

CONTINUOUS TIME CONTROL OF MAKE-TO-STOCK PRODUCTION SYSTEMS

A DISSERTATION SUBMITTED TO
THE DEPARTMENT OF INDUSTRIAL ENGINEERING
AND THE INSTITUTE OF ENGINEERING AND SCIENCES
OF BILKENT UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

By
Önder Bulut
July, 2010

I certify that I have read this thesis and that in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of doctor of philosophy.

Asst. Prof. M. Murat Fadilođlu (Supervisor)

I certify that I have read this thesis and that in my opinion it is fully adequate, in scope and in quality, as a thesis for the degree of doctor of philosophy.

Prof. Nesim Erkip

I certify that I have read this thesis and that in my opinion it is fully adequate, in scope and in quality, as a thesis for the degree of doctor of philosophy.

Asst. Prof. Z. Pelin Bayındır

I certify that I have read this thesis and that in my opinion it is fully adequate, in scope and in quality, as a thesis for the degree of doctor of philosophy.

Asst. Prof. Seil Savařaneril

I certify that I have read this thesis and that in my opinion it is fully adequate, in scope and in quality, as a thesis for the degree of doctor of philosophy.

Asst. Prof. Osman Alp

Approved for the Institute of Engineering and Sciences:

Prof. Levent Onural

Director of Institute of Engineering and Sciences

ABSTRACT

Continuous Time Control of Make-to-Stock Production Systems

Önder Bulut

Ph.D. in Industrial Engineering

Supervisor: Asst. Prof. M. Murat Fadılođlu

July, 2010

We consider the problem of production control and stock rationing in a make-to-stock production system with multiple servers –parallel production channels–, and several customer classes that generate independent Poisson demands. At decision epochs, in conjunction with the stock allocation decision, the control specifies whether to increase the number of operational servers or not. Previously placed production orders cannot be cancelled. We both study the cases of exponential and Erlangian processing times and model the respective systems as $M/M/s$ and $M/E_k/s$ make-to-stock queues. We characterize properties of the optimal cost function, and of the optimal production and rationing policies. We show that the optimal production policy is a state-dependent base-stock policy, and the optimal rationing policy is of state-dependent threshold type. For the $M/M/s$ model, we also prove that the optimal ordering policy transforms into a bang-bang type policy when we relax the model by allowing order cancellations. Another model with partial order-cancellation flexibility is provided to fill the gap between the no-flexibility and the full-flexibility models. Furthermore, we propose a dynamic rationing policy for the systems with uncapacitated replenishment channels, i.e., exogenous supply systems.

Such systems can be modeled by letting s --the number of replenishment channels-- go to infinity. The proposed policy utilizes the information on the status of the outstanding replenishment orders.

This work constitutes a significant extension of the literature in the area of control of make-to-stock queues, which considers only a single server. We consider an arbitrary number of servers that makes it possible to cover the spectrum of the cases from the single server to the infinite servers. Hence, our work achieves to analyze both the exogenous and endogenous supply leadtimes.

Keywords: Inventory; Production; Stock Rationing; Dynamic Rationing; Make-to-Stock; Multiple Servers; Multiple Demand Classes; Optimal Control; Simulation.

ÖZET

Stoğa-Üretim Sistemlerinin Sürekli Zamanda Kontrolü

Önder Bulut

Endüstri Mühendisliği, Doktora

Tez Yöneticisi: Yrd. Doç. Dr. M. Murat Fadiloğlu

Temmuz, 2010

Bu çalışmada paralel üretim kanalları ve birden çok müşteri sınıfına sahip stoğa-üretim sistemleri için üretim ve stok tayinleme kontrol problemleri ele alınmaktadır. Çalışmada, bağımsız Poisson talep süreçleri varsayılmıştır. Karar anlarında kontroller, stok paylaşırma kararı ile birlikte aktif olan üretim kanal sayısının artırılıp arttırılmayacağını belirtir. Daha önce verilen üretim siparişleri iptal edilemez. Üssel ve Erlang dağılımına sahip üretim zamanlı sistemler sırasıyla $M/M/s$ ve $M/E_k/s$ stoğa-üretim kuyruk modelleri olarak incelenmiştir. Maliyet fonksiyonun ve en iyi üretim ve tayinleme politikalarının özellikleri belirlenmiştir. En iyi üretim politikasının duruma-bağımlı temel-düzey politikası ve en iyi tayinleme politikasının duruma-bağımlı eşik tipi politika olduğu gösterilmiştir. $M/M/s$ modeli için, herhangi bir siparişin iptal edilebilmesine izin verildiğinde, en iyi üretim politikasının ya hep-ya hiç tarzı bir politikaya dönüştüğü ispatlanmıştır. Tam iptal esnekliği olan ve hiç esnekliği olmayan modellerin yanısıra, kısmi sipariş iptal esnekliğini içeren model de incelenmiştir. Çalışmada ayrıca kapasite kısıtı olmayan tedarik kanalına sahip sistemler için, beklenen siparişlerin ulaşmalarına kalan zamanı kullanan, yeni bir devingen tayinleme politikası önerilmektedir.

Bu çalışmayla, stoğa-üretim sistemlerini şimdiye kadar tek bir üretim kanalıyla modelleyen çalışmaları içeren teknik yazına önemli katkıda bulunulmaktadır. Çalışmada ele alınan modellerde üretim kanalı sayısı herhangi bir değeri alabilecek şekilde serbest bırakılarak, kapasite kısıtı olan ve olmayan sistemlerin aynı anda incelenmesi sağlanmıştır.

Anahtar Kelimeler: Envanter; Üretim; Stok Tayınlama; Devingen Tayınlama; Stoğa-Üretim; Paralel Üretim Kanalları; Çoklu Müşteri Sınıfları; Eniyi Kontrol; Benzetim.

Acknowledgement

I would like to express my sincere gratitude to my advisor Asst. Prof. M. Murat Fadilođlu for all the trust and encouragement.

I am indebted to Prof. Nesim Erkip, Asst. Prof. Z. Pelin Bayındır, Asst. Prof. Seçil Savaşaneril, Asst. Prof. Osman Alp, Assoc. Prof. Oya Ekin Karaşan and Asst. Prof. Banu Yüksel Özkaya for accepting to read and review this thesis, and for their invaluable suggestions.

I would also like to thank TÜBİTAK, The Scientific and Technological Research Council of Turkey, for providing financial support to my PhD study under grant #2211.

CONTENTS

1	Introduction	1
2	Literature Review	9
3	<i>M/M/s</i> Model with multiple-demand-classes and lost sales	21
3.1	Primary Model	22
3.1.1	Model Formulation	22
3.1.2	Characterization of the Optimal Production and Rationing Policies ...	26
3.2	Variations on the Primary Model	30
3.2.1	Model with Full Order-Cancellation Flexibility	30
3.2.2	Model with Partial Order-Cancellation Flexibility	32
3.2.3	Model with Fixed Production and Order Cancellation Costs	34
3.3	Average Cost Criterion	35
3.4	Stationary Analysis	36
3.4.1	Stationary Analysis under Base-stock	36
3.4.2	Stationary Analysis under Bang-Bang Policy	40
3.5	Numerical Study	42

3.6	Appendix	55
4	M/E _k /s Model with multiple-demand-classes and lost sales	74
4.1	Model Formulation	74
4.2	Characterization of the Optimal Production and Rationing Policies	78
4.3	Appendix	84
5	A Dynamic Rationing Policy for Continuous-Review Inventory Systems	89
5.1	Dynamic Rationing	92
5.2	Rationing with Continuous Replenishment Flow	95
5.3	Performance Evaluation of <i>RERF</i> via Simulation	101
5.3.1	Backordering Case	104
5.3.2	Lost Sales Case	113
6	Conclusion	115

LIST OF TABLES

2.1: Related Literature on Continuous-time Systems	20
3.1: Optimal Production Policy under Discounted Cost Criterion	43
3.2: Optimal Rationing Policy For Class 2 under Discounted Cost Criterion	43
3.3: Optimal Production Policy under Average Cost Criterion	44
3.4: Optimal Rationing Policy For Class 2 under Average Cost Criterion	45
3.5: Optimal Rationing Policy versus FCFS policy: impact of demand mix	51
4.1: Optimal Production Policy under Discounted Cost Criterion	82
4.2: Optimal Rationing Policy For Class 2 under Discounted Cost Criterion	83
5.1: Comparison of Policies for all possible shortage cost pairs	105
5.2: Comparison of Policies for $\lambda=25$, $A=0$	111
5.3: Comparison of Policies for $\lambda=25$, $A=2$	112
5.4: Lost sales: Comparison of Policies for all possible shortage cost pairs	114

LIST OF FIGURES

3.1: Birth-Death Process for (S, K) Policy	37
3.2: Optimal production policy vs. Base-stock policy: Impact of number of available server.....	46
3.3: Optimal production policy vs. Base-stock policy: Impact of demand rate	48
3.4: Optimal production policy vs. Base-stock policy: Constant traffic intensity.....	50
3.5: Value of Order Cancellation Flexibility	50
3.6: Optimal rationing policy vs. FCFS Policy: Impact of demand rate	53
3.7: Optimal production and Base-stock policies with or without rationing.....	53
5.1: Exponential flow function ($L = 1, \lambda = 1$).....	100
5.2: Impact of demand mix	107
5.3: Impact of total demand rate ($p_1 = 0.5$).....	108
5.4: Impact of total demand rate ($p_1 = 0.9$).....	109

Chapter 1

Introduction

In this thesis, we study the problem of production control and stock rationing of a single-item, make-to-stock facility with parallel production channels and several demand classes. In a production system that keeps inventory to instantaneously satisfy random demand originating from distinct customer classes, the decision maker should develop a strategy in order to efficiently use system resources and allocate inventory among different customer classes.

In order to better understand the problem, let us consider the following example. We have a company that produces spare parts for a large car manufacturer. There are two types of demand for the parts. The first one is the demand from the car manufacturer and the second one is the demand from different spare part distributors. We are obliged to provide the parts demanded by the manufacturer instantaneously or pay a hefty fine due to our contract. We do not have such an obligation with the distributors, although the sale is lost. Hence the manufacturer's demand has higher priority. Our production facility has s parallel channels such that each can process one part at a time. Given the state of the system, we would like to determine how

many production channels to utilize and which types of demand to satisfy so as to operate the system optimally with respect to a predetermined cost function.

For the most general case of this problem, the generally distributed customer inter-arrival times and non-identical servers (production channels) with controllable generally distributed processing times should be considered. However, in order to make the analytical analysis tractable, we assume that each customer class generates demand according to a stationary Poisson process independent from the other classes, and the servers are identical with independent processing times. Moreover, we restrict the analysis to Markov policies by assuming exponential or Erlangian processing times. Therefore, the system considered in this research is modeled as an $M/M/s$ or an $M/E_k/s$ make-to-stock queue.

In the most general setting, characterization of the optimal strategy may not be analytically tractable because the decision maker should continuously adjust production and stock allocation decisions based on the current status of the production (the age information for all the outstanding production orders) and the current inventory level. The existing literature on the production control problem in make-to-stock systems does not even include the more tractable single server models with general production times. Moreover, the studies those consider the rationing problem in the classical inventory setting can only provide approximate results even for the static policies. The reader is directed to Chapter 2 for a more detailed discussion on the literature.

In our setting, at any point in time, the decision maker determines the number of active servers and makes a rationing decision for an arriving customer demand. For the lost sales case, we characterize structural properties of the optimal cost function, and the optimal production and stock rationing policies. In addition to these, we also propose a dynamic rationing policy for the systems with uncapacitated replenishment channels, which corresponds to the continuous-review inventory systems and can be

Chapter 1 Introduction

obtained by letting s --the number of replenishment channels-- go to infinity. The proposed policy utilizes the age information for all the outstanding replenishment orders.

Modeling the supply process of the inventory systems with an arbitrary number of replenishment channels is instrumental in filling the gap between the two main streams of studies in the existing literature. One of these streams considers standard inventory models where supply leadtimes are exogenous, i.e., the replenishment channel is uncapacitated. In this setting, the optimal policy is not fully characterized under lost sales. Recently, Zipkin (2008a) reformulates the standard periodic-review lost-sales inventory problem with a new approach based on discrete convex analysis. He shows that the optimal policy is state-dependent, i.e., the ages and the quantities of all outstanding orders have an effect on the optimal ordering decision. On the other hand, for the backordering case, Erhardt (1984) shows that if the replenishment orders do not cross in time (e.g., deterministic leadtimes), the optimal ordering policy is independent from the status of the outstanding orders. For such settings, simple base-stock, i.e., order-up-to, policy is optimal. Most of the studies in this stream assume deterministic leadtimes and provide analyses under simple base-stock policies irrespective of the shortage dynamics. However, for the lost sales systems, numerical results of Zipkin (2008b) manifest that simple (state-independent) base-stock policies do not perform well. In many settings, it is even worse than the constant-order policy, which orders the same amount at fixed intervals.

The other stream of studies considers production-inventory systems. These systems are characterized by capacitated replenishment channels. With the exception of the work of Zipkin (2000), all the works in this stream model endogenous supply leadtimes with a single server. For basic single server models (models without additional sources of information such as advance demand and assembly component inventory levels) the optimal production policy is a simple base-stock policy defined in

terms of a constant produce-up-to level. It is optimal to produce up to a certain inventory level and then stop the production. This holds for both of the lost sales and the backordering cases (see Ha, 1997a; Ha, 1997b; and Gayon et al., 2009b). In fact, simple base-stock is the only meaningful policy that can be considered for the single server case. In the manufacturing systems literature, this policy is known as the production authorization mechanism (see Buzacott and Shantikumar (1993, pp 103)). A natural extension of the single production channel is parallel production channels. Buzacott and Shantikumar (1993, pp. 43) call systems with parallel production channels “single-stage systems.” In these systems, “a job can be processed by any one of the machines, but only one machine is required to complete the required tasks.” Zipkin (2000, pp 244) calls the same kind of systems “parallel processing systems” and provides an analysis for such systems with independent, stochastic leadtimes under the base-stock policy. But the base-stock policy is not optimal for this setting. There is flexibility to utilize different number of servers at different inventory levels, which gives rise to the state-dependent policies. Identifying optimal production policies for “parallel processing systems” and quantifying the optimality gap left by the base-stock policy are among the main issues addressed in this PhD dissertation.

Our work achieves the analyses of both the exogenous and endogenous supply leadtimes within a single model. This is made possible by considering an arbitrary number of supply channels so as to cover the spectrum from the single server to the infinite servers. Our model allows analysis of single location continuous-review inventory systems with exogenous exponential leadtimes (i.e., uncapacitated replenishment channel) by letting the number of replenishment channels go to infinity. On the other hand, having $s = 1$ corresponds to the single server, capacitated production model, which is the subject of most of the literature on the control of make-to-stock queues. Furthermore, as Zipkin (2000, pp. 246) discusses, no real supply system has infinitely many processing channels. Therefore, realistic models should consider finite processing capacity. In this context, the $M/M/\infty$ and $M/E_k/\infty$ models

should be considered as the limiting cases of the $M/M/s$ and $M/E_k/s$ models, respectively. It should also be noted that the models give the exact solution of the $M/M/\infty$ and $M/E_k/\infty$ when s is selected to be sufficiently large, since the optimal number of servers to be utilized is bounded. The existence of such a bound (beyond which the system is equivalent to a system with exogenous leadtimes) is discussed in Section 3.5 and illustrated in Figure 3.2. Furthermore, an algorithm is provided to calculate this bound under the average cost criterion.

In our numerical study, we investigate the performance of the base stock policy in comparison with the optimal policy. If the number of servers is limited, i.e., production capacity is tight, base-stock performs well. When there is ample capacity, base-stock results in dramatic loss.

Furthermore, increasing the number of servers while keeping the traffic intensity constant, undermines the base-stock's performance. In this setting, as the available number of servers increases, the control space becomes more finely discretized. Consequently, the control problem resembles to the one that Mayorga et al. (2006) consider in which the service rate of a single server is controlled over a continuous set.

Another issue addressed in the thesis is the problem of allocating a common stock pool among different customer classes, which is known as the stock rationing problem in the literature. It allows differentiating customer classes in order to provide different service levels and to operate the system more cost-effectively. The stock rationing policy stops serving lower priority classes when the on-hand inventory drops below a certain threshold level. Under the threshold level, only the demands from higher priority classes are satisfied. There is a threshold rationing level for each customer class. The threshold levels could change dynamically according to the status of the production process.

Chapter 1 Introduction

Inventory systems subject to multiple demand classes for the same item are frequently observed in real life. Consider a spare parts inventory system. A part can be demanded in order to repair different end products of different importance and criticality. Considering the fact that all demands may not be satisfied instantaneously, demands for spare parts should be prioritized. Moreover, the system may experience urgent orders in case of system breakdowns. The unit shortage cost experience under such a scenario is to be dramatically higher compared to the unit shortage cost of the orders due to the planned maintenance activities. Another example would be a two-echelon inventory system consisting of a warehouse and many retailers. In case of stockout, retailers may place urgent, more critical orders to the warehouse. Furthermore, it may be beneficial to better serve certain retailers that constitute a larger portion of the warehouse's business. In multi-echelon systems, intershipments between the inventory locations in the same echelon may be allowed. However, for any inventory location, direct customer orders have precedence over the intershipment orders that are placed by the other locations.

Customer differentiation is also very important in service sectors. Hotel or airline companies ration their limited capacity according to the priorities of their different customer classes. In this setting, in addition to the rationing decision, another key concern is deciding the prices to be charged to individual customer classes.

The rest of the thesis is organized as follows. In Chapter 2 we review the related literature and then provide our models and analyses in subsequent chapters. Chapter 3 is devoted to the analysis of $M/M/s$ model. We first introduce our primary model and provide the dynamic programming formulation. The primary model assumes no setup cost and it is not allowed to cancel the previously placed production orders. That is, at decision epochs, in conjunction with the stock allocation decision, the control specifies whether to increase the number of operational servers or not. The objective is to minimize the infinite horizon expected discounted cost. We character-

ize the properties of the optimal cost function, and of the optimal production and rationing policies. We show that the optimal production policy is a state-dependent base-stock policy, and the optimal rationing policy is of threshold type. We also prove that the optimal production and rationing policies are monotone in the inventory level and the number of operational servers. We consider variations on the primary model in Section 3.2. It is shown that the optimal ordering policy transforms into a bang-bang type policy when we relax the model by allowing order cancellations. Another model with partial order-cancellation flexibility is provided to fill the gap between the no-flexibility and the full-flexibility models. The effects of the setup and the order cancellation costs are also discussed. In Section 3.3, we generalize our results and show that there exists an average cost optimal stationary policy that possesses all the structural properties of the optimal policy under the discounted cost criterion. Section 3.4 provides the stationary analysis of the system under both base-stock and bang-bang policies. Next, in Section 3.5, with a numerical study, we quantify the additional gain that the optimal policy provides over the –suboptimal— base-stock policy proposed in the literature, along with the value of the flexibility to cancel production orders. We also compare the optimal rationing policy with the first-come-first-served (FCFS) policy. The proofs of the lemmas and theorems presented in Chapter 3 are provided in the Appendix of the chapter (Section 3.6).

We discuss the system with Erlangian servers in Chapter 4. In order to get insights and develop a general method for the analysis of the $M/E_k/s$ model, we first consider the $M/E_k/2$ model and then discuss the generalization. We state some conjectures on the structure of the optimal cost function. Based on these properties, we characterize the optimal production and rationing policies. We have succeeded in proving some theoretical statements but some others remain conjectures. We present the model formulation in Section 4.1 and the analysis in Section 4.2. The proofs are presented in the Appendix (Section 4.3).

The analysis of the $M/E_k/s$ model is a direct extension of the $M/M/s$ model. Moreover, it provides clues for the analysis of the model with deterministic production times, because the deterministic production times are the limit of Erlangian production times as the number of Erlang stages increases. Thus, based on the discussion related to the system with Erlangian servers, it is possible to propose better performing rationing policies than the static one for the continuous-review inventory systems with deterministic lead time and uncapacitated replenishment channels (exogenous supply can be modeled by letting s --the number of replenishment channels-- go to infinity). The analysis of $M/E_k/s$ model shows that the outstanding that completes more Erlang stages has more value in terms of the rationing decision. Based on this fact, in Chapter 5, we propose a new class of dynamic rationing policies for continuous-review inventory systems with multiple customer classes. The new class of policies is based on the idea of rationing the inventory as if the outstanding replenishments were flowing into the system in a continuous fashion. The age information for all the outstanding orders is used to modify the inventory level dynamically. Upon a discussion delineating the effect of the flow function on the inventory dynamics (Section 5.1), we suggest a policy that assumes exponential flow of the replenishment orders (Section 5.2). For both backordering and lost sales environments, in Section 5.3, we conduct simulation studies to compare the performance of the dynamic policy with the static critical level and the FCFS (common stock) policies and quantify the gain obtained. We also propose two new bounds on the performance of the --unknown-- optimum dynamic rationing policy that enables us to tell how much of the potential gain the proposed dynamic policy realizes. We discuss the conditions under which stock rationing --both dynamic and static-- is beneficial and assess the value of the dynamic policy. Finally, we provide concluding remarks and discuss future research directions in Chapter 6.

Chapter 2

Literature Review

In this chapter we review the literature on stock rationing and production control for make-to-stock systems. We classify the stock rationing literature according to the assumed rationing policy and its dynamics (static or state-dependent), and by the clearing mechanism for the backorders that defines how to handle the arriving replenishment orders. Similar to the other stochastic inventory problems, stock rationing literature can also be categorized based on the review policy (continuous or periodic) and on the consequence of shortages (backorders or lost sales). There is also a parallel literature on the production environment. We review all the important works that address rationing problem in these different settings.

The other area of subject that we achieve to make contribution is the control of make-to-stock production systems. Contrary to the inventory systems, production systems have capacitated channels, i.e., the number replenishment channels is limited. All the works in this stream model the capacitated channel with a single server. However, they assume different processing time and shortage cost structures. Some of the works also consider settings in which there are other sources of information such as advanced demand and assembly component inventory levels.

Chapter 2 Literature Review

Veinott (1965) is the first to study the rationing problem. He considers a zero leadtime backordering model in the periodic review setting with exogenous supply. He introduces the concept of threshold rationing levels that are used to allocate the on-hand inventory among different customer classes. For the same setting, using dynamic programming Topkis (1968) shows that a time remembering rationing policy is optimal. He also considers the lost sales case. He divides the review periods into sub-periods and finds the threshold rationing levels (for all classes) at each sub-period that depend on the remaining time to the next review.

For the infinite horizon multi-period problem, under the static rationing policy, i.e., the threshold rationing levels are state-independent and fixed; Nahmias and Demmy (1981) derive approximate expressions for the expected number of backorders for each customer class. They assume that the stock is replenished according to the (s, S) policy, leadtime is zero, and demand is realized at the end of each review period. Cohen et al. (1989) also consider a periodic review (s, S) policy with lost sales, deterministic leadtimes and two demand classes. At the end of each period, after the demand realizations, they use the on hand stock to meet the demands of customer classes in the order of priorities. They propose a greedy heuristic to minimize the expected system cost under the service level constraints.

Frank et al. (2003) analyze a periodic review model with two demand classes. While high priority class experiencing deterministic demand, the demand for the other class is stochastic. The deterministic demand must be met immediately in each period and any unsatisfied stochastic demand is lost. They show that the optimal replenishment and rationing policies are complex in structure of the optimal policy and propose the simpler (s, S) replenishment policy under static rationing. In the study, it is assumed that the orders arrive instantaneously. Therefore, stock rationing is used to gain from fixed ordering cost instead of saving stock for future deterministic demand.

Chapter 2 Literature Review

The study of Nahmias and Demmy (1981) is the first in the literature that considers the stock rationing problem in continuous time. They consider a setting with unit Poisson arrivals, two demand classes, constant leadtime, and backordering. They derive approximate expressions for the expected number of backorders and for the fill rates for both classes under a (Q, r, K) policy, which is a (Q, r) policy with the fixed threshold rationing level K . Their approximation is based on the at-most-one-order-outstanding assumption. Moon and Kang (1998) extend this work by considering compound Poisson demand. They analyze the system with a simulation model.

In backordering environments, to completely define the stock rationing policy the way that the backorders are cleared should also be defined. The clearing mechanism specifies how the replenishment orders should be allocated between increasing the stock level and clearing the backorders. Nahmias and Demmy (1981) derive approximate service levels without taking the effects of clearing mechanisms on the system dynamics into the consideration. They totally ignore the clearing issue and analyze the system within a single replenishment leadtime window.

The natural way to perform the clearing is to employ the same threshold levels, which are used to control the demand traffic, for clearing the backorders. That is, the backorders for a certain customer class are not cleared until the inventory level reaches to the threshold level which is associated for that customer class. This mechanism is referred as the priority clearing in the literature. The related literature either ignores the clearing issue (Nahmias and Demmy (1981), Dekker et al. (1998) and Kocaga and Sen (2007)) or resorts to tractable clearing mechanisms that eliminate interaction between consecutive leadtime periods (Deshpande et al. (2003) and Arslan et al. (2007)). Deshpande et al. (2003) address this issue on page 684 of their study: “The optimal scheme is to always clear higher-priority customers first. However, this “priority-clearing” scheme is intractable because it does not allow closed-

Chapter 2 Literature Review

form expressions for the stockout levels, and average number of demand in backlog, for each demand class. To overcome this problem we introduce a tractable “threshold clearing” scheme to approximate the systems dynamics”.

In this thesis, in Chapter 5, we propose a dynamic rationing policy together with the associated dynamic clearing mechanism. We conduct a simulation study to evaluate the performance of the proposed policy. Since the analytical evaluation of the policy is not tractable without simplifying assumptions –our policy is a state-dependent one which also captures the dynamics of priority clearing--, simulation is the only available tool. In the existing literature, the only exception that provides analysis under priority clearing is the work of Fadiloglu and Bulut (2008), which will be detailed below together with the other cited works.

Dekker et al. (1998) consider the same setting with Nahmias and Demmy (1981) with $Q = 1$. Without specifying any clearing mechanism, they derive the exact fill rate expression for the non-critical demand class and make an approximation for the critical class fill rate by conditioning on the time that stock level hits the critical level. They test their approximation under three different clearing mechanisms using simulation. Kocaga and Sen (2007) extend the approximation of Dekker et al. (1998) to accommodate a demand lead-time for non-critical orders. They conduct a simulation study to assess the performance of the approximation

Without any restriction on the number of outstanding orders, Deshpande et al. (2003) work on the stock rationing problem for the same setting that Nahmias and Demmy (1981) consider. They introduce the threshold clearing mechanism that allows clearing low priority backorders before clearing all class 1 backorders and raising the inventory above the threshold rationing level. Threshold clearing makes it possible to obtain close form expressions for the desired performance measures of the system. Deshpande et al. (2003) derive the expected system cost under threshold clearing and provide an algorithm to obtain the policy parameters that minimizes this

Chapter 2 Literature Review

cost. It is interesting that their analysis yields the same results with Dekker et al. (1998) when $Q=1$. They simulate the rationing policy with priority clearing mechanism to compare the performance of threshold clearing with the performance of priority clearing. They also propose a lower bound on the cost of the unknown optimal policy. In Chapter 5, we propose another lower bound and show with numerical examples that our bound is tighter than the bound of Deshpande et al. (2003).

Zhao et al. (2005) analyze a game theoretical model of a decentralized dealer network in which each dealer can share its inventory with the others. They use the threshold clearing mechanism that Deshpande et al. (2003) introduce. Arslan et al. (2007) analyze the multiple demand-classes extension of the same setting Nahmias and Demmy (1981) and Deshpande et al. (2003) consider. They construct an equivalent multi stage serial system of the original single location system. For the ease of analysis, they assume that in each stage the backorders are cleared in the order of occurrence. Thereby, they derive approximate results for the priority clearing mechanism. Their results are exact under the clearing mechanism introduced by Deshpande et al. (2003).

Fadiloglu and Bulut (2008) also consider the same setting. They propose a method which captures the priority clearing dynamics for continuous-review inventory systems with backordering under static rationing policy. They assume two demand classes with Poisson arrivals and constant lead-time. They sample the continuous system at multiples of the lead time and show that the state of the system evolves according to an embedded Markov chain. They provide a recursive procedure to obtain the transition probabilities of the embedded chain and obtain the steady-state probabilities of interest with desired accuracy by considering a truncated version of the chain. This is the only work in the literature that analyzes the static rationing policy under priority clearing mechanism.

Chapter 2 Literature Review

Dekker et al. (2002) consider a lot-per-lot continuous-review setting with the static rationing in a lost sales environment. The clearing mechanism is not relevant in the lost sales case. They provide exact expressions for service levels under general stochastic lead-time and multiple demand classes. Their results are adapted from the analysis of $M/G/\infty$ queue under state-dependent arrival rates. Melchiors et al. (2000) also analyze the lost sales case with static rationing. They assume (Q, r) replenishment policy, deterministic leadtimes, and at most one outstanding order. This analysis is exact when $r < Q$.

There are only two studies in the existing literature that consider dynamic rationing policies for continuous review systems with exogenous leadtimes. Melchiors (2003) extends the work of Melchiors et al. (2000) by considering a time remembering rationing policy that allows different threshold rationing levels for different time slots between the placement of an order and its arrival. Threshold levels are set according to the age of the outstanding order. This is a restricted dynamic policy because threshold levels are assumed to be constant over predetermined time intervals. However, the unknown optimal policy should allow the threshold levels change at any point in time. Teunter and Haneveld (2008) also consider a time remembering policy for the backordering case. They aim to determine the set of critical remaining lead-time values (L_1, L_2, \dots) for the rationing decision. If the remaining lead-time is less than L_1 they do not ration the stock, if it is between L_1 and $(L_1 + L_2)$ one item is reserved for the high priority class and so on. Under the at-most-one-outstanding-order assumption, they approximate the optimal critical remaining lead-time values. Using two examples, they demonstrate that the dynamic policy outperforms the static policy.

The study presented in Chapter 5 of this thesis provides an extension to the analysis of the dynamic rationing policies. We propose a new dynamic policy and discuss the conditions under which stock rationing – static or dynamic—is beneficial

Chapter 2 Literature Review

and assess the value of the proposed policy. We also compare our results with the numerical results of Teunter and Haneveld (2008) and show that our policy outperforms theirs.

The other stream of studies that is related to our PhD study considers the problem of production control in make-to-stock production systems. The studies of this stream all assume Poisson demand arrivals and exponential or Erlangian processing times. The rationale behind these assumptions is twofold. First one is for the ease of analysis. Considering memoryless inter-demand-arrival and process stage completion times (Erlangian processing times are generalizations of exponential processing times since they are composed of exponential stages) enable to characterize the evolution of the system without keeping track of the history of the state variable(s). That is, the optimal policy for the system under consideration is a Markovian one. Second, by changing the number of Erlang stages, a wide range of systems with different processing time variances can be analyzed since the variance is decreasing in the number of Erlang stages. Exponential processing times have single stages and deterministic processing times are the limit of Erlangian processing times as the number of stages goes to infinity. Most of the studies of this stream also address the stock rationing problem.

One of the other commonalities of the studies that consider production control in make-to-stock systems is about the modeling perspective. In almost all the studies, the capacitated production system is modeled using only a single channel. The considered settings either constitute $M/M/1$ or $M/E_k/1$ make-to-stock queue models depending on the distribution of the processing times. Characterization of optimal production policies for multi-channel systems is for the first time provided in chapters 3 and 4 of this thesis. We allow arbitrary number of processing channels (servers) and develop $M/M/s$ and $M/E_k/s$ models in Chapter 3 and 4, respectively. The only analysis that we have come across for multiple replenishment channels with

Chapter 2 Literature Review

stochastic leadtimes is presented in Zipkin (2000) for the $M/G/s$ system with a single customer class and lost sales. Zipkin provides a performance analysis for the system under –suboptimal-- simple base-stock policy. Multi-channel settings are also identified by Buzacott and Shantikumar (1993), but again base-stock policy is proposed. In this thesis, while characterizing the optimal policy, we show that base-stock policy is not optimal.

All the analyses for make-to-stock production systems, including ours, rely on the uniformization technique proposed by Lippman (1975). Lippman shows that a Continuous Time Markov Chain (CTMC) is statistically equivalent to another CTMC that is obtained from the original chain by allowing fictitious self-transitions to obtain a uniform, i.e., state-independent, transition rate. Here, statistical equivalence means that starting from an initial state the probability that the process will be in a specific state at time t is identical for the original process and its uniform version. Uniformization can be applied by choosing an arbitrary uniform rate that is greater than all the rates of the original exponential transition times. Once this result is on hand, continuous-time Markovian control problems can be converted to equivalent discrete-time problems (to Markov Decision Processes) as outlined in Volume II-Chapter 5 of the book of Bertsekas (2000).

Ha (1997a) is the first to study the production control problem in continuous-time systems with capacitated replenishment channel. He considers a make-to-stock production facility with a single exponential server, zero setup cost, multiple demand classes and lost sales. He formulated the queueing control model using a single state variable that keeps track of the inventory level. He shows that base-stock policy is optimal for production control and static threshold level policy is optimal for stock rationing. Ha (1997a) also performs a stationary analysis of the system with two demand classes and distinguish the cases where the optimal average cost is convex.

Chapter 2 Literature Review

In a numerical study, he also compares the performance of the optimal rationing policy with the performance of FCFS policy under different scenarios.

Ha (1997b) analyzes the same problem in a backordering environment with two customer classes. Since it is possible to have both on-hand inventory and class 2 backorders at the same time, he defines the state of the system with two variables: the inventory level, negative part of which corresponds to the number of class 1 backorders, and the number of class 2 backorders. He characterizes optimal policies by a single monotone switching curve; optimal production decision is determined by a base stock policy and optimal rationing decision is determined by a rationing level that is decreasing in the number of backorders of class 2. Vericourt et al. (2002) extends the work of Ha (1997b) to multiple demand classes. They provide the full characterization of the optimal rationing policy and present an efficient algorithm to compute the optimal policy parameters, i.e., the optimal rationing levels for all classes.

Ha (2000) and Gayon et al. (2009b) consider Erlangian production times in the lost sales and backordering environments, respectively. Ha (2000) shows that for the lost sales case the optimal production and rationing policies can be characterized with a single-state variable called work storage level. Work storage level is the number of completed Erlang stages for the items in the system and it captures all the information regarding to the inventory level and the status of the production. Ha (2000) proves that optimal production and rationing policies are threshold work storage level policies. It is optimal to produce until a target work storage level (a base-stock policy) and for each customer class, there exists a threshold work storage level such that it is optimal to satisfy the demand of this class only above this level. On the other hand, in addition to the work storage level, Gayon et al. (2009b) also keep track of the number of backorders from each class. For the setting that Gayon et al. (2009b) consider, the full characterization of the optimal policies is problematic be-

Chapter 2 Literature Review

cause of the curse of dimensionality. Therefore, they provide a partial characterization. However, they succeed to show that optimal rationing policy is a work-storage type policy when excess production can be diverted to a salvage market.

Huang and Iravani (2008) extend the findings of Ha (1997a) to a random batch demand setting. They also consider the backorder case with two demand classes and fixed (deterministic) order sizes. Mayorga et al. (2006) consider a two-class make-to-stock system with adjustable exponential service rate. They study both the finite and infinite horizon problems. They show that optimal service rate adjustment and rationing policies are monotone in the inventory and backorder levels. They also prove that the optimal production policy turns into a bang-bang type policy when the production cost is a concave function of the production rate. That is, it is optimal to produce with full capacity (with the highest available service rate) until a threshold inventory level is reached and then leave the server idle (set the service rate to zero).

There are other works in the literature that consider settings in which other sources of information such as advanced demand and assembly component inventory levels are available. For these settings, additional information is incorporated to the state definition and it is shown that the optimal policies are state-dependent. Iravani et al. (2007) study the production control and rationing problem when advance demand information is available. They prove that the optimal policies are monotone with respect to the critical customer's order quantity, which is a random variable. Gayon et al. (2009a) also consider a setting with advance demand information. However, in their case, the advance demand information is imperfect because the customers who announce their orders ahead of the due date may decide to order prior to or later than the due date or to cancel the order altogether. In addition to the production decision, the controller must specify whether or not to satisfy an order that becomes due from on-hand inventory. Gayon et al. (2009a) show that the optimal production and rationing policies are monotone in the number of announced orders. Benjaafar

Chapter 2 Literature Review

and ElHafsi (2006) analyze a single-item, multi-component assemble-to-order system. They assume that the assembly is instantaneous and each assembly unit requires one unit of each component. If any of the components is out of stock, demands to the assembly are lost. The authors show that the optimal policies for each component, i.e., a state-dependent base-stock and a state-dependent threshold rationing policy, are monotone in the inventory level of all the other components.

Apart from the control literature on production-inventory systems, there is also a vast literature that considers queueing control problems, which involve mechanisms such as admission control, capacity control and pricing. Queueing control problems find application in the areas of service, telecommunication, and make-to-order manufacturing systems. We direct the reader to the recent works of Cil et al. (2009) and Gans and Savin (2007) for the related literature. We would like to point out that multiple-exponential-server models are also used for service systems as in Gans and Savin (2007).

All the above mentioned studies in production environments assume a single server replenishment channel. Our works presented in Chapter 3 and 4 are mostly related to the works of Ha (1997a) and Ha (2000), respectively. We extend their studies to the multi-server cases. We characterize the properties of the optimal production and rationing policies and show that they are state-dependent and monotone in the state variables.

In this thesis, we work on the continuous time control of production-inventory systems. Hence, we would like to conclude this chapter with Table 2.1 which provides the list of main works that consider the problem of production control and/or stock rationing in continuous-time. In the table, we classify the studies on the basis of the shortage cost structure and the capacity of the processing (replenishment) channel. While the classical inventory systems assume uncapacitated replenishment

Chapter 2 Literature Review

channels, i.e., infinitely many parallel servers, the production systems are capacitated and all the works in the literature handle this capacity constraint with a single server.

Table 2.1 Related Literature on Continuous-Time Systems

	Stock Rationing in Continuous-Review Inventory Systems	Production Control and Stock Rationing in Make-to-Stock Production Systems
Backordering	<p>Nahmias and Demmy (1981) Dekker et al. (1998) Moon and Kang (1998) Deshpande et al. (2003) Melchiors (2003) Zhao et al. (2005) Kocaga and Sen (2007) Arslan et al. (2007) Teunter and Haneveld (2008) Fadiloglu and Bulut (2008)</p>	<p>Ha (1997b) Vericourt et al. (2002) Gayon et al. (2006) Mayorga et al. (2006) Gayon et al. (2009b)</p>
Lost Sales	<p>Melchiors et al. (2000) Dekker et al. (2002) Melchiors (2003)</p>	<p>Ha (1997a) Ha (2000) Benjaafar and ElHafsi (2006) Iravani et al. (2007) Hung and Iravani (2008) Gayon et al. (2009a)</p>

Chapter 3

M/M/s Model with multiple-demand classes and lost sales

In this chapter, we characterize structural properties of the optimal cost function, and the optimal production and stock rationing policies for a single-item, multi-exponential-server make-to-stock production system with multiple-customer-classes and lost sales. We assume that each customer class generates demand according to a stationary Poisson process independent of the other classes, and the servers – parallel production channels-- have independent exponential processing times with identical rates. In effect, we model the production system as an $M/M/s$ make-to-stock queue.

To the best of our knowledge, with the exception of Zipkin (2000), there is no work in the literature that considers production control and/or stock rationing in a multi-server make-to-stock production system. Our work is most related to the work of Ha (1997a). We extend his study to the multi-server case using a two-dimensional state space. Along with a characterization of the optimal rationing policy, we also provide properties of the optimal production control policy and show that the optimal policy is not a base-stock policy for the multiple-servers case.

Section 3.1 introduces our primary model and, provides the dynamic programming formulation and the characterization of optimal policies under discounted cost

criterion. Section 3.2 discusses variations on the primary model including the existence of order-cancellation flexibility and the fixed costs. In Section 3.3, we consider average cost criterion then in Section 3.4 provide the stationary analysis of the system under both base-stock and bang-bang policies. Section 3.5 is devoted to a numerical study in which we quantify the benefit of the optimal policy and assess the value of order-cancellation flexibility. The chapter concludes with an Appendix (Section 3.6) that includes the proofs of the lemmas and theorems presented in the chapter.

3.1 Primary Model

3.1.1 Model Formulation

Consider a single-item make-to-stock production system with s identical servers having exponential production times with mean $1 / \mu$. Demand is generated by $n \geq 2$ customer classes according to independent Poisson processes with rates λ_i , $i \in \{1, 2, \dots, n\}$. We suppose that any unmet demand is lost and a lost sales cost of c_i is incurred for each unit of class i demand that is lost. Without loss of generality, we assume that $c_1 \geq c_2 \dots \geq c_n$. Let h be the inventory holding cost rate, p be the production cost rate and α be the discount rate. This setting can also be applied to a retailer environment where the retailer has s identical suppliers and can order single units from any of them. For such an inventory setting, it would be more appropriate to assume that p is zero.

The state of the system is defined with two variables. Let $X(t)$ be the inventory level at time t and $Y(t)$ be the number of the operational servers at the time of the last event occurrence prior to time t . The events are production completion and demand arrivals for each customer class. $Y(t)$ can also be considered as the number of the

outstanding replenishment orders at the time of last event occurrence prior to time t . This elaborate state definition is necessary to eliminate instantaneous state transitions at decision epochs, which causes problems at the application of the uniformization technique.

At any time point t , the control specifies whether to keep the number of active servers at the same level or to increase it. We denote the production decision at time t as $u_p(t)$ where $u_p(t) \in \{Y(t), Y(t)+1, \dots, s\}$. When a class i demand arrives at time t , the control specifies whether to satisfy the demand or not. We denote the rationing decision for class i at time t as $u_{r_i}(t)$ such that $u_{r_i}(t) \in \{0, 1\}$, $i \in \{1, 2, \dots, n\}$. If $u_{r_i}(t) = 0$, an arriving i^{th} class demand is rejected, otherwise it is satisfied. The complete policy for our model can be represented as $\{(u_p(t), u_{r_1}(t), \dots, u_{r_n}(t)) | t \geq 0\}$. Since the model is Markovian, optimal policies are also Markovian. Thus, it is sufficient to consider the set of admissible Markovian policies, i.e., $u_p(t) = u_p(X(t), Y(t))$ and $u_{r_i}(t) = u_{r_i}(X(t), Y(t))$, $i \in \{1, \dots, n\}$.

Starting at state (x, y) , under the policy π , the infinite horizon expected discounted system cost is

$$E_{(x,y)}^\pi \left[\int_0^\infty e^{-\alpha t} (h(X^\pi(t)) + p(Y^\pi(t))) dt + \sum_{i=1}^n \int_0^\infty e^{-\alpha t} c_i dN_i^\pi(t) \right] \quad (1)$$

where $N_i^\pi(t)$ be the number of class i customers who have been rejected up to time t , $i \in \{1, 2, \dots, n\}$. Given a control policy π , the process $\{(X^\pi(t), Y^\pi(t)) | t \geq 0\}$ is a continuous time Markov chain where the transition rate at state (x, y) is

$$V_{(x,y)} = \sum_{i=1}^n \lambda_i u_{r_i} + u_p \mu.$$

Using the uniformization technique proposed by Lippman (1975) we can analyze our model within the discrete-time framework. Uniformization is a conversion method that provides a discrete-time version of the original continuous-time optimization model. The conversion is achieved by first constructing a new Markov process that is statistically equivalent to the original one. This new process should have a uniform transition rate, which is same for all states and controls, and allow fictitious self transitions. The uniform rate can be anything that is greater than or equal to the maximum of the transition rates of the original process. Here statistical equivalence means that for any given policy π , initial state $(X(0), Y(0))$, time t , and state (x, y) $P\{X^\pi(t) = x, Y^\pi(t) = y | X(0), Y(0)\}$ is identical for the original process and its uniform version. Once the uniform process is obtained, continuous-time control problem can be easily converted to an equivalent discrete-time problem as explicitly explained in Volume II-Chapter 5 of the book of Bertsekas (2000).

Let us define the uniform rate as $\nu = \sum_{i=1}^n \lambda_i + s\mu$. Without loss of generality, we rescale the time and assume that $\alpha + \nu = 1$. Then, the optimal cost-to-go function can be expressed as

$$J(x, y) = \min_{s \geq u \geq y} \{hx + pu + (s - u)\mu J(x, y) + u\mu J(x + 1, u - 1) + T_R(x, u)\} \quad (2)$$

where $T_R(x, y) = \sum_{i=1}^n T_{R_i}(x, y)$ and for $i \in \{1, 2, \dots, n\}$,

$$T_{R_i}(x, y) = \begin{cases} \lambda_i \min \{J(x - 1, y), c_i + J(x, y)\} & , x > 0 \\ \lambda_i (c_i + J(0, y)) & , x = 0 \end{cases} \quad (3)$$

In (2), the minimization operation corresponds to the production decision, i.e., deciding the number of operational servers when there are x units on hand and y servers

are operational. The term $(s-u)\mu J(x, y)$ corresponds to the fictitious self-transitions due to uniformization, while the term $u\mu J(x+1, u-1)$ corresponds to production completion event that takes place at one of the u active production channels with probability $u\mu$. The minimization operator $T_R(x, y)$ corresponds to the rationing decision for class i . With probability λ_i a demand arrival event for class i occurs and at this decision epoch the controller should decide whether to satisfy or reject the arriving demand. At the boundary, when there is no stock on-hand, all the arriving demands are lost.

For notational purposes, we provide equations (4), (5) and (6) below. In (5), $u^*(x, y)$ is defined as the optimal number of operational production channels for the given state (x, y) . Equation (6) defines a base-stock level for each inventory level x .

$$f(x, y, u) = hx + pu + (s-u)\mu J(x, y) + u\mu J(x+1, u-1) + T_R(x, u) \quad (4)$$

$$u^*(x, y) = \arg \min_{s \geq u \geq y} f(x, y, u) \quad (5)$$

$$S_x = x + u^*(x, 0) \quad (6)$$

We also define the following operators on a function $v(x, y)$:

$$\Delta^y v(x, y) = v(x, y+1) - v(x, y)$$

$$\Delta^x v(x, y) = v(x+1, y) - v(x, y)$$

3.1.2 Characterization of the Optimal Production and Rationing Policies

This section provides a detailed characterization of the optimal production and rationing policies in three theorems and a corollary. The three theorems are proven via the methodology formalized by Porteus (1982). This approach is based on identifying a set of structural properties and then showing that these properties are preserved under the optimization operator. For the *M/M/s* model described in the previous section, the optimization operator is

$$T(J(x, y)) = \min_{s \geq u \geq y} f(x, y, u). \quad (7)$$

We define \mathcal{G} as a set of functions on the integers such that if $v \in \mathcal{G}$, then

$$\Delta^x v(x, y+1) \geq \Delta^x v(x, y) \quad (8)$$

$$\Delta^x v(x, y) \geq \Delta^x v(x-1, y+1) \quad (9)$$

Note that, (8) can also be written as $\Delta^y v(x+1, y) \geq \Delta^y v(x, y)$.

In Theorem 3.1 and its corollary, we characterize the behavior of the optimal cost function and the optimal production policy with respect to the number of operational servers.

Theorem 3.1. If $J \in \mathcal{G}$, for given inventory level x ,

- i. $J(x, 0) = \dots = J(x, u^*(x, 0))$
- ii. For $y \geq u^*(x, 0)$, $J(x, y)$ is a convex-increasing function of y . That is,

$$\Delta^y J(x, y+1) \geq \Delta^y J(x, y) > 0.$$

Corollary 3.1. For given state (x, y) ,

$$\text{i. } u^*(x, y) = \begin{cases} u^*(x, 0), & y \leq u^*(x, 0) \\ y, & y > u^*(x, 0) \end{cases}$$

$$\text{ii. } u^*(x, y) = \min \{y' : J(x, y'+1) - J(x, y') > 0, s \geq y' \geq y\}$$

$$\text{iii. } u^*(x, y+1) = \begin{cases} u^*(x, y), & u^*(x, y) \geq y+1 \\ y+1, & u^*(x, y) = y \end{cases}$$

Theorem 3.1 implies that the optimal cost function is constant with respect to the number of operational servers in the region where the number of operational servers is less than or equal to the optimal number of operational servers at state $(x, 0)$. On the other hand, in the complementary region, the optimal cost function is convex-increasing in the number of operational servers.

The first part of Corollary 3.1 indicates that if the current inventory position, $x + y$, is less than the base-stock level S_x , then it is optimal to increase the number of operational servers to $u^*(x, 0)$ in order to raise the inventory position to the base-stock level. Otherwise, it is optimal not to change the number of operational production channels. It is optimal to set the number of operational servers to $u^*(x, 0)$ when it is possible, i.e., the number of currently operational servers is less or equal to $u^*(x, 0)$. Hence, the optimal costs for all the states in which $u^*(x, 0)$ are feasible are the same as stated in the first part of Theorem 3.1.

As the numerical study in the next section illustrates, $u^*(x, 0) = u^*(x+1, 0) + 1$ does not hold in general. Consequently, a single order-up-to level that is independent of the inventory position is not optimal and the optimal production policy is a state-dependent base-stock policy. As stated in the literature (Erhardt, 1984), the op-

tinality of a simple base-stock policy cannot be guaranteed when replenishment orders cross in time. In our model, order crossing is possible due to parallel production channels. In the single server case, order crossing does not occur since stock-units are produced one by one. In this case, base-stock policy is optimal as shown by Ha (1997a).

The second part of Corollary 3.1 provides an alternative definition for the optimal number of operational servers. It is optimal to increase the number of operational servers until the optimal cost function starts to increase. Finally, the last part of the corollary exhibits how the optimal number of operational channels changes with the number of currently operational servers. The optimal number of servers at state $(x, y + 1)$ is either equal to the optimal number of servers at state (x, y) or one more.

In Theorem 3.2, we characterize the behavior of the optimal cost function, and the optimal production and rationing policies with respect to the inventory level. We also characterize the effect of the number of operational servers on the optimal rationing policy.

Theorem 3.2. If $J \in \mathcal{G}$, then

- i. $J(x, y)$ is convex in x .
- ii. $u^*(x + 1, y) \leq u^*(x, y)$.
- iii. $\Delta^x J(x - 1, y) = J(x, y) - J(x - 1, y) \geq -c_1$, and so $T_{R_1}(x, y) = \lambda_1 J(x - 1, y)$.

That is, it is always optimal to satisfy a class 1 demand when there is stock on hand.

- iv. There exists a threshold inventory level $K_x^i(y)$ for class $i \geq 2$, which is a function of operational servers, y , such that it is optimal to satisfy a class i

demand above $K_x^i(y)$ and reject it otherwise. Moreover, $K_x^n(y) \geq K_x^{n-1}(y) \geq \dots \geq K_x^2(y) \geq 0$, and $K_x^i(y+1) \leq K_x^i(y)$ for $i \in \{2, \dots, n\}$.

- v. There exists a threshold number of operational servers $K_y^i(x)$ for class $i \geq 2$ as a function of inventory level such that it is optimal to satisfy a class i demand above $K_y^i(x)$ and reject it otherwise. Moreover, $K_y^n(x) \geq K_y^{n-1}(x) \geq \dots \geq K_y^2(x) \geq 0$, and $K_y^i(x+1) < K_y^i(x)$ for $i \in \{2, \dots, n\}$.

Theorem 3.2 states that the optimal cost function is x -convex and the optimal number of production channels that should be used is non-increasing in the inventory level. The last three parts of the theorem characterize the optimal rationing policy, which is of threshold type. If there is stock on hand, it is always optimal to satisfy an arriving class 1 demand independent of the observed state. For each of the other customer classes, given the number of the operational servers, there exists a rationing inventory level, which is non-decreasing in the class index. Similarly, for each class, given the inventory level, there exists a rationing level for the number of operational servers, which is non-decreasing in the class index. Furthermore, the rationing inventory levels are non-increasing in the number of operational channels, and the rationing levels for the number of operational production channels are decreasing in the inventory level. The latter statement means that if it is optimal to satisfy an arriving class i demand at state $(x, y+1)$, then it is optimal to satisfy an arriving class i demand at state $(x+1, y)$. Moreover, a class i demand arriving at state $(x, K_y^i(x))$ should be rejected, but it is optimal to satisfy an arriving class i demand at state $(x+1, K_y^i(x))$.

Theorem 3.3 states that $J \in \mathcal{G}$, that is, the optimal cost function is an element of the function space characterized by (8) and (9). Since $J \in \mathcal{G}$ is the hypothesis of the previous two theorems, Theorem 3.3 is needed to ensure that the results of the previous two theorems apply to our model without any restriction.

Theorem 3.3. $J \in \mathcal{G}$, that is

- i. $\Delta^x J(x, y+1) \geq \Delta^x J(x, y)$
- ii. $\Delta^x J(x, y) \geq \Delta^x J(x-1, y+1)$

3.2 Variations on the Primary Model

3.2.1 Model with Full Order-Cancellation Flexibility

In this section, we consider a variation on the previous model in which cancellation of all previously placed production orders is permitted. The rationale behind this model is twofold. Firstly, this model enables us to characterize the optimal policy for make-to-stock queues where the outstanding orders can be cancelled at a negligible cost. Secondly, this model permits us to quantify the value of the full flexibility to cancel orders. The difference between the performances of the primary model and this model is the value of the full order-cancellation flexibility. In many systems, order cancellations are only possible at a cost. In such cases the cost of canceling orders should be compared with the value of order-cancellation flexibility.

Given that order cancellation is possible, at each decision epoch the number of operational servers can be chosen from the set $\{0, 1, \dots, s\}$. As previously discussed, for the primary model the feasible set is $\{Y(t), Y(t)+1, \dots, s\}$ where $Y(t)$ is the number of operational servers at the time of the last event occurrence prior to time t . There-

fore, for the model with order cancellation, there is no need to keep track of the number of operational servers and it is possible to model the system evolution with a single state variable, which is the inventory level.

For this model, the optimal cost-to-go function of this model can be expressed as

$$J(x) = \min_{s \geq u \geq 0} \{hx + pu + (s - u)\mu J(x) + u\mu J(x + 1)\} + T_R(x) \quad (10)$$

where $T_R(x) = \sum_{i=1}^n T_{R_i}(x)$ and for $i \in \{1, 2, \dots, n\}$,

$$T_{R_i}(x) = \begin{cases} \lambda_i \min(J(x-1), c_i + J(x)) & , x > 0 \\ \lambda_i (c_i + J(0)) & , x = 0 \end{cases} \quad (11)$$

Let $u^*(x)$ be the optimal number of operational production channels at state x .

Then,

$$u^*(x) = \arg \min_{s \geq u \geq 0} \{u(p + \mu\Delta J(x))\} \quad (12)$$

It should also be noted that for $s = 1$ the model is the same with the one analyzed in Ha (1997a). Thus, the below theorem that provides the properties of the cost function and the optimal policy extends the results presented in Ha (1997a) to a multiple-servers setting.

Theorem 3.4.

- i. $u^*(x)$ is either s or 0
- ii. $J(x)$ is a convex function of x

- iii. There exists a threshold inventory level K^i for class i such that it is optimal to satisfy a class i demand above K^i and reject it otherwise.
- iv. For $x > 0$, $\Delta J(x-1) \geq -c_1$, and so $T_{R_i}(x) = \lambda_i J(x-1)$. That is, $K^1 = 0$ and it is always optimal to satisfy a class 1 demand when there is stock on hand.
- v. Optimal production policy is a bang-bang type policy. That is, there exists a threshold inventory level \bar{x} such that $u^*(x) = s$, for $x \in \{0, 1, \dots, \bar{x}\}$ and $u^*(x) = 0$, for $x \in \{\bar{x} + 1, \bar{x} + 2, \dots\}$.

Theorem 3.4 states that the optimal production policy when order cancellations are permitted is a bang-bang policy. Up to a certain inventory level the policy prescribes using all available production channels. Beyond that level all of the production channels are idled. A static stock-rationing policy is employed for stock allocation. The rationing levels are threshold inventory levels and fixed, because we characterize the system only with the inventory level.

3.2.2 Model with Partial Order-Cancellation Flexibility

In this section, we consider a variation on the primary model in which cancellation of only a limited number of previously placed production orders is permitted. This model allows us to quantify how much value can be captured via order cancellation, when there is a limitation on the number of orders that can be cancelled. Although a manufacturer may desire to reduce its operating costs by cancelling some orders, it may not be willing to use this flexibility in an indiscriminate fashion. The optimal policy under full flexibility is a bang-bang policy that utilizes all available production channels until the inventory reaches a threshold level. When one of the orders is completed at that level, the rest of the orders have to be cancelled. For many manufacturers, this drastic cancellation practice would not be desirable. If it is possible to

capture most of the value available via order cancellation while espousing a smoother production policy that involves fewer cancellations, this could be the avenue of choice for those manufacturers.

We define the cancellation flexibility index f as the maximum number of servers that can be shut down. At each decision epoch, the control selects the number of active servers from the set $\{(Y(t) - f)^+, \dots, s\}$ where $Y(t)$ is the number of active servers at the time of the last event occurrence prior to t . This model has the versatility to cover both the primary model and the model with full-cancellation flexibility. The primary model and the model with full-cancellation flexibility can be obtained by setting $f = 0$, and $f = s$, respectively. The optimal cost-to-go function for this model is a straight-forward extension of the one for the primary model given in (2) and can be expressed as

$$J(x, y) = \min_{s \geq u \geq (y-f)^+} \{hx + pu + (s - u)\mu J(x, y) + u\mu J(x + 1, u - 1) + T_r(x, u)\} \quad (13)$$

The optimal policy for the model with partial order-cancellation flexibility fully conforms to the characterization provided for the primary model.

3.2.3 Model with Fixed Production and Order Cancellation Costs

We now consider a setting under the existence of fixed costs: the fixed costs of activating an idle production channel (setup) and deactivating a busy channel (order cancellation). Let us define w_1 and w_2 as the setup and cancellation costs, respectively. Then, the optimal cost function can be expressed as

$$J(x, y) = \min_{s \geq u \geq 0} \left\{ \begin{array}{l} hx + pu + w_1(u - y)^+ + w_2(y - u)^+ + (s - u)\mu J(x, y) \\ + u\mu(\min\{J(x + 1, u - 1), J(x + 1, u)\}) + T_R(x, u) \end{array} \right\} \quad (14)$$

In (14), it is possible to choose any number of active servers, i.e., $s \geq u \geq 0$, because we allow order cancellation. In addition to this flexibility and the fixed cost terms, compared to the cost function of the primary model, (14) has one more minimization operator. The new operator, $\min\{J(x + 1, u - 1), J(x + 1, u)\}$, ensures that the activation cost is not paid if it is optimal to continue keeping the channel active at which production of a job has just finished.

When $w_1 = 0$ and w_2 is very large (goes to infinity), (14) gives the same results with the primary model. And obviously if we set $w_1 = w_2 = 0$, we get the model with full-order-cancellation flexibility. For the other cases, the optimal production policy is not monotone in the number of operational channels, because $(w_1(u - y)^+ + w_2(y - u)^+)$ is not a monotone function of y .

3.3 Average Cost Criterion

This section provides the analysis of the primary model and its variations with partial and full order-cancellation flexibility under the infinite horizon expected average cost criterion. Let $g^\pi(x, y)$ be the average cost function under policy π starting from state (x, y) . Then,

$$g^\pi(x, y) = \lim_{\tau \rightarrow \infty} \frac{E_{(x,y)}^\pi \left[\int_0^\tau (h(X^\pi(t)) + p(Y^\pi(t))) dt + \sum_{i=1}^n \int_0^\tau c_i dN_i^\pi(t) \right]}{\tau} \quad (15)$$

The following theorem characterizes the structural properties of the optimal control policy under the average cost criterion.

Theorem 3.5. Consider the primary model and its variations with partial and full order-cancellation flexibility. Under infinite horizon average cost criterion, for each of these models:

- i. There exists an optimal stationary policy and the optimal average cost is finite and independent from the initial state.
- ii. The optimal stationary policy possesses all the structural properties of the optimal policy under the discounted cost criterion.


Theorem 3.5 implies that the optimal production and rationing policies that minimize expected average cost are also state-dependent base-stock and state-dependent threshold type policies, respectively (see Corollary 3.1 and Theorem 3.2). Moreover, the optimal production policy turns into a bang-bang policy under full order cancellation flexibility as in the case of discounted cost criterion (see Theorem 3.4).

We prove Theorem 3.5 using the results outlined by Cavazos-Cadena and Sennott (1992). In this work, the authors compare the conditions proposed by Borkar (1984), Borkar (1989), Weber and Stidham (1987), Cavazos-Cadena (1989), and Sennott (1989) for the existence of an average cost optimal stationary policy. All these works follow the approach of analyzing average cost models as “the limit of discounted models”. Here “the limit of discounted models” stands for the analysis of models as the discount rate goes to zero. Cavazos-Cadena and Sennott (1992) show that the conditions proposed by Sennott (1989) are the weakest, i.e., they are implied by the other condition sets. In the proof we opt to use the two necessary conditions by Cavazos-Cadena and Sennott that are equivalent to the larger set of conditions of Weber and Stidham (1987). These conditions, which are explicitly stated within the proof provided in the appendix, are the existence of a stationary policy that induces an irreducible and ergodic Markov chain with finite average cost, and the finiteness of the number of states for which stage cost is less than this average cost.

3.4 Stationary Analysis

3.4.1 Stationary Analysis under Base-stock

In Sections 3.1 and 3.3 we show that the optimal production and stock allocation policies are state-dependent under both discounted and average cost criteria. There is no fixed target inventory level and the rationing levels are dependent on the number of operational servers. Yet, if we operate under a simple base-stock policy with a fixed target level S , the number of operational servers is known for each inventory level, i.e., $\min\{(S-x)^+, s\}$. In this case, the stock allocation decision is solely determined by the inventory level. In this section, we provide the average cost of the system for this setting.


 Figure 3.1 Birth-Death Process for (S, \bar{K}) Policy

Let (S, \bar{K}) denote the base-stock policy with rationing whose fixed target inventory level is S and fixed threshold rationing levels are $\bar{K} = (K_1, K_2, \dots, K_n)$. Under (S, \bar{K}) policy the inventory level, x , evolves according to a Birth-Death Process depicted in Figure 3.1. Births correspond to production completions whose rate is $\mu \min\{(S-x)^+, s\}$; whereas deaths correspond to inventory depletion by demand whose rate is $\sum_{i=1}^m \lambda_i$ if $K_m < x \leq K_{m+1}$ for $m \in \{1, \dots, n\}$. Note that K_{n+1} is set to S for convenience of notation. The stationary probability that the inventory level is j where $K_m < j \leq K_{m+1}$ for $m \in \{1, \dots, n\}$ can be expressed as

$$P_j = \frac{\mu^j \prod_{i=0}^{j-1} \min(S-i, s)}{\left(\sum_{i=1}^m \lambda_i\right)^{j-K_m} \prod_{r=1}^{m-1} \left(\sum_{i=1}^r \lambda_i\right)^{K_{r+1}-K_r}} P_0 \quad (16)$$

where

$$P_0 = \left(1 + \sum_{j=1}^S \frac{\mu^j \prod_{i=0}^{j-1} \min(S-i, s)}{\left(\sum_{i=1}^m \lambda_i\right)^{j-K_m} \prod_{r=1}^{m-1} \left(\sum_{i=1}^r \lambda_i\right)^{K_{r+1}-K_r}} \right)^{-1}.$$

Hence, the expected inventory level, the fill rate for each class and the expected cost of the system are respectively

$$E[X] = \sum_{j=0}^S jP_j, \quad (17)$$

$$\beta_i = \sum_{j=K_i+1}^S P_j, \quad 1 \leq i \leq n \quad (18)$$

$$C(S, \vec{K}) = hE[X] + \sum_{i=1}^n \lambda_i c_i (1 - \beta_i). \quad (19)$$

In order to find the optimal policy parameters, for each fixed S , we can first find the optimal \vec{K} vector that minimizes $C(S, \vec{K})$ by performing an exhaustive search on each K_i (starting from $i = 2$) over $\{K_{i-1}, \dots, S\}$ for $2 \leq i \leq n$. Note that Theorem 3.2 states it is always optimal to satisfy a class 1 demand when there is stock on hand, therefore K_1 is set to 0. Let us denote $\vec{K}^*(S)$ as the vector of optimal rationing levels for a fixed S . Starting from $S = 0$ we can search for the optimal S , denoted as S^* , that minimizes $C(S, \vec{K}^*(S))$. We suggest to perform an extensive search on S while keeping track of the first difference of the expected cost function, $C(S+1, \vec{K}^*(S+1)) - C(S, \vec{K}^*(S))$, until $C(S, \vec{K}^*(S))$ is sufficiently larger than the current minimum and the first difference continues to remain positive over a large range. Ha (1997a) also proposes a similar algorithm for the case where $s = 1$ and shows that the cost function is not convex in general. The non-convexity result applies to our problem which is a generalization of Ha (1997a).

The discussion above outlines a general method for the analysis of the system under base-stock. We now consider a special case with infinitely many servers and

denote the optimal base-stock for this special case as S_{inf}^* . This would provide us a bound on the number of servers beyond which the system is equivalent to a system with exogenous leadtimes, i.e., uncapacitated replenishment channel. If the number of available servers is greater or equal than S_{inf}^* , then the system never utilizes more than S_{inf}^* servers because it is the optimal base-stock level for the problem with no constraint on s . We can also find an upper bound on the value of S_{inf}^* by considering the $(S, \vec{0})$ policy, i.e., the stock allocation is performed on FCFS basis. The optimal base-stock level for the $(S, \vec{0})$ policy would obviously constitute an upper bound on S_{inf}^* , since we do not ration the inventory and the effective demand increases. By letting $\lambda = \sum_{i=1}^n \lambda_i$, the performances measures for an $(S, \vec{0})$ policy can be expressed as

$$P_0 = \left(\sum_{j=0}^S \frac{S!}{(S-j)!} \frac{\mu^j}{\lambda^j} \right)^{-1}, \quad \beta = 1 - P_0, \quad E[X] = S - \frac{\lambda}{\mu} \beta \quad (20)$$

$$C(S, \vec{0}) = h \left(S - \frac{\lambda}{\mu} \beta \right) + (1 - \beta) \sum_{i=1}^n \lambda_i c_i \quad (21)$$

Lemma 3.1. $C(S, \vec{0})$ is a convex function of S .

Using Lemma 3.1, it is easy to find the optimal base-stock for the $(S, \vec{0})$ policy, which is the smallest S that satisfies $C(S+1, \vec{0}) - C(S, \vec{0}) \geq 0$. It is also interesting that the cost function, which is not convex for finite number of servers, becomes convex when the number of servers tends to infinity. The reader may also refer to Jaarsveld and Dekker (2009) for a discussion on different algorithms proposed in the

literature for finding S_{inf}^* and the corresponding optimal rationing levels. We also would like to point out that in this setting with ample servers the provided analysis is valid for general service times as well due to Palm's theorem (Takacs 1962, p.160).

3.4.2 Stationary Analysis under Bang-Bang Policy

In Section 3.2 (for the discounted cost objective) and Section 3.3 (for the average cost objective), we show that the optimal policy is of bang-bang type when order cancellations are allowed. In this section, we provide stationary analysis for the optimal bang-bang policy. Note that this policy is only possible when there is flexibility to cancel all outstanding orders.

The bang-bang policy is characterized with a threshold inventory level \bar{x} below which all servers are utilized. Once this target level is reached, all the servers are shut down. Under this policy, the inventory level, x , corresponds to the number of customers in an $M/M/1/\bar{x}$ queue where the arrival rate (the rate of production completion) is $s\mu$ at levels and the departure rate (the rate of inventory depletion) is $\sum_{i=1}^m \lambda_i$ when $K_m < x \leq K_{m+1}$ for $m \in \{1, \dots, n\}$. Here, K_{n+1} is set to \bar{x} for convenience of notation. The stationary probability of having j units of on-hand inventory where $K_m < j \leq K_{m+1}$ for $m \in \{1, \dots, n\}$ is

$$P_j = \frac{(s\mu)^j}{\left(\sum_{i=1}^m \lambda_i\right)^{j-K_m} \prod_{r=1}^{m-1} \left(\sum_{i=1}^r \lambda_i\right)^{K_{r+1}-K_r}} P_0 \quad (22)$$

where

$$P_0 = \left(1 + \sum_{j=1}^{\bar{x}} \frac{(s\mu)^j}{\binom{m}{i=1}^{j-K_m} \prod_{r=1}^{m-1} \binom{r}{i=1}^{K_{r+1}-K_r}} \right)^{-1}.$$

Hence, the expected inventory level, the fill rate for each class and the expected cost of the system are respectively

$$E[X] = \sum_{j=0}^{\bar{x}} jP_j, \quad (23)$$

$$\beta_i = \sum_{j=K_i+1}^{\bar{x}} P_j, \quad 1 \leq i \leq n \quad (24)$$

$$C(\bar{x}) = hE[X] + \sum_{i=1}^n \lambda_i c_i (1 - \beta_i). \quad (25)$$

In order to find the optimal policy parameters, the methodology described in Section 3.4.1 is directly applicable. Furthermore, the system under bang-bang policy is equivalent to the single server system of Ha (1997a) when $s\mu$ is set to be service rate. This is due to the fact that for single server systems base-stock and bang-bang policy are equivalent.

3.5 Numerical Study

In this section, we illustrate the results obtained in the previous sections, and quantify the impact of the system parameters on the performance of the optimal production and rationing policies. Under different scenarios, we compare the optimal production policies (for the primary model and its variations) with the base-stock policy, which is the one suggested in the literature for systems with a limited number of processing channels (see Zipkin (2000) pp. 261-263), and the optimal rationing policy with the first-come-first-served (FCFS) policy. We quantify the benefit of the optimal production and rationing policies as the percent cost reduction obtained by operating the system under the optimal policies instead of the base-stock and FCFS policies. We also present a graph illustrating the impact of the cancellation flexibility index on the performance of the optimal production policy. We obtain the numerical results presented in this section via a value iteration algorithm coded in MATLAB.

In order to illustrate the properties of the optimal policies, let us consider a two-class system with $(s, \mu, \lambda_1, \lambda_2, h, p, c_1, c_2) = (15, 1, 5, 1, 1, 1, 4, 1)$. For the discounted cost criterion with $\alpha = 0.6$, Table 3.1 and Table 3.2 show the optimal production and rationing policies, respectively. A continuous discount rate of α corresponds to a

periodic discount factor of $\nu/(\alpha + \nu) = \left(\sum_{i=1}^n \lambda_i + s\mu \right) / \left(\alpha + \sum_{i=1}^n \lambda_i + s\mu \right)$, once unifor-

mization is performed. Thus, the corresponding discount factor is 0.97. In the tables, rows indicate the on-hand inventory level (0 to 4) and columns indicate the current number of operational servers (0 to 15). In Table 3.1, the value corresponding to the state (x, y) is the optimal number of operational servers, $u^*(x, y)$, which is bounded below by y (the current number of operational servers) and above by s (the total number of available servers). Table 3.1 is in agreement with Corollary 3.1, which states, if $y \leq u^*(x, 0)$, then $u^*(x, y) = u^*(x, 0)$, otherwise $u^*(x, y) = y$. Moreover,

Chapter 3 $M/M/s$ Model With Multiple Demand Classes and Lost Sales

in parallel with Theorem 3.2, the optimal number of operational servers at state $(x, 0)$ decreases by one or more units, for each unit increase in the on-hand inventory level, and then it remains constant at 0. It is also observed that the base-stock level for each inventory level, $S_x = x + u^*(x, 0)$, varies with the inventory level. The reader should also note that the region $\{(x, y) | x \leq S_0, y \leq u^*(0, 0)\}$ is recurrent whereas all the other states in the state space are transient. Since it is optimal not to activate any of the servers beyond the target inventory level, which is the base-stock level at $x = 0$, even if the system starts to operate in the transient region it will definitely visit the recurrent region and never turns back.

Table 3.1 Optimal Production Policy under Discounted Cost Criterion

$x \backslash y$	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	10	10	10	10	10	10	10	10	10	10	10	11	12	13	14	15
1	6	6	6	6	6	6	6	7	8	9	10	11	12	13	14	15
2	1	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
3	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
4	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15

Table 3.2 Optimal Rationing Policy For Class 2 under Discounted Cost Criterion

$x \backslash y$	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1
2	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
3	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1
4	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1

Theorem 3.2 states that it is always optimal to satisfy an arriving class 1 demand as long as there is inventory on-hand. Therefore, we only provide the optimal rationing policy for class 2 (Table 3.2). A zero in the cell corresponding to state (x, y) indicates that an arriving class 2 demand should be rejected at state (x, y) , and a one indicates that the demand should be satisfied. As can be easily observed from the table, the threshold inventory rationing levels for class 2, which are non-increasing in the number of operational servers, are: $K_x^2(0) = K_x^2(1) = 3$, $K_x^2(2) = \dots = K_x^2(7) = 2$, $K_x^2(8) = \dots = K_x^2(14) = 1$, and $K_x^2(15) = 0$. The rationing thresholds for the number of operational servers, which are decreasing in the on-hand inventory level until hitting -1 as shown in Theorem 3.2, are: $K_y^2(1) = 14, K_y^2(2) = 7, K_y^2(3) = 1$, and $K_y^2(x) = -1$ for $x \geq 4$.

For the same setting, Table 3.3 and Table 3.4 illustrate the optimal policies under average cost criterion, i.e., $\alpha = 0$. As discussed in Section 3.3 the structure of the average cost optimal policies are the same with the ones exhibited in Tables 3.1 and 3.2.

Table 3.3 Optimal Production Policy under Average Cost Criterion

$x \backslash y$	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	9	9	9	9	9	9	9	9	9	9	10	11	12	13	14	15
1	6	6	6	6	6	6	6	7	8	9	10	11	12	13	14	15
2	3	3	3	3	4	5	6	7	8	9	10	11	12	13	14	15
3	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
4	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15

Table 3.4 Optimal Rationing Policy For Class 2 under Average Cost Criterion

$x \backslash y$	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1
3	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1
4	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1

In this study we primarily consider the infinite-horizon discounted cost as the objective function. Therefore, in the remaining part of the section, we compare the performance of different policies under this criterion. Since the optimal discounted cost depends on the initial state, the reader should note that the comparisons will be for state $(0,0)$.

In order to assess the value of the optimal production policy relative to the optimal base-stock policy, we suppress the effect of rationing and consider a setting with a single customer class by letting $(\lambda, \mu, h, p, c, \alpha) = (10, 2, 0.2, 0.2, 10, 0.6)$. In this setting, Figure 3.2 exhibits the effect of the number of available production channels s , on the cost reduction. As seen from the figure, the optimal policy does not provide any cost reduction for small values of s . This is due to the fact that when available processing channels are scarce, both of the optimal production and the optimal base-stock policies, try to use all of the limited capacity. For $s = 1$, Ha (1997a) already showed that the optimal policy is base-stock. However, for moderate values of s , the benefit of the optimal policy over the base-stock policy increases rapidly with s , because the optimal policy has flexibility to adjust the number of operational servers at each inventory level. In contrast, the base-stock policy dictates a fixed production target for all states. The percent cost reduction stabilizes after s exceeds the

optimal number of operational processing channels at state $(0,0)$, which is 23 for the considered setting. When $s > 23$, the system is equivalent to a system with exogenous exponential leadtimes, i.e., uncapacitated replenishment channel, because it always operates with less than the available s processing channels and therefore additional channels do not provide further gain.

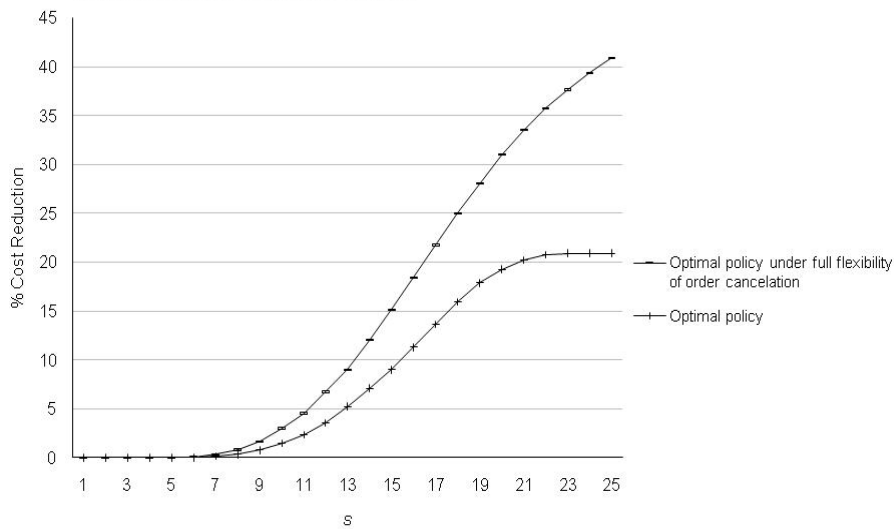


Figure 3.2 Optimal production policy vs. Base-stock policy: Impact of number of available servers

When we allow cancellation of previously placed production orders, the system operates under full flexibility. For small values of s , all optimization models (base-stock, primary, and models with order cancellation) try to use the whole capacity and, thereby, their performances are indistinguishable. For larger values of s , the savings obtained by using the optimal policy under order cancellation instead of using the best base-stock policy, or the optimal policy for the primary model, is positive and grows until hundred-percent with s . This is due to the fact that the optimal policy under order cancellation is a bang-bang type policy (as discussed in the pre-

vious section). As s tends to infinity, it is optimal to utilize all available servers at the time of each demand arrival in order to satisfy the arriving demand instantaneously (because the replenishment rate tends to infinity) and then to cancel all the other production orders after the first unit is produced.

In Figure 3.3, we use the same setting with Figure 3.2 with the exception that s is fixed to 46 and the demand rate is a variable. The figure illustrates the impact of traffic intensity on the cost reduction provided by the optimal policy. As the demand rate increases the benefit of the optimal policy over the best base-stock policy first increases and then decreases all the way down to zero. In parallel with the discussion related to Figure 3.2, the percent cost reduction increases until the optimal number of operational servers (at zero inventory level) hits s , which is observed when demand rate is 28. As demand rate increases beyond 28, the optimal policy is unable to open more servers at lower inventory levels, which would be needed to realize its full potential. Therefore, when the demand rate is sufficiently large, both policies start to behave in a similar manner by utilizing all the available capacity at most of the inventory levels, and thereby the cost of the optimal policy converges to the cost of the base-stock policy. As discussed above, the optimal policy under the flexibility of the order cancellation, outperforms both the optimal policy of the primary model and the best base-stock policy for small to moderate demand rate values, and it also becomes identical with the other policies for high demand rates. For sufficiently small s values, one would not observe the region in which the cost reduction increases. Thus, for small values of s the optimal base-stock policy can be used as a good approximation of the optimal policy and this approximation performs better at high demand rates.




Figure 3.3 Optimal production policy vs. Base-stock policy: Impact of demand rate

Figures 3.2 and 3.3 illustrate the effect of the number of available servers and the arrival rate on cost reduction, respectively. Figure 3.4 investigates the joint effect of these two system parameters. For different s and traffic intensity values –traffic intensity is denoted by ρ and is equal to $\lambda/s\mu$ –, the figure compares the optimal production policy with the best base-stock policy. The arrival rate and the number of servers are scaled up proportionally such that ρ remains constant while the number of servers increases. The figure exhibits the results for the case where μ is fixed to unity and $(h, p, c, \alpha) = (0.2, 0.2, 10, 0.6)$. For all ρ values, the percent cost reduction increases with s (and λ). That is, the benefit of having additional production channels (illustrated in Figure 3.2) is much more pronounced than the detriment of heavier traffic (illustrated in Figure 3.3) when traffic intensity is kept constant. The main advantage of the optimal policy over the base-stock policy is its flexibility in adjusting the number of operational servers at each inventory level and this flexibility increases with the number of available servers. The figure also illustrates that, for small values of s , the cost reduction is more significant at lower ρ values. Further-

more, there exists an s value for each traffic intensity beyond which the cost reduction at this intensity is higher than the cost reductions achieved at lower intensities. When the traffic intensity is kept constant, the number of servers used specifies at what proportions of the traffic intensity, the system provides service, i.e., it effectively discretizes the control space. The higher the number of available servers, the finer is the discretization. This ability to adjust the instantaneous utilization more precisely is especially instrumental when the system capacity is tight.

Figures 3.2 and 3.3 compare the performance of the optimal policies for the primary model ($f = 0$) and the model with full-cancellation flexibility ($f = s$) with the optimal base stock policy. Figure 3.5, which uses the same setting with Figure 3.2 with the exception that s is a variable, illustrates the impact of cancellation flexibility index. As stated in the discussion related to Figure 3.2, the benefit of order cancellation under full flexibility increases with the number of available servers. Figure 3.5 reveals that this is also true at any given cancellation flexibility index. It is obvious that more flexibility is better as manifested in the figure. Moreover, as the flexibility index increases, the rate of increase in the percent cost reduction obtained via order cancellation decreases sharply all the way down to zero. Thus, a little flexibility goes a long way and captures most of the value that can be realized via order cancellation. For $s = 26$, while a 27% cost reduction can be obtained with full order cancellation flexibility, having the flexibility of cancelling only one of the previously placed orders captures 67% of this potential gain. Moreover, at $f = 6$, 93% of the potential is captured. As the number of available servers increase, more flexibility is necessary to secure most of the potential gain. However, for all s values, having a little flexibility --compared to full flexibility ($f = s$)-- is sufficient to obtain a significant cost reduction as seen in the figure. As the flexibility index increases, the optimal policy becomes jitterier, i.e., it frequently shuts down operating servers. Due to this tra-

deoff, a manufacturer is likely to opt for little flexibility that enables a significant reduction in operating costs while keeping the production relatively smooth.




Figure 3.4 Optimal production policy vs. Base-stock policy:
Constant traffic intensity




Figure 3.5 Value of Order Cancellation Flexibility

Chapter 3 *M/M/s* Model With Multiple Demand Classes and Lost Sales

For different demand mixes and number of available servers, Table 3.5 compares the optimal rationing policy with the FCFS policy. The comparison is performed under the respective optimal production policies. The table assumes that there are three customer classes and $(\lambda_1 + \lambda_2 + \lambda_3, \mu, h, p, c_1, c_2, c_3, \alpha) = (30, 2, 0.2, 0.2, 10, 6, 2, 0.6)$. The demand mix is specified using the class 1 and class 2 ratios which are $p_1 = \lambda_1 / (\lambda_1 + \lambda_2 + \lambda_3)$ and $p_2 = \lambda_2 / (\lambda_1 + \lambda_2 + \lambda_3)$, respectively.

Table 3.5 Optimal Rationing Policy versus FCFS policy: impact of demand mix

$\begin{matrix} p_2 \\ p_1 \end{matrix}$	0.00	0.17	0.33	0.50	0.67	0.83	1.00	$\begin{matrix} p_2 \\ p_1 \end{matrix}$	0.00	0.17	0.33	0.50	0.67	0.83	1.00
0.00	0.00	15.55	23.40	22.29	14.96	6.97	0.00	0.00	0.00	3.99	10.27	13.30	13.22	9.61	0.00
0.17	28.58	31.43	27.63	19.48	11.27	4.22		0.17	13.06	15.91	17.96	16.60	11.45	0.00	
0.33	37.38	31.60	22.98	14.85	8.05			0.33	21.70	21.68	19.36	12.97	0.79		
0.50	31.94	23.19	15.55	9.10				0.50	25.10	21.68	15.04	3.31			
0.67	19.82	12.97	7.16					0.67	23.96	16.67	3.73				
0.83	8.81	3.70						0.83	16.83	2.84					
1.00	0.00							1.00	0.00						
							<i>s</i> = 6								<i>s</i> = 16
$\begin{matrix} p_2 \\ p_1 \end{matrix}$	0.00	0.17	0.33	0.50	0.67	0.83	1.00	$\begin{matrix} p_2 \\ p_1 \end{matrix}$	0.00	0.17	0.33	0.50	0.67	0.83	1.00
0.00	0.00	0.01	0.00	1.36	2.09	1.54	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.25	0.00
0.17	0.00	1.82	3.15	3.10	1.97	0.00		0.17	0.00	0.00	0.00	0.52	0.55	0.00	
0.33	4.23	4.68	3.97	2.33	0.00			0.33	0.00	0.77	1.08	0.80	0.00		
0.50	6.00	4.72	2.83	0.00				0.50	1.64	1.61	1.03	0.00			
0.67	5.37	3.35	0.00					0.67	2.07	1.23	0.00				
0.83	3.81	0.00						0.83	1.44	0.00					
1.00	0.00							1.00	0.00						
							<i>s</i> = 26								<i>s</i> = 36

The table exhibits that the cost reduction obtained via the optimal rationing policy is much more pronounced at smaller values of s , which in contrast with our observations for the production policy. When s is small, base-stock policy provides a good approximation for the optimal production policy and cost reduction is achieved mainly through rationing. On the other hand, when production capacity is not scarce, the optimal production policy utilizes a large number of servers at lower inventory levels, and thereby replenishes the inventory rapidly. In such settings, the optimal rationing policy reserves limited stock for the customers from more critical classes and satisfies arriving demands on a FCFS basis at most of the inventory levels. Hence, the stock allocation policy is not very critical at larger values of s . Table 3 also shows that when the demand rate of one of the classes is zero, i.e., a two-class system is under consideration, the gap between the performance of the FCFS policy and the optimal rationing policy is maximized when the demand rates of the remaining classes are close to each other. This is an expected result since for a two-class system when one class dominates the other and the value of class differentiation diminishes. Moreover, the benefit of the optimal rationing policy over the FCFS policy is maximized when only the classes with the highest and lowest lost sales costs are present, i.e., when the medium class vanishes. When the differential between the lost sales costs is significant, there is more value to be captured via class differentiation.

We also provide Figure 3.6 in order to better illustrate the performance of optimal rationing policy for a two-class system. The two classes considered in the figure are the class 1 and class 3 of Table 3.5, i.e., the graph assumes $(\lambda_1 + \lambda_2, \mu, h, p, c_1, c_2, \alpha) = (30, 2, 0.2, 0.2, 10, 2, 0.6)$.

Chapter 3 $M/M/s$ Model With Multiple Demand Classes and Lost Sales

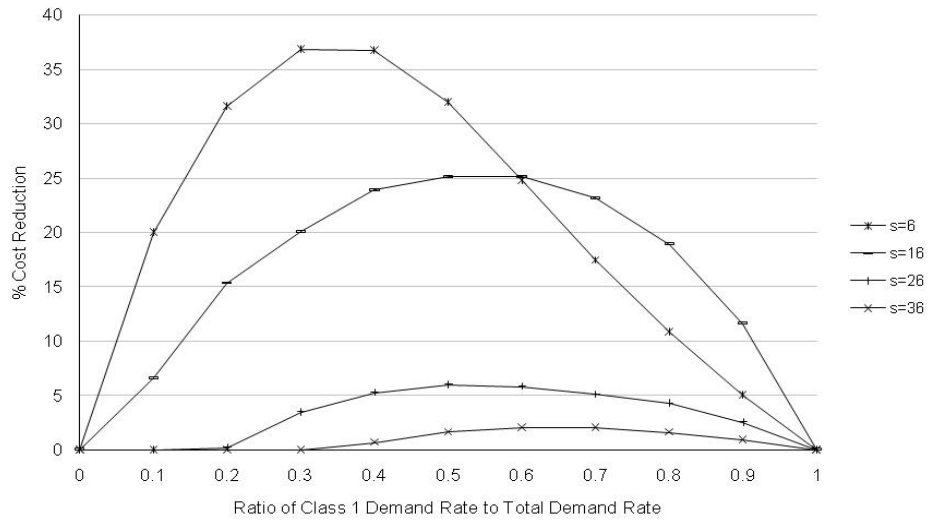


Figure 3.6 Optimal rationing policy vs. FCFS Policy: Impact of demand rate




Figure 3.7 Optimal production and Base-stock policies with or without rationing Impact of total demand rate

Figure 3.7 compares the performance of the production/inventory system under four different policy combinations related to production control and stock allocation over a range of values for the total demand rate. For production control, it considers the optimal base stock and the optimal production policy derived in this paper; whereas, for stock allocation, it considers the FCFS and the optimal rationing policies. The results in the graph pertain to the case where class 1 ratio is fixed to 0.5 and $(s, \mu, h, p, c_1, c_2, \alpha) = (56, 2, 0.2, 0.2, 10, 2, 0.6)$. As expected, the combination of the optimal production policy with the optimal rationing policy provides the lowest cost at all total demand rates. The base stock policy with the FCFS policy yields the highest cost. For a given production policy, the performance gap between the FCFS and the optimal rationing policies grows with the total demand rate. Thus, the value of rationing increases with the traffic intensity. For a given stock allocation policy, as total demand rate increases the cost reduction obtained via the optimal production policy first increases and then goes down all the way to zero. Under high traffic, the optimal production policy behaves in a similar fashion to the optimal base stock policy. The figure shows that the costs of the optimal base stock and the optimal production policies converge as the traffic intensity increases irrespective of the stock allocation policy. The figure also illustrates that, for small to moderate values of the total demand rate, the optimal production policy yields more significant cost savings compared to the optimal rationing policy. However, the opposite is true for higher values of the total demand rate.

3.6 Appendix

Proof of Theorem 3.1: In order to prove that parts i and ii of Theorem 3.1 hold, it is enough to show that the optimization operator preserves these structural properties. Thence, suppose parts i and ii hold. We will first show that

$$u^*(x, y) = \begin{cases} u^*(x, 0), & y \leq u^*(x, 0) \\ y, & y > u^*(x, 0) \end{cases}.$$

Now, for $y \leq u^*(x, 0)$ assume that $u^*(x, y) \neq u^*(x, 0)$. Then, by the hypothesis and the assumption we have

$$\begin{aligned} f(x, 0, u^*(x, 0)) &= f(x, y, u^*(x, 0)) \\ &> f(x, y, u^*(x, y)) \\ &= f(x, 0, u^*(x, y)) \\ &> f(x, 0, u^*(x, 0)) \end{aligned}$$

which is a contradiction. Therefore, $u^*(x, y) = u^*(x, 0)$, for $y \leq u^*(x, 0)$.

In order to show that $u^*(x, y) = y$, for $y > u^*(x, 0)$, we will first show that

$\Delta^u f(x, u^*(x, 0) + 1, u^*(x, 0) + 1) \geq \Delta^u f(x, u^*(x, 0), u^*(x, 0)) > 0$ holds where

$$\begin{aligned} \Delta^u f(x, u^*(x, 0) + 1, u^*(x, 0) + 1) &= f(x, u^*(x, 0) + 1, u^*(x, 0) + 2) - f(x, u^*(x, 0) + 1, u^*(x, 0) + 1) \\ &= p + \mu \Delta^x J(x, u^*(x, 0) + 1) + (u^*(x, 0) + 1) \mu \Delta^y J(x + 1, u^*(x, 0)) \\ &\quad + \Delta^y T_R(x, u^*(x, 0) + 1) \end{aligned}$$

and

$$\begin{aligned} \Delta^u f(x, u^*(x, 0), u^*(x, 0)) &= f(x, u^*(x, 0), u^*(x, 0) + 1) - f(x, u^*(x, 0), u^*(x, 0)) \\ &= p + \mu \Delta^x J(x, u^*(x, 0)) + u^*(x, 0) \mu \Delta^y J(x + 1, u^*(x, 0) - 1) \\ &\quad + \Delta^y T_R(x, u^*(x, 0)) \end{aligned}$$

We can immediately say that $\Delta^u f(x, u^*(x, 0), u^*(x, 0)) > 0$, because

$u^*(x, u^*(x, 0)) = u^*(x, 0)$ as shown just above. Besides,

$\mu\Delta^x J(x, u^*(x, 0) + 1) \geq \mu\Delta^x J(x, u^*(x, 0))$ holds by property (8),

$(u^*(x, 0) + 1)\mu\Delta^y J(x + 1, u^*(x, 0)) \geq u^*(x, 0)\mu\Delta^y J(x + 1, u^*(x, 0) - 1)$ holds by the hy-

pothesis. For each $i \in \{2, \dots, n\}$, $\Delta^y T_{R_i}(x, u^*(x, 0) + 1) \geq \Delta^y T_{R_i}(x, u^*(x, 0))$ is shown by

considering all the possible cases:

Case1. $\Delta^y T_{R_i}(x, u^*(x, 0) + 1) = \lambda_i \Delta^y J(x - 1, u^*(x, 0) + 1)$,

$\Delta^y T_{R_i}(x, u^*(x, 0)) = \lambda_i \Delta^y J(x - 1, u^*(x, 0))$

Then, $\Delta^y T_{R_i}(x, u^*(x, 0) + 1) \geq \Delta^y T_{R_i}(x, u^*(x, 0))$ holds by the hypothesis.

Case2. $\Delta^y T_{R_i}(x, u^*(x, 0) + 1) = \lambda_i \Delta^y J(x - 1, u^*(x, 0) + 1)$, and

$\Delta^y T_{R_i}(x, u^*(x, 0)) = \lambda_i (J(x - 1, u^*(x, 0) + 1) - J(x, u^*(x, 0)) - c_i)$.

Then, $\lambda_i \Delta^y J(x, u^*(x, 0)) \geq \Delta^y T_{R_i}(x, u^*(x, 0))$ by the definition of $T_{R_i}(x, u^*(x, 0) + 1)$,

and $\Delta^y J(x - 1, u^*(x, 0) + 1) \geq \Delta^y J(x, u^*(x, 0))$ by the property

$\Delta^y v(x - 1, y + 1) \geq \Delta^y v(x, y)$. Thus, $\Delta^y T_{R_i}(x, u^*(x, 0) + 1) \geq \Delta^y T_{R_i}(x, u^*(x, 0))$.

Case3. $\Delta^y T_{R_i}(x, u^*(x, 0) + 1) = \lambda_i (J(x - 1, u^*(x, 0) + 2) - J(x, u^*(x, 0) + 1) - c_i)$, and

$\Delta^y T_{R_i}(x, u^*(x, 0)) = \lambda_i \Delta^y J(x, u^*(x, 0))$.

Then, $\Delta^y T_{R_i}(x, u^*(x, 0) + 1) \geq \lambda_i \Delta^y J(x - 1, u^*(x, 0) + 1)$ by the definition of

$T_{R_i}(x, u^*(x, 0) + 1)$, and $\Delta^y J(x - 1, u^*(x, 0) + 1) \geq \Delta^y J(x, u^*(x, 0))$ by the property

$\Delta^y v(x - 1, y + 1) \geq \Delta^y v(x, y)$. Thus, $\Delta^y T_{R_i}(x, u^*(x, 0) + 1) \geq \Delta^y T_{R_i}(x, u^*(x, 0))$.

Case4. $\Delta^y T_{R_i}(x, u^*(x, 0) + 1) = \lambda_i \Delta^y J(x, u^*(x, 0) + 1)$,

$\Delta^y T_{R_i}(x, u^*(x, 0)) = \lambda_i \Delta^y J(x, u^*(x, 0))$

Then, $\Delta^y T_{R_i}(x, u^*(x, 0) + 1) \geq \Delta^y T_{R_i}(x, u^*(x, 0))$ holds by the hypothesis.

The results of above four cases imply that

$$\sum_{i=1}^n \Delta^y T_{R_i}(x, u^*(x, 0) + 1) \geq \sum_{i=1}^n \Delta^y T_{R_i}(x, u^*(x, 0)), \text{ i.e.,}$$

$$\Delta^y T_R(x, u^*(x, 0) + 1) \geq \Delta^y T_R(x, u^*(x, 0)).$$

(Here it should be noted that following the steps of the above four cases, we can also conclude that $\Delta^y T_R(x, y + 1) \geq \Delta^y T_R(x, y)$ holds for any y . That is, $T_R(x, y)$ is y -convex.)

Thus, $\Delta^u f(x, u^*(x, 0) + 1, u^*(x, 0) + 1) \geq \Delta^u f(x, u^*(x, 0), u^*(x, 0)) > 0$ holds. Moreover, having $J(x, u)$ is a convex-increasing function of u and $T_R(x, u)$ is u -convex implies that $f(x, y, u) = hx + pu + (s - u)\mu J(x, y) + u\mu J(x + 1, u - 1) + T_R(x, u)$ is a convex function of u . Therefore, for any $u \geq u^*(x, 0) + 2$, $\Delta^u f(x, u^*(x, 0) + 1, u) \geq \Delta^u f(x, u^*(x, 0) + 1, u^*(x, 0) + 1) > 0$, that is $u^*(u^*(x, 0) + 1) = u^*(x, 0) + 1$. Following the same logic used in the proof of $u^*(u^*(x, 0) + 1) = u^*(x, 0) + 1$, it is easy to show also that $u^*(x, y) = y$, for $y \geq u^*(x, 0) + 2$. Hence, we conclude that $u^*(x, y) = \begin{cases} u^*(x, 0), & y \leq u^*(x, 0) \\ y, & y > u^*(x, 0) \end{cases}$

holds. Using this fact, we will show that the optimization operator preserves the properties stated in Theorem 3.1. Now, first consider

$$\begin{aligned} \Delta^y T(J(x, y)) &= T(J(x, y + 1)) - T(J(x, y)) \\ &= f(x, y + 1, u^*(x, y + 1)) - f(x, y, u^*(x, y)) \end{aligned}$$

For $y < u^*(x, 0)$, $u^*(x, y) = u^*(x, y+1) = u^*(x, 0)$ (as shown above) and

$J(x, y) = J(x, y+1)$ (by the hypothesis). Then,

$$\Delta^y T(J(x, y)) = f(x, y, u^*(x, y)) - f(x, y, u^*(x, y)) = 0.$$

For $y \geq u^*(x, 0)$, $u^*(x, y) = y$, $u^*(x, y+1) = y+1$ and $J(x, y) < J(x, y+1)$. Then,

$$\Delta^y T(J(x, y)) = f(x, y+1, y+1) - f(x, y, y) > f(x, y, y+1) - f(x, y, y). \quad \text{Moreo-}$$

ver, since $f(x, y, u)$ is u -convex and $u^*(x, y) = y$, $f(x, y, y+1) - f(x, y, y) \geq 0$.

Hence, $\Delta^y T(J(x, y)) > 0$.

We will finally show that for $y \geq u^*(x, 0)$, $\Delta^y T(J(x, y+1)) \geq \Delta^y T(J(x, y))$, i.e.,

$J(x, y)$ is y -convex. For $y \geq u^*(x, 0)$, $u^*(x, y) = y$, $u^*(x, y+1) = y+1$ and

$$u^*(x, y+2) = y+2.$$

Therefore,

$$\begin{aligned} & \Delta^y T(J(x, y+1)) - \Delta^y T(J(x, y)) \\ &= (s - (y+2)\mu)(\Delta^y J(x, y+1) - \Delta^y J(x, y)) \\ & \quad + y\mu(\Delta^y J(x+1, y) - \Delta^y J(x+1, y-1)) + 2\mu(\Delta^y J(x+1, y) - \Delta^y J(x, y)) \\ & \quad + \Delta^y T_R(x, y+1) - \Delta^y T_R(x, y) \end{aligned}$$

Where $(s - (y+2)\mu)(\Delta^y J(x, y+1) - \Delta^y J(x, y))$ and

$y\mu(\Delta^y J(x+1, y) - \Delta^y J(x+1, y-1))$ are greater or equal to zero by the hypothesis,

and $2\mu(\Delta^y J(x+1, y) - \Delta^y J(x, y)) \geq 0$ by the property (8). Moreover, it is shown

above that $\Delta^y T_R(x, y+1) - \Delta^y T_R(x, y) \geq 0$. Thus, we conclude that

$$\Delta^y T(J(x, y+1)) - \Delta^y T(J(x, y)) \geq 0, \text{ i.e., } J(x, y) \text{ is } y\text{-convex.}$$

Proof of Corollary 3.1:

i. It is shown in the proof of Theorem 3.1 that $u^*(x, y) = \begin{cases} u^*(x, 0), & y \leq u^*(x, 0) \\ y, & y > u^*(x, 0) \end{cases}$

holds.

ii. Theorem 3.1 indicates that

$J(x, 0) = \dots = J(x, u^*(x, 0)) < J(x, u^*(x, 0) + 1) < \dots < J(x, s)$. Using this result and part-i of Corollary 3.1, we can alternatively define the optimal number of operational servers at state (x, y) as the number of operational servers beyond which the optimal cost function starts to increase.

iii. It is apparent from part-i of Corollary 3.1.

Proof of Theorem 3.2:

i. Immediate conclusion from properties (8) and (9).

ii. By definition of $u^*(x, y)$, $J(x, u^*(x, y) + 1) - J(x, u^*(x, y)) > 0$. From property (8) we have

$J(x + 1, u^*(x, y) + 1) - J(x + 1, u^*(x, y)) \geq J(x, u^*(x, y) + 1) - J(x, u^*(x, y))$. Therefore, $J(x + 1, u^*(x, y) + 1) - J(x + 1, u^*(x, y)) > 0$. Thus, we conclude that $u^*(x + 1, y) \leq u^*(x, y)$.

iii. It will be enough to show that the optimization operator preserves the property

$\Delta^x J(x - 1, y) \geq -c_1$. Suppose $\Delta^x J(x - 1, y) \geq -c_1$. Now, for any $u \geq y$,

$$f(x, y, u) - f(x - 1, y, u) = h + s\mu\Delta^x J(x - 1, y) + u\mu(\Delta^x J(x, u - 1) - \Delta^x J(x - 1, y)) + \sum_{i=1}^n \Delta^x T_{R_i}(x - 1, u)$$

In the above equation $h \geq 0$ and $u\mu(\Delta^x J(x, u-1) - \Delta^x J(x-1, y)) \geq 0$ by properties

(8) and (9). Moreover, by the fact that $s\mu = 1 - \sum_{i=1}^n \lambda_i$ and the hypothesis,

$$s\mu\Delta^x J(x-1, y) \geq -c_1 \left(1 - \sum_{i=1}^n \lambda_i\right). \text{ Therefore,}$$

$$h + s\mu\Delta^x J(x-1, y) + u\mu(\Delta^x J(x, u-1) - \Delta^x J(x-1, y)) \geq -c_1 \left(1 - \sum_{i=1}^n \lambda_i\right).$$

By the hypothesis, $\Delta^x T_{R_i}(x-1, u) = \lambda_i \Delta^x J(x-2, u) \geq -\lambda_i c_1$. For each $i \in \{2, \dots, n\}$,

$\Delta^x T_{R_i}(x-1, u) \geq -\lambda_i c_1$ can be shown by considering all the possible cases:

Case1. $T_{R_i}(x, u) = \lambda_i J(x-1, u)$, $T_{R_i}(x-1, u) = \lambda_i J(x-2, u)$

Then, $\Delta^x T_{R_i}(x-1, u) = \lambda_i (J(x-1, u) - J(x-2, u)) \geq -\lambda_i c_1$ by the hypothesis.

Case2. $T_{R_i}(x, u) = \lambda_i J(x-1, u)$, $T_{R_i}(x-1, u) = \lambda_i (c_i + J(x-1, u))$

Then, $\Delta^x T_{R_i}(x-1, u) = \lambda_i (J(x-1, u) - J(x-1, u) - c_i) = -\lambda_i c_i \geq -\lambda_i c_1$ because $c_1 \geq c_i$

Case3. $T_{R_i}(x, \bar{y}) = \lambda_i (c_i + J(x, u))$, $T_{R_i}(x-1, u) = \lambda_i (c_i + J(x-1, u))$

Then, $\Delta^x T_{R_i}(x-1, u) = \lambda_i (J(x, u) - J(x-1, u)) \geq -\lambda_i c_1$ by the hypothesis.

Hence, we have $\sum_{i=1}^n \Delta^x T_{R_i}(x-1, u) \geq -c_1 \sum_{i=1}^n \lambda_i$, and so $f(x, y, u) - f(x-1, y, u) \geq -c_1$.

Having $f(x, y, u) - f(x-1, y, u) \geq -c_1$, $u \geq y$, implies that

$$f(x, y, u^*(x, y)) - f(x-1, y, u^*(x, y)) \geq -c_1. \text{ Since}$$

$$f(x-1, y, u^*(x, y)) \geq f(x-1, y, u^*(x-1, y)), \text{ we conclude that}$$

$$f(x, y, u^*(x, y)) - f(x-1, y, u^*(x-1, y)) \geq -c_1, \text{ i.e.,}$$

$$T(J(x, y)) - T(J(x-1, y)) \geq -c_1. \text{ Thus, } \Delta^x J(x-1, y) \geq -c_1.$$

iv. For $i \in \{2, \dots, n\}$, let us define $K_x^i(y) = \min\{x : \Delta^x J(x, y) \geq -c_i\}$. Since

$$-c_n \leq -c_{n-1} \dots \leq -c_2 \text{ and } \Delta^x J(x, y) \text{ is non-decreasing in } x \text{ (from part-i),}$$

$$K_x^n(y) \geq K_x^{n-1}(y) \geq \dots \geq K_x^2(y) \geq 0 \text{ holds.}$$

From property (8), $\Delta^x J(K_x^i(y), y+1) \geq \Delta^x J(K_x^i(y), y) \geq -c_i$, and since

$$\Delta^x J(x, y) \text{ is non-decreasing in } x \text{ (from part-i) } K_x^i(y+1) \leq K_x^i(y) \text{ holds.}$$

v. For $i \in \{2, \dots, n\}$, let us define $K_y^i(x) = \min\{y : \Delta^x J(x-1, y) \geq -c_i\} - 1$. Having

$$K_y^i(x) = -1 \text{ implies that it is optimal to satisfy an arriving class } i \text{ demand at all } y$$

$$\text{levels when the on-hand inventory level is } x. \text{ Since } -c_n \leq -c_{n-1} \dots \leq -c_2 \text{ and}$$

$$\Delta^x J(x, y) \text{ is non-decreasing in } y, K_y^n(x) \geq K_y^{n-1}(x) \geq \dots \geq K_y^2(x) \geq 0 \text{ holds.}$$

Property (9) implies that if $\Delta^x J(x, y+1) \geq -c_i$, then $\Delta^x J(x+1, y) \geq -c_i$. That is, if

$$\Delta^x J(x+1, y) \geq \Delta^x J(x, y+1) \geq -c_i \text{ and } \Delta^x J(x, y) < -c_i \text{ then } K_y^i(x) = y+1 \text{ and}$$

$$K_y^i(x+1) < y. \text{ Therefore, we can conclude that } K_y^i(x+1) < K_y^i(x).$$

Proof of Theorem 3:

Suppose $\Delta^x J(x, y+1) \geq \Delta^x J(x, y)$ and $\Delta^x J(x, y) \geq \Delta^x J(x-1, y+1)$. We will show that the optimization operator T preserves this structure.

We will first show that

$$\Delta^x f(x, y+1, y+1) \geq \Delta^x f(x, y, y) \tag{26}$$

$$\Delta^x f(x, y, y) \geq \Delta^x f(x-1, y+1, y+1) \tag{27}$$

then $\Delta^x T(J(x, y+1)) \geq \Delta^x T(J(x, y)) =$

$$\begin{aligned} & f(x+1, u^*(x+1, y+1), u^*(x+1, y+1)) - f(x, u^*(x, y+1), u^*(x, y+1)) \\ & \geq f(x+1, u^*(x+1, y), u^*(x+1, y)) - f(x, u^*(x, y), u^*(x, y)) \end{aligned} \quad (28)$$

and $\Delta^x T(J(x, y)) \geq \Delta^x T(J(x-1, y+1)) =$

$$\begin{aligned} & f(x+1, u^*(x+1, y), u^*(x+1, y)) - f(x, u^*(x, y), u^*(x, y)) \\ & \geq f(x, u^*(x, y+1), u^*(x, y+1)) - f(x-1, u^*(x-1, y+1), u^*(x-1, y+1)) \end{aligned} \quad (29)$$

From the hypothesis and part-iii of Theorem2;

$$\begin{aligned} \Delta^x f(x, y+1, y+1) &= h + (s - y - 1) \mu \Delta^x J(x, y+1) \\ &+ (y+1) \mu \Delta^x J(x+1, y) + \lambda_1 \Delta^x J(x-1, y+1) \\ &+ \sum_{i=2}^n T_{R_i}(x+1, y+1) - \sum_{i=2}^n T_{R_i}(x, y+1) \\ &\geq h + (s - y - 1) \mu \Delta^x J(x, y) + y \mu \Delta^x J(x+1, y-1) \\ &+ \mu \Delta^x J(x+1, y-1) + \lambda_1 \Delta^x J(x-1, y) \\ &+ \sum_{i=2}^n T_{R_i}(x+1, y+1) - \sum_{i=2}^n T_{R_i}(x, y+1) \\ &= h + (s - y) \mu \Delta^x J(x, y) + y \mu \Delta^x J(x+1, y-1) \\ &+ \mu \Delta^x J(x+1, y-1) + \lambda_1 \Delta^x J(x-1, y) \\ &+ \sum_{i=2}^n T_{R_i}(x+1, y+1) - \sum_{i=2}^n T_{R_i}(x, y+1) + \mu (\Delta^x J(x+1, y-1) - \Delta^x J(x, y)) \\ &\geq h + (s - y) \mu \Delta^x J(x, y) + y \mu \Delta^x J(x+1, y-1) \\ &+ \mu \Delta^x J(x+1, y-1) + \lambda_1 \Delta^x J(x-1, y) \\ &+ \sum_{i=2}^n T_{R_i}(x+1, y+1) - \sum_{i=2}^n T_{R_i}(x, y+1) \end{aligned}$$

In order to conclude that the right hand side of the above inequality is greater than $\Delta^x f(x, y, y)$, we need to show that

$$\forall i \in \{2, \dots, n\}, T_{R_i}(x+1, y+1) - T_{R_i}(x, y+1) \geq T_{R_i}(x+1, y) - T_{R_i}(x, y),$$

i.e., $\Delta^x T_{R_i}(x, y+1) \geq \Delta^x T_{R_i}(x, y)$. We will consider three cases in order to show that the inequality holds.

Case1. $K_x^i(y+1) \leq x-1$. Then, $\Delta^x T_{R_i}(x, y+1) = \lambda_i \Delta^x J(x-1, y+1)$

i. $K_x^i(y) \leq x-1$

$$\Delta^x T_{R_i}(x, y) = \lambda_i \Delta^x J(x-1, y). \text{ Thus, } \Delta^x T_{R_i}(x, y+1) \geq \Delta^x T_{R_i}(x, y).$$

ii. $K_x^i(y) = x$

$$\Delta^x T_{R_i}(x, y) = \lambda_i (J(x, y) - c_i - J(x, y)) = -\lambda_i c_i.$$

$J(x-1, y+1) \geq -c_i$ because $K_x^i(y+1) \leq x-1$. Thus,

$$\Delta^x T_{R_i}(x, y+1) \geq \Delta^x T_{R_i}(x, y).$$

iii. $K_x^i(y) \geq x+1$

This case is not possible, because the hypothesis

$\Delta^x J(x, y) \geq \Delta^x J(x-1, y+1)$ implies that if $\Delta^x J(x-1, y+1) \geq -c_i$, then

$\Delta^x J(x, y) \geq -c_i$. In words, if we satisfy an arriving class i demand

when there are x units of inventory and $(y+1)$ active servers, we

should satisfy an arriving class i demand when there are $(x+1)$ units

of inventory and y active servers.

Therefore, $K_x^i(y) \leq x$ whenever $K_x^i(y+1) \leq x-1$.

Case2. $K_x^i(y+1) = x$. Then,

$$\Delta^x T_{R_i}(x, y+1) = \lambda_i (J(x, y+1) - c_i - J(x, y+1)) = -\lambda_i c_i$$

i. $K_x^i(y) = x$

$$\Delta^x T_{R_i}(x, y) = \lambda_i (J(x, y) - c_i - J(x, y)) = -\lambda_i c_i. \text{ Thus,}$$

$$\Delta^x T_{R_i}(x, y+1) = \Delta^x T_{R_i}(x, y).$$

ii. $K_x^i(y) = x + 1$

$$\Delta^x T_{R_i}(x, y) = \lambda_i (c_i + J(x+1, y) - c_i - J(x, y)) = \lambda_i \Delta^x J(x, y).$$

And, $-\lambda_i c_i \geq \lambda_i \Delta^x J(x, y)$ because $K_x^i(y) = x + 1$. Thus,

$$\Delta^x T_{R_i}(x, y+1) \geq \Delta^x T_{R_i}(x, y).$$

Case3. $K_x^i(y+1) \geq x + 1$.

Then,

$$\Delta^x T_{R_i}(x, y+1) = \lambda_i (c_i + J(x+1, y+1) - c_i - J(x, y+1)) = \lambda_i \Delta^x J(x, y+1).$$

Since we have $K_x^i(y+1) \leq K_x^i(y)$, it is true that

$$\Delta^x T_{R_i}(x, y) = \lambda_i (c_i + J(x+1, y) - c_i - J(x, y)) = \lambda_i \Delta^x J(x, y). \text{ Then by the hypothesis } \Delta^x T_{R_i}(x, y+1) \geq \Delta^x T_{R_i}(x, y).$$

We have just shown that (14) holds. Showing (15) is equivalent to show that

$f(x+1, y, y) - f(x, y+1, y+1) \geq f(x, y, y) - f(x-1, y+1, y+1)$, where

$$\begin{aligned} f(x+1, y, y) - f(x, y+1, y+1) &= h - p + (s - y - 1) \mu (J(x+1, y) - J(x, y+1)) \\ &\quad + y \mu (J(x+2, y-1) - J(x+1, y)) + \lambda_1 (J(x, y) - J(x-1, y+1)) \\ &\quad + \sum_{i=2}^n T_{R_i}(x+1, y) - \sum_{i=2}^n T_{R_i}(x, y+1) \end{aligned}$$

and

$$\begin{aligned} f(x, y, y) - f(x-1, y+1, y+1) &= h - p + (s - y - 1) \mu (J(x, y) - J(x-1, y+1)) \\ &\quad + y \mu (J(x+1, y-1) - J(x, y)) + \lambda_1 (J(x-1, y) - J(x-2, y+1)) \\ &\quad + \sum_{i=2}^n T_{R_i}(x, y) - \sum_{i=2}^n T_{R_i}(x-1, y+1) \end{aligned}$$

Using the hypothesis it is easy to show that the terms of the right hand side of the first equation are greater or equal to the respective terms of the right hand side of the

second equation expect the terms related to the rationing decision for each class other than class 1. Therefore, we need to show that

$$\forall i \in \{2, \dots, n\}, \Delta^x T_{R_i}^i(x, y) \geq \Delta^x T_{R_i}^i(x-1, y+1).$$

Case1. $K_x^i(y+1) \leq x-2$. Then, $\Delta^x T_{R_i}^i(x-1, y+1) = \lambda_i \Delta^x J(x-2, y+1)$

In this case we have $\Delta^x T_{R_i}^i(x, y) = \lambda_i \Delta^x J(x-1, y)$. Thus, by the hypothesis

$$\Delta^x T_{R_i}^i(x, y) \geq \Delta^x T_{R_i}^i(x-1, y+1)$$

Case2. $K_x^i(y+1) = x-1$. Then,

$$\Delta^x T_{R_i}^i(x-1, y+1) = \lambda_i (J(x-1, y+1) - c_i - J(x-1, y+1)) = -\lambda_i c_i$$

i. $K_x^i(y) = x-1$

$$\Delta^x T_{R_i}^i(x, y) = \lambda_i \Delta^x J(x-1, y), \text{ and } \Delta^x J(x-1, y) \geq -c_i \text{ since}$$

$$K_x^i(y) = x-1. \text{ So, } \Delta^x T_{R_i}^i(x, y) \geq \Delta^x T_{R_i}^i(x-1, y+1).$$

ii. $K_x^i(y) = x$

$$\Delta^x T_{R_i}^i(x, y) = \lambda_i (J(x, y) - c_i - J(x, y)) = -\lambda_i c_i = \Delta^x T_{R_i}^i(x-1, y+1)$$

Case3. $K_x^i(y+1) \geq x$. Then,

$$\Delta^x T_{R_i}^i(x-1, y+1) = \lambda_i (c_i + J(x, y+1) - c_i - J(x-1, y+1)) = \lambda_i \Delta^x J(x-1, y+1)$$

i. $K_x^i(y) = x$ (possible if $K_x^i(y+1) = x$)

$$\Delta^x T_{R_i}^i(x, y) = -\lambda_i c_i, \text{ and } \Delta^x J(x-1, y+1) \leq -c_i \text{ because we do not sa-}$$

tisfy class 2 demand at state $(x, y+1)$. Thus,

$$\Delta^x T_{R_i}^i(x, y) \geq \Delta^x T_{R_i}^i(x-1, y+1).$$

ii. $K_x^i(y) \geq x+1$

$\Delta^x T_{R_i}(x, y) = \lambda_i \Delta^x J(x, y)$, and by the hypothesis

$\Delta^x J(x, y) \geq \Delta^x J(x-1, y+1)$. Thus, $\Delta^x T_{R_i}(x, y) \geq \Delta^x T_{R_i}(x-1, y+1)$

We have also shown (15). Now, we need to show that (16) holds in order to complete the proof of part-i of the theorem.

From Corollary 3.1 and part-ii of Theorem 3.2, we have

- $u^*(x, y+1) \geq u^*(x, y) \geq u^*(x+1, y)$
- $u^*(x, y+1) \geq u^*(x+1, y+1) \geq u^*(x+1, y)$
- $u^*(x, y+1) \leq u^*(x, y) + 1$ and $u^*(x+1, y+1) \leq u^*(x+1, y) + 1$

Case1. $u^*(x, y+1) \geq u^*(x+1, y+1) \geq u^*(x, y) \geq u^*(x+1, y)$

i. $u^*(x+1, y) = u^*(x, y) = u^*(x+1, y+1) = u^*(x, y+1)$

Trivial case.

ii. $u^*(x+1, y) = u^*(x, y) = u^*(x+1, y+1) = u^*(x, y+1) - 1$

Since $u^*(x+1, y) = u^*(x+1, y+1)$, we should have $u^*(x+1, 0) \geq y+1$ and so $u^*(x, 0) \geq y+1$. Therefore, $u^*(x, y) = u^*(x, y+1)$ should hold. Thus, this case is not possible.

iii. $u^*(x+1, y) = u^*(x, y) = u^*(x+1, y+1) - 1 = u^*(x, y+1) - 1$

Then, $u^*(x+1, 0) \leq u^*(x, 0) \leq y$. So, $u^*(x+1, y) = u^*(x, y) = y$ and

$u^*(x+1, y+1) = u^*(x, y+1) = y+1$. Thus, (16) holds because (14) holds.

iv. $u^*(x+1, y) = u^*(x, y) - 1 = u^*(x+1, y+1) - 1 = u^*(x, y+1) - 1$

Then, $u^*(x+1,0) \leq y$, $y^*(x+1,y) = y$. And,

$$u^*(x,y) = u^*(x+1,y+1) = u^*(x,y+1) = y+1.$$

Therefore, the left hand side of (16) becomes

$f(x+1,y+1,y+1) - f(x,y+1,y+1)$ and the right hand side becomes

$f(x+1,y,y) - f(x,y+1,y+1)$. Since $u^*(x+1,0) \leq y$, due to Theorem 3.1

we have $f(x+1,y+1,y+1) \geq f(x+1,y,y)$ and (16) holds.

Case2. $u^*(x,y+1) \geq u^*(x,y) \geq u^*(x+1,y+1) \geq u^*(x+1,y)$

i. $u^*(x+1,y) = u^*(x+1,y+1)$ and $u^*(x,y) = u^*(x,y+1)$

Trivial case.

ii. $u^*(x+1,y) = u^*(x+1,y+1) - 1$ and $u^*(x,y) = u^*(x,y+1) - 1$

In this case (16) holds due to (14).

iii. $u^*(x+1,y) = u^*(x+1,y+1)$ and $u^*(x,y) = u^*(x,y+1) - 1$

Then, $u^*(x+1,y) = u^*(x+1,y+1)$ implies that

$u^*(x,0) \geq u^*(x+1,0) \geq y+1$ and so $u^*(x,y) = u^*(x,y+1)$. Thus, this case is not possible.

iv. $u^*(x+1,y) = u^*(x+1,y+1) - 1$ and $u^*(x,y) = u^*(x,y+1)$

Then, $u^*(x+1,0) \leq y$, and so $u^*(x+1,y) = y$, $u^*(x+1,y+1) = y+1$.

Therefore, the left hand side of (16) becomes

$f(x+1,y+1,y+1) - f(x,u^*(x,0),u^*(x,0))$ and the right hand side be-

comes $f(x+1, y, y) - f(x, u^*(x, 0), u^*(x, 0))$. Since $u^*(x+1, 0) \leq y$, due to Theorem 3.1 we have $f(x+1, y+1, y+1) \geq f(x+1, y, y)$ and (16) holds.

Finally we will prove (17) where

$$\begin{aligned} \Delta^x T(J(x, y)) &= f(x+1, y, u^*(x+1, y)) - f(x, y, u^*(x, y)) \\ &\geq f(x+1, y, u^*(x+1, y)) - f(x, y, u^*(x+1, y)) \\ &= h + (s - u^*(x+1, y))\mu\Delta^x J(x, y) + u^*(x+1, y)\mu\Delta^x J(x+1, u^*(x+1, y) - 1) \text{ and} \\ &\quad + \sum_{i=1}^n \Delta^x T_{R_i}(x, u^*(x+1, y)) \end{aligned}$$

$$\begin{aligned} \Delta^x T(J(x-1, y+1)) &= f(x, y+1, u^*(x, y+1)) - f(x-1, y+1, u^*(x-1, y+1)) \\ &\leq f(x, y+1, u^*(x-1, y+1)) - f(x-1, y+1, u^*(x-1, y+1)) \\ &= h + (s - u^*(x-1, y+1))\mu\Delta^x J(x-1, y) + u^*(x-1, y+1)\mu\Delta^x J(x, u^*(x-1, y+1) - 1) \\ &\quad + \sum_{i=1}^n \Delta^x T_{R_i}(x-1, u^*(x-1, y+1)) \end{aligned}$$

We multiply the second inequality by -1 and add to the first one. Then,

$$\begin{aligned} \Delta^x T(J(x, y)) - \Delta^x T(J(x-1, y+1)) &\geq (s - u^*(x+1, y))\mu\Delta^x J(x, y) - (s - u^*(x-1, y+1))\mu\Delta^x J(x-1, y) \\ &\quad + u^*(x+1, y)\mu\Delta^x J(x+1, u^*(x+1, y) - 1) - u^*(x-1, y+1)\mu\Delta^x J(x, u^*(x-1, y+1) - 1) \\ &\quad + \sum_{i=1}^n \Delta^x T_{R_i}(x, u^*(x+1, y)) - \sum_{i=1}^n \Delta^x T_{R_i}(x-1, u^*(x-1, y+1)) \end{aligned}$$

For $i \in \{1, 2, \dots, n\}$, we have

$\Delta^x T_{R_i} J(x-1, u^*(x-1, y+1)) = \Delta^x T_{R_i} J(x-1, u^*(x+1, y+1))$ due to part-i of Theorem 3.1 and part-iii of Corollary 3.1.

Therefore, by the hypothesis

$$\left(\begin{aligned} &\lambda_1 \Delta^x J(x-1, u^*(x+1, y)) + \sum_{i=2}^n \lambda_i \Delta^x T_{R_i} J(x, u^*(x+1, y)) \\ &- \left(\lambda_1 \Delta^x J(x-2, u^*(x-1, y+1)) + \sum_{i=2}^n \lambda_i \Delta^x T_{R_i} J(x-1, u^*(x-1, y+1)) \right) \end{aligned} \right) \geq 0.$$

Moreover,

$$\begin{aligned} &(s - u^*(x+1, y)) \mu \Delta^x J(x, y) \\ &= (s - u^*(x-1, y+1)) \mu \Delta^x J(x, y) + (u^*(x-1, y+1) - u^*(x+1, y)) \mu \Delta^x J(x, y) \end{aligned}$$

and $(s - u^*(x-1, y+1)) \mu (\Delta^x J(x, y) - \Delta^x J(x-1, y)) \geq 0$. The remaining terms of the right hand side is greater or equal to the zero which can be shown by applying the operator technique once more.

Proof of Theorem 3.4:

- i.** If $p + \mu \Delta J(x) < 0$, then it is optimal to have $u^*(x) = s$ in order to minimize cost function as much as possible. Otherwise, $u^*(x)$ should be zero because having a positive number of operational servers would inflate the cost.
- ii.** Under maximization objective, Cil et al. (2009) show that both the production and the rationing operators are concave. It is easy to adapt their results to our case (where the value functions are cost-to-go functions) and show that $J(x)$ is a convex function of x .
- iii.** Since $\Delta J(x)$ is non-decreasing (from part ii), there exists a threshold rationing inventory level $K^i := \min \{x : \Delta J(x) \geq -c_i\}$

- iv. For a specific value of L , Cil et al. (2009) show that $v(x-1) - v(x) \geq L$ where $v(x)$ is the reward function. Letting $J(x) = -v(x)$ and $L = -c_1$ is sufficient to show that $\Delta J(x-1) \geq -c_1$.
- v. Immediately follows from parts i and ii.

Proof of Lemma 3.1:

We have $P_0(S) = \left(\sum_{j=0}^S \frac{S!}{(S-j)!} \left(\frac{\mu}{\lambda} \right)^j \right)^{-1}$,

$$\Delta C(S, \vec{0}) = C(S+1, \vec{0}) - C(S, \vec{0}) = h + \left(\frac{h\lambda}{\mu} + \sum_{i=1}^n \lambda_i c_i \right) \Delta P_0(S), \text{ and}$$

$$\Delta^2 C(S, \vec{0}) = \Delta C(S+1, \vec{0}) - \Delta C(S, \vec{0}) = \left(\frac{h\lambda}{\mu} + \sum_{i=1}^n \lambda_i c_i \right) \Delta^2 P_0(S). \text{ To conclude that}$$

$\Delta^2 C(S, \vec{0}) > 0$ we first need to compute $\Delta P_0(S)$ and then $\Delta^2 P_0(S)$:

$$\Delta P_0(S) = P_0(S+1) - P_0(S) = \frac{1}{\sum_{j=0}^{S+1} \frac{(S+1)!}{(S+1-j)!} \left(\frac{\mu}{\lambda} \right)^j} - \frac{1}{\sum_{j=0}^S \frac{S!}{(S-j)!} \left(\frac{\mu}{\lambda} \right)^j}$$

$$= \frac{\sum_{j=0}^S \left(\frac{S!}{(S-j)!} - \frac{(S+1)!}{(S+1-j)!} \right) \left(\frac{\mu}{\lambda} \right)^j - (S+1)! \left(\frac{\mu}{\lambda} \right)^{S+1}}{\left(\sum_{j=0}^{S+1} \frac{(S+1)!}{(S+1-j)!} \left(\frac{\mu}{\lambda} \right)^j \right) \left(\sum_{j=0}^S \frac{S!}{(S-j)!} \left(\frac{\mu}{\lambda} \right)^j \right)}. \text{ That is,}$$

$$\begin{aligned} \Delta P_0(S) &= \frac{\sum_{j=0}^S \left(\frac{S!}{(S-j)!} - \frac{(S+1)(S)!}{(S+1-j)(S-j)!} \right) \left(\frac{\mu}{\lambda} \right)^j - (S+1)! \left(\frac{\mu}{\lambda} \right)^{S+1}}{\left(\sum_{j=0}^{S+1} \frac{(S+1)!}{(S+1-j)!} \left(\frac{\mu}{\lambda} \right)^j \right) \left(\sum_{j=0}^S \frac{S!}{(S-j)!} \left(\frac{\mu}{\lambda} \right)^j \right)} \\ &= \frac{\sum_{j=0}^{S+1} \frac{S!}{(S-j)!} (-j) \left(\frac{\mu}{\lambda} \right)^j}{\left(\sum_{j=0}^{S+1} \frac{(S+1)!}{(S+1-j)!} \left(\frac{\mu}{\lambda} \right)^j \right) \left(\sum_{j=0}^S \frac{S!}{(S-j)!} \left(\frac{\mu}{\lambda} \right)^j \right)} \end{aligned}$$

$$\Delta^2 P_0(S) = \Delta P_0(S+1) - \Delta P_0(S)$$

$$\begin{aligned} &\left(\sum_{j=0}^S \frac{S!}{(S-j)!} \left(\frac{\mu}{\lambda} \right)^j \right) \left(\sum_{j=0}^{S+2} \frac{(S+1)!}{(S+1-j)!} (-j) \left(\frac{\mu}{\lambda} \right)^j \right) - \\ &\left(\sum_{j=0}^{S+2} \frac{(S+2)!}{(S+2-j)!} \left(\frac{\mu}{\lambda} \right)^j \right) \left(\sum_{j=0}^{S+1} \frac{S!}{(S-j)!} (-j) \left(\frac{\mu}{\lambda} \right)^j \right) \\ &= \frac{\left(\sum_{j=0}^{S+2} \frac{(S+2)!}{(S+2-j)!} \left(\frac{\mu}{\lambda} \right)^j \right) \left(\sum_{j=0}^{S+1} \frac{(S+1)!}{(S+1-j)!} \left(\frac{\mu}{\lambda} \right)^j \right) \left(\sum_{j=0}^S \frac{S!}{(S-j)!} \left(\frac{\mu}{\lambda} \right)^j \right)}{\left(\sum_{j=0}^S \frac{S!}{(S-j)!} \left(\frac{\mu}{\lambda} \right)^j \right) \left(\sum_{j=0}^{S+1} \frac{(S+1)!}{(S+1-j)!} (-j) \left(\frac{\mu}{\lambda} \right)^j - (S+2)! \left(\frac{\mu}{\lambda} \right)^{S+2} \right)} \\ &- \frac{\left(\sum_{j=0}^S \frac{(S+2)!}{(S+2-j)!} \left(\frac{\mu}{\lambda} \right)^j + (S+2)! \left(\frac{\mu}{\lambda} \right)^{S+1} + (S+2)! \left(\frac{\mu}{\lambda} \right)^{S+2} \right) \left(\sum_{j=0}^{S+1} \frac{S!}{(S-j)!} (-j) \left(\frac{\mu}{\lambda} \right)^j \right)}{\left(\sum_{j=0}^{S+2} \frac{(S+2)!}{(S+2-j)!} \left(\frac{\mu}{\lambda} \right)^j \right) \left(\sum_{j=0}^{S+1} \frac{(S+1)!}{(S+1-j)!} \left(\frac{\mu}{\lambda} \right)^j \right) \left(\sum_{j=0}^S \frac{S!}{(S-j)!} \left(\frac{\mu}{\lambda} \right)^j \right)} \end{aligned}$$

$$\begin{aligned} \Delta^2 P_0(S) = & \frac{\left[\left(\sum_{j=0}^S \frac{(S+2)(S+1)S!}{(S+2-j)(S+1-j)(S-j)!} \left(\frac{\mu}{\lambda}\right)^j \right) \left(\sum_{j=0}^{S+1} \frac{(S)!}{(S-j)!} (j) \left(\frac{\mu}{\lambda}\right)^j \right) \right. \\ & - \left. \left(\sum_{j=0}^S \frac{(S)!}{(S-j)!} \left(\frac{\mu}{\lambda}\right)^j \right) \left(\sum_{j=0}^{S+1} \frac{(S+1)S!}{(S+1-j)(S-j)!} (j) \left(\frac{\mu}{\lambda}\right)^j \right) \right] \\ & + \left[(S+2)! \left(\frac{\mu}{\lambda}\right)^{S+2} \left(\sum_{j=0}^{S+1} \frac{(S)!}{(S-j)!} (j) \left(\frac{\mu}{\lambda}\right)^j - \sum_{j=0}^S \frac{(S)!}{(S-j)!} \left(\frac{\mu}{\lambda}\right)^j \right) \right] \\ & + \left[(S+2)! \left(\frac{\mu}{\lambda}\right)^{S+1} \sum_{j=0}^{S+1} \frac{(S)!}{(S-j)!} (j) \left(\frac{\mu}{\lambda}\right)^j \right] \\ & \left. \left(\sum_{j=0}^{S+2} \frac{(S+2)!}{(S+2-j)!} \left(\frac{\mu}{\lambda}\right)^j \right) \left(\sum_{j=0}^{S+1} \frac{(S+1)!}{(S+1-j)!} \left(\frac{\mu}{\lambda}\right)^j \right) \left(\sum_{j=0}^S \frac{S!}{(S-j)!} \left(\frac{\mu}{\lambda}\right)^j \right) \right] \end{aligned}$$

It can be easily concluded that all the terms in brackets are positive, i.e.

$\Delta^2 P_0(S) > 0$. Thus, $C(S, \bar{0})$ is a convex function of S .

Proof of Theorem 3.5:

We provide the proof only for the primary model, because the same steps are applicable for the proof of the models with order-cancellation flexibility.

In order to prove the theorem we will first show that the following two conditions hold (see Cavazos-Cadena and Sennott (1992)):

- 1) there exists a stationary policy π inducing an irreducible and ergodic Markov chain with finite average cost g^π
- 2) the set $G = \{(x, y) \mid \text{there exists a decision } u \in \{y, \dots, s\} \text{ such that } hx + pu < g^\pi\}$ is finite.

Let us consider (S, \bar{K}) policy described in Section 3.4.1 where S , which is finite, denotes the fixed target inventory level and $\bar{K} = (K_1, K_2, \dots, K_n)$ denotes the fixed three-

shold rationing levels. The corresponding continuous time process is depicted in Figure 1. If we apply the uniformization technique and obtain an equivalent discrete time model, there will be self-transitions and the chain will be irreducible positive recurrent and aperiodic with the stationary distribution given in (16). Moreover, the average cost of (S, \vec{K}) policy given in (19) is finite for finite system parameters. Hence, condition 1 holds. It is easy to verify that condition 2 holds since the stage cost $(hx + pu)$ is increasing in x and u . Therefore, the cardinality of the set G is finite.

Under the above conditions, Weber and Stidham (1987) show that the optimal average cost g^* , which is the same for all initial states, and a function $r(x, y)$ satisfy Bellman's equation for the average cost problem (w.l.o.g assume $x > 0$):

$$g^* + r(x, y) = \min_{s \geq u \geq y} \left\{ \begin{array}{l} hx + pu + (s - u)\mu r(x, y) + u\mu r(x + 1, u - 1) \\ + \sum_{i=1}^n \lambda_i \min \{r(x - 1, u), c_i + r(x, u)\} \end{array} \right\}$$

The minimizer of the right-hand side is an optimal stationary policy (see page 386 of Bertsekas (2000)). Hence, the structural properties of the optimal policy under average cost criterion are determined through $r(x, y)$ just in the same way as the properties of discounted cost optimal policy are determined through $J(x, y)$ (see (2) and (3)). Thus, the optimal stationary policy under average cost criterion possesses all the properties of the optimal policy under discounted cost criterion.

Chapter 4

M/ E_k /s Model with multiple-demand classes and lost sales

In this chapter we study the Erlangian-servers extension of the $M/M/s$ model. This extension also allows us to derive conclusions for the systems with deterministic processing times. In order to develop a general method for the analysis of the $M/E_k/s$ model, we first consider the $M/E_k/2$ model. Using the results for the two-server case, we then extend the model to the multi-server case. We also code the value iteration algorithm and obtain numerical results and insights. For the two-server case, we present the model formulation in Section 4.1 and the analysis in Section 4.2. In Section 4.2, we also discuss how the formulation for $M/E_k/2$ model can be used to obtain the formulation for the general model. The main body of the chapter concludes with a numerical study. The proofs of the theoretical results presented in this chapter are provided in the Appendix (Section 4.3).

4.1 Model Formulation

The setting considered in this part is same as the one stated in Section 2.1 with the exceptions that $s = 2$ and the production times consist of $k \geq 2$ identical and independent exponentially distributed stages. The expected production time of a part is

$\frac{1}{\mu}$, i.e., the expected length of each stage is $\frac{1}{k\mu}$. The state of the system is defined with three variables. As in Section 2, $X(t)$ is the inventory level at time t . Let $S_1(t)$ and $S_2(t)$ denote the stage of the current production at servers 1 and 2, respectively, at time t . We have $S_i(t) \in \{0, 1, 2, \dots, k\}$, $i \in \{1, 2\}$, where stage 0 denotes that there is no production at the server at time t . At state $(X(t), S_1(t), S_2(t))$, the production control specifies whether to initiate production at server- i if $S_i(t) = 0$, and the rationing control specifies whether to satisfy an arriving demand or not. Ha (2000) studies the $M/E_k/I$ model and formulate the problem with a single state variable called the work storage level which is the total number of completed production stages. Specifically, $Y(t) = I(t) + kX(t)$ is the work storage level defined in Ha (2000), where $I(t)$ is the number of stages completed for the job under production at time t (a part in the inventory already completed all of the k stages of the production). However, in a setting with parallel servers, the controller should make the production decision for a specific server by considering the current states (stages) of the other servers. Thus, we keep track of the state of each server.

Following the same approach discussed in Section 2.1, we obtain an equivalent discrete-time formulation of the problem using the uniformization technique. Let us define the uniform rate as $\nu = \sum_{i=1}^n \lambda_i + 2k\mu$ and (without loss of generality) rescale the time and assume that $\alpha + \nu = 1$. Then, the optimal cost-to-go function can be expressed as

$$J(x, s_1, s_2) = hx + k\mu T_{P_1}(x, s_1, s_2) + k\mu T_{P_2}(x, s_1, s_2) + T_R(x, s_1, s_2) \quad (1)$$

$$\text{where } T_{P_1}(x, s_1, s_2) = \begin{cases} \min \{J(x, 0, s_2), p + J(x, 2, s_2)\} & , s_1 = 0 \\ J(x, s_1 + 1, s_2) + p & , k > s_1 > 0, \\ J(x + 1, 0, s_2) + p & , s_1 = k \end{cases}$$

$$T_{P_2}(x, s_1, s_2) = \begin{cases} \min \{J(x, s_1, 0), p + J(x, s_1, 2)\} & , s_2 = 0 \\ J(x, s_1, s_2 + 1) + p & , k > s_2 > 0, \\ J(x + 1, s_1, 0) + p & , s_2 = k \end{cases}$$

$$T_R(x, s_1, s_2) = \sum_{i=1}^n T_{R_i}(x, s_1, s_2) \text{ and for } i \in \{1, 2, \dots, n\},$$

$$T_{R_i}(x, s_1, s_2) = \begin{cases} \lambda_i \min \{J(x - 1, s_1, s_2), c_i + J(x, s_1, s_2)\} & , x > 0 \\ \lambda_i (c_i + J(0, s_1, s_2)) & , x = 0 \end{cases}$$

The operators T_{P_1} and T_{P_2} correspond to the production decisions at servers 1 and 2, respectively. When a specific production channel is idle, the controller decides whether to initiate production or keep the channel idle. On the other hand, when there is a job under production, the controller has nothing to do; cancellation of previously placed production orders is not allowed. As in Chapter 3, the operator T_R corresponds to the rationing decision for class i .

We show in Lemma 4.1 that the optimal cost function is symmetric with respect to the production stages of the servers.

Lemma 4.1. For any given inventory level x and for all s_1, s_2 values, we have

$$J(x, s_1, s_2) = J(x, s_2, s_1).$$

Lemma 4.1 implies that $T_{P_2}(x, s_1, s_2) = T_{P_1}(x, s_2, s_1)$ because,

$$T_{P_2}(x, s_1, s_2) = \begin{cases} \min\{J(x, s_1, 0), p + J(x, s_1, 2)\} \\ = \min\{J(x, 0, s_1), p + J(x, 2, s_1)\} & , s_2 = 0 \\ J(x, s_1, s_2 + 1) + p = J(x, s_2 + 1, s_1) + p & , k > s_2 > 0 \\ J(x + 1, s_1, 0) + p = J(x + 1, 0, s_1) + p & , s_2 = k \end{cases}$$

which is equal to $T_P(x, s_2, s_1)$. Then, by defining a single production operator $T_P(x, s_1, s_2) = T_{P_1}(x, s_1, s_2) + T_{P_2}(x, s_2, s_1) = T_{P_1}(x, s_1, s_2) + T_{P_2}(x, s_2, s_1)$, we can rewrite (1) as

$$J(x, s_1, s_2) = hx + k\mu T_P(x, s_1, s_2) + T_R(x, s_1, s_2) \quad (2)$$

where

$$T_P(x, s_1, s_2) = \begin{cases} 2 \min\{J(x, 0, 0), p + J(x, 2, 0)\} & , s_1 = s_2 = 0 \\ \min\{J(x, s_i, 0), p + J(x, s_i, 2)\} \\ + p + J(x, s_i + 1, 0) & , k > s_i > 0, s_j = 0, i \in \{1, 2\}, j \in \{1, 2\} \setminus \{i\} \\ \min\{J(x, k, 0), p + J(x, k, 2)\} \\ + p + J(x + 1, 0, 0) & , s_i = 0, s_j = k, i \in \{1, 2\}, j \in \{1, 2\} \setminus \{i\} \\ J(x, s_1 + 1, s_2) + J(x, s_2 + 1, s_1) \\ + 2p & , k > s_1 > 0, k > s_2 > 0 \\ J(x, s_i + 1, k) + J(x + 1, s_i, 0) \\ + 2p & , k > s_i > 0, s_j = k, i \in \{1, 2\}, j \in \{1, 2\} \setminus \{i\} \\ 2J(x + 1, k, 0) + 2p & , s_1 = s_2 = k \end{cases}$$

4.2 Characterization of Optimal Production and Rationing Policies

As in Section 3.1.2, we identify some structural properties of the optimal cost function. Using these properties we characterize the optimal production and rationing policies. Let us first define the following operators on a function $v(x, s_1, s_2)$ such that $v(x, s_1, s_2) = v(x, s_2, s_1)$:

$$\Delta^x v(x, s_1, s_2) = v(x+1, s_1, s_2) - v(x, s_1, s_2)$$

$$\Delta^{s_1} v(x, s_1, s_2) = v(x, s_1+1, s_2) - v(x, s_1, s_2)$$

$$\Delta^{xx} v(x, s_1, s_2) = \Delta^x v(x+1, s_1, s_2) - \Delta^x v(x, s_1, s_2)$$

$$\Delta^{xs_1} v(x, s_1, s_2) = \Delta^{s_1} v(x+1, s_1, s_2) - \Delta^{s_1} v(x, s_1, s_2)$$

$$\Delta^{s_2 s_1} v(x, s_1, s_2) = \Delta^{s_1} v(x, s_1, s_2+1) - \Delta^{s_1} v(x, s_1, s_2)$$

We define \mathcal{G} as a set of functions on the integers such that if $v \in \mathcal{G}$, then for any (x, s_1, s_2)

$$\Delta^{s_2 s_1} v(x, s_1, s_2) \geq 0, \quad (3)$$

$$\Delta^{s_1} v(x+1, s_1, 0) - \Delta^{s_1} v(x, s_1, 0) \geq 0, \quad (4)$$

$$v(x+2, s_1, s_2) - v(x+1, s_1, s_2) \geq v(x+1, \bar{s}_1, \bar{s}_2) - v(x, \bar{s}_1, \bar{s}_2), \text{ for any } \bar{s}_1 \text{ and } \bar{s}_2, \quad (5)$$

$$\Delta^x v(x, s_1+1, s_2) \geq \Delta^x v(x, s_1, s_2+1), \text{ for } s_1 \geq s_2 > 0. \quad (6)$$

Note that by symmetry (3) is equivalent to saying $\Delta^{s_1 s_2} v(x, s_1, s_2) \geq 0$ and (4) is to $\Delta^{s_2} v(x+1, 0, s_2) - \Delta^{s_2} v(x, 0, s_2) \geq 0$.

At this stage, we conjecture that the optimal cost function is an element of the function space \mathcal{G} .

Conjecture 4.1. $J \in \mathcal{G}$, that is $J(x, s_1, s_2)$ satisfies (3), (4), (5) and (6).

In the Appendix (Section 4.3) we prove that $J(x, s_1, s_2)$ satisfies (3). However, we are still working on the proofs of the other properties. Based on Conjecture 4.1, we state Theorem 3.4.1 that characterizes the optimal policies.

Theorem 4.1. If $J \in \mathcal{G}$, then

- i. $\Delta^{xx}J(x, s_1, s_2) \geq 0$, i.e., $J(x, s_1, s_2)$ is x -convex.
- ii. For an idle server- i , $i \in \{1, 2\}$, there exists a threshold inventory level $N_x^i(s_j)$, which is a function of the state of server- j , $j \in \{1, 2\} \setminus \{i\}$, such that it is optimal to trigger the production at server- i below $N_x^i(s_j)$ and leave server idle otherwise. Moreover, $N_x^i(s_j)$ is non-increasing in s_j and $N_x^1(s) = N_x^2(s)$.
- iii. For an idle server- i , $i \in \{1, 2\}$, there exists a threshold state of server- j , $j \in \{1, 2\} \setminus \{i\}$, $N_{s_j}^i(x)$, which is a function of the inventory level, such that it is optimal to trigger the production at server- i below $N_{s_j}^i(x)$ and leave server idle otherwise. Moreover, $N_{s_j}^i(x)$ is non-increasing in x and $N_{s_2}^1(x) = N_{s_1}^2(x)$.
- iv. It is always optimal to satisfy a class 1 demand when there is stock on hand. That is, $T_{R_1}(x, y) = \lambda_1 J(x-1, y)$.

- v. There exists a threshold inventory level $K_x^d(s_1, s_2)$ for class $d \geq 2$, which is a function of the states of the servers, such that it is optimal to satisfy a class- d demand above $K_x^d(s_1, s_2)$ and reject it otherwise. $K_x^d(s_1, s_2)$ is non-increasing in s_1 and s_2 , and for $s_1 \geq s_2$, $K_x^d(s_1 + 1, s_2) \leq K_x^d(s_1, s_2 + 1)$, which means that the outstanding order that is more close to arrive has more value in terms of the rationing decision. Moreover,
- $$K_x^n(s_1, s_2) \geq K_x^{n-1}(s_1, s_2) \geq \dots \geq K_x^2(s_1, s_2) \quad \text{and} \quad K_x^d(s_1, s_2) = K_x^d(s_2, s_1) \quad \text{also hold.}$$
- vi. There exists a threshold stage of server- i , $i \in \{1, 2\}$, $K_{s_i}^d(x, s_j)$ for class $d \geq 2$, which is a function of the inventory level and stage of server- j , $j \in \{1, 2\} \setminus \{i\}$, such that it is optimal to satisfy a class- d demand above $K_{s_i}^d(x, s_j)$ and reject it otherwise. $K_{s_i}^d(x, s_j)$ is non-increasing in x and s_j . Moreover, $K_{s_i}^n(x, s_j) \geq K_{s_i}^{n-1}(x, s_j) \geq \dots \geq K_{s_i}^2(x, s_j)$, $K_{s_1}^d(x, s) = K_{s_2}^d(x, s)$.

Theorem 4.1 implies that the optimal production and rationing policies are state-dependent and monotone. These results are also applicable to multi-server case, i.e., to the $M/E_k/s$ model, because the model formulation given in Section 4.1 can be directly extended to this case. For the general case with $s \geq 2$ servers, with uniform rate $\nu = \sum_{i=1}^n \lambda_i + sk\mu$, the optimality equation - the modified version of Equation 4.1- can be expressed as

$$J(x, s_1, \dots, s_s) = hx + \sum_{i=1}^s k\mu T_{P_i}(x, s_1, \dots, s_s) + T_R(x, s_1, \dots, s_s) \quad (7)$$

$$\text{where } T_{P_i}(x, s_1, \dots, s_i, \dots, s_s) = \begin{cases} \min \left\{ \begin{array}{l} J(x, s_1, \dots, s_{i-1}, 0, \dots, s_s), \\ p + J(x, s_1, \dots, s_{i-1}, 2, \dots, s_s) \end{array} \right\} & , s_i = 0 \\ J(x, s_1, \dots, s_{i-1}, s_i + 1, \dots, s_s) + p & , k > s_i > 0 \\ J(x, s_1, \dots, s_{i-1}, 0, \dots, s_s) + p & , s_i = k \end{cases}$$

Moreover, by following the same steps of the proof of Lemma 4.1 we can conclude that the optimal cost function is symmetric with respect to the production stages of the servers, i.e.,

$$J(x, s_1, \dots, s_{i-1}, s_i, s_{i+1}, \dots, s_{j-1}, s_j, s_{j+1}, \dots, s_s) = J(x, s_1, \dots, s_{i-1}, s_j, s_{i+1}, \dots, s_{j-1}, s_i, s_{j+1}, \dots, s_s).$$

Therefore, we can define a single production operator and rewrite the optimality equation with this operator as we did at the end of Section 4.1. For this formulation, we can conjecture on properties 4.3 through 4.6 for any two of the servers.

We illustrate the properties of the optimal policies for a two-class system with $(s, k, \mu, \lambda_1, \lambda_2, h, p, c_1, c_2) = (3, 3, 4, 4, 1, 1, 1, 10, 2)$ in Tables 4.1 and 4.2. For the discounted cost criterion with $\alpha = 0.6$, Table 4.1 and Table 4.2 show the optimal production and rationing policies, respectively. For this three-server setting, Table 4.1 illustrates the optimal production decisions at each state $(x, s_1, s_2, 0)$ for $x \leq 4$. While a zero cell value indicating that it is optimal to leave server-3 idle, one indicates that server-3 should be activated. Table shows that the production decision at server-3 depends on the stages of the other servers and the inventory level (1 corresponds to activating the server and 0 to leaving idle). Due to the symmetry, we would also observe the same policy for any of the other two servers.

Table 4.2 is drawn for $x = 3$. A zero in the cell corresponding to anyone of the states indicates that an arriving class 2 demand should be rejected, and a one indi-

cates that the demand should be satisfied. We see that threshold rationing levels are non-increasing in the stages of the servers. Furthermore, while an arriving class 2 demand is not satisfied at states $(x, s_1, s_2, s_3) = (3, 1, 2, 1)$ and $(x, s_1, s_2, s_3) = (3, 2, 2, 1)$, it is satisfied at $(x, s_1, s_2, s_3) = (3, 1, 3, 1)$. That means, a unit increase in the stage of server 1 does not change the rationing level. However, a unit increase at stage of server 2 (instead of server 1) makes it possible to satisfy the demand. This fact, which is also stated in Theorem 4.1, supports the rationale behind the dynamic rationing policy that we propose in Chapter 5: the outstanding order that is more close to arrive has more value in terms of the rationing decision.

It should be also noted that by following almost the same steps of the proof of Theorem 3.5, it is straightforward to show that the properties of the optimal policies that are stated in Theorem 4.1 are also preserved under the average cost criterion.

Table 4.1 Optimal Production Policy under Discounted Cost Criterion

$s_1 \backslash s_2$	0	1	2	3	$s_1 \backslash s_2$	0	1	2	3
0	1	1	1	1	0	1	1	1	1
1	1	1	1	1	1	1	1	1	1
2	1	1	1	1	2	1	1	1	1
3	1	1	1	1	3	1	1	1	0
$x = 0$					$x = 1$				
$s_1 \backslash s_2$	0	1	2	3	$s_1 \backslash s_2$	0	1	2	3
0	1	0	0	0	0	0	0	0	0
1	0	0	0	0	1	0	0	0	0
2	0	0	0	0	2	0	0	0	0
3	0	0	0	0	3	0	0	0	0
$x = 2$					$x = 3$				

Table 4.2 Optimal Rationing Policy under Discounted Cost Criterion for $x = 3$.

$s_1 \backslash s_2$	0	1	2	3	$s_1 \backslash s_2$	0	1	2	3
0	0	0	0	0	0	0	0	0	0
1	0	0	0	0	1	0	0	0	1
2	0	0	0	0	2	0	0	0	1
3	0	0	0	0	3	0	1	1	1
$s_3 = 0$					$s_3 = 1$				
$s_1 \backslash s_2$	0	1	2	3	$s_1 \backslash s_2$	0	1	2	3
0	0	0	0	0	0	0	0	0	0
1	0	0	0	1	1	0	1	1	1
2	0	0	1	1	2	0	1	1	1
3	0	1	1	1	3	0	1	1	1
$s_3 = 2$					$s_3 = 3$				

4.3 Appendix

Proof of Lemma 4.1: Let T be the optimization operator. That is,

$$T(J(x, s_1, s_2)) = hx + k\mu T_p(x, s_1, s_2) + T_R(x, s_1, s_2).$$

In order to show that the property holds, it is enough to show that the optimization operator preserves the property. Suppose $J(x, s_1, s_2) = J(x, s_2, s_1)$ holds and let us define $\delta = T(J(x, s_1, s_2)) - T(J(x, s_2, s_1))$. We will show that $\delta = 0$.

$$\delta = k\mu(T_p(x, s_1, s_2) - T_p(x, s_2, s_1)) + \sum_{i=1}^n (T_{R_i}(x, s_1, s_2) - T_{R_i}(x, s_2, s_1)).$$
 Let us write

$$\delta = \delta_p + \sum_{i=1}^n \delta_{R_i} \quad \text{where } \delta_p = k\mu(T_p(x, s_1, s_2) - T_p(x, s_2, s_1)),$$

$\delta_{R_i} = (T_{R_i}(x, s_1, s_2) - T_{R_i}(x, s_2, s_1))$. Now, we will separately show that δ_p and δ_{R_i} are both zero.

Let us start with showing $\delta_p = 0$:

For the cases $k > s_1 > 0, k > s_2 > 0$; $k > s_i > 0, s_j = k, i \in \{1, 2\}, j \in \{1, 2\} \setminus \{i\}$; and $s_1 = s_2 = k$ the production decision operator $T_p(x, s_1, s_2)$ only contains the optimal cost functions and the unit production cost (there is no minimization operator in contrast to the remaining cases). Therefore, using the hypothesis (supposition) we directly conclude that $\delta_p = 0$.

For the other cases the definition of $T_p(x, s_1, s_2)$ includes a minimization operator, but due to the hypothesis and the symmetric region (with respect to s_1, s_2) of each case we again directly conclude that $\delta_p = 0$.

Now, we will show that $\delta_{R_i} = 0$:

By the hypothesis,
$$T_{R_i}(x, s_1, s_2) = \begin{cases} \lambda_i \min \{J(x-1, s_1, s_2), c_i + J(x, s_1, s_2)\} & , x > 0 \\ \lambda_i (c_i + J(0, s_1, s_2)) & , x = 0 \end{cases}$$

is equivalent to
$$T_{R_i}(x, s_1, s_2) = \begin{cases} \lambda_i \min \{J(x-1, s_2, s_1), c_i + J(x, s_2, s_1)\} & , x > 0 \\ \lambda_i (c_i + J(0, s_1, s_2)) & , x = 0 \end{cases}$$

which is exactly $T_{R_i}(x, s_2, s_1)$. Therefore, $\delta_{R_i} = 0$.

Thus, we conclude that $J(x, s_1, s_2) = J(x, s_2, s_1)$.

Proof of Property (3) – A partial proof of Conjecture 4.1: We will show that the optimization operator preserves $\Delta^{s_1 s_2} J(x, s_1, s_2) \geq 0$. Suppose $\Delta^{s_1 s_2} J(x, s_1, s_2) \geq 0$ holds. Now,

$$\Delta^{s_1 s_2} T(J(x, s_1, s_2)) = k\mu \Delta^{s_1 s_2} T_p(x, s_1, s_2) + \sum_{i=1}^n \Delta^{s_1 s_2} T_{R_i}(x, s_1, s_2).$$

Showing that $\Delta^{s_1 s_2} T_p(x, s_1, s_2) \geq 0$:

Case1. $s_1 = s_2 = 0$.

In this case,

$$\Delta^{s_1 s_2} T_p(x, 0, 0) = 2 \left(\begin{array}{l} J(x, 2, 1) - J(x, 2, 0) - \min \{J(x, 0, 1), p + J(x, 2, 1)\} \\ + \min \{J(x, 0, 0), p + J(x, 2, 0)\} \end{array} \right). \text{ Then, we}$$

should analyze the following sub-cases:

- i. Produce at server-1 at $s_2 = 0$ and at $s_2 = 1$:

Then,

$$\Delta^{s_1 s_2} T_p(x, 0, 0) = 2(J(x, 2, 1) - J(x, 2, 0) - p - J(x, 2, 1) + p + J(x, 2, 0)) = 0$$

- ii. Do not produce neither at $s_2 = 0$ and nor at $s_2 = 1$:

Then, $\Delta^{s_1, s_2} T_p(x, 0, 0) = 2(J(x, 2, 1) - J(x, 2, 0) - J(x, 0, 1) + J(x, 0, 0)) \geq 0$

holds by the hypothesis, because property (3) implies

$$J(x, s_1 + 2, s_2 + 1) - J(x, s_1 + 2, s_2) \geq J(x, s_1, s_2 + 1) - J(x, s_1, s_2)$$

iii. Produce at $s_2 = 0$, do not produce at $s_2 = 1$:

Then,

$$\begin{aligned} \Delta^{s_1, s_2} T_p(x, 0, 0) &= 2(J(x, 2, 1) - J(x, 2, 0) - J(x, 0, 1) + p + J(x, 2, 0)) \\ &= 2(J(x, 2, 1) + p - J(x, 2, 0)) \end{aligned}$$

which is greater or equal to zero because it is optimal not to produce at $s_2 = 1$.

Here note that, property (3), i.e., $\Delta^{s_1, s_2} J(x, s_1, s_2) \geq 0$, guarantees that it is not possible (optimal) to produce at any $s_2' > s_2$ level while not producing at s_2 .

Case2. $s_1 = 0, k - 1 > s_2 > 0$.

In this case,

$$\begin{aligned} \Delta^{s_1, s_2} T_p(x, 0, s_2) &= J(x, 2, s_2 + 1) + J(x, s_2 + 2, 1) - J(x, 2, s_2) - J(x, s_2 + 1, 1) \\ &\quad - \left(\begin{aligned} &J(x, s_2 + 2, 0) - J(x, s_2 + 1, 0) \\ &+ \min \{ J(x, 0, s_2 + 1), p + J(x, 2, s_2 + 1) \} \\ &- \min \{ J(x, 0, s_2), p + J(x, 2, s_2) \} \end{aligned} \right) \end{aligned}$$

By the hypothesis, we

have $J(x, s_2 + 2, 1) - J(x, s_2 + 2, 0) - J(x, s_2 + 1, 1) + J(x, s_2 + 1, 0) \geq 0$. Therefore, we only need to show

$$\left(\begin{aligned} &J(x, 2, s_2 + 1) - J(x, 2, s_2) \\ &- \min \{ J(x, 0, s_2 + 1), p + J(x, 2, s_2 + 1) \} + \min \{ J(x, 0, s_2), p + J(x, 2, s_2) \} \end{aligned} \right) \geq 0. \quad \text{It}$$

is easy to show that this inequality holds by following the same steps of ‘‘Case 1’’.

Case3. $k-1 > s_1 > 0, k-1 > s_2 > 0$.

In this case,

$$\begin{aligned} \Delta^{s_1, s_2} T_p(x, s_1, s_2) &= (J(x, s_1 + 2, s_2 + 1) + J(x, s_2 + 2, s_1 + 1) - J(x, s_1 + 2, s_2) - J(x, s_2 + 1, s_1 + 1)) \\ &\quad - (J(x, s_1 + 1, s_2 + 1) + J(x, s_2 + 2, s_1) - J(x, s_1 + 1, s_2) - J(x, s_2 + 1, s_1)) \end{aligned}$$

is greater or equal to zero by the hypothesis (group first two and last two terms of each parenthesis).

For the other cases $\Delta^{s_1, s_2} T_p(x, s_1, s_2) \geq 0$ also holds due to the symmetric definition of $T_p(x, s_1, s_2)$.

Showing that $\Delta^{s_1, s_2} T_R(x, s_1, s_2) \geq 0$:

For $x > 0$,

$$\begin{aligned} \Delta^{s_1, s_2} T_R(x, s_1, s_2) &= \left(\min \{ J(x-1, s_1 + 1, s_2 + 1), c_i + J(x, s_1 + 1, s_2 + 1) \} \right) \\ &\quad - \left(\min \{ J(x-1, s_1 + 1, s_2), c_i + J(x, s_1 + 1, s_2) \} \right) \\ &\quad - \left(\min \{ J(x-1, s_1, s_2 + 1), c_i + J(x, s_1, s_2 + 1) \} \right) \\ &\quad - \left(\min \{ J(x-1, s_1, s_2), c_i + J(x, s_1, s_2) \} \right) \end{aligned}$$

Case1. Satisfy an arriving demand at levels s_1 and s_2 .

This means it is optimal to satisfy at all levels $(\bar{s}_1, \bar{s}_2) \geq (s_1, s_2)$ by the property (4).

Then,

$$\Delta^{s_1, s_2} T_R(x, s_1, s_2) = \Delta^{s_1, s_2} J(x-1, s_1, s_2) \geq 0 \text{ holds by the hypothesis.}$$

Case2. Do not satisfy an arriving demand at levels $(s_1 + 1)$ and $(s_2 + 1)$.

This means it is optimal not to satisfy at all levels $(\bar{s}_1, \bar{s}_2) \leq (s_1, s_2)$ by the property

(4). Then,

Chapter 4 $M/E_k/s$ Model With Multiple Demand Classes and Lost Sales

$\Delta^{s_1, s_2} T_{R_i}(x, s_1, s_2) = \Delta^{s_1, s_2} J(x, s_1, s_2) \geq 0$ holds by the hypothesis.

Therefore, we conclude that $\Delta^{s_1, s_2} T_{R_i}(x, s_1, s_2) \geq 0$ and so $\Delta^{s_1, s_2} T(J(x, s_1, s_2)) \geq 0$.

Chapter 5

Dynamic Rationing Policy for Continuous Review Inventory Systems

For inventory systems with distinct customer classes demanding the same item, stock rationing is a well-known tool to differentiate customer classes. Different customers may have different service level requirements or different shortage costs. In such cases, stock rationing allows prioritization of demand classes in order to provide different levels of service and to achieve higher operational efficiency. It is possible to maintain high service levels for certain demand classes while keeping inventory costs at bay by providing lower service levels to certain other demand classes. Demand classes are categorized on the basis of their shortage costs. The highest priority class is the one with the largest unit shortage cost, and the lowest priority class has the smallest unit shortage cost.

It is possible to come up with different rationing policies. Yet, the mechanism through which any rationing policy is implemented is to stop serving a lower priority class when the on-hand inventory drops below a certain critical level. The unsatisfied demands are either backordered or lost depending on the nature of the system. Under the critical level only higher priority classes are served and this results in higher service levels for those classes. If there are more than two demand classes, then there

has to be more than a single critical level. The critical level for the highest priority class may be assumed to be zero. The critical levels may change dynamically according to the number and the ages of outstanding orders or static threshold levels may be used.

In backordering environments, to completely define the stock rationing policy the way that the backorders are cleared should also be defined. The clearing mechanism specifies how the replenishment orders should be allocated between increasing the stock level and clearing the backorders. The natural way to perform the clearing is to employ the same critical levels for clearing the backorders, i.e., the backorders for a certain customer class are not cleared until the inventory level reaches to the critical level which is associated for that customer class. This mechanism is referred as the priority clearing in the literature.

In the continuous review setting, there are only a couple of studies that consider a dynamic adjustment of the critical levels for rationing. Since the analysis of rationing systems is complicated even for the static policy, this is a difficult setting. Under the at-most-one-outstanding-order assumption, Teunter and Haneveld (2008) consider a dynamic rationing policy for the backordering environment and Melchioris (2003) considers a so-called time remembering policy for the lost sales case. Except these works, the common practice in the literature is to assume static (time invariant) rationing levels with clearing mechanisms other than the priority clearing. Priority clearing and/or adjusting the critical levels dynamically complicates the analysis considerably.

In this chapter, we propose a dynamic rationing policy together with the associated dynamic priority clearing mechanism for continuous-review backordering systems with constant lead-time and unit Poisson demands for two demand classes. The replenishment orders are placed according to the (Q, r) policy. The (Q, r) policy dictates that a batch of Q units is ordered whenever inventory position hits the reorder

level r . The proposed policy uses the age information for all the outstanding orders in order to decide whether a lower priority demand should be satisfied instantaneously or should be backordered (or lost depending on the setting). More specifically, the policy incorporates the outstanding replenishment orders into the on-hand inventory as if they arrive continuously within the lead-time. With the currently available information technologies, it is easy to monitor the status of the outstanding orders and to incorporate the information that they carry into the decision mechanism. Therefore, the proposed policy should not suffer from implementation issues in today's environment.

We conduct a simulation study to evaluate the performance of the proposed policy. Since the analytical evaluation of the policy is not tractable without simplifying assumptions, simulation is the only available tool. In spite of the popularity of simulation in the broader area of operations research and operations management, there is no work in the stock rationing literature that uses simulation to analyze complex, dynamic policies, which outperform static ones. One should also point out that simulation can also be used to estimate any long-run performance measures of the inventory systems under static rationing that are discussed in the literature. Still, authors usually resort to simulation not for direct performance analysis of their models but only for testing their results (e.g., Dekker et al. (1998), Deshpande et al. (2003)). However, it is possible to obtain the performance measures in any desired confidence interval via simulation, and thus conduct the performance evaluation of the inventory policies under scrutiny.

The rest of the chapter is organized as follows. In Section 5.1, we discuss the characteristics that the optimal policy should exhibit and proposes two new lower bounds on the optimal policy. In Section 5.2, we introduce a new class of dynamic rationing policies and discuss the properties a good dynamic rationing policy should exhibit. Based on these properties, we develop a new policy called *Rationing with*

Exponential Replenishment Flow (RERF). In Section 5.3, we compare the performance of *RERF* with the static rationing policy and quantify the gain obtained through *RERF* under different scenarios by simulation.

5.1 Dynamic Rationing

In the stock rationing literature, it is well documented that the performance of the static rationing policy can be improved by utilizing the information on the status of the outstanding replenishment orders. For a given set of policy parameters, although the static policy would not allow it, one should prefer to satisfy an arriving lower priority customer instantaneously if an outstanding order is about to arrive and increase the inventory level. A similar scenario can also be entertained for the back-order clearing mechanism of the rationing policy; it may be profitable to clear some backorders of lower priority class before the inventory level is increased above the threshold rationing level. Therefore, it should be clear that the optimal rationing policy should be a dynamic policy that allows the threshold rationing level to change in time depending on the number and ages of outstanding orders. However, the characterization of this optimal policy structure would be quite hard, if possible at all. Such a policy would have to depend on the time-to-arrive of all outstanding orders, which is a random variable. Thereby, it is very hard to analyze, if possible, any dynamic policy without simplifying assumptions.

Teunter and Haneveld (2008) consider dynamic stock rationing but the analysis is based on the assumption of at-most-one-outstanding-order and is computationally tractable only for limited settings. The difficulty in the analysis of dynamic policies mainly arises from the fact that one should incorporate the ages (or the time-to-arrive values) of all outstanding orders into the system state definition. Moreover, since the number of outstanding orders changes in time depending on the realizations of the demand processes and the policy parameters, the size of the state vector itself is a

random variable. Thus, we evaluate the performance of the policy we propose, which is called *Rationing with Exponential Replenishment Flow (RERF)*, via simulation.

Although the structure of the optimal rationing policy is unknown, we would like to tell something about the quality of the policy we suggest with respect to the optimal policy. Since it is not possible to do the performance evaluation of an unknown policy, we develop two lower bounds for any possible rationing policy (static or dynamic). Using these bounds, we are able to assess how much our policy realizes out of the maximum possible potential for the rationing policies. Moreover, these bounds also point out the settings at which there is no meaningful gain to be obtained by applying any kind of dynamic policy.

Consider the inventory system with parameters $\lambda_i, \pi_i, \hat{\pi}_i, i \in \{1, 2\}$, and h, A , i.e., the arrival rates, the unit backordering costs, the time-dependent backordering costs, the holding cost rate and the fixed ordering cost. Without loss of generality we assume $\pi_1 > \pi_2$ and $\hat{\pi}_1 > \hat{\pi}_2$, which means class 1 is the higher priority demand class. We denote this system with O . Based on this system; we can construct two related inventory systems, each of which is subject to a single demand class in order to obtain lower bounds. The first system we propose is the one with demand rate $\lambda_1 + \lambda_2$, and backorder costs $\pi_2, \hat{\pi}_2$. The second one is another single-demand class system with demand rate λ_1 and backorder costs $\pi_1, \hat{\pi}_1$. We denote the first system with N_1 , and the second one with N_2 . For each of these new systems we assume that the holding cost rate and the fixed ordering cost are the same with the original two demand class system. Since N_1 is constructed using the minimum of the backorder costs of the two demand classes (while keeping the total demand rate constant), the optimal long-run average cost of N_1 is a lower bound on the long-run average cost of O using an optimal dynamic rationing policy. N_2 constitutes another lower bound, since

class 2 is completely disregarded and it is not possible for any rationing policy to operate without experiencing any cost related to class 2. Note that for both systems, since there is only a single demand class, there is no need for any kind of rationing. To obtain the above mentioned lower bounds, one only needs to optimize N_1 and N_2 with respect to the parameters Q and r .

Any one of the two bounds may be tighter than the other for a given problem instance. Therefore, the strategy is computing both of the bounds and taking the maximum as the lower bound on the cost of the optimal policy. However, it is more likely that N_1 provides a tighter bound when a high proportion of the total demand is due to class 2. Note that for a given total demand rate $\lambda = \lambda_1 + \lambda_2$, the lower bound provided by N_1 is independent of the demand mix, i.e., it is the same for all $p_1 = \lambda_1/\lambda$ values. On the other hand, when $p_1 = 1$ the lower bound provided by N_2 takes its maximum value (which is certainly greater than the lower bound obtained by N_1 , because when $p_1 = 1$ both of the systems experience the same demand rate but N_2 assumes higher backorder costs) and it decreases down to zero as p_1 goes to zero. Therefore, depending on the values of the total demand rate and the cost parameters, the lower bounds obtained by the systems N_1 and N_2 should intersect at a p_1 value in $[0, 1]$. Thus, for the values of p_1 between zero and the intersection point the lower bound provided by N_1 should be tighter and for the other possible values of p_1 (from the intersection point to 1) the lower bound obtained by N_2 should be tighter.

At this point, it should also be noted that Deshpande et al. (2003) provides a different lower bound on the cost of any rationing policy. Their bound is based on an approximate analysis for the priority clearing mechanism. However, the simulation study considered in Section 5.3 illustrates that our bound is tighter than the one suggested by Deshpande et al. (2003).

5.2 Rationing with Continuous Replenishment Flow

The current level of sophistication in information and computer technologies enables us to consider more elaborate policies. Although the analyses of these elaborate policies are difficult and mostly intractable, it is still possible to estimate the steady state behavior of the system with simulation in a reasonable amount of time with the current computer speeds. Hence, we propose a dynamic rationing policy that makes use of all the available system information at any point in time and evaluate the performance of the policy by a simulation study.

In this section, we introduce our policy with its dynamic priority clearing mechanism in a continuously reviewed inventory setting where the replenishment orders are placed according to the (Q, r) policy. The (Q, r) policy dictates that a batch of Q units is ordered whenever inventory position hits the reorder level r . Inventory position is the sum of inventory level (on-hand inventory minus the number of back-orders) and ordered units that are in the replenishment channel. This means that the inventory position increases at ordering points, whereas inventory level increases when the orders arrive. We assume two customer classes and a deterministic supply lead-time, $L > 0$. Demand arrivals are generated by two independent stationary Poisson processes with rates λ_1 and λ_2 , respectively, for class 1 and class 2 customers.

Static rationing policies, which are easier to analyze and implement, only utilize part of the information on system state. In the static policies, the replenishment and rationing decisions are based on the inventory position and the on-hand inventory level, respectively. However, unless the replenishment lead-time is memoryless, i.e., exponential, there is additional information available about the arrival times of the orders in the replenishment channel. In the case of deterministic lead-time we consider, the arrival times are exactly known once the orders are placed. Thus, a “good” rationing policy should exploit this information to extract value.

Even if the on-hand inventory level is below the (static) critical level, it is better to instantaneously satisfy a class 2 demand if the outstanding replenishment orders are about to increase the on-hand inventory level. If a replenishment arrival will happen after a short duration, the likelihood of a class 1 demand arrival before the replenishment is relatively small. Note that the higher is the class 1 arrival rate, the shorter the duration has to be. Thus, the critical level should be lowered as the ages of the replenishment orders increase.

Instead of defining the critical level as a function of the ages of outstanding orders, we opt to use a constant critical level K and a modified on-hand stock level, which are adjusted dynamically utilizing the information on the outstanding replenishment orders. Note that although the two approaches are basically equivalent, the latter better lends itself to interpretation.

We define a new class of policies called *Rationing with Continuous Replenishment Flow (RCRF)* that uses a constant critical rationing level on the modified on-hand inventory, which incorporates the outstanding orders to the on-hand inventory as if they are arriving continuously within the lead-time. The only difference between *RCRF* and the static rationing policy is the variable on which the rationing mechanism is defined.

Let $X(t)$ denote the number of outstanding replenishment orders at time t , $a_i(t)$ denote the age of i^{th} oldest outstanding replenishment order at time t where $0 \leq a_i(t) \leq L$, $1 \leq i \leq X(t)$, and $OH(t)$ denote the on-hand inventory level at time t . We define the modified on-hand inventory level at time t , $OH_m(t)$, as

$$OH_m(t) = \begin{cases} OH(t) + Q \sum_{i=1}^{X(t)} f(a_i(t)), & OH(t) > 0 \\ 0, & OH(t) = 0 \end{cases} \quad (1)$$

We do not modify the on-hand inventory when it is zero, since any incoming demand has to be backordered. Different *RCRF* policies correspond to different families of $f(\cdot)$ functions, satisfying the following properties:

1. $f(t) : [0, L] \rightarrow [0, 1]$,
2. $f(t)$ is an increasing function of t ,
3. $f(t) = 1$ for $t = L$.

(1) implies that the contribution of i^{th} outstanding order to the modified on-hand inventory level at time t is $Qf(a_i(t))$. Thus, $f(a_i(t))$ should be considered as the fraction of the replenishment order quantity that is incorporated to the on-hand inventory when the age of the outstanding order is $a_i(t)$. Obviously, as the age of the outstanding order increases this fraction should increase up-to 1 (it should be 1 when the outstanding order actually arrives, i.e., when $a_i(t) = L$).

According to *RCRF*, if there is no outstanding order at time t , we have $OH_m(t) = OH(t)$. In this circumstance, we compare the on hand stock level with the critical level to make the rationing decision as in the static critical level policy. If $OH(t) > K$, then $OH_m(t)$ is certainly greater than K . Thus an arriving class 2 demand is satisfied. If $0 < OH(t) \leq K$, the arriving class 2 demand is satisfied instantaneously provided that $OH_m(t) > K$. For the other cases, class 2 demand is backordered. If the arriving demand belongs to class 1, it is satisfied instantaneously if $OH(t) > 0$ as in the case of the static policy. Otherwise, it is backordered.

In order to complete the definition of *RCRF*, the modified priority clearing mechanism should also be specified. Suppose a replenishment order arrives at time t . After clearing class 1 backorders (if there is any), the remaining replenishment order quantity (if any remains) first used to increase the modified on-hand level up-to K and then class 2 backorders are cleared.

Any flow function $f(t)$ that satisfies the stated three properties can be used to modify the on-hand inventory according to (1). However, at this stage, we take a further step and propose to consider only the *RCRF* policies that assume strictly convex flow function. Because, the value of the information gained from the outstanding orders should diminish as we go from the oldest to the youngest order. In other words, the change in the impact of the orders on modified on-hand as we move in time, should be greater for those orders that are closer to arrive. Consequently, the difference between $OH_m(t)$ and $OH(t)$ should be mostly due to the outstanding orders that are to arrive in the very near future. This fact also emerges in Part-v of Theorem 4.1 stated in Section 4.2.

The above reasoning can be clarified with an example. Suppose that the flow function has the form $f(a_i(t)) = \left(\frac{a_i(t)}{L}\right)^n$ and at time t' there are two outstanding orders, $X(t') = 2$. Under these conditions, consider the following two cases: In the first case, $\frac{a_1(t')}{L} = 0.9$, $\frac{a_2(t')}{L} = 0.1$ and in the second one $\frac{a_1(t')}{L} = 0.6$, $\frac{a_2(t')}{L} = 0.4$. When $n = 1$, (1) returns the same $OH_m(t')$ values for both cases, because

$\sum_{i=1}^{X(t')} f(a_i(t')) = 1$ for both of them. However, when $n = 2$, i.e., when the flow function is strictly convex, $\sum_{i=1}^{X(t')} f(a_i(t'))$ is 0.82 for the first case and it is 0.52 for the second

one. Thus, $OH_m(t')$ is larger for the first case and the decision maker is more eager to satisfy a class 2 demand that arrives at t' . The unknown optimal policy would also distinguish the two cases and would be less conservative in satisfying an arriving class 2 demand for the first case than for the second case. Because, in the first case, there is an outstanding order that is close to arrive (compared to the other outstanding orders considered in the example), it is very likely that class 1 backorders in the near future can be avoided.

Another important criterion in choosing an appropriate policy from the *RCRF* class is the value of the flow function for an outstanding order that is just placed. For the flow function discussed in the above example, we have $\lim_{a_i(t) \rightarrow 0^+} f(a_i(t)) = 0$. However, the flow function should assume a positive value right at the onset, because the information that there is an outstanding order should have a nonzero value. There may be situations where we would choose to satisfy an arriving class 2 demand if we know a replenishment order has just been placed and will arrive after a lead-time period (for example if the lead-times are extremely small).

In the light of the above discussion about the intuitive criteria, we propose (2) as the flow function and define the policy *Rationing with Exponential Replenishment Flow (RERF)*. As a member of *RCRF* class *RERF* modifies the on-hand inventory according to (1).

$$f(a_i(t)) = e^{-\lambda_1(L-a_i(t))n} \quad (2)$$

In (2), $(L - a_i(t))$ is the remaining time to arrive for the i^{th} oldest outstanding order and $e^{-\lambda_1(L-a_i(t))}$ is the probability that there will be no class 1 demand arrivals until the arrival of the i^{th} oldest outstanding order. Since (2) is a decreasing function of λ_1 , i.e., the probability of zero class 1 demand arrival in a specific length of time de-

creases as λ_1 increases, if class 1 arrival rate is high the modified on-hand inventory is close to the real on-hand level and not many extra class 2 demands are satisfied compared to the static rationing. On the contrary, if λ_1 is low then much more class 2 demands can be satisfied. That is to say, *RERF* updates the on-hand inventory by considering the risk of having class 1 backorders in the future.




Figure 5.1 Exponential flow function ($L = 1, \lambda = 1$)

Power n is a parameter that is used to fine-tune the policy. It is a parameter to be optimized together with Q , r and K . As n increases the value of the information gained from the outstanding orders diminishes and the diminishing rate is faster for the younger orders. Figure 5.1 illustrates this situation for $L = 1, \lambda = 1$. As $n \rightarrow \infty$, *RERF* does not utilize any information that the outstanding orders carry. Thus, it becomes identical with the static critical level policy, i.e., $\lim_{n \rightarrow \infty} OH_m(t) = OH(t)$ for all t . This observation implies that *RERF* would perform at least as good as the static policy. For the same system parameters, the optimal long-run average cost of the

static policy should be an upper bound for the optimal long-run average cost of *RERF*.

It should be also noted that, for the same (Q, r, K) values, *RERF* increases class 2 fill rate and decreases class 1 fill rate compared to the static policy. Even though the real on-hand stock level is at or below K , *RERF* satisfies class 2 customers if the modified on-hand level is above K . Thus, the reserve stock allocated for class 1 demands decreases.

In this study, we only consider the case of two customer classes. However, it is straightforward to adapt *RERF* to the case with m customer classes by defining a different flow function for each class with the exception of class 1 (class 1 demands are always satisfied whenever there is stock on-hand). For $j \in \{2, 3, \dots, m\}$, (3) defines the flow function for class j that consider the risk of having backorders from the higher priority customer classes within the remaining lead-time of i^{th} outstanding order

$$f_j(a_i(t)) = e^{-\left(\sum_{k=1}^{j-1} \lambda_k\right)(L-a_i(t))n} \quad (3)$$

This function is employed in computing the modified on-hand inventory level (see (1)) to be used for the rationing decision at the arrival time of j^{th} class customer. There are $m-1$ critical levels for classes other than class 1. The modified inventory level is compared with the critical level for class j to decide whether to ration or not.

5.3 Performance Evaluation of *RERF* via Simulation

Simulation is one of the best available tools for the analysis of complex systems for which analytical techniques are not tractable. It enables us to compute any performance measure of interest for systems under the policy of consideration. For any parameter setting, it is possible to find the optimal policy parameters using simulation-

optimization. Hence, to evaluate the performance of the policy proposed in the previous section, we developed a simulation model and programmed in C. We identified the optimum policy parameters and quantified the gain obtained through *RERF* under different scenarios by comparing the static critical level policy. To distinguish the cases where stock rationing is beneficial and to assess the relative value of *RERF*, we also simulated the common stock policy (FCFS policy), which provides an upper bound on the costs of the rationing policies, and compared its performance with the static policy. Moreover, to characterize the performance of *RERF* relative to all possible rationing policies and to identify the conditions under which dynamic stock rationing is valuable, for some problem instances we simulated the single-class systems N_1 and N_2 (described in Section 5.1) and obtained two lower bounds on the performance of the unknown optimal rationing policy.

For each problem instance, we identified the policy parameters ((Q, r, K, n) for *RERF*, (Q, r, K) for the static policy and (Q, r) for the common stock policy) that provide the minimum long-run average cost estimate. We run the simulation model of each inventory system for 600,000 customer arrivals to ensure the stability of the estimates. To verify our results, we compared the exact class 2 fill rate attained under the static policy, which is provided by Deshpande et al. (2003), with the simulation estimate for each instance of our problem set. The largest deviation of simulation estimate from the exact value is 0.0044, i.e., class 2 fill rate estimate is exact to the second significant digit. Therefore, we concluded that for each problem instance “600,000 customer arrivals” provide a reasonable run-length to observe the steady state conditions of the inventory systems named above.

For given policy parameters, the long-run average cost expression of any policy can be written as

$$C = A \frac{\lambda_1 + \lambda_2}{Q} + h\bar{I} + \pi_1(1 - \beta_1)\lambda_1 + \pi_2(1 - \beta_2)\lambda_2 + \hat{\pi}_1\gamma_1\lambda_1 + \hat{\pi}_2\gamma_2\lambda_2, \quad (4)$$

where A is the fixed ordering cost, h is the holding cost rate, π_i is the unit back-order cost and $\hat{\pi}_i$ is the time-dependent backordering cost. In addition, \bar{I} denotes the average inventory, β_i denotes the fill rate for class i and γ_i denotes the average backorder time per customer for class i ($i = 1, 2$). Let \tilde{C}_{REF} , \tilde{C}_{SP} and \tilde{C}_{CS} denote the long-run average cost estimate of *REF*, the static policy and the common stock policy correspondingly. Then, the performance gain of *REF* as the percent cost reduction obtained by operating the system under *REF* instead of the static rationing policy with the optimal policy parameters can be defined as

$$G_{RCRF} = 100 \frac{\left(\begin{array}{l} \min_{(Q,r,K)} \{ \tilde{C}_{SP} : r + Q > K \geq 0, Q \geq 1, r \geq 0, \text{ all integer} \} \\ - \min_{(Q,r,K,n)} \{ \tilde{C}_{RCRF} : r + Q > K \geq 0, Q \geq 1, r \geq 0, n \geq 1, \text{ all integer} \} \end{array} \right)}{\min_{(Q,r,K)} \{ \tilde{C}_{SP} : r + Q > K \geq 0, Q \geq 1, r \geq 0, \text{ all integer} \}} \quad (5)$$

Similarly, the performance gain of the static rationing policy as the percent cost reduction relative to the common stock policy is

$$G_{SP} = 100 \frac{\left(\begin{array}{l} \min_{(Q,r)} \{ \tilde{C}_{CS} : Q \geq 1, r \geq 0, \text{ all integer} \} \\ - \min_{(Q,r,K)} \{ \tilde{C}_{SP} : r + Q > K \geq 0, Q \geq 1, r \geq 0, \text{ all integer} \} \end{array} \right)}{\min_{(Q,r)} \{ \tilde{C}_{CS} : Q \geq 1, r \geq 0, \text{ all integer} \}} \quad (6)$$

Since the simulation runs are very fast, we find the optimal policy parameters by complete enumeration under the constraints defined in (5) and (6). To obtain G_{RCRF} and G_{SP} values for the settings that reflect the main trade-offs we generated 288 problem instances varying in total demand rate, ratio of class 1 demand rate to the total rate, unit and time dependent backorder costs of both classes and the setup cost. For the lost sales case, we used the same problem instances excluding the time dependent shortage costs. Since we considered different values of the other cost parameters, we fixed the unit holding cost to 5 without loss of generality. We also fixed the lead-time to the unit-time without loss of generality. The problem instances considered in the simulation study were formed from the elements of the following sets:

$$\lambda = (\lambda_1 + \lambda_2) \in \{25, 5\}, \quad p_1 = \frac{\lambda_1}{\lambda} \in \{0.1, 0.5, 0.9\}, \quad \pi_1 \in \{2, 10\}, \quad r_1 = \frac{\pi_1}{\pi_2} \in \{5, 1.25\},$$

$$\hat{\pi}_1 \in \{1, 5\}, \quad \hat{r}_1 = \frac{\hat{\pi}_1}{\hat{\pi}_2} \in \{5, 1.25\}, \quad A \in \{0, 2, 10\}, \quad L \in \{1\}, \quad h \in \{5\}.$$

5.3.1 Backordering Case

Tables 5.1, 5.2 and 5.3 compare the performance of the policies for the backordering case. The data on each row of Table 5.1 (the maximum and the average percent gain) were obtained using 16 different settings varying in $(\pi_1, \pi_2, \hat{\pi}_1, \hat{\pi}_2)$ values. As seen from Table 5.1, for any given total demand rate and the setup cost, the benefit of the static rationing over the common stock policy is maximized when the total demand rate is evenly distributed among the classes, i.e., $\lambda_1 = \lambda_2$. This is an expected result, because as p_1 goes to 0 or 1, the value of the static rationing should diminish. As p_1 decreases, i.e., as the portion of class 2 demand gets higher and higher, the static policy lowers the threshold rationing level towards 0 in order to prevent large numbers of class 2 backorders. When $p_1 = 0.1$, almost for all the cases the optimal rationing

level is 0. Therefore, cost reduction is obtained only through the priority clearing mechanism, which decreases γ_1 and increases γ_2 compared to the first-come-first-served basis clearing. On the other extreme, as p_1 goes to 1, K is increased and most of the stock is reserved for the class 1 customers. When $p_1 = 0$ or $p_1 = 1$, the static rationing policy, *REF* and the common stock policy all become identical, because in these cases there is only one customer class in the system and any kind of class differentiation is irrelevant.

Table 5.1 Comparison of Policies for all possible shortage cost pairs

			G_{REF}		G_{SP}	
			Max	Avg.	Max	Avg.
$\lambda = 25$	$A = 0$	$p_1 = 0.10$	10.10	5.43	4.49	1.26
		0.50	8.95	4.63	29.76	12.01
		0.90	1.35	0.79	12.45	6.39
	$A = 2$	$p_1 = 0.10$	3.67	1.46	2.90	0.58
		0.50	8.73	3.06	13.51	4.71
		0.90	1.13	0.55	6.84	2.99
	$A = 10$	$p_1 = 0.10$	1.49	0.52	3.47	0.83
		0.50	2.84	0.84	7.40	2.65
		0.90	0.46	0.19	3.27	1.29
$\lambda = 5$	$A = 0$	$p_1 = 0.10$	2.21	0.51	6.88	1.18
		0.50	2.61	0.94	12.48	4.42
		0.90	0.64	0.23	5.55	2.34
	$A = 2$	$p_1 = 0.10$	0.51	0.25	1.72	0.28
		0.50	2.75	0.83	8.49	2.38
		0.90	0.88	0.32	3.60	1.41
	$A = 10$	$p_1 = 0.10$	0.21	0.18	1.98	0.53
		0.50	0.87	0.32	4.32	1.56
		0.90	0.17	0.09	2.02	0.89

As expected, *REF* outperforms the static rationing policy at each parameter setting. Parallel to the above discussion about the static rationing, for given λ and A , the benefit of *REF* over the static rationing policy appears to be maximized when

p_1 is far away from the extreme cases and close to 0.5. The only exception is observed when $\lambda = 25$ and $A = 0$. In this case, *RERF* provides the most significant cost reduction, 10.10% in the best case and 5.43% on the average, when the large proportion of the demand is from class 2. This shows that although the static rationing seems to be a valuable policy when the demand rates of the classes are close to each other, rationing the stock dynamically can provide substantial extra benefit for such cases (especially when the class 1 backorder costs are much more higher than the class 2 backorder costs). In addition to that, dynamic rationing is also a valuable tool when λ_2 is high compared to λ_1 . In general, when p_1 is lower, the static policy does not perform well as discussed in the previous paragraph. However, for such cases *RERF* provides considerable additional savings especially when total demand rate is high. This is due to the fact that *RERF* has the capability of increasing β_2 and decreasing γ_2 , i.e., decreasing β_1 and increasing γ_1 , values that attained under the static policy. Naturally, this capability provides significant gain when λ_2 is high compared to λ_1 .

On the other hand, when p_1 is high, *RERF* does not provide noteworthy additional savings. This can be explained with the fact that when the demand is mostly generated by class 1, keeping γ_1 and β_1 values at the levels dictated by static policy is much more beneficial than decreasing them by satisfying some class 2 demands earlier. Hence, assuming continuous replenishment flow and modifying the inventory level accordingly (to satisfy some class 2 demands earlier) is not useful when p_1 is high. In such settings, the information that outstanding orders provide does not have any significance. However, this is true not only for *RERF* but also for all possible dynamic policies as demonstrated by the bounds in Figure 5.2.

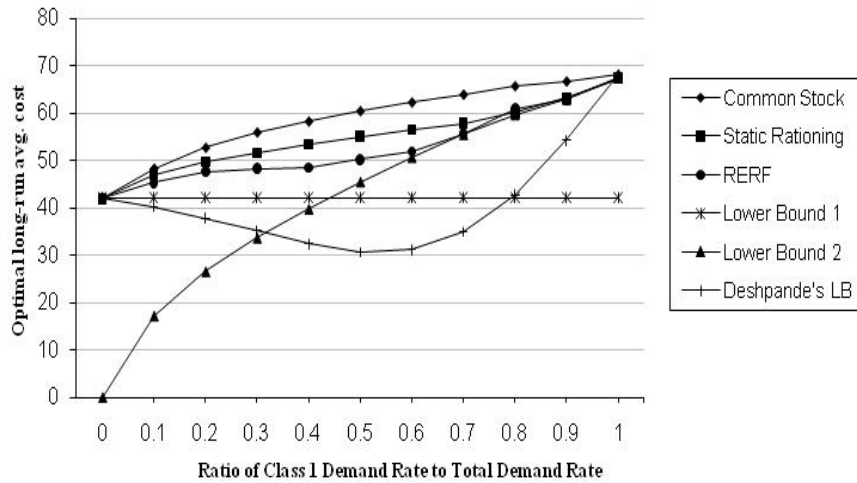


Figure 5.2 Impact of demand mix

Figure 5.2 summarizes all the above discussions related to Table 1 for the case $\lambda = 25, \pi_1 = 10, \pi_2 = 2, \hat{\pi}_1 = 1, \hat{\pi}_2 = 0.2, h = 5, A = 2$. It clarifies the impact of p_1 on the performance of *RERF* and the static policy. In the figure, in addition to the optimal costs of common stock, static and *RERF* policies, the lower bounds obtained from the systems N_1, N_2 are exhibited. Lower Bound 1 and Lower Bound 2 are the optimal long-run average costs of systems N_1 and N_2 , respectively. For each case, the lower bound we propose should be considered as the maximum of Lower Bound 1 and Lower Bound 2. It can be seen from Figure 2 that when $p_1 \geq 0.7$ none of the dynamic rationing policies can provide any extra gain over the static policy, which explains why *RERF* could not provide meaningful savings when $p_1 = 0.9$. For the other cases, i.e., $p_1 < 0.7$, the extra benefit of *RERF* (the benefit of *RERF* over the static policy) is almost same as the benefit of the static policy over the common stock policy. The lower bound proposed by Deshpande et al. (2003) is also provided in Figure 5.2. It is apparent that the lower bound proposed in this paper is tighter than the one proposed by Deshpande et al. (2003).

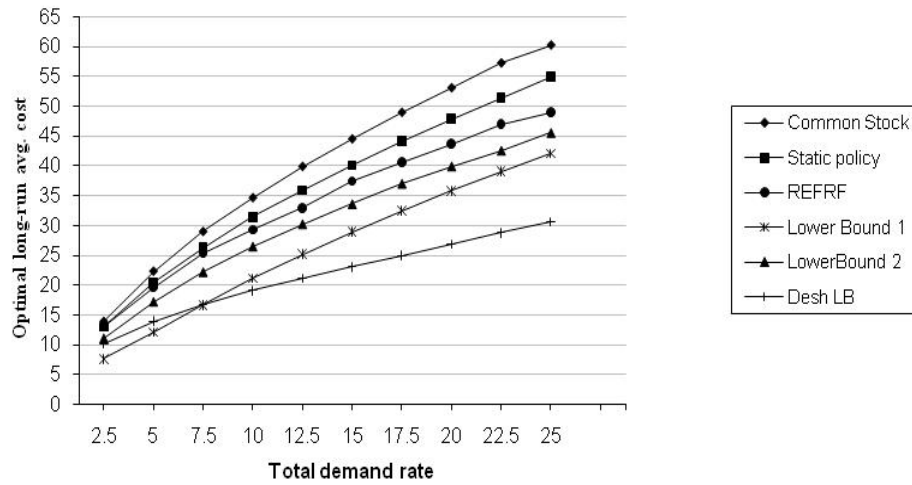


Figure 5.3 Impact of total demand rate ($p_1 = 0.5$)

Table 5.1 also shows that as the total traffic rate increases both the static rationing and *RERF* provides considerable cost savings. Especially when p_1 is 0.1 or 0.5 and the setup cost is low, i.e., A is 0 or 2, and $\lambda = 25$, *RERF* results in remarkable additional savings. Increasing the total traffic rate sharpens the trade-off between holding and shortage costs. Since the value of rationing is based on this trade-off, it increases with the traffic rate. Moreover, we observe that as λ increases, $\lambda L/Q$, which is the expected number of outstanding orders, increases in most cases. Since *RERF* assumes continuous flow of the outstanding replenishment orders, it has more capability to re-optimize the parameters and to increase the cost saving as the number of outstanding orders increases. Figure 5.3 generalizes this discussion by considering ten different total demand rate values for the

case $p_1 = 0.5, \pi_1 = 10, \pi_2 = 2, \hat{\pi}_1 = 1, \hat{\pi}_2 = 0.2, h = 5, A = 2$. As λ increases, both the benefit of static policy over the common stock policy and the benefit of *RERF* over the static policy increase. Moreover, the gap between *RERF* and Lower Bound 2 is not very large, which means that our policy realizes most of the existing potential for the rationing policies. Here it should be also noted that the actual gap between *RERF*

and the unknown optimal policy should be much less than the gap between *RERF* and the lower bounds. If we consider the case where $p_1 = 0.9$, the benefit of unknown optimal rationing policy over the static policy is almost zero for all λ values as shown in Figure 5.4, i.e., the curves corresponding to the static policy and Lower Bound 2 overlap.




Figure 5.4 Impact of total demand rate ($p_1 = 0.9$)

Setup cost is the other important cost parameter that affects the performance of the policies. In an environment with less setup cost the effect of the service level differentiation is more important. As the setup cost increases, we observe that the average ordering cost appears to dominate the average holding and shortage costs. Since the rationing policies derive benefit from the trade-off between the average holding and shortage costs, the relative effect of rationing diminishes as the setup cost increases. Moreover, as the batch size increases (due to the increase in the setup cost) the expected number of outstanding orders decreases and the benefit of *RERF* diminishes.

To summarize the results obtained from Table 5.1, we can say that rationing is a valuable tool when the total demand rate is high and the setup cost is low. In addition, if the demand rates of the customer classes are close to each other, the benefit of the static rationing over the common stock policy increases. In all such environments, *RERF* provides significant additional cost savings. Moreover, when the high portion of the demand is from class 2, the benefit of *RERF* over the static rationing policy is relatively high compared to the benefit of the static rationing over the common stock policy. Tables 5.2 and 5.3 detail the cases where *RERF* and the static policy provide high cost savings. Table 2 compares the policies for $\lambda = 25, A = 0$ and Table 5.3 for $\lambda = 25, A = 2$. For $A = 0$, the benefit of *RERF* over the static policy is maximized with the cost reduction of 10.10% when $p_1 = 0.1, \pi_1 = 10, \pi_2 = 2$ and $\hat{\pi}_1 = 1, \hat{\pi}_2 = 0.2$; and the benefit of the static rationing over the common stock is greatest with the reduction of 29.76% when $p_1 = 0.5, \pi_1 = 2, \pi_2 = 0.4$ and $\hat{\pi}_1 = 1, \hat{\pi}_2 = 0.2$. For $A = 2$, cost savings are maximized when $p_1 = 0.5$. In this case, 8.73% and 13.51% cost reductions are observed when $\pi_1 = 10, \pi_2 = 2, \hat{\pi}_1 = 1, \hat{\pi}_2 = 0.2$ and $\pi_1 = 2, \pi_2 = 0.4, \hat{\pi}_1 = 1, \hat{\pi}_2 = 0.2$ correspondingly for *RERF* and the static rationing policy.

As expected, cost reductions are greatest when $r_1 = \frac{\pi_1}{\pi_2}$ and $\hat{r}_1 = \frac{\hat{\pi}_1}{\hat{\pi}_2}$ are at their maximum value because service differentiation is meaningful when class 1 shortage costs are high compared to the class 2 shortage costs. But, in both cases *RERF* provides the maximum cost reduction over the static policy when $\pi_2 = 2$ although the static policy provides when $\pi_2 = 0.4$. This can be explained by the tendency of *RERF* to decrease the class 2 shortage costs. When the backorder cost of class 2 is high, *RERF* performs better. This fact can also be seen by comparing the cases $(\pi_1 = 10, \hat{\pi}_1 = 5)$ and $(\pi_1 = 2, \hat{\pi}_1 = 5)$ in Tables 5.2 and 5.3.

As discussed in the previous section, *RERF* assumes the exponential flow function $f(a_i(t)) = e^{-\lambda_1(L-a_i(t))n}$ where n is another policy parameter to be optimized with (Q, R, K) values. In our simulation study we consider only integer values of n to obtain the optimal parameters in a reasonable amount of time. Interestingly, for both $\lambda = 25$ and $\lambda = 5$, $n^* \in [3, 11]$ when $p_1 = 0.1$, $n^* \in [1, 4]$ when $p_1 = 0.5$, $n^* \in [1, 3]$ when $p_1 = 0.9$, where n^* stands for the optimum value of the power n . Although n varies in small ranges, we observe that it is not effective to use an appropriate constant n value and optimize the policy only over (Q, R, K) values because changes in n results in considerable cost savings.

Table 5.2 Comparison of Policies for $\lambda=25, A=0$

		$p_1 = 0.10$		$p_1 = 0.50$		$p_1 = 0.90$	
		G_{RERF}	G_{SP}	G_{RERF}	G_{SP}	G_{RERF}	G_{SP}
$\pi_1 = 10,$ $\hat{\pi}_1 = 5$	$r_1 = 5, \hat{r}_1 = 5$	8.81	2.53	7.79	18.60	0.76	8.08
	$r_1 = 5, \hat{r}_1 = 1.25$	7.18	0.36	3.58	14.87	1.35	6.87
	$r_1 = 1.25, \hat{r}_1 = 5$	3.23	0.09	3.96	1.69	1.00	1.97
	$r_1 = 1.25, \hat{r}_1 = 1.25$	3.12	0.00	2.87	2.06	0.79	1.06
$\pi_1 = 10,$ $\hat{\pi}_1 = 1$	$r_1 = 5, \hat{r}_1 = 5$	10.10	4.10	8.95	19.31	0.61	9.23
	$r_1 = 5, \hat{r}_1 = 1.25$	8.65	3.70	8.49	17.86	0.40	8.92
	$r_1 = 1.25, \hat{r}_1 = 5$	2.87	0.00	3.91	1.96	1.00	2.28
	$r_1 = 1.25, \hat{r}_1 = 1.25$	2.47	0.51	3.17	3.02	0.33	2.77
$\pi_1 = 2,$ $\hat{\pi}_1 = 5$	$r_1 = 5, \hat{r}_1 = 5$	5.12	4.49	3.60	25.02	0.63	10.96
	$r_1 = 5, \hat{r}_1 = 1.25$	0.82	0.79	0.42	9.87	1.14	3.90
	$r_1 = 1.25, \hat{r}_1 = 5$	5.80	0.60	5.97	8.83	0.79	6.63
	$r_1 = 1.25, \hat{r}_1 = 1.25$	2.19	0.52	2.40	2.94	0.97	2.37
$\pi_1 = 2,$ $\hat{\pi}_1 = 1$	$r_1 = 5, \hat{r}_1 = 5$	9.86	1.92	3.67	29.76	1.03	12.45
	$r_1 = 5, \hat{r}_1 = 1.25$	4.42	0.00	3.18	21.22	0.32	10.17
	$r_1 = 1.25, \hat{r}_1 = 5$	7.15	0.00	7.11	8.29	0.74	8.00
	$r_1 = 1.25, \hat{r}_1 = 1.25$	5.10	0.57	5.02	6.83	0.85	6.57

Table 5.3 Comparison of Policies for $\lambda=25, A=2$

		$p_l = 0.10$		$p_l = 0.50$		$p_l = 0.90$	
		G_{REF}	G_{SP}	G_{REF}	G_{SP}	G_{REF}	G_{SP}
$\pi_1 = 10,$ $\hat{\pi}_1 = 5$	$r_1 = 5, \hat{r}_1 = 5$	3.40	1.69	7.06	9.01	1.04	4.57
	$r_1 = 5, \hat{r}_1 = 1.25$	3.13	0.46	3.98	7.05	0.95	3.51
	$r_1 = 1.25, \hat{r}_1 = 5$	1.21	0.21	2.35	0.00	0.94	0.53
	$r_1 = 1.25, \hat{r}_1 = 1.25$	1.23	0.00	2.01	0.00	1.08	0.00
$\pi_1 = 10,$ $\hat{\pi}_1 = 1$	$r_1 = 5, \hat{r}_1 = 5$	3.45	2.90	8.73	9.52	0.53	5.84
	$r_1 = 5, \hat{r}_1 = 1.25$	3.67	1.62	7.60	8.49	0.57	5.03
	$r_1 = 1.25, \hat{r}_1 = 5$	0.98	0.00	2.10	0.37	0.56	0.24
	$r_1 = 1.25, \hat{r}_1 = 1.25$	1.38	0.00	1.95	0.09	1.13	0.00
$\pi_1 = 2,$ $\hat{\pi}_1 = 5$	$r_1 = 5, \hat{r}_1 = 5$	1.92	1.33	1.29	12.53	0.20	5.91
	$r_1 = 5, \hat{r}_1 = 1.25$	0.39	0.00	0.96	3.02	0.00	2.20
	$r_1 = 1.25, \hat{r}_1 = 5$	0.38	0.35	3.30	1.95	0.42	2.86
	$r_1 = 1.25, \hat{r}_1 = 1.25$	0.18	0.00	0.80	0.33	0.37	0.77
$\pi_1 = 2,$ $\hat{\pi}_1 = 1$	$r_1 = 5, \hat{r}_1 = 5$	0.90	0.88	0.57	13.51	0.17	6.84
	$r_1 = 5, \hat{r}_1 = 1.25$	0.41	0.00	1.05	8.38	0.04	4.60
	$r_1 = 1.25, \hat{r}_1 = 5$	0.24	0.00	3.49	0.94	0.44	2.84
	$r_1 = 1.25, \hat{r}_1 = 1.25$	0.43	0.00	1.78	0.19	0.27	2.06

As a final test of performance, we compare the results of Teunter and Haneveld (2008) with the results obtained through *RERF*. Under the at-most-one-outstanding-order assumption, Teunter and Haneveld (2008) find the optimal critical remaining lead-time values. In their numerical analysis, they consider two cases. In Example 1, they assume $L = \frac{13}{24}, \lambda_1 = 0.222, \lambda_2 = 1.444, h = 1.4, \tilde{\pi}_1 = 150, \tilde{\pi}_2 = 6.5, A = 0.42$. For this case, they propose that only if the remaining lead-time is greater than 0.248, one item should be reserved for class 1 demand. The average total cost for their suggested policy is 3951. However, for the same setting the minimum average total cost obtained through *RERF* is 3805 (the optimal parameters are $(Q^*, r^*, K^*, n^*) = (2, 1, 1, 15)$). As the second example they consider a setting in which $L = 1, \lambda_1 = 4, \lambda_2 = 10, h = 1, \tilde{\pi}_1 = 100, \tilde{\pi}_2 = 10, A = 0.025$. In this setting, the authors

conclude that it is optimal to reserve at most five items for class 1 demand arrivals. The reserved number decreases with the remaining lead-time. The cost of applying this policy with the optimum reorder point, order quantity and rationing times is 8777. Similar to Example 1, *RERF* provides a lower cost. With the optimal parameters $(Q^*, r^*, K^*, n^*) = (2, 19, 2, 5)$, the cost of *RERF* is 8479.

5.3.2 Lost Sales Case

We performed a similar simulation study also for the lost sales case with the same data set. Since there is no time dependent lost sale cost, for this case we generated 72 cases varying in $\lambda, p_1 = \frac{\lambda_1}{\lambda}, \pi_1, r_1 = \frac{\pi_1}{\pi_2}$. The findings are summarized in Table 5.4, which is the counterpart of Table 5.1 for the lost sales case.

It seems that the dynamic policy does not provide savings comparable to the ones in backordering environment. The benefit of *RERF* over the static policy is greatest with the reduction of 4.30% when $\lambda = 25, A = 0, p_1 = 0.5$. As in the backordering case, the cost reduction obtained by *RERF* increases with λ and decreases with A . Similarly, the highest reductions are observed when p_1 is low or close to 0.5. On the other hand, it is not possible to say similar things to the backordering case for the behavior of the static policy. Table 5.4 illustrates that the static policy performs better when the setup cost increases. Here, the rationing policy is mostly effective as a demand admission control mechanism. Moreover, the highest cost reduction observed when $\lambda = 25$ and $p_1 = 0.1$. Contrary to the backordering case, inventory position does not change with the demands that are not satisfied (lost). In addition, there is no clearing issue in the lost sales case and so all the units of an arriving replenishment order are used to increase the stock. Due to these facts it is not easy to characterize the behavior of the performance of rationing strategy in the lost sales case.

Table 5.4 Lost sales: Comparison of Policies for all possible shortage cost pairs

			G_{REF}		G_{SP}	
			Max	Avg.	Max	Avg.
$\lambda = 25$	$A = 0$	$p_l = 0.10$	3.73	1.67	0.15	0.04
		0.50	4.30	2.05	8.93	3.71
		0.90	0.80	0.43	3.20	1.34
	$A = 2$	$p_l = 0.10$	3.79	1.15	0.46	0.22
		0.50	2.68	0.86	7.04	3.22
		0.90	0.46	0.13	2.32	1.11
	$A = 10$	$p_l = 0.10$	1.45	0.53	8.58	2.65
		0.50	2.13	0.71	6.04	2.53
		0.90	0.52	0.26	1.65	0.92
$\lambda = 5$	$A = 0$	$p_l = 0.10$	0.47	0.31	0.08	0.02
		0.50	1.66	0.51	4.03	1.04
		0.90	0.38	0.16	1.48	0.45
	$A = 2$	$p_l = 0.10$	0.60	0.34	0.10	0.03
		0.50	0.64	0.26	4.85	1.23
		0.90	0.43	0.15	0.83	0.34
	$A = 10$	$p_l = 0.10$	0.25	0.16	21.27	9.96
		0.50	0.17	0.10	6.27	3.58
		0.90	0.38	0.19	1.37	0.62

6. Conclusion

This work constitutes a significant extension of the literature in the area of control of make-to-stock queues, which considers only a single server. We allow an arbitrary number of servers in our model. We model the multi-exponential-server system as an $M/M/s$ make-to-stock queue and show that the optimal production policy is a state-dependent base-stock policy. Furthermore, the optimal rationing policy is of state-dependent threshold type. We also prove that the optimal production and rationing policies are monotone in the inventory level and the number of operational servers. We compare the optimal policy with the previously suggested base-stock policy and demonstrate there are settings where the optimality gap is significant. We also provide three variations on the primary model. The first two variations handle the partial and full order-cancellation flexibility. Our experiments demonstrate that a little flexibility goes a long way and captures most of the value that can be realized via order cancellation. We also discuss a setting with fixed production setup and order cancellation costs as the third variation. In such settings, the optimal policy is no more monotonic in the number of operational channels.

We then consider Erlangian production times as an extension of the $M/M/s$ model. We postulate for this case that the optimal policies are state dependent. The production decision for a specific server depends on the current states (stages) of the other servers. Furthermore, the outstanding production order that is more close to finish (the one that completes more stages) has more value in terms of the rationing decision.

As the number of available servers increases, the optimal policy stops changing beyond a certain number of servers. Therefore, our work also handles the case of infinitely-many servers, i.e., exogenous supply system. Moreover, based on the results of $M/E_k/s$ model we propose a new dynamic rationing policy for the exogenous

supply systems (continuous-review inventory systems) with deterministic leadtimes. For both backordering and lost sales environments, we conduct simulation studies to compare the performance of the dynamic policy with the static critical level and the common stock policies and quantify the gain obtained. We also propose two new bounds on the optimum dynamic rationing policy that enable us to tell how much of the potential gain the proposed dynamic policy realizes.

The multiple server extension provided by this work to the control of make-to-stock queues has potential to open new research avenues. There is a rich and well-established literature in the area of make-to-stock queues. It would be interesting to see how the previous findings in the literature would apply to this more general production setting that considers multiple production channels. This setting should enable us to address important issues such as the effect of pipeline inventory in rationing decisions.

One can consider applying the semi-Markov control approach to analyze the systems with general processing times. However, this would also be problematic in our setting with multiple servers, for the same reason that the analysis of $M/G/s$ queues are. Furthermore, the current literature does not even include the more tractable single server models with general production times, which was suggested as future research by Ha(1997a).

The state-of-the-art in control of make-to-stock queues does not also address the set-up times/cost issue. Extending our model to incorporate this set-up dynamics and costs would be a worthwhile effort. However, in such a setting the monotonicity results of our work no longer hold as discussed in Section 3.2.

References

- Arslan, H., S.C. Graves, T. Roemer. 2007. A single product inventory model for multiple demand classes. *Management Science* **53** 1486-1500.
- Benjaafar, S., M. ElHafsi. 2006. Production and inventory control of a single product assemble-to-order systems with multiple customer classes. *Management Science* **52** 1896-1912.
- Bertsekas, D.P. 2000. Dynamic programming and optimal control (Volumes I and II). Athena Scientific.
- Borkar, V.S. 1984. On minimum cost per unit time control of Markov Chains. *SIAM J. Control Optimization* **22** 965-978.
- Borkar, V.S. 1989. Control of Markov Chains with long-run average cost criterion: The dynamic programming equations. *SIAM J. Control Optimization* **27** 642-657.
- Buzacott, J.A., J.G. Shanthikumar. 1993. Stochastic models of manufacturing systems. Prentice Hall.
- Cavazos-Cadena, R 1989. Weak conditions for the existence of optimal stationary policies in average Markov decision chains with unbounded costs. *Kybernetika* **25** 145-156.
- Cil, E.B., E.L. Ormeci, F. Karaesmen. 2009. Effects of system parameters on the optimal control structure in a class of queueing control problems. *Queueing Systems* DOI 10.1007/s11134-009-9109-x.
- Cohen, M.A., P. Kleindorfer, H.L. Lee. 1989. Service constrained (s, S) inventory systems with priority demand classes and lost sales. *Management Science* **34** 482-499.
- Dekker, R., M.J. Kleijn, P.J. de Rooij. 1998. A spare parts stocking policy based on equipment criticality. *International Journal of Production Economics* **56-57** 69-77.

- Dekker, R., R.M. Hill, M.j. Kleijn, R.H. Teunter. 2002. On the $(S-1, S)$ lost sales inventory model with priority demand classes. *Naval Research Logistics* **49** 593-610.
- Deshpande, V., M.A. Cohen, K. Donohue. 2003. A threshold inventory rationing policy for service-differentiated demand classes. *Management Science* **49** 683-703.
- Erhardt, R. 1984. (s,S) policies for a dynamic inventory model with stochastic lead times. *Operations Research* **32** 121-132.
- Fadiloglu, M.M., O. Bulut 2008. An embedded Markov chain approach to the analysis of inventory systems with backordering under rationing. Working Paper, Department of Industrial Engineering, Bilkent University, Ankara, Turkey.
- Fadiloglu, M.M., O. Bulut 2010. A dynamic rationing policy for continuous-review inventory systems. *European Journal of Operational Research* **202** 675-685.
- Frank, K.C., R.Q. Zhang, I. Duenyas. 2003. Optimal policies for inventory systems with priority demand classes. *Operations Research* **51** 993-1002.
- Gans, N., S. Savin. 2007. Pricing and capacity rationing for rentals with uncertain durations. *Management Science* **53** 390-407.
- Gayon, J.P, S. Benjaafar, F.D. Vericourt. 2009a. Using imperfect advance demand information in production-inventory systems with multiple customer classes. *Manufacturing & Service Operations Management* **11** 128-143.
- Gayon, J.P, F.D. Vericourt, F. Karaesmen. 2009b. Stock rationing in an $M/E_r/1$ multi-class make-to-stock queue with backorders. *IIE Transactions* **41** 1096-1109.
- Ha, A.Y. 1997a. Inventory rationing in a make-to-stock production system with several demand classes and lost sales. *Management Science* **43** 1093-1103.
- Ha, A.Y. 1997b. Stock rationing policy for a make-to-stock production system with two priority classes and backordering. *Naval Research Logistics* **43** 458-472.
- Ha, A.Y. 2000. Stock Rationing in an $M/E_k/1$ make-to-stock queue. *Management Science* **46** 77-87.

- Huang, B., S.M.R Iravani. 2008. A make-to-stock system with multiple customer classes and batch ordering. *Operations Research* **56** 1312-1320.
- Iravani, S.M.R, T. Liu, K.L. Luangkesorn, D.S. Levi. 2007. A produce-to-stock system with advance demand information and secondary customers. *Naval Research Logistics* **54** 331-345.
- Jaarsveld, W., R. Dekker. 2009. Finding the optimal policies in the $(S-1, S)$ lost sales inventory model with multiple demand classes. Technical Report, Econometric institute, Erasmus University Rotterdam, The Netherlands.
- Lippman, S. 1975. Applying a new device in the optimization of exponential queuing systems. *Operations Research* **23** 687-710.
- Mayorga, M., E., H. Ahn, J.G. Shanthikumar. 2006. Optimal control of a make-to-stock system with adjustable service rate. *Probability in the Engineering and Informational Sciences* **20** 609-634.
- Melchior, P., R. Dekker, M.J. Kleijn. 2000. Inventory rationing in an (s, Q) inventory model with lost sales and two demand classes. *Journal of Operational Research Society* **51** 111-122.
- Melchior, P. 2003. Restricted time remembering policies for the inventory rationing problem. *International Journal of Production Economics* **81** 461-468.
- Moon, I., S. Kang. 1998. Rationing policies for some inventory systems. *Journal of Operational Research Society* **49** 509-518.
- Nahmias, S., W. Demmy. 1981. Operating characteristics of an inventory system with rationing. *Management Science* **27** 1236-1245.
- Porteus, E.L. 1982. Conditions for characterizing the structure of optimal strategies in infinite horizon dynamic programs. *Journal of Optimization Theory and Applications* **36** 419-432.
- Sennott, L.I. 1989. Average cost optimal stationary policies in infinite state Markov decision processes with unbounded costs. *Operations Research* **37** 626-633.

- Takacs, L. 1962. Introduction to the theory of queues. Oxford University Press.
- Teunter, R.H., W.K. Klein Haneveld. 2008. Dynamic inventory rationing strategies for inventory systems with two demand classes, Poisson demand and backordering. *European Journal of Operational Research* **190** 156-178.
- Topkis, D.M. 1968. Optimal ordering and rationing policies in a non-stationary dynamic inventory model with n demand classes. *Management Science* **15** 160-176.
- Veinott, A.F. 1965. Optimal policy in a dynamic, single product, non-stationary inventory model with several demand classes. *Operations Research* **13** 761-778.
- Vericourt, F.D, F. Karaesmen, Y.Dallery. 2002. Optimal stock allocation for a capacitated supply system. *Management Science* **48** 1486-1501.
- Weber, R.R, S. Stidham. 1987. Optimal control of service rates in networks of queues. *Advances in Appl. Probability* **19** 202-218.
- Zhao, H, V. Deshpande, J.K. Ryan. 2005. Inventory Sharing and rationing in decentralized dealer networks. *Management Science* **51** 531-547.
- Zipkin, P. 2000. Foundations of inventory management. McGraw-Hill.
- Zipkin, P. 2008a. On the structure of lost-sales inventory models. *Operations Research* **56** 937-944.
- Zipkin, P. 2008a. Old and new methods for lost-sales inventory models. *Operations Research* **56** 1256-1263.