOM 2000. The Nineteenth Conference on Computer Communications*, volume 1, pages 118–127. IEEE.
- Guo, P. and Zipkin, P. (2007). Analysis and comparison of queues with different levels of delay information. *Management Science*, 53(6):962–970.
- Guok, C., Robertson, D., Thompson, M., Lee, J., Tierney, B., and Johnston, W. (2006). Intra and interdomain circuit provisioning using the oscar reservation system. In *3rd International Conference on Broadband Communications, Networks and Systems, 2006. BROADNETS 2006*, pages 1–8. IEEE.
- Halmos, P. R. (1953). Measure theory, new york, 1950. *Mathematical Reviews (MathSciNet): MR11: 504d*.
- Hassin, R. (2007). Information and uncertainty in a queuing system. *Probability in the Engineering and Informational Sciences*, 21(03):361–380.
- Hassin, R. (2016). *Rational Queueing*. CRC Press.

- Hassin, R. and Roet-Green, R. (2011). Equilibrium in a two dimensional queueing game: When inspecting the queue is costly. Technical report, Working paper, Tel Aviv University, Israel.
- Hassin, R. J. and Haviv, M. (2003). *To Queue or Not to Queue: Equilibrium Behaviour in Queueing Systems*. Kluwer Academic Publishers.
- Haviv, M., Kella, O., and Kerner, Y. (2010). Equilibrium strategies in queues based on time or index of arrival. *Probability in the Engineering and Informational Sciences*, 24(1):13.
- Haviv, M. and Roughgarden, T. (2007). The price of anarchy in an exponential multi-server. *Operations Research Letters*, 35(4):421–426.
- Heydenreich, B., Müller, R., and Uetz, M. (2007). Games and mechanism design in machine scheduling an introduction. *Production and Operations Management*, 16(4):437–454.
- Hurwicz, L. (1973). The design of mechanisms for resource allocation. *The American Economic Review*, 63(2):1–30.
- Jain, R., Juneja, S., and Shimkin, N. (2011). The concert queueing game: to wait or to be late. *Discrete Event Dynamic Systems*, 21(1):103–138.
- Kaushik, N. R., Figueira, S. M., and Chiappari, S. A. (2006). Flexible time-windows for advance reservation scheduling. In *14th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer Telecommunication Systems, 2006. MASCOTS 2006*, pages 218–225. IEEE.
- Kendall, D. G. (1953). Stochastic processes occurring in the theory of queues and their analysis by the method of the imbedded markov chain. *The Annals of Mathematical Statistics*, 24(3):338–354.
- Koutsoupias, E. and Papadimitriou, C. (1999). Worst-case equilibria. In *STACS 99*, pages 404–413. Springer.
- Lakshmivarahan, S. (1981). *Learning algorithms theory and applications*. Springer-Verlag New York, Inc.
- Lieberman, V. and Yechiali, U. (1978). On the hotel overbooking problem-an inventory system with stochastic cancellations. *Management Science*, 24(11):1117–1126.
- Littman, M. L. (1994). Markov games as a framework for multi-agent reinforcement learning. In *WW Cohen and H. Hirsh (eds.) Machine Learning: Proceedings of the Eleventh International Conference, (pp. 157-163)*. New Brunswick, NJ: Rutgers University., volume 157, pages 157–163.

- Liu, Q. and van Ryzin, G. (2011). Strategic capacity rationing when customers learn. *Manufacturing & Service Operations Management*, 13(1):89–107.
- Milgrom, P. and Roberts, J. (1991). Adaptive and sophisticated learning in normal form games. *Games and economic Behavior*, 3(1):82–100.
- Myerson, R. B. (1998). Population uncertainty and poisson games. *International Journal of Game Theory*, 27(3):375–392.
- Naor, P. (1969). The regulation of queue size by levying tolls. *Econometrica*, 37(1):15–24.
- Nasiry, J. and Popescu, I. (2012). Advance selling when consumers regret. *Management Science*, 58(6):1160–1177.
- Niu, D., Feng, C., and Li, B. (2012). Pricing cloud bandwidth reservations under demand uncertainty. In *ACM SIGMETRICS Performance Evaluation Review*, volume 40, pages 151–162. ACM.
- Oh, J. and Su, X. (2012). Pricing restaurant reservations: Dealing with no-shows. Available at SSRN 2169567.
- Reiman, M. I. and Wang, Q. (2008). An asymptotically optimal policy for a quantity-based network revenue management problem. *Mathematics of Operations Research*, 33(2):257–282.
- Ross, K. (1995). *Multiservice loss networks for broadband telecommunications networks*. Springer-Verlag NY.
- Ross, S. M. et al. (1996). *Stochastic processes*, volume 2. John Wiley & Sons New York.
- Shone, R., Knight, V. A., and Williams, J. E. (2013). Comparisons between observable and unobservable M/M/1 queues with respect to optimal customer behavior. *European Journal of Operational Research*, 227(1):133–141.
- Simhon, E., Cramer, C., Lister, Z., and Starobinski, D. (2015). Pricing in dynamic advance reservation games. In *2015 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pages 546–551. IEEE.
- Smith, W., Foster, I., and Taylor, V. (2000). Scheduling with advanced reservations. In *14th International Parallel and Distributed Processing Symposium, 2000. IPDPS 2000. Proceedings.*, pages 127–132. IEEE.
- Sotomayor, B., Montero, R. S., Llorente, I. M., and Foster, I. (2009). Virtual infrastructure management in private and hybrid clouds. *IEEE Internet Computing*, 13(5):14–22.

- Srivastava, V., Neel, J. O., MacKenzie, A. B., Menon, R., DaSilva, L. A., Hicks, J. E., Reed, J. H., and Gilles, R. P. (2005). Using game theory to analyze wireless ad hoc networks. *IEEE Communications Surveys and Tutorials*, 7(1-4):46–56.
- Syed, A. A., Ye, W., and Heidemann, J. (2008). T-lohi: A new class of mac protocols for underwater acoustic sensor networks. In *IEEE INFOCOM 2008. The 27th Conference on Computer Communications*. IEEE.
- Takács, L. (1962). *Introduction to the Theory of Queues*. Oxford University Press New York.
- Tan, M. (1993). Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the tenth international conference on machine learning*, pages 330–337.
- Thorlund-Petersen, L. (1990). Iterative computation of cournot equilibrium. *Games and Economic Behavior*, 2(1):61–75.
- Trench, W. F. (2003). *Introduction to real analysis*. Prentice Hall/Pearson Education Upper Saddle River, NJ.
- Virtamo, J. T. (1992). A model of reservation systems. *IEEE Transactions on Communications*, 40(1):109–118.
- Von Neumann, J. and Morgenstern, O. (1947). *Theory of games and economic behavior*. Princeton university press.
- Xie, D., Ding, N., Hu, Y. C., and Kompella, R. (2012). The only constant is change: incorporating time-varying network reservations in data centers. *ACM SIGCOMM Computer Communication Review*, 42(4):199–210.
- Yessad, S., Nait-Abdesselam, F., Taleb, T., and Bensaou, B. (2007). R-mac: Reservation medium access control protocol for wireless sensor networks. In *32nd IEEE Conference on Local Computer Networks, 2007. LCN 2007.*, pages 719–724. IEEE.
- Zohar, E., Mandelbaum, A., and Shimkin, N. (2002). Adaptive behavior of impatient customers in tele-queues: Theory and empirical support. *Management Science*, 48(4):566–583.

CURRICULUM VITAE

Eran Simhon

Mr. Simhon was born in Haifa, Israel in 1979. He received his B.Sc. and M.Sc. degrees in Industrial Engineering and Management from The Technion - Israel Institute of Technology, Haifa, Israel in 2006 and 2012, respectively.

Mr. Simhon joined the Ph.D. program at Boston University, Systems Engineering Division in 2012. His Ph.D. research focused primarily on the economic outcomes of queueing systems that support advance reservations. He has authored a number of well-received papers on the topic. He was awarded The Hariri Institute Award for Innovative Computing Models, Algorithms and Systems at Boston University Scholars Day in 2014 and was a finalist at the Bell Labs Best Paper Award for Smart Data Pricing in 2015. During his studies at the Technion and Boston University, he was awarded teaching and research fellowships. He also has four years of work experience as a supply-chain engineer.