Provably Near-Optimal Algorithms for Multi-stage Stochastic Optimization Models in Operations Management

by

Cong Shi

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2012

Author ....
Sloan School of Management
August 7, 2012

Certified by ....
Retsef Levi
J. Spencer Standish (1945) Professor of Management
Thesis Supervisor

Accepted by ....
Dimitris Bertsimas
Co-Director, Operations Research Center

Provably Near-Optimal Algorithms for Multi-stage Stochastic Optimization Models in Operations Management

by

Cong Shi

Submitted to the Sloan School of Management on August 7, 2012, in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Operations Research

Abstract

Many if not most of the core problems studied in operations management fall into the category of *multi-stage stochastic optimization models*, whereby one considers multiple, often correlated decisions to optimize a particular objective function under uncertainty on the system evolution over the future horizon. Unfortunately, computing the optimal policies is usually computationally intractable due to *curse of dimensionality*. This thesis is focused on providing provably near-optimal and tractable policies for some of these challenging models arising in the context of inventory control, capacity planning and revenue management; specifically, on the design of *approximation algorithms* that admit *worst-case performance guarantees*.

In the first chapter, we develop new algorithmic approaches to compute provably near-optimal policies for *multi-period stochastic lot-sizing inventory models* with positive lead times, general demand distributions and dynamic forecast updates. The proposed policies have worst-case performance guarantees of 3 and typically perform very close to optimal in extensive computational experiments. We also describe a 6-approximation algorithm for the counterpart model under uniform capacity constraints.

In the second chapter, we study a class of revenue management problems in systems with *reusable resources* and *advanced reservations*. A simple control policy called the *class selection policy* (CSP) is proposed based on solving a knapsack-type linear program (LP). We show that the CSP and its variants perform provably near-optimal in the *Halfin-Whitt* regime. The analysis is based on modeling the problem as *loss network systems* with *advanced reservations*. In particular, asymptotic upper bounds on the blocking probabilities are derived.

In the third chapter, we examine the problem of capacity planning in joint ventures to meet stochastic demand in a *newsvendor-type* setting. When resources are heterogeneous, there exists a unique revenue-sharing contract such that the corresponding Nash Bargaining Solution, the Strong Nash Equilibrium, and the system optimal solution coincide. The optimal scheme rewards every participant proportionally to her marginal cost. When resources are homogeneous, there does not exist a revenue-sharing scheme which induces the system optimum. Nonetheless, we propose provably good revenue-sharing contracts which suggests that the reward should be inversely proportional to the marginal cost of each participant.

Thesis Supervisor: Retsef Levi
Title: J. Spencer Standish (1945) Professor of Management

Acknowledgments

My deepest gratitude goes to my advisor Professor Retsef Levi whose enthusiasm and scientific rigor are extremely motivating. This thesis work could not have been possible without his enormous support and guidance throughout my five years at MIT. I have been very fortunate to have him as my advisor who cares so much about my research, and who always responds to my questions and queries so promptly. These experiences will benefit me tremendously in both my career and life.

I was also very lucky to have Georgia Perakis and Vivek Farias as my two other committee members, for their continuous support, stimulating discussions and constructive feedback that has drastically improved this thesis work. I am particularly indebted to Georiga who shared her invaluable expertise and insights that have led to numerous important results in our joint project.

My special thanks goes to Wei Sun, who has been a wonderful collaborator and also a great friend. A complete chapter of this thesis is the result of countless hours of hard work together.

I would also like to thank my classmates Allison Chang, Phillip Keller, Matthew Fontana, Shubham Gupta for their great companionship. A special debt of gratitude is due to my two best friends at ORC, Yuan Zhong and Yehua Wei for supporting me in tough times and sharing my joy in good times. My five years at ORC were extremely joyful and memorable in large part due to them.

This thesis is dedicated to my parents Xiping Shi and Qiong Yao who sacrified a great deal to get me to where I stand today. They are honest and diligent who serve as my role models in life. This thesis is also dedicated to my wife Miao Ni for her love and endless support. She is the source of encouragement, enabling me to continue in rough times or to do the difficult thing in trying times. They are the people I love the most, and they are my true inspirations.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Multi-stage stochastic optimization models have been prevalent in the field of operations management, whereby operations manager considers multiple, often correlated decisions to optimize a particular objective function under uncertainty on the system evolution over the remaining time horizon. Unfortunately, for most of these models computing the optimal solutions is usually computationally intractable due to *curse of dimensionality.* Alternatively, one may resort to designing heuristics that can generate efficient solutions with possibly good quality. Some of the most successful attempts include exact and approximate dynamic programming, stochastic approximation algorithms, sampling-based methods and robust optimization.

This thesis is focused on constructing provably near-optimal and tractable policies to several core models in operations management, in particular, in the areas of inventory control, revenue management and capacity management. These algorithms are computationally tractable and admit *worst-case performance guarantees.* The notion of worst-case performance guarantees has been used extensively in computer science in the analysis of *approximation algorithms* for combinatorial NP-hard problems (Vazirani (2001)). Put formally, an algorithm is called an α-approximation algorithm or is said to have a *worst-case guarantee* of α (for some constant $\alpha > 1$) if it is a polynomial time algorithm, and for any instance of the problem the algorithm is guaranteed to provide a solution with cost that is at most α times the optimal cost.

Traditionally, approximation algorithm techniques have been applied primarily to deterministic combinatorial optimization problems. The work on approximation algorithms for stochastic combinatorial problems goes back to the work on stochastic scheduling problem of Möhring et al. (1984a) and Möhring et al. (1984b) and the more recent work of Möhring et al. (1999). Recently, there has been a growing stream of approximation results for several two-stage stochastic combinatorial problems. For a comprehensive literature review, we refer the readers to Stougie and van der Vlerk (2003), Dye et al. (2003) and Shmoys and Swamy (2004, 2006a). In contrast, this thesis is focused on the relatively harder multistage stochastic optimization models, for which there has been relatively little work (for example, see Dean et al. (2004), Shmoys and Swamy (2006b), Chan and Farias (2009), Levi et al. (2005, 2007, 2008a,d) and Levi and Radovanovic (2010)).

The concept of approximation algorithms has been applied to several problems in operations management, but again primarily to deterministic problems; for examples, see Silver and Meal (1973), Roundy (1993), and Levi et al. (2006, 2008b,c). Until recently, there have been relatively few examples of worst-case analysis of heuristics for stochastic optimization models within operations management (Chen (1999)). In fact, with relatively few exceptions (e.g. Gallego and van Ryzin (1994), Lu et al. (2006), Halman et al. (2009), Chu and Shen (2010)), most of the heuristics and algorithms that have been proposed for operations management models were evaluated merely through computational experiments on randomly generated instances. This does not necessarily provide strong indications that the proposed heuristics are good in general, beyond the instances that were actually tested. In contrast, worst-case performance analysis has the advantage that it provides a priori and posteriori guarantees on the quality of the solution produced by the algorithm. Moreover, the performance analysis provides insights on how to design algorithms that have good typical (empirical) performance, which in most cases is significantly better than the worst-case analysis.

In this thesis, we will present some of the recent work to develop provably near-optimal approximation algorithms for operations management models. We shall describe the respective algorithms and their theoretical (worst-case) and typical (computational) performance analysis. In addition, we shall highlight some of the central techniques that have been used, and point out interesting future research directions. As will be demonstrated, the respective techniques span ideas from many disciplines, such as optimization, computer science, and stochastic analysis. The discussion in this thesis is focused on three classes of models, specifically, stochastic lot-sizing problems and their capacitated counterparts (Levi and Shi (2009, 2010)), loss network systems with advanced reservations (Levi and Shi (2011b)) and modeling joint ventures in operations management (Levi et al. (2011)).

Chapter 1: Stochastic Lot-sizing Problems

We address several classical stochastic inventory control models in the presence of fixed costs. We develop the first provably near-optimal *randomized* algorithms for stochastic lot-sizing problems and capacitated stochastic lot-sizing problems which are core problems inventory theory. The goal is to coordinate a sequence of orders of a single commodity, aiming to supply stochastic demands over a discrete finite horizon with minimum total expected cost, including fixed, ordering, holding and backlogging costs.

These models capture two very important aspects of managing inventory in practice, the first being *uncertainty* and the second being *economies of scales*. First, uncertainty is a significant aspect in modeling real life situations. However, modeling uncertainty in inventory models usually makes them significantly harder to solve compared to their deterministic counterparts. Our models allow the most general exogeneous demand processes including auto-correlated and non-stationary demands as well as dynamic forecast updates. Secondly, in stochastic lot-sizing models, we also need to consider fixed cost that arises in many real-life scenarios. Fixed cost reflects the fact that ordering, production and transportation in large quantities lead to economies of scales.

The models. Stochastic inventory theory provides streamlined models with the following common setting. The goal is to coordinate a sequence of orders over a planning horizon of finitely many discrete periods, aiming to satisfy a sequence of random demands with minimum expected cost. The cost consists of a *fixed ordering cost* incurred in each period, in which a strictly positive quantity of supply units is ordered regardless of the size of the order; a per-unit *holding cost* for carrying excess inventory from one period to the next; and a per-unit *backlogging penalty cost* that is incurred in each period for each unit of unsatisfied demand. Specifically, at the beginning of each period one needs to decide how many units to order. If an order is placed then the fixed ordering cost is incurred and the order arrives after a given lead time of several periods. Then the period demand is observed and satisfied to the maximum extent possible from the inventory on hand. Excess supply or unsatisfied demand are carried to the next period incurring appropriate holding and backlogging costs, respectively. The goal is to find an ordering policy that minimizes the overall expected costs over the entire horizon. The models studied in this work capture very general demand structures. In particular, demands in different periods can be auto-correlated and the information about the joint distribution of future demands can evolve over time as more information becomes available to the decision maker. Allowing general demand structures captures many important aspects, such as forecast updates. However, it usually gives rise to very complex models since the underlying state space becomes multidimensional, even in simpler models without fixed ordering costs.

Our contributions. First, we propose a new policy that can be applied under very general assumptions, i.e., with positive lead times and general demand structures. The policy is called *randomized cost-balancing policy* and has a worst-case performance guarantee of 3. That is, the expected cost of the policy is guaranteed to be at most 3 times the optimal expected cost, regardless of the specific instance. We also propose a similar policy for a related model. This model is called the stochastic lot-sizing problem with uniform capacity constraints. The worst-case performance guarantee for this model is 6. One of the novel aspects of these policies is the use of randomized decision rules. Specifically, the policy randomly chooses among different ordering quantities. While randomized algorithms have been used extensively for many optimization problems, we are not aware of any applications to inventory control models. The worst-case analysis of these algorithms employs several novel ideas that provide new insights on the respective stochastic lot-sizing models; we believe that this will contribute to the future research on these models.

Secondly, we show how these policies can be parameterized to create a broader class of policies. A simulation based optimization is used to find the 'best' parameters per a given instance of the problem. This clearly preserves the same worst-case guarantees. Moreover, computational experiments that we conducted indicate that it can lead to near-optimal policies that perform empirically within few percentages of optimal, significantly better than the worst-case performance guarantees.

Chapter 2: Revenue Management of Reusable Resources with Advanced Reservations

We consider a class of revenue management problems that arise in systems with *reusable resources* and *advanced reservations*. The work is motivated by both traditional and emerging application domains, such as *hotel room management*, *car rental management* and *workforce management*. For instance, in hotel industries, customers make requests to book a room in the future for a specified number of days. This is called advanced reservation. Rooms are allocated to customers based on their requests, and after one customer used a room it becomes available to serve other customers. One of the major issues in these systems is how to manage capacitated pool of reusable resources over time in a dynamic environment with many uncertainties. In particular, one wishes to choose the most profitable customers to maximize the resulting revenue.

Models with reusable resources and advanced reservations are typically very hard to analyze, particularly due to the existence of advanced reservations. There has been relatively little related work both on finding provably good policies for these important models and structural properties of optimal or even practically good policies. In this chapter, we analyze the performance of conceptually and computationally simple policies. We show that they perform provably near-optimal in the *Halfin-Whitt* (see Halfin and Whitt (1981)) heavy-traffic regime. That is, the expected long-run revenue of the policy is guaranteed to obtain at least a constant fraction of the optimal revenue regardless of the input instance. Moreover, the analysis builds upon novel approaches to analyze the important class of *loss network models with advanced reservations.* The latter class of models is fundamental in the analysis of many applications in operations management, communication networks and other domains. There has been very little known about the structural properties of models with advanced reservations, and we believe that our work could open new opportunities to analyze additional models.

The models. There is a single pool of resources of integer capacity C that is used to satisfy the demands of M different classes of customers. The customers of each class arrive according to an independent Poisson process with a specific class-dependent rate. Each customer requests to reserve one unit of the capacity for a specified *service time interval* in the future according to her class.

Consider a customer of class-k arrives at the system at some random time, requesting to reserve a service time interval in the future. The time between her arrival and her requested start of service is distributed according to a *reservation distribution*, while her service time is distributed according to a *service distribution.* In this model, we assume that the reservation distribution and the service distribution are arbitrary discrete distributions that could be correlated per each customer, but are independent of the arrival process and between customers. If the request is accommodated, then upon the arrival of each customer a decision is made whether to accommodate the request. During the time a customer is served, the requested unit cannot be used by any other customer; after the service is over, the unit becomes available again to serve other customers. If the resource is reserved, the customer pays a class-dependent revenue rate per unit of service time. The resource can be reserved for an arriving customer only if upon her arrival there is at least one unit of capacity that is available (i.e., not reserved) throughout her requested service interval in the future. Specifically, a customer's request can be satisfied if the maximum number of already reserved resources over the requested service interval is smaller than the capacity C. However, customers can be rejected even if there is available capacity. Rejecting a customer now possibly enables serving more profitable customers in the future. Customers whose request is not reserved upon arrival are *lost* and leave the system. The goal is to find a feasible admission policy that maximizes the expected long-run revenue rate.

Like many stochastic optimization models, one can formulate this problem using a dynamic programming approach. However, even in special cases (e.g., no advanced reservations allowed and with exponentially distributed service times), the resulting dynamic program seems computationally intractable because the corresponding state space grows very fast. This is known as the *curse of dimensionality.* Thus, finding provably good policies is a very challenging task.

Our contributions. The contributions of this chapter are two-fold. First, we employ a simple knapsack linear program (LP) to devise a conceptually simple policy that is called the *class selection policy* (CSP). The optimal solution of the LP guides the policy regarding which classes of customers should be admitted service and which ones should be declined service. A similar policy has been analyzed before by Levi and Radovanovic (2010) for *models without advanced reservations* that are significantly easier. In fact, the analysis in Levi and Radovanovic (2010) does not carry through to models with advanced reservations. Instead, we develop an entirely new analysis that shows the policy performs provably near-optimal in the Halfin-Whitt heavy-traffic regime ($C = \rho + \beta\sqrt{\rho} + o(\rho) \to \infty$, where $\beta > 0$ is a scaling factor.) In particular, the CSP is guaranteed to obtain at least $\Phi(\beta)$ of the optimal long-run revenue in the Halfin-Whitt regime, respectively. (Note that $\Phi(\cdot)$ is the cumulative density function of a standard normal. Thus, $\Phi(\beta)$ approaches 1 when β is large.) Moreover, we propose a modified version of CSP that is guaranteed to asymptotically obtain $1 - \epsilon$ fraction of the optimal revenue, for every fixed $\epsilon > 0$.

Secondly, the analysis approaches we develop are based on modeling the problem as a loss network system with advanced reservations (specifically, a $M/G/C/C$

loss system with advanced reservations). These models are concerned with the setting in which customers arrive to the system according to a stochastic process and are being served as long as there is available capacity. Customers who find a fully utilized system are lost (see, for example, the survey paper by Kelly (1991)). We are able to derive explicit upper bounds on the steady state *blocking probability*, i.e., the probability that a random customer at steady state will find a fully utilized system, and analyze them asymptotically in the above regimes. To the best of our knowledge, there have been very few successful attempts to characterize the blocking probabilities for loss network models with advanced reservations (see, for example, Coffman-Jr et al. (1999) and Lu and Radovanovic (2007a) that studied several special cases). The assumptions in our model are fairly general: a time-homogeneous Poisson arrival process. a general finite discrete service distribution and a general finite discrete reservation distribution. Models with advanced reservations are significantly harder to analyze than those without advanced reservations. One of the major difficulties in models with advanced reservations is the fact that a randomly arriving customer effectively observes a nonhomogeneous Poisson process that is induced by the already reserved service intervals. Moreover, analyzing the blocking probability of an arriving customer requires considering the entire requested service interval instead of the instantaneous load of the system. Analyzing the load over an interval immediately introduces correlation that is challenging to analyze. The upper bound on the blocking probability is obtained by considering an identical system with infinite capacity, where all customers are admitted (a $M/G/\infty$ system with advanced reservations). The probability of having more than C customers reserved in the infinite capacity system provides an upper bound on the blocking probability in the original system; we call this the *virtual blocking probability.* Through an innovative reduction to a random walk setting, we obtain an exact analytical expression for this virtual blocking probability and then analyze it asymptotically. The analysis of the virtual blocking probability is tight and constitutes a contribution for the analysis of $M/G/\infty$ systems with advanced reservations.

Chapter 3: Joint-ventures in Operations Management

A proliferation of joint ventures has been witnessed across the globe since last decade (see Bamford et al. (2004)). A joint venture is when two or more business partners pool their resources and expertise to achieve a particular goal for a contractual period of time. Joint ventures stand in the middle ground between non-cooperative competition and merging. They provide companies with the opportunities to gain new capacity and expertise, enter related businesses or new geographic markets, gain new technological knowledge access to greater resources, and share risks with other venture partners.

The models. We consider settings where n players take part in a joint venture of capacity pooling seeking to satisfy random demand. Each player contributes one type of resource. We distinguish two types of resource pooling in joint ventures, depending on whether the resources are *heterogeneous* or *homogeneous*. When resources are heterogeneous, they are not substitutable. Thus, the effective capacity of a joint venture is limited to the the minimum level of resource contributed among all the players. In other words, the lowest contribution by one player becomes the bottleneck in planning the capacity for the joint venture. On the other hand, when resources are homogeneous, the resources pooled from all the entities are perfectly substitutable and the overall effective capacity of the joint venture is determined by summing up the individual contributions.

Consider n players building capacity (according to the different resource pooling schemes) to meet stochastic demand in a *newsvendor-type* setting. That is, stochastic demand is satisfied by the pooled capacity to the maximum extent possible. Each satisfied unit of demand incurs a revenue. Revenue-sharing contracts are very common in practice, whereby each player receives a fixed fraction of the expected collective revenue. The profit of each player is the fraction of the revenue allocated to her minus the cost. In addition, each player incurs a cost that is convex and increasing in her investment level.

For a pre-fixed revenue-sharing contract, we examine the capacity investment problem by using the concepts of Nash equilibrium from non-cooperative game theory and also Nash Bargaining Solution from cooperative game theory. These are compared to the system optimum that is obtained if all the players would act as one centrally coordinated unit. We are interested in finding an optimal revenue-sharing contract that aligns the selfish objectives and incentives of the n separate players and induces the system optimum. We also quantity the loss of efficiency (*price of anarchy*) if such a contract does not exist. Finally we study the setting in which both the revenue sharing and capacity investments are negotiated simultaneously.

Our contributions We have shown that in joint ventures with heterogeneous resource pooling, any Nash equilibrium induces an equal contribution from every player, despite of them being asymmetric. The intuition is that since the revenue received by each player depends solely on the bottleneck capacity (minimum capacity contributed by some single player) when resource-sharing is heterogeneous, any further investment beyond the bottleneck capacity only increases her cost and decreases her profit.

Although multiple Nash equilibria could exist, we show that there always exists a unique Strong Nash equilibrium. Next, we focus on a Nash Bargaining model which is a natural framework to define and design fair assignment of the capacity investment levels between multiple players. We conclude that there exists a unique revenue sharing contract such that the corresponding Nash Bargaining Solution, the Strong Nash equilibrium, and the system optimal solution coincide. This revenue sharing contract indicates that the award each party receives must be equal to the ratio of her marginal cost to the total marginal cost bore by all partners evaluated at the optimal investment level.

For joint ventures with homogeneous resource pooling, we first prove some structural properties on the effective capacity under any demand distribution with convex costs. The analysis is challenging as the investment of each player could only be determined by solving a system of implicit equations. We show that joint venture always *underinvests* as the effective capacity is always lower than that of a coordinated setting.

We then focus on quadratic-linear cost functions and show that, through an intercept-argument, the effective capacity in a joint venture with respect to any revenue sharing ratio is *at least* $1/n$ of the optimal level. Moreover, the ratio between the capacity level could be upper bounded in terms of the cost asymmetry between the two players and the revenue sharing ratio. While we show that there does not exist a fixed marginal revenue sharing contract which can coordinate the players, we propose an interval for the revenue sharing ratio which induces an outcome that is guaranteed to achieve *at least 50%* of the optimal profit for a 2-player model. This interval depends on the cost asymmetry between the two players and the demand concentration.

Next, we consider general convex cost in the homogeneous resource pooling model with an arbitrary number of asymmetric players. We show that a lower bound to the efficiency of the original setting with the nonlinear convex costs is that of a modified setting with linear costs, where the coefficients are equal to the marginal cost of each player evaluated at the Nash equilibrium of the original problem. As a result, we show that the comparative analysis on profit can be reduced to analyze the joint investment level made in the Nash and the system in the setting with linear costs.

Chapter 2

Stochastic Lot-sizing Problems

2.1 Introduction

In this paper, we develop new provably near-optimal algorithms for stochastic inventory control models with *fixed costs*, general demand distributions and dynamic forecast updates. Fixed costs arise in many real-life scenarios, and reflect the fact that ordering, production and transportation in large quantities lead to economies of scales. Specifically, we study several general variants of the classical *stochastic lot-sizing problem*. Finding optimal policies in these settings is often computationally intractable. Instead, we develop new algorithmic approaches that yield a 3-*approximation*, i.e., they have a *worst-case performance guarantee of 3*. This implies that the algorithms are guaranteed to have expected cost at most three times the optimal expected cost, regardless of the input instance.

2.1.1 Contributions

The new algorithmic and performance analysis approaches that are developed in this paper depart from the previous work of Levi et al. (2007), and provide multi-fold contributions to the study of stochastic inventory control as well as more generally to the design and analysis of randomized algorithms. The paper extends the recent stream of work to develop *cost-balancing* algorithmic techniques for computationally challenging multi-period stochastic inventory control problems. This stream of work has been initiated by Levi et al. (2007) and subsequent work (Levi et al. (2005, 2008a, 2007, 2008d)), which primarily studied stochastic inventory control problems with *no* fixed costs. The conceptual idea underlying cost-balancing based algorithms is a repeated attempt to balance opposing costs, for example, in models without fixed ordering cost one seeks to balance the cost of over-ordering (*holding cost*) and the cost of under-ordering (*backlogging cost*) based on the notion of *marginal cost accounting schemes* (Levi et al. (2005, 2007, 2008d)) (see also the discussion in Section 2.4.1).

The existence of fixed costs adds a third *nonlinear* component to the cost, and makes the cost balancing more subtle. Levi et al. (2007) did study a very special case of the model studied in this paper, in which orders arrive instantaneously and demand in each period is known deterministically at the beginning the period *before* the ordering decision is made. They proposed the *triple-balancing policy* that aims to balance the fixed ordering cost, the holding cost and the backlogging cost over each *time interval* between consecutive orders. Their policy is a 3-approximation. However, the algorithm and the worst-case analysis can be applied effectively only to models, in which there is no lag, commonly called *lead time*, from when an order is placed until it arrives. In fact, in models with positive lead times the assumption in Levi et al. (2007) is equivalent to knowing deterministically the cumulative demand over the lead time. This is clearly a *very* restrictive assumption, since in many scenarios forecasting the demand over the lead time is the major challenge. Moreover, in Section 2.3.2, we show that if this assumption does not hold, the triple-balancing policy can perform arbitrarily worse than an optimal policy. This stands in contrast to most of the analytical work done on inventory models with backlogged demand, for which the extensions from models with no lead time to models with positive lead time are often immediate.

To address the nonlinearity induced by the fixed costs, a novel *randomized decision rule* is employed to balance the expected fixed ordering costs, holding costs and backlogging costs, in each *period*. In particular, the order quantity in each period is decided based on a carefully designed randomized rule that chooses among various possible order quantities with carefully chosen probabilities. To the best of our knowledge, this is the first randomized policy proposed for stochastic inventory control policies. Levi et al. (2007) used a straightforward randomized rule for the model with no fixed costs, but merely as a 'rounding' technique to address the constraint to order in integer quantities. Unlike the triple-balancing policy that balances the costs over intervals, the newly randomized policy balances the costs in each period. Like the triple-balancing policy, the randomized cost-balancing policy proposed in this paper has a worst-case guarantee of 3, but this holds under very general assumptions, i.e., general demand distributions and positive lead times . The worst-case performance analysis of the randomized policy employs several fundamental new ideas that depart from the previous work of Levi et al. (2007). Like the previous work, the analysis is based on an *amortization* of the cost incurred by the balancing policy against the cost of an optimal policy. However, all of the previous work is entirely based on *sample-path* arguments. In contrast, the analysis in this paper is based on more subtle *averaging* arguments. We believe that the new algorithmic and analysis techniques developed in this paper will turn out to be effective in the design of provably near-optimal algorithms for other stochastic inventory control problems.

Our proposed randomized policies can be parameterized to create a broader class of policies. A simulation based optimization is used to find the 'best' parameters for a given instance of the problem. This preserves the same worst-case guarantees. Moreover, relatively extensive computational experiments that we conducted indicate that it typically leads to near-optimal policies that perform empirically within few percentages of optimal, significantly better than the worst-case performance guarantees.

In addition, the work in this paper contributes to the body of work on *randomized algorithms*. The last two decades have witnessed a tremendous growth in the area of randomized algorithms. During this period, randomized algorithms went from being a tool in computational number theory to finding widespread applications in other fields, such as data structures, geometric algorithms, graph algorithms, number theory, enumeration, parallel algorithms, approximation algorithms and online algorithms. Part of the reason why randomized algorithms are attractive is the fact that they are usually conceptually simple and computationally fast. Randomized decision rules have been used extensively to obtain approximation algorithms with worst-case guarantees for many deterministic NP-hard optimization problems, including several examples of deterministic inventory management problems (see for example, Teo and Bertsimas (1996); Levi et al. (2008c)). In addition, randomized decision rules are very common in the field of *online algorithms* (see Borodin and El-Yaniv (1998)), in which there are used to obtain algorithms with *competitive ratios*. However, in spite of the increasing use of randomized algorithms, there have been relatively few successful attempts to incorporate randomized decision rules to obtain algorithms for multistage stochastic control problems. Rust (1997) proposed random versions of successive approximations and multi-grid algorithms for computing approximate solutions to Markovian decision problems. Prandini et al. (1999) designed a randomized algorithm to obtain an estimate of the probability of aircraft conflict. Bouchard et al. (2005) studied a maturity randomization technique for approximating optimal control problems to price American put options. Shmoys and Talwar (2008) proposed a randomized 4-approximation algorithm of the *a priori Traveling Salesman Problem*. Shmoys and Swamy (2006b) gave a fully polynomial randomized approximation scheme for solving 2-stage stochastic integer optimization problems. However, the techniques developed in this paper are different and we believe they have a promising potential to apply in other multistage stochastic optimization models.

2.1.2 Literature review

The dominant paradigm in most of the existing literature has been to formulate stochastic inventory control problems (including the models studied in this paper) using a dynamic programming framework. This approach turned out to be effective in characterizing the structure of optimal policies. For many of these models, it can be shown that *state-dependent* (s, S) *policies* are optimal. The ordering decision in each period is driven by two thresholds. Specifically, an order is placed if and only if the inventory level falls below the threshold s. In addition, if an order is placed the inventory level is brought up to the threshold S. The thresholds s and S are determined based on the state of the system at the beginning of the period. Scarf (1960) and Veinott (1966) have established the optimality of (s, S) policies in models with independent demands. Cheng and Sethi (1997) have extended the optimality proof to exogenous Markov-modulated demands that capture cycles and seasonality to some extent. Gallego and Özer (2001) have shown that (s, S) policies are optimal under *advance demand information*, a demand model that allows correlation and forecast updates.

Unfortunately, the rather simple forms of these optimal policies do not usually lead to efficient algorithms for computing the optimal policies. There are very few cases, in which there are efficient algorithms to compute the optimal policies. Federgruen and Zipkin (1984) proposed an algorithm to compute the optimal stationary (s, S) policy in a model with infinite horizon and independent and identically distributed demands. Federgruen and Zheng (1991) described a simple and efficient algorithm to compute the infinite horizon optimal policy in a continuous-reviewed system with demand that is generated by a renewal process. (In this setting, (s, S) policies are equivalent to (R, Q) policies, in which one places an order of Q units, whenever the inventory level drops below R.) For other more complex variants of the model, there are currently no known exact algorithms, but only heuristics. Bollapragada and Morton (1999) proposed a simple myopic policy, assuming that the demands in different periods have the same form of distribution function with the same coefficient of variation but with different means. Gavirneni (2001) designed an efficient heuristic to compute (s, S) policies for nonstationary and capacitated model. Song and Zipkin (1993) considered uncapacitated models with exogenous Markov-modulated Poisson demand. They developed an algorithm to compute the optimal (s, S) policy using a modified value iteration approach. However, they impose strong assumptions on the structure and the size of the state space of the underlying Markov process. Gallego and Özer (2001) and Özer and Wei (2004) considered uncapacitated and capacitated inventory models with advance demand information, respectively. They proposed backward induction algorithms to numerically solve problems with a relatively short planning horizon, and conducted computational experiments to study the impact of advance demand information on the optimal policy. (In the computational experiments in Section 5, we have applied the newly proposed policies to the instances they considered.) Guan and Miller (2008b) proposed an exact and polynomial-time algorithm for the uncapacitated stochastic economic lot-sizing problem if the stochastic programming scenario tree is polynomially representable. Guan and Miller (2008a) extended these algorithms to allow backlogging. Huang and Küçükyavuz (2008) considered similar problems with random lead times. These models allow stochastic and correlated demands. The main limitation comes from the fact that the number of nodes in the stochastic programming scenario tree (the size of input) is likely to be exponentially large in the size of the planning horizon. To the best of our knowledge, all of the existing heuristics and algorithms, either lack any performance guarantees or can be applied under restrictive assumptions on the demand distributions or the input size.

2.2 The Periodic-Review Stochastic Lot-Sizing Inventory Control Problem

In this section, we provide the mathematical formulation of the stochastic lot-sizing inventory control problem. We consider a finite planning horizon of T periods indexed $t = 1, \ldots, T$. The demands over these periods are random variables, denoted by $D_1, \ldots, D_T$, and the goal is to coordinate a sequence of orders over the planning horizon to satisfy these demands with minimum cost. As a general convention, from now on we will refer to a random variable and its realization using capital and lower case letters, respectively. Script font is used to denote sets.

In each period $t = 1, \ldots T$, four types of costs are incurred, a per-unit ordering cost c_t for ordering any number of units at the beginning of period t, a per-unit holding cost h_t for holding excess inventory from period t to $t + 1$, a per-unit backlogging penalty b_t that is incurred for each unsatisfied unit of demand at the end of period t, and a *fixed ordering* cost K that is incurred in each period with strictly positive ordering quantity. Unsatisfied units of demand are usually called *backorders*. Each unit of unsatisfied demand incurs a per-unit backlogging penalty cost b_t in each period t until it is satisfied. In addition, we consider a model with a lead time of L periods between the time an order is placed and the time at which it actually arrives. We assume that the lead time is a known integer L. Following the discussion in Levi et al. (2007), we assume without loss of generality that the discount factor is equal to 1, and that $c_t = 0$ and h_t, $b_t \geq 0$, for each t.

At the beginning of each period s, we observe what is called an *information set* denoted by f_s. The information set f_s contains all of the information that is available at the beginning of time period s. More specifically, the information set f_s consists of the realized demands $d_1, \ldots, d_{s-1}$ over the interval $[1, s)$, and possibly some exogenous information denoted by $(w_1, \ldots, w_s)$. The information set f_s in period s is one specific realization in the set of all possible realizations of the random vector $F_s = (D_1, \ldots D_{s-1}, W_1, \ldots, W_s)$. The set of all possible realizations is denoted by $\mathcal{F}_s$. The observed information set f_s induces a given conditional joint distribution of the future demands $(D_s, \ldots, D_T)$. For ease of notation, D_t will always denote the random demand in period t according to the conditional joint distribution in some period $s \leq t$, where it will be clear from the context to which period s it refers. The index t will be used to denote a general time period, and s will always refer to the current period. The only assumption on the demands is that for each $s = 1, \ldots, T$, and each $f_s \in \mathcal{F}_s$, the conditional expectation $E[D_t \mid f_s]$ is well defined and finite for each period $t \geq s$. In particular, we allow non-stationary and correlation between the demands in different periods.

The goal is to find an ordering policy that minimizes the overall expected discounted fixed ordering cost, holding cost and backlogging cost. We consider only policies that are *nonanticipatory*, i.e., at time s, the information that a feasible policy can use consists only of f_s and the current inventory level. The superscripts PL and OPT will be used to refer to a given feasible policy PL and an optimal policy, respectively.

Given a feasible policy PL, the dynamics of the system are described using the following notation. Let $D_{[s,t]}$ to denote the cumulative demand over the interval $[s,t]$, i.e., $D_{[s,t]} = \sum_{j=s}^{t} D_j$.In addition, let NI_t denote the *net inventory* at the end of period t. Thus, $NI_t^+ = \max(NI_t, 0)$ and $NI_t^- = \max(-NI_t, 0)$ are net holding inventory and net backlog quantities in period t, respectively. Since there is a lead time of L periods, one also considers the *inventory position* of the system, which is the sum of all outstanding orders plus the current net inventory. Let X_t be the *inventory position* at the beginning of period t before the order in period t is placed, i.e., $X_t := NI_{t-1} + \sum_{j=t-L}^{t-1} Q_j$ (for $t = 1, \ldots, T$), where Q_j denotes the number of units ordered in period j. Similarly, let Y_t be the inventory position *after* the order in period t is placed, i.e., $Y_t = X_t + Q_t$. Note that for every possible policy PL, once the information set $f_t \in \mathcal{F}_t$ is given, the values ni_{t-1}, x_t and y_t are known, where these are the realizations of NI_{t-1}, X_t and Y_t, respectively. At the end of each period t, the costs incurred are $h_t NI_t^+$ holding cost and $b_t NI_t^-$ backlogging cost. In addition, if the order quantity $Q_t > 0$, then the fixed ordering cost K is incurred. Thus, the total cost of a feasible policy PL is

$$\mathscr{C}(PL) = \sum_{t=1}^{T} \left(h_t NI_t^{+,PL} + b_t NI_t^{-,PL} + K \cdot \mathbb{1}(Q_t^{PL} > 0) \right). \tag{2.1}$$

2.3 Triple-Balancing Policy - Bad Example

In this section, we briefly discuss the triple-balancing policy proposed by Levi et al. (2007) for a special case of the stochastic lot-sizing problem. The discussion sheds light on the limitation of this policy, and motivates the newly proposed randomized cost-balancing policy discussed in section 5. Levi et al. (2007) considered a model in which in each period $t = 1, \ldots, T$, conditioning on some information set $f_t \in \mathcal{F}_t$, the conditional distribution of future demands $(D_t, \ldots, D_T)$ is such that the demand D_t is known deterministically (i.e., with probability one). This implies that the order in period t is placed *after* the demand in that period is already known. The underlying assumption here is that at the beginning of period t, our forecast for the demand in that period is sufficiently accurate, so that we can assume that it is given deterministically. A primary example is make-to-order systems. However, this assumption does not hold if there is a positive lead time and one considers D_{t+L} instead.

2.3.1 Description of the policy

First we briefly discuss the original triple-balancing policy in Levi et al. (2007), denoted by TB. This policy is based on the following two rules.

(I) When to order. At the beginning of period t, let s be the last period in which an order is placed before t. An order is placed in period t if and only if by not placing it in period t, the cumulative backlogging cost over the interval $(s, t]$ will exceed K. Once a new order is placed, s is updated to be equal to t. Observe that since, at the beginning of each period t, the conditional joint distribution of future demands is such that D_t is known deterministically, this procedure is well-defined. Notice that an optimal policy will never incur any backlogging costs in a period when an order is placed, since the cumulative backlog quantities are known prior to placing the order.

(II) How much to order. Suppose that an order is placed in period $t < T$. Focus on the holding cost incurred by the units ordered in period t over the interval $[t, T]$. The order is set to the maximum quantity q_t^{TB}, such that the conditional expected marginal holding cost incurred does not exceed K. (The exact definition of marginal holding cost is provided in Section 4.1.)

Worst-case Analysis. The analysis in Levi et al. (2007) showed that the triple-balancing policy has a worst-case performance guarantee of 3. In particular, one can show that, for each time interval between two consecutive orders of the triple-balancing policy, the expected cost incurred by an optimal policy over that interval is at least one-third of the expected cost incurred by the triple-balancing policy over the same interval. However, this is only valid under the restrictive assumptions of no lead times and period demand known at the beginning of the period.

If the period demand is not known at the beginning of the period (or there is a positive lead time), then *(I)* above is enforced on expectation. It turns out that this policy can perform arbitrarily bad compared to an optimal policy and does not have a worst-case performance guarantee where the assumptions are dropped. As a result this policy may not be applicable in more general and realistic settings. The example that shows this fact is discussed in section 3.2.

2.3.2 A bad example

The *triple-balancing policy* can be applied in general settings and one might hope to obtain a worst-case performance guarantee in general. However, the following example shows that such guarantee fails to exist in general. Consider the following instance with infinite horizon $T = \infty$, let $h_t = h = 0$, $b_t = b = 1$, $\forall t \in \mathbb{Z}^+$, $L = 1$ and $K \in \mathbb{Z}^+$, and

$$D_t = \begin{cases} \lambda K & \text{with probability } \frac{1}{\lambda} - \frac{\epsilon}{\lambda K} \\ 0 & \text{otherwise} \end{cases}, \tag{2.2}$$

where ϵ is a positive number satisfying $0 < \epsilon < K$. Moreover, the demand drops to 0 in all periods after the first positive demand. Note that the per-unit holding cost is $h = 0$, and therefore there is no penalty for holding extra units in the inventory. The optimal policy orders λK units at the beginning of period 1. The demand λK will eventually come in some period with probability 1. Thus, the optimal cost incurs fixed ordering K only. However, if no demand has arrived, the cumulative backlogging cost is 0, and the expected backlogging cost upon not ordering is $K - \epsilon$. This implies that the policy does not place any orders before the positive demand λK occurs. Thus, the policy incurs a cost of $K + \lambda K$. If we let $\lambda \to \infty$, the cost ratio goes to ∞, indicating that the triple-balancing policy can perform arbitrarily bad compared to the optimal cost, and does not admit a worst-case guarantee. This example illustrates that the policy fails to make a good ordering decision, when there is a potential impulse in demand with a positive but small probability. Thus, the policy may incur potentially a very high backlogging cost.

2.4 Randomized Cost-Balancing Policy

One of the difficulties in the stochastic lot-sizing problem is the need to balance the nonlinear fixed ordering cost against the backlogging cost that may have large spikes because of the variability of the demands. The new policy we propose aims to strike a better balance between these costs by *randomization*. The policy is called *randomized cost-balancing policy*. To strike this balance the policy employs randomized decision rules. That is, in each period, the decision whether to order and how much to order is based on a suitably chosen randomized decision rule; the policy chooses among various order quantities with certain respective probabilities. Before the description of the new policy, we briefly discuss a marginal cost accounting scheme that is used to employ the policy. This cost accounting scheme was introduced by Levi et al. (2007).

2.4.1 Marginal cost accounting scheme

Following Levi et al. (2007), we next describe an alternative cost accounting scheme that is called *marginal cost accounting scheme*. Unlike (2.1) that decomposes the cost by periods, the main idea underlying this approach is to decompose the cost by decisions. That is, the decision in period t is associated with all costs that, after that decision is made, become unaffected by any future decision, and are only affected by future demands. This may include costs in subsequent periods.

Focus first on the holding costs and assume, without loss of generality, that units in inventory are consumed on a *first-ordered first-consumed* basis. This implies that the overall holding cost of the q_s units ordered in period s (i.e., the holding cost they incur over the entire horizon $[s,T]$) is a function only of future demands, and is unaffected by any future decisions. Specifically, the total marginal holding cost associated with the decision to order q_s units in period s is defined to be $\sum_{j=s+L}^{T} h_j \left(q_s - (D_{[s,j]} - x_s)^+\right)^+$. Note that at the time the order q_s is made, the inventory position x_s is already known and indeed the marginal holding cost is just a function of future demands. In addition, once the order in period s is determined, the backlogging cost a lead time ahead in period $s+L$, i.e., $b_{s+L}\left(D_{[s,s+L]} - (x_{s+L} + q_s)\right)^+$, is also affected only by the future demands. This leads to a marginal cost accounting scheme. For each feasible policy PL, let H_t^{PL} be the holding cost incurred by the Q_t^{PL} units ordered in period t (for $t = 1, \ldots, T$) over the interval $[t, T]$, and let Π_t^{PL} be the backlogging cost associated with period t, i.e., the cost incurred a lead time ahead in period $t + L$ ($t = 1 - L, \ldots, T - L$). That is,

$$\begin{aligned} H_t^{PL} &= H_t(Q_t^{PL}) = \sum_{j=t+L}^{T} h_j \left(Q_t^{PL} - (D_{[t,j]} - X_t)^+\right)^+, \qquad (2.3)\\ \Pi_t^{PL} &= \Pi_t(Q_t^{PL}) = b_{t+L} \left(D_{[t,t+L]} - (X_t + Q_t^{PL})\right)^+. \qquad (2.4) \end{aligned}$$

Let $\mathscr{C}(PL)$ be again the cost of the policy PL. Clearly, we have

$$\mathscr{C}(PL) = \sum_{t=1-L}^{0} \Pi_t^{PL} + H_{(-\infty,0]} + \sum_{t=1}^{T-L} \left(K \cdot \mathbb{1}(Q_t^{PL} > 0) + H_t^{PL} + \Pi_t^{PL}\right), \qquad (2.5)$$

where $H_{(-\infty,0]}$ denotes the total expected holding cost incurred by units ordered before period 1. We note that the first two expressions $\sum_{t=1-L}^{0} \Pi_t^{PL}$ and $H_{(-\infty,0]}$ are not affected by any decision (i.e., they are the same for any feasible policy and each realization of the demands) and, therefore, we will omit them. Since they are nonnegative, this will not affect our approximation results. Also, observe that without loss of generality, we can assume that $Q_t^{PL} = H_t^{PL} = 0$ for any policy PL in each period $t = T - L + 1, \ldots, T$, since nothing ordered in these periods can be used within the given planning horizon. We now can write the effective cost of a policy PL as

$$\mathscr{C}(PL) = \sum_{t=1}^{T-L} \left(K \cdot \mathbb{1}(Q_t^{PL} > 0) + H_t^{PL} + \Pi_t^{PL}\right). \qquad (2.6)$$

2.4.2 Description of the policy

To describe the new policy, we modify the definition of the information set f_t to also include the randomized decisions of the randomized balancing policy up to period $t - 1$. Thus, given the information set f_t, the inventory position at the beginning of period t is known. However, the order quantity in period t is still *unknown* because the policy randomizes among various order quantities. We denote the randomized cost-balancing policy by RB. The decision in each period, whether to order and how much to order, is based on the following quantities.

- Compute the *balancing quantity* $\hat{q}_t$ which balances the conditional expected marginal holding cost incurred by the units ordered against the conditional expected backlogging cost in period $t + L$. That is, $\hat{q}_t$ solves

$$E\left[H_t^{RB}(\hat{q}_t) \mid f_t\right] = E\left[\Pi_t^{RB}(\hat{q}_t) \mid f_t\right], \tag{2.7}$$

  where H_t^{RB} and Π_t^{RB} are defined as in Section 4.1, respectively. Let $\theta_t = \theta_t(f_t) \triangleq E[H_t^{RB}(\hat{q}_t) \mid f_t] = E[\Pi_t^{RB}(\hat{q}_t) \mid f_t]$ denote the *balancing cost.* The solution to (2.7) is unique and can be computed efficiently via bi-section search (Levi et al. (2007)).

- Compute the *holding-cost-K quantity* $\tilde{q}_t$ that solves $E[H_t^{RB}(\tilde{q}_t) \mid f_t] = K$, i.e., $\tilde{q}_t$ is the order quantity that brings the conditional expected marginal holding cost to K. Note that $\tilde{q}_t$ can be computed readily since $E[H_t^{RB}(\cdot) \mid f_t]$ is monotonically increasing.

- Compute $E[\Pi_t^{RB}(\tilde{q}_t) \mid f_t]$, i.e., the resulting conditional expected backlogging cost in period $t + L$ if one orders the holding-cost-K quantity $\tilde{q}_t$ units in period t.

- Compute $E[\Pi_t^{RB}(0) \mid f_t]$, i.e., the conditional expected backlogging cost in period $t + L$ resulting from not ordering in period t.

Based on the above quantities computed, the following randomized rule is used in each period t. Let P_t denote our ordering probability which is a priori random. With the observed information set f_t, the ordering probability $p_t = P_t \mid f_t$ in period t is defined differently in the two cases below.

Case (I)

If the balancing cost exceeds K, i.e., $\theta_t \geq K$, the RB policy orders the balancing quantity $q_t^{RB} = \hat{q}_t$ with probability $p_t = 1$. The intuition is that when $\theta_t \geq K$, the fixed ordering cost K is less dominant compared to marginal holding and backlogging costs. Moreover, if the RB policy does not place an order, the conditional expected backlogging cost is potentially large. Thus, it is worthwhile to order the balancing quantity $q_t^{RB} = \hat{q}_t$ with probability $p_t = 1$.

Case (II)

If the balancing cost is less than K, i.e., $\theta_t < K$, the RB policy orders the holding-cost-K quantity (i.e., $q_t^{RB} = \tilde{q}_t$) with probability p_t and nothing with probability $1 - p_t$. That is,

$$q_t^{RB} = \begin{cases} \tilde{q}_t, & \text{with probability } p_t \\ 0, & \text{with probability } 1 - p_t \end{cases} . \tag{2.8}$$

The probability p_t is computed by solving the following equation

$$p_t K = p_t \cdot E[\Pi_t^{RB}(\tilde{q}_t) \mid f_t] + (1 - p_t) \cdot E[\Pi_t^{RB}(0) \mid f_t]. \tag{2.9}$$

The underlying reason behind the choice of this particular randomization in (2.9) is that the policy perfectly balances the three types of costs, namely, the marginal holding cost, the marginal backlogging cost and the fixed ordering cost associated with the period t. In particular, since we order the holding-cost-K quantity with probability p_t and nothing with probability $1 - p_t$, the conditional expected marginal holding cost in this case is

$$E[H_t^{RB}(q_t^{RB}) \mid f_t] = p_t E[H_t^{RB}(\tilde{q}_t) \mid f_t] + (1 - p_t) E[H_t^{RB}(0) \mid f_t] = p_t K. \tag{2.10}$$

By the construction of p_t in (2.9), the conditional expected backlogging cost is

$$E[\Pi_t^{RB}(q_t^{RB}) \mid f_t] = p_t E[\Pi_t^{RB}(\tilde{q}_t) \mid f_t] + (1 - p_t) E[\Pi_t^{RB}(0) \mid f_t] = p_t K. \tag{2.11}$$

Since p_t is the ordering probability in Case (II), the expected fixed ordering cost is $p_t K$. It can be shown that (2.9) has the following solution,

$$0 \leq p_t = \frac{E[\Pi_t^{RB}(0) \mid f_t]}{K - E[\Pi_t^{RB}(\tilde{q}_t) \mid f_t] + E[\Pi_t^{RB}(0) \mid f_t]} < 1. \tag{2.12}$$

The inequalities in (2.12) follows from the fact that $\theta_t < K$ and $\tilde{q}_t > \hat{q}_t$, which implies that $E[\Pi_t^{RB}(\tilde{q}_t) \mid f_t] < E[\Pi_t^{RB}(\hat{q}_t) \mid f_t] = \theta_t < K$. Figure 2-1 illustrates how the *RB* policy computes the ordering probability p_t in Case (II) where $\theta_t < K$.

This concludes the description of the *RB* policy. In the next section, we shall show that the *RB* policy has an expected worst-case performance guarantee of 3.

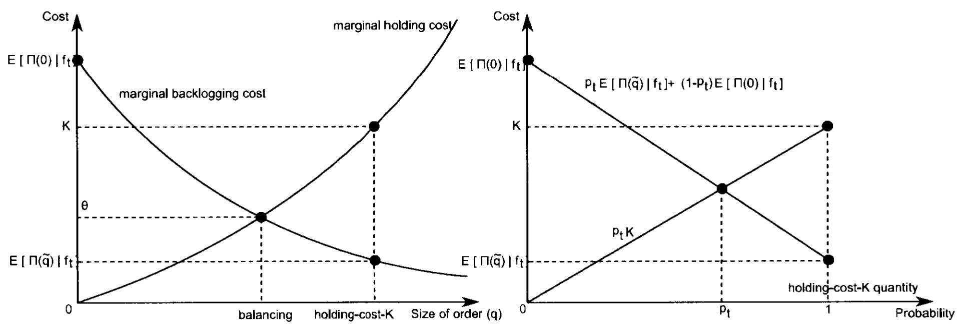

Figure 2-1: A graphical depiction of how the *RB* policy computes the probability of ordering p_t when the balancing cost θ is below the fixed ordering cost K (Case (II)).

2.4.3 Worst-case analysis

To obtain a 3-approximation, one wishes to show that on expectation the cost of an optimal policy can 'pay' for at least one-third of the expected cost of the randomized cost-balancing policy. The periods are decomposed into subsets in which we will define explicitly. For certain well-behaved subsets, we want to show that the holding and backlogging costs incurred by an optimal policy can 'pay' for one-third of the cost incurred by the *RB* policy. The difficulty arises in analyzing the remaining subset of *problematic periods*, for which it is not a priori clear how to 'pay' for their cost. These problematic periods are further partitioned into intervals defined by each two consecutive orders placed by the optimal policy. It can be shown that the total expected cost incurred by the RB policy in problematic periods within each interval, does not exceed $3K$. This implies that the fixed ordering cost incurred by an optimal policy can 'pay' on expectation one-third of the cost incurred by the randomized cost-balancing policy in problematic periods. Next we discuss the details of this approach, and we defer all proofs to Electronic Companion for ease of presentation.

Let Z_t^{RB} be a random variable defined as

$$Z_t^{RB} := E[H_t^{RB}(Q_t^{RB}) \mid F_t] = E[\Pi_t^{RB}(Q_t^{RB}) \mid F_t]. \tag{2.13}$$

Note that Z_t^{RB} is a random variable that is realized with the information set in period t. Observe that by the construction of the RB policy, the random variable Z_t^{RB} is well-defined since the expected marginal holding costs and the expected marginal backlogging costs are always balanced. That is, the conditional expected marginal holding cost is always equal to the conditional expected backlogging cost. In the following lemma we show that the expected cost of the RB policy can be upper bounded using the Z_t^{RB} variables defined in (2.13).

Lemma 2.4.1 *Let $\mathscr{C}(RB)$ be the total cost incurred by the RB policy. Then we have,*

$$E[\mathscr{C}(RB)] \leq 3 \cdot \sum_{t=1}^{T-L} E[Z_t^{RB}]. \tag{2.14}$$

To complete the worst-case analysis, we would like to show that the expected cost of an optimal policy denoted by OPT is at least $\sum_{t=1}^{T-L} E[Z_t^{RB}]$. This will be done by amortizing the cost of OPT against the cost of the RB policy. In particular, we shall show that on expectation OPT pays for a large fraction of the cost of the RB policy. In the subsequent analysis, we will use a random partition of periods $t = \{1, 2, \dots T - L\}$ to the following sets:

The set $\mathscr{T}_{1H} \triangleq \left\{t : \Theta_t \geq K \text{ and } Y_t^{OPT} > Y_t^{RB}\right\}$ consists of periods in which the balancing cost Θ_t exceeds K and the optimal policy had higher inventory position than that of the RB policy after ordering (recall that if $\Theta_t \geq K$ then the RB policy orders the balancing quantity with probability 1 and the value Y_t^{RB} is known deterministically (i.e., realized) with F_t).

The set $\mathscr{T}_{1\Pi} \triangleq \{t : \Theta_t \geq K \text{ and } Y_t^{OPT} \leq Y_t^{RB}\}$ consists of periods in which the balancing cost exceeds K and the inventory position of the optimal policy does not exceed that of the RB policy after ordering (see the comment above regarding $\mathscr{T}_{1H}$).

The set $\mathscr{T}_{2H} \triangleq \left\{t : \Theta_t < K \text{ and } Y_t^{OPT} \geq X_t^{RB} + \tilde{Q}_t^{RB}\right\}$ consists of periods in which the balancing cost is less than K and, in such periods, the inventory position of the RB policy after ordering would be either X_t^{RB} if no order was placed, or $X_t^{RB} + \tilde{Q}_t^{RB}$ if the *holding-cost-K quantity* is ordered, depending on the randomized decision of the RB policy. However, the inventory position of OPT after ordering exceeds even $X_t^{RB} + \tilde{Q}_t^{RB}$. (Note again that the quantity $\tilde{Q}_t^{RB}$ is known deterministically (i.e., realized) with F_t.)

Analogous to $\mathscr{T}_{2H}$, the set $\mathscr{T}_{2\Pi} \triangleq \{t : \Theta_t < K \text{ and } X_t^{RB} \geq Y_t^{OPT}\}$ consists of periods in which the inventory position of OPT after ordering is below X_t^{RB}.

The set $\mathscr{T}_{2M} \triangleq \left\{t : \Theta_t < K \text{ and } X_t^{RB} < Y_t^{OPT} < X_t^{RB} + \tilde{Q}_t^{RB}\right\}$ consists of periods in which the balancing cost is less than K and the inventory position of OPT after ordering is within $(X_t^{RB}, X_t^{RB} + \tilde{Q}_t^{RB})$. Thus, whether the RB policy or OPT has more inventory depends on whether the RB policy placed an order.

Note that the sets $(\mathscr{T}_{1H} - \mathscr{T}_{2M})$ are disjoint and the union makes a complete set. Conditioning on f_t, it is already known which part of the partition period t belongs.

Next we will show that the total holding cost incurred by OPT is higher than the marginal holding cost incurred by the RB policy in periods that belong to $\mathscr{T}_{1H} \bigcup \mathscr{T}_{2H}$, and that the total backlogging cost incurred by OPT is higher than the backlogging cost incurred by the RB policy associated with periods within $\mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi}$.

Lemma 2.4.2 *The overall holding cost and backlogging cost incurred by OPT are denoted by H^{OPT} and Π^{OPT}, respectively. Then we have, with probability 1,*

$$H^{OPT} \geq \sum_t H_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H}), \ \Pi^{OPT} \geq \sum_t \Pi_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi}). \quad (2.15)$$

Note that the periods in the set $\mathscr{T}_{2M}$ introduce some uncertainties in the relation between the inventory positions after ordering of the RB policy and OPT. Thus, we are unable to carry out an analysis similar to Lemma 2.4.2. For this reason, we call $\mathscr{T}_{2M}$ a *problematic* set of periods. Naturally, we also define the *non-problematic* set of periods to be $\mathscr{T}_N = \mathscr{T}_{1H} \bigcup \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi} \bigcup \mathscr{T}_{2H}$. The analysis of the problematic periods in the set $\mathscr{T}_{2M}$ will be done in two steps. In the first step, we will conceptually create a *bank account* A that will be used to pay some of the cost of the RB policy in these problematic periods. In particular, for each period $t \in \mathscr{T}_{2M}$, we *borrow* an amount of Z_t^{RB} from the bank account. Thus, the total amount of borrowing from the bank is given by $A = \sum_{t \in \mathscr{T}_{2M}} Z_t^{RB}$, and so $E[A] = E\left[\sum_t Z_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{2M})\right]$.

The following lemma shows that, with the borrowed amount A from the bank, the overall holding cost and backlogging cost incurred by OPT exceed $\sum_{t=1}^{T-L} E[Z_t^{RB}]$. The next step will be to show that $E[A]$ is at most the expected fixed ordering cost incurred by OPT. That is,

$$E[A] \leq E\left[\sum_{t=1}^{T-L} K \cdot \mathbb{1}(Q_t^{OPT} > 0)\right]. \tag{2.16}$$

Lemma 2.4.3 *The expected holding cost and backlogging cost incurred by OPT plus the expected amount borrowed from the bank account A are at least $\sum_{t=1}^{T-L} E[Z_t^{RB}]$. That is, the following inequality holds*

$$E\left[\left(H^{OPT} + \Pi^{OPT}\right) + A\right] \geq \sum_{t=1}^{T-L} E[Z_t^{RB}]. \tag{2.17}$$

By Lemmas 2.4.1 and 2.4.3, the overall holding and backlogging costs incurred by OPT, *plus* the borrowed amount A from the bank, account on expectation for *one-third* of the overall expected costs incurred by the RB policy. To complete the worst-case analysis, we will show in Lemma 2.4.4 that the expected amount borrowed from the bank account does not exceed the expected fixed ordering cost incurred by OPT, i.e., $E\left[\sum_{t=1}^{T-L} K \cdot \mathbb{1}(Q_t^{OPT} > 0)\right]$. We will highlight the key steps involved in proving this lemma. We decompose the problematic periods in the set $\mathscr{T}_{2M}$ into intervals between ordering points of OPT, and we want to show that, for each such interval, the fixed ordering cost K incurred by OPT will cover the expected amount borrowed from the bank in periods that belong to set $\mathscr{T}_{2M}$. Conditioning on f_T^- (the entire evoluation of the system excluding the randomized decisions of the RB policy), we construct a decision tree based on the randomized decisions of the RB policy. We then show that, by a tree traversal argument and Lemma 2.4.5, the expected borrowing from the problematic nodes (which belong to the set $\mathscr{T}_{2M}$) within an interval between ordering points of OPT does not exceed K.

Lemma 2.4.4 *The following inequality holds*

$$E\left[A\right] \leq E\left[\sum_{t=1}^{T-L} K \cdot \mathbb{1}(Q_t^{OPT} > 0)\right]. \tag{2.18}$$

In other words, the expected borrowing $E[A]$ is less than the total expected fixed ordering cost incurred by OPT.

Lemma 2.4.5 *Let $\{p_l\}_{l=1}^{\infty}$ satisfy the condition $0 \leq p_l \leq 1$ for all l. Then the following inequality holds,*

$$p_1^2 + \sum_{l=2}^{\infty} \left\{ \left(\prod_{s=1}^{l-1} (1-p_s) \right) p_l \left(\sum_{k=1}^{l} p_k \right) \right\} \leq 1. \tag{2.19}$$

As an immediate consequence of Lemmas 2.4.3 and 2.4.4, we obtain the following lemma and theorem.

Lemma 2.4.6 *Let $\mathscr{C}(OPT)$ be the total cost incurred by the cost-balancing policy RB. Then we have,*

$$E[\mathscr{C}(OPT)] \geq \sum_{t=1}^{T-L} E[Z_t^{RB}]. \tag{2.20}$$

Theorem 2.4.7 *For each instance of the stochastic lot-sizing problem, the expected cost of the randomized cost-balancing policy RB is at most three times the expected cost of an optimal policy OPT, i.e.,*

$$E[\mathscr{C}(RB)] \leq 3 \cdot E[\mathscr{C}(OPT)]. \tag{2.21}$$

2.5 Numerical Experiments

The randomized cost-balancing policies described above can be parameterized to obtain general classes of policies, respectively. The worst-case analysis discussed above can then be viewed as choosing parameter values that perform well against any possible instance. In contrast, find the 'best' parameter values, for each given instance. This gives rise to policies that have at least the same worst-case performance guarantees, but are likely to work better empirically, since we can refine the parameters according to the specific instance being solved. Using simulation based optimization, we have implemented this approach and tested the empirical performance of the resulting policies. The policies were tested under the model of *advanced demand information* proposed by Gallego and Özer (2001) and Özer and Wei (2004). To the best of our knowledge, these are the few papers that report computational results (by brute-force backward induction algorithm) on the stochastic lot-sizing problem with correlated demands.

2.5.1 Parameterized policies.

We describe a class of parameterized policies involving parameters β, γ and η where β controls the holding-cost-βK quantity, γ controls the ratio of marginal holding costs and backlogging costs and η controls the level of expected backlogging cost resulting from not ordering.

- The balancing quantity $\hat{q}_t$ that solves $E[H_t^{RB}(\hat{q}_t) \mid f_t] = \gamma \cdot E[\Pi_t^{RB}(\hat{q}_t) \mid f_t] := \theta_t$.
- The holding-cost-βK quantity $\tilde{q}_t$ that solves $E[H_t^{RB}(\tilde{q}_t) \mid f_t] = \beta \cdot K$.
- Compute $E[\Pi_t^{RB}(\tilde{q}_t) \mid f_t]$, and $\eta \cdot E[\Pi_t^{RB}(0) \mid f_t]$.

(I) If $\theta_t \geq \beta \cdot K$, the RB policy orders $q_t^{RB} = \hat{q}_t$ with probability $p_t = 1$ in period t.

(II) If $\theta_t < \beta \cdot K$, the RB policy orders $q_t^{RB} = \tilde{q}_t$ with probability p_t and order nothing with probability $1 - p_t$ in period t, where the probability $p_t =$

$$\frac{\eta \cdot E[\Pi_t^{RB}(0) \mid f_t]}{\beta \cdot K - E[\Pi_t^{RB}(\tilde{q}_t) \mid f_t] + \eta \cdot E[\Pi_t^{RB}(0) \mid f_t]}.$$

Since T is relatively small, we also introduce an end-of-horizon rule. Suppose we are in period t, we estimate the total expected cumulative backlogging cost (assuming no orders are placed) over the interval $[t, T]$. If this amount is less than K, we do not order in period t.

2.5.2 Experiment design

Under advance information model, the demand vector in each period t is observed as $D_t = (D_{t,t}, \ldots, D_{t,t+N})$ where $D_{t,s}$ represents order placed by customers during period t for future periods $s \in \{t, \ldots, t+N\}$ and N is the length of the information horizon over which we have advance demand information. Note that D_t is a random vector and is realized only at the end of period t. At the beginning of period t, the demand to prevail in a future period s ($s \geq t$) can be divided into two parts: the observed demand vector $\sum_{r=s-N}^{t-1} D_{r,s}$ and the unobserved demand vector $\sum_{r=t}^{s} D_{r,s}$. As a result, this introduces a correlation between period demands (however the conditional joint distribution of the future demands is known in each period t). The state space of the proposed dynamic programming formulation contains the inventory position and the observed demand vector which explodes exponentially with the length of the information horizon N when $N > L + 2$. Gallego and Özer (2001) verified some structural properties of the dynamic program via numerical studies for a number of small instances. The experiments that we performed expand their numerical studies by incorporating non-zero lead times as well as longer planning horizons. Following the methodology of Aviv and Federgruen (2001), we generated a total of 90 instances to test the quality of the randomized-balancing heuristics compared to the optimal cost. The instances we used have the following combination of parameters: $T = 12, 15$, $L = 0, 1, 2$, $N = L + 2$, $K = 0.5, 50, 100$, $h = 1, 2, 3, 6$, $p = 1, 3, 6, 9$ and $(D_{t,t}, D_{t,t+1}, D_{t,t+2})$ are modeled by Poisson random variables with mean $\lambda_0, \lambda_1, \lambda_2$.

2.5.3 Algorithmic complexity

We describe the procedures of finding the optimal parameters for a specific instance of the problem. First, assume that there exists a positive constant U such that the optimal parameters β^*, γ^*, η^* are upper bounded by U. In addition, we discretize U with some step-size Δ, i.e., $\beta, \gamma, \eta \in [0, U]$ can only take values as integer multiples of Δ. Then we conduct an exhaustive search on a cube of $U \times U \times U$ for the parameters β, γ and η. In our numerical studies, $U = 10$ and $\Delta = 0.1$ are chosen to be the upper bound and the resolution for discretization, respectively. The algorithm runs on every point on this cube, simulates the cost of each parameterized policy and returns the best possible (β^*, γ^*, η^*) that minimize the cost. Secondly, assume that there exists a positive constant $\hat{U}$ that serves an upper bound on the balancing and hold-cost-K quantities. For each $t = 1, \ldots, T$, the complexity for evaluating marginal holding cost is $O(T)$ and the complexity for carrying out bisection search is $O(\log \hat{U})$. The algorithm runs in $O(T^2 \log \hat{U})$, for each set of parameters (β, γ, η). Hence, the algorithm that returns both the optimal parameters and the lowest cost runs in $O(U^3 \Delta^{-3} T^2 \log \hat{U}) \approx O(T^2)$ since $U^3 \Delta^{-3} \log \hat{U}$ is some positive constant. For all tested instances with $T = 12$, the average CPU time per test instance on a Pentium 1.58GHz PC is 233s. In contrast, the dynamic programming algorithm takes 1840s on average per test instance.

2.5.4 Numerical results

The numerical results with $(T, L) = (12, 0)$, $(T, L) = (12, 2)$ and $(T, L) = (15, 0)$ are tabulated in Table A.1, Table A.2 and Table A.3, respectively (refer to Electronic Companion). The $(*)$ in both tables indicates that the designated parameters can take arbitrary numbers without affecting the optimal values of the parameterized policy. It is observed that $(\beta^*, \eta^*) = (*, *)$ in all instances where $K = 0$, since the holding-cost-$\beta^* K$ quantity is trivially 0 and therefore the algorithm only considers the balancing quantities. In some instances where K is relatively large and the holding-cost-$\beta^* K$ quantity is near-optimal, it is observed that $\gamma^* = (*)$ implying that the algorithm only orders the holding-cost-$\beta^* K$ quantities. For the rest of instances, the algorithm uses both the balancing quantity and the holding-cost-K quantity.

In the case where $L = 0$, on average the parameterized RB policy performs within 4.6% and always within 7% of the optimal cost for $T = 12, 15$. The numerical results show that the performance of the parameterized RB policy is insensitive to the planning horizon T. Moreover, the optimal parameters in the parameterized RB policy are intuitive: β controls the quantity of each order; γ controls the ratio in which the marginal holding cost is balanced against the marginal backlogging cost; η controls the weight put on the do-nothing backlogging cost resulted from not ordering. The optimal $\eta^* = 9$ coincides with the ratio of p to h, which implies that more weight should be put on backlogging cost so that the ordering probability can be increased. The optimal $\gamma^* = 2$ suggests that the marginal holding cost should be twice the backlogging cost. The optimal β^* is close to 1 when K is large, implying that using the holding-cost-K quantity is near optimal. The unparameterized RB policy (i.e., $(\beta, \gamma, \eta) = (1, 1, 1)$) performs on average within 27% and always within 50% error of optimal cost, which is significantly better than the theoretical worst-case performance guarantee of 3. The cost ratio is observed to be decreasing in the magnitude of fixed ordering cost K. In the case where $L = 2$, the parameterized RB policy performs on average within 10% and always within 16% error of the optimal cost. The optimal parameters are similar to those in $L = 0$. The deviation from the optimal cost is resulted from stocking more inventory units by the RB policy, as the lead time induces more uncertainty in future demands. The unparameterized RB policy performs within 50% (on average 29%) error of optimal cost. It is also noted that the average CPU time of running the RB policy is insensitive to the planning horizon T.

2.6 Capacitated Stochastic Lot-sizing Problem

We develop new algorithmic approaches to compute provably near-optimal policies for multiperiod, uniform-capacitated, stochastic lot-sizing inventory models with stochastic, non-stationary and correlated demands that evolve over time. Our approach is computationally efficient and guaranteed to produce a policy with worst-case performance guarantee of 6. We also characterize a class of parameterized policies based on this approach. Empirical studies show that these policies perform close to optimal in computational experiments, which is significantly better than the worst-case guarantees.

2.6.1 Marginal backlogging cost accounting

In capacitated model, it is no longer true that a mistake of ordering too little in the current period can always be fixed by decisions made in the future periods. Levi, Roundy, Shmoys and Truong Levi et al. (2008d) proposed a new backlogging cost accounting that associates with decision of how many units to order in period t what is called *forced backlogging cost* resulting from this decision in future periods.

Consider some period t. Suppose that x_t is the inventory position at the beginning of period t and that the number of units ordered in period is $q_t < u$. Let $\bar{q}_t$ be the resulting unused *slack capacity* in period t, i.e., $\bar{q}_t = u - q_t > 0$. Focus now on some future period $s \geq t + L$ when this order arrives and becomes available. Suppose that for some realization of the demands. We have that

$$d_{[t,s]} - (x_t + q_t + \sum_{j \in (t, s-L]} u) > 0. \tag{2.22}$$

This implies that there exists a shortage in period s, and moreover, even if in each period after period t and until period $s-L$ the orders placed were up to the maximum available capacity, this part of the shortage in period s would still exist and incur the corresponding backlogging cost. The actual shortage may be even bigger and equal to

$$d_{[t,s]} - (x_t + q_t + \sum_{j \in (t, s-L]} q_j) > 0, \tag{2.23}$$

(recall that $q_j \leq u$ for each period j). In other words, given our decision in period t, this part of the shortage could not be avoided by any decision made over the interval

$(t, s-L]$ (clearly, any order placed after period $s-L$ will not be available by time s). We conclude that, if more units had been ordered in period t, then at least some of the shortage in period s could have been avoided. More precisely, the maximum number of units of shortage that could have been avoided by ordering more units in period t is equal to

$$\min\left\{\bar{q}_t, \left[d_{[t,s]} - (x_t + q_t + \sum_{j\in(t,s-L]} u)\right]^+\right\}. \tag{2.24}$$

The intuition is that by ordering more units in period t, we could have averted part of the shortage in period s, but clearly not more than the unused slack capacity $\bar{q}_t$, since we could not have ordered in period t more than additional $\bar{q}_t$ units. In this case, we would say that this part of the backlogging cost in period s was forced by the decision in period t, and hence period t is associated with a backlogging penalty of

$$b_s \min\left\{\bar{q}_t, \left[d_{[t,s]} - (x_t + q_t + \sum_{j\in(t,s-L]} u)\right]^+\right\}. \tag{2.25}$$

This is significantly different from the traditional backlogging cost accounting, in which this cost would be associated with period $s-L$. Denote $W_{s,[1,t]}$ as the backlogging cost in period s associated with periods $[1, t]$. Then we can write

$$\begin{aligned} W_{ts} &= \min\left\{W_{s,[1,t]}, p_s(u - q_t)\right\} \qquad (2.26)\\ &= \min\left\{p_s\left(D_{[t,s]} - (X_t + Q_t + \sum_{j=t+1}^{s} u)\right)^+, p_s(u - Q_t)\right\}. \end{aligned}$$

The marginal backlogging costs that are incurred by any feasible policy P is given by

$$\bar{\Pi}_t^P = \sum_{s=t+L}^{T} W_{ts}^P. \tag{2.27}$$

2.6.2 Description of the policy

We consider the forced marginal backlogging cost accounting and the corresponding cost it associates with period s. let $E[\bar{\Pi}_s^{RB}(q_s^{RB})|f_s]$ be the expected backlogging cost associated with period s by the forced marginal backlogging cost accounting scheme described above, again conditioned on the observed information set f_s.

The decision in each period, whether to order and how much to order, is based on the following quantities.

- Compute the *balancing quantity* $\hat{q}_t$ which balances the conditional expected marginal holding cost incurred by the units ordered against the conditional expected backlogging cost in period $t + L$. That is, $\hat{q}_t$ solves

$$E[H_t^{RB}(\hat{q}_t) \mid f_t] = E[\bar{\Pi}_t^{RB}(\hat{q}_t) \mid f_t] \triangleq \theta_t, \tag{2.28}$$

  where H_t^{RB} and $\bar{\Pi}_t^{RB}$ are defined as in (2.4) and (2.27), respectively. Let $\theta_t = \theta_t(f_t) \triangleq E[H_t^{RB}(\hat{q}_t) \mid f_t] = E[\bar{\Pi}_t^{RB}(\hat{q}_t) \mid f_t]$ denote the *balancing cost.* Since $\bar{\Pi}_t^{RB}(u) = 0$, it follows that the quantity $\hat{q}_t \leq u$. The existence and uniqueness of solution to (2.7) have been shown in Levi et al. (2008d). It has also been shown in Levi et al. (2008d) that $\hat{q}_t$ can be computed efficiently via bi-section search.

- Compute the *holding-cost-K quantity* $\tilde{q}_t$ that solves

$$E[H_t^{RB}(\tilde{q}_t) \mid f_t] = K. \tag{2.29}$$

  That is, $\tilde{q}_t$ is the order quantity that brings the conditional expected marginal holding cost to K. Since $E[H_t^{RB}(q_t) \mid f_t]$ is monotone and continuous and goes to infinity as q_t goes to infinity, it is straightforward to compute $\tilde{q}_t$. Note that in computing $\tilde{q}_t$, we temporarily ignore the capacity constraint u in each period t.

- Compute the resulting conditional expected marginal backlogging cost in period

$t + L$ if one orders the minimum of $\tilde{q}_t$ and the capacity u in period t, denoted by ϕ_t. That is,

$$\phi_t = E[\bar{\Pi}_t^{RB}(\min\{\tilde{q}_t, u\}) \mid f_t]. \tag{2.30}$$

- Compute the conditional expected marginal backlogging cost in period $t + L$ resulting from not ordering in period t, denoted by ψ_t. That is,

$$\psi_t = E[\bar{\Pi}_t^{RB}(0) \mid f_t]. \tag{2.31}$$

Based on the above quantities computed, the following randomized rule is used in each period t (we assume f_t is the observed information set).

(I) If the balancing cost exceeds K, i.e., $\theta_t \geq K$, the RB policy orders the balancing quantity $q_t^{RB} = \hat{q}_t$ with probability $p_t = 1$.

(II) If the balancing cost is less than K, i.e., $\theta_t < K$, the RB policy orders in period t the holding-cost-K quantity (i.e., $q_t^{RB} = \tilde{q}_t$) with probability p_t and nothing with probability $1 - p_t$. That is,

$$q_t^{RB} = \begin{cases} \min\{\tilde{q}_t, u\}, & \text{with probability } p_t \\ 0, & \text{with probability } 1 - p_t \end{cases}. \tag{2.32}$$

The probability $p_t = p_t \mid f_t$ is computed by solving the following equation

$$p_t K = p_t \phi_t + (1 - p_t)\psi_t. \tag{2.33}$$

It can be shown that Equation (2.33) has the following solution,

$$0 \leq p_t = \frac{\psi_t}{K - \phi_t + \psi_t} < 1. \tag{2.34}$$

The inequalities in Equation (2.34) follows from the fact that $\theta_t < K$ and $\tilde{q}_t > \hat{q}_t$, which implies that $\phi_t < \theta_t < K$. Note that p_t is a priori random and is realized with the information set $f_t \in \mathcal{F}_t$. Following our convention we will use P_t to denote this a priori random probability.

This concludes the description of the RB policy. In the next section, we shall show that the above the RB policy has an expected worst-case performance guarantee of 6. Following the same argument in the uncapacitated case, this randomized decision rule *almost* balances, up to the uniform capacity constraint, the three types of costs associated with the period.

2.6.3 Worst-case analysis

Let Z_t^{RB} be a random variable defined as

$$Z_t^{RB} \triangleq \begin{cases} \Theta_t, & \text{if } \Theta_t \geq K \\ P_t K, & \text{otherwise} \end{cases} \tag{2.35}$$

Note that Z_t^{RB} is a random variable that is realized with the information set in period t. In the following lemma we show that the expected cost of the RB policy can be upper bounded using the Z_t^{RB} variables defined in (2.35).

Lemma 2.6.1 *Let $\mathscr{C}(RB)$ be the total cost incurred by the RB policy. Then we have,*

$$E[\mathscr{C}(RB)] \leq 3 \cdot \sum_{t=1}^{T-L} E[Z_t^{RB}]. \tag{2.36}$$

To complete the worst-case analysis, we would like to show that *twice* of the expected cost of an optimal policy denoted by OPT is at least $\sum_{t=1}^{T-L} E[Z_t^{RB}]$. This will be done by armotizing the cost of OPT against the cost of the RB policy. In particular, we shall show that on expectation OPT pays for a large fraction of the cost of the RB policy. In the subsequent analysis, we will use the following random partition of the periods $t = \{1, 2, \ldots T - L\}$ to the following sets:

$$\mathscr{T}_{1H} = \left\{ t : \Theta_t \geq K \text{ and } Y_t^{OPT} > Y_t^{RB} \right\}. \tag{2.37}$$

The set $\mathscr{T}_{1H}$ consists of periods in which the balancing cost Θ_t exceeds K and the optimal policy had higher inventory position than that of the RB policy after ordering (recall that if $\Theta_t \geq K$ then the RB policy orders the balancing quantity with probability 1 and the value Y_t^{RB} is realized with F_t).

$$\mathscr{T}_{1\mathrm{II}} = \left\{t : \Theta_t \geq K \text{ and } Y_t^{OPT} \leq Y_t^{RB}\right\}. \tag{2.38}$$

The set $\mathscr{T}_{1\mathrm{II}}$ consists of periods in which the balancing cost exceeds K and the inventory position of the optimal policy does not exceed that of the RB policy after ordering (see the comment above regarding $\mathscr{T}_{1H}$).

$$\mathscr{T}_{2H} = \left\{t : \Theta_t < K \text{ and } Y_t^{OPT} \geq X_t^{RB} + \min\{\tilde{Q}_t^{RB}, u\}\right\}. \tag{2.39}$$

The set $\mathscr{T}_{2H}$ consists of periods in which the balancing cost is less than K and, in such a period, the inventory position of the RB policy after ordering would be either X_t^{RB} if no order was placed, or $X_t^{RB} + \min\{\tilde{Q}_t^{RB}, u\}$ if minimum of the *holding-cost-K quantity* and the uniform capacity u is ordered, depending on the randomized decision of the RB policy. However, the inventory position of OPT after ordering exceeds even $X_t^{RB} + \min\{\tilde{Q}_t^{RB}, u\}$. (Note again that the quantity $\tilde{Q}_t^{RB}$ is known deterministically (i.e., realized) with F_t.)

$$\mathscr{T}_{2\mathrm{II}} = \left\{t : \Theta_t < K \text{ and } X_t^{RB} \geq Y_t^{OPT}\right\}. \tag{2.40}$$

Analogous to $\mathscr{T}_{2H}$, the set $\mathscr{T}_{2\mathrm{II}}$ consists of periods in which the inventory position of OPT after ordering is below X_t^{RB}.

$$\mathscr{T}_{2M} = \left\{t : \Theta_t < K \text{ and } X_t^{RB} < Y_t^{OPT} < X_t^{RB} + \min\{\tilde{Q}_t^{RB}, u\}\right\}. \tag{2.41}$$

$\mathscr{T}_{2M}$ consists of periods in which the balancing cost is less than K and the inventory position of OPT after ordering is within $(X_t^{RB}, X_t^{RB} + \min\{\tilde{Q}_t^{RB}, u\})$. Thus, whether the RB policy or OPT have more inventory depends on whether the RB policy placed an order.

Note that the sets (2.37) $-$ (2.41) are disjoint and the union makes a complete set. It is also straightforward to check that conditioning on f_t, it is already known which part of the partition period t belongs.

Lemma 2.6.2 *The overall holding cost and backlogging cost incurred by OPT are denoted by H^{OPT} and Π^{OPT} respectively. Then we have*

$$\begin{aligned} E[H^{OPT}] &\geq E\left[\sum_t Z_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{1H})\right] && (2.42)\\ E[\Pi^{OPT}] &\geq E\left[\sum_t Z_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi})\right] && (2.43)\\ E\left[\sum_t K \cdot \mathbb{1}(Q_t^{OPT} > 0)\right] &\geq E\left[\sum_t Z_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{2M})\right] && (2.44)\\ E\left[H^{OPT} + \sum_t K \cdot \mathbb{1}(Q_t^{OPT} > 0)\right] &\geq E\left[\sum_t Z_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{2H})\right]. && (2.45) \end{aligned}$$

It can be readily verified by summing up inequalities (2.42) to (2.45) that

$$2 \cdot E[\mathscr{C}(OPT)] \geq \sum_{t=1}^{T-L} E[Z_t^{RB}]. \qquad (2.46)$$

Following (2.46) and Lemma 2.6.1, we have established that the RB policy has an expected worst-case performance guarantee of 6.

Theorem 2.6.3 *For each instance of the stochastic lot-sizing problem under uniform capacity constraint, the expected cost of the randomized cost-balancing policy RB is at most six times the expected cost of an optimal policy OPT, i.e.,*

$$E[\mathscr{C}(RB)] \leq 6 \cdot E[\mathscr{C}(OPT)]. \qquad (2.47)$$

2.6.4 Numerical Experiments

The policies were tested using the demand model of *advance demand information* proposed by Gallego and Özer (2001) and Özer and Wei (2004), similar to the counterpart model without capacity constraints.

Parameterized policies. We describe a class of parameterized policies involving parameters β, γ and η, where β controls the holding-cost-βK quantity, γ controls the ratio of marginal holding costs and backlogging costs and η controls the level of expected backlogging cost resulting from not ordering. The parameterized policy first computes several quantities.

- The balancing quantity $\hat{q}_t$ that solves $E[H_t^{RB}(\hat{q}_t) \mid f_t] = \gamma \cdot E[\bar{\Pi}_t^{RB}(\hat{q}_t) \mid f_t] := \theta_t$.

- The holding-cost-βK quantity $\tilde{q}_t$ that solves $E[H_t^{RB}(\tilde{q}_t) \mid f_t] = \beta \cdot K$.

- The resulting conditional expected backlogging cost if one orders $\min\{\tilde{q}_t, u\}$ units in period t, denoted by ϕ_t. That is, $\phi_t = E[\bar{\Pi}_t^{RB}(\min\{\tilde{q}_t. u\}) \mid f_t]$.

- The conditional expected backlogging cost resulting from not ordering in period t, denoted by ψ_t. That is, $\psi_t = \eta \cdot E[\bar{\Pi}_t^{RB}(0) \mid f_t]$.

Based on the above quantities computed, the following randomized rule is used in each period t.

(I) If $\theta_t \geq \beta \cdot K$, the RB policy orders $q_t^{RB} = \hat{q}_t$ with probability $p_t = 1$ in period t.

(II) If $\theta_t < \beta \cdot K$, the RB policy orders $q_t^{RB} = \min\{\tilde{q}_t. u\}$ with probability p_t and order nothing with probability $1 - p_t$ in period t. That is,

$$q_t^{RB} = \begin{cases} \min\{\tilde{q}_t, u\}, & \text{with probability } p_t \\ 0, & \text{with probability } 1 - p_t \end{cases}, \qquad (2.48)$$

where the probability

$$0 \leq p_t = \frac{\psi_t}{\beta \cdot K - \phi_t + \psi_t} < 1. \tag{2.49}$$

End-of-horizon rule. To prevent the policy from over-ordering too much near the end of horizon, we also incorporate the following end-of-horizon rule. Suppose we are in period t, we estimate the total expected cumulative backlogging cost (assuming no orders are placed) over the interval $[t, T]$. If the amount is less than K, we do not order with probability 1 in period t.

Numerical results. We have conducted computational experiments under the cost and demand structure used in Gallego and Özer (2001) and Özer and Wei (2004). In this section, we focus on the uniform capacitated model, and the empirical results are tabulated in Table A.4. The RB policies perform around 30% of the error from the optimal cost, which significantly better than the worst-case performance guarantee. The parameters embedded in the capacitated model have the same intuitive interpretations as in the uncapacitated case.

Chapter 3

Revenue Management of Reusable Resources with Advanced Reservations

3.1 Introduction

In this chapter, we consider a class of revenue management problems that arise in systems with *reusable resources* and *advanced reservations.* This work is motivated by both traditional and emerging application domains, such as *hotel room management, car rental management* and *workforce management.* For instance, in hotel industries, customers make requests to book a room in the future for a specified number of days. This is called advanced reservation. Rooms are allocated to customers based on their requests, and after one customer used a room it becomes available to serve other customers. One of the major issues in these systems is how to manage capacitated pool of reusable resources over time in a dynamic environment with many uncertainties. In particular, one wishes to choose the most profitable customers to maximize the resulting revenue.

Models with reusable resources and advanced reservations are typically very hard to analyze, particularly due to the existence of advanced reservations. There has been relatively little related work both on finding provably good policies for these important models and structural properties of optimal or even practically good policies. In this chapter, we analyze the performance of conceptually and computationally simple policies. We show that they perform provably near-optimal in the *Halfin-Whitt* (see Halfin and Whitt (1981)) heavy-traffic regime. That is, the expected long-run revenue of the policy is guaranteed to obtain at least a constant fraction of the optimal revenue regardless of the input instance. Moreover, the analysis builds upon novel approaches to analyze the important class of *loss network models with advanced reservations.* The latter class of models is fundamental in the analysis of many applications in operations management, communication networks and other domains. There has been very little known about the structural properties of models with advanced reservations, and we believe that our work could open new opportunities to analyze additional models.

3.1.1 The model

This chapter is focused on models concerning the revenue management of a single pool of reusable resources used to serve multiple classes of customers through advanced reservations. The details of the model are as follows. There is a single pool of resources of integer capacity $C < \infty$ that is used to satisfy the demands of M different classes of customers. The customers of each class $k = 1, \ldots, M$, arrive according to an independent Poisson process with respective rate λ_k. Each class-k customer requests to reserve one unit of the capacity for a specified *service time interval* in the future.

Let D_k be the reservation distribution of a class-k customer, and S_k be the respective service distribution with mean μ_k (see Figure 3-1). In particular, upon an arrival of a class-k customer at some random time t, the customer requests to reserve the service time interval $[t + d, t + d + s]$, where d is distributed according to D_k and s is distributed according to S_k. Note that D_k and S_k are independent of the arrival process and between customers; however, per customer, D_k and S_k can be correlated. (We assume that both D_k and S_k are finite discrete distributions.) During the time a customer is served (i.e., $[t + d, t + d + s]$), the requested unit cannot be used by any other customer; after the service is over, the unit becomes available again to serve other customers. If the resource is reserved, the customer pays a class-specific rate of r_k dollars per unit of service time. The resource can be reserved for an arriving customer only if upon arrival there is at least one unit of capacity that is available (i.e., not reserved) throughout the entire requested interval $[t + d, t + d + s]$. Specifically, a customer's request can be satisfied if the maximum number of already reserved resources throughout the requested service interval is smaller than the capacity C. However, customers can be rejected even if there is available capacity. Rejecting a customer now possibly enables serving more profitable customers in the future. Customers whose requests are not reserved upon arrival are *lost* and leave the system. The goal is to find a feasible admission policy that maximizes the expected long-run average revenue. Specifically, if $\mathcal{R}_\pi(T)$ denotes the revenue achieved by policy π over the interval $[0, T]$, then the expected long-run average revenue of π is defined as $\mathcal{R}(\pi) \triangleq \liminf_{T\to\infty}(\mathbb{E}[\mathcal{R}_\pi(T)]/T)$, where the expectation is taken with respect to the probability measure induced by π.

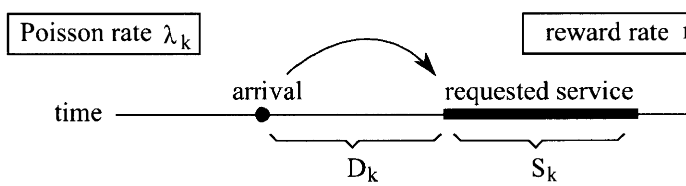

Figure 3-1: Reservation distributions and service distributions

Like many stochastic optimization models, one can formulate this problem using dynamic programming approach. However, even in special cases (e.g., no advanced reservations allowed and with exponentially distributed service times), the resulting dynamic programs seem computationally intractable because the corresponding state space grows very fast. This is known as the *curse of dimensionality.* Thus, finding provably good policies is a challenging task.

3.1.2 Our Contributions

The contributions of this chapter are two-fold. First, we employ a simple knapsack linear program (LP) to devise a conceptually simple policy that is called the *class* *selection policy* (CSP). The optimal solution of the LP guides the policy regarding which classes of customers should be admitted service and which ones should be declined service. A similar policy has been analyzed before by Levi and Radovanovic (2010) for *models without advanced reservations* that are significantly easier. In fact, the analysis in Levi and Radovanovic (2010) does not carry through to models with advanced reservations. Instead, we develop an entirely new analysis that shows the policy performs provably near-optimal under the Halfin-Whitt heavy-traffic regimes. In the Halfin-Whitt regime $C = \rho + \beta\sqrt{\rho} + o(\rho) \to \infty$, where $\beta \geq 0$ is a scaling factor. In particular, the CSP is guaranteed to obtain at least $\Phi(\beta) \geq \frac{1}{2}$ of the optimal long-run revenue in the Halfin-Whitt regime. (Note that $\Phi(\cdot)$ is the cumulative density function of a standard normal. Thus, $\Phi(\beta)$ approaches 1 when β is large.) Moreover, we propose a modified version of CSP that is guaranteed to asymptotically obtain $1 - \epsilon$ fraction of the optimal revenue, for every fixed $\epsilon > 0$.

Secondly, the analysis approaches we develop are based on modeling the problem as a loss network system with advanced reservations (specifically, a $M/G/C/C$ loss system with advanced reservations). These models are concerned with the setting in which customers arrive to the system according to a stochastic process and are being served as long as there is available capacity. Customers who find a fully utilized system are lost (see, for example, the survey paper by Kelly (1991)). We are able to derive explicit upper bounds on the steady state *blocking probability*, i.e., the probability that a random customer at steady state will find a fully utilized system, and analyze them asymptotically in the above regimes. To the best of our knowledge, there have been very few successful attempts to characterize the blocking probabilities for loss network models with advanced reservations (see, for example, Coffman-Jr et al. (1999) and Lu and Radovanovic (2007a) that studied several special cases). The assumptions in our model are fairly general: a time-homogeneous Poisson arrival process, a general finite discrete service distribution and a general finite discrete reservation distribution. Models with advanced reservations are significantly harder to analyze than those without advanced reservations. One of the major difficulties in models with advanced reservations is the fact that a randomly arriving customer effectively observes a nonhomogeneous Poisson process that is induced by the already reserved service intervals. Moreover, analyzing the blocking probability of an arriving customer requires considering the entire requested service interval instead of the instantaneous load of the system. Analyzing the load over an interval immediately introduces correlation that is challenging to analyze. The upper bound on the blocking probability is obtained by considering an identical system with infinite capacity, where all customers are admitted (a $M/G/\infty$ system with advanced reservations). The probability of having more than C customers reserved in the infinite capacity system provides an upper bound on the blocking probability in the original system; we call this the *virtual blocking probability*. Through an innovative reduction to a random walk setting, we obtain an exact analytical expression for this virtual blocking probability and then analyze it asymptotically. The analysis of the virtual blocking probability is tight and constitutes a contribution for the analysis of $M/G/\infty$ systems with advanced reservations.

The analysis approaches that are developed in this chapter significantly depart from previous work, and provide multi-fold contributions to queueing theory. We believe that these new approaches will be very effective in analyzing other important models in operations management and other application domains.

3.1.3 Literature Review

Levi and Radovanovic (2010) used a simple knapsack-type linear program (LP) to devise a conceptually simple admission control policy called class selection policy (CSP) for the model *in the absence* of advanced reservations (i.e., customers start service upon arrival). The optimal solution obtained by solving the LP guides the policy to select the more profitable classes of customers. The LP provides an upper bound on the optimal expected long-run average revenue and can be used to analyze the performance of CSP. The analysis is based on the fact that the CSP induces a stochastic process that can be reduced to a classical loss network model without advanced reservations. They were able to develop explicit expressions for the resulting blocking probabilities induced by the CSP, and then showed that the CSP is guaranteed to achieve at least half of the optimal long-run revenue. Also, the CSP was shown to be asymptotically optimal when the capacity goes to infinity with no other assumptions.

The knapsack-type LP considered by Levi and Radovanovic (2010) has been previously discussed by several other researchers (see, for example, Key (1990) and Hunt and Laws (1997)). In fact, a variant of the CSP has been discussed by Key (1990) and Kelly (1991), who analyzed the *randomized thinning policy*. Moreover, Key (1990) has shown that the variant of the CSP for the single resource case without advanced reservations is asymptotically optimal in the critically loaded regime. Iyengar and Sigman (2004) have also used an identical LP to devise a heuristic called *exponential penalty function control* for a finite-horizon variant. All of these works have considered models without advanced reservations.

Loss network models without advanced reservations are well known; they were introduced over four decades ago and have been studied extensively, primarily in the context of communication networks (see, for example, the survey paper by Kelly (1991)) and recently other application domains. Two of the major issues in the literature on loss networks have been the *study and design of heuristics* for admission control (see, for example, Miller (1969), Ross and Tsang (1989), Key (1990), Kelly (1991), Hunt and Laws (1997), Puhalskii and Reiman (1998), Fan-Orzechowski and Feinberg (2006)), and the development of *approximations and bounds* as well as *sensitivity analysis* of loss (blocking) probabilities with respect to input parameters and resource capacities (see, for example, Erlang (1917), Sevastyanov (1957), Kaufman (1981), Burman et al. (1984), Whitt (1985), Kelly (1991), Ross and Yao (1990), Zachary (1991), Louth et al. (1994), Kumar et al. (1998) and Adelman (2006)).

However, there have been relatively few successful attempts to characterize the blocking probabilities for the loss network models with advanced reservations. In particular, all the results mentioned above do not carry through. Coffman-Jr et al. (1999) derived explicit formulas for the limiting blocking probabilities in several special cases, for instance, in a setting where the reservation distribution is uniform and all requested intervals have unit length. They extended the result to more general reservation distributions by relating the problem to an on-line interval packing problem. Lu and Radovanovic (2007a) studied the asymptotic blocking probabilities when the capacity of the system approaches infinity with subexponential resource requirements. Some papers are devoted to study the transient behavior or approximations of blocking probabilities for the $M_t/G/\infty$ queue as well as $M_t/G/C/C$ loss systems (without advanced reservations) where the arrival process is nonhomogeneous Poisson (see Eick et al. (1993b), Eick et al. (1993a)), Massey (1985) for the details on some of the results along these lines). The deterministic counterpart of this problem with advanced reservations has been considered in the scheduling and parallel computing literature, which is not the main focus of this chapter.

3.2 An LP-based Approach

In this section, we describe a simple linear program (LP) that provides an upper bound on the achievable expected long-run average revenue. The LP conceptually resembles to the one used by Levi and Radovanovic (2010), Key (1990) and Iyengar and Sigman (2004) who study models without advanced reservations. It is also similar in spirit to the one used by Adelman (2007) in the queueing networks framework with unit resource requirements again without advanced reservations. We shall show how to use the optimal solution of the LP to construct a simple admission control policy that is called *class selection policy* (CSP). This type of policy was first analyzed by Levi and Radovanovic (2010) in models without advanced reservations.

At any point of time t, the state of the system is specified by the entire booking profile consisting of the class, reservation and service information of each customer in the booking system as well as the customers currently served. Without loss of generality, we restrict attention to *state-dependent policies.* Note that each state-dependent policy induces a Markov process over the state-space. Moreover, by following similar arguments as in Lu and Radovanovic (2007b) and Sevastyanov (1957), one can show that the induced Markov process has a unique stationary distribution which is Ergodic. (Though it is not the main focus of this chapter, for completeness the detailed proof of Ergodicity is provided in the electronic companion.) Since any state-dependent policy induces a Markov process on the state-space of the system that is ergodic, for a given state-dependent policy π, there exists a long-run stationary probability α^{π}_{ijk} for accepting a class-k customer who wishes to start service in i units of time for j units of time, which is equal to the long-run proportion of accepted customers of this type while running the policy π. In other words, any state-dependent policy π is associated with the stationary probabilities α^{π}_{ijk} for all possible reservation time i, service time j and class k. Let $\lambda_{ijk} \triangleq \lambda_k \mathbb{P}(D_k = i, S_k = j)$ be the arrival rate of class-k customers with reservation time i and service time j. Therefore the mean arrival rate of accepted class-k customers with reservation time i and service time j is $\alpha^{\pi}_{ijk}\lambda_{ijk}$. By applying Little's Law and PASTA property (see Gallager (1996)), the expected number of class-k customers with reservation time i and service time j being served in the system under state-dependent policy π is $\alpha^{\pi}_{ijk}\lambda_{ijk}j$. It follows that under policy π the expected long-run average number of resource units being used to serve customers can be expressed as $\sum_{k=1}^{M}\sum_{i,j}\alpha^{\pi}_{ijk}\lambda_{ijk}j$. This gives rise to the following *knapsack* LP:

$$\max_{\alpha_{ijk}} \quad \sum_{k=1}^{M}\sum_{i,j} r_k \alpha^{\pi}_{ijk}\lambda_{ijk}j, \tag{3.1}$$

$$\text{s.t.} \quad \sum_{k=1}^{M}\sum_{i,j} \alpha^{\pi}_{ijk}\lambda_{ijk}j \leq C, \tag{3.2}$$

$$0 \leq \alpha^{\pi}_{ijk} \leq 1, \qquad \forall i, j, k. \tag{3.3}$$

Note that for each feasible state-dependent policy π, the vector $\alpha^{\pi} = \{\alpha^{\pi}_{ijk}\}$ is a feasible solution for the LP with objective value equal to the expected long-run average revenue of policy π. In fact, the LP enforces the capacity constraint (3.2) of the system only in expectation, whereas in the original problem this constraint has to hold, for each sample path. It follows that the LP relaxes the original problem and provides an upper bound on the best obtained expected long-run average revenue. The LP can be solved optimally by applying the following greedy rule. Without loss of generality, assume that classes are re-numbered such that $r_1 \geq r_2 \geq \ldots \geq r_M$. Then, for each

$k = 1, \ldots, M$, we sequentially set $\alpha_{ijk} = 1$ for all i and j as long as constraint (3.2) is satisfied. If there exists a class $M' \leq M$ such that

$$C = (1 - \chi) \sum_{k=1}^{M'-1} \sum_{i,j} \lambda_{ijk} j + \chi \sum_{k=1}^{M'} \sum_{i,j} \lambda_{ijk} j$$

for some $\chi \in (0, 1)$, we set $\alpha_{ijM'} = \chi$ for all i and j. Note that for each class k, the values of α_{ijk} are all equal regardless of i and j. We abuse the notation and drop the subscripts i and j of α_{ijk}. Then the optimal solution reduces to the following structure: for $k = 1, \ldots, M' - 1$, $\alpha_k = 1$; $\alpha_{M'} = \chi$; and for $k = M' + 1, \ldots, M$, we have $\alpha_k = 0$.

Next, we shall use the optimal solution of the knapsack LP to construct a very simple admission policy. Let $\alpha^* = \{\alpha_k^*\}$ be the optimal solution of the knapsack LP. We propose a simple policy that is called *class selection policy* (CSP). Consider an arrival of a class-k customer ($k = 1, \ldots, M$). For each $k = 1, \ldots, M' - 1$, *accept* the customer upon arrival (regardless of the reservation time and the service time) as long as there is sufficient unreserved capacity throughout the requested service interval. If $k = M'$, *accept* the customer with probability χ (regardless of the reservation time and the service time) and as long as there is sufficient unreserved capacity throughout the requested service interval. For each $k = M' + 1, \ldots, M$, *reject.*

The CSP has a very simple structure. It always admits customers from the classes for which the corresponding value α_k^* in the optimal LP solution equals to one as long as capacity permits. It never admits customers from classes for which the corresponding value α_k^* equals to zero, and it flips a coin for the possibly one class with fractional value $\alpha_{M'}^* = \chi$. The CSP is conceptually very intuitive in that it splits the classes into profitable and nonprofitable that should be ignored. In fact, we can assume, without loss of generality, that there is no fractional variable in the optimal solution α^*, i.e., for each $k = 1, \ldots, M'$, $\alpha_k^* = 1$. (If $\alpha_{M'}^* = \chi$ is fractional, we think of class M' as having an arrival rate $\lambda'_{M'} = \chi \lambda_{M'}$ and then eliminate the fractional variable from α^*.)

3.3 Performance Analysis of the CSP

In this section, we discuss performance analysis of the CSP under models with advanced reservations. The CSP induces a well-structured stochastic process called loss networks with advanced reservations (i.e., a $M/G/C/C$ loss system with advanced reservations). Each class $k = 1, \dots, M$ induces a Poisson arrival stream with respective rate $\alpha_k^* \lambda_k$, $1 \leq k \leq M$. Thus, for each class k with $\alpha_k^* = 1$, the arrival process is identical to the original process, and each class k with $\alpha_k^* = 0$ can be ignored. For each class $k = 1, \dots, M'$, let S_k (discrete with finite support $[1, v_k]$) and D_k (discrete with finite support $[1, u_k]$) be the service and reservation distributions of class-k customers, respectively. We are interested in characterizing the long-run *blocking probability* of class-k customers with reservation time i and service time j under the CSP, i.e., the stationary probability that a class-k customer with reservation time i and service time j arrives at a random time to the system and is rejected by the CSP because there is no available capacity at some point within the requested service interval. For each $k = 1, \dots, M'$, let Q_{ijk} be the stationary probability of blocking a class-k customers with reservation time i and service time j under the CSP. Since the corresponding stochastic process is Ergodic, Q_{ijk} is well-defined. Thus, the expected long-run average revenue of the CSP can be expressed as $\sum_{k=1}^{M'} \sum_{i,j} r_k \lambda_{ijk} j (1 - Q_{ijk})$. However, $\sum_{k=1}^{M'} \sum_{i,j} r_k \lambda_{ijk} j$ is the optimal value of the LP, which is an upper bound on the best achievable expected long-run average revenue, denoted by $\mathcal{R}(OPT)$. Thus, a key aspect of the performance analysis of the CSP is to obtain a lower bound on the probabilities $(1 - Q_{ijk})$'s or, equivalently, an upper bound on the probabilities Q_{ijk}'s. Specifically, if $1 - Q_{ijk} \geq \xi$, for each i, j, and k, it follows that

$$\mathcal{R}(CSP) = \sum_{k=1}^{M'} \sum_{i,j} r_k \lambda_{ijk} j (1 - Q_{ijk}) \geq \sum_{k=1}^{M'} \sum_{i,j} r_k \lambda_{ijk} j \xi \geq \xi \mathcal{R}(OPT). \tag{3.4}$$

We want to obtain upper bounds on the probabilities Q_{ijk}'s and analyze their asymptotic behavior under the Halfin-Whitt regime. The *traffic intensity*

$$\rho \triangleq \sum_{k=1}^{M} \sum_{i,j} \lambda_{ijk} j = \sum_{k=1}^{M} \lambda_k \mu_k.$$

Under the Halfin-Whitt regime (see Halfin and Whitt (1981)), the capacity C and the arrival rates λ_k as well as the traffic intensity ρ are scaled together to infinity while keeping the service and the reservation distributions fixed (i.e., $C = \rho + \beta\sqrt{\rho} + o(\sqrt{\rho}) \to \infty$, for some scaling factor $\beta \geq 0$). We next formally state one of the main theorems of this chapter.

Theorem 3.3.1 *Consider the revenue management model with a single pool of capacitated reusable resources and advanced reservations under the CSP. Let $\Phi(\cdot)$ be the cumulative density function of a standard Normal distribution. Then:*

(a) For each k and j, the blocking probability Q_{ijk} has the following asymptotic upper bound,

$$\lim_{\rho\to\infty} Q_{0jk} \leq \Phi(-\beta); \qquad \lim_{\rho\to\infty} Q_{ijk} = 0, \qquad \forall i \geq 1.$$

(b) The CSP is guaranteed to obtain at least half of the optimal expected long-run average revenue in the critically loaded limit, and at least $1 - \Phi(-\beta) = \Phi(\beta) \geq \frac{1}{2}$ in the Halfin-Whitt heavy-traffic limit, where $\beta > 0$ is the scaling factor in the Halfin-Whitt regime.

For the sake of clear exposition, we merge the M'-class arrival process, and the merged arrival process has an aggregate rate $\lambda = \sum_{k=1}^{M'} \lambda_k$. A customer upon arrival has probability of λ_k/λ to be a class-k customer. Define $v = \max_k v_k$ and $u = \max_k u_k$. Let S (discrete with finite support $[1, v]$ and mean μ) and D (discrete with finite support $[1, u]$) be the 'merged' service and reservation distributions. The joint density of S and D is $f_{D,S}(i,j) \triangleq \mathbb{P}(D = i, S = j) = \sum_{k=1}^{M} \frac{\lambda_k}{\lambda} \cdot \mathbb{P}(D_k = i, S_k = j)$, for $i \in [0, u]$ and $j \in [1, v]$. Similarly, the marginal density functions of S and D are $f_S(j) \triangleq \mathbb{P}(S = j) = \sum_{k=1}^{M} \frac{\lambda_k}{\lambda} \cdot \mathbb{P}(S_k = j)$ and $f_D(i) \triangleq \mathbb{P}(D = i) = \sum_{k=1}^{M} \frac{\lambda_k}{\lambda} \cdot \mathbb{P}(D_k = i)$, respectively, for $i \in [0, u]$ and $j \in [1, v]$.

An arriving customer at some random time t requests to reserve the service interval $[t + d, t + d + s]$, where d and s are drawn according to joint density of D and S. Thus, $t + d$ is the *starting service time.* This customer will be blocked if and only if the maximum reserved capacity over the requested service interval $[t + d, t + d + s]$ just prior to time t is already C. This system captures the stochastic process induced by the CSP, and we are interested in deriving upper bounds on the induced blocking probabilities of this loss network system.

3.3.1 Main Challenges

We consider the counterpart system with infinite capacity (i.e., a $M/G/\infty$ system with advanced reservations) while keeping all other problem parameters fixed. In this counterpart system, all customers are admitted since there is an infinite number of resources. Also it is readily verified that, for each sample path and each time t, the admitted customers reserved to get service in the capacitated system are a subset of those reserved in the infinite capacity counterpart system. Consider now a customer arriving at some random time t in the counterpart system with infinite capacity requesting service interval $[t + d, t + d + s]$. Define the *virtual blocking probability* to be the probability that the maximum reserved capacity over the requested service interval $[t + d, t + d + s]$ just prior to time t is larger than C. Since the set of served customers in the infinite capacity system is always a superset of that served in the original capacitated system, it follows that the virtual blocking probability is in fact an upper bound on the blocking probability in the capacitated system. Next we shall analyze the asymptotic behavior of the virtual blocking probabilities under the Halfin-Whitt regime. In turn, this will provide asymptotic upper bounds on the blocking probabilities in the original capacitated system.

The major challenges in analyzing the blocking probabilities in loss network systems with advanced reservations lie in the fact that we need a complete characterization of the booking profile (the pre-reserved arrival and departure processes) to obtain the maximum reserved capacity over a particular requested service interval. As shown in Figure 3-2 (the capacity $C = 2$), in the models without advanced reser-

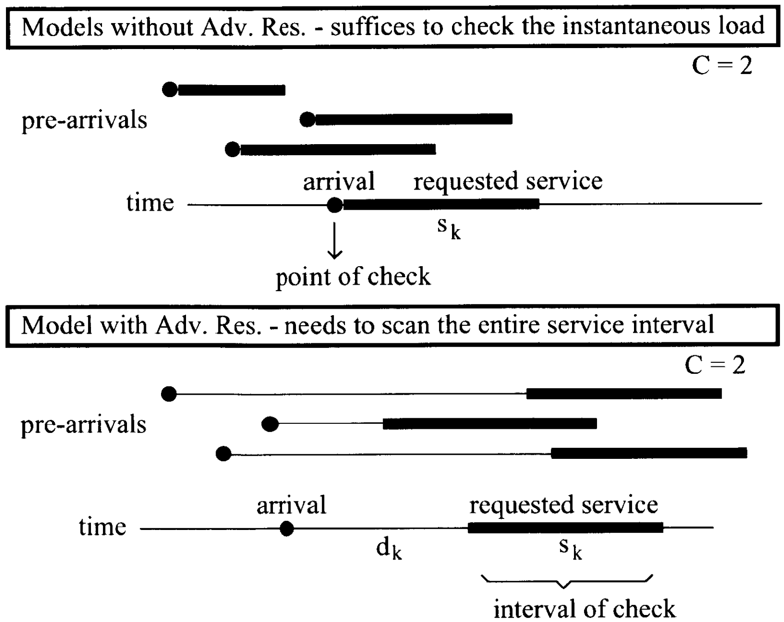

Figure 3-2: Challenges in analyzing the blocking probabilities in loss network systems with advanced reservations

vation, it suffices to check the instantaneous load of the system upon arrival of a customer. However, in the models with advanced reservation, we cannot guarantee one's request by merely checking the instantaneous load of the system at her starting service time upon her arrival, because her request may be potentially blocked by reserved slots of those customers who booked prior to her but will start services after her. This introduces much difficulties in handling this correlation issue between the incoming requests and the booking profiles.

One may suggest regarding the original system with advanced reservation as a tandem queueing model of two stations, where the first station has infinite capacity and the second station has finite capacity, customers first enter the system from the first station, but if the second station is full when customers arrive, they will be rejected or lost. When we relax the finite capacity assumption on the second station, the blocking probability can seemingly be approximated by the probability that the number of customers in the second station is bigger than C (see Boxma (1984) and Schmidt (1987)). However, the dynamics of tandem queues of two stations is very different from the loss network systems with advanced reservation. From the example in Figure 3-2, the incoming request will be accepted in the infinite tandem queueing model of two stations while the request will be virtually blocked in our model. Moreover, the tandem queueing model of two stations checks the instantaneous load of the second server upon arrival of a customer and therefore cannot be used to upper bound the blocking probabilities in our system. It may serve as an approximation of the blocking probabilities but we are unsure how good the approximation is, since the stationary distribution can no longer be expressed as a product-form.

3.3.2 The Simplest Non-trivial Case

We will start the asymptotic analysis with the simplest non-trivial case, and then extend it gradually to the more general case. Suppose that S takes only one value $s = 1$ deterministically. Then the traffic intensity $\rho = \lambda\mu = \lambda$. In addition, assume that D follows a two-point distribution,

$$D = \begin{cases} 0 & \text{w.p. } \gamma, \\ 1 & \text{w.p. } 1-\gamma, \end{cases}$$

i.e., $f_D(0) = \gamma$ and $f_D(1) = 1 - \gamma$. That is, an arriving customer either wants to start the service immediately or in 1 unit of time. Consider the counterpart system with an infinite number of servers in steady state (note that the steady state exists due to the induced semi-Markov process). Upon a customer arrival to the system at some time t, all the starting service times of the customers who had arrived prior to t are already known in the booking profile. For ease of exposition, we call these starting service times *pre-arrivals*. Similarly, we call all the starting service times of the customers, who will arrive after t *post-arrivals*. (Note that the pre-arrivals and post-arrivals are always defined with respect to the current time.) It is important to note that the virtual blocking probability at time t (as well as the blocking probability in the original capacitated system) is independent of *post-arrivals*. Without loss of generality, we can assume $t = 0$.

Lemma 3.3.2 below characterizes the pre-arrival processes observed by a customer arriving at time 0 in steady state.

Lemma 3.3.2 *Consider the counterpart system with an infinite number of servers, then a customer arriving at the system at time* 0 *in steady state, observes that the pre-arrivals follow a non-homogeneous Poisson process with piecewise rate* $\eta(r)$ *at time* r

$$\eta(r) = \begin{cases} \lambda, & \textit{if} \quad r \le 0, \\ (1-\gamma)\lambda, & \textit{if} \quad 0 < r \le 1, \\ 0, & \textit{if} \quad r > 1. \end{cases}$$

Proof of Lemma 3.3.2. If $r \le 0$, we focus on the interval $(\lceil r \rceil - 1, \lceil r \rceil]$ and its preceding interval $(\lceil r \rceil - 2, \lceil r \rceil - 1]$. The arrival process in $(\lceil r \rceil - 2, \lceil r \rceil - 1]$ follows a Poisson process with rate λ. Each arrival has γ probability of starting services immediately in $(\lceil r \rceil - 2, \lceil r \rceil - 1]$, and $1-\gamma$ probability of starting services in 1 unit of time in $(\lceil r \rceil - 1, \lceil r \rceil]$. By the splitting argument, the pre-arrivals in $(\lceil r \rceil - 1, \lceil r \rceil]$ follow a Poisson process with rate $(1-\gamma)\lambda$. By a similar argument, the pre-arrival process in $(\lceil r \rceil - 1, \lceil r \rceil]$ induced by customers arriving to the system in $(\lceil r \rceil - 1, \lceil r \rceil]$ follows a Poisson process with rate $\gamma\lambda$. Note that these two processes are independent of each other since they are generated by customers arriving in disjoint intervals. Now merge these two pre-arrival processes, and the resulting pre-arrival process in $(\lceil r \rceil - 1, \lceil r \rceil]$ follows a Poisson process with rate $(1-\gamma)\lambda + \gamma\lambda = \lambda$.

If $0 < r \le 1$, focus on the interval $(0, 1]$ and its preceding interval $(-1, 0]$. By the similar argument above, there is a Poisson process of pre-arrivals with rate $(1-\gamma)\lambda$ induced by customers arriving in $(-1, 0]$. There is also a Poisson process with rate $\gamma\lambda$ induced by customers arriving in $(0, 1]$. However, the latter process consists of post-arrivals. Thus, the resulting pre-arrivals at time 0 over $(0, 1]$ follow a Poisson process with rate $(1-\gamma)\lambda$.

Since the maximum reservation time is 1, it is impossible for customers arriving prior to 0 to start service at any time greater than 1. Thus, the rate of pre-arrivals from 1 onwards is 0. $\square$

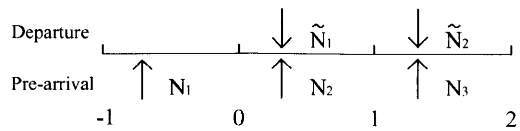

Figure 3-3: One-class departure and pre-arrival processes

Let B be the event that a customer arriving at time 0 in steady state is virtually blocked. The conditional long-run virtual blocking probability $P_i \triangleq \mathbb{P}(B \mid D = i)$, for each $i = 0, 1$. In Lemma 3.3.3 below, we show how to obtain exact analytical expressions to P_0 and P_1. Moreover, we analyze the asymptotic behavior of these expressions under the Halfin-Whitt regime.

Let N_i $(i = 1, 2, 3)$ denote the Poisson counting process (see Gallager (1996)) induced by the pre-arrivals over $[i-2, i-1]$ as seen from time 0. Also, let $N_i(r)$, for $r \in [0, 1]$, be the number of events over $[0, 1]$. Next, we introduce the notion of *mirror image* of a Poisson counting process. The mirror image of a Poisson counting process N, denoted by $\tilde{N}$, is a backward counting process of N. Let $N(r)$ and $\tilde{N}(r)$ be the number of counted events over $[0, r]$ of the respective processes ($r \leq 1$). More formally, if N is a Poisson counting process from time 0 to 1, then $\tilde{N}(r) = N(1) - N(1-r)$ for each $r \in [0, 1]$. It is evident that $\tilde{N}$ is also a Poisson process with the same rate as N.

Lemma 3.3.3 below characterizes P_0 and P_1 based on the counting processes introduced above. More generally, we will use $N(\cdot; \lambda)$ to denote a Poisson counting process with rate λ.

Lemma 3.3.3 *Consider the counterpart system with an infinite number of servers, if a customer arrives at time 0 in steady state and requests service $S = 1$ deterministically to commence in D units of time ($D = 0$ or 1 with probabilities γ and $1 - \gamma$, respectively), the conditional virtual blocking probabilities are given by*

$$P_0 \triangleq \mathbb{P}(B \mid D = 0) \triangleq \mathbb{P}\left(\max_{r \in [0,1]} \left\{\tilde{N}_1(1-r) + N_2(r)\right\} \geq C\right), \tag{3.5}$$

$$P_1 \triangleq \mathbb{P}(B \mid D = 1) \triangleq \mathbb{P}\left(\max_{r \in [0,1]} \left\{\tilde{N}_2(1-r) + N_3(r)\right\} \geq C\right), \tag{3.6}$$

where the process N_i *(*$i = 1, 2, 3$*) is a Poisson counting process with respective rate* λ_i*, with* $\lambda_1 = \lambda$*,* $\lambda_2 = (1-\gamma)\lambda$ *and* $\lambda_3 = 0$*. The process* $\tilde{N}_i$ *(*$i = 1, 2$*) is the mirror image of* N_i *with rate* λ_i*.*

Proof of Lemma 3.3.3. Suppose that a customer arrives at time 0 in steady state and requests the service to commence immediately ($D = 0$), i.e., requesting the service interval $(0, 1]$. Focus solely on the pre-arrivals as seen from 0. By Lemma 3.3.2, the pre-arrivals over the time interval $(-1, 0]$ follow a Poisson process with rate λ, denoted by N_1. However, this implies that, over the time interval $(0, 1]$, the customers depart the system following a Poisson process with rate λ (a shift of N_1 by 1 unit of time). Let $\tilde{N}_1$ be the mirror image of the departure process induced by N_1 over $(0, 1]$. (See Figure 3-3.) By Lemma 3.3.2, we also know that the pre-arrivals over $(0, 1]$ (namely, customers starting service within the interval) follow a Poisson process with rate $(1-\gamma)\lambda$. We denote this pre-arrival process by N_2.

Consider now the number of customers in the system at some time r. These fall exactly into one of the two types; customers that started service over $(0, r]$ and customers that started service over $(r-1, 0]$ and will depart over $(r, 1]$. It follows that the number of customers in service at time $r \in (0, 1]$ can be expressed as $\tilde{N}_1(1-r) + N_2(r)$. Specifically, in time r the number of departures over $(r, 1]$ (equal to $\tilde{N}_1(1-r)$) captures customers starting service before 0, and still in the system at time r. In addition, the number of pre-arrivals over $(0, r]$ (equal to $N_2(r)$) captures customers arriving before 0, starting service over $(0, r]$ and still being served in time r. The sum of the two is exactly equal to the total number of customers in the system at time r. Note that by Poisson splitting arguments it follows that $\tilde{N}_1$ and N_2 are independent of each other. The virtual blocking probability is expressed in terms of the maximum of the sum of these two Poisson counting processes running towards each other (see Figure 3-4), i.e., $P_0 \triangleq \mathbb{P}\left(\max_{r\in[0,1]} \left\{\tilde{N}_1(1-r) + N_2(r)\right\} \geq C\right)$.

Consider now the case that the arriving customer requests the service to commence in $D = 1$ unit of time, i.e., the service will cover the interval $(1, 2]$. The departure process in $(1, 2]$ is a shift of the pre-arrival process N_2 in $(0, 1]$ by 1 unit of time, and its mirror image is denoted by $\tilde{N}_2$. Moreover, by Lemma 3.3.2, the pre-arrival process

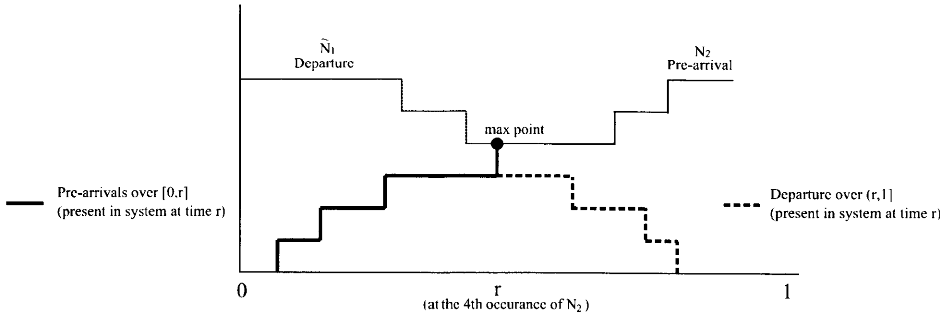

Figure 3-4: Two Poisson counting processes running towards each other

in $(1, 2]$ has rate 0. Thus,

$$P_1 \triangleq \mathbb{P}\left(\max_{r\in[0,1]}\left\{\tilde{N}_2(1-r)+N_3(r)\right\}\geq C\right)=\mathbb{P}\left(\max_{r\in[0,1]}\left\{\tilde{N}_2(1-r)\right\}\geq C\right).$$

The second equality follows from Lemma 3.3.2 above. $\square$

Observe that P_0 and P_1 are expressed through rather complex random variables. However, in the next lemmas, we show how to analyze the limits P_0 and P_1 under the Halfin-Whitt regime. We first assume that the probability that an arriving customer seeks to start service immediately is positive (i.e., $\gamma > 0$), and then relax this assumption.

Assuming that $\gamma > 0$, we shall show that under the Halfin-Whitt regime where $C = \lambda + \beta\sqrt{\lambda} + o(\sqrt{\lambda}) \to \infty$ and $\beta \geq 0$, the conditional virtual blocking probabilities P_0 and P_1 have the following asymptotic limits $\lim_{\lambda\to\infty} P_0 = \Phi(-\beta)$ and $\lim_{\lambda\to\infty} P_1 = 0$. In fact, we shall prove a more general statement that will be useful in the analysis of the general case.

Theorem 3.3.4 *Let N_1, N_2 and N_3 be Poisson counting processes (mutually independent) with rates λ, $\theta_1\lambda$ and $\theta_2\lambda$, respectively, where $1 > \theta_1 \geq \theta_2 \geq 0$ are fixed constants. Let $\tilde{N}_1$ and $\tilde{N}_2$ be the mirror images of N_1 and N_2, respectively. Let*

$$\begin{aligned} X &\triangleq \max_{r\in[0,1]}\left\{\tilde{N}_1(1-r;\lambda)+N_2(r;\theta_1\lambda)\right\}, \\ Y &\triangleq \max_{r\in[0,1]}\left\{\tilde{N}_2(1-r;\theta_1\lambda)+N_3(r;\theta_2\lambda)\right\}. \end{aligned}$$

For each $\beta > 0$, we have $\lim_{\lambda \to \infty} \mathbb{P}(X \geq C) = \Phi(-\beta)$ *and* $\lim_{\lambda \to \infty} \mathbb{P}(Y \geq C) = 0$.

Note that P_0 and P_1 can be obtained by setting $\theta_1 = (1 - \gamma)$ and $\theta_2 = 0$ in X and Y above. We observe that only customers with zero reservation time are blocked. This stems from the fact that for a given time slot and unit capacity, these customers are the last to arrive. Since we have a system where accepted demand is close to supply, these are the customers that will most likely be blocked.

To prove Theorem 3.3.4, we provide an alternative characterization of X and Y above based on a downward-drifting asymmetric random walk that takes a down-step, for each departure, and an up-step, for each pre-arrival. We would like to show that in the asymptotic heavy-traffic regimes, the maximum level of the random walk stays relatively close to its starting position by showing that the rate of the random walk going up is sublinear in $\sqrt{\lambda}$.

Consider the merged process induced on $[0, 1]$ by the two Poisson counting processes $\tilde{N}_1$ and N_2. Let $\mathcal{N} = \tilde{N}_1(1; \lambda) + N_2(1; \theta\lambda)$ denote the total number of occurrences over $[0, 1]$ of the two independent Poisson counting processes of $\tilde{N}_1$ and N_2. Note that since $\tilde{N}_1$ and N_2 are independent of each other, $\mathcal{N}$ is a Poisson random variable with rate $(1 + \theta)\lambda$. Conditioning on $\mathcal{N} = n$, the induced merged process has n points uniformly distributed over the interval $[0.1]$. By splitting argument applied to the merged process, each of these n points has independent probability $p = \frac{\theta\lambda}{(1+\theta)\lambda} = \frac{\theta}{1+\theta} < \frac{1}{2}$ to be from the process N_2 and probability $q = 1 - p$ from the process $\tilde{N}_1$. If we associate $+1$ with each point from N_2 , and -1 with each point from $\tilde{N}_1$, then each configuration of these n points induces a downward-drifting asymmetric random walk of length n. The random walk starts at the origin 0, with up probability p and down probability q. Let $\mathcal{R}_n$ denote the corresponding random walk of length n, M_n denote the maximum level attained by $\mathcal{R}_n$, and G_n denote the overall number of down-steps taken by $\mathcal{R}_n$. Also let $X_n \triangleq (X \mid \mathcal{N} = n)$. Then, we claim that, almost surely, $X_n = G_n + M_n$. Note again that for each $r \in [0, 1]$, $\tilde{N}_1(1 - r; \theta_1\lambda) + N_2(r; \lambda)$ is equal to the number of occurrences of N_1 over $(1 - r, 1]$ and the number of occurrences of N_2 over $[0, r)$. Also observe that the value of X is obtained either at time 0 or upon on occurrence of N_2. Now condition on $\mathcal{R}_n = \omega_n$

(a specific realization of the random walk $\mathcal{R}_n$), and consider the l^{th} occurrence of N_2 ($l \in \{0, \ldots, n\}$), at time, say r. Then we have (see Figure 3-4),

$$\begin{aligned} & \tilde{N}_1(1-r;\lambda) + N_2(r;\theta_1\lambda) \\ = \; & (\text{\# up-steps before and including } l + \text{\# down-steps after } l) \\ = \; & (\text{\# up-steps before and including } l - \text{\# down-steps before and including } l) \\ & +(\text{\# down-steps before and including } l + \text{\# down-steps after } l). \end{aligned}$$

The first term is exactly the location of the random walk after l steps and the second expression is exactly G_n. Since X is the maximum of the above sum over all arrivals $l = 0, 1, \ldots, n$, it follows that indeed $X_n \mid (\mathcal{R}_n = \omega_n) = (G_n + M_n) \mid (\mathcal{R}_n = \omega_n)$, from which the claim follows. However, it should be noted that M_n and G_n are correlated.

To address the correlation between M_n and G_n, we will replace M_n by M_∞. However, first we would like to obtain an expression for the hitting probability of a downward-drifting asymmetric random walk. This is done in Lemma 3.3.5 given below. (Lawler (2006) provided a proof in Chapter 2, Section 2.2; for completeness, we present a shorter proof in the electronic companion.)

Lemma 3.3.5 *Consider a random walk defined by a sequence of independent random variables* $E_i = 1$ *with probability* p *and* -1 *with probability* $q = 1 - p$. *Let* $S_n = \sum_{i=1}^{n} E_i$. *Define* $M_\infty \in [0, \infty) \bigcup \{\infty\}$ *to be maximum level attained by the random walk (i.e.,* $M_\infty = \max_n S_n$*). Given that* $0 \leq p < q \leq 1$ *(downward drifting), then the probability that the random walk ever hits above level* b *is* $\mathbb{P}(M_\infty \geq b) = (p/q)^b$.

We are now ready to prove Theorem 3.3.4.

Proof of Theorem 3.3.4. First we shall prove that, for $\beta \geq 0$, $\lim_{\lambda \to \infty} \mathbb{P}(X \geq C) = \Phi(-\beta)$. Let M_∞ be the maximum level attained by the infinite-step random walk defined above. Since the random walk has a negative drift, it follows from Lemma 3.3.5 above that $\mathbb{P}(M_\infty \geq -\log\lambda/\log\theta) \leq 1/\lambda$. (Note that $\theta < 1$, so $-\log\lambda/\log\theta > 0$.)

Now, we have

$$
\begin{aligned}
& \mathbb{P}(X_n \geq C) & (3.7)\\
= \; & \mathbb{P}\left(G_n + M_n \geq C\right) \\
= \; & \mathbb{P}\left(G_n + M_n \geq C \bigcap M_n \geq -\frac{\log \lambda}{\log \theta}\right) + \mathbb{P}\left(G_n + M_n \geq C \bigcap M_n < -\frac{\log \lambda}{\log \theta}\right) \\
\leq \; & \mathbb{P}\left(M_n \geq -\frac{\log \lambda}{\log \theta}\right) + \mathbb{P}\left(G_n \geq C + \frac{\log \lambda}{\log \theta}\right) \\
\leq \; & \mathbb{P}\left(M_\infty \geq -\frac{\log \lambda}{\log \theta}\right) + \mathbb{P}\left(G_n \geq C + \frac{\log \lambda}{\log \theta}\right) \\
\leq \; & \frac{1}{\lambda} + \mathbb{P}\left(G_n \geq C + \frac{\log \lambda}{\log \theta}\right).
\end{aligned}
$$

The third inequality follows from the fact that $M_\infty \geq M_n$ almost surely. The fourth inequality follows from Lemma 3.3.5 above. Since G_n is distributed as $(\tilde{N}_1(1;\lambda) \mid \mathcal{N} = n)$, we get from (3.7) that,

$$
\begin{aligned}
\mathbb{P}(X \geq C) &= \sum_{n=1}^{\infty} \mathbb{P}(X_n \geq C)\mathbb{P}(\mathcal{N} = n) \leq \frac{1}{\lambda} + \sum_{n=1}^{\infty} \mathbb{P}\left(G_n \geq C + \frac{\log \lambda}{\log \theta}\right) \mathbb{P}(\mathcal{N} = n) \\
&= \frac{1}{\lambda} + \mathbb{P}\left(\tilde{N}_1(1;\lambda) \geq C + \frac{\log \lambda}{\log \theta}\right) = \frac{1}{\lambda} + \mathbb{P}\left(\text{Poisson}(\lambda) \geq C + \frac{\log \lambda}{\log \theta}\right).
\end{aligned}
$$

By virtue of Central Limit Theorem, we have

$$
\begin{aligned}
\lim_{\lambda \to \infty} \mathbb{P}(X \geq C) &\leq \lim_{\lambda \to \infty} \mathbb{P}\left(\text{Poisson}(\lambda) \geq C + \frac{\log \lambda}{\log \theta}\right) & (3.8)\\
&= \lim_{\lambda \to \infty} \mathbb{P}\left(\text{Poisson}(\lambda) \geq \lambda + \beta\sqrt{\lambda} + o(\sqrt{\lambda}) + \frac{\log \lambda}{\log \theta}\right) & (3.9)\\
&= \lim_{\lambda \to \infty} \mathbb{P}\left(\text{Poisson}(\lambda) \geq \lambda + \beta\sqrt{\lambda} + o(\sqrt{\lambda})\right) \leq \Phi(-\beta). & (3.10)
\end{aligned}
$$

On the other hand, from the definition of X, under the Halfin-Whitt regime, we have $\lim_{\lambda \to \infty} \mathbb{P}(X \geq C) \geq \lim_{\lambda \to \infty} \mathbb{P}(N_1(1;\lambda) \geq C) = \lim_{\lambda \to \infty} \mathbb{P}(\text{Poisson}(\lambda) \geq \lambda + \beta\sqrt{\lambda} + o(\sqrt{\lambda})) = \Phi(-\beta)$.

Now, we are ready to prove the second part of Theorem 3.3.4, i.e., $\lim_{\lambda \to \infty} \mathbb{P}(Y \geq C) = 0$. Since $\theta_1 \in [0, 1)$, we can always find a $\bar{\theta}_1$ such that $\theta_1 < \bar{\theta}_1 < 1$. Then $\bar{\theta}_1 >$

$\theta_1 \geq \theta_2$, and define $\bar{Y} = \max_{r\in[0,1]} \left\{ \bar{N}_2(1-r;\bar{\theta}_1\lambda) + N_3(r;\theta_2\lambda) \right\}$. It is easy to see that $\bar{Y}$ stochastically dominates Y. Therefore, without loss of generality, we simply drop the bar of $\bar{\theta}_1$ and $\bar{Y}$, and assume that $\theta_2/\theta_1 = \theta < 1$. Following the same argument as in the first part of the proof, we have $\mathbb{P}(Y \geq C) \leq \frac{1}{\lambda} + \mathbb{P}\left(\text{Poisson}(\theta\lambda) \geq \lambda + \frac{\log\lambda}{\log\theta}\right)$, and again by virtue of Central Limit Theorem, we have

$$
\begin{aligned}
\lim_{\lambda\to\infty} \mathbb{P}(Y \geq C) &\leq \lim_{\lambda\to\infty} \mathbb{P}\left(\text{Poisson}(\theta\lambda) \geq C + \frac{\log\lambda}{\log\theta}\right) &(3.11)\\
&= \lim_{\lambda\to\infty} \mathbb{P}\left(\text{Poisson}(\theta\lambda) \geq \lambda + \beta\sqrt{\lambda} + o(\sqrt{\lambda}) + \frac{\log\lambda}{\log\theta}\right) &(3.12)\\
&= \lim_{\lambda\to\infty} \mathbb{P}\left(\text{Poisson}(\theta\lambda) \geq \lambda + \beta\sqrt{\lambda} + o(\sqrt{\lambda})\right) = 0. &(3.13)
\end{aligned}
$$

This completes the proof. □

The observation that $\lim_{\lambda\to\infty} P_1 = 0 \leq \lim_{\lambda\to\infty} P_0$ suggests us that a customer, who requests to start service immediately than in the future in these asymptotic regimes, is more likely to be blocked. We have shown that if $0 < \gamma \leq 1$, $\lim_{\lambda\to\infty} P_0 \leq \Phi(-\beta)$, and $\lim_{\lambda\to\infty} P_1 = 0$. Next consider the case where $\gamma = 0$ implying that no arriving customer will start service immediately. Thus, we have $\lim_{\lambda\to\infty} P_1 = \Phi(-\beta)$.

3.3.3 Arbitrary Finite Discrete Reservation Distributions

Next we extend the simple model to allow an arbitrary finite discrete reservation distribution D with marginal probability mass function $f_D(i)$. We still assume that the service distribution remains fixed at $S = 1$, deterministically. Now let $f_D(i) = \gamma_i$ for $i \in [0, u]$, $0 \leq \gamma_i \leq 1$ and $\sum_{i=1}^{\infty} \gamma_i = 1$. Lemma 3.3.6 below is a generalization of Lemma 3.3.2. (The proof is given in the electronic companion.)

Lemma 3.3.6 *Consider the counterpart system with an infinite number of servers, a customer arriving at the system at time 0 in steady state, observes that the pre-arrivals follow a non-homogeneous Poisson input process with piecewise rate $\eta(r)$ at* *time* r

$$\eta(r) = \begin{cases} \lambda, & \text{if} \quad r \leq 0, \\ \lambda(1 - F_D(\lceil r \rceil - 1)), & \text{if} \quad r > 0, \end{cases}$$

where F_D *is the cumulative probability mass function of* D *and* $F_D(\lceil r \rceil - 1)) = \sum_{i=0}^{\lceil r \rceil - 1} f_D(i) = \sum_{i=0}^{\lceil r \rceil - 1} \gamma_i$.

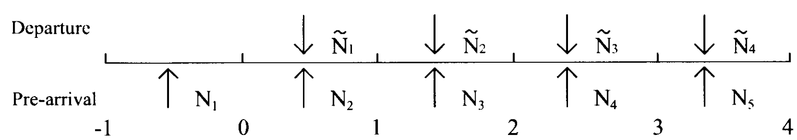

Figure 3-5: One-class departure and pre-arrival processes with general reservation distribution

Define N_i (for $i \in [1, u]$) to be the process of pre-arrivals prior to t over $(i-2, i-1]$. This process induces a departure process over the interval $(i-1, i]$, and let $\tilde{N}_i$ denote its mirror image. (An example is shown in Figure 3-5.) The conditional virtual blocking probabilities are given in Lemma 3.3.7 below, which is a generalization of Lemma 3.3.3. (The proof is given in the electronic companion.)

Lemma 3.3.7 *Consider the counterpart system with an infinite number of servers, if a customer comes at time* 0 *in steady state and requests service* ($S = 1$) *deterministically to commence in* D *units of time* ($D \in [0, u]$)*, the conditional virtual blocking probability is given by, for all* $i \in [0, u]$,

$$P_i \triangleq \mathbb{P}(B \mid D = i) \triangleq \mathbb{P}\left(\max_{r \in [0,1]} \left\{ \tilde{N}_{i+1}(1 - r; \lambda_{i+2}) + N_{i+2}(r; \lambda_{i+3}) \right\} \geq C \right),$$

where N_i *is a Poisson counting process with rate* $\lambda_i = \lambda(1 - F_D(i-2))$*, and* $\tilde{N}_i$ *is a mirror image of* N_i *with the same rate.*

Theorem 3.3.8 is a generalization of Theorem 3.3.4 with general reservation distribution. (The proof is given in the electronic companion.) The traffic intensity $\rho = \lambda$ since the service distribution $S = 1$ deterministically.

Theorem 3.3.8 *The conditional long-run virtual blocking probabilities have the following asymptotic upper bounds: for each $i \in [0, u]$ (the service distribution $S = 1$ deterministically), $P_i \leq \Phi(-\beta)$.*

3.3.4 Arbitrary Finite Discrete Service Distributions

Next we extend the model further to allow an arbitrary finite discrete service distribution. The total arrival rate is λ, and the reservation distribution D is defined on $[0, u]$ defined as in Section 3.2. Now assume that the service time S is a general finite discrete distribution on $[1, v]$. More specifically, let $f_S(\cdot)$ be the marginal service distribution with $f_S(j) = \mathbb{P}(S = j) = \kappa_j$, where $\sum_{j=1}^{v} \kappa_j = 1$ and $0 \leq \kappa_j \leq 1$, for each $j \in [1, v]$.

We partition the arriving customers according to their requested service time, i.e., the customers are partitioned into v disjoint sets numbered $1, \ldots, v$ according to their requested service time. For each $j \in [1, v]$, the arrival process of customers in set j follows a thinned Poisson process with rate $\kappa_j \lambda$. Moreover, these processes are independent of each other. Now, for each set $j \in [1, v]$, let the conditional reservation distribution be $f_D^j(i) = \mathbb{P}(D = i \mid S = j) = \gamma_i^j$ for $i \in [0, u]$. Note that $\sum_{i=0}^{u} \gamma_i^j = 1$, for each $j \in [1, v]$.

Consider the counterpart system with an infinite number of servers, if a customer of set j ($j \in [1, v]$) arrives at time 0 in steady state and requests j units of service time to commence after i units of time ($i \in [0, u]$), the conditional virtual blocking probability is defined as $P_i^j \triangleq \mathbb{P}(B \mid D = i, S = j)$. In addition, the traffic intensity is $\rho = \sum_{j=1}^{v} j\kappa_j \lambda = \mu\lambda$, where $\mu = \sum_{j=1}^{v} j\kappa_j$ is the mean service time.

Let N_i^j (for $j \in [1, v]$ and $i \in [1, u]$) denote the pre-arrival process of set-j customers over (i.e., customers requesting j units of service time) the interval $(i - j - 1, i - j]$. This induces a departure process over the interval $(i - 1, i]$, and let $\tilde{N}_i^j$ denote its mirror image. The rate of N_i^j (and $\tilde{N}_i^j$) is given in Lemma 3.3.9 below. (The proof is given in the electronic companion.)

Lemma 3.3.9 *Let N_i^j and $\tilde{N}_i^j$ be defined as above. Then, for each $j \in [1, v]$ and each*

$i \in [1, u]$, N_i^j and $\tilde{N}_i^j$ are Poisson processes with the same rate

$$\lambda_i^j = \kappa_j \lambda \left(1 - \sum_{l=0}^{i-j-1} \gamma_l^j \right) = \lambda_1^j (1 - F_D^j(i - j - 1)). \tag{3.14}$$

Moreover, N_i^j is independent of $N_{i'}^{j'}$ for $i \neq i'$ or $j \neq j'$.

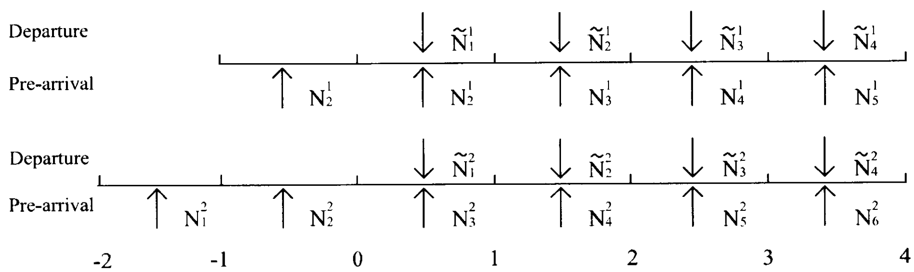

Figure 3-6: Two-service-set departure and pre-arrival processes

First assume that $\exists j \in [1, v]$ such that $\gamma_0^j > 0$, i.e., the probability of an arriving customer requesting to start the service immediately upon arrival is strictly positive. Later we will show that this assumption can be relaxed. Let A_i be the maximum number of customers in the system over the interval $(i, i+1]$ for $i \in [0, u]$. In fact, one can derive an exact mathematical expression of each A_i for $i \in [0, u]$,

$$A_i = \sum_{j=2}^{v} \sum_{l=i+2}^{i+j} N_l^j(1; \lambda_l^j) + \max_{r \in [0,1]} \left\{ \sum_{j=1}^{v} \tilde{N}_{i+1}^j (1 - r; \lambda_{i+1}^j) + \sum_{j=1}^{v} N_{i+j+1}^j (r; \lambda_{i+j+1}^j) \right\}. \tag{3.15}$$

For $r \in [0, 1]$, the term $\sum_{j=1}^{v} \tilde{N}_{i+1}^j (1 - r; \lambda_{i+1}^j)$ captures all the departures over $(i + r, i+1]$, the term $\sum_{j=1}^{v} N_{i+j+1}^j (r; \lambda_{i+j+1}^j)$ captures all the pre-arrivals over $(i, i+r]$, and the term $\sum_{j=2}^{v} \sum_{l=i+2}^{i+j} N_l^j(1; \lambda_l^j)$ captures all the customers being served over $(i, i+1]$. The sum captures exactly all the customers being served at time $i+r$. It is important to note that since $\tilde{N}_l^j$ and N_l^j do not appear together in A_i, for each $l \in [1, u]$ and $j \in [1, v]$, all the Poisson counting processes in the expression of A_i are independent of each other (see Lemma 3.3.9).

We shall further explain (3.15) by providing the following example when $v = 2$

(refer to Figure 3-6),

$$\begin{aligned}
A_0 &= N_2^2(1;\lambda_2^2) + \max_{r\in[0,1]} \left\{ \tilde{N}_1^1(1-r;\lambda_1^1) + \tilde{N}_1^2(1-r;\lambda_1^2) + N_2^1(r;\lambda_2^1) + N_3^2(r;\lambda_3^2) \right\}, \\
A_1 &= N_3^2(1;\lambda_3^2) + \max_{r\in[0,1]} \left\{ \tilde{N}_2^1(1-r;\lambda_2^1) + \tilde{N}_2^2(1-r;\lambda_2^2) + N_3^1(r;\lambda_3^1) + N_4^2(r;\lambda_4^2) \right\}, \\
&\quad \ldots\ldots
\end{aligned}$$

More specifically, A_0 represents the maximum customers in the system over the interval $(0,1]$ (refer to Figure 3-6). At time $r \in (0,1]$, the number of departures over $(r,1]$ is equal to $\tilde{N}_1^1(1-r;\lambda_1^1) + \tilde{N}_1^2(1-r;\lambda_1^2)$, capturing customers in both sets starting before 0 and still in the system at time r. (Note that the service time is at least 1.) In addition, the number of pre-arrivals over $(0,r]$ is equal to $N_2^1(r;\lambda_2^1) + N_3^2(r;\lambda_3^2)$, capturing pre-arrivals of customers with service time 1 and 2, respectively, starting service over $(0,r]$. Finally, N_2^2 captures set-2 customers with service time 2 who started service within $(-1,0]$. These customers will continue service over the entire interval $(0,1]$. Therefore $N_2^2(1;\lambda_2^2)$ appears in the expression A_0 outside the max. The same reasoning applies to A_i for each $i \in [1,u]$.

Now for each $i \in [0,u]$ and $j \in [1,v]$, we have $P_i^j = \mathbb{P}(\max(A_i,\ldots,A_{i+j-1}) \geq C)$. It should be noted that A_i and $A_{i'}$ can be correlated. To analyze the limiting behavior of P_i^j, we first analyze the limits of $\mathbb{P}(A_i \geq C)$, for each $i \in [0,u]$.

Lemma 3.3.10 *Assume that there exists $j \in [1,v]$ such that $\gamma_0^j > 0$. Let A_i be defined as in (3.15). The traffic intensity is $\rho = \sum_{j=1}^v j\kappa_j\lambda = \sum_{j=1}^v j\lambda_1^j$. Under the Halfin-Whitt regime,*

$$\lim_{\lambda\to\infty} \mathbb{P}(A_0 \geq C) = \Phi(-\beta); \qquad \lim_{\lambda\to\infty} \mathbb{P}(A_i \geq C) = 0, \qquad i \in [1,u]. \tag{3.16}$$

Proof of Lemma 3.3.10. The assumption $\gamma_0^j > 0$ for some $j \in [1,v]$ implies that in the interval $(0,1]$, the total departure rate is strictly greater than the total pre-arrival rate, i.e., $\sum_{j=1}^v \lambda_1^j > \sum_{j=1}^v \lambda_{1+j}^j$. For subsequent intervals $(i,i+1]$ for $i \geq 1$, we have $\sum_{j=1}^v \lambda_{1+i}^j \geq \sum_{j=1}^v \lambda_{1+j+i}^j$. Therefore the conditions of Theorem 3.3.4 are satisfied.

Theorem 3.3.4 implies that

$$
\begin{aligned}
& \lim_{\lambda\to\infty} \mathbb{P}(A_0 \geq C) \qquad (3.17)\\
= \;& \lim_{\lambda\to\infty} \mathbb{P}\left(\left(\sum_{j=2}^{v}\sum_{l=2}^{j} N_l^j(1;\lambda_l^j) + \sum_{j=1}^{v} \tilde{N}_1^j(1;\lambda_1^j)\right) \geq C\right)\\
= \;& \lim_{\lambda\to\infty} \mathbb{P}\left(\text{Poisson}\left(\sum_{j=2}^{v}\sum_{l=2}^{j} \lambda_l^j + \sum_{j=1}^{v} \lambda_1^j\right) \geq C\right)\\
= \;& \lim_{\lambda\to\infty} \mathbb{P}\left(\text{Poisson}\left(\sum_{j=1}^{v} j\lambda_1^j\right) \geq C\right)\\
= \;& \lim_{\lambda\to\infty} \mathbb{P}(\text{Poisson}(\rho) \geq C) = \Phi(-\beta),
\end{aligned}
$$

and similarly $\lim_{\lambda\to\infty} \mathbb{P}(A_i \geq C) = 0$ for $i \in [1,u]$. The third equality of (3.17) follows from (3.14) in Lemma 3.3.9. $\square$

Theorem 3.3.11 *Assume that there exists $j \in [1,v]$, such that $\gamma_0^j > 0$. The traffic intensity is $\rho = \sum_{j=1}^{v} j\kappa_j\lambda = \sum_{j=1}^{v} j\lambda_1^j$. Then we have, for each $j \in [1,v]$,*

$$\lim_{\lambda\to\infty} P_0^j = \Phi(-\beta); \qquad \lim_{\lambda\to\infty} P_i^j = 0, \qquad i \in [1,u]. \tag{3.18}$$

Proof of Theorem 3.3.11. For each $j \in [1,v]$, by union bound and Lemma 3.3.10, we have

$$
\begin{aligned}
& \lim_{\lambda\to\infty} P_0^j = \lim_{\lambda\to\infty} \mathbb{P}(\max(A_0,\ldots,A_{j-1}) \geq C) \leq \lim_{\lambda\to\infty} \mathbb{P}\Big(\bigcup_{i=0}^{j-1} A_i \geq C\Big) \qquad (3.19)\\
\leq \;& \lim_{\lambda\to\infty} \sum_{i=0}^{j-1} \mathbb{P}(A_i \geq C) = \lim_{\lambda\to\infty} \mathbb{P}(A_0 \geq C) = \Phi(-\beta).
\end{aligned}
$$

On the other hand, it is obvious that, for each $j \in [1,v]$,

$$\lim_{\lambda\to\infty} P_0^j = \lim_{\lambda\to\infty} \mathbb{P}(\max(A_0,\ldots,A_{j-1}) \geq C) \geq \lim_{\lambda\to\infty} \mathbb{P}(A_0 \geq C) = \Phi(-\beta),$$

Similarly, $\lim_{\lambda\to\infty} P_i^j = 0$ for each $i \in [1,u]$ and $j \in [1,v]$. $\square$

As discussed previously, we can relax the assumption that $\gamma_0^j > 0$ for some $j \in$

$[1, v]$. Suppose now $\gamma_0^j = 0$ for all $j \in [1, v]$. This implies that no arriving customers at time 0 will start the service over $(0, 1]$, and hence we can ignore the blocking probability over this interval. Let i' be the minimal index such that $\gamma_{i'}^j > 0$ for some $j \in [1, v]$ and $\gamma_i^j = 0$ for $i < i'$ and all $j \in [1, v]$. Observe that no arriving customers at time 0 will start the service over $(0, i' + 1]$. For the subsequent interval $(i' + 1, i' + 2]$, the total departure rate is strictly greater than the pre-arrival rate, i.e., $\sum_{j=1}^{v} \lambda_{j+i'}^j > \sum_{j=1}^{v} \lambda_{1+j+i'}^j$. Thus, it suffices to show that for each $j \in [1, v]$, $\lim_{\lambda \to \infty} P_{i'}^j = \lim_{\lambda \to \infty} \mathbb{P}(\max(A_{i'}, \ldots, A_{i'+j-1}) \geq C) = \Phi(-\beta)$ and $\lim_{\lambda \to \infty} P_i^j = 0$ for each $i \in [i' + 1, u]$ and $j \in [1, v]$. The same arguments in Theorem 3.3.11 carry through.

3.4 An Improved Policy

In this section, we propose a variant of the CSP that improves the asymptotic worst-case performance guarantee. Fix a small $\epsilon < \min_k(\lambda_k \mu_k C^{-1})$, the variant solves the following LP:

$$\begin{aligned} \max_{\alpha_{ijk}} \quad & \sum_{k=1}^{M} \sum_{i,j} r_k \alpha_{ijk}^{\pi} \lambda_{ijk} j, \\ \text{s.t.} \quad & \sum_{k=1}^{M} \sum_{i,j} \alpha_{ijk}^{\pi} \lambda_{ijk} j \leq (1 - \epsilon) C, \qquad 0 \leq \alpha_{ijk} \leq 1, \qquad \forall i, j, k. \end{aligned} \tag{3.20}$$

Note this LP defined in (3.20) differs from the original LP defined in (3.1) by changing the right hand side of the capacity constraint to $(1-\epsilon)C$. Suppose the optimal solution of the original LP defined in (3.1) is $\{\alpha_{ijk}^*\}$. Now consider $\{\tilde{\alpha}_{ijk}\} = \{(1 - \epsilon)\alpha_{ijk}^*\}$. Since

$$\sum_{k=1}^{M} \sum_{i,j} \tilde{\alpha}_{ijk} \lambda_{ijk} j = (1 - \epsilon) \sum_{k=1}^{M} \sum_{i,j} \alpha_{ijk}^* \lambda_{ijk} j \leq (1 - \epsilon) C.$$

$\{\tilde{\alpha}_{ijk}\}$ is a feasible solution to (3.20). Let $\{\hat{\alpha}_{ijk}\}$ be the optimal solution of (3.20). Then we have

$$\sum_{k=1}^{M}\sum_{i,j} r_k \hat{\alpha}_{ijk}\lambda_{ijk} j \geq \sum_{k=1}^{M}\sum_{i,j} r_k \tilde{\alpha}_{ijk}\lambda_{ijk} j = (1-\epsilon)\sum_{k=1}^{M}\sum_{i,j} \alpha^*_{ijk}\lambda_{ijk} j = (1-\epsilon)\mathcal{R}(OPT)$$

It is easy to verify that the optimal solution (again omitting i and j) to (3.20) is $\hat{\alpha}_1 = \ldots = \hat{\alpha}_{M'-1} = 1$, $\hat{\alpha}_{M'} = 1 - \epsilon C(\lambda_{M'}\mu_{M'})^{-1}$, $\hat{\alpha}_{M'+1} = \ldots = \hat{\alpha}_M = 0$. This solution then gives rise to the modified CSP: for each $k = 1, \ldots, M'-1$, accept the customer upon arrival if there is sufficient unreserved capacity throughout the requested service interval. If $k = M'$, accept with probability $1 - \epsilon C(\lambda_{M'}\mu_{M'})^{-1}$ if there is sufficient unreserved capacity throughout the requested service interval. For $k = M'+1, \ldots, M$, reject.

Theorem 3.4.1 *Consider the revenue management model with a single pool of capacitated reusable resources and advanced reservations under the modified CSP. For any* ϵ *such that* $0 < \epsilon < \min_k(\lambda_k \mu_k C^{-1})$,

(a) The blocking probabilities Q_{ijk}*'s are asymptotically zero, i.e.,* $\lim_{\rho\to\infty} Q_{ijk} = 0$.

(b) The CSP is guaranteed to obtain at least $1-\epsilon$ *of the optimal long-run expected revenue in the Halfin-Whitt heavy-traffic limit.*

3.5 Price-Driven Customer Arrivals

In this section, we follow Levi and Radovanovic (2010) and consider an extension of the model discussed in previous sections, in which the arrival rates of the different classes of customers are affected by prices. Specifically, consider a two-stage decision. At the first stage, we set the respective prices $r_1, \ldots, r_M$ for each class. This determines the respective arrival rates $\lambda_1(r_1), \ldots, \lambda_M(r_M)$. (The rate of class-$i$ customers is affected only by price r_i.) Then, given the arrival rates, we wish to find the optimal admission policy that maximizes the expected long-run revenue rate. In particular, we assume that $\lambda_i(r_i)$ is nonnegative, differentiable, and decreasing in r_i

for each $1 \leq i \leq M$. In addition, we assume that all prices are nonnegative real numbers and that there exists a price r_∞ such that, for each $i = 1, \ldots, M$, we have $\lambda_i(r_\infty) = 0$. (The latter condition is required to guarantee that the problem has an optimal solution.)

Using arguments analogous to the discussion in Section 2, we construct an upper bound on the achievable expected long-run revenue rate through a nonlinear program, and then use it to construct a similar policy with the same performance guarantees. The detailed discussion and analysis can be found in the electronic companion.

Using some of the techniques developed in this chapter, in follow up work (Levi and Shi (2011a)), we study a dynamic pricing model and derive provably near-optimal policies.

3.6 Numerical Experiments

In this section, we conduct some numerical experiments to find out the empirical blocking probabilities in the original capacitated system. We run tests on four different reservation distributions and six different service distributions. The four reservation distributions considered are as follows,

$$\begin{aligned} D_1 &\sim \text{Uniform}(10)/1000, & D_2 &\sim \text{Binomial}(10, 0.5)/1000, \\ D_3 &\sim \text{Poisson}(10)/1000, & D_4 &\sim \text{Hypergeometric}(100, 40, 50)/1000. \end{aligned}$$

The six service distributions considered are as follows,

$$\begin{aligned} S_1 &\sim 1, & S_2 &\sim \text{Exponential}(1), \\ S_3 &\sim \text{Uniform}(10)/10, & S_4 &\sim \text{Binomial}(10, 0.5)/10. \\ S_5 &\sim \text{Poisson}(10)/10, & S_6 &\sim \text{Hypergeometric}(100, 40, 50)/10. \end{aligned}$$

The arrival process follows a Poisson process with rate λ ranging from 1 to 100, and the capacity is pegged with the traffic intensity. The total number of experiments ran is $100 \times 4 \times 6 = 2400$. The computational results are shown in Figure 3-7. We observe that the blocking probabilities in the original capacitated system go to zero in the two heavy-traffic limits. This implies that the blocking probabilities converge to zero asymptotically under the Halfin-Whitt regime (see graphs in the appendix), and perhaps the asymptotic analysis could be tightened in the future.

3.7 Conclusion and Future Directions

This chapter derives asymptotic upper bounds on the blocking probabilities in loss network systems with *advanced reservations* under the Halfin-Whitt regime. The theoretical results find applications in a class of revenue management problems in systems with reusable resources and advanced reservations. A simple control policy called the class selection policy (CSP) is proposed based on solving a knapsack-type linear program (LP). It is shown that the CSP and its variants perform provably near-optimal under the Halfin-Whitt regime.

There are several issues that still remain open. From the comparison of the upper bounds and the simulation results, it is clear that there is a gap between the empirical blocking probabilities and the theoretical bounds. This gap is due to the approximation using infinite capacity systems. It opens an opportunity to tighten the upper bound using another fictitious system between the original capacitated system and the infinite capacity counterpart.

There are also several plausible extensions into pricing models. The follow-up work (Levi and Shi (2011a)) will study both the static and dynamic pricing model of reusable resources with advanced reservations. The static pricing model allows the arrival rates being affected by prices. Specifically, consider a two-stage decision. At the first stage, we set the respective prices $r_1, \dots, r_M$ for each class. This determines the respective arrival rates $\lambda_1(r_1), \dots, \lambda_M(r_M)$. (The rate of class-$i$ customers is affected only by price r_i.) Then, given the arrival rates, we wish to find the optimal admission policy that maximizes the expected long-run revenue rate. In particular, we assume that $\lambda_i(r_i)$ is nonnegative, differentiable, and decreasing in r_i for each

$1 \leq i \leq M$. In addition, we assume that all prices are nonnegative real numbers and that there exists a price r_∞ such that, for each $i = 1, \ldots, M$, we have $\lambda_i(r_\infty) = 0$. (The latter condition is required to guarantee that the problem has an optimal solution.) We construct an upper bound on the achievable expected long-run revenue rate through a nonlinear program, and then use it to construct a similar policy with the same performance guarantees.

In the dynamic pricing model, consider a single-class time-homogenous Poisson arrival process with rate λ. Each customer's reservation and service-time are drawn from D and S, respectively. The system offers a price from a fixed price menu $[r_1, \ldots, r_n]$ to an arriving customer with d and s, depending on the current state. The state is characterized by the booking profile, d, and s. Moreover, we introduce a reservation price distribution denoted by R. The customer only accepts the offer if the price offered falls below the reservation price. We construct a different linear program, and use it to obtain provably near-optimal randomized policies.

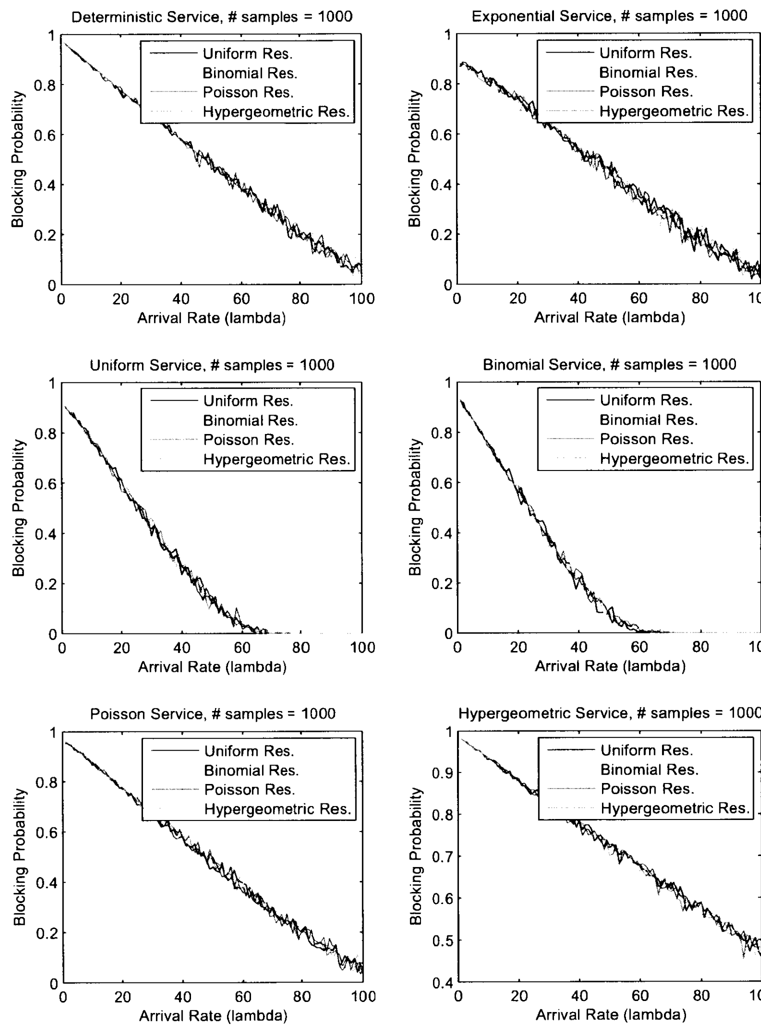

Figure 3-7: Computational Results

Chapter 4

Joint-ventures in Operations Management

4.1 Introduction

A proliferation of joint ventures has been witnessed across the globe in the recent years (see Bamford et al. (2004)). A joint venture takes place when two or more business partners pool their resources and expertise to achieve a particular goal for a contractual period of time. Joint ventures stand in the middle ground between non-cooperative competition and merging. They provide companies with the opportunities to gain new capacity and expertise, enter related businesses or new geographic markets, gain new technological knowledge access to greater resources, and share risks with other venture partners.

In this work, we consider a setting where multiple entities take part in a joint venture and each of them contributes one type of resources. We distinguish two types of resource pooling in joint ventures, depending on whether the resources are *heterogeneous* or *homogeneous*. When resources are heterogeneous, they are not fully substitutable. Thus, the effective capacity of a joint venture is limited to the minimum level of an individual contribution. In other words, the lowest contribution by one partner becomes the bottleneck in planning the capacity for the joint venture. This is in contrast with homogeneous resource pooling, where the resources are perfectly substitutable and the overall capacity of a joint venture is determined by aggregating all individual contributions.

One example that demonstrates the success of a joint venture with heterogeneous resource sharing is Massachusetts Eye and Ear Infirmary (MEEI), a hospital specialized in providing patient care for disorders of the eye, ear, nose, throat, head and neck in downtown Boston. With the vast majority of its services is outpatient in nature, MEEI experiences lower profit margins than a regular hospital and has been pressured to increase its patient volume so as to strengthen its financial status. Since 2005, MEEI has established five satellite clinics through joint ventures by collaborating with community hospitals in the suburbs. A typical agreement specifies that MEEI provides expertise (physicians and nurses) and its brand name[1] while the community hospital is responsible for providing facility and other necessary hardware. The two types of resources, i.e., expertise and facility, are not interchangeable. The maximum number of services that can be supported in such a satellite clinics is limited by MEEI's input as well as the space constraints such as the number of operating rooms available in the new location.

In 2003, US-based car rental firm Avis has set up a joint venture in Shanghai, China. The new company named Anji Car Rental and Leasing, 50-50 owned by Avis Europe and Shanghai Automotive Industry Sales Corporation, takes over the existing fleet of 1,000 vehicles from Shanghai Anji Car Rental and operates it under the Avis brand name. The venture expects to establish more than 70 outlets nationwide. This is a typical joint venture with homogeneous resource sharing, where the capacity in the new company is supported by aggregating the number of vehicles from the two companies.

Besides the healthcare industry and car-rental industry, another sector which has been a flurry to establish joint ventures is the airline industry. An airline alliance is an agreement between multiple independent partners to collaborate in various activities to streamline costs while expanding global reach and market penetration. The presence of alliances in the airline industry has followed an increasing trend since

[1]The satellite clinic located within the community hospital is labeled as a MEEI branch.

the first large airline alliance was formed in 1989 between Northwest and KLM. By March 2009, the three major alliances (Star, Sky Team and Oneworld) combined flew around 73% of all passengers worldwide (Hu et al., 2012). On the cost side, there are strong incentives for airlines to operate large networks as the evidences on economies of scale have been well documented (Caves et al., 1984, Brueckner and Spiller, 1994, Keeler and Formby, 1994, etc). On the revenue side, one of the fundamental attractions of an airline alliance is the ability to offer codeshare fights. Code sharing is an agreement between two carriers whereby one carrier allows a different carrier to market and sell seats on some of its flights. Based on empirical evidences, Brueckner (2003) conclude that codesharing among Star Alliance partners yielded an annual benefit of around $20 million. Morever, the information comes with codesharing can be tremendously beneficial. Jain (2010) show that sharing information on bid prices yields higher revenues of the order of $100 million for every big partnering carrier in the alliance.

4.1.1 Results and Contributions

In this work, we study both types of joint venture models and address some issues pertinent to the success of joint ventures. When several companies agree to a partnership, disparate interests often exist as each participant is more concerned with his or her own gain. Given the misalignment in incentives and uncertainties in demand, we are interested in measuring the performance of a joint venture by quantifying the difference in the investment level and the total profit attained with respect to a system optimal outcome.

We have shown that in joint ventures with heterogeneous resource pooling, any Nash equilibrium induces an equal contribution from every player, despite of them being asymmetric. The intuition is that since the revenue received by each player depends solely on the bottleneck capacity (minimum capacity contributed by some single player) when resource-sharing is heterogeneous, any further investment beyond the bottleneck capacity only increases her cost and decreases her profit.

Although multiple Nash equilibria could exist, we show that there always exists a unique Strong Nash equilibrium. Next, we focus on a Nash Bargaining model which is a natural framework to define and design fair assignment of the capacity investment levels between multiple players. We conclude that there exists a unique revenue sharing contract such that the corresponding Nash Bargaining Solution, the Strong Nash equilibrium, and the system optimal solution coincide. This revenue sharing contract indicates that the award each player receives must be equal to the ratio of her marginal cost to the total marginal cost bore by all partners evaluated at the optimal investment level.

For joint ventures with homogeneous resource pooling, we first prove some structural properties on the effective capacity under any demand distribution with convex costs. The analysis is challenging as the investment of each player could only be determined by solving a system of implicit equations. We show that joint venture always *underinvests* as the effective capacity is always lower than that of a coordinated setting.

We then focus on quadratic-linear cost functions and show that, through an intercept-argument, the effective capacity in a joint venture with respect to any revenue sharing ratio is *at least* $1/n$ of the optimal level. Moreover, the ratio between the capacity level could be upper bounded in terms of the cost asymmetry between the two players and the revenue sharing ratio. While we show that there does not exist a fixed marginal revenue sharing contract which can coordinate the players, we propose an interval for the revenue sharing ratio which induces an outcome that is guaranteed to achieve *at least 50%* of the optimal profit for a 2-player model. This interval depends on the cost asymmetry between the two players and the demand concentration.

Next, we consider general convex cost in the homogeneous resource pooling model with an arbitrary number of asymmetric players. We show that a lower bound to the efficiency of the original setting with the nonlinear convex costs is that of a modified setting with linear costs, where the coefficients are equal to the marginal cost of each player evaluated at the Nash equilibrium of the original problem. As a result, we show that the comparative analysis on profit can be reduced to analyze the joint investment level made in the Nash and the system in the setting with linear costs.

The rest of the chapter is organized as follows. We begin with a review on related literature in Section 4.1.2. Section 4.2 describes the two models and assumptions. We analyze and present the main results on capacity sharing and substitution model in Section 4.3 and Section 4.4 respectively.

4.1.2 Related Literature

This paper studies strategic capacity management under uncertainty. In the operations management literature, there is a vast body of work using the classic newsvendor model or some variations to capture uncertainties. Federgruen and Zipkin (1986) is the classic reference for capacitated inventory management. Papers including Kapuscinski and Tayur (1998), Angelus and Porteus (2002), Bradley and Glynn (2002), Van Mieghem and Rudi (2002) consider capacity investment decisions in capacitated Newsvendor networks. Van Mieghem and Dada (1999) take a different approach at capacity management and address how the relative timing of the decisions on capacity, inventory, and price impact the sensitivity and profitability. We refer readers to Van Mieghem (2003) for an excellent survey paper on the recent development on capacity management. In this work, the capacity of a joint venture depends on the contribution of multiple participants. Depending on the nature of the resources, the effective capacity can be the minimum or the sum of individual contributions.

In many settings, capacity-investment decisions are the results after interacting with other economic agents. Thus, it seems natural for capacity investment models to incorporate the strategic behavior of self-interested agents. Cachon and Lariviere (1999) consider the manufacturer's capacity investment and allocation decisions to several downstream retailers that have private information. Caldentey and Wein (2003) present contracts that are linear in backorder, inventory, and capacity levels to coordinate a manufacturer and retailer production-inventory system, including the capacity decision. Examples on single-resource, multiple-agent also include Carr and Lovejoy (2000), Porteus and Whang (1991), Kouvelis and Lariviere (2000), etc. Bassok et al. (1999) and Netessine and Rudi (2003) explore the impact of substitution in an inventory context, and its effects are likely to be similar in capacity problems.

In this work, the strategic behavior of participants involving in a joint ventures is captured in a noncooperative game, as each entity determines his level of contribution with the goal to maximize his profit. While the revenue each party receives depends on the effective capacity of the joint venture, the incentive of each entity might not be correctly aligned to one which maximizes the collective return. We consider a fixed rate revenue sharing contract described in Cachon and Lariviere (2005) to split revenue among the participants. To capture the high capital investment incurred in joint ventures in the healthcare industry, we consider general convex cost function so as to capture the diminishing returns, in contrast to linear cost function which is common in the operations management literature (e.g., Bernstein and Federgruen, 2007, Cachon, 2003, Corbett et al., 2005, Martinez-de Alberniz and Simchi-Levi, 2009). In this setting, we show that an "optimal" coordinating contract which enables the parties with self-interests to behave as a coordinated entity does not necessarily exist with homogeneous resources. We then propose a range for fixed revenue sharing ratio which induces reasonably good outcomes.

Standing in the middle ground between non-cooperative competition and merging, one of the most fundamental building blocks of joint ventures is negotiation. Empirical studies suggest that "the power of a joint venture is only as strong as the negotiation behind it" (Luo and Shenkar (2002), Lin and Germain, 1998). The topic on negotiation has gained a lot of attraction in the economics literature since Nash (1950) (e.g., Myerson, 1979, Binmore et al., 1986, Rubinstein, 1982, etc). In the fast few years, more results on negotiation have become known in the field of operations management (see for example, Reyniers and Tapiero, 1995, Miller, 1992, Chod and Rudi, 2006, etc). Nagarajan and Sosic (2008) present an excellent survey paper on cooperative game theory in the field of supply chain management. In this work, utilizing the bargaining model, we propose a revenue sharing scheme which induces an outcome which is coincides with the system optimum.

Our work which measures the performance of an unregulated setting with respect to a centralized system is related to a stream of literature on *price of anarchy*, popularized by Koutsoupias and Papadimitriou (1999). It compares the performance of the worst-case Nash equilibrium with respect to the centralized system. The concept has been used in transportation networks (Roughgarden and Tardos (2002), Correa et al. (2004, 2007), Roughgarden (2005)), network pricing (Acemoglu and Ozdaglar (2007), Weintraub et al. (2010)), oligopolistic pricing games in a single tier (Farahat and Perakis (2010a,b)), and supply chain games with exogenous pricing (Perakis and Roels (2007), Martinez-de Alberniz and Simchi-Levi (2009); Martinez-de Alberniz and Roels (2010)).

4.2 Model Formulation

In this section, we first present the model for a joint venture with n players as an uncoordinated game. As a benchmark, we also present the model in the system setting, i.e., n entities were merged and coordinated as a single entity with the goal to maximize the total return.

4.2.1 Joint-venture: an uncoordinated game

Consider a joint venture with n profit-maximizing players with asymmetric cost functions. The joint venture generates a joint revenue $R(p, \mathbf{K})$ where p is the fixed price and $\mathbf{K} = (K_1, \ldots, K_n)$ captures the resources contributed by each player. A revenue-sharing contract dictates that player i receives revenue $\beta_i R(p, \mathbf{K})$. Let $f_i(K_i)$ be the convex cost associated with investing K_i resources by player i. Based on a pre-negotiated revenue-sharing ratio $\beta = (\beta_1, \ldots, \beta_n)$, player i tries to maximize her profit $\pi_i(\beta) \triangleq \beta_i R(p. \mathbf{K}) - f_i(K_i)$ by choosing her own investment level K_i, which leads to a Nash equilibrium (NE).

4.2.2 Merger: the system optimum

Consider the centralized system in which n players are merged and coordinated as a single player. The merger generates the highest possible profit $\pi_T^* \triangleq R(p, \mathbf{K}) -$

$\sum_{i=1}^{n} f_i(K_i)$ by collectively choosing the resource investment $\mathbf{K}$. This yields the first-best or system optimal solution.

4.2.3 Resource-sharing models

We consider two types of resource-sharing models depending on the nature of the resources pooled from different players. The nature of the resources determines the effective capacity in a joint venture, which in turn affects the revenue function $R(p, \mathbf{K})$. We formally define them as follows:

Definition Heterogeneous resource-sharing. The aggregate revenue generated by the joint venture is given by $R(p, \mathbf{K}) = p\mathbb{E}(\min(D, \min_i(K_i)))$.

The type of resource provided by each player is heterogeneous and not fully substitutable. A service can only be performed with a complete portfolio of resource types. The effective capacity supported by the joint venture is therefore limited to the minimum capacity level invested by the players.

Definition Homogeneous resource-sharing. The aggregate revenue generated by the joint venture is given by $R(p, \mathbf{K}) = p\mathbb{E}(\min(D, \sum_{i=1}^{n}(K_i)))$.

The type of resource provided by each player is homogeneous to each other and hence fully substitutable. A service can be performed by using the resource contributed by any (possibly single) player. The effective capacity supported by the joint venture is therefore the sum of capacity level invested by each player.

In the next two sections, we will study both types of resource-sharing models and present the differences in the capacity investment and the total profit generated in a joint venture to those in a system optimum.

4.3 Heterogeneous Resource-sharing Models

With heterogeneous resources, the effective capacity is limited by the minimum capacity invested among all players, which becomes the *bottleneck capacity*. Consider the merger setting, the central planner tries to maximize the aggregate revenue by collectively choosing the capacity investment $\mathbf{K}$, i.e.,

$$\pi_T^* \triangleq \max_{K,K_i} p\mathbb{E}[\min(K,D)] - \sum_{i=1}^{n} f_i(K_i), \text{ s.t. } K \leq K_i, \ i = 1,\ldots,n. \tag{4.1}$$

Let K^* and $K_1^*,\ldots,K_n^*$ be the system optimal solution.

Lemma 4.3.1 *At system optimality, the capacity invested by each player is the same, i.e., $K^* = K_i^*$ for all $i = 1,\ldots,n$, where K^* solves $\mathbb{P}(D \leq K^*) = 1 - \sum_{i=1}^{n} f_i'(K^*)/p$.*

Proof of Lemma 4.3.1. Without loss of generality, if there exists a pair of players i and j such that $K_i^* < K_j^*$, we can decrease the capacity invested by player j from K_j^* to K_i^*. By doing so, the profit increases by reducing the cost while maintaining the same revenue. Hence, we reach a contradiction. At system optimality, $K^* = K_i^*$ for all $i = 1,\ldots,n$, and (4.1) reduces to a single variable optimization in which K^* can be obtained by the first-order condition. $\square$

In the system optimum, each individual capacity investment K_i^* must be reduced to the bottleneck capacity K^* when resource-sharing is heterogeneous, since any further investment beyond the bottleneck capacity only increases the total cost and decreases the total profit.

In a joint-venture with a pre-negotiated revenue-sharing contract β, player i tries to maximize her profit by choosing her profit-maximizing capacity investment level K_i based on other players' strategies K_{-i}, which leads to a Nash equilibrium, i.e.,

$$\pi_i^N(\beta) \triangleq \max_{K,K_i;K_{-i}} \beta_i p\mathbb{E}[\min(K,D)] - f_i(K_i), \text{ s.t. } K \leq K_j, \ j = 1,\ldots,n,$$

Now, let K^N and $K_1^N,\ldots,K_n^N$ be the Nash equilibrium solutions.

Lemma 4.3.2 *In joint-ventures, any*

$$K^N(\beta) = K_1^N(\beta) = \ldots = K_n^N(\beta) \leq \min_{1\leq k\leq n}(A_k)$$

are Nash Equilibria, where A_k solves

$$\mathbb{P}(D \leq A_k) = 1 - \frac{f_k'(A_k)}{\beta_k p}.$$

In particular, $K^{SN}(\beta) = K_1^{SN}(\beta) = \ldots = K_n^{SN}(\beta) = \min_{1\leq k\leq n}(A_k)$ is a unique Strong Nash equilibrium.

Proof of Lemma 4.3.2. Without loss of generality, if there exists a pair of players i and j such that $K_i^N(\beta) < K_j^N(\beta)$, player j can decrease its capacity investment from $K_j^N(\beta)$ to $K_i^N(\beta)$ lowering her cost and improving her profit. Thus, at Nash equilibrium, all players must have the same capacity investment level, i.e., $K^N(\beta) = K_i^N(\beta)$ for all $i = 1, \ldots, n$.

Now assume that $\min_{1\leq k\leq n}(A_k) = A_m$. Now if $A_m < K^N(\beta) = K_m^N(\beta)$, player m always has incentives to unilaterally lower her investment level to A_m since A_m is her profit-maximizer. This forces all players to invest at A_m. Any capacity investment level $\tilde{A}_m$ such that $0 \leq \tilde{A}_m \leq A_m$ is also a Nash equilibrium since no player has incentives to unilaterally deviate from $\tilde{A}_m$. In particular, $K^{SN}(\beta) = K_1^{SN}(\beta) = \ldots = K_n^{SN}(\beta) = A_m$ is a unique *Strong Nash equilibrium* in which no coalition, taking the actions of its complements as given, can cooperatively deviate in a way that benefits all of its members. $\square$

Lemma 4.3.2 indicates that the capacity invested by each player must be the same in a joint venture. Since the revenue received by player i depends solely on the bottleneck capacity $K^N(\beta)$ when resource-sharing is heterogeneous, any further investment beyond the bottleneck capacity only increases her cost and decreases her profit. Lemma 4.3.2 also implies that A_k is the profit-maximizing capacity for player k. Since the resource-sharing is heterogeneous, the player m with the lowest profit-maximizing capacity (i.e., $A_m = \min_{1\leq k\leq n}(A_k)$) can unilaterally choose to invest at her profit maximizing capacity, forcing all other players to invest at the same capacity level. Note that any capacity investment level no greater than A_m is a Nash equilibrium whereas any capacity investment level above A_m is not. As a result, it is easy to see that with the existence of multiple Nash equilibria, it is possible for a joint venture to achieve an arbitrarily bad outcome compared to the system optimum.

So far, we have modeled the decision making process in a joint venture as a Nash Equilibrium. Next, we will propose an alternative model where the players participate a Nash bargaining game to determine their respective investment decisions for a given revenue sharing ratio β.

Nash Bargaining Solution (NBS). The Nash Bargaining Solution (see Appendix C) is a natural framework that allows us to define and design fair assignment of the capacity investment levels between n players, which can derive desirable properties such as Pareto efficiency and proportional fairness. Based on a particular revenue sharing contract β, n players choose their capacity investment levels according to a Nash Bargaining game, i.e.,

$$\max_{K,K_i} \prod_{i=1}^{n} \pi_i^B(\beta), \text{ s.t. } K \leq K_j, \; j = 1, \ldots, n,$$

which is equivalent to solving

$$\max_{K,K_i} \log \sum_{i=1}^{n} \pi_i^B(\beta), \text{ s.t. } K \leq K_j, \; j = 1, \ldots, n. \tag{4.2}$$

Let K^B and $K_1^B, \ldots, K_n^B$ be the Nash Bargaining Solution from solving (4.2).

Theorem 4.3.3 *There exists a unique revenue sharing contract,*

$$\beta_i^* = \frac{f_i'(K^*)}{\sum_{j=1}^{n} f_j'(K^*)}, \qquad i = 1, \ldots, n,$$

such that the Nash Bargaining Solution, the unique Strong Nash equilibrium, and the system optimal solution coincide, i.e., $K^B(\beta^*) = K^{SN}(\beta^*) = K^*$.

Proof of Theorem 4.3.3. Observe that (4.2) is equivalent to a single variable optimization,

$$\max_{K} \log \sum_{i=1}^{N} \left(\beta_i p \mathbb{E}[\min(K, D)] - f_i(K)\right). \tag{4.3}$$

The first-order condition gives us

$$\sum_{i=1}^{N} \frac{\beta_i p \mathbb{P}(D \geq K^B) - f_i'(K^B)}{(\beta_i p \mathbb{E}[\min(K^B, D)] - f_i(K^B))} = 0. \tag{4.4}$$

By Lemma 4.3.2, at Nash Equilibrium, $K^N \leq \min_{1 \leq k \leq n}(A_k)$, where A_k solves

$$\mathbb{P}(D \geq A_k) = \frac{f_k'(A_k)}{\beta_k p}.$$

This implies that

$$\beta_i p \mathbb{P}(D \geq K^N) - f_i'(K^N) \geq 0, \text{ for all } i = 1, \ldots, n. \tag{4.5}$$

Suppose that there exists a solution γ to both the Nash Bargaining game and the Nash equilibrium, i.e., $\gamma = K^N(\beta) = K^B(\beta)$. Then γ must satisfy (4.4) and (4.5) simultaneously, implying that

$$\beta_i p \mathbb{P}(D \geq \gamma) - f_i'(\gamma) = 0 \text{ for all } i = 1, \ldots, n. \tag{4.6}$$

If such γ exists, $\gamma = K^{SN}(\beta)$, i.e. γ is the unique Strong Nash equilibrium since $\gamma = A_1 = \ldots = A_n = \min_{1 \leq k \leq n}(A_k)$ by (4.6).

Now summing (4.6) over all players and $\sum_{i=1}^{n} \beta_i = 1$, we have

$$p \mathbb{P}(D \geq \gamma) - \sum_{i=1}^{n} f_i'(\gamma) = 0. \tag{4.7}$$

By (4.6) and (4.7), we know that β must be of the following form,

$$\beta_i = \frac{f_i'(\gamma)}{\sum_{j=1}^{n} f_j'(\gamma)}, \qquad i = 1, \ldots, n,$$

Moreover, note that by Lemma 4.3.1, (4.7) implies that $\gamma = K^*$. Since K^* is the unique system optimal solution, there exists a unique revenue sharing contract

$$\beta_i^* = \frac{f_i'(K^*)}{\sum_{j=1}^{n} f_j'(K^*)}, \qquad i = 1, \ldots, n,$$

such that $\gamma = K^* = K^{SN}(\beta^*) = K^B(\beta^*)$. $\square$

Theorem 4.3.3 shows that when resources are heterogeneous, there is a way to rely on the revenue sharing contract to eliminate the incentive misalignment among the players and induce the system optimal outcome. In addition, the way to do so is the same when the players' behavior is predicted by a Nash equilibrium as well as the Nash bargaining solution.

In addition, besides inducing the efficient decision, the optimal revenue sharing contract in Theorem 4.3.3 also embodies the notion of *proportional fairness*. For an investment level K^*, player i bears a marginal cost $f_i'(K^*)$ and the aggregate marginal cost is given by summing up the marginal cost of every player participating in the joint venture, $\sum_j f_j'(K^*)$. Theorem 4.3.3 specifics that the marginal revenue ratio which player i is entitled to receive (β_i) should be equal to the proportion of his marginal cost to the aggregate marginal cost ($f_i'(K^*)/\sum_j f_j'(K^*)$). In simple words, "fairness" in this context suggests that every participant in a joint venture should be awarded "proportionally" to the risk (cost) she has to undertake.

4.3.1 Numerical Examples

We conduct numerical studies to compare our approach with the existing approach adpoted by some joint-ventures (such as MEEI). In the existing model, joint-ventures set the their capacity investment level according to the long-run average demand, i.e. $K^{EX} = \mathbb{E}[D]$. In addition, they split the revenue based on how much each party invests in *total* capacity investment. More specifically, they set the revenue sharing parameter to be

$$\beta_i^{EX} = \frac{f_i(K^{EX})}{\sum_{j=1}^{n} f_j(K^{EX})}, \qquad i = 1, \ldots, n.$$

We consider a 2-player game with unit service price $p = 1200$. Assume that the demand follows a normal distribution, and the cost functions to be quadratic, i.e. $f_i(K_i) = a_i K_1^2/2 + b_i K_i + c_i$ for $i = 1, 2$. Without loss of generality, we let $a_1 = 1$, $a_2 = 0.5$, $b_1 = b_2 = 100$ and $c_1 = c_2 = 0$. Table 4.1 shows the simulation results.

Demand	Player 1					Player 2					Total		
	Share		Profit ($\times 10^5$)			Share		Profit			Profit ($\times 10^5$)		
	RS	EX	RS	EX	%	RS	EX	RS	EX	%	RS	EX	%
N(800,100)	63.8%	62.5%	2.19	1.70	29%	36.2	37.5	1.11	1.01	8.8%	3.30	2.72	21%
N(800,200)	63.5%	62.5%	2.06	1.40	47%	36.5	37.5	1.06	0.84	26%	3.12	2.24	39%
N(800,300)	63.3%	62.5%	1.86	1.10	69%	36.7	37.5	0.96	0.66	46%	2.83	1.77	60%
N(700,100)	63.5%	62.1%	2.12	1.77	20%	36.5	37.9	1.09	1.08	1.2%	3.21	2.84	13%
N(700,200)	63.3%	62.1%	1.92	1.47	31%	36.7	37.9	1.01	0.90	12%	2.93	2.37	23%
N(700,300)	63.0%	62.1%	1.69	1.17	45%	37.0	37.9	0.89	0.72	24%	2.58	1.89	37%

Table 4.1: Numerical results comparing the revenue-sharing contract (RS) with the existing contract (EX).

The simulation results show that our approach outperforms the existing approach by increasing the profit of both players. The profit increases in the variability of the demand distribution. Moreover, we observe that the proportional sharing scheme based on marginal costs (our approach) gives slightly more weight to the less cost-effective player as compared to the proportional sharing scheme based on total costs (the existing approach).

4.4 Homogeneous Resource-sharing Models

When resources are homogeneous, they are completely substitutable for one another. The effective capacity is therefore the sum of the individual capacity invested by each player. The alliances among airlines and car rental companies are some of the applications of this model.

In a merger (system), the central planner tries to maximize the aggregate revenue by collectively choosing the capacity investment $\mathbf{K}$, i.e.,

$$\pi_T^* \triangleq \max_{K_i} p\mathbb{E}[\min(L, D)] - \sum_{i=1}^{n} f_i(K_i).$$

where the total capacity investment L is the sum of all K_i's, i.e., $L \triangleq \mathbf{K}\mathbf{e}$ with $\mathbf{e}$

being the column vector with all one's.

Lemma 4.4.1 *Define an auxiliary function*

$$g(\hat{L}) \triangleq \max_{K_i} p\mathbb{E}[\min(L, D)] - \sum_{i=1}^{n} f_i(K_i), \quad s.t. \quad L \leq \hat{L}.$$

Then $g(\hat{L})$ is concave in $\hat{L}$ where $\hat{L}$ is the budget on total capacity investment.

Proof of Lemma 4.4.1. Suppose L^* is the optimal solution to the system problem. It is easy to see that for all $\hat{L} \geq L^*$, $g(\hat{L}) = \pi_T^*$. For all $L < L^*$, the budget constraint becomes tight. It suffices to show that

$$h(\hat{L}) = \min_{K_i} \sum_{i=1}^{n} f_i(K_i), \text{ s.t. } \sum_{i=1}^{n} K_i = \hat{L}$$

is convex in $\hat{L}$. For any $\lambda \in [0, 1]$,

$$\begin{aligned} h(\lambda \hat{L}_1 + (1-\lambda)\hat{L}_2) = \quad & \min_{K_i, K_i'} \sum_{i=1}^{n} f_i(\lambda K_i + (1-\lambda)K_i') \\ & \text{s.t. } \sum_{i=1}^{n} K_i = \hat{L}_1, \; \sum_{i=1}^{n} K_i' = \hat{L}_2, \end{aligned}$$

and

$$\begin{aligned} \lambda h(\hat{L}_1) + (1-\lambda)h(\hat{L}_2) = \quad & \min_{K_i, K_i'} \lambda \sum_{i=1}^{n} f_i(K_i) + (1-\lambda) \sum_{i=1}^{n} f_i(K_i') \\ & \text{s.t. } \sum_{i=1}^{n} K_i = \hat{L}_1, \; \sum_{i=1}^{n} K_i' = \hat{L}_2. \end{aligned}$$

By convexity of function f_i for $i = 1, \ldots, n$, for any K_i, we know that

$$f_i(\lambda K_i + (1-\lambda)K_i') \geq \lambda f_i(K_i) + (1-\lambda)f_i(K_i').$$

Taking the minimum with respect to the same constraints preserves the inequality, we have

$$h(\lambda \hat{L}_1 + (1-\lambda)\hat{L}_2) \geq \lambda h(\hat{L}_1) + (1-\lambda) h(\hat{L}_2).$$

This completes the proof. $\square$

In a joint-venture with a pre-negotiated revenue-sharing contract β, player i tries to maximize her profit by choosing her profit-maximizing capacity investment level K_i based on other players' strategy K_{-i}, i.e.,

$$\pi_i^N(\beta) \triangleq \max_{K, K_i; K_{-i}} \beta_i p \mathbb{E}[\min(L, D)] - f_i(K_i),$$

which leads to a Nash equilibrium.

Lemma 4.4.2 *The total capacity investment level in a joint-venture is no greater than that in a merger (system), i.e.,* $\sum_{i=1}^n K_i^N \leq \sum_{i=1}^n K_i^*$.

Proof of Lemma 4.4.2. Suppose that, without loss of generality, $K_1^N \geq K_1^*$. Then we have

$$\frac{\beta p - f_1'(K_1^N)}{\beta p} \leq \frac{\beta p - f_1'(K_1^*)}{\beta p} \leq \frac{p - f_1'(K_1^*)}{p}.$$

Take F^{-1} on both sides (F^{-1} is monotonely increasing, so the sign does not change), then we have $\sum_{i=1}^n K_i^N \leq \sum_{i=1}^n K_i^*$. $\square$

The result in Lemma 4.4.2 does not depend on demand distribution or symmetry among the players. It shows that the effective capacity in a joint venture is always lower compared to a system optimum. However, when the players have asymmetric costs, it is likely that some players over-invest as compared to their counterparts in the optimal setting. In particular, the individual contribution depend on the revenue sharing ratio β.

In contrast to the heterogeneous resource sharing case where an optimal revenue sharing method exists, one can show that there does not exist a fixed revenue sharing method which will induce the system optimal actions in the Nash equilibrium. In other words, there does not exist β such that $\pi_T^N(\beta) = \pi_T^*$.

In the rest of the section, we will investigate the following questions: (1) For a fixed revenue sharing ratio β, how is performance in a joint venture compared to the optimum. (2) How to choose β such that we can have some performance guarantee. We will first restrict ourselves to linear quadratic costs. We begin with a 2-player game and extend our results to a n-player setting. In the end of this section, we will consider n-player setting with general convex costs.

4.4.1 2-player game with linear-quadratic cost functions

Assume that the cost functions are linear-quadratic, i.e.,

$$f_1(K_1) = \frac{a_1(K_1 + b_1)^2}{2} + c_1, \qquad f_2(K_2) = \frac{a_2(K_2 + b_2)^2}{2} + c_2.$$

Without loss of generality, assume that $a_1 \geq a_2$. Now define $\bar{K}_1 = K_1 + b_1$ and $\bar{K}_2 = K_2 + b_2$, and their corresponding modified total capacity investment levels,

$$\bar{L}^N = L^N + b_1 + b_2, \qquad \bar{L}^* = L^* + b_1 + b_2.$$

Lemma 4.4.3 *For a 2-player game with any demand distribution D and linear-quadratic cost functions, for all $\beta_1 \leq 0.5$, the ratio of the total capacity investment level in the system to that in the joint-venture is upper and lower bounded by*

$$1 \geq \frac{\bar{L}^N}{\bar{L}^*} \geq \frac{\beta_1 a_2 + \beta_2 a_1}{a_1 + a_2} \geq \frac{1}{2}.$$

Proof of Lemma 4.4.3. The lower bound is proven by Lemma 4.4.2. Now we show how to obtain an upper bound by utilizing an intercept argument. By optimality conditions, we have

$$\mathbb{P}(D \leq K_1^N + K_2^N) = \frac{\beta_1 p - a_1(K_1^N - b_1)}{\beta_1 p} = \frac{\beta_2 p - a_2(K_2^N - b_2)}{\beta_2 p}.$$

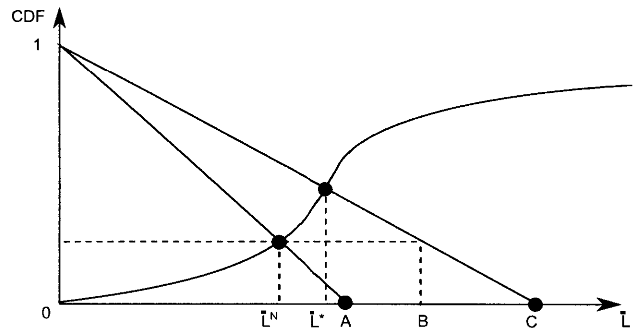

Figure 4-1: A graphical proof for Lemma 4.4.3.

By changing of variables,

$$\mathbb{P}(D + b_1 + b_2 \leq \bar{K}_1^N + \bar{K}_2^N) = \frac{\beta_1 p - a_1 \bar{K}_1^N}{\beta_1 p} = \frac{\beta_2 p - a_2 \bar{K}_2^N}{\beta_2 p}.$$

Then $\beta_2 a_1 \bar{K}_1^N = \beta_1 a_2 \bar{K}_2^N$ and we have

$$\bar{L}^N = \frac{\beta_1 a_2 + \beta_2 a_1}{\beta_1 a_2} \bar{K}_1^N, \qquad \text{or} \qquad \bar{L}^N = \frac{\beta_1 a_2 + \beta_2 a_1}{\beta_2 a_1} \bar{K}_2^N.$$

Thus, we have

$$\mathbb{P}(D + b_1 + b_2 \leq \bar{L}^N) = 1 - \frac{1}{p} \left(\frac{a_1 a_2}{\beta_1 a_2 + \beta_2 a_1} \right) \bar{L}^N. \tag{4.8}$$

By the similar transformation of the first-order condition in the system optimal, we have

$$\mathbb{P}(D + b_1 + b_2 \leq \bar{L}^*) = 1 - \frac{1}{p} \left(\frac{a_1 a_2}{a_1 + a_2} \right) \bar{L}^* \tag{4.9}$$

As shown in Figure 4-1, the horizontal axis is the modified total capacity investment level and the vertical axis is the cumulative distribution function of the demand. The upward sloping curve (cumulative distribution function) represents the left hand sides of (4.8) and (4.9), and the two downward sloping lines represent the right hand sides of (4.8) and (4.9). Thus, $\bar{L}^N$ and $\bar{L}^*$ can be solved graphically. We also observe that

$$\frac{\bar{L}^*}{\bar{L}^N} \leq \frac{B}{\bar{L}^N} = \frac{C}{A} = \frac{a_1 + a_2}{\beta_1 a_2 + \beta_2 a_1}.$$

where the points C and A are the x-intercepts which can be evaluated from (4.8) and (4.9). $\square$

Lemma 4.4.3 shows that for a 2-player game with linear-quadratic costs, the effective modified capacity in a joint venture depends on both the cost asymmetry as well as the revenue sharing ratio. However, the worst case, $\bar{L}^* = 2\bar{L}^N$, can happen under two circumstances: (1) equal revenue sharing ($\beta_1 = \beta_2$) and independent of cost asymmetry, and/or (2) with symmetric players ($a_1 = a_2$) and independent of revenue sharing contracts (with the assumption that $\beta_1 \leq 0.5$. Intuitively, dividing revenue equally among asymmetric entities sounds like a bad idea. It is surprising to see that having symmetric players in a joint venture could lead to the worst outcome, and having different revenue sharing contracts might not mitigate its impact. Note that when $\beta_1 > 0.5$, it is easy to construct examples that worst case becomes unbound.

Lemma 4.4.3 also highlights a notable difference between the homogeneous and the heterogeneous resource pooling. Note that in Theorem 4.3.3 for the heterogeneous resources, we have shown that the optimal revenue sharing rule suggests that every player should be compensated proportionally to his share of the marginal cost to the aggregate marginal cost. That if, if $a_1 \geq a_2$, the optimal way to share revenue must follow that $\beta_1 \geq \beta_2$. Lemma 4.4.3 implies the exact opposite, i.e., in order to have the worst case performance guarantee, given $a_1 \geq a_2$, then $\beta_1 \leq \beta_2$!

The intuition is that for heterogeneous resource pooling, the effective capacity of the entire system is constrained by a bottleneck capacity due to certain key players. To induce these players to produce at K^*, they have to be awarded such that they are willing to produce at K^* but not lower. Now consider homogeneous resource pooling, every player can contribute to the effective capacity, the only difference is the cost. Therefore, one should encourage the cost efficient player to produce more and discourage those with higher cost. It is captured by a lower revenue sharing ratio for the player with higher marginal cost.

This observation on a 2-player game can be generalized to a n-player game as shown in the following proposition.

Proposition 4.4.4 *Consider a n-player game with cost structure $a_1 \geq a_2 \cdots \geq a_n$ and revenue sharing contract $\beta_1 \leq \beta_2 \cdots \leq \beta_n$. Under any demand distribution D and any linear-quadratic cost functions, the ratio of the total capacity investment level in the system to that in the joint-venture is upper and lower bounded by*

$$1 \geq \frac{\bar{L}^N}{\bar{L}^*} \geq \frac{\sum_{i=1}^n \frac{\beta_i}{a_i}}{\sum_{i=1}^n \frac{1}{a_i}} \geq \frac{1}{n}.$$

With n-players, the worst case in terms of the effective capacity is $\bar{L}^* = n\bar{L}^N$, i.e., the worst case of a joint venture decreases as the number of participants increases. The result is intuitive as with more parties involved, it becomes increasingly challenging to coordinate the joint venture. Similar to the 2-player game studied earlier, the worst case occurs with symmetric players and/or equal sharing of the revenue when players are asymmetric.

In the next theorem, we will show that the profit generated in a joint venture can be bounded by the optimal profit.

Theorem 4.4.5 *For a 2-player game with any linear-quadratic cost functions and any demand distribution with mode m, we have*

$$\frac{\pi_T^N(\beta)}{\pi_T^*} \geq \frac{1}{2}, \qquad \textit{for all } a_1 \geq a_2 \textit{ and } \beta_1 \in \left[\frac{mp+1}{2mp + (a_1/a_2+1)}, \frac{1}{2}\right].$$

Moreover, the optimal β_1^* that maximizes the total joint-venture profit falls in the following interval,

$$\beta_1^* \in \left[\frac{1}{a_1/a_2+1}, \frac{mp+1}{2mp+(a_1/a_2+1)}\right].$$

Proof of Theorem 4.4.5. By Lemma 4.4.3, we know that

$$\bar{L}^N = \frac{\beta_1 a_2 + \beta_2 a_1}{\beta_1 a_2} \bar{K}_1^N, \qquad \text{or} \qquad \bar{L}^N = \frac{\beta_1 a_2 + \beta_2 a_1}{\beta_2 a_1} \bar{K}_2^N.$$

The Nash profit functions can be expressed as functions of $\bar{L}^N$ i.e.,

$$\pi_T^N(\beta) = p\mathbb{E}[\min(\bar{L}^N - b_1 - b_2, D)] - \left(\frac{a_1 a_2^2 \beta_1^2 + a_2 a_1^2 \beta_2^2}{2(a_2\beta_1 + a_1\beta_2)^2}\right) \bar{L}^{N2} - c_1 - c_2. \quad (4.10)$$

If we impose a budget constraint $L \leq \bar{L}^N$ on the system optimal, the budget-constrained system optimal profit can also be expressed as functions of $\bar{L}^N$ i.e.,

$$g(\bar{L}^N) \quad = \quad p\mathbb{E}[\min(\bar{L}^N - b_1 - b_2, D)] - \left(\frac{a_1 a_2}{2(a_2 + a_1)}\right) \bar{L}^{N2} - c_1 - c_2,$$

Observe that $g(\bar{L}^N) = \pi_T^N(\beta)$ when $\beta_1 = \frac{1}{2}$. By Lemma 4.4.3, we know that for all $\beta_1 \leq \frac{1}{2}$, $\frac{\bar{L}^*}{\bar{L}^N} \leq 2$. In addition, $g(\bar{L})$ is concave in $\bar{L}$ by Lemma 4.4.1. Thus, we have

$$\frac{\pi_T^N(\beta_1 = \frac{1}{2})}{\pi_T^*} = \frac{\pi_T^N(\beta_1 = \frac{1}{2})}{g(\bar{L}^*)} \geq \frac{\pi_T^N(\beta_1 = \frac{1}{2})}{2g(\bar{L}^*/2)} \geq \frac{\pi_T^N(\beta_1 = \frac{1}{2})}{2g(\bar{L}^N)} = \frac{1}{2}.$$

Now let $\bar{D} = D + b_1 + b_2$. Since

$$\mathbb{P}(\bar{D} \leq \bar{L}^N) = 1 - \frac{1}{p}\left(\frac{a_1 a_2}{\beta_1 a_2 + (1-\beta_1)a_1}\right) \bar{L}^N$$

$$\Rightarrow \quad \frac{\mathbf{d}\bar{L}^N}{\mathbf{d}\beta_1} = \frac{-\frac{\bar{L}^N}{p}\left(\frac{a_1 a_2 (a_1 - a_2)}{(\beta_1 a_2 + (1-\beta_1)a_1)^2}\right)}{f_{\bar{D}}(\bar{L}^N) + \frac{1}{p}\left(\frac{a_1 a_2}{\beta_1 a_2 + (1-\beta_1)a_1}\right)}.$$

We have by (4.10),

$$
\begin{aligned}
&\frac{\mathbf{d}\pi_T^N(\beta_1)}{\mathbf{d}\beta_1} \\
= \; & \frac{-(1-\mathbb{P}(\bar{D} \le \bar{L}^N))\bar{L}^N\left(\frac{a_1a_2(a_1-a_2)}{(\beta_1a_2+(1-\beta_1)a_1)^2}\right)}{f_{\bar{D}}(\bar{L}^N)+\frac{1}{p}\left(\frac{a_1a_2}{\beta_1a_2+(1-\beta_1)a_1}\right)} \\
& + \frac{a_1^2a_2^2(1-2\beta_1)}{(a_2\beta_1+a_1(1-\beta_1))^3}\bar{L}^{N2} - \left(\frac{a_1a_2^2\beta_1^2+a_2a_1^2(1-\beta_1)^2}{(\beta_1a_2+(1-\beta_1)a_1)^2}\right)\bar{L}^N\frac{\mathbf{d}\bar{L}^N}{\mathbf{d}\beta_1} \\
= \; & \frac{\bar{L}^{N2}}{(\beta_1a_2+(1-\beta_1)a_1)^3}\left(-\frac{a_1^2a_2^2(a_1-a_2)(a_1+a_2)\beta_1'(1-\beta_1)}{p(\beta_1a_2+(1-\beta_1)a_1)f_{\bar{D}}(\bar{L}^N)+a_1a_2}+a_1^2a_2^2(1-2\beta_1)\right).
\end{aligned}
$$

If the mode of D is m, then $\pi_T^N(\beta_1)$ is decreasing in β_1 for all

$$\beta \in \left[\frac{mp+a_2}{2mp+a_1+a_2}, \frac{1}{2}\right],$$

and $\pi_T^N(\beta_1)$ is increasing in β_1 for all

$$\beta \in \left[0, \frac{a_2}{a_1+a_2}\right].$$

Thus, the optimal β_1^* lies in the following interval

$$\beta^* \in \left[\frac{a_2}{a_1+a_2}, \frac{mp+a_2}{2mp+a_1+a_2}\right].$$

This completes the proof. □

In Theorem 4.4.5, we propose an interval for which the aggregate Nash profit is guaranteed to achieve at least half of the optimal profit. The interval depends on the cost asymmetry between the two players and the mode of demand. In particular, the interval shrinks as the two players have more similar cost structure, i.e., with two fully symmetric players, the best revenue sharing ratio asks for an equal division of the revenue. On the other hand, the interval widens as the mode of demand increases, i.e., if the demand distribution is flatter, our proposed revenue sharing contracts have more rooms for error in capturing the peak demand.

For a n-player game, we show that an equal revenue sharing scheme could guarantee a worst case performance of at least $1/n$ of the optimal profit as shown in the following proposition.

Proposition 4.4.6 *For a n-player game with any linear-quadratic cost functions and any demand distribution, if we choose $\beta_i = 1/n$, i.e., dividing the aggregate revenue equally among all the players, we have*

$$\frac{\pi_T^N(\beta)}{\pi_T^*} \geq \frac{1}{n}.$$

Proof of Proposition 4.4.6. From a 2-player setting, one can see that the profit functions can be expressed as functions of $\bar{L}$,

$$\pi_T(\bar{L}) \quad = \quad p\mathbb{E}[\min(\bar{L} + b_1 + b_2, D)] - \left(\frac{a_1}{2(a_2 + a_1)}\right)\bar{L}^2 - c_1 - c_2.$$

Note that it is equivalent to $\pi_T^N(\beta_1, \beta_2)$ when $\beta_1 = \beta_2 = 0.5$, where

$$\pi_T^{\mathrm{N}}(\beta) \quad = \quad p\mathbb{E}[\min(\bar{L}^{\mathrm{N}} + b_1 + b_2, D)] - \left(\frac{a_1 a_2 \beta_1^2 + a_1^2 \beta_2^2}{2(a_2\beta_1 + a_1\beta_2)^2}\right)\bar{L}^{\mathrm{N}2} - c_1 - c_2.$$

In Lemma 4.4.1, we have shown the concavity of $\pi_T(\bar{L})$. Then by making use of the bound on investment level as shown in Proposition 4.4.4, we obtain the desired result. $\square$

4.4.2 n-player game with general convex costs

We consider n-player games with asymmetric convex cost functions. Denote $f = (f_i(K_i))_{i=1}^n$ as general convex cost functions. Let $\pi^N(f)$ and $\pi^*(f)$ be the Nash and system profit of n players with respect to the general cost f, respectively. Define the Price of Anarchy with respect to f as

$$POA(f) = \frac{\pi^N(f)}{\pi^*(f)}.$$

We first show that $POA(f)$ can be lower bounded by $POA(\bar{f})$ where $\bar{f}$ is a set of modified *linear* cost functions.

Proposition 4.4.7 *The price of anarchy on the total profit of a joint venture is lower bounded by*

$$POA(f) = \frac{\pi^N(f)}{\pi^*(f)} \geq \frac{\pi^N(\bar{f})}{\pi^*(\bar{f})} = POA(\bar{f}),$$

where $\bar{f} = (\bar{f}_1, \ldots, \bar{f}_n)$ *are linear cost functions such that* $\bar{f}_i = \alpha_i \cdot K_i$ *where* $\alpha_i = f_i'(K_i^N)$.

Proof of Proposition 4.4.7. By convexity of f_i for all $i = 1, \ldots, n$, we know that

$$f_i(K_i^*) \geq f_i(K_i^N) + f_i'(K_i^N)(K_i^* - K_i^N).$$

Therefore

$$\begin{aligned} POA(f) &= \frac{p\mathbb{E}[\min(L^N, D)] - \sum_{i=1}^n f_i(K_i^N)}{p\mathbb{E}[\min(L^*, D)] - \sum_{i=1}^n f_i(K_i^*)} \\ &\geq \frac{p\mathbb{E}[\min(L^N, D)] - \sum_{i=1}^n f_i(K_i^N)}{p\mathbb{E}[\min(L^*, D)] - \sum_{i=1}^n (f_i(K_i^N) + f_i'(K_i^N)(K_i^* - K_i^N))}. \end{aligned} \tag{4.11}$$

Since

$$0 = f_i(0) \geq f_i(K_i^N) + f_i'(K_i^N)(-K_i^N) \Rightarrow f_i(K_i^N) - f_i'(K_i^N)(K_i^N) \leq 0, \tag{4.12}$$

we add (4.12) onto both the numerator and denominator of (4.11),

$$POA(f) \geq \frac{p\mathbb{E}[\min(L^N, D)] - \sum_{i=1}^n f_i'(K_i^N)(K_i^N)}{p\mathbb{E}[\min(L^*, D)] - \sum_{i=1}^n f_i'(K_i^N)(K_i^*)}$$

Now let $\tilde{K}_i^N$ and $\tilde{K}_i^*$ be the Nash Equilibrium solution and the system optimal solution with respect to the same problem but with the modified linear cost functions such that $\bar{f}_i = \alpha_i \cdot K_i$ where $\alpha_i = f_i'(K_i^N)$. Correspondingly, $\tilde{L}^N = \sum_{i=1}^n \tilde{K}_i^N$ and $\tilde{L}^* = \sum_{i=1}^n \tilde{K}_i^*$.

Since $\tilde{K}_i^N = K_i^N$ (having the same set of first-order conditions), we have

$$p\mathbb{E}[\min(L^N, D)] - \sum_{i=1}^{n} f_i'(K_i^N)(K_i^N) = p\mathbb{E}[\min(\tilde{L}^N, D)] - \sum_{i=1}^{n} \alpha_i \tilde{K}_i^N.$$

Because $\tilde{K}_i^*$ is the optimal capacity investment level for the modified problem, it implies that

$$p\mathbb{E}[\min(L^*, D)] - \sum_{i=1}^{n} f_i'(K_i^N)(K_i^*) \leq p\mathbb{E}[\min(\tilde{L}^*, D)] - \sum_{i=1}^{n} \alpha_i \tilde{K}_i^*.$$

Thus, we have

$$POA(f) \geq \frac{p\mathbb{E}[\min(\tilde{L}^N, D)] - \sum_{i=1}^{n} \alpha_i \tilde{K}_i^N}{p\mathbb{E}[\min(\tilde{L}^*, D)] - \sum_{i=1}^{n} \alpha_i \tilde{K}_i^*} \geq \frac{\pi^N(\bar{f})}{\pi^*(\bar{f})} = POA(\bar{f}).$$

This completes the proof. $\square$

By making use of Proposition 4.4.7, we can obtain a lower bound on the profit by using the cost asymmetry factor and the ratio between the investment levels in the Nash and the system optimum.

Lemma 4.4.8 *Price of anarchy on the total profit of a joint venture is lowered bounded by*

$$POA(\bar{f}) = \frac{\pi^N(\bar{f})}{\pi^*(\bar{f})} \geq \tilde{\alpha}\frac{\tilde{L}^N}{\tilde{L}^*},$$

where the cost asymmetry factor is given by

$$\tilde{\alpha} = \frac{\min_i \alpha_i}{\max_i \alpha_i} \leq 1.$$

Proof of Lemma 4.4.8. Assume that, without loss of generality, $\alpha_m = \alpha_1 \leq \alpha_2 \leq \ldots \leq \alpha_n = \alpha_M$. Define the set $P = \{j \mid \alpha_j = \alpha_m\}$. If $|P| = s$, s symmetric players

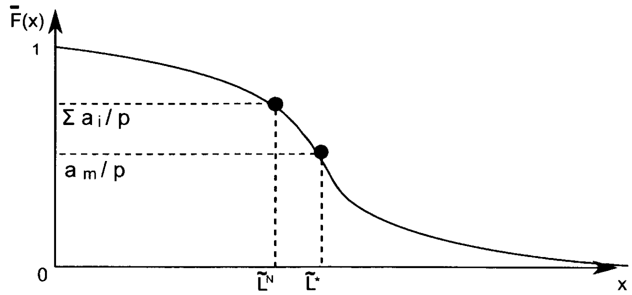

Figure 4-2: A graphical proof for Lemma 4.4.8.

invest in the system optimal solution and therefore $\tilde{L}^* = s\tilde{K}_j^*$ for $i \in P$.

$$
\begin{aligned}
POA(\bar{f}) &= \frac{p\int_0^{\tilde{L}^N} \bar{F}_D(x)dx - \sum_{i=1}^N \alpha_i \tilde{K}_i^N}{p\int_0^{\tilde{L}^N} \bar{F}_D(x)dx + p\int_{\tilde{L}^N}^{\tilde{L}^*} \bar{F}_D(x)dx - \alpha_m \tilde{L}^N} \\
&\geq \frac{\sum_{i=1}^N \alpha_i \tilde{L}^N - \sum_{i=1}^N \alpha_i \tilde{K}_i^N}{\sum_{i=1}^N \alpha_i \tilde{L}^N + \sum_{i=1}^N \alpha_i (\tilde{L}^* - \tilde{L}^N) - \alpha_m \tilde{L}^*} \\
&\geq \frac{\sum_{i=1}^N \alpha_i (\tilde{L}^N - \tilde{K}_i^N)}{\sum_{i=1}^N \alpha_i \tilde{L}^* - \alpha_m \tilde{L}^*} \\
&\geq \frac{\alpha_m (n-1)\tilde{L}^N}{\alpha_M (n-1)\tilde{L}^*} \geq \tilde{\alpha}\frac{\tilde{L}^N}{\tilde{L}^*}.
\end{aligned}
$$

where the cost asymmetry factor $\tilde{\alpha} = \alpha_m/\alpha_M \leq 1$. This completes the proof. □

Note that equal revenue sharing induces equal marginal costs for every player in a Nash equilibrium, since $\beta_i = \alpha_i / \sum_{j=1}^n \alpha_j$. Therefore, $\tilde{\alpha} = 1$, and the comparison between the profit can be reduced to a comparison between the total investment level, i.e, $\frac{\pi^N(\bar{f})}{\pi^*(\bar{f})} \geq \frac{\tilde{L}^N}{\tilde{L}^*}$

Next, we will present the how the profit in a joint venture can be bounded from below by the system optimum. Define the demand spread

$$\tilde{\theta} = \frac{\theta_m}{\theta_M} = \frac{\max f_D(x)}{\min f_D(y)},$$

where $x \leq \tilde{L}^N \leq y \leq \tilde{L}^*$.

Theorem 4.4.9

$$POA(f) \geq \tilde{\alpha}\frac{1-n\bar{r}}{1-n\bar{r}+(n-1)\bar{r}\tilde{\theta}},$$

where $\bar{r} = \max_i \alpha_i/p$, *and* $\tilde{\theta} \geq 1$ *measures the demand spread.*

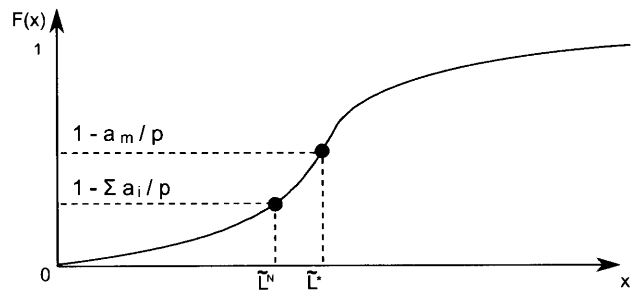

Figure 4-3: A graphical proof for Theorem 4.4.9.

Proof of Theorem 4.4.9. First we lower bound the ratio of $\tilde{L}^N$ to $\tilde{L}^*$.

$$\begin{aligned}
\frac{\tilde{L}^N}{\tilde{L}^*} &\geq \frac{\tilde{L}^N}{\tilde{L}^N + (\sum_{i=1}^n \alpha_i - \alpha_m)/(\theta_m p)} \\
&\geq \frac{(1-\sum_{i=1}^n \alpha_i/p)/\theta_M}{(1-\sum_{i=1}^n \alpha_i/p)/\theta_M + (\sum_{i=1}^n \alpha_i - \alpha_m)/(\theta_m p)} \\
&= \frac{p-\sum_{i=1}^n \alpha_i}{p-\sum_{i=1}^n \alpha_i + (\sum_{i=1}^n \alpha_i - \alpha_m)\tilde{\theta}} \\
&\geq \frac{p-n\alpha_M}{p-n\alpha_M+(n-1)\alpha_M\tilde{\theta}} \\
&= \frac{1-n\bar{r}}{1-n\bar{r}+(n-1)\bar{r}\tilde{\theta}},
\end{aligned}$$

where $\bar{r} = \alpha_M/p$. This result then follows from Lemme 4.4.8. □

Note that when D is uniform, the demand spread $\tilde{\theta} = 1$, we have

$$POA(f) \geq \tilde{\alpha}\frac{1-n\bar{r}}{1-\bar{r}}.$$

Figure 4-4, 4-5 and 4-6 show the lower bounds on POA with uniform demand, normal demand $N(400, 100)$ and exponential demand $\exp(400)$, respectively. The lower bound on POA decreases as the number of players increases or the marginal cost to price ratio increases. We also observe that the lower bound on POA has a steeper rate of decrease when the demand spead is higher. Note that in our simulation, the exponential demand has the highest demand spread ($\tilde{\theta} = 7.35$), followed by the normal demand ($\tilde{\theta} = 3.86$) and then the uniform demand ($\tilde{\theta} = 1$).

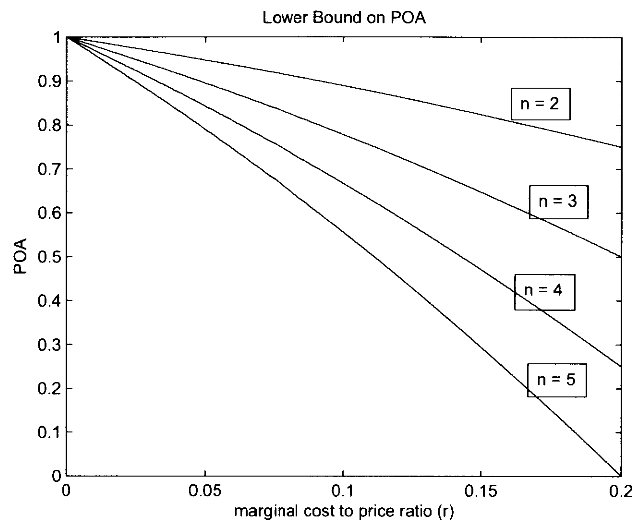

Figure 4-4: Lower Bound on Price of Anarchy for Uniform Demand.

4.5 Conclusion

In this work, we study resource pooling and capacity planning in joint ventures under uncertainties. We distinguish two types of resources pooling, based on whether the resources are heterogeneous or homogeneous. When resources are heterogeneous, the effective capacity in a joint venture is constrained by the lowest level of contribution from one participant. We have shown that every participant is committed to make an equal contribution in a joint venture with heterogeneous resources. We have also shown that, there exists a same efficient and fair revenue sharing scheme in both Nash equilibrium and Nash Bargaining solution. The optimal scheme rewards every participant proportionally to his marginal cost. When resources are homogeneous, however, there does not exist a revenue sharing scheme which induces actions to

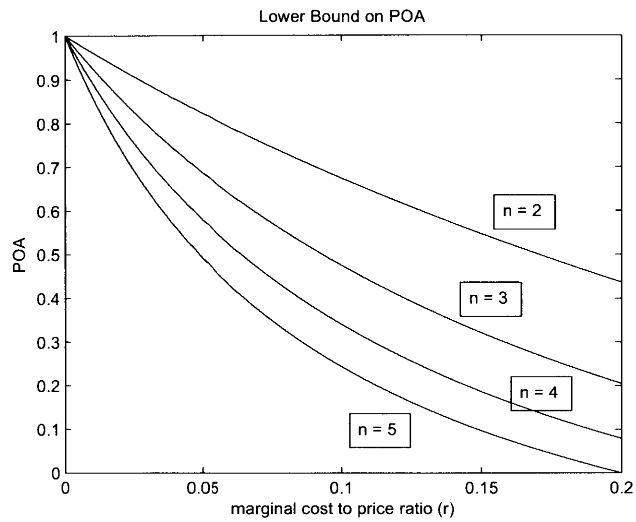

Figure 4-5: Lower Bound on Price of Anarchy for Normal Demand.

achieve the optimum. Nonetheless, we propose some methods to share revenue with the worst case performance guarantee. The methods suggest that the reward should be inversely proportional to the marginal cost of each participant with homogeneous resources.

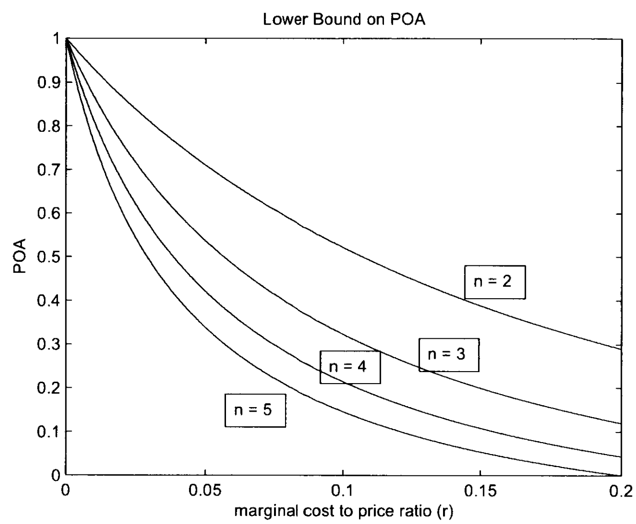

Figure 4-6: Lower Bound on Price of Anarchy for Exponential Demand.

Chapter 5

Conclusions

In this thesis, we have surveyed several recent work to develop provably near-optimal approximation algorithms for operations management models. We would like to summarize our results and point out possible future research directions.

Stochastic Lot-sizing Problems

We have developed new algorithmic approaches to compute provably near-optimal policies for multi-period stochastic lot-sizing inventory models with positive lead times, general demand distributions and dynamic forecast updates. The goal is to coordinate a sequence of orders of a single commodity, aiming to supply stochastic demands over a discrete finite horizon with minimum total expected cost, including fixed ordering, holding and backlogging costs. The policies that are developed have worst-case performance guarantees of 3 and typically perform very close to optimal in extensive computational experiments. We also propose a 6-approximation algorithm for the counterpart model under uniform capacity constraints.

We believe that these ideas will be effective to develop new near-optimal algorithms to various core stochastic multi-echelon and multi-item inventory control models. In particular, we attempt to tackle the stochastic joint-replenishment problem. In the joint-replenishment problem, we have a cross-docking warehouse that is not allowed to hold any inventory, and multiple retailers each facing stochastic demands.

A major set-up cost is incurred whenever the warehouse places an order from some exogenous supplier, and a minor set-up cost is incurred whenever an order is shipped to a retailer. Any order made by the warehouse has to be distributed among retailers immediately since the warehouse is not allowed to hold any inventory. The goal is to coordinate a sequence of orders over a discrete finite horizon to minimize the system-wide expected cost, including set-up, ordering, holding and backlogging costs.

Another important future research direction is to study the performance of dual-balancing or, more generally, cost-balancing policies under various assumptions on the underlying demand distributions. As much as it is powerful to establish general worst-case analysis, it is equally important to refine this analysis to various parametric regimes of the underlying demand distributions and other key parameters of the problem. We call this *parametric worst-case analysis.*

Revenue Management of Reusable Resources with Advanced Reservations

We have studied a class of revenue management problems in systems with *reusable resources* and *advanced reservations.* A simple control policy called the *class selection policy* (CSP) is proposed based on solving a knapsack-type linear program (LP). We show that the CSP and its variants perform provably near-optimal in the *Halfin-Whitt* regime. The analysis is based on modeling the problem as *loss network systems* with *advanced reservations.* In particular, asymptotic upper bounds on the blocking probabilities are derived.

There are several issues that still remain open. From the comparison of the upper bounds and the simulation results, it is clear that there is a gap between the empirical blocking probabilities and the theoretical bounds. This gap is due to the approximation using infinite capacity systems. It opens an opportunity to tighten the upper bound using another fictitious system between the original capacitated system and the infinite capacity counterpart.

We believe that these ideas can be effective to analyze dynamic pricing models with advanced reservation. Consider a single-class time-homogeneous arrival process of customers. Depending on the current state of the booking profile, the system offers each arriving customer a price from a fixed menu. The customer accepts the offer only if the price offered falls below her reservation price. A similar knapsack-type linear program can be used to guide the system to dynamically decide what prices to offer.

Joint-ventures in Operations Management

We have examined the problem of capacity planning in joint ventures to meet stochastic demand in a *newsvendor-type* setting. When resources are heterogeneous, there exists a unique revenue-sharing contract such that the corresponding Nash Bargaining Solution, the Strong Nash Equilibrium, and the system optimal solution coincide. The optimal scheme rewards every participant proportionally to her marginal cost. When resources are homogeneous, there does not exist a revenue-sharing scheme which induces the system optimum. Nonetheless, we propose provably good revenue-sharing contracts which suggests that the reward should be inversely proportional to the marginal cost of each participant.

We will explore different cost structures and generalize the results as much as possible. A viable approach is to use piecewise linear or quadratic functions to approximate any convex cost functions. We are also interested in *mechanism design* that aligns the incentives of both parties to achieve system optimum profit. Another important element of this study is test the validity and quality of our models using empirical data.

Appendix A

Appendix for Chapter 2

A.1 Proofs of Technical Lemmas and Theorems

LEMMA 2.4.1. *Let $\mathscr{C}(RB)$ be the total cost incurred by the RB policy. Then we have,*

$$E[\mathscr{C}(RB)] \leq 3 \cdot \sum_{t=1}^{T-L} E[Z_t^{RB}]. \tag{A.1}$$

Proof of Lemma 2.4.1. Using the marginal cost accounting in Equation (2.6) and standard arguments of conditional expectations, we express

$$\begin{aligned} E[\mathscr{C}(RB)] &= \sum_{t=1}^{T-L} E[H_t^{RB}(Q_t^{RB}) + \Pi_t^{RB}(Q_t^{RB}) + K \cdot \mathbb{1}(Q_t^{RB} > 0)] \qquad \text{(A.2)} \\ &= \sum_{t=1}^{T-L} E\left[E[H_t^{RB}(Q_t^{RB}) + \Pi_t^{RB}(Q_t^{RB}) + K \cdot \mathbb{1}(Q_t^{RB} > 0) \mid F_t]\right] \\ &= \sum_{t=1}^{T-L} E[2Z_t^{RB} + P_t K] \leq 3 \sum_{t=1}^{T-L} E[Z_t^{RB}]. \end{aligned}$$

The third equality follows directly from (2.13). To establish the first inequality in (A.2) above, we shall show that $Z_t \geq P_t K$ almost surely. That is, for each $f_t \in F_t$, $z_t \geq p_t K$. Given any information set f_t, all the quantities x_t, θ_t, ψ_t, ϕ_t and p_t defined above are known deterministically. We split the analysis into two cases:

1. If $\theta_t \geq K$, then $q_t^{RB} = \hat{q}_t$ (the *balancing quantity*) with probability $p_t = 1$

implying $z_t = \theta_t \geq K$. The claim follows.

2. If $\theta_t < K$, then $q_t^{RB} = \tilde{q}_t$ (the *holding-cost-K quantity*) with probability p_t and $q_t^{RB} = 0$ with $1 - p_t$. Thus, by Equations (2.10) and (2.11), we have $z_t = p_t K$, and the claim follows.

This concludes the proof of the lemma. □

LEMMA 2.4.2. *The overall holding cost and backlogging cost incurred by OPT are denoted by H^{OPT} and Π^{OPT}, respectively. Then we have, with probability* 1,

$$H^{OPT} \geq \sum_t H_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H}), \ \Pi^{OPT} \geq \sum_t \Pi_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi}). \quad (A.3)$$

Proof of Lemma 2.4.2. The proof is identical to Lemmas 4.2 and 4.3 in Levi et al. (2007). □

LEMMA 2.4.3. *The expected holding cost and backlogging cost incurred by OPT plus the expected amount borrowed from the bank account A are at least $\sum_{t=1}^{T-L} E[Z_t^{RB}]$. That is, The following inequality holds*

$$E\left[\left(H^{OPT} + \Pi^{OPT}\right) + A\right] \geq \sum_{t=1}^{T-L} E[Z_t^{RB}]. \quad (A.4)$$

Proof of Lemma 2.4.3. Using linearity of expectation, it suffices to show

$$E\left[H^{OPT} + \Pi^{OPT}\right] \geq \sum_{t=1}^{T-L} E\left[\mathbb{1}(t \in \mathscr{T}_N) \cdot Z_t^{RB}\right]. \quad (A.5)$$

Using Lemma 2.4.2 and standard arguments of condition expectations, we have

$$\begin{aligned} E[H^{OPT}] &\geq E\left[\sum_t H_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H})\right] \quad (A.6) \\ &= E\left[E\left[\sum_t H_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H}) \mid F_t\right]\right] \\ &= E\left[\sum_t Z_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{1H} \bigcup \mathscr{T}_{2H})\right]. \end{aligned}$$

Similarly, we also have

$$E[\Pi^{OPT}] \;\geq\; E\left[\sum_t Z_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi})\right]. \tag{A.7}$$

Equation (A.5) follows from summing up Equations (A.6) and (A.7). □

LEMMA 2.4.4. *The following inequality holds*

$$E\left[A\right] \leq E\left[\sum_{t=1}^{T-L} K \cdot \mathbb{1}(Q_t^{OPT} > 0)\right]. \tag{A.8}$$

In other words, the expected borrowing $E[A]$ *is less than the total expected fixed ordering cost incurred by* OPT.

Proof of Lemma 2.4.4. First we define the *reduced information set* f_t^- to be the information up to period t *excluding* the randomized decisions of the RB policy over $[1, t-1]$. In particular, given the entire evolution of demand f_T^-, the sequence of orders placed by OPT is known deterministically. Let $1 \leq t_1 < t_2 < \ldots < t_n \leq T-L$ be the periods in which OPT placed $n = n \mid f_T^-$ orders sequentially. Let $t_0 = 0$ and $t_{n+1} = T - L + 1$. We shall show that there are no problematic periods within (t_0, t_1) and that, for each $i = 1, \ldots n$, the expected borrowing within the interval $[t_i, t_{i+1})$ does not exceed K. That is,

$$(t_0, t_1) \bigcap \mathscr{T}_{2M} \;=\; \emptyset, \tag{A.9}$$

$$E\left[\sum_{t \in [t_i, t_{i+1}) \bigcap \mathscr{T}_{2M}} Z_t^{RB} \mid f_T^-\right] \;\leq\; K. \tag{A.10}$$

It is important to note that f_T^- does not include the randomized decisions of the RB policy. Thus, the set $\mathscr{T}_{2M}$ is still random and so is the amount borrowed from the bank. In particular, the expectation in Equation (A.10) is taken with respect to the randomized decisions of the RB policy. Equations (A.10) and (A.9) imply that, for each f_T^-,

$$E\left[\sum_{t\in\mathscr{T}_{2M}} Z_t^{RB} \mid f_T^-\right] \leq K\cdot n \mid f_T^- = K\cdot n, \tag{A.11}$$

and therefore

$$E[\Lambda] \leq K\cdot E[N] = E\left[\sum_{t=1}^{T-L} K\cdot \mathbb{1}(Q_t^{OPT}>0)\right]. \tag{A.12}$$

Thus, it suffices to prove Equations (A.10) and (A.9). Figure A-1 gives a graphical interpretation of Equation (A.10), i.e., we want to show that the fixed ordering cost K incurred by OPT in period t_i will cover the expected amount borrowed from the bank in periods that belong to set $\mathscr{T}_{2M}$ within the interval $[t_i, t_{i+1})$.

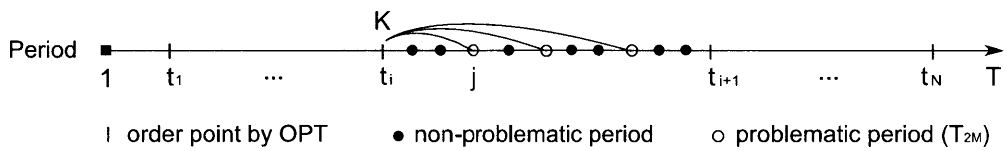

Figure A-1: Decomposition of the problematic periods in the set $\mathscr{T}_{2M}$ into intervals between ordering points of OPT

Proof of Equation (A.9). We first show that Equation (A.9) holds. Recall the definition $\mathscr{T}_{2M} = \left\{t : \Theta_t < K \text{ and } X_t^{RB} < Y_t^{OPT} \leq X_t^{RB} + \tilde{Q}_t^{RB}\right\}$. Since at the beginning of the planning horizon, it is assumed that every feasible policy will have the same initial inventory position, it follows that if period t is in $\mathscr{T}_{2M}$, OPT must have placed an order and overtaken the inventory position of the RB policy. (The two policies face the same sequence of demands.) However, (t_0, t_1) denotes the set of periods in which OPT has not placed any order yet. Thus, the intersection of these two sets is empty.

Proof of Equation (A.10). Next we show that Equation (A.10) holds. Recall that f_T^- denotes an entire evolution of the system excluding the randomized decisions of the RB policy. Given the entire evolution of demands f_T^-, construct a decision tree based on the randomized decisions of the RB policy. The root node corresponding to period 1 contains the information set $f_1 = f_1^- \in f_T^-$. The tree is built in layers, each corresponding to a period, where the number of nodes in layer t is 2^{t-1} numbered $l = 1, \ldots, 2^{t-1}$. In particular, a node l in period (layer) t corresponds to some information set $f_t \in \mathscr{F}_t$ which includes the realized reduced information set $f_t^- \subseteq f_T^-$, and the realized randomized decisions up to period $t-1$ of the RB policy. Therefore it is known whether under this state period t belongs to the set $\mathscr{T}_{2M}$ or not.

The edges in the tree represent the different (randomized) decisions that the RB policy may make with their respective probabilities. Each path from the root to a specific node corresponds to a sequence of realized randomized ordering decisions made by the RB policy. For example, consider again some node l in period (layer) t in which the RB policy will order $\tilde{q}_{tl}^{RB}$ units with probability p_{tl} and nothing with probability $1 - p_{tl}$; then the node l in period t (denoted by tl) will have two edges to two children nodes in the next period $t+1$ each containing its distinctive ordering information. Conceptually one can think about the decision tree as a collection of independent coins, each corresponding to a node in the tree. The coin corresponding to node l at layer (period) t has probability of success (ordering) p_{tl}.

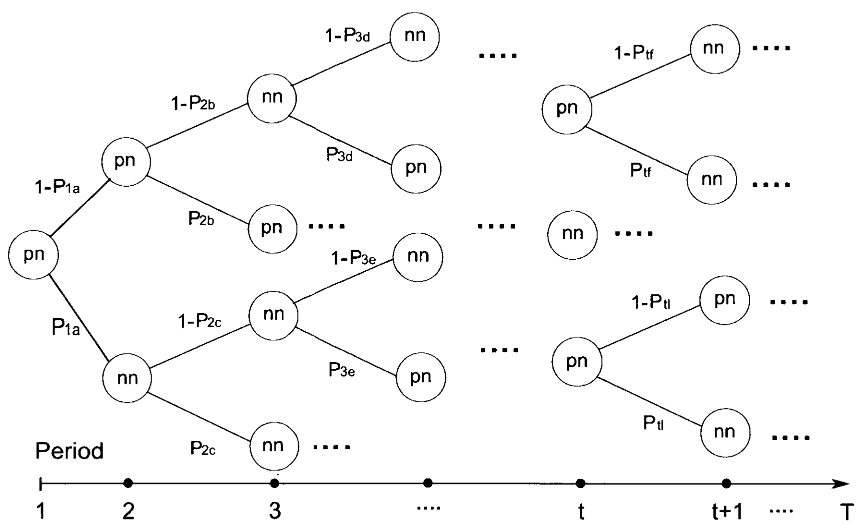

Figure A-2: An example of a general decision tree

Next we partition the nodes in the tree into *problematic* nodes (pn nodes), i.e., nodes that correspond to a pair (t, f_t) for which $t \in \mathscr{T}_{2M}$, and *non-problematic* nodes (nn nodes). An example of a general decision tree is illustrated in Figure A-2.

Focus now on a specific time interval $[t_i, t_{i+1})$. Suppose we have constructed the tree from period 1 to T; the number of nodes and paths are clearly finite (possibly exponential). Let the set $\mathscr{G}$ to be the set of all possible outcomes of the randomized decisions in all nodes in layers within the interval $[1, t_i - 1]$ and in all the *nn* nodes within the interval $[1, T]$. In particular, each $g \in \mathscr{G}$ corresponds to a specific set of outcomes in all nodes in layers (periods) within the interval $[1, t_i - 1]$ and in all the *nn* nodes in the tree. Using the terminology of coins proposed before, g corresponds to the outcome of the respective subset of coins corresponding to all nodes within $[1, t_i - 1]$ and all *nn* nodes within $[1, T]$.

Conditioning on some $g \in \mathscr{G}$ induces a path from the root of the tree (in period 1) up to the earliest *pn* node, say j, where j corresponds to the period (layer) of that node. Here we abuse the notation ignoring the index of the node within layer j. (Namely, the exact value will be je for some e.) It is straightforward to see that $j \geq t_i$. If j falls outside the interval $[t_i, t_{i+1})$, i.e., $j \geq t_{i+1}$, it follows that there are no *pn* nodes within the interval $[t_i, t_{i+1})$, and there is no borrowing over the interval. Assume now that j falls within the interval $[t_i, t_{i+1})$ (j can possibly be in period (layer) t_i). We will show that the expected borrowing does not exceed K. That is,

$$E\left[\sum_{s \in [j, t_{i+1}) \bigcup \mathscr{T}_{2M}} Z_s^{RB} \mid f_T^-, g\right] \leq K. \tag{A.13}$$

The proof of Equation (A.10) will then follow.

Recall that node j corresponds to some information set $f_j \in \mathscr{F}_j$. It follows that the starting inventory position x_j^{RB} and the corresponding *holding-cost-K* quantity $\tilde{q}_j^{RB}$ are known deterministically. Conditioning on g, the only uncertainty in the evolution of the system depends on the randomized decisions made in *pn* nodes within $[j, t_{i+1})$. Consider the sub-tree induced by conditioning on g. The non-problematic nodes (*nn* nodes) in the sub-tree have only one outgoing edge that corresponds to the decision (order/no-order) specified by g to that node. The problematic nodes (*pn* nodes) have two outgoing edges corresponding to the order/no-order decisions, respectively. (Recall that g does not specify the decisions in these nodes.) Moreover, each *pn* node

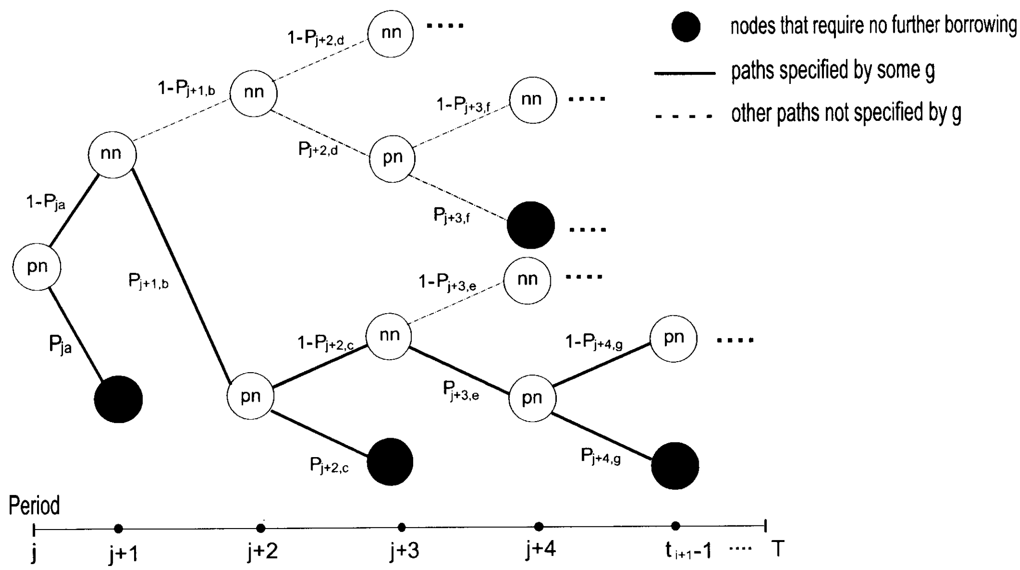

Figure A-3: An example of a decision subtree: focus on the interval $[t_i, t_{i+1})$ and some $g \in \mathscr{G}$, j is the earliest period in which a problematic node (pn) occurs. According to g, there are two possible outcomes whenever a problematic node (pn) is reached, and there is only one possible outcome whenever a non-problematic node (nn) is reached. If a problematic node (pn) orders, there will not be further borrowing until the next order of OPT in period t_{i+1}.

$s \in [j, t_{i+1})$ is associated with the probability p_s of ordering. (We again abuse the notation introduced before and omit the index e of the node within the layer/period.) An example of a decision subtree specified by some $g \in \mathscr{G}$ is illustrated in Figure A-3. Any sequence of randomized outcomes corresponding to the decisions in the pn nodes induces a path of evolution of the system. The resulting cumulative borrowing from the bank account A, corresponding to this path, is equal to K times the sum of probabilities associated with the pn nodes in this path. (For each pn node s in the path, the borrowing is equal to $p_s K = z_s$.)

Next we claim that the sub-tree defined above includes at most one pn node in each layer (period). This follows from the fact that any path between two pn nodes r, s such that $j \leq r < s < t_{i+1}$ in the tree includes only no-ordering edges of pn nodes. To see why the latter is true, observe that if an order is placed by the RB policy in a pn node, the resulting inventory position of the RB policy is higher than OPT. Since both policies face the same sequence of demands, the RB policy will not have higher inventory position than OPT at least until the next order placed by OPT. This excludes the existence of pn nodes in subsequent periods until OPT places another order, i.e., beyond period $t_{i+1}-1$.

In light of the latter observation, we re-number all the pn nodes in the sub-tree as $1, 2, \ldots, M$ (where 1 corresponds to j, specified before). Moreover, it follows that the probability to arrive at node $m = 1, \ldots, M$ and borrow $p_m K$ is equal to $\prod_{s=1}^{m-1}(1-p_s)$. (This probability corresponds to no-ordering decisions in all the pn nodes prior to m.) The total expected borrowing is then

$$K \cdot \left\{ p_1^2 + \sum_{m=2}^{M} \left\{ \left(\prod_{s=1}^{m-1} (1-p_s) \right) p_m \left(\sum_{k=1}^{m} p_k \right) \right\} \right\}. \tag{A.14}$$

Observe that the probability to borrow exactly $K \cdot \sum_{k=1}^{m} p_k$ is equal to

$$\left(\prod_{s=1}^{m-1} (1-p_s) \right) p_m. \tag{A.15}$$

Moreover, we have already shown that the expression in (A.14) is bounded above by K (see Lemma 2.4.5). This concludes the proof of the lemma. □

LEMMA 2.4.5. *Let $\{p_l\}_{l=1}^{\infty}$ satisfy the condition $0 \le p_l \le 1$ for all l. Then the following inequality holds,*

$$p_1^2 + \sum_{l=2}^{\infty} \left\{ \left(\prod_{s=1}^{l-1} (1-p_s) \right) p_l \left(\sum_{k=1}^{l} p_k \right) \right\} \le 1. \tag{A.16}$$

Proof of Lemma 2.4.5. We construct an increasing sequence $\{a_m\}$ where

$$a_m = p_1^2 + \sum_{l=2}^{m} \left\{ \left(\prod_{s=1}^{l-1} (1-p_s) \right) p_l \left(\sum_{k=1}^{l} p_k \right) \right\}. \tag{A.17}$$

For each m, if we replace p_m by 1, we get

$$\bar{a}_m = p_1^2 + \sum_{l=2}^{m-1} \left\{ \left(\prod_{s=1}^{l-1} (1-p_s) \right) p_l \left(\sum_{k=1}^{l} p_k \right) \right\} + \left(\prod_{s=1}^{m-1} (1-p_s) \right) \left(1 + \sum_{k=1}^{m-1} p_k \right), \tag{A.18}$$

such that $a_m \leq \bar{a}_m$. Next we will show by induction that $\bar{a}_m \leq 1$ for all m from which the proof of the lemma follows. It is straightforward to verify $\bar{a}_1, \bar{a}_2 \leq 1$. Assume that $\bar{a}_m \leq 1$ for some $m \in \mathbb{Z}^+$, we will show that $\bar{a}_{m+1} \leq 1$.

$$\begin{aligned}
\bar{a}_{m+1} &= p_1^2 + \sum_{l=2}^{m} \left\{ \left(\prod_{s=1}^{l-1} (1-p_s) \right) p_l \left(\sum_{k=1}^{l} p_k \right) \right\} + \left(\prod_{s=1}^{m} (1-p_s) \right) \left(1 + \sum_{k=1}^{m} p_k \right) \\
&= a_{m-1} + \left(\prod_{s=1}^{m-1} (1-p_s) \right) p_m \left(\sum_{k=1}^{m} p_k \right) + \left(\prod_{s=1}^{m} (1-p_s) \right) \left(1 + \sum_{k=1}^{m} p_k \right) \\
&= a_{m-1} + \left(\prod_{s=1}^{m-1} (1-p_s) \right) \left[\left(1 + \sum_{k=1}^{m} p_k \right) (1-p_m) + p_m \sum_{k=1}^{m} p_k \right] \\
&= a_{m-1} + \left(\prod_{s=1}^{m-1} (1-p_s) \right) \left(1 + \sum_{k=1}^{m-1} p_k \right) = \bar{a}_m \leq 1.
\end{aligned}$$

Hence the claim follows by induction. $\square$

LEMMA 2.6.1. *Let $\mathscr{C}(RB)$ be the total cost incurred by the RB policy. Then we have,*

$$E[\mathscr{C}(RB)] \leq 3 \cdot \sum_{t=1}^{T-L} E[Z_t^{RB}]. \tag{A.19}$$

Proof of Lemma 2.6.1. Using the standard arguments of conditional expectations, we express

$$\begin{aligned}
E[\mathscr{C}(RB)] &= \sum_{t=1}^{T-L} E[H_t^{RB}(Q_t^{RB}) + \bar{\Pi}_t^{RB}(Q_t^{RB}) + K \cdot \mathbb{1}(Q_t^{RB} > 0)] \qquad (A.20) \\
&= \sum_{t=1}^{T-L} E\left[E[H_t^{RB}(Q_t^{RB}) + \bar{\Pi}_t^{RB}(Q_t^{RB}) + K \cdot \mathbb{1}(Q_t^{RB} > 0) \mid F_t] \right] \\
&\leq \sum_{t=1}^{T-L} E[2Z_t^{RB} + P_t K] \\
&\leq 3 \sum_{t=1}^{T-L} E[Z_t^{RB}].
\end{aligned}$$

To establish the two inequalities in (A.20), we shall show that $Z_t^{RB} \geq E[H_t^{RB}(Q_t^{RB}) \mid F_t]$, $Z_t^{RB} = E[\bar{\Pi}_t^{RB}(Q_t^{RB}) \mid F_t]$ and $Z_t^{RB} \geq P_t K$ almost surely. Given any information set f_t, we know the inventory level x_t and all the quantities θ_t, ψ_t, ϕ_t, p_t defined above are also known deterministically. We split the analysis into two cases:

1. If $\theta_t \geq K$, then $q_t^{RB} = \hat{q}_t$ (the *balancing quantity*) with probability $p_t = 1$ implying $z_t^{RB} = \theta_t \geq K$. In addition, we have $z_t^{RB} = E[H_t^{RB}(\hat{q}_t) \mid f_t] = E[\bar{\Pi}_t^{RB}(\hat{q}_t) \mid f_t]$. The claim follows.

2. If $\theta_t < K$, then $q_t^{RB} = \min\{\tilde{q}_t, u\}$ with probability p_t and $q_t^{RB} = 0$ with $1 - p_t$. Thus, by the construction of the probability p_t, we have $z_t^{RB} = p_t K = E[\bar{\Pi}_t^{RB}(q_t^{RB}) \mid f_t]$ and $z_t^{RB} = p_t K = E[H_t^{RB}(\tilde{q}_t^{RB}) \mid f_t] \geq E[H_t^{RB}(q_t^{RB}) \mid f_t]$, and the claim follows.

This completes the proof of the lemma. $\square$

LEMMA 2.6.2. *The overall holding cost and backlogging cost incurred by OPT are denoted by H^{OPT} and Π^{OPT} respectively. Then we have*

$$
\begin{aligned}
E[H^{OPT}] &\geq E\left[\sum_t Z_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{1H})\right] \\
E[\Pi^{OPT}] &\geq E\left[\sum_t Z_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{1\Pi} \bigcup \mathscr{T}_{2\Pi})\right] \\
E\left[\sum_t K \cdot \mathbb{1}(Q_t^{OPT} > 0)\right] &\geq E\left[\sum_t Z_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{2M})\right] \\
E\left[H^{OPT} + \sum_t K \cdot \mathbb{1}(Q_t^{OPT} > 0)\right] &\geq E\left[\sum_t Z_t^{RB} \cdot \mathbb{1}(t \in \mathscr{T}_{2H})\right].
\end{aligned}
$$

Proof of Lemma 2.6.2. The proof of (2.42) and (2.43) is identical to Lemma 4.2 in Levi et al. (2007) and Lemma 2 in Levi et al. (2008d) respectively. The proof of (2.44) is identical to Lemma 5 in Levi and Shi (2009). Next we shall show that (2.45) holds true. Recall that

$$\mathscr{T}_{2H} = \left\{t : \Theta_t < K \text{ and } Y_t^{OPT} > X_t^{RB} + \min\{\tilde{Q}_t^{RB}, u\}\right\}. \tag{A.21}$$

In other words, $\mathscr{T}_{2H}$ consists of periods in which the balancing cost is less than K and the inventory position of OPT after ordering exceeds even $X_t^{RB} + \min\{\tilde{Q}_t^{RB}, u\}$. We split the analysis into two cases.

1. If $\tilde{Q}_t^{RB} \leq u$, the RB policy will order the *holding-cost-K quantity* $\tilde{Q}_t^{RB}$ incurring exactly K expected marginal holding cost. Since OPT has more physical inventory than the RB policy, OPT had already ordered these units before thus incurring more holding cost.

2. On the other hand, if $\tilde{Q}_t^{RB} > u$, the RB policy will order the capacity u incurring less than K marginal holding cost. We shall show that the *fixed ordering costs* incurred by OPT cover this cost. It suffices to show that the number of orders placed by OPT over the interval $[1, t]$ is at least the number of orders in which RB orders up to capacity u over $[1, t]$. We prove the claim by contradiction. Suppose otherwise, the number of orders placed by OPT over the interval $[1, t]$ is m and the number of orders in which RB orders up to capacity u over $[1, t]$ is n and $m < n$. The maximum inventory position of OPT in period t is $x_1 + m \cdot u$, whereas the minimum inventory position of the RB policy in period t is $x_1 + n \cdot u$. This contradicts to the fact that OPT has higher inventory position than the RB policy in period t where $t \in \mathscr{T}_{2H}$. Hence the claim holds.

This completes the proof of the lemma. □

A.2 Performance of the proposed algorithms

The first two columns specify the test instances, namely, fixed ordering cost K, per-unit holding cost h, per-unit backlogging cost p and demand rate vector λ. The third column shows the cost incurred by the optimal policy. The fourth column shows the optimal parameters of parametrized RB policy. The fifth column shows the cost incurred by the parameterized RB policy. The sixth column shows the cost ratio of the parameterized RB policy to the optimal policy. The seventh column shows the cost of unparameterized RB policy (i.e., the original policy without parameter optimization). The eighth columns shows the cost ratio of the unparameterized RB policy to the optimal policy.

(K,h,p)	Demands $(\lambda_0, \lambda_1, \lambda_2)$	Cost of OPT	Optimal $(\beta^*, \gamma^*, \eta^*)$	Cost of param. RB	Cost Ratio	Cost of unparam. RB	Cost Ratio
(0,1,9)	(4,1,4)	46.85	(*,2,*)	49.18	1.0497	58.30	1.2444
(0,1,9)	(4,1,2)	46.39	(*,2,*)	49.30	1.0627	55.24	1.1908
(0,1,9)	(4,1,1)	46.20	(*,2,*)	47.81	1.0348	54.26	1.1745
(0,1,9)	(3,1,2)	41.02	(*,2,*)	41.41	1.0095	49.40	1.2043
(0,1,9)	(2,1,3)	32.88	(*,2,*)	34.42	1.0468	41.51	1.2625
(0,1,9)	(1,1,4)	24.74	(*,2,*)	26.40	1.0671	31.40	1.2692
(5,1,9)	(4,1,1)	102.66	(0.2,2,9)	108.28	1.0547	135.37	1.3186
(5,1,9)	(1,1,4)	86.47	(0.2,2,9)	90.70	1.0489	128.70	1.4884
(5,1,1)	(4,1,1)	71.35	(0.4,1,1)	75.42	1.0570	84.13	1.1791
(100,1,9)	(5,1,0)	427.81	(0.9,*,9)	451.68	1.0558	605.10	1.4144
(100,1,9)	(4,1,1)	424.81	(0.9,*,9)	449.65	1.0585	601.29	1.4154
(100,1,9)	(3,1,2)	421.76	(0.9,*,9)	443.12	1.0506	595.10	1.4110
(100,1,9)	(2,1,3)	418.63	(0.9,*,9)	443.64	1.0597	611.48	1.4607
(100,1,9)	(1,1,4)	415.49	(0.8,*,9)	437.36	1.0526	618.36	1.4883
(100,1,9)	(0,1,5)	412.29	(0.8,*,9)	435.65	1.0567	593.88	1.4404

Table A.1: Numerical results with lead time $L = 0$ and finite horizon $T = 12$.

(K,h,p)	Demands $(\lambda_0, \lambda_1, \lambda_2)$	Cost of OPT	Optimal $(\beta^*, \gamma^*, \eta^*)$	Cost of param. RB	Cost Ratio	Cost of unparam. RB	Cost Ratio
(0,1,9)	(4,1,4)	93.81	(*,2,*)	98.32	1.0481	120.14	1.2807
(0,1,9)	(4,1,2)	88.27	(*,2,*)	94.25	1.0677	108.24	1.2262
(0,1,9)	(4,1,1)	85.48	(*,2,*)	90.21	1.0553	93.97	1.0993
(0,1,9)	(3,1,2)	80.04	(*,2,*)	89.73	1.1211	90.40	1.1294
(0,1,9)	(2,1,3)	73.98	(*,1.5,*)	84.42	1.1411	90.99	1.2625
(0,1,9)	(1,1,4)	70.96	(*,1.5,*)	81.40	1.1471	87.60	1.2345
(5,1,9)	(4,1,1)	137.66	(0.2,2,9)	153.97	1.1185	161.10	1.1703
(5,1,9)	(1,1,4)	121.47	(0.2,2,9)	140.26	1.1525	148.47	1.2223
(5,1,1)	(4,1,1)	78.18	(0.4,1,1)	90.42	1.1566	97.47	1.2467
(100,1,9)	(5,1,0)	434.30	(0.9,*,9)	479.03	1.1030	614.17	1.4142
(100,1,9)	(4,1,1)	431.87	(0.9,*,9)	466.33	1.0798	611.96	1.4170
(100,1,9)	(3,1,2)	429.41	(0.9,*,9)	453.24	1.0555	551.00	1.2832
(100,1,9)	(2,1,3)	426.86	(0.9,*,9)	451.17	1.0570	644.13	1.5090
(100,1,9)	(1,1,4)	424.25	(0.9,*,9)	466.43	1.0994	623.56	1.4698
(100,1,9)	(0,1,5)	421.56	(0.9,*,9)	461.65	1.0951	595.40	1.4124

Table A.2: Numerical results with lead time $L = 2$ and finite horizon $T = 12$.

(K,h,p)	Demands $(\lambda_0, \lambda_1, \lambda_2)$	Cost of OPT	Optimal $(\beta^*, \gamma^*, \eta^*)$	Cost of param. RB	Cost Ratio	Cost of unparam. RB	Cost Ratio
(0,1,9)	(4,1,4)	57.71	(*,2,*)	58.23	1.0090	61.92	1.0730
(0,1,9)	(4,1,2)	57.71	(*,2,*)	58.36	1.0113	60.94	1.0560
(0,1,9)	(4,1,1)	57.71	(*,2,*)	58.30	1.0102	60.38	1.0463
(0,1,9)	(3,1,2)	50.19	(*,2,*)	51.49	1.0259	53.62	1.0683
(0,1,9)	(2,1,3)	41.27	(*,2,*)	41.96	1.0167	43.63	1.0572
(0,1,9)	(1,1,4)	30.55	(*,2,*)	30.88	1.0108	31.66	1.0363
(5,1,9)	(4,1,1)	128.17	(0.2,2,9)	133.91	1.0448	166.10	1.2959
(5,1,9)	(1,1,4)	101.70	(0.2,2,9)	107.34	1.0555	148.85	1.4636
(5,1,1)	(4,1,1)	86.07	(0.4,1,1)	90.51	1.0516	104.24	1.2111
(100,1,9)	(5,1,0)	535.14	(1.1,*,9)	566.23	1.0581	663.61	1.2401
(100,1,9)	(4,1,1)	533.51	(1.1,*,9)	570.65	1.0696	659.29	1.2358
(100,1,9)	(3,1,2)	529.77	(1.1,*,9)	566.09	1.0686	682.76	1.2888
(100,1,9)	(2,1,3)	523.94	(1.1,*,9)	555.57	1.0604	729.15	1.3917
(100,1,9)	(1,1,4)	520.03	(1.0,*,9)	550.36	1.0583	744.45	1.4316
(100,1,9)	(0,1,5)	516.05	(1.0,*,9)	550.65	1.0670	711.22	1.3782

Table A.3: Numerical results with lead time $L = 0$ and finite horizon $T = 15$.

No.	Demands $(\lambda_0, \lambda_1, \lambda_2)$	OPT	(β, γ, η)	RB	Ratio	OPT	(β, γ, η)	RB	Ratio
	Capacity		3				6		
1	(6,0,0)	1027.0	$(1, 1, 50)$	1141.5	1.111	275.9	$(1, 1, 10)$	326.1	1.182
2	(3,3,0)	840.9	$(1, 1, 50)$	914.9	1.088	207.2	$(1, 1, 10)$	251.8	1.215
3	(0,6,0)	681.9	$(1, 1, 50)$	756.9	1.111	180.5	$(1, 1, 10)$	237.9	1.318
4	(0,3,3)	552.2	$(1, 1, 50)$	584.4	1.058	169.9	$(1, 1, 10)$	218.2	1.284
5	(0,0,6)	446.7	$(1, 1, 50)$	474.3	1.061	154.6	$(1, 1, 10)$	205.1	1.327
	Capacity		9				12		
1	(6,0,0)	166.4	$(1, 1, 1)$	243.4	1.462	162.8	$(1, 1, 1)$	223.4	1.372
2	(3,3,0)	154.6	$(1, 1, 1)$	223.6	1.446	153.3	$(1, 1, 1)$	216.8	1.414
3	(0,6,0)	132.1	$(1, 1, 1)$	210.7	1.595	129.7	$(1, 1, 1)$	185.9	1.433
4	(0,3,3)	129.4	$(1, 1, 1)$	197.9	1.529	129.2	$(1, 1, 1)$	180.0	1.393
5	(0,0,6)	119.8	$(1, 1, 1)$	174.7	1.458	118.9	$(1, 1, 1)$	168.4	1.416
	Capacity		3				6		
6	(6,0,0)	1279.0	$(1, 1, 50)$	1606.1	1.256	526.7	$(1, 1, 10)$	625.0	1.186
7	(3,3,0)	1093.0	$(1, 1, 50)$	1224.8	1.121	450.3	$(1, 1, 10)$	537.1	1.192
8	(0,6,0)	934.5	$(1, 1, 50)$	1104.5	1.182	410.6	$(1, 1, 10)$	513.3	1.250
9	(0,3,3)	804.7	$(1, 1, 50)$	857.6	1.065	377.7	$(1, 1, 10)$	486.2	1.287
10	(0,0,6)	698.6	$(1, 1, 50)$	761.0	1.089	343.3	$(1, 1, 10)$	437.9	1.275
	Capacity		9				12		
6	(6,0,0)	351.9	$(1, 1, 10)$	474.6	1.349	305.8	$(1, 1, 10)$	404.2	1.321
7	(3,3,0)	327.3	$(1, 1, 10)$	403.1	1.231	286.6	$(1, 1, 10)$	392.3	1.368
8	(0,6,0)	287.3	$(1, 1, 10)$	389.0	1.354	250.5	$(1, 1, 10)$	321.3	1.282
9	(0,3,3)	273.2	$(1, 1, 10)$	367.2	1.344	241.3	$(1, 1, 10)$	311.1	1.289
10	(0,0,6)	249.6	$(1, 1, 10)$	406.0	1.626	220.0	$(1, 1, 10)$	284.4	1.292

Table A.4: Numerical results for the capacitated model: $h = 1$, $p = 9$, $c = 2$, $T = 10$; $K = 10$ for experiments 1 to 5, $K = 50$ for experiments 6 to 10

Appendix B

Appendix for Chapter 3

B.1 Proof of Ergodicity

In this section, we prove the existence and uniqueness of stationary distribution for the Markov chain induced by the class selection policy (CSP). Let requests for resources from a common resource pool of capacity $C \leq \infty$ arrive at time points $\{\tau_n, -\infty < n < \infty\}$. By observing the system at the moments of request arrivals, we define a discrete time process $I_n \triangleq (N_n^{(C)}, L_i, D_i, S_i, i = 1, 2, \ldots N_n^{(C)})$ where $N_n^{(C)}$ is the number of active (reserved) requests in the system at the moment of nth arrival τ_n, L_i is the elapsed time from the arrival of the i^{th} request to τ_n, D_i and S_i represent the reservation time (between arrival and actual service) and service time of the i^{th} request, respectively. Note that $L_i \leq D_i + S_i$ for $i = 1, 2, \ldots N_n^{(C)}$. The discrete-time Markov chain I_n describes the entire booking profile at the moment of nth arrival τ_n. We use a discrete version of Theorem 1 in Sevastyanov (1957) to prove the existence of a unique stationary distribution for $\{I_n\}$, which we state next for completeness.

Theorem B.1.1 *A Markov chain homogeneous in time has a unique stationary distribution which is ergodic if, for any $\epsilon > 0$, there exists a measurable set S, a probability distribution R on Ω, and $n_1 > 0$, $k > 0$, $K > 0$ such that*

- $kR(A) < P_{n_1}(x, A)$ *for all points $x \in H$ and measurable sets $A \subset H$; for any initial distribution P_0 there exists n_0 such that for any $n \geq n_0$,*

- $P_n(H) \geq 1 - \epsilon$,

- $P_n(A) \leq KR(A) + \epsilon$ *for all measurable sets* $A \subset H$.

Proof of Ergodicity. The proof follows similar arguments as in Lu and Radovanovic (2007b) and Sevastyanov (1957). Define set $H(a, b, c, d)$ as

$$H(a,b,c,d) \triangleq \left\{N_n^{(C)} \leq a, 0 \leq L_i \leq b, 0 \leq D_i \leq c, 0 \leq S_i \leq d\right\}, \tag{B.1}$$

for some positive finite constants a, b, c, d. Now we show that for any $\epsilon > 0$, there exists $H(a, b, c, d) \in \Omega$, such that for any initial distribution P_0 there exists n_0 such that for all $n \geq n_0$,

$$P_n(H(a, b, c, d)) \geq 1 - \epsilon. \tag{B.2}$$

Note that

$$\begin{aligned} P_n(\bar{H}(a,b,c,d)) \;\leq\; & \mathbb{P}\left[N_{0,n}^{(C)} \geq a\right] + \mathbb{P}\left[\bigcup_{i \in N_{0,n}^{(C)}} \{L_i > b\}, N_{0,n}^{(C)} \leq a\right] \qquad \text{(B.3)} \\ & + \mathbb{P}\left[\bigcup_{i \in N_{0,n}^{(C)}} \{D_i > c\}, N_{0,n}^{(C)} \leq a\right] \\ & + \mathbb{P}\left[\bigcup_{i \in N_{0,n}^{(C)}} \{S_i > d\}, N_{0,n}^{(C)} \leq a\right] \\ \leq\; & \mathbb{P}\left[N_{a,n}^{(C)} + N_{0,n}^{0} > a\right] + a\mathbb{P}\left[L_i > b\right] + a\mathbb{P}\left[D_i > c\right] + a\mathbb{P}\left[S_i > d\right], \end{aligned}$$

where $N_{a,n}^{(C)}$ represents the number of active requests at τ_n that originated from n arrivals at $\tau_0, \dots, \tau_{n-1}$, and the rest of active requests at τ_n, $N_{0,n}^0 = N_{0,n}^{(C)} - N_{a,n}^{(C)}$ are those that were active at the initial point τ_0 and are still active in the system at the moment of nth arrival. Next, since

$$
\begin{aligned}
& \mathbb{P}\left[N_{a,n}^{(C)} + N_{0,n}^{0} > a\right] \qquad\qquad (\text{B.4})\\
\leq\ & \mathbb{P}\left[N_{a,n}^{(C)} > \frac{a}{2}\right] + \mathbb{P}\left[N_{0,n}^{0} > \frac{a}{2}\right]\\
\leq\ & \mathbb{P}\left[N_{n}^{(\infty)} > \frac{a}{2}\right] + \mathbb{P}\left[\sum_{i=1}^{N_{0,n}^{0}} \mathbb{1}\left[D_i^0 + S_i^0 > \tau_n - \tau_0\right] > \frac{\psi}{2}\right]\\
\leq\ & \mathbb{P}\left[N_{n}^{(\infty)} > \frac{a}{2}\right]\\
& + \sum_{m=0}^{\infty} \mathbb{P}\left[N_{0,n}^{0} = m\right] \mathbb{P}\left[\sum_{i=1}^{m} \mathbb{1}\left[D_i^0 + S_i^0 > (1-\epsilon_1)n\mathbb{E}(\tau_1 - \tau_0)\right] > \frac{a}{2}\right]\\
& + \mathbb{P}\left[\tau_n - \tau_0 < (1-\epsilon_1)n\mathbb{E}(\tau_1 - \tau_0))\right],
\end{aligned}
$$

where $0 < \epsilon_1 < 1$ is an arbitrary constant and we used $N_n^{(\infty)} \geq N_{a,n}^{(C)}$ a.s. where $N_n^{(\infty)}$ is the active requests under infinite capacity system.

Next we prove that there exists $a = a_0$ large enough such that (B.4) is bounded by $\epsilon/4$. By virtue of Little's Law, we know that $\mathbb{E}N_n^{(\infty)} < \infty$ and therefore, uniformly for all $n > 0$,

$$\lim_{a\to\infty} \mathbb{P}\left[N_n^{(\infty)} > \frac{a}{2}\right] \to 0. \qquad (\text{B.5})$$

Next, note that $\mathbb{1}\left[D_i^0 + S_i^0 > (1-\epsilon_1)n\mathbb{E}(\tau_1 - \tau_0)\right] \leq \mathbb{1}\left[D_i^0 + S_i^0 > (1-\epsilon_1)\mathbb{E}(\tau_1 - \tau_0)\right]$ a.s., and that for any fixed m,

$$
\begin{aligned}
& \mathbb{P}\left[\sum_{i}^{m} \mathbb{1}\left[D_i^0 + S_i^0 > (1-\epsilon_1)n\mathbb{E}(\tau_1 - \tau_0)\right] > \frac{a}{2}\right] \qquad (\text{B.6})\\
\leq\ & \mathbb{P}\left[\sum_{i}^{m} \mathbb{1}\left[D_i^0 + S_i^0 > (1-\epsilon_1)\mathbb{E}(\tau_1 - \tau_0)\right] > \frac{a}{2}\right] \downarrow 0 \qquad \text{as} \qquad a \to \infty,
\end{aligned}
$$

which by the monotone convergence theorem implies that, uniformly for all $n > 0$,

$$\lim_{a\to\infty} \sum_{m=0}^{\infty} \mathbb{P}\left[N_{0,n}^{0} = m\right] \mathbb{P}\left[\sum_{i=1}^{m} \mathbb{1}\left[D_i^0 + S_i^0 > (1-\epsilon_1)n\mathbb{E}(\tau_1 - \tau_0)\right] > \frac{a}{2}\right] = 0. \quad (\text{B.7})$$

Finally, by the Weak Law of Large Numbers, for all n large enough,

$$\mathbb{P}\left[\tau_n - \tau_0 < (1-\epsilon_1)n\mathbb{E}(\tau_1 - \tau_0))\right] \leq \epsilon/12. \tag{B.8}$$

Thus, by (B.5) and (B.6), for an arbitrary $0 < \epsilon < 1$, there exists $n_0 < \infty$ and $a_0 < \infty$ large enough such that for all $n \geq n_0$,

$$\mathbb{P}\left[N_n^{(\infty)} > \frac{a_0}{2}\right] \leq \epsilon/12, \tag{B.9}$$

and

$$\sum_{m=0}^{\infty} \mathbb{P}\left[N_{0,n}^0 = m\right] \mathbb{P}\left[\sum_{i=1}^{m} \mathbb{1}\left[S_i^0 + D_i^0 > (1-\epsilon_1)n\mathbb{E}(\tau_1 - \tau_0)\right] > \frac{a_0}{2}\right] \leq \epsilon/12. \tag{B.10}$$

Now since $\mathbb{E}L_i < \infty$, $\mathbb{E}D_i < \infty$ and $\mathbb{E}S_i < \infty$, there exists b_0, c_0 and d_0 such that

$$\mathbb{P}\left[L_i > b_0\right] \leq \frac{\epsilon}{4a_0}, \qquad \mathbb{P}\left[D_i > c_0\right] \leq \frac{\epsilon}{4a_0}, \qquad \mathbb{P}\left[S_i > d_0\right] \leq \frac{\epsilon}{4a_0}. \tag{B.11}$$

Thus, by (B.8), (B.9), (B.10) and (B.11), we have

$$P_n(\bar{H}(a,b,c,d)) \leq \epsilon \qquad \Rightarrow \qquad P_n(H(a,b,c,d)) \geq 1 - \epsilon. \tag{B.12}$$

Next, we show that there exists $n_1 > 0$ and $k > 0$ such that for all points $x \in H(a_0, b_0, c_0, d_0)$ and measurable sets $A \in H(a_0, b_0, c_0, d_0)$, the following inequality holds

$$P_{n_1}(x, A) \geq kR(A). \tag{B.13}$$

Let $F_V(v)$ denote a cumulative distribution function of a random duration V, i.e. $\mathbb{P}[V \leq v]$. Next, for any n_1,

$$P_{n_1}(x, A) \geq P_1(x, \omega_0) P_{n_2}(\omega_0, A), \tag{B.14}$$

where $n_2 = n_1 - 1$. Let $x = (m, l_1, \ldots, l_m, d_1, \ldots, d_m, s_1, \ldots, s_m) \in H(a_0, b_0, c_0, d_0)$.

Then,

$$P_1(x,\omega_0) \geq \mathbb{P}\left[\tau_1-\tau_0 \geq \Delta, \text{all } m \text{ requests depart in } (\tau_0,\tau_1)\right] \tag{B.15}$$
$$\geq \mathbb{P}\left[\tau_1-\tau_0 > c_0+d_0\right] = 1-F_a(c_0+d_0) = e^{-\lambda(c_0+d_0)},$$

where $F_a(u)$ represents cumulative inter-arrival distribution of a renewal process $\{\tau_n\}$, i.e. $F_a(u)=\mathbb{P}\left[\tau_1-\tau_0\leq u\right]$. Next, we derive a lower bound for $P_{n_2}(\omega_0,A)$ for some n_2 large enough such that

$$\mathbb{P}\left[\tau_{n_2}-\tau_0 \geq c_0+d_0\right] \geq 1-\frac{\epsilon}{2}. \tag{B.16}$$

Note that the condition imposed on n_2 is possible due to the Weak Law of Large Numbers, since for any $\epsilon>0$ and all n_2 large enough with $c_0+d_0<(1-\epsilon)\mathbb{E}(\tau_{n_2}-\tau_0)$,

$$\mathbb{P}\left[\tau_{n_2}-\tau_0\geq c_0+d_0\right] \geq \mathbb{P}\left[\tau_{n_2}-\tau_0 \geq (1-\epsilon)\mathbb{E}(\tau_{n_2}-\tau_0)\right] \geq 1-\frac{\epsilon}{2}. \tag{B.17}$$

Next, pick any $x'=(m',l'_1,\dots l'_{m'},d'_1,\dots d'_{m'},s'_1,\dots,s'_{m'})\in A$. Define $x'+dx' \triangleq (m',l'_1+dl'_1,\dots l'_{m'}+dl'_{m'},d'_1+dd'_1,\dots d'_{m'}+dd'_{m'},s'_1+ds'_1,\dots,s'_{m'}+ds'_{m'})$ where $dl'_1,\dots,dl'_{m'},dd'_1,\dots,dd'_{m'},ds'_1,\dots,ds'_{m'}$ are infinitesimal elements. Then the transition probability into state $(x',x'+dx')$ starting from ω_0 can be lower bounded by the probability of the event that there are exactly m' arrivals between τ_1 and τ_{n_2} whose arrivals times are determined by $(\tau_{n_2}-l'_i-dl'_i,\tau_{n_2}-l'_i)$ for $i=1,\dots,m'$, and none of these m' arrivals concluded at time τ_{n_2} and there were no other arrivals. Therefore,

$$\mathbb{P}_{n_2}(\omega_0,(x'+dx')) \geq e^{-\lambda\tau_{n_2}}\frac{\lambda^{m'}}{m'!}\prod_{i=1}^{m'}\left[1-F_{D+S,D,S}(l_i,d_i,s_i)\right]. \tag{B.18}$$

where $F_{D+S,D,S}(\cdot)$ is the joint cumulative probability mass function. Now define probability distribution

$$R(A)\triangleq \nu\int_{x'\in A}\frac{\lambda^{m'}}{m'!}\prod_{i=1}^{m'}\left[1-F_{D+S,D,S}(l_i,d_i,s_i)\right], \tag{B.19}$$

where ν is a normalization constant. Thus, we have

$$P_{n_2+1}(x, A) \geq e^{-\lambda(c_0+d_0+\tau_{n_2})}\nu^{-1}R(A). \tag{B.20}$$

Finally, it is left to show that there exists $K > 0$ such that for every initial distribution P_0, for all n large and for any measurable set $A \subset S(a_0, b_0, c_0, d_0)$,

$$P_n(A) \leq KR(A) + \epsilon. \tag{B.21}$$

By (B.16), for all $n \geq n_2$,

$$\begin{aligned} P_n(A) &\leq \mathbb{P}\left[H_n \in A, \tau_n - \tau_0 > c_0 + d_0\right] + \mathbb{P}\left[\tau_n - \tau_0 \leq c_0 + d_0\right] \qquad \text{(B.22)} \\ &\leq \mathbb{P}\left[H_n \in A, \tau_n - \tau_0 > c_0 + d_0\right] + \frac{\epsilon}{2} \\ &\leq \int_{x' \in A} \left\{ \frac{\lambda^{m'}}{m'!} \prod_{i=1}^{m'} \left[1 - F_{D+S,D,S}(l_i, d_i, s_i)\right] \right\} + \epsilon \\ &\leq \nu^{-1}R(A) + \epsilon. \end{aligned}$$

We have verified the conditions stated in Theorem B.1.1 and thus the process $\{H_n\}$ has a unique stationary distribution as well implying the existence of the stationary blocking probability. $\square$

B.2 Proofs of Technical Lemmas and Theorems in Section 3

LEMMA 3.3.5. *Consider a random walk defined by a sequence of independent random variables $E_i = 1$ with probability p and -1 with probability $q = 1 - p$. Let $S_n = \sum_{i=1}^{n} E_i$. Define $M_\infty \in [0, u] \bigcup \{\infty\}$ to be maximum level attained by the random walk (i.e., $M_\infty = \max_n S_n$). Given that $0 \leq p < q \leq 1$ (downward drifting), then the probability that the random walk ever hits above level b is $\mathbb{P}\left(M_\infty \geq b\right) = \left(p/q\right)^b$.*

Proof of Lemma 3.3.5. Define the stopping time τ as follows,

$$\tau = \inf\{t \geq 1 : S_t \leq -a \text{ or } S_t \geq b\}.$$

It is straightforward to check the following two conditions,

$$\mathbb{E}(\tau) \leq \infty, \qquad \mathbb{E}(|E_{t+1} - E_t| \mid \mathcal{F}_t) \leq 2, \qquad \forall t \in \tau. \tag{B.23}$$

The Wald's identity (see Gallager (1996))

$$G_n(\theta) \triangleq \frac{e^{\theta S_n}}{[\phi(\theta)]^n} \tag{B.24}$$

is a martingale where the moment generating function $\phi(\theta) \triangleq \mathbb{E}(e^{\theta Y}) \geq 1$. First we compute $\hat{\theta}$ that solves the equation $\mathbb{E}(e^{\hat{\theta} Y}) = 1$, i.e.,

$$\mathbb{E}(e^{\hat{\theta} Y}) = pe^{\hat{\theta}} + qe^{\hat{\theta}} = 1 \qquad \Rightarrow \qquad e^{\hat{\theta}} = \frac{q}{p}. \tag{B.25}$$

By Optional Sampling Theorem (see Gallager (1996)),

$$\mathbb{E}\left[\frac{e^{\hat{\theta} S_\tau}}{[\phi(\hat{\theta})]^\tau}\right] = \mathbb{E}\left[e^{\hat{\theta} S_\tau}\right] = \mathbb{E}\left[e^{\hat{\theta} S_0}\right] = 1. \tag{B.26}$$

This leads to

$$\mathbb{P}(S_\tau \geq b)\underbrace{\mathbb{E}(e^{\hat{\theta} S_\tau} \mid S_\tau \geq b)}_{E_b} + (1 - \mathbb{P}(S_\tau \geq b))\underbrace{\mathbb{E}(e^{\hat{\theta} S_\tau} \mid S_\tau \leq -a)}_{E_a} = 1. \tag{B.27}$$

Thus, we have

$$\mathbb{P}(S_\tau \geq b) = \frac{1 - E_a}{E_b - E_a} = \frac{1 - e^{-\hat{\theta} a}}{e^{\hat{\theta} b} - e^{-\hat{\theta} a}} = \frac{1 - \left(\frac{q}{p}\right)^{-a}}{\left(\frac{q}{p}\right)^{b} - \left(\frac{q}{p}\right)^{-a}}. \tag{B.28}$$

Let $S_\tau^a \triangleq S_\tau$ be the stopping time location of the process. Let B_a be the event that the random walk hits b before $-a$. Observe that $\mathbb{P}(B_a) = \mathbb{P}(S_\tau^a \geq b)$ and also note that $B_i \subset B_{i+1}$ for all i. Define $B = \bigcup_{i=1}^{\infty} B_i$, i.e., there exists an i that the random walk hits b before $-i$. Therefore $\mathbb{P}(M_\infty \geq b) = \mathbb{P}(B)$. By properties of probability measures, we have

$$\mathbb{P}(M_\infty \geq b) \;=\; \mathbb{P}(\bigcup_{i=1}^{\infty} B_i) = \lim_{a\to\infty} \mathbb{P}(B_a) = \lim_{a\to\infty} \left(\frac{1-\left(\frac{q}{p}\right)^{-a}}{\left(\frac{q}{p}\right)^{b} - \left(\frac{q}{p}\right)^{-a}} \right) = \left(\frac{p}{q}\right)^{b} \tag{B.29}$$

This completes the proof. □

LEMMA 3.3.6. *Consider the counterpart system with an infinite number of servers, a customer arriving at the system at time* 0 *in steady state, observes that the pre-arrivals follow a non-homogeneous Poisson input process with piecewise rate* $\eta(r)$ *at time* r

$$\eta(r) = \begin{cases} \lambda, & \text{if} \quad r \leq 0, \\ \lambda(1 - F_D(\lceil r\rceil - 1)), & \text{if} \quad r > 0, \end{cases}$$

where F_D *is the cumulative probability mass function of* D *and* $F_D(\lceil r\rceil - 1)) = \sum_{i=0}^{\lceil r\rceil-1} f_D(i) = \sum_{i=0}^{\lceil r\rceil-1} \gamma_i$.

Proof of Lemma 3.3.6. Lemma 3.3.6 is a generalized version of Lemma 3.3.2. For $r \leq 0$, consider the time interval $(\lceil r\rceil - 1, \lceil r\rceil]$. By arguments similar to those used in Lemma 3.3.2, for each $l \in [0, u]$, the interval $(\lceil r\rceil - 1 - l, \lceil r\rceil - l]$ generates a stream of pre-arrivals over $(\lceil r\rceil - 1, \lceil r\rceil]$ that follow a Poisson process of rate $\gamma_l\lambda$. These processes are independent of each other and the overall merged process has rate $\lambda = \gamma_0\lambda + \gamma_1\lambda + \ldots + \gamma_u\lambda$.

For $\lceil r\rceil = i$ for $i \in [1, u]$, then the pre-arrivals prior to t over $(\lceil r\rceil - 1, \lceil r\rceil]$ are induced by arriving customers over the intervals $(\lceil r\rceil - l - 1, \lceil r\rceil - l]$, for $l \in [i, u]$, and the total rate is $\gamma_i\lambda + \gamma_{i+1}\lambda + \ldots + \gamma_u\lambda$. Note again that the rate $\gamma_i\lambda$ is induced from the Poisson arrival stream of customers over $(\lceil r\rceil - l - 1, \lceil r\rceil - l]$ who wish to start in l units of time. Since we only consider pre-arrivals prior to t, the terms $\gamma_{i-1}\lambda, \gamma_{i-2}\lambda, \ldots, \gamma_0\lambda$ are missing. □

LEMMA 3.3.7. *Consider the counterpart system with an infinite number of servers, if a customer comes at time* 0 *in steady state and requests service* ($S = 1$) *deterministically to commence in* D *units of time* ($D \in [0, u]$), *the conditional virtual blocking probability is given by, for all* $i \in [0, u]$,

$$P_i \triangleq \mathbb{P}(B \mid D = i) \triangleq \mathbb{P}\left(\max_{r\in[0,1]} \left\{\tilde{N}_{i+1}(1 - r; \lambda_{i+2}) + N_{i+2}(r; \lambda_{i+3})\right\} \geq C\right),$$

where N_i *is a Poisson counting process with rate* $\lambda_i = \lambda(1 - F_D(i - 2))$, *and* $\tilde{N}_i$ *is a mirror image of* N_i *with the same rate.*

Proof of Lemma 3.3.7. By Lemma 3.3.6, for each $i \in [1, u]$, the pre-arrival process N_i over the interval $(i - 2, i - 1]$ follows a Poisson process with rate $\lambda_i = \eta(i - 1) = \lambda(1 - F_D(i - 2))$. This implies that over the interval $(i - 1, i]$, the customers depart the system following a Poisson process with rate λ_i (a shift of N_i by 1 unit of time). Let $\tilde{N}_i$ be the mirror image of the departure process induced by N_i over $(i - 1, i]$, and therefore $\tilde{N}_i$ has the same rate λ_i. The rest of arguments is identical to that of Lemma 3.3.3. $\square$

THEOREM 3.3.8. *The conditional long-run virtual blocking probabilities have the following asymptotic upper bounds: for each* $i \in [0, u]$ *(the service distribution* $S = 1$ *deterministically),* $P_i \leq \Phi(-\beta)$.

Proof of Theorem 3.3.8. *First we assume that* $\gamma_0 > 0$. *By Lemma 3.3.7, we have that* $\lambda_1 > \lambda_2$ *and* $\lambda_i \geq \lambda_{i+1}$ *for each* $i \in [2, u]$. *By Theorem 3.3.4, it follows that,*

$$\lim_{\lambda\to\infty} P_0 = \Phi(-\beta); \qquad \lim_{\lambda\to\infty} P_i = 0, \qquad \text{for } i \in [1, u].$$

Therefore $P_i \leq \Phi(-\beta)$ *holds given that* $\gamma_0 > 0$. *In fact, we can relax the assumption of* $\gamma_0 > 0$. *If* $\gamma_0 = 0$, *it implies that over the interval* $(0, 1]$ *(recall that the customer arrives at time* 0 *in steady state), the departure rate is equal to the pre-arrival rate, i.e.,* $\lambda_1 = \lambda_2 = \lambda$. *Theorem 3.3.4 cannot be applied under this case. However, the fact that* $\gamma_0 = 0$ *implies that no arriving customers will start the service right away.*

Therefore, we do not have to consider the probability P_0 in the expression of P. Let i' be the minimal index such that $\gamma_{i'} > 0$ and $\gamma_i = 0$ for $i \in [0, i'-1]$, by the same argument, we can ignore the probabilities $P_0, \ldots, P_{i'-1}$. Instead, again by Theorem 3.3.4, we have

$$\lim_{\lambda\to\infty} P_{i'} = \Phi(-\beta); \qquad \lim_{\lambda\to\infty} P_i = 0, \qquad i \in [i'+1, u].$$

Thus, $P_i \leq \Phi(-\beta)$ still holds and this completes the proof. $\square$

LEMMA 3.3.9. *Let N_i^j and $\tilde{N}_i^j$ be defined as above. Then, for each $j \in [1, v]$ and each $i \in [1, u]$, N_i^j and $\tilde{N}_i^j$ are Poisson processes with the same rate*

$$\lambda_i^j = \kappa_j \lambda \left(1 - \sum_{l=0}^{i-j-1} \gamma_l^j\right) = \lambda_1^j (1 - F_D^j(i-j-1)). \tag{B.30}$$

Moreover, N_i^j is independent of $N_{i'}^{j'}$ for $i \neq i'$ or $j \neq j'$.

Proof of Lemma 3.3.9. *For each set $j \in [1, v]$, and $i \in [1, u]$, the pre-arrivals prior to t (i.e., N_i^j) over $(i-j-1, i-j]$ are induced by arriving customers over the intervals $(i-j-l-1, i-j-l]$ for $l = \max(0, i-j), \ldots, u$, and the total rate λ_i^j is therefore*

$$\gamma_{\max(0,i-j)}^j \lambda_1^j + \ldots + \gamma_u^j \lambda_1^j = \kappa_j \lambda \left(1 - \sum_{l=0}^{i-j-1} \gamma_l^j\right) = \lambda_1^j (1 - F_D^j(i-j-1)). \tag{B.31}$$

Note again that the rate γ_l^j is induced from the Poisson arrival stream of customers over $(i-j-l-1, i-j-l]$ who wish to start in l units of time. It follows from Poisson splitting arguments that N_i^j and $N_{i'}^{j'}$ are independent of each other for $(i, j) \neq (i', j')$. Note that they are generated by pre-arrivals in disjoint intervals. $\square$

B.3 Analysis of Price-driven Customer Arrivals

For the model with price-driven demand we use the following nonlinear program (NLP1):

$$\max_{\alpha_{ijk}, r_k} \quad \sum_{k=1}^{M} \sum_{i,j} r_k \alpha_{ijk} \lambda_{ijk}(r_k) j, \tag{B.32}$$

$$\text{s.t.} \quad \sum_{k=1}^{M} \sum_{i,j} \alpha_{ijk} \lambda_{ijk}(r_k) j \leq C, \tag{B.33}$$

$$0 \leq \alpha_{ijk} \leq 1, \qquad \forall i, j, k,$$

$$0 \leq r_k \leq 1, \qquad \forall k.$$

In particular, it can be verified that any optimal solution of (NLP1) has only non-negative prices. Also, observe that for any fixed prices $r_1, \ldots, r_M$*, the corresponding solution of* $\{\alpha_{ijk}\}$ *has the same knapsack structure defined in Section 2 above. Let* $(r^*. \alpha^*) = \{r_k, \alpha_{ijk}\}$ *be the corresponding optimal solution. Note that if one can solve (NLP1) and obtain the solution* (r^*, α^*) *then one can construct a similar CSP that will be amenable to the same performance analysis discussed in Section 3 above. However, solving (NLP1) directly may be computationally hard. Next, we show that under relatively mild assumptions imposed on the functions* $\lambda_1(r_1), \ldots, \lambda_M(r_M)$*, one can reduce (NLP1) to an equivalent nonlinear program that is more tractable; we denote it by (NLP2). (By equivalent we mean that they have the same set of optimal solutions.) Consider (NLP2) as follows:*

$$\max_{r_k} \quad \sum_{k=1}^{M} \sum_{i,j} r_k \lambda_{ijk}(r_k) j, \tag{B.34}$$

$$\text{s.t.} \quad \sum_{k=1}^{M} \sum_{i,j} \lambda_{ijk}(r_k) j \leq C. \tag{B.35}$$

$$0 \leq r_k \leq 1, \qquad \forall k.$$

It can be readily verified that as long as $\lambda_{ijk}(r_k)$ *is nonnegative (and decreasing) it is always optimal to have nonnegative prices, so the nonnegativity constraints can be* *dropped.*

Theorem B.3.1 *The programs (NLP1) and (NLP2) are equivalent.*

Proof of Theorem B.3.1. *First, we show that for each solution $\{r_k\}$ of (NLP2), we can construct a solution of (NLP1) with the same objective value. Specifically, consider solution $\{r'_k, \alpha'_{ijk}\}$ such that $r'_k = r_k$ and $\alpha'_{ijk} = 1$ if and only if $\sum_{i,j} \lambda_{ijk}(r_k)j > 0$. It can be verified that the resulting solution is feasible for (NLP1) and has the same objective value.*

Next, we show how to map optimal solution $\{r^*_k, \alpha^*_{ijk}\}$ of (NLP1) to a feasible solution of (NLP2) with the same objective function. For each $i = 1, \ldots, M' - 1$, set $r_k = r^*_k$, and for each $i = M' + 1, \ldots, M$ set set $r_k = r_\infty$. It is clear that, for each $i \neq M' - 1$, the resulting contributions to the objective value and constraint (3.2) are the same as in (NLP1). Consider now possibly fractional α'_M. The respective contribution of class M' to the objective value is $\sum_{i,j} r^*_{M'}\alpha^*_{ijM'}\lambda_{ijM'}(r^*_{M'})j$. Similarly, the contribution to constraint (3.2) $\sum_{i,j} \alpha^*_{ijM'}\lambda_{ijM'}(r^*_{M'})j$. Thus, it is sufficient to show that there exists a price $r_{M'}$ such that $\sum_{i,j} r_{M'}\lambda_{ijM'}(r_{M'})j \geq \sum_{i,j} r^*_{M'}\alpha^*_{ijM'}\lambda_{ijM'}(r^*_{M'})j$ and $\sum_{i,j} \lambda_{ijM'}(r_{M'})j \leq \sum_{i,j} \alpha_{ijM'}\lambda_{ijM'}(r^*_{M'})j$.

Since $\sum_{i,j} r^*_{M'}\lambda_{ijM'}(r^*_{M'})j \geq \sum_{i,j} r^*_{M'}\alpha^*_{ijM'}\lambda_{ijM'}(r^*_{M'})j$, we know that there exists $\bar{r} \in [r_{M'}, r_\infty)$ such that $\sum_{i,j} \bar{r}\lambda_{ijM'}(\bar{r})j = \sum_{i,j} r^*_{M'}\alpha^*_{ijM'}\lambda_{ijM'}(r^*_{M'})j$ by the properties of $\lambda_{ijM'}(r_{M'})$. Note that $\bar{r} \geq r^*_{M'}$, and therefore, we obtain $\sum_{i,j} r^*_{M'}\lambda_{ijM'}(\bar{r})j \leq \sum_{i,j} \bar{r}\lambda_{ijM'}(\bar{r})j = \sum_{i,j} r^*_{M'}\alpha^*_{ijM'}\lambda_{ijM'}(r^*_{M'})j$. Therefore, we have that $\sum_{i,j} \lambda_{ijM'}(\bar{r})j \leq \sum_{i,j} \alpha^*_{ijM'}\lambda_{ijM'}(r^*_{M'})j$, which concludes the proof of this theorem. $\square$

Theorem B.3.1 implies that instead of solving (NLP1) we can solve (NLP2). However, (NLP2) is computationally more tractable and can be solved relatively easy in many scenarios. Specifically, Lagrangify (dualize) constraint 3.2 with some Lagrange multiplier Θ and consider the unconstraint problem

$$\max_{r_k \in [\Theta, r_\infty)} \sum_{1 \leq k \leq M} \sum_{i,j} (r_k - \Theta)\lambda_{ijk}(r_k)j,$$

which is separable in $r_1, \dots, r_{M'}$. In fact, one aims to find the minimal Θ for which the resulting solution satisfies constraint 3.2. This can be done by applying bi-section search on the interval $[0, p_\infty]$. The complexity of this procedure depends on the complexity of maximizing $\sum_{1 \le k \le M} \sum_{i,j} (r_k - \Theta)\lambda_{ijk}(r_k)j$ for each $1 \le k \le M$. It is not hard to check that there are at least two tractable cases: (i) $\lambda_{ijk}(r_k)$ is a concave function on $[0, r_\infty)$, for each $1 \le k \le M$; (ii) $\lambda_{ijk}(r_k)$ is convex, but $r_k\lambda_{ijk}(r_k)$ is concave function on $[0, r_\infty)$, for each $1 \le k \le M$.

Appendix C

Appendix for Chapter 4

C.1 Nash Bargaining Game

A n-person Nash Bargaining game consists of a pair $(\mathcal{N}, \omega)$*, where* $\mathcal{N} \subseteq \mathbb{R}^n_+$ *is a compact and convex set and* $\omega \in \mathcal{N}$*. Set* $\mathcal{N}$ *is the feasible set and its elements give utilities that the* n *players can simultaneously accrue. Point* ω *is the disagreement point - it gives the utilities that the* n *players obtain if they decide not to cooperate. Game* $(\mathcal{N}, \omega)$ *is said to be feasible if there is a point* $\upsilon \in \mathcal{N}$ *such that* $\upsilon_1 > \omega_1$ *and* $\upsilon_2 > \omega_2$*. The solution to a feasible game is the point that satisfies the following four axioms,*

1. *Pareto optimality: No point in* $\mathcal{N}$ *can weakly dominate* υ*.*
2. *Invariance under affine transformation of utilities*
3. *Symmetry: The numbering of the players should not affect the solution.*
4. *Independence of irrelevant alternatives: If* υ *is the solution for* $(\mathcal{N}, \omega)$*, and* $\mathcal{S} \subseteq \mathbb{R}^n_+$ *is a compact and convex set satisfying* $\omega \in \mathcal{S}$ *and* $\upsilon \in S \subseteq \mathcal{N}$*, then* υ *is also the solution for* $(\mathcal{S}, \omega)$*.*

Nash Bargaining Solution (NBS) *If game* $(\mathcal{N}, \omega)$ *is feasible then there is a unique point in* $\mathcal{N}$ *satisfying the axioms stated above. This is also the unique point that maximizes* $\prod_{i=1}^{n}(\upsilon_i - \omega_i)$ *over all* $\upsilon \in \mathcal{N}$*.*

Bibliography

Acemoglu, D., A. Ozdaglar. 2007. Competition and efficiency in congested markets. Mathematics of Operations Research **32***(1) 1–31.*

Adelman, D. 2006. A simple algebraic approximation to the Erlang loss system. Operations Research Letters **36***(4) 484–491.*

Adelman, D. 2007. Price-directed control of a closed logistics queuing network. Operations Research **55***(6) 1022–1038.*

Angelus, A., E. L. Porteus. 2002. Simultaneous capacity and production management of short-life-cycle, produce-to-stock goods under stochastic demand. Management Science **48***(3) pp. 399–413.*

Aviv, Y., A. Federgruen. 2001. Capacitated multi-item inventory systems with random and seasonally fluctuating demands: Implications for postponement strategies. Management Science **47***(4) 512–531.*

Bamford, J., D. Ernst, D. G. Fubini. 2004. Launching a world-class joint venture. Harvard Business Review .

Bassok, Y., R. Anupindi, R. Akella. 1999. Single-period multiproduct inventory models with substitution. Operations Research **47***(4) pp. 632–642.*

Bernstein, F., A. Federgruen. 2007. Coordination Mechanisms for Supply Chains Under Price and Service Competition. Manufacturing and Service Operations Management **9***(3) 242–262.*

Binmore, K., A. Rubinstein, A. Wolinsky. 1986. The nash bargaining solution in economic modelling. The RAND Journal of Economics **17***(2) pp. 176–188.*

Bollapragada, S., T. E. Morton. 1999. A simple heuristic for computing nonstationary (s,S) policies. Operations Research **47***(4) 576–584.*

Borodin, A., R. El-Yaniv. 1998. Online Computation and Competitive Analysis. Cambridge University Press.

Bouchard, B., N. E. Karoui, N. Touzi. 2005. Maturity randomization for stochastic control problems. The Annals of Applied Probability **15***(4) 2575–2605.*

Boxma, O. J. 1984. $M/G/\infty$ tandem queues. Stochastic Processes and their Applications **18** *153–164.*

Bradley, J. R., P. W. Glynn. 2002. Managing capacity and inventory jointly in manufacturing systems. Management Science **48***(2) pp. 273–288.*

Brueckner, J. K., P. T. Spiller. 1994. Economies of traffic density in the deregulated airline industry. Journal of Law and Economics **37***(2) pp. 379–415.*

Brueckner, Jan K. 2003. The benefits of codesharing and antitrust immunity for international passengers, with an application to the star alliance. Journal of Air Transport Management **9***(2) 83 – 89.*

Burman, D. Y., J. P. Lehoczky, Y. Lim. 1984. Insensitivity of blocking probabilities in a circuit-switching network. J. Appl. Probab. **21***(4) 850–859.*

Cachon, G. P. 2003. Handbook in Operations Research and Management Science: Supply Chain Management, chap. Supply chain coordination with contracts. Elsevier, Amsterdam, 229–339.

Cachon, G. P., M. A. Lariviere. 1999. Capacity choice and allocation: Strategic behavior and supply chain performance. Management Science **45***(8) pp. 1091–1108.*

Cachon, G. P., M. A. Lariviere. 2005. Supply chain coordination with revenue-sharing contracts: Strengths and limitations. Management Science **51***(1) 30–44.*

Caldentey, R., L. M. Wein. 2003. Analysis of a decentralized production-inventory system. Manufacturing and Service Operations Management **5***(1) 1.*

Carr, S., W. Lovejoy. 2000. The inverse newsvendor problem: Choosing an optimal demand portfolio for capacitated resources. Management Science **46***(7) pp. 912–927.*

Caves, D. W., L. R. Christensen, M. W. Tretheway. 1984. Economies of density versus economies of scale: Why trunk and local service airline costs differ. The RAND Journal of Economics **15***(4) pp. 471–489.*

Chan, C. W., V. F. Farias. 2009. Stochastic depletion problems: Effective myopic policies for a class of dynamic optimization problems. Mathematics of Operations Research **34** *333–350.*

Chen, F. 1999. 94%-effective policies for a two-stage serial inventory system with stochastic demand. Management Science **45***(12) 1679–1696.*

Cheng, F., S. Sethi. 1997. Optimality of (s, S) policies in inventory models with markovian demand. Operations Research **45***(6) 931–939.*

Chod, J., N. Rudi. 2006. Strategic investments, trading, and pricing under forecast updating. Management Science **52***(12) pp. 1913–1929.*

Chu, L., Z. Shen. 2010. A power-of-two ordering policy for one-warehouse multiretailer systems with stochastic demand. Operations Research **58***(2) 492–502.*

Coffman-Jr, E. G., P. Jelenkovic, B. Poonen. 1999. Reservation probabilities. Adv. Perf. Anal. **2** *129–158.*

Corbett, C. J., G. A. DeCroix, A. Y. Ha. 2005. Optimal shared-savings contracts in supply chains: Linear contracts and double moral hazard. European Journal of Operational Research **163***(3) 653 - 667.*

Correa, J. R., A. S. Schulz, N. E. Stier-Moses. 2004. Selfish routing in capacitated networks. Mathematics of Operations Research **29***(4) 961–976.*

Correa, J. R., A. S. Schulz, N. E. Stier-Moses. 2007. Fast, Fair, and Efficient Flows in Networks. Operations Research **55***(2) 215–225.*

Dean, B. C., M. X. Goemans, J. Vondrák. 2004. Approximating the stochastic knapsack problem: The benefit of adaptivity. Proceedings of the 45th Annual IEEE Symposium on the Foundations of Computer Science 208–217IEEE Computer Society, Washington, DC.

Dye, S., L. Stougie, A. Tomasgard. 2003. The stochastic single resource service provision problem. Naval Research Logistics **50** *869–887.*

Eick, S. G., W. A. Massey, W. Whitt. 1993a. Infinite-server approximations for multi-server loss models with time-dependent arrival rates. Working Paper .

Eick, S. G., W. A. Massey, W. Whitt. 1993b. The physics of the $M_t/G/\infty$ queue. Operations Research **41***(4) 731–742.*

Erlang, A. K. 1917. Solution of some problems in the theory of probabilities of significance in automatic telephone exchanges. Elektrotkeknikeren **13** *5–13.*

Fan-Orzechowski, X., E. A. Feinberg. 2006. Optimality of randomized trunk reservation for a problem with a single constraint. Adv. Appl. Probab. **38***(1) 199–220.*

Farahat, A., G. Perakis. 2010a. The anarchy of price competition. Working paper .

Farahat, A., G. Perakis. 2010b. A nonnegative extension of the affine demand function and equilibrium analysis for multiproduct price competition. Operations Research Letters **38***(4) 280–286.*

Federgruen, A., Y. S. Zheng. 1991. An efficient algorithm for computing an optimal (r, Q) policy in continuous review stochastic inventory systems. Operations Research **40***(4) 808–813.*

Federgruen, A., P. Zipkin. 1986. An inventory model with limited production capacity and uncertain demands ii. the discounted-cost criterion. Mathematics of Operations Research **11***(2) pp. 208–215.*

Federgruen, A., P. H. Zipkin. 1984. An efficient algorithm for computing optimal (s, S) policies. Operations Research **32***(6) 1268–1285.*

Gallager, R. G. 1996. Discrete Stochastic Processes. Kluwer.

Gallego, G., Ö. Özer. 2001. Integrating replenishment decisions with advance demand information. Management Science **47***(10) 1344–1360.*

Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demands over finite horizons. Management Science **40** *999–1020.*

Gavirneni, S. 2001. An efficient heuristic for inventory control when the customer is using a (s, S) policy. Operations Research Letters **28***(4) 187–192.*

Guan, Y., A. J. Miller. 2008a. A polynomial time algorithm for the stochastic uncapacitated lot-sizing problem with backlogging. IPCO 2008, volume 5035 of Lecture Notes in Computer Science 450–462.

Guan, Y., A. J. Miller. 2008b. Polynomial time algorithms for stochastic uncapacitated lot-sizing problems. Operations Research **56***(5) 1172–1183.*

Halfin, S., W. Whitt. 1981. Heavy-traffic limits for queues with many exponential servers. Operations Research **29***(3) 567–588.*

Halman, N., D. Klabjan, M. Mostagir, J. Orlin, D. Simichi-Levi. 2009. A fully polynomial time approximation scheme for single-item stochastic lot-sizing problems with discrete demand. Mathematics of Operations Research **34***(3) 674–685.*

Hu, X., R. Caldentey, G Vulcano. 2012. Revenue sharing in airline alliances. Forthcoming in Management Science .

Huang, K., S. Küçükyavuz. 2008. On stochastic lot-sizing problems with random lead times. Operations Research Letters **36***(3) 303–308.*

Hunt, P. J., C. N. Laws. 1997. Optimization via trunk reservation in single resource loss systems under heavy traffic. Ann. Appl. Probab. **7***(4) 1058–1079.*

Iyengar, G., K. Sigman. 2004. Exponential penalty function control of loss networks. Ann. Appl. Probab. **14***(4) 1698–1740.*

Jain, H. 2010. Alliance revenue management in practice : techniques and simulation analysis. Master's thesis, Massachusetts Institute of Technology, USA.

Kapuscinski, R., S. Tayur. 1998. A capacitated production-inventory model with periodic demand. Operations Research **46***(6) pp. 899–911.*

Kaufman, J. S. 1981. Blocking in a shared resources environment. IEEE Trans. Comm. **29** *1474–1481.*

Keeler, J P., J P. Formby. 1994. Cost economies and consolidation in the us airline industry. International Journal of Transport Economics **21***(1) pp. 21– 45.*

Kelly, F. P. 1991. Effective bandwidths at multi-class queues. Queueing Systems **9***(1-2) 5–16.*

Key, P. 1990. Optimal control and trunk reservation in loss networks. Probab. Engrg. Informs. Sci. **4** *203–242.*

Koutsoupias, E., C. Papadimitriou. 1999. Worst-case equilibria. Proceedings of 16th Sympos. Theoretical Aspects of Computer Science. 404–413.

Kouvelis, P., M. A. Lariviere. 2000. Decentralizing cross-functional decisions: Coordination through internal markets. Management Science **46***(8) pp. 1049–1058.*

Kumar, S., R. Srikant, P. R. Kumar. 1998. Bounding blocking probabilities and throughput in queueing networks with buffer capacity constraints. Queueing Systems **28***(1-3) 55–77.*

Lawler, G. F. 2006. Introduction to Stochastic Processes. Chapman & Hall.

Levi, R., G. Janakiraman, M. Nagarajan. 2008a. A 2-approximation algorithm for stochastic inventory control models with lost-sales. Mathematics of Operations Research **33***(2) 351–374.*

Levi, R., A. Lodi, , M. Sviridenko. 2008b. Approximation algorithms for the multi-item capacitated lot-sizing problem via flow-cover inequalities. Mathematics of Operations Research **33***(2) 461–474.*

Levi, R., M. Pál, R. O. Roundy, D. B. Shmoys. 2007. Approximation algorithms for stochastic inventory control models. Mathematics of Operations Research **32***(2) 284–302.*

Levi, R., G. Perakis, C. Shi, W. Sun. 2011. Efficiency of revenue-sharing contracts in joint ventures in operations management. Working Paper.

Levi, R., A. Radovanovic. 2010. Technical note: Provably near-optimal LP-based policies for revenue management in systems with reusable resources. Operations Research **58***(2) 503–507.*

Levi, R., R. O. Roundy, D. B. Shmoys. 2006. Primal-dual algorithms for deterministic inventory problems. Mathematics of Operations Research **31** *267–284.*

Levi, R., R. O. Roundy, D. B. Shmoys, M. Sviridenko. 2008c. A constant approximation algorithm for the one-warehouse multi-retailer problem. Management Science **54***(4) 763–776.*

Levi, R., R. O. Roundy, D. B. Shmoys, V. A. Truong. 2008d. Approximation algorithms for capacitated stochastic inventory models. Operations Research **56***(5) 1184–1199.*

Levi, R., R. O. Roundy, V. A. Truong. 2005. Provably near-optimal balancing policies for multi-echelon stochastic inventory control models. Working Paper.

Levi, R., C. Shi. 2009. Approximation algorithms for the stochastic lot-sizing problem. Working Paper.

Levi, R., C. Shi. 2010. Approximation algorithms for the stochastic lot-sizing problem under uniform capacity constraints. Working Paper.

Levi, R., C. Shi. 2011a. Dynamic pricing models with reusable resources and advanced reservations. Working Paper .

Levi, R., C. Shi. 2011b. Revenue management of reusable resources with advanced reservations. Working Paper.

Lin, X., R. Germain. 1998. Sustaining satisfactory joint venture relationships: The role of conflict resolution strategy. Journal of International Business Studies **29***(1) pp. 179–196.*

Louth, G., M. Mitzenmacher, F. Kelly. 1994. Bounding blocking probabilities and throughput in queueing networks with buffer capacity constraints. Theoret. Comput. Sci. **125** *45–59.*

Lu, X., J. S. Song, A. C. Regan. 2006. Inventory planning with forecast updates: Approximate solutions and cost error bounds. Operations Research **54** *1079–1097.*

Lu, Y., A. Radovanovic. 2007a. Asymptotic blocking probabilities in loss networks with subexponential demands. J. Appl. Probab. **44***(4) 1088–1102.*

Lu, Y., A. Radovanovic. 2007b. Asymptotic blocking probabilities in loss networks with subexponential demands. Working Paper Preprint. Available at http://arxiv.org/abs/0708.4059.

Luo, Y., O. Shenkar. 2002. An empirical inquiry of negotiation effects in cross-cultural joint ventures. Journal of International Management **8***(2).*

Martinez-de Alberniz, V., G. Roels. 2010. Competing for shelf space. Forthcoming in Production and Operation Management .

Martinez-de Alberniz, V., D. Simchi-Levi. 2009. Competition in the supply option market. Operation Research **57***(5) 1082–1097.*

Massey, W. A. 1985. Asymptotic analysis of the time dependent $M/M/1$ *queue. Mathematics of Operations Research* **10***(2) 305–327.*

Miller, B. 1969. A queueing reward system with several customer classes. Management Science **16***(3) 234–245.*

Miller, K. D. 1992. A framework for integrated risk management in international business. Journal of International Business Studies **23***(2) pp. 311–331.*

Möhring, R. H., F. J. Radermacher, G. Weiss. 1984a. Stochastic scheduling problems I: General strategies. Zeitschrift für Operations Research **28** *193–260.*

Möhring, R. H., F. J. Radermacher, G. Weiss. 1984b. Stochastic scheduling problems II: Set strategies. Zeitschrift für Operations Research **29** *65–104.*

Möhring, R. H., A. Schulz, M. Uetz. 1999. Approximation in stochastic scheduling: The power of LP-based priority policies. Journal of the ACM **46** *924–942.*

Motwani, R., P. Raghavan. 1995. Randomized algorithms. Cambridge University Press .

Myerson, R. B. 1979. Incentive compatibility and the bargaining problem. Econometrica **47***(1) pp. 61–73.*

Nagarajan, M., G. Sosic. 2008. Game-theoretic analysis of cooperation among supply chain agents: Review and extensions. European Journal of Operational Research **187***(3) 719 – 745.*

Nash, J. F. 1950. The bargaining problem. Econometrica **18***(2) pp. 155–162.*

Netessine, Serguei, Nils Rudi. 2003. Centralized and competitive inventory models with demand substitution. Oper. Res. **51***(2) 329–335.*

Özer, Ö., W. Wei. 2004. Inventory control with limited capacity and advance demand information. Operations Research **52***(6) 988–1000.*

Perakis, G., G. Roels. 2007. The price of anarchy in supply chains: Quantifying the efficiency of price-only contracts. Management Science **53***(8) 1249 – 1268.*

Porteus, E. L., S. Whang. 1991. On manufacturing/marketing incentives. Management Science **37***(9) pp. 1166–1181.*

Prandini, M., J. Lygeros, A. Nilim, S. Sastry. 1999. Randomized algorithms for probabilistic aircraft conflict detection. IEEE Conference on Decision and Control **3** *2444–2449.*

Puhalskii, A. A, M. I. Reiman. 1998. A critically loaded multirate link with trunk reservation. Queueing Systems **28***(1-3) 157–190.*

Reyniers, D. J., C. S. Tapiero. 1995. The delivery and control of quality in supplier-producer contracts. Management Science **41***(10) pp. 1581–1589.*

Ross, K., D. Tsang. 1989. The stochastic knapsack problem. Management Science **37***(7) 740–747.*

Ross, K., D. Yao. 1990. Monotonicity properties for the stochastic knapsack. IEEE Trans. Inform. Theory **36***(5) 1173–1179.*

Roughgarden, T. 2005. Selfish Routing and the Price of Anarchy. MIT Press.

Roughgarden, T., E. Tardos. 2002. How bad is selfish routing? Journal of ACM **49** *236–259.*

Roundy, R. O. 1993. Efficient, effective lot-sizing for multi-product, multi-stage production systems. Operations Research **41** *371–386.*

Rubinstein, A. 1982. Perfect equilibrium in a bargaining model. Econometrica **50***(1) pp. 97–109.*

Rust, J. 1997. Using randomization to break the curse of dimensionality. Econometrica **65***(3) 487–516.*

Scarf, H. 1960. The optimality of (S,s) policies in the dynamic inventory problem. Mathematical Methods in Social Sciences .

Schmidt, V. 1987. On joint queue-length characteristics in infinite-server tandem queues with heavy traffic. Adv. Appl. Prob. **19** *474–486.*

Sevastyanov, B. A. 1957. An ergodic theorem for markov processes and its application to telephone systems with refusals. Theory of Probability and its Applications **2***(1) 104–112.*

Shmoys, D., K. Talwar. 2008. A constant approximation algorithm for the a priori traveling salesman problem. Integer Programming and Combinatorial Optimization **5035** *331–343.*

Shmoys, D. B., C. Swamy. 2004. Stochastic optimization is (almost) as easy as deterministic optimization. Proceedings of the 45th Annual IEEE Symposium on the Foundations of Computer Science **50** *228–237. IEEE Computer Society, Washington, DC.*

Shmoys, D. B., C. Swamy. 2006a. Approximation algorithms for 2-stage stochastic optimization problems. ACM SIGACT News Archive **37** *33–46.*

Shmoys, D. B., C. Swamy. 2006b. An approximation scheme for stochastic linear programming and its application to stochastic integer programs. Journal of the ACM **53***(6) 978–1012.*

Silver, E. A., H. C. Meal. 1973. A heuristic selecting lot-size requirements for the case of a deterministic time varying demand rate and discrete opportunities for replenishment. Production and Inventory Management **14** *64–74.*

Song, J., P. H. Zipkin. 1993. Inventory control in a fluctuating demand environment. Operations Research **41***(2) 351–370.*

Stougie, L., M. H. van der Vlerk. 2003. Approximation in stochastic integer programming SOM Research Report 03A14, Eindhoven University of Technology, Eindhoven, The Netherlands.

Teo, C. P., D. Bertsimas. 1996. Improved randomized approximation algorithms for lot-sizing problems. Proceedings of the 5th International IPCO Conference on Integer Programming and Combinatorial Optimization 359–373.

Van Mieghem, J. A. 2003. Capacity management,investment, and hedging:review and recent developments. Manufacturing and Service Operations Management **5***(4) 269 – 302.*

Van Mieghem, J. A., M. Dada. 1999. Price versus production postponement: Capacity and competition. Management Science **45***(12) pp. 1631–1649.*

Van Mieghem, J. A., N. Rudi. 2002. Newsvendor networks: Inventory management and capacity investment with discretionary activities. Manufacturing and Service Operations Management **4***(4) 313.*

Vazirani, V. J. 2001. Approximation Algorithms. Springer.

Veinott, A. 1966. On the optimality of (s,S) inventory policies: new conditions and a new proof. SIAM J. Appl. Math **14***(5) 1067–1083.*

Weintraub, G. Y., R. Johari, B. Van Roy. 2010. Investment and market structure in industries with congestion. Forthcoming in Operations Research .

Whitt, W. 1985. Blocking when service is required from several facilities simultaneously. AT&T Tech. J. **64** *1807–1856.*

Zachary, S. 1991. On blocking in loss networks. Adv. Appl. Probab. **23***(2) 355–372.*