

## Abstract

# Outsourcing Warranty Repairs: Models for the Allocation of Failed Items to Multiple Vendors

*Li Ding*

Doctor of Philosophy  
School of Management  
University of Edinburgh

2007



## Abstract

We consider a scenario in which several external service vendors are contracted to repair purchased items which fail under warranty. We develop and analyze various allocation models concerning how the repair work should be distributed among the vendors in a cost-effective manner. Furthermore, we depart from previous work by arguing the importance of approaches to the modelling of goodwill costs which penalize long waits experienced by individual customers.

We firstly study a simple static allocation model with a fixed warranty population. Both theoretical considerations and numerical results show that a simple greedy approach to the distribution of items under static models works outstandingly well. However, such a static formulation ignores the stochastic nature of the warranty population. Hence, we develop a second allocation model in which new equipment purchases are made according to a compound Poisson process. As in the static allocation model, the current information regarding the repair queue at each vendor is not available to the decision maker. The resulting stochastic dynamic optimization problem is non-standard. We develop an effective allocation procedure to this non-standard problem using a dynamic programming policy improvement approach. We report representative results from a simulation investigation to evaluate the status of the allocation heuristic developed in comparison to two simpler heuristics suggested by static models. Thirdly, we propose a dynamic allocation model which utilizes data on the queue length at each vendor for decisions on the routing of real-time failures to the vendors. Due to the problem size and state space in practice, traditional stochastic dynamic programming is not a realistic and computationally viable option. Hence, Whittle's restless bandits approach is deployed to develop the index-based heuristic for this dynamic allocation problem. A crucial theoretical result in this part of the study is that the system considered is indeed *indexable*. All the numerical results reported show that the performance of the derived index policy from the restless bandit is superior to that of a range of alternatives.

## Acknowledgements

This journey has been accompanied and supported by many people. Without them, I would have wandered lonely in the dark. The first person to whom I am deeply indebted is Professor Kevin Glazebrook. He guided me all the way through the research and the writing of this thesis. I owe him lots of gratitudes for having me work with him. His integrity and enthusiasm for producing high-quality work sets a great example for me to follow.

I want to thank you Dr Jamal Ouenniche, Dr Tom Archibald, Professor Johnathan Crook and Professor Jake Ansell in Management School and Economics at University of Edinburgh, and Dr Christopher Kirkbride at Lancaster University for their help and valuable advice during my PhD period. I also gratefully acknowledge financial supports both from University of Edinburgh and from an Overseas Research Student Award.

I am very grateful for all the supports from my family and friends in the U.K. and back to China. Especially, I would like to pay my special thanks to my husband Pingchuan Ma. Without his love and understanding, I cannot imagine how I get where I am today.

## Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.



Examined by: Yingchen Ding and Wanyou Zhou

# Table of Contents

1	Introduction	1
1.1	Introduction to the chapters	1
1.2	Professor background	4
1.2.1	Warranty issues	7
1.2.2	The research agenda	7
1.2.3	Multiple vendor choice	4
1.3	Research problems	5
1.4	References	7
1.4.1	To my parents: Yingchun Ding and Wenwen Zhang	7
1.4.2	Dynamic programming	8
1.4.3	DP for inventory control	10
1.4.4	Value iteration	12
1.4.5	Multi-Agent Reinforcement Learning	13
1.4.6	Reinforcement Learning	15
1.5	Thesis outline	16
1.6	Contributions	18
2	Three Approaches to Formulating the Goodwill Cost	29
2.1	Introduction to the chapters	29
2.2	A general review of goodwill cost models	31
2.3	The Breakdown and Repair Process at a Single Vendor	32
2.4	Mathematical formulation for the goodwill cost	34
2.5	Properties of the proposed cost models	36

# Table of Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Introduction to the chapter . . . . .	1
1.2	Problem background . . . . .	1
1.2.1	Warranty repairs . . . . .	2
1.2.2	The outsourcing trend . . . . .	3
1.2.3	Multiple vendor choice . . . . .	4
1.3	Research problem . . . . .	5
1.4	Related research and methodologies . . . . .	7
1.4.1	Related research for the problem . . . . .	7
1.4.2	Stochastic dynamic programming . . . . .	8
1.4.3	DP policy improvement . . . . .	10
1.4.4	Value iteration . . . . .	12
1.4.5	Multi-Armed Bandit problems . . . . .	13
1.4.6	Restless bandits . . . . .	15
1.5	Thesis outline . . . . .	16
1.6	Contributions . . . . .	18
<b>2</b>	<b>Three Approaches to Formulating the Goodwill Cost</b>	<b>20</b>
2.1	Introduction to the chapter . . . . .	20
2.2	A succinct review of goodwill cost models . . . . .	21
2.3	The Breakdown and Repair Process at a Single Vendor . . . . .	22
2.4	Mathematical formulations for the goodwill cost . . . . .	24
2.5	Properties of the proposed cost models . . . . .	26

2.6	Conclusions to the chapter . . . . .	27
<b>3</b>	<b>Static Allocations for a Fixed Warranty Population</b>	<b>29</b>
3.1	Introduction to the Chapter . . . . .	29
3.2	The Formulation of a Cost Rate Function at a Single Vendor . . . . .	30
3.3	Discrete Resource Allocation Model . . . . .	33
3.4	Heuristic Solution: Greedy Algorithm . . . . .	34
3.5	Numerical Results and Related Analysis . . . . .	45
3.6	Conclusions to the Chapter . . . . .	55
<b>4</b>	<b>Allocation Models and Heuristics for a Variable Warranty Population</b>	<b>56</b>
4.1	Introduction to the chapter . . . . .	56
4.2	The System . . . . .	58
4.3	DP Policy Improvement Approach by Using Approximate Costs Model	61
4.4	Simulation Study . . . . .	73
4.5	Conclusions to the chapter . . . . .	86
<b>5</b>	<b>Dynamic Allocations upon Real-time Breakdowns</b>	<b>87</b>
5.1	Introduction to the chapter . . . . .	87
5.2	Model Formulation . . . . .	89
5.3	The Restless Bandit Approach . . . . .	92
5.4	Computational Study . . . . .	101
5.4.1	The application of the value-iteration algorithm . . . . .	102
5.4.2	Simulation study . . . . .	106
5.5	Conclusions to the Chapter . . . . .	118
<b>6</b>	<b>Conclusions and Future Research</b>	<b>119</b>
6.1	Answers to the research questions of interest . . . . .	119
6.2	Implications and recommendations of the study . . . . .	123
6.2.1	Implications of academic interest . . . . .	123
6.2.2	Recommendations of managerial relevance . . . . .	124
6.3	Future research . . . . .	125

<b>A Detailed Results from the Value Iteration Algorithm</b>	<b>126</b>
<b>B The published paper</b>	<b>142</b>
<b>Bibliography</b>	<b>143</b>

## List of Figures

2.1 Plot of the guaranteed costs against the response time $\tau$ for Models 1,2 and 3 respectively with $\lambda^1 = 1, \lambda^2 = \lambda^3 = 10, k = 1$ and $r = 0.04$ .	24
3.1 Plot of $\rho^i(t)$ (the guaranteed costs) and $\bar{y}^i(t)$ (the expected costs) for Models 1,2,3 against various population $\lambda$ for cost models 1,2 and 3 respectively.	30
3.4 A flow chart of the resolution for the full system.	34
4.1 Diagram of a queue system with a variable population.	50
4.2 Diagram of a queue system with a fixed population.	51

# List of Figures

2.1	Plots of the goodwill costs against the response time $r$ for Models 1,2 and 3 respectively with $d^1 = 1, d^2 = d^3 = 10, h = 1$ and $\tau = 0.04$ . . . . .	24
3.1	Plots of $g^I(k)$ (the goodwill costs rate) and $f^I(k)$ (the expected costs rate), $I = 1, 2, 3$ against vendor population $k$ for cost models 1,2 and 3 respectively	40
4.1	A flow chart of the simulation for the full system . . . . .	74
5.1	Dynamic routing system with a variable population . . . . .	90
5.2	Dynamic routing system with a closed population . . . . .	91

# List of Tables

3.1	Convexity boundary points for cost model 1 when $c = 0$ , $c = 2$ and $d^1 = 10$ .	47
3.2	Convexity boundary points for cost model 2 when $c = 0$ , $c = 2$ and $d^2 = 1000$	48
3.3	Greedy allocation and associated cost rates for cost model 1 when $c_v = 1$ , $1 \leq v \leq 4$ , and $d^1 = 10$ . . . . .	50
3.4	Greedy allocation and associated cost rates for cost model 1 when repair costs $c_v$ , $1 \leq v \leq 4$ are drawn independently from a $U(0.90, 1.10)$ distribution, and $d^1 = 10$ . . . . .	50
3.5	Greedy allocation and associated cost rates for cost model 2 when $c_v = 1$ , $1 \leq v \leq 4$ , and $d^2 = 1000$ . . . . .	52
3.6	Greedy allocation and associated cost rates for cost model 2 when repair costs $c_v$ , $1 \leq v \leq 4$ are drawn independently from a $U(0.90, 1.10)$ distribution, and $d^2 = 1000$ . . . . .	52
3.7	Greedy allocation and associated cost rates for cost model 2 when $c_v = 1$ , $1 \leq v \leq 4$ , and $d^2 = 10,000$ . . . . .	53
3.8	Greedy allocation and associated cost rates for cost model 1 when $c_v = 1$ , $1 \leq v \leq 4$ , $K = 500$ , and the total service rate is reduced by half. . . . .	53
3.9	the difference between GA and H1, H2 and H3 in terms of overall cost rates for cost model 2 when repair costs $c_v$ , $1 \leq v \leq 4$ are drawn independently from a $U(0.90, 1.10)$ distribution, and $d^2 = 1000$ . . . . .	55
4.1	Means and standard deviations of $N$ , the total number of items under warranty, for examples in subsequent tables. . . . .	76
4.2	Results of a simulation study of the comparative performance of four allocation heuristics when orders are singletons. See above text for further details.	78

4.3	Results of a simulation study of the comparative performance of four allocation heuristics when order sizes are random. See above text for further details.	81
4.4	Equivalent Results of a simulation study for Table 4.3(a) and Table 4.3(b) but under <i>one-item-at-a time</i> rule. See the following text for further details. . . .	82
4.5	Estimated means and standard deviations of inter-allocation times at vendor 1 generated by three heuristics for the scenarios in Table 4.3(a) . . . . .	83
5.1	The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for $K = 100$ and cost model 3. See above text for further details. . . . .	105
5.2	The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for $K = 200$ and cost model 3. See above text for further details. . . . .	105
5.3	Comparative performance of three alternative heuristics under light traffic for cost Model 3. See above text for further details. . . . .	108
5.4	Comparative performance of three alternative heuristics under medium traffic for cost Model 3. See above text for further details. . . . .	109
5.5	Comparative performance of three alternative heuristics under high traffic for cost Model 3. See above text for further details. . . . .	109
5.6	Comparative performance of three alternative heuristics for low inequality of the service rates for cost Model 3. See above text for further details. . . . .	111
5.7	Comparative performance of three alternative heuristics for medium inequality of the service rates for cost Model 3. See above text for further details. . . . .	111
5.8	Comparative performance of three alternative heuristics for high inequality of the service rates for cost Model 3. See above text for further details. . . . .	112
5.9	Problem scaled-up from $K = 500$ under the three options available for cost Model 3. See above text for further details. . . . .	115
5.10	Comparative performance of three alternative heuristics for cost Model 3 and $K = 10,000$ . See above text for further details. . . . .	117

A.1	The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for $K = 100$ , cost model 3 and $TS = 140$ in Table 5.1 See text in Chapter 5 for further details. . . . .	126
A.2	The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for $K = 100$ , cost model 3 and $TS = 170$ in Table 5.1 See text in Chapter 5 for further details. . . . .	127
A.3	The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for $K = 100$ , cost model 3 and $TS = 200$ in Table 5.1 See text in Chapter 5 for further details. . . . .	128
A.4	The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for $K = 100$ , cost model 3 and $TS = 230$ in Table 5.1 See text in Chapter 5 for further details. . . . .	129
A.5	The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for $K = 100$ , cost model 3 and $TS = 260$ in Table 5.1 See text in Chapter 5 for further details. . . . .	130
A.6	The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for $K = 100$ , cost model 3 and $TS = 290$ in Table 5.1 See text in Chapter 5 for further details. . . . .	131
A.7	The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for $K = 100$ , cost model 3 and $TS = 320$ in Table 5.1 See text in Chapter 5 for further details. . . . .	132
A.8	The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for $K = 100$ , cost model 3 and $TS = 350$ in Table 5.1 See text in Chapter 5 for further details. . . . .	133
A.9	The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for $K = 200$ , cost model 3 and $TS = 170$ in Table 5.2 See text in Chapter 5 for further details. . . . .	134
A.10	The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for $K = 200$ , cost model 3 and $TS = 200$ in Table 5.2 See text in Chapter 5 for further details. . . . .	135

A.11 The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 200$ , cost model 3 and  $TS = 230$  in Table 5.2 See text in Chapter 5 for further details. . . . . 136

A.12 The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 200$ , cost model 3 and  $TS = 260$  in Table 5.2 See text in Chapter 5 for further details. . . . . 137

A.13 The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 200$ , cost model 3 and  $TS = 290$  in Table 5.2 See text in Chapter 5 for further details. . . . . 138

A.14 The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 200$ , cost model 3 and  $TS = 320$  in Table 5.2 See text in Chapter 5 for further details. . . . . 139

A.15 The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 200$ , cost model 3 and  $TS = 350$  in Table 5.2 See text in Chapter 5 for further details. . . . . 140

A.16 The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 200$ , cost model 3 and  $TS = 380$  in Table 5.2 See text in Chapter 5 for further details. . . . . 141

## 1.2 Problem background

Optimization has become a widespread phenomenon in almost every region of contemporary business practice since the late 1990s. As a major component of the post-war service to the manufacturing industry, warranty models have experienced a rising prominence in the 2000s. However, the rising role of manufacturing brings not only opportunities but also challenges to manufacturers, especially related to the effective management of the maintenance activities reported in Pallas (2005). Despite some negative views in the media and journals to focus attention towards manufacturing activities, the rapid evolution of scientific and technical knowledge, availability of data and

# Chapter 1

## Introduction

### 1.1 Introduction to the chapter

We begin this thesis by providing an overview of the problem background, within which the research questions are addressed. We then proceed to introduce the broad themes of this research area and associated methodologies, which are of great relevance to the solution methods developed for the problem concerned. A thesis outline and a summary of our contributions to this area follow and conclude this introductory chapter.

### 1.2 Problem background

Outsourcing has become a widespread phenomenon in almost every region of contemporary business practice since the late 1990s. As a major component of the post-sale service in the manufacturing industry, warranty repairs have experienced a rising outsourcing tide too. However, the rising tide of outsourcing brings not only opportunities but also challenges to manufacturers, especially related to the effective management of the outsourced vendors reported in Patton (2005). Despite some negative views from the media and journalists in recent articles towards outsourcing activities, the trend continues relentlessly and with growing awareness of challenges and risks.

Wong (2006) summarizes the top ten challenges in the outsourcing practice. Patton (2005) points out that by contracting with *multiple* service vendors large equipment manufacturers can reduce the risk associated with the sole dependence on a single vendor. In the rest of the section, the problem background will be introduced under three headings: warranty repairs, the outsourcing trend and the multiple-vendors choice respectively.

### 1.2.1 Warranty repairs

Warranty repairs belong to a range of post-sale services which when provided, are of great value in sustaining customer relations. The use of warranties serves many purposes. According to Murthy and Djamaludin (2002), these include protection for manufacturer and buyer, a signaling of product quality, forming an important element of marketing strategy and assuring buyers against items which do not perform as promised. Meanwhile, Byrne (2004) presents some important figures to draw attention to the cost factor in providing warranty repair services. In the United States, the 25 largest manufacturers spent a total of around \$15 billion on warranty claims in 2004. Warranty claims processing is estimated to consume 2.5% ~ 4.5% of revenues for companies across all industries. In light of such impressive numbers, serious considerations of effective approaches to warranty management is crucial. Upon purchase of a product, rational customers choose between brands on the basis of product reliability together with the quality of the associated warranty and maintenance service in the event that product prices and the functional features offered by competing brands are similar. As a result, product warranty plays an increasingly vital role in promoting products and commercial transactions. Hence, besides the primary incentive of cost reduction, any improvement in warranty management can also boost sales and revenues, enhance customer satisfaction and loyalty, and even drive up product quality.

In general, a product warranty is an agreement offered by a manufacturer to a customer to repair a faulty item, or to partially or fully reimburse the customer in the event of a failure during the prescribed length of the warranty period. There are several forms of reimbursement, such as a lump-sum rebate, a free repair warranty and a pro-rata

warranty (see Chukova and Johnston (2006)), among which a free repair of the failed item is the type considered throughout the thesis. Hence, a warranty is actually a liability to the manufacturer in selling a product. For this reason, predicting and reducing warranty costs is of principal interest to the manufacturer. According to Chukova and Johnston (2006), underestimating true warranty costs results in losses for a company. On the other hand, overestimating them will lead to uncompetitive product prices, and as a result product sales could decrease. Regarding the mechanism of warranty coverage, there are two types of warranty policies used in standard business practice and studied in the literature. These are a renewable warranty for which the warranty period is variable and a non-renewable warranty where the warranty period is fixed. To simplify the modelling of the warranty costs of interest, the length of the warranty period in the thesis is assumed to be predetermined and fixed. Further, the recent trend of outsourcing warranty repair services makes the management of warranty repairs even more complicated and challenging.

### 1.2.2 The outsourcing trend

Outsourcing occurs when a vendor contracts with a company to provide services or produce products for a major function or activity. Work that is traditionally done internally is shifted to an external provider, and the employees of the original organization are often transferred to the service provider. Recently, the practice of outsourcing has grown in both scope and sophistication. Major advantages of outsourcing are in helping companies focus on their core competencies and costs reduction. According to Gartner Group, the IT outsourcing market alone totalled over US \$70 billion in 2001 and is projected to grow to \$160 billion by 2005. Moreover, more than 70% of Fortune 500 companies use outsourcing at various levels of their business activities, including software development and warranty repairs.

Many manufacturers are seeking an outsourcing solution for warranty repairs and hoping that the subcontractors offer at least one of the following benefits: staff availability, special expertise or cutting costs. Meantime, subcontractors are the co-beneficiaries in this blooming outsourcing market. Opp et al. (2003) themselves quote a recent re-

port by Merrill Lynch (2001) to the effect that warranty services represent a 100 billion dollar opportunity for both manufacturers and subcontractors. However, we adopt the view of manufacturers for discussions regarding the outsourcing of warranty repairs in this thesis. In the upcoming chapters, we are motivated to investigate effective approaches to managing the distribution of warranty repairs among a collection of outsourced vendors.

### 1.2.3 Multiple vendor choice

Typically a manufacturer outsources warranty repairs to multiple vendors. Such a situation is not uncommon and can arise because of a range of factors: First, the volume of the company's output may be such that there may not be a single vendor (whether with an existing contract or not) capable of handling the entire repair volume. Business Outsourcing Corporation (2006) cited a case study in which a major computer company contracted **two** service vendors to undertake a large volume of emergency repairs to computers under warranty. Second, even should such a highly capable vendor exist, it may be that they would ask for a higher price for repairs. Hence the company may have an incentive to contract with several less capable but cheaper vendors with the consideration of cost reduction as well as the flexibility of shifting work among them. McDougal (2006) stated that "a company wants enough service providers familiar with the company and its business so that it is possible to shift work among them and keep all the vendors competing for new work." Above all, Briskman (2005) states that "multi-vendor situations can lower risk . . . . Certainly there are cases where more than one vendor is selected for a single service". A Deloitte Consulting study (2005) reports the effect that "73% of the participants are working with multiple vendors to reduce vendor dependency". It is certainly the case that in an area as sensitive for customer relations as warranty repairs a manufacturer may not wish to take the risk of being totally reliant on a single service vendor. See, for example, the related comments in Aberdeen Group's 2005 study "Best Practice in Strategic Service Management" and cited by Violino (2006). Another lesson learned from vendor dependency is the British Airways (BA) catering crisis in 2005 for which BA paid a heavy price. This was rooted in the fact that BA outsourced its whole catering operation to a single firm in 1997.

### 1.3 Research problem

This thesis addresses important research questions concerning decisions related to the management (particularly, the workload distribution) of warranty repairs among a collection of service vendors, of the kind likely to confront a large manufacturer when the manager/decision-maker in the organization is seeking cost-effective solutions to the problem considered. Maintaining or even improving the goodwill from valued customers by avoiding long delays in receipt of their warranty repair services is a primary concern which should influence the manufacturer's decision-making. This thesis contains comprehensive analyses which aim to shed light on a range of important questions relating to the decision-making process for the manufacturer. These include:

1. What level of service capacity among the contracted vendors needs to be available to meet the anticipated demand for the manufacturer's post-sales repair service effectively?
2. Given that the manufacturer's currently contracted vendors do possess sufficient service capacity, how should the repair work be best distributed among them?
3. How much might the manufacturer have saved by employing a central decision controller to determine to which vendor a warranty repair is sent every time an item fails rather than preassigning them to the service vendors when they are purchased?
4. How much might the manufacturer be losing (economically and in customer goodwill) by maintaining an existing suboptimal approach to workload distribution?

Attempts to uncover the truth underlying such questions led to the formulation of the outsourcing warranty problems with an objective of the minimization of the average cost rate (or come close to doing so). The following assumptions are stated to define boundaries for the research problem formulated in this thesis:

1. This thesis only studies one single type of repairs. Repairs of interest under warranty have identical repair cost structures. We also assume that the service vendors and customers concerned are not geographically dispersed and the logistics costs for transporting items between vendors and customers are identical for all items of interest;
2. A non-renewable warranty with fixed warranty period is assumed;
3. The manufacturer is responsible for all the costs incurred associated with the repairs under warranty. The warranty coverage to customers is called a free repair warranty;
4. Multiple vendors are deemed to be available in provision of warranty repair services;
5. Both *fixed* and *variable* warranty populations are considered. The assumption of a fixed warranty population in which the number of items under warranty is a constant helps to simplify the model formulation for the allocation problems studied. This model is a useful approximation to circumstances where the warranty population is subject to light to moderate variation. However, minimizing the warranty repair cost arising from a variable warranty population is more often of primary interest to the manufacturer. Specifically, the number of items under warranty is dynamic with an item entering the system when it is newly purchased and leaving the system when its warranty expires;
6. Based on decision time and system state information available, we have two allocation approaches, namely, allocating upon purchase and allocating upon real-time breakdown. Under the former allocation approach, the vendor assigned to an item upon failure within the warranty period is decided at the time of its purchase and in the absence of any system state information regarding repair queues at the vendors. The latter approach delays allocation decisions until the time of an item's breakdown when information concerning the repair queue length and the assigned warranty population at each vendor is available to be utilized in support of decision-making;

7. We assume that the type of repair service provided is of single-priority. Items under warranty awaiting repairs are treated impartially and served on a first-come-first-served (FCFS) base.

On the basis of the delimitations of the research problem described above, we study the outsourcing of warranty repairs within three modelling frameworks. We start with the simple static allocation model in which it is supposed that there is a fixed number of items under warranty all the time. This is formulated as a static optimization problem in Chapter 3 and the fixed warranty population is to be divided between a collection of vendors. Chapter 4 gives rise to the non-standard stochastic optimization problem in which the warranty population considered is generalised to be constantly changing. However, the information of repair queue lengths is not available for decision-making at that stage. Finally, the dynamic allocation model in which continuous information of all repair queues becomes observable to the decision maker is formulated as a Markov decision problem in Chapter 5. In the following section, we summarize the pertinent research and solution approaches within the above modelling frameworks to tackle this research problem.

## **1.4 Related research and methodologies**

### **1.4.1 Related research for the problem**

The research concerning the optimal allocation for the outsourcing warranty repairs was originally stimulated by contacts between colleagues in the Department of Statistics and Operations Research at the University of North Carolina at Chapel Hill and a large equipment manufacturer with which several external service vendors were contracted to undertake warranty repairs. As a result, Opp et al. (2003) were the first to address this question and to utilize resource allocation and queueing models to formulate the problem under the assumptions of allocating-upon-purchase with a fixed warranty population. Note that they took a traditional linear-holding-costs approach

in modelling goodwill costs. Such an approach results in a separable and convex objective function (under simply stated conditions), and the application of a greedy algorithm solves the problem (see the proof in Gross (1956)). We argue that in the current context it is important to take explicit account of delays experienced by individual customers when considering the formulation of goodwill costs. More detailed discussion concerning our proposed goodwill cost models are contained in Chapter 2.

In a later paper by Opp, Glazebrook and Kulkarni (2005), an allocating-upon-breakdown model for the outsourcing of warranty repairs is studied. Other assumptions remain as in Opp et al. (2003). Here the stochastic optimization problem of interest is formulated as a Markov decision process (MDP). The authors realize that solving a problem of realistic size is not numerically tractable by using a conventional stochastic dynamic programming approach. Thus they develop two index-based heuristics deploying the ideas of DP policy improvement and Whittle's restless bandits respectively, as near-optimal solutions for this dynamic allocation problem.

Several researchers in this area have been interested in extending the allocation problems studied by Opp et al. (2003) and Opp, Glazebrook and Kulkarni (2005) by allowing multi-prioritized warranty repairs. Buczkowski, Kulkarni and Hartmann (2005) develop an efficient algorithm for the multiple-priority and allocating-upon-purchase problem under a fixed warranty population. Chen and Kulkarni (forthcoming) on the other hand consider a more complicated version of the problem with a model incorporating multi-prioritized customer classes, allocating-upon-breakdown and a variable-warranty-population. This problem is tackled effectively by attaching calibrating indices to each vendor. This follows similar approaches proposed in Opp, Glazebrook and Kulkarni (2005).

### 1.4.2 Stochastic dynamic programming

We consider a class of Markov decision problems (MDPs), which are described as follows. At each decision epoch, a decision maker observes the current system state  $i \in \mathcal{S}$  and chooses an action, say  $a \in A(i)$  from a set of available actions. As a result of the

action choice  $a$ , an instantaneous cost, say  $c_i(a)$  is incurred (or an immediate reward is earned) and a (possibly) new system state is determined according to the transition probabilities  $p_{ij}(a)$ ,  $j \in S$ . Both costs (or rewards) and transition probabilities depend on the current state of the system and the action selected by the decision maker in that state. We seek an optimal policy, namely a rule for choosing actions which will minimize the total expected cost incurred (or maximize the total expected reward earned) over some horizon. Such MDP models have seen much development due to their relevance to a broad range of applications, including, for example, queueing control and inventory management. Dynamic programming (DP) which is primarily based on recursive backward induction was popularized by Bellman (1957) and Howard (1960) in tackling MDPs over a finite horizon. Such a conventional approach solves the stochastic dynamic optimization problem by computing optimal values from a DP recursive equation in a finite number of stages. This has the following form:

$$V_t(i) = \min_{a \in A(i)} [c_i(a) + \sum_{j \in S} p_{ij}(a)V_{t-1}(j)], \quad i \in S \quad (1.1)$$

where  $p_{ij}(a) \geq 0$  and

$$\sum_{j \in S} p_{ij}(a) = 1, \quad a \in A(i), \quad i \in S$$

In (1.1),  $V_t(i)$  denotes the optimal expected total cost incurred over  $t$  time periods when the initial system state is  $i$ . The optimal choice of actions is the minimizing choices in (1.1). In comparison to exhaustive enumeration, dynamic programming lowers the computational burden considerably. The theory of DP has evolved a great deal in the last five decades, and a large literature now exists on this topic. In fact, it continues to be the main pillar for solving stochastic dynamic optimization problems. Suppose now that a policy  $u$  from the stationary class  $U$  (of policies which choose actions on the basis of the current state only) is applied to the system. Equation (1.1) is now modified to

$$V_t^u(i) = [c_i(a) + \sum_{j \in S} p_{ij}(a)V_{t-1}^u(j)], \quad i \in S \quad (1.2)$$

where  $p_{ij}(a) \geq 0$  and

$$\sum_{j \in S} p_{ij}(a) = 1, \quad a \equiv u_i, \quad i \in S.$$

We use  $u^* \in U$  for an optimal policy which satisfies

$$V_t^{u^*}(i) = \sup_{u \in U} V_t^u(i), \quad i \in S. \quad (1.3)$$

In (1.2),  $V_t^u(i)$  denotes the expected cost over  $t$  time periods when the initial system state is  $i$  and policy  $u$  is used. Action  $a$  is the one determined by policy  $u$  when the system is in state  $i$ .  $V_t^{u^*}(i)$  in (1.3) is the optimal expected total cost over a horizon of length  $t$  obtained by executing the optimal policy  $u^*$  when the initial system state is  $i$ . Some researchers attempt to answer questions such as when does an optimal policy exist, whether it has a particular structure, and how to determine and compute an optimal policy efficiently for a general model. See, for example, Puterman (1994).

In this thesis we confine our discussions regarding MDPs to the optimality criterion of average expected cost rate with finite state and action space over an infinite planning horizon. Suppose  $g_i(u)$  to be the average expected cost per unit time under stationary policy  $u$  when the initial system state is  $i$ . Technically,  $g_i(u)$  is the limit of  $V_t^u(i)$  divided by the time length  $t$  when  $t \rightarrow \infty$  and the existence of the limit for stationary policies is ensured according to Tijms (1994) under given conditions. Further, the subscript of initial state  $i$  in  $g_i(u)$  could be dropped when the problem considered is unichain. We can then replace  $g_i(u)$  unambiguously by  $g(u)$ .

As the optimizing problem is over an infinite horizon, we are unlikely to proceed with a traditional (backwards induction) dynamic programming approach to solve the problem of interest which features computational intractability. Hence, in the following two subsections we will discuss two widely used DP approaches to tackle such MDPs, DP policy improvement and value iteration respectively.

### 1.4.3 DP policy improvement

The origins of the policy improvement approach may be traced back to the 1960s. At that time Howard (1960) pioneered a policy-iteration algorithm for solving probabilistic sequential decision processes over an infinite planning horizon by using principles from Markov chain theory and dynamic programming. A theoretic foundation

of Howard's policy-iteration approach was given in Blackwell (1962). The basic idea underlying policy iteration is to start with an arbitrary stationary policy, under which the average cost such as  $g(u)$  and its relative values<sup>1</sup> are computed, and then monotonically move to a better policy in every iteration. This process will continue until no further improvement is possible and/or it achieves optimality. The relative values play a key role for constructing a new (better) policy in each step. See Tijms (1994) for a detailed discussion. Following this pioneering work, the policy-iteration method has been applied to a range of Markov decision problems, especially in the domain of queueing control. See Glazebrook et al. (2004), Ansell et al. (2001), Tijms (1994), Krishnan (1990) and Krishnan and Ott (1986,1987).

Under given conditions, after performing a finite number of policy improvement steps, one should obtain an optimal policy upon convergence of the policy-iteration algorithm. However, given the complexity of general problems (e.g. multi-dimensional MDPs), executing several policy improvement iterations is numerically and computationally infeasible. Norman (1972) was the first to propose the heuristic approach of applying a single policy-improvement step to such MDPs. According to empirical experience (see, for example, Krishnan and Ott (1986,1987) and Wijngaard (1979)), the most improvement is usually secured at the first step. We can then follow Krishnan and Ott (1987) in developing a "one policy-iteration step" approach by applying a single DP policy improvement step to an initial optimal static (i.e. state independent) policy with the expectation of obtaining a near-optimal policy.

We consider this DP heuristic approach to the outsourcing warranty repairs problem considered in Chapter 4. Our approach is to design an optimal static policy at the first stage. Then at the second stage a single DP recursion is applied to this initial static policy. The resulting policy is index-based. Items under warranty are allocated to the vendor with the smallest associated state-based index.

---

<sup>1</sup>the bias values associated with a value function, which indicate transient effects when starting in different states

### 1.4.4 Value iteration

Much development of the value-iteration method for Markov decision models was also carried out in the 1960s. Blackwell (1965) stated a set of sufficient conditions, namely, discounting and monotonicity to ensure the convergence of the sequence of value functions obtained under value iteration based on the famous Banach fixed point theorem. Value-iteration bounds for the discounted optimal costs were proposed in MacQueen (1966). On the other hand, White (1963) gave the first proof of the geometric convergence of the undiscounted value-iteration algorithm for MDPs under a recurrence condition. Odoni (1969) and Hastings (1971) introduced lower and upper bounds on value functions. They also extended MacQueen's original work for the discounted models to the undiscounted cases. Due to its conceptual simplicity, and the ease with which it may be coded and implemented, value iteration has become one of the most widely used and best understood algorithms for solving Markov decision problems, although it leads to a numerical rather than analytical solution. Its popularity also extends to other scientific disciplines, such as reinforcement learning in artificial intelligence.

The rest of this subsection describes basic ideas behind one form of the value-iteration algorithm for the undiscounted MDPs with average cost optimality criteria following the book of Tijms (1994).

We use equation (1.1) for illustration. One computes and updates the value functions in (1.1) recursively at each iterative step. Specifically, one starts with the value function  $V_0(i), i \in S$ , which is usually assumed to be identically zero.  $V_t(i)$  is then computed recursively using  $V_{t-1}(i)$  on the r.h.s. of (1.1) for  $t \geq 1$ . Given the functions  $V_t$  and  $V_{t-1}$ , upper and lower bounds (denoted as  $M_t$  and  $m_t$  respectively) for the optimal average cost rate are given by

$$M_t = \max_{i \in S} \{V_t(i) - V_{t-1}(i)\}, \quad t = 1, 2, \dots$$

$$m_t = \min_{i \in S} \{V_t(i) - V_{t-1}(i)\}, \quad t = 1, 2, \dots$$

Generally, the value-iteration algorithm will stop when the difference between  $M_t$  and  $m_t$  is less than some given tolerance value. Tijms (1994) stated that under the *Weak*

*Unichain Assumption*<sup>2</sup>, the following inequality exists

$$m_t \leq g^* \leq M_t, \quad t = 1, 2, \dots$$

where  $g^*$  denotes the optimal average cost per unit time over an infinite horizon. Further, the sequence  $\{m_t, t \geq 1\}$  is nondecreasing and the sequence  $\{M_t, t \geq 1\}$  is nonincreasing. Hence, the optimal average cost rate can then be accurately estimated.

In Chapter 5, we adapt the above value-iteration algorithm to our problem context to compute the optimal average cost rate for our dynamic allocation model. Moreover, a suitable version of this algorithm can also yield the cost rate associated with a specified policy.

However, the two traditional approaches of policy and value iteration discussed above are computationally infeasible in instances of our problems where the state space is of realistic size. Hence, there is a strong demand for the development of efficient approaches and techniques for generating optimal or near-optimal solutions to complicated stochastic dynamic optimization problems of special structure.

### 1.4.5 Multi-Armed Bandit problems

The multi-armed bandit problem derives its name from a gambling decision problem, in which a gambler has to choose a sequence of plays on  $N$  slot machines (bandits) to maximize the total expected gain when the winning probability of the  $i^{\text{th}}$  machine ( $i = 1, 2, \dots, N$ ) is unknown. Playing on machine  $i$  will enable the gambler to obtain information on the associated winning probability. In making choices between machines there may be a trade-off between gaining information which can be used later (with the prospect of better future rewards) and exploiting the information already available (to secure high current returns). With such a dilemma along with the 'curse of dimensionality', the multi-armed bandit problems had been found to be prohibitively difficult for many years.

---

<sup>2</sup>For each average cost optimal stationary policy the associated Markov Chain has no two disjoint closed sets.

We describe a bandit process by its state and the reward yielded upon activation. In general, the state evolution of a bandit when operated is not necessarily Markovian. The multi-armed bandit problem consists of  $N(N > 1)$  independent bandit processes (also referred to as machines or arms or projects) and one decision controller. At each decision epoch, the controller chooses only one bandit to play. Note that the states of all other unselected bandits remain unchanged. The multi-armed bandit problem aims to find an optimal policy to maximize some measure of total reward. Following Mahajan and Teneketzis (2007), there are four key features which characterize the classical multi-armed bandit problem:

1. Exactly one machine is operated at each decision stage;
2. Machines that are not being activated remain frozen;
3. Machines evolve independently of each other;
4. Passive machines yield no reward.

The groundbreaking paper of Gittins and Jones (1974) successfully solved a discounted version of the multi-armed bandit problem by decomposing it into  $N$  single-armed bandit problems. They derived a *dynamic allocation index* (later referred to as the Gittins' index) and so reduced the dimensionality of the problem. A Gittins' index is attached to each machine and is a function of its current state. The policy which at all epochs selects a machine for which the associated Gittins' index is maximal is optimal. Furthermore, an optimal policy can be obtained by a process of forwards induction. This is ensured by the four features of general multi-armed bandit problems described above. It is known that a forwards induction procedure is computationally more efficient than a backwards induction one. See Mahajan and Teneketzis (2007). Applications of the multi-armed bandit model to gambling, stochastic scheduling, sequential clinical trials and optimal search are discussed and formulated by Gittins (1979). Inspired by this breakthrough, a large quantity of research work followed: Whittle (1980) used a dynamic programming approach to verify the optimality of Gittins' index policies; Whittle (1981) also produced an optimality proof of Gittins' index result for an *open* version of the problem with arrivals. A variety of simpler proofs of the optimality of

Gittins' index policy have emerged subsequently; See Weber(1992), Tsitsiklis (1994), Garbe and Glazebrook (1996) and Frostig and Weiss (1999). See also Whittle (1982), Glazebrook (1983), Ross (1983), Berry and Fristedt (1985), Katehakis and Veinott (1985), Lai and Robbins (1985), Varaiya et al. (1985), Ishikida and Varaiya (1994) and Kaspi and Mandelbaum (1998) for a wide range of discussions regarding this area of bandits decision problems.

#### 1.4.6 Restless bandits

Restless bandits were introduced by Whittle (1988) and yield a generalisation of multi-armed bandit problems to allow for machines to evolve even when passive. They constitute an intractable class of decision processes, which have been showed to be PSPACE-hard by Papadimitriou and Tsitsiklis (1999). In formulating such problems, Whittle was concerned with the development of index policies by considering a Lagrangian relaxation of the original problem. The indices derived generalize Gittins' indices to the restless case.

In Whittle (1988), the assumptions for the classic multi-armed bandit model are relaxed such that bandits will evolve whether chosen for play or not though according to different stochastic dynamics. Instead of operating a single machine at one time, the model now is generalized to allow  $m > 1$  machines to operate at every decision stage. With such modifications, one cannot apply the original Gittins' solution. Alternatively, Whittle uses a Lagrangian approach with multiplier  $\nu$  to relax the problem. From an economic point of view,  $\nu$  could be interpreted as a 'subsidy for passivity', whose value should be set to ensure that  $m$  bandits are active on average in an optimal policy for the Lagrangian relaxation. It turns out that a natural index for bandit  $i$  in state  $x_i$ ,  $v_i(x_i)$  is the value of subsidy  $\nu$  which makes both active and passive actions optimal for bandit  $i$  in state  $x_i$ . He clarifies 'indexability' with the following statement: for any bandit to be indexable, if it is optimal for the bandit not to be selected under a  $\nu$ -subsidy policy, then this remains the case under a  $\nu'$ -subsidy policy where  $\nu' > \nu$ . Note that by a  $\nu'$ -subsidy policy, Whittle means an optimal policy for a single bandit when  $\nu'$  is the subsidy for passivity. However, there are no simple sufficient conditions for indexability given in

his paper. Whittle also warns that the index policy he develops for the restless bandit problem is unlikely to be optimal, but may well be very close to being so.

There is a developing research literature centered around the theoretical study of restless bandits models. Niño-Mora (2001) explored a polyhedral approach to the elucidation of simple sufficient conditions for Whittle's indexability by utilizing partial conservation laws (PCL). Niño-Mora (2002) extended Whittle's ideas and indicated conditions under which the derived indices for restless bandits can be successfully applied to dynamic routing problems. See also Weber and Weiss (1990), Whittle (1996), Glazebrook, Niño-Mora, and Ansell (2002), and Ansell et al (2003). Further, it has been shown that the restless bandit formulation is relevant for a wide range of applications. The performance of Whittle's index policy has appeared highly promising especially in the area of queueing control. See, for example, Glazebrook, Mitchell and Ansell(2005) and Glazebrook, Kirkbride, and Ouenniche (2005). The approach to developing index policies for restless bandits proposed by Whittle is deployed to obtain an index heuristic for a dynamic allocation model arising in the outsourcing warranty repairs problem in Chapter 5.

## 1.5 Thesis outline

This thesis contains an extensive investigation concerning the effective allocation of warranty repairs among a collection of outsourced vendors, which is organized in the following way. Chapter 1 presents an overview of the problem background and indicates the motivation and significance of the research questions addressed. Aiming to explore solution methods for the problem in a cost-effective manner, literature which is of relevance to this thesis have been reviewed. Stochastic dynamic optimization is discussed as a broad theme which embraces the approaches utilized in this thesis. Specifically, these include DP policy improvement, value iteration and restless bandit approaches.

Chapter 2 formulates the breakdown/repair process at a single vendor as a finite population queueing system and proposes three approaches to the modelling of the goodwill costs encountered in our problem, all of which take explicit account of long delays experienced by individual customers.

Chapter 3 addresses the warranty repair problem under the assumption of allocating-upon-purchase for a fixed warranty population. The resulting static allocation problem is formulated as a resource allocation problem in which the objective function is additive and separable. We elucidate the appropriate form of the cost rate function associated with an individual vendor using the cost models and stochastic breakdown/repair process formulated in Chapter 2. Though no claim to the optimality of simple heuristics can be made in any generality due to the special cost structure, a greedy heuristic is proposed and evaluated both theoretically and numerically with impressively strong performance.

Chapter 4 studies a non-standard allocation problem for a variable warranty population where the associated system state is partially observed. We develop a two-step allocation procedure which deploys the approach of DP policy improvement so as to minimize the overall cost rate or come close to doing so. The resulting allocation procedure makes thorough use of the partial system information from previous allocations at each vendor which considers both the number of items allocated and their durations of unexpired warranties. We then derive a vendor-specific calibrating index which is a function of system current state information. The incoming order (whether a singleton or in bulk) is allocated to the vendor with the smallest index. An extensive simulation study is conducted to attest the status of this sophisticated allocation policy in comparison to three other simpler heuristics.

Chapter 5 considers the dynamic allocation problem upon real-time breakdowns, which is formulated as a Markov decision process. However, the intractability of this dynamic routing problem in a multi-vendor context makes the direct application of traditional stochastic dynamic programming infeasible for problems of realistic size. Following

Whittle's approach, we develop an index-based heuristic using a restless bandit approach. This results in a policy which uses information regarding the length of the repair queue at each vendor. Both the value-iteration algorithm and a Monte Carlo simulation study are used to verify the superior status of the restless bandits heuristic developed in comparison with other simpler heuristics.

Chapter 6 concludes the findings of the thesis and suggests some possible directions of future research.

## 1.6 Contributions

First, we investigate the allocation problem regarding the outsourcing of warranty repairs within a systematic framework. Both static allocation and dynamic allocation models for fixed or dynamic warranty populations are comprehensively studied. We develop efficient and near-optimal heuristics, namely the greedy algorithm, the DP policy improvement heuristic, and an index policy from the restless bandit approach, by exploiting system information available at different levels. When the heuristics developed do not have any simple explicit form for the models we consider, theoretical discussions have been carried out to show that they nevertheless have sensible properties, such as monotonicity and convexity.

Second, we argue the importance of approaches to the modelling of goodwill costs which penalize long waits experienced by individual customers. If these are ignored, the manufacturer might put valued customers' loyalty at risk. Unlike the conventional linear holding cost model used in many application areas of queueing control, we appreciate the importance of taking explicit account of a service time threshold into the formulation of the goodwill cost. Such formulations complicate the resulting optimization problem and make the development of near-optimal heuristics even more challenging. No such modelling approaches for the goodwill cost have been taken in this area to date to our knowledge.

Third, we undertake extensive computational studies to assess the status of the heuristics we develop for the outsourcing problems. We present a wide range of comparisons between alternative heuristics proposed. Overall, the numerical results reported show that the heuristics we develop for both static and dynamic models throughout the thesis work outstandingly well in comparison to the empirical equivalences. Though these computational investigations are model-based, we attempt to construct our numerical study in a realistic manner with sensibly chosen parameters. Hence, the associated efforts have been made: consistent results are obtained under either numerical or simulation-based computational studies for problems of small size in the dynamic allocation model. Hence, we are confident of the validity of our results under the simulation study alone for problems of realistic size when a numerical study is not a computationally viable option. We conduct sensitivity analyses to investigate the influence of key parameters on the results obtained for both static and dynamic models to provide business insights and managerial relevance for large manufacturers.

Note that one paper regarding the simple static allocation model has been published in the *Journal of the Operational Research Society*, see Ding and Glazebrook (2005). Further two based on the work in Chapters 4 and 5 are submitted to the highly regarded international journals, *Management Science* and *Probability in the Engineering and Informational Sciences* respectively. The *Management Science* paper is in press while the latter one is in the process of revision.

## Chapter 2

# Three Approaches to Formulating the Goodwill Cost

### 2.1 Introduction to the chapter

In principle, this thesis will develop index-based heuristics under either static or dynamic allocation models for the best utilization of the outsourced vendors in undertaking the workload of warranty repairs. In other words, we are concerned to distribute the workload in a cost-effective manner in which the overall cost rate incurred at a collection of vendors is (close to) minimal. Therefore, fundamental to our study in the upcoming chapters is the formulation of an objective function which consists of both goodwill costs (in some literatures, these are referred to as waiting costs) and physical repair costs. We argue that the former's contribution to the total cost rate should predominate. In comparison to goodwill costs, the modelling of physical repair costs, namely, labor, parts and so on is simpler and less contentious. As we will show, this chapter primarily describes three different approaches to formulating goodwill cost models, which penalise the delays experienced by individual customers when broken items under warranty are awaiting repairs. The difference between the three approaches lies in the ways of calculating the penalty actually incurred due to long waits which reflect different perspectives towards the role played by a given service time threshold.

The rest of the chapter is organized as follows. Section 2.2 presents a succinct summary of approaches to the modelling of the goodwill cost in the relevant literature. Section 2.3 introduces a stochastic process for breakdowns and repairs at a single vendor. This is a premise for the mathematical approaches to the modelling of goodwill costs which follow. Accordingly, three different goodwill cost models are formulated in Section 2.4 to take explicit account of individual long waits. We then explore some important properties of the proposed cost models in Section 2.5. Section 2.6 contains some concluding remarks.

## 2.2 A succinct review of goodwill cost models

Linear holding costs for waiting are conventionally used in such areas as inventory management and queueing control. See, for example, Naor (1969), Stidham (1970), Stidham (1985), Sennott (1999) and Opp et al. (2003). Opp et al. (2003) take such an approach and assume that a breakdown awaiting repair incurs a fixed holding cost per unit time. This linear cost formulation under the static allocation model (to be presented in Chapter 3) has the consequence that the resulting optimization problem for the distribution of warranty repairs then becomes a separable convex resource allocation problem (under simply stated conditions), for which greedy solutions are known to be optimal from Gross (1956).

However, such an approach to the modelling of goodwill costs has been brought into question in application settings concerned with after-sale repair services. Hopp and Sturgis (2000) point out that the linear cost approach does not take adequate account of the features of service level guarantees. Taylor (1994) shows that long waits affect the overall service evaluations obtained from customers. See also the discussions in van Meigham (1995) and Ansell et al. (2003) regarding the unrealistic use of linear costs in the queueing control field. In the current context, it is realistic to assume that the manufacturer has a code of service level guarantees related to response times for

warranty repairs, of which the customers are well informed. There is also a substantial literature that studies customer psychology in waiting situations. Naor (1969) conducts an investigation of how customers' decisions are influenced by delays in queue. A discussion of the impact of uncertainty can be found in Maister (1984). Kumar et al. (1997) examine the impact of waiting time guarantees on customers' waiting experiences. Hence, growth of the goodwill cost beyond a guaranteed threshold is reflected in the cost modelling approaches proposed in the following sections. Meanwhile, from the point of view of the manufacturer, the allocation of warranty work to a collection of service vendors should be made in such a way that delays greater than any given service time threshold become improbable. Encouraged by the common interest of both customers and the manufacturer, we propose three approaches to the modelling of goodwill costs in this outsourcing warranty repair problem.

## 2.3 The Breakdown and Repair Process at a Single Vendor

Before we proceed to the mathematic formulations of the three cost models, it is important to introduce the stochastic process for modelling breakdowns and repairs at a single service vendor assumed to be responsible for a fixed number of warranty items:

For fixed population size  $k$ , we develop a finite population queueing model for the breakdown and repair process at a single vendor. Gross and Harris (1998) confirm machine repair as an archetypical application of finite population queueing models. We make the following assumptions:

1. A single server queue is assumed in the study, which in our context is a good approximation to more general models with  $s$  servers working in parallel;
2. We make Markovian assumptions for the inter-arrival time and repair service time of any broken item. To be specific, individual item up times follow an

exponential distribution with rate  $\lambda$  while repair times follow an exponential distribution with rate  $\mu$ ;

3. All the item up times and repair service times are independent;
4. Items are served on a FCFS basis. There are no privileged customers.

Therefore, an  $M/M/1/\infty/k$  queueing system is used to model the repair process at a single vendor with arrival rate  $\lambda$ , service rate  $\mu$ , a single server, infinite buffer space and finite population  $k$ . We categorize the  $k$  items into two groups, namely DOWN items and UP items. Suppose at some arbitrary time point  $x$  items are queued at the vendor awaiting repairs. Consequently the number of items which are functioning properly is  $k - x$ . As time evolves, any particular item from the group of  $x$  DOWN items will rejoin the other group when its repair is completed, while on the contrary the  $k - x$  UP items are all subject to breakdown. From standard results, the time to the next breakdown (requiring repairs) has an exponential distribution with rate  $\lambda(k - x)$  where  $\lambda$  is the individual breakdown rate. It appears that a birth-death process (which is a simple special case of a continuous time Markov process) is appropriate for the modelling of the repair process at the vendor, where the birth rates and the death rates in state  $x$  are given by

$$\begin{aligned} \lambda_x &= \lambda(k - x), \quad 0 \leq x \leq k, \\ \text{and} \\ \mu_x &= \mu, \quad 1 \leq x \leq k. \end{aligned} \tag{2.1}$$

It is standard that the process is ergodic. An ergodic process has the property that in the long run it reaches a stationary distribution, irrespective of the initial state. The equilibrium distribution  $\{\Pi_x(k), 0 \leq x \leq k\}$  is expressed by

$$\Pi_x(k) = \hat{\rho}^x \left\{ \prod_{r=0}^{x-1} (k - r) \right\} \Pi_0(k), \quad 0 \leq x \leq k, \tag{2.2}$$

where  $\hat{\rho} = \lambda/\mu$  in (2.2) and

$$\Pi_0(k) = \left\{ \sum_{x=0}^k \hat{\rho}^x \left[ \prod_{r=0}^{x-1} (k - r) \right] \right\}^{-1}. \tag{2.3}$$

## 2.4 Mathematical formulations for the goodwill cost

In this section, We shall consider three different approaches to the modelling of goodwill costs, labeled Models 1, 2 and 3. With the above comments in mind, all three approaches capture the feature of the penalty arising from long waits experienced by individual customers. The distinctive characteristics between the three approaches to the modelling of the goodwill cost as a function of the response time  $r$  (time between a breakdown and the ensuing completion of the corresponding repair) are illustrated in Figure 2.1 to give a graphical perception prior to the mathematical formulations.




Figure 2.1: Plots of the goodwill costs against the response time  $r$  for Models 1,2 and 3 respectively with  $d^1 = 1$ ,  $d^2 = d^3 = 10$ ,  $h = 1$  and  $\tau = 0.04$ .

**Model 1** Under Model 1, a goodwill cost  $d^1$  is incurred whenever an item requiring repair spends more than time  $\tau$  at the vendor. It is natural to assume that all customers

share a common goodwill cost parameter  $d^1$ . Use  $Y$  for the waiting time (including service time) of items and  $L(x)$  for the random variable of the queue length in steady state. When an item arriving at the vendor finds a queue of length  $x$ , its waiting time will have a gamma  $\Gamma(x+1, \mu)$  distribution. We express this by writing

$$Y|L(k) = x \sim \Gamma(x+1, \mu).$$

It follows from standard properties of the gamma distribution that

$$\begin{aligned} d^1 P(Y > \tau | L(k) = x) &= d^1 \sum_{y=0}^x (\mu\tau)^y (y!)^{-1} \exp(-\mu\tau) \\ &\equiv a^1(x), \text{ say.} \end{aligned} \quad (2.4)$$

The quantity  $a^1(x)$  in (2.4) is the expected goodwill cost for a failure just arriving for repair when there are  $x$  items queued ahead of it. The expected goodwill cost  $a^1(x)$  under Model 1 has such a feature that it will quickly reach the limit of  $d^1$  as we increase the vendor population  $k$ . This follows from the standard properties of the Poisson distribution.

**Model 2** It is probably more reasonable to suppose that a cost of  $d^2(w-\tau)^+$  is incurred when an item spends time  $w$  at the vendor. As in Model 1, it is natural to assume that all customers share a common  $d^2$ . Hence, as before, no penalty is incurred when an item spends less time than  $\tau$  at the vendor, but now we suppose that the goodwill cost is proportional to the excess of the response time over  $\tau$ . With a queue of length  $x$ , it follows from standard properties of the gamma distribution and of expectation that

$$\begin{aligned} d^2 E\{(Y-\tau)^+ | L(k) = x\} &= d^2 \int_{\tau}^{\infty} P(Y > w | L(k) = x) dw \\ &= d^2 (\mu)^{-1} \sum_{y=0}^x (x+1-y) (\mu\tau)^y (y!)^{-1} \exp(-\mu\tau) \\ &\equiv a^2(x), \text{ say.} \end{aligned} \quad (2.5)$$

The quantity  $a^2(x)$  in (2.5) is the expected goodwill cost under Model 2 for an arriving failure which joins a repair queue of length  $x$ .

**Model 3** As in Model 2, a goodwill cost of  $d^3(w - \tau)^+$  is incurred if the waiting time of an item is more than  $\tau$ . However, we now consider the introduction of a linear holding cost  $h$  per unit of time when an item spends less than threshold  $\tau$ . In order to reflect a rapid growth in the goodwill cost after  $\tau$ , the linear holding cost  $h$  should be much smaller than the long-waits penalty cost  $d^3$ . Accordingly, we also assume that all customers share a common  $h$  and a common  $d^3$ . The queue-length dependent goodwill cost under Model 3 is mathematically expressed by

$$\begin{aligned}
 hE\{\Upsilon|L(k) = x\} &+ (d^3 - h)E\{(\Upsilon - \tau)^+|L(k) = x\} \\
 &= (d^3 - h) \int_{\tau}^{\infty} P(\Upsilon > w|L(k) = x) dw + h(x + 1)(\mu)^{-1} \\
 &= (d^3 - h)(\mu)^{-1} \sum_{y=0}^x (x + 1 - y)(\mu\tau)^y (y!)^{-1} \exp(-\mu\tau) + h(x + 1)(\mu)^{-1} \\
 &\equiv a^3(x), \text{ say.} \tag{2.6}
 \end{aligned}$$

The quantity  $a^3(x)$  in (2.6) is the expected goodwill cost under Model 3 for an arriving failure when there are  $x$  items ahead of it.

## 2.5 Properties of the proposed cost models

Following the three mathematical formulations of the queue-length dependent goodwill costs in (2.4), (2.5) and (2.6), this section is concerned to present some important theoretical results regarding the cost sequences  $\{a^I(0), a^I(1), a^I(2), \dots\} \equiv \{a^I(x), x \in \mathbb{N}, I = 1, 2, 3\}$ . Note that we say that some function  $f: \mathbb{N} \rightarrow \mathbb{R}$  is convex/concave if there is some  $B \in \mathbb{Z}^+$  such that  $f$  is convex over the range  $[0, 1, \dots, B]$  and concave over  $[B, B + 1, \dots]$ .

### Lemma 2.1 (Monotonicity and convexity of the cost sequences)

- (i) Model 1: The sequence  $\{a^1(x), x \in \mathbb{N}\}$  is increasing and convex/concave;
- (ii) Model 2: The sequence  $\{a^2(x), x \in \mathbb{N}\}$  is increasing and convex.
- (iii) Model 3: The sequence  $\{a^3(x), x \in \mathbb{N}\}$  is increasing and convex.

*Proof.* For (i), see (2.4) and note that it is a well known property of the Poisson distribution that the difference

$$a^1(x+1) - a^1(x) = d^1(\mu\tau)^{x+1}\{(x+1)!\}^{-1}\exp(-\mu\tau) \geq 0$$

is positive and increasing/decreasing in  $x$ . The increasing and convex/concave nature of  $\{a^1(x), x \in \mathbb{N}\}$  follows.

For (ii), observe (2.5) that

$$a^2(x+1) - a^2(x) = d^2(\mu)^{-1} \sum_{y=0}^{x+1} (\mu\tau)^y (y!)^{-1} \exp(-\mu\tau)$$

is increasing in  $x$  and is always positive. The conclusions in (ii) follow easily.

Similarly for (iii), use (2.6) that

$$a^3(x+1) - a^3(x) = (d^3 - h)(\mu)^{-1} \sum_{y=0}^{x+1} (\mu\tau)^y (y!)^{-1} \exp(-\mu\tau) + h(\mu)^{-1}$$

is increasing in  $x$  and is always positive. The conclusions in (iii) follow easily.  $\square$

Lemma 2.1 gives important properties of the goodwill cost sequences under three models. These have great relevance to the theoretical findings in the upcoming chapters, in particular, those reported in Proposition 3.3 and Lemma 5.3.

## 2.6 Conclusions to the chapter

In this chapter, three approaches have been proposed to the modelling of goodwill costs, which should predominate in the overall costs. All of the three cost models discussed take explicit account of the long waits experienced by individual customers when a failure with a warranty is awaiting repair. Such approaches in the formulation of goodwill costs have not been seen in the related area to our knowledge. The primary difference between them lies in the different means by which the penalty incurred by long waits is actually computed. An encouraging result of this chapter is mathematical

elucidations concerning important properties (i.e. monotonicity and convexity) for the three cost models.

## Chapter 3

# Static Allocations for a Fixed Warranty Population

### 3.1 Introduction to the Chapter

In this chapter, we will study the problems of the subcontracting of warranty repairs within the framework of a static allocation model. We consider the following scenario. A large manufacturer sells a fixed amount of identical items,  $A$  say, with a warranty period  $\tau$  allocated upon purchase among  $V$  vendors to which the warranty responsibility is outsourced. If any failure of a purchased item happens in the duration of its warranty, the customer will call his allocated vendor directly to arrange for collection. The broken item then joins a queue to get repaired at the assigned vendor where the cost of this service will be actually incurred. In general, the manufacturer absorbs the entire repair costs. From the point of view of the manufacturer the objective of the problem is therefore seeking minimal overall costs by allocating items to the collection of outsourced vendors in an optimal manner. Clearly, the modelling of the cost structure in such a study is absolutely crucial. Customers' loyalty plays an important role in contemporary business life. Without the consideration of goodwill costs, the manufacturer might allocate all the items to the cheapest vendor to save costs. Under such a situation, customers are very likely to be dissatisfied with delayed services provided, which might discourage spent business in future times. Hence, the overall costs of the

## Chapter 3

# Static Allocations for a Fixed Warranty Population

### 3.1 Introduction to the Chapter

In this chapter, we will study the problem of the outsourcing of warranty repairs within the framework of a static allocation model. We consider the following scenario. A large manufacturer sells a fixed number of identical items,  $K$  say, with a warranty. Items are allocated upon purchase among  $V$  vendors to which the warranty repair service is outsourced. If any failure of a purchased item happens in the duration of its warranty, the customer will call its allocated vendor directly to arrange for collection. The broken item then joins a queue to get repaired at the assigned vendor where the cost of this service will be actually incurred. In general, the manufacturer absorbs the entire repair costs. From the point of view of the manufacturer the objective of the problem is therefore seeking minimal overall costs by allocating items to the collection of outsourced vendors in an optimal manner. Plainly, the modelling of the cost structure in such a study is absolutely crucial. Customers' loyalty plays an important role in contemporary business life. Without the consideration of goodwill costs, the manufacturer might allocate all the items to the cheapest vendor to cut costs. Under such a situation, customers are very likely to be dissatisfied with delayed service received, which might discourage repeat business in future sales. Hence, the overall costs of the

problem are a combination of the physical costs (e.g. labour, parts and etc.) and the invisible goodwill costs which should in our view predominate in the cost structure. Departing from the linear holding cost model proposed by Opp et al. (2003), the goodwill costs which impose penalties for unacceptably long waits have been formulated in three modelling approaches by Chapter 2.

In outline, the rest of the chapter has the following structure. Section 3.2 will describe suitable forms of a cost rate function  $f(k)$  by utilizing the cost models and the stochastic repair process at a single vendor to which  $k$  items are allocated. Then, the static allocation model is formulated as a resource allocation problem with integer variables in Section 3.3. The status of a greedy solution to this problem is discussed in Section 3.4. Section 3.5 presents a detailed numerical study with some sensitivity analysis. This chapter concludes with some remarks as usual.

## 3.2 The Formulation of a Cost Rate Function at a Single Vendor

Suppose  $k_v$  items (the decision variable of the model of interest) are allocated to vendor  $v$ ,  $1 \leq v \leq V$ . Any failure of these  $k_v$  items would be taken care of by this specific vendor  $v$ . The quantity  $f_v(k_v)$  denotes a consequential cost per unit time incurred at vendor  $v$  with the allocated workload  $k_v$ . In this section, we will focus on a single vendor for the ultimate goal of the formulation of the cost rate function  $f_v(k_v)$ . The vendor subscript  $v$  is redundant during this discussion and hence will be dropped throughout this section.

The breakdown and repair process at a single vendor has been modelled as a finite population queueing system  $M/M/1/\infty/k$  with arrival rate  $\lambda$ , service rate  $\mu$ , a single server, infinite buffer space and finite population  $k$  in Chapter 2. The associated equilibrium distribution  $\{\Pi_x(k), 0 \leq x \leq k\}$  for that ergodic process can be calculated more efficiently by the following transformation.

The following will describe an efficient means of calculating the equilibrium distribution  $\{\Pi_x(k), 0 \leq x \leq k\}$ , and the expected queue length (denoted  $\bar{L}(k)$ ) recursively. Firstly, a recursion for  $\Pi_0(k)$  in (2.3) is given by

$$\begin{aligned}\Pi_0^{-1}(0) &= 1; \\ \Pi_0^{-1}(k) &= 1 + \hat{\rho}k\Pi_0^{-1}(k-1), k \in \mathbb{N}.\end{aligned}\quad (3.1)$$

where  $\hat{\rho} = \lambda/\mu$ . By far, the equilibrium distribution in (2.2) could be computed with less effort by using (3.1). Nevertheless, we may make the further simplifications by re-expressing (2.2) as:

$$\Pi_x(k) = A_x(k)\Pi_0(k), 0 \leq x \leq k,$$

where the quantity

$$A_x(k) = \hat{\rho}^x \left\{ \prod_{r=0}^{x-1} (k-r) \right\}$$

satisfies the recursion

$$A_x(k) = \begin{cases} 1, & x = 0; \\ \hat{\rho}(k-x+1)A_{x-1}(k), & 0 < x \leq k. \end{cases}$$

The computation of mean queue length  $\bar{L}(k)$  is related closely to this way of computing the equilibrium distribution  $\{\Pi_x(k), 0 \leq x \leq k\}$ . We write

$$\begin{aligned}\bar{L}(k) &= \sum_{x=0}^k x\Pi_x(k) \\ &= \sum_{x=0}^k xA_x(k)\Pi_0(k) \\ &= B(k)\Pi_0(k)\end{aligned}$$

It is straightforward to show algebraically that  $B(k)$  also satisfies the recursion

$$B(k) = \begin{cases} 0, & k = 0; \\ \hat{\rho}k\Pi_0^{-1}(k-1) + \hat{\rho}kB(k-1), & k \in \mathbb{N}.\end{cases}$$

By using (2.2) and the above recursion for  $B(k)$ ,  $\bar{L}(k)$  could be recursively expressed by

$$\begin{aligned} \bar{L}(k) &= [1 - \Pi_0(k)][1 + \bar{L}(k - 1)], k \in \mathbb{N} \\ \bar{L}(0) &= 0 \end{aligned} \tag{3.2}$$

With the simple birth and death rates of the breakdown and repair process (2.1) introduced in Chapter 2 and the efficient way of computing the equilibrium distribution as shown above, we are able to proceed to the formulation of the cost rate function  $f(k)$ . We have considered three approaches to the modelling of the goodwill costs in Chapter 2, and hence we will show the development of the cost rate functions  $f(k) \equiv \{f^I(k), I = 1, 2, 3\}$  for three formulations, where  $I$  is the cost model index. The goodwill cost rate functions  $g(k) \equiv \{g^I(k), I = 1, 2, 3\}$  and the repair cost rates  $r(k) \equiv \{r^I(k), I = 1, 2, 3\}$  are introduced individually since these two components of the costs rate  $f(k)$  tend to have opposite properties in terms of convexity which is important and will be discussed later in this chapter.

Recalling goodwill costs model  $I, I = 1, 2, 3$  in Chapter 2, the goodwill cost rate in steady state for a single vendor with a fixed population of  $k$  is given by

$$\sum_{x=0}^k \Pi_x(k) \lambda(k-x) a^I(x) \equiv g^I(k), I = 1, 2, 3. \tag{3.3}$$

Additionally, the repair cost rate in steady state when the single vendor population is  $k$  and a common fixed repair cost per item  $c$  is applied is given by

$$\sum_{x=0}^k \Pi_x(k) c \lambda(k-x) \equiv r^I(k), I = 1, 2, 3. \tag{3.4}$$

Overall, the cost rate in steady state with population size  $k$  at the vendor is the sum of (3.3) and (3.4) and may be expressed as

$$\sum_{x=0}^k \Pi_x(k) \lambda(k-x) \{c + a^I(x)\} \equiv f^I(k), I = 1, 2, 3. \tag{3.5}$$

Before moving to next the section, we shall also present the linear holding cost model of Opp et al. (2003) with the associated cost rate function  $f^L$  as

$$\sum_{x=0}^k \Pi_x(k) \{\lambda(k-x)c + hx\} =$$

$$\lambda kc + (h - \lambda c)\bar{L}(k) \equiv f^L(k) \quad (3.6)$$

### 3.3 Discrete Resource Allocation Model

With the formulation of cost rate functions  $f^I(k)$ ,  $I = 1, 2, 3$  in place, we now restore the vendor subscript  $v$ ,  $1 \leq v \leq V$  and try to decide the best allocation  $k_v$  to each contracted vendor in order to achieve minimal total cost rate. The resource allocation problem addressed in the standard literature is an optimization problem with a single constraint.

Given a fixed amount of *resource*, one is asked to determine its best allocation to a number of *activities* so that the objective function under consideration is optimized.

Such a problem is described by Ibaraki (1988). This simply structured resource allocation model has a wide variety of applications, and the static allocation model for the outsourcing of warranty repairs is one of them. The *resource* now becomes the population of purchased items, whose size is fixed, say  $K$ , and the *activities* are substituted by  $V$  contracted vendors. We seek a minimal overall cost rate incurred by the resulting allocation. Happily, the resource allocation problem that we encounter here is an even simpler version with discrete variables and a separable objective function where each of the additive components  $f_v^I$  only depends on one decision variable  $k_v$ . The optimal allocation problem for cost model  $I$  may be expressed by

$$\begin{aligned} \min \quad & \sum_{v=1}^V f_v^I(k_v) \\ \text{s.t.} \quad & \sum_{v=1}^V k_v = K, \\ & k_v \in \mathbb{N}, \quad v = 1, \dots, V. \end{aligned} \quad (3.7)$$

where  $I = 1, 2, 3$ . Dynamic programming (DP) is a conventional approach to solving the problem (3.7) in  $O(VK^2)$  time, given that the evaluation of each of the  $f_v^I$ 's is done in constant time, see Ibaraki (1988). When we increase the population size to realistic levels the computational complexity will escalate much faster than when we increase the number of vendors. We use  $F^I(v, \beta)$  to denote the best expected overall cost rate

when allocating  $\beta$  items to  $v$  vendors in the order  $1, \dots, v$ , where  $\beta = 0, 1, 2, \dots, K$  and  $v = 1, \dots, V$ . The DP recursion for problem (3.7) is given as follows:

$$F^I(v, \beta) = \begin{cases} \min_{k_v: k_v \leq \beta} (f_v^I(k_v) + F^I(v-1, \beta - k_v)), & 1 < v \leq V \\ f_v^I(\beta), & v = 1 \end{cases}$$

The recurrence will continue till  $F^I(V, K)$  is obtained. The DP procedure to solve the problem (3.7) is formally presented as

### DP Procedure for (3.7)

Step 1: Let  $F^I(1, \beta) = f_1^I(\beta)$  for  $\beta = 0, 1, \dots, K$ . Let  $v=2$  and go to Step 2.

Step 2: If  $v = V$ , go to Step 3, otherwise compute

$$F^I(v, \beta) = \min_{k_v} [f_v^I(k_v) + F^I(v-1, \beta - k_v) | k_v = 0, 1, \dots, \beta],$$

for  $\beta = 0, 1, \dots, K$ . Let  $v = v + 1$  and return to the beginning of Step 2.

Step 3: If  $v = V$ , compute  $F^I(V, K)$  by

$$F^I(V, K) = \min_{k_V} [(f_V^I(k_V) + F^I(V-1, K - k_V) | k_V = 0, 1, \dots, K]$$

and halt.  $F^I(V, K)$  gives the optimal objective value of the static allocation problem considered.

Though DP guarantees optimal solutions for the problem (3.7), its backward induction is exhaustive. Further, its computational complexity depends on the the number of vendors and the size of population. Hence, the DP procedure is not a computationally efficient option for problems with realistic size. We therefore seek some heuristic alternatives which may be implemented efficiently in the following section.

## 3.4 Heuristic Solution: Greedy Algorithm

The greedy algorithm (GA) is straightforward and simple in the sense that it makes decisions at each stage which are locally optimal without any consideration of future

consequences. This is the main difference between GA and DP. DP is exhaustive and is guaranteed to find an optimal solution. For solving multi-stage problems, all the decisions made in the previous stages by DP become the basis for ensuing decision stages, while GA makes the best decision for the current stage only, and will never reconsider its old decisions. Indeed, GA often finds a local rather than global optimum for optimization problems in general. Nevertheless, it is easy to show by means of a pairwise interchange argument that when each of the  $f_i^I$ 's is increasing and convex then a greedy algorithm will yield a global optimum to the problem (3.7). This was first found by Gross (1956), see also Fox (1966). The greedy algorithm when applied to problem (3.7) operates as follows

### Greedy Algorithm for (3.7)

Step 0: Set  $k_v = 0, 1 \leq v \leq V$ ;

Step 1: Choose any  $j \in \operatorname{argmin}_{1 \leq v \leq V} [f_v^I(k_v + 1) - f_v^I(k_v)], I = 1, 2, 3$ ;

Step 2: Set  $k_j = k_j + 1$ ;

Step 3: If  $\sum_{v=1}^V k_v < K$ , go to Step 1; otherwise stop.

We now proceed to explore two possible properties that cost rate functions  $f_v^I$ 's may possess, namely monotonicity and convexity. To simplify notation, we will drop the vendor subscript  $v$  again for further discussion on these two issues. Let  $L(k)$  (repair queue length) be a random variable such that  $L(k) \sim \Pi(k)$ , the stationary distribution where  $\Pi(k) \equiv \{\Pi_x(k), 0 \leq x \leq k\}$ .

**Definitions** If  $X$  is a Poisson random variable with parameter (the mean)  $\hat{\rho}^{-1}$  and  $Y$  is a random variable whose distribution is that of  $X$  conditioned on the event  $X \leq k$ , we then say that  $Y$  has the Poisson distribution with parameter  $\hat{\rho}^{-1}$ , truncated at  $k$ .

**Lemma 3.1** *The number of up items,  $k - L(k)$ , has the Poisson distribution with parameter  $\hat{\rho}^{-1}$ , truncated at  $k$ .*

*Proof.* Fix  $x$  in the range  $0 \leq x \leq k$ . From (2.2) and (2.3) we have that

$$\begin{aligned}
 P(k-L(k) = x) &= P(L(k) = k-x) = \Pi_{k-x}(k) \\
 &= \hat{\rho}^{k-x} \left\{ \prod_{r=0}^{k-x-1} (k-r) \right\} \left\{ \sum_{y=0}^k \hat{\rho}^y \left[ \prod_{r=0}^{y-1} (k-r) \right] \right\}^{-1} \\
 &= \hat{\rho}^{-x} (x!)^{-1} \left[ \sum_{y=0}^k \hat{\rho}^{-y} (y!)^{-1} \right]^{-1} \\
 &= \hat{\rho}^{-x} (x!)^{-1} \exp(-\hat{\rho}^{-1}) \left[ \sum_{y=0}^k \hat{\rho}^{-y} (y!)^{-1} \exp(-\hat{\rho}^{-1}) \right]^{-1}
 \end{aligned}$$

The results follows. □

**Comment** The above result implies that when the vendor population  $k$  is large, the number of up items has a distribution which is close to being  $k$ -independent. More specifically, it is close to having a Poisson distribution with parameter  $\hat{\rho}^{-1}$ . It follows that when  $k$  is large, the effect in steady state of increasing the vendor population size by a single item is very close to increasing the queue length at the vendor by one. Lemma 3.1 also contributes to the proof of the monotonicity of cost rate functions  $f^l$ s in Proposition 3.3.

**Definitions** If  $X$  is a random variable with distribution function  $F$  and  $Y$  a random variable with distribution function  $G$  we say that  $X$  is **stochastically larger** than  $Y$ , written  $X \geq_{st} Y$ , if  $F(x) \leq G(x), x \in \mathbb{R}$ . The sequence of random variables  $\{X_k, k \in \mathbb{N}\}$  is **stochastically increasing** if  $X_{k+1} \geq_{st} X_k, k \in \mathbb{N}$ .

**Lemma 3.2** *The sequences of the number of up items  $\{k-L(k), k \in \mathbb{N}\}$  and the number of down items  $\{L(k), k \in \mathbb{N}\}$  are both stochastically increasing.*

*Proof.* The claim in the statement of the Lemma in relation to the sequence  $\{k-L(k), k \in \mathbb{N}\}$  is a trivial consequence of Lemma 3.1. In considering the number of down items, write  $F_k$  for the distribution function of  $L(k)$ . Using Lemma 3.1 we have that, for any  $x$  in the range  $0 \leq x \leq k$ ,

$$1 - F_k(x) = P\{k-L(k) \leq k-x\};$$

$$\begin{aligned}
&= \left[ \sum_{y=0}^{k-x} \hat{\rho}^{-y} (y!)^{-1} \exp(-\hat{\rho}^{-1}) \right] \left[ \sum_{y=0}^k \hat{\rho}^{-y} (y!)^{-1} \exp(-\hat{\rho}^{-1}) \right]^{-1} \\
&= A(k-x) [A(k)]^{-1},
\end{aligned}$$

where

$$A(k) \equiv \sum_{y=0}^k \hat{\rho}^{-y} (y!)^{-1} \exp(-\hat{\rho}^{-1}).$$

It is straightforward to show that

$$A(k-x) [A(k)]^{-1} \leq A(k+1-x) [A(k+1)]^{-1}$$

if and only if

$$\left[ \sum_{y=0}^{k-x} \hat{\rho}^{-y} (y!)^{-1} \right] \hat{\rho}^{-k-1} [(k+1)!]^{-1} \leq \left[ \sum_{y=0}^k \hat{\rho}^{-y} (y!)^{-1} \right] \hat{\rho}^{-k-1+x} [(k+1-x)!]^{-1}$$

if and only if

$$\sum_{y=0}^{k-x} \hat{\rho}^{-k-1-y} \{y!(k+1)!\}^{-1} \leq \sum_{y=0}^{k-x} \hat{\rho}^{-k-1-y} \{(x+y)!(k+1-x)!\}^{-1}.$$

However, the latter inequality is a trivial consequence of the fact that

$$(x+y)!(k+1-x)! \leq (k+1)!y!, 0 \leq y \leq k-x.$$

It follows from the above calculation that

$$1 - F_k(x) \leq 1 - F_{k+1}(x) \Rightarrow F_k(x) \geq F_{k+1}(x), 0 \leq x \leq k,$$

and hence trivially that

$$F_k(x) \geq F_{k+1}(x), x \in \mathbb{R}.$$

It now follows that  $L(k+1) \geq_{st} L(k)$ , as required.  $\square$

**Comment** It follows from Lemma 3.2 that, for our models, **both** of the contributing components to the overall cost rate will increase as the population size  $k$  grows. These components are (i) the repair cost rate, which is related to  $k - L(k)$ , the number of up items (hence vulnerable to breakdown), and (ii) the goodwill cost rate, which is related

to  $L(k)$ , the number of down items.

By using Lemma 3.2 and Lemma 2.1 we are able to show that the cost rates  $f^I$ s for cost models  $I = 1, 2, 3$  are increasing in the vendor population  $k$ . This result is formally presented in Proposition 3.3.

It follows by a standard result that if  $X \geq_{st} Y$  then  $E\{\Phi(X)\} \geq E\{\Phi(Y)\}$  for any increasing function  $\Phi$ .

**Proposition 3.3** *Cost rate functions  $f^I(k)$  are increasing in  $k, I = 1, 2, 3$ .*

*Proof.* First, note from (2.2) that

$$\Pi_x(k)\lambda(k-x) = \Pi_{x+1}(k)\mu, \quad 0 \leq x \leq k-1. \quad (3.8)$$

Hence, from (3.5), the formula for  $f^I(k)$  may be rewritten

$$\mu c \sum_{x=0}^{k-1} \Pi_{x+1}(k) + \mu \sum_{x=0}^{k-1} \Pi_{x+1}(k) a^I(x) = \mu E[\Phi_1\{L(k)\} + \Phi_2\{L(k)\}],$$

where

$$\Phi_1(x) = \begin{cases} c, & x \geq 1; \\ 0, & x < 1, \end{cases}$$

and

$$\Phi_2(x) = \begin{cases} a^I(x-1), & x \geq 1; \\ 0, & x < 1. \end{cases}$$

where  $I = 1, 2, 3$ .

However, from Lemma 2.1, both  $\Phi_1$  and  $\Phi_2$  are non-decreasing. The result now follows from Lemma 3.2 and the property of stochastic ordering stated before the statement of the Proposition.  $\square$

With the knowledge of convexity of  $\bar{L}(k)$  for the finite population queueing model, which is shown by Opp et al. (2003), we rewrite the repair costs rate element (3.4) as

$$\sum_{x=0}^k \Pi_x(k) \lambda(k-x)c = \lambda kc - \lambda c \bar{L}(k) \equiv r^I(k), I = 1, 2, 3 \quad (3.9)$$

which is concave in  $k$ . In order to show the degree of the convexity which the cost rates  $f^I$ s may possess, the goodwill cost rate element  $g^I$  shall be discussed separately from the repair cost rate elements  $r^I$ . We also argue (with support of Opp et al. (2003)) that the goodwill cost contributions to the expected cost rate functions  $f^I$  should predominate. The following Figure 3.1 and Proposition 3.4 provide a firm basis to the claim of strong convexity of the  $g^I(k), I = 1, 2, 3$  over a range of population  $k$  of interest.

Consider a vendor with service rate  $\mu = 62.5$  per year, service quality threshold  $\tau = 0.04$  per year and individual item breakdown rate  $\lambda = 1.2$  per year for the whole population  $K = 100$ . Plots of  $g^I$  drawn against vendor population  $k$  for Models 1, 2 and 3 respectively are shown in the left side of Figure 3.1. We choose the goodwill cost parameter  $d^I = 10, I = 1, 2, 3$  for all three models and the linear holding cost rate  $h = 1$  for cost model 3 only.

Note that  $g^2$  and  $g^3$  are convex throughout the range  $[0, 100]$  while  $g^1$  is convex up to 52 and concave beyond.

**Definitions** If an increasing function  $g : \mathbb{N} \rightarrow \mathbb{R}$  is convex over the range  $[0, 1, \dots, B]$  and  $B$  is maximal in this regard, we shall refer  $B$  as its *convexity boundary point*.

Next we incorporate the concave element and set the repair cost parameter  $c = 2, I = 1, 2, 3$ , the plots of  $f^I$  against vendor population  $k$  for cost models 1, 2 and 3 are shown in the right hand of Figure 3.1. As expected, the *convexity boundary point* is reduced by the influence of the concavity of the repair costs function  $r^I$ .

In the above setting, the *convexity boundary point*  $B = 51$  for cost model 1 while  $f^2$  and  $f^3$  are still convex throughout the range. Nevertheless, we reiterate the argument

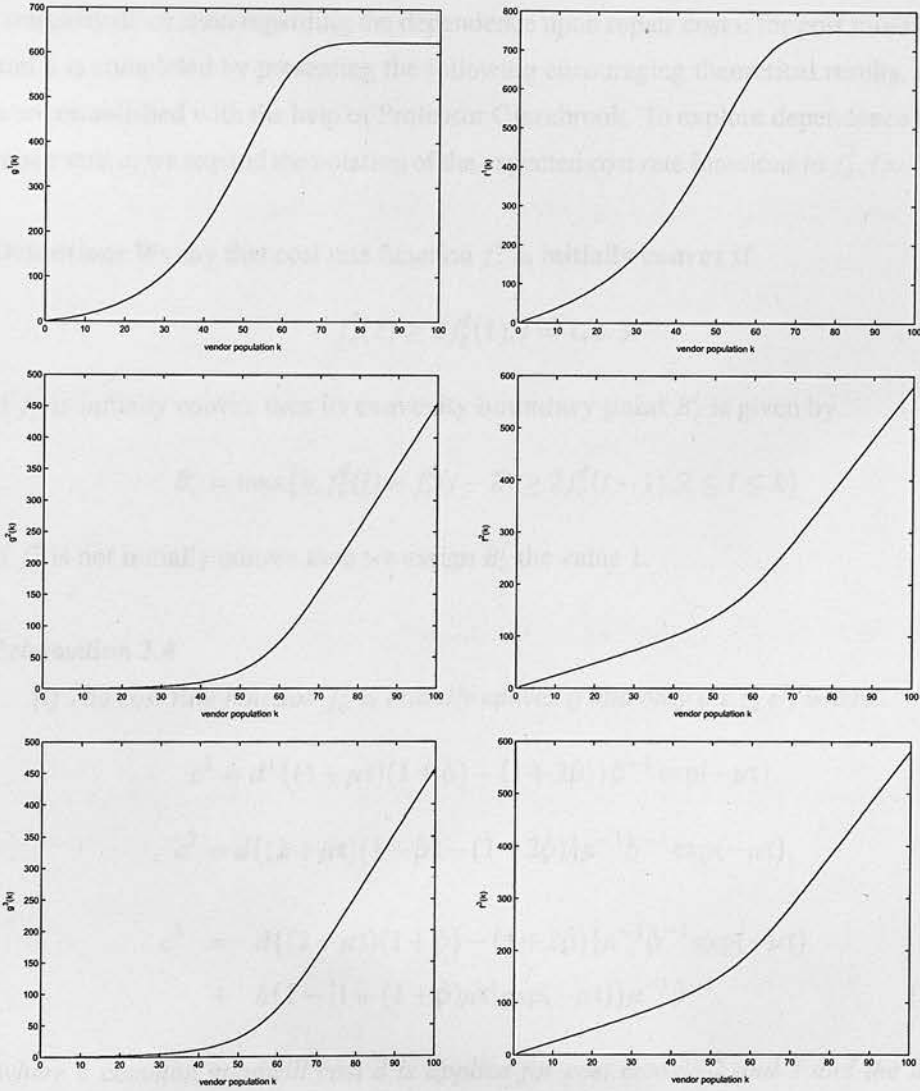


Figure 3.1: Plots of  $g^I(k)$ (the goodwill costs rate) and  $f^I(k)$  (the expected costs rate),  $I = 1, 2, 3$  against vendor population  $k$  for cost models 1,2 and 3 respectively

that the goodwill cost should predominate over the repair cost. The approach that we took in modelling the goodwill cost aims to make any long delay beyond the time threshold  $\tau$  very unlikely. Hence, if the parameters of  $c$  and  $d^I$  are chosen reasonably, we will have a large convexity range instead of full convexity for cost rate function  $f^1$ , while  $f^2$  and  $f^3$  should be fully convex over the range of practical interest. The

convexity discussion regarding the dependence upon repair cost  $c$  for cost models 1, 2 and 3 is completed by presenting the following encouraging theoretical results, which were established with the help of Professor Glazebrook. To explore dependence on the repair cost  $c$ , we expand the notation of the expected cost rate functions to  $f_c^I, I = 1, 2, 3$ .

**Definitions** We say that cost rate function  $f_c^I$  is **initially convex** if

$$f_c^I(2) \geq 2f_c^I(1), I = 1, 2, 3.$$

If  $f_c^I$  is initially convex then its **convexity boundary point**  $B_c^I$  is given by

$$B_c^I = \max\{k; f_c^I(l) + f_c^I(l-2) \geq 2f_c^I(l-1), 2 \leq l \leq k\}$$

If  $f_c^I$  is not initially convex then we assign  $B_c^I$  the value 1.

**Proposition 3.4**

(i) The cost rate function  $f_c^I$  is initially convex if and only if  $c \leq c^I$ , where

$$c^1 = d^1 \{(1 + \mu\tau)(1 + \hat{\rho}) - (1 + 2\hat{\rho})\} \hat{\rho}^{-1} \exp(-\mu\tau), \tag{3.10}$$

$$c^2 = d \{(2 + \mu\tau)(1 + \hat{\rho}) - (1 + 2\hat{\rho})\} \mu^{-1} \hat{\rho}^{-1} \exp(-\mu\tau), \tag{3.11}$$

$$c^3 = d \{(2 + \mu\tau)(1 + \hat{\rho}) - (1 + 2\hat{\rho})\} \mu^{-1} \hat{\rho}^{-1} \exp(-\mu\tau) + h \{1 - [1 + (1 + \hat{\rho})\mu\tau] \exp(-\mu\tau)\} \mu^{-1} \hat{\rho}^{-1}. \tag{3.12}$$

where a common goodwill cost  $d$  is applied for cost models 2 and 3 and the linear holding cost  $h < d$ .

(ii) The convexity boundary point  $B_c^I$  is decreasing in  $c$ .

*Proof.* (i) For definiteness consider cost model 1. Cost rate function  $f_c^1$  is initially convex if and only if  $f_c^1(2) \geq 2f_c^1(1)$ , which becomes, upon evaluation of these quantities using (3.5), (2.2) and (2.3)

$$\begin{aligned} \{2\hat{\rho}[c + a(0)] + 2\hat{\rho}^2[c + a(1)]\}(1 + 2\hat{\rho} + 2\hat{\rho}^2)^{-1} &\geq 2\hat{\rho}[c + a(0)](1 + \hat{\rho})^{-1} \\ \Leftrightarrow [c + a(1)][c + a(0)]^{-1} &\geq (1 + 2\hat{\rho})(1 + \hat{\rho})^{-1} \end{aligned} \tag{3.13}$$

This becomes, upon utilization of the form of  $a^1(x)$  in (2.4) from Chapter 2,

$$[d + d\mu\tau + c \exp(\mu\tau)][d + c \exp(\mu\tau)]^{-1} \geq (1 + 2\hat{\rho})(1 + \hat{\rho})^{-1}. \quad (3.14)$$

The conclusion in (i) for the  $I = 1$  case now follows trivially from (3.14). The case  $I = 2, 3$  is dealt with similarly. Note that expression (3.12) partitions  $c^3$  between  $c^2$  and the quantity multiplied by the linear holding cost  $h$  given that a common goodwill cost  $d$  is applied to cost models 2 and 3.

(ii) Consider two distinct repair costs  $\tilde{c} < \bar{c}$ . Firstly note that it follows from (i) that if  $\bar{c} > c^I$  then we must have  $B_{\bar{c}}^I = 1$  and  $B_{\tilde{c}}^I \geq B_{\bar{c}}^I$  trivially. Suppose now that  $\bar{c} \leq c^I$  and hence that  $B_{\tilde{c}}^I \geq 2$ . Explicitly, we have that

$$f_{\tilde{c}}^I(k) + f_{\tilde{c}}^I(k-2) \geq 2f_{\tilde{c}}^I(k-1), 2 \leq k \leq B_{\tilde{c}}^I, I = 1, 2, 3. \quad (3.15)$$

But, from the definition of the cost rate functions in Section 3.2 it follows that

$$f_{\tilde{c}}^I(k) = f_{\bar{c}}^I(k) + (\tilde{c} - \bar{c}) \sum_{x=0}^k \Pi_x(k) \lambda(k-x), I = 1, 2, 3. \quad (3.16)$$

where from (3.4), the second term in the right hand side of (3.16) is concave in  $k$ . It then follows from (3.15) and (3.16) that

$$\begin{aligned} f_{\tilde{c}}^I(k) + f_{\tilde{c}}^I(k-2) &= f_{\bar{c}}^I(k) + f_{\bar{c}}^I(k-2) \\ &\quad + (\tilde{c} - \bar{c}) \sum_{x=0}^k \Pi_x(k) \lambda(k-x) + (\tilde{c} - \bar{c}) \sum_{x=0}^{k-2} \Pi_x(k-2) \lambda(k-2-x) \\ &\geq 2f_{\bar{c}}^I(k-1) + 2(\tilde{c} - \bar{c}) \sum_{x=0}^{k-1} \Pi_x(k-1) \lambda(k-1-x) \\ &= 2f_{\bar{c}}^I(k-1), 2 \leq k \leq B_{\tilde{c}}^I, I = 1, 2, 3. \end{aligned} \quad (3.17)$$

It now follows from (3.17) that  $B_{\tilde{c}}^I \geq B_{\bar{c}}^I, I = 1, 2, 3$ . and the result follows.  $\square$

Proposition 3.4 gives a formal mathematical expression to a range of ideas which have an important bearing on the convexity of our cost rate functions and consequential likely status of greedy heuristics. First, one should note that the conditions expressed in (3.10), (3.11) and (3.12) will be met comfortably for sensibly chosen model parameters and hence our cost rate functions can be safely assumed to be at least initially

convex. Further, by comparing expressions (3.11) with (3.12), it is straightforward to infer the fact that the cost rate  $f^3$  enjoys stronger convexity properties than does  $f^2$  when  $\hat{p}$  in the quantity of the last term in the right-hand side of (3.12) is realistically chosen. Proposition 3.4(ii) points to the fact that the initial range in which the cost rate functions are increasing, convex will decrease as the repair cost element increases. For the linear holding cost model, Opp et al.(2003) were able to insist that the repair cost  $c$  should be sufficiently small that the resulting overall cost rate  $f^L$  remain increasing, convex and hence that greedy heuristics are guaranteed to provide optimal solutions to the allocation problem. This is not possible for our more complex models. Rather, we can comparably require that repair costs  $c$  for our cost models 1, 2 and 3 are sufficiently small that the corresponding convexity boundary points for the cost rates are large (see Proposition 3.4(ii)) and hence, following Lemma 3.5 below, that greedy heuristics are likely to perform well for the static allocation problem. Our numerical results indicate that this happens most of the time.

In light of the findings of Gross (1956), we conclude this section with another motivating theoretical result. We restore the vendor subscript  $v$  and drop the cost model superscript  $I$  for this aspect.

Let  $B_v$  be the convexity boundary point for the increasing function  $f_v, 1 \leq v \leq V$ . Suppose that  $\sum_{v=1}^V B_v \geq K$ . Consider the following variant of optimization problem (3.7):

$$\begin{aligned} \min \quad & \sum_{v=1}^V f_v(k_v) \\ \text{s.t.} \quad & \sum_{v=1}^V k_v = K, \\ & k_v \leq B_v, \\ & k_v \in \mathbb{N}, \quad v = 1, \dots, V. \end{aligned} \tag{3.18}$$

A greedy solution to problem (3.18) proceeds as follows:

### Greedy Algorithm for (3.18)

Step 0: Set  $k_v = 0, 1 \leq v \leq V$ ;

Step 1: Choose any  $j \in \operatorname{argmin}_{1 \leq v \leq V} [f_v(k_v + 1) - f_v(k_v)]$  when the *argmin* is taken over those  $v$  for which  $k_v < B_v$ ;

Step 2: Set  $k_j = k_j + 1$ ;

Step 3: If  $\sum_{v=1}^V k_v < K$ , go to Step 1; otherwise stop.

**Lemma 3.5** *The above greedy algorithm solves optimization problem (3.18).*

*Proof.* We construct function  $\bar{f}_v$  such that it equals  $f_v$  when  $k_v$  is in the range of  $[0, B_v]$  and remains convex for  $k_v \geq B_v$  by adding an artificial straight line with slope  $\Psi$  beyond the boundary point  $B_v$ :

$$\bar{f}_v(k_v) = \begin{cases} f_v(k_v), & 0 \leq k_v \leq B_v; \\ f_v(B_v) + (k_v - B_v)\Psi, & k_v \geq B_v + 1. \end{cases} \quad (3.19)$$

where

$$\Psi = \max_{1 \leq v \leq V} [f_v(B_v) - f_v(B_v - 1)].$$

Since  $\bar{f}_v$  and  $f_v$  are equal over the convex range  $[0, B_v]$ , it must be true that any solution to the following optimization problem for which  $k_v \leq B_v$ ,  $1 \leq v \leq V$ , must solve (3.18) as well.

$$\begin{aligned} \min \quad & \sum_{v=1}^V \bar{f}_v(k_v) \\ \text{s.t.} \quad & \sum_{v=1}^V k_v = K, \\ & k_v \in \mathbb{N}, \quad v = 1, \dots, V. \end{aligned} \quad (3.20)$$

By construction, the  $\bar{f}_v$ s are increasing convex for the entire range and for all  $v$ ,  $1 \leq v \leq V$ . Hence (3.20) is solved by the greedy algorithm for (3.7) above. Note that

$$\begin{aligned} k_v \geq B_v & \Rightarrow \bar{f}_v(k_v + 1) - \bar{f}_v(k_v) = \Psi \\ & = \max_{1 \leq v \leq V} \{f_v(B_v) - f_v(B_v - 1)\} \\ & = \max_{1 \leq v \leq V} \{ \max_{k_v \leq B_v} [f_v(k_v) - f_v(k_v - 1)] \} \\ & = \max_{1 \leq v \leq V} \{ \max_{k_v \leq B_v} [\bar{f}_v(k_v) - \bar{f}_v(k_v - 1)] \}. \end{aligned} \quad (3.21)$$

It follows from (3.21) that the greedy algorithm designed for (3.20) needs never allocate further items to any vendor  $v$  for which  $k_v \geq B_v$ . Hence this algorithm exactly coincides with the above form for (3.18). This concludes the proof.  $\square$

**Comment** It is known that the search process under the greedy algorithm is very likely to terminate at a local minimum. Nevertheless, if we have proved that the search beyond the convexity boundary points  $B_v$ s will not generate a better solution than the one we find within the convexity range, then the solution given by the greedy algorithm is not just a local minimum but a global one. Further, our experimental results in the next section verify these theoretical findings numerically.

### 3.5 Numerical Results and Related Analysis

A wide range of numerical investigations have been conducted for the static allocation problem addressed in earlier sections for the primary goal of assessing the status of the greedy algorithm as a sound heuristic. To achieve this, every GA solution has been compared with a DP solution to check its optimality. The numerical study also includes the following subsidiary goals,

- To explore if the optimal solutions are bounded by the convexity boundary points;
- To carry out an analysis on the sensitivity of both convexity boundary points and optimal allocations to the choice of key parameters;
- To compare the greedy solutions with some *ad hoc* heuristics to see how much can be saved by using the GA.

Some results from our numerical study will be presented in the tables below to shed light on the above issues. First of all, the choice of essential parameters are made in accordance with Opp et al. (2003) and are based on the information acquired from an industrial context in which the items under considerations are PCs. We suppose that the manufacturer contracted with four vendors,  $V=4$ , to undertake the warranty repair

work for PCs, where the individual PC failure rate is  $\lambda = 1.2$  breakdowns per year. The service time threshold  $\tau$  is taken to be 10 working days (i.e.  $\tau = 0.04$  years). We also consider six profiles with regard to the distribution of service capacities among the four vendors. From profile 1 through to profile 6, the degree of inequality between the service capacities among the four vendors is increasing. Specifically, the service rate  $\mu_{vj}$  for vendor  $v$  in profile  $j$  is given by

$$\mu_{v1} = 0.625K, \quad 1 \leq v \leq 4$$

and

(3.22)

$$\mu_{vj} = 2.5Ky_j^{v-1}(1-y_j)(1-y_j^4)^{-1}, \quad 1 \leq v \leq 4, 2 \leq j \leq 6$$

where

$$y_j = 1 - 0.1(j-1), \quad 2 \leq j \leq 6.$$

Hence, with the above setting all vendors have equal service rates in profile 1 while in profile 6 vendor 1 has a service rate which is around eight times that for vendor 4. Nevertheless, the overall service rates of four vendors always sum to  $2.5K$  throughout the six profiles, which is more than twice the approximating arrival stream rate ( $< \lambda K$ ). This setting is to make sure that there are sufficient service capacities across the four vendors to accommodate the entire population of  $K$  PCs under warranty constantly.

The computational experiments that have been carried out are for cost models 1 and 2 only, because the plots for cost models 2 and 3 have a very similar shape, see Figure 3.1. Besides, cost model 3 should have strong convexity properties with the inclusion of the initial linear holding costs factor  $h$ . Hence, we are confident that the analyses and conclusions obtained from cost model 2 will give insight for cost model 3. Further, we have shown that for cost model 1 the pattern of  $f_v^1(k_v)$ s is increasing convex/concave in  $k_v$  and is asymptotically constant when the expected goodwill cost  $a^1(x)$  in (2.4) approaches  $d^1$ .

Vendor	Profile	$B^1(c=0)$	$OA^1(c=0)$	$B^1(c=2)$	$OA^1(c=2)$
1	1	52	25	51	25
1	2	61	31	60	31
1	3	71	38	71	38
1	4	84	47	83	47
1	5	98	57	97	57
1	6	100 <sup>+</sup>	68	100 <sup>+</sup>	68
4	1	52	25	51	25
4	2	44	19	43	19
4	3	35	14	34	14
4	4	27	9	26	9
4	5	19	4	18	4
4	6	12	0	11	0

Table 3.1: Convexity boundary points for cost model 1 when  $c = 0$ ,  $c = 2$  and  $d^1 = 10$ 

Table 3.1 and Table 3.2 present the convexity boundary points  $B^l$  and their correspondent optimal allocations  $OA^l$  with great relevance to the first two subsidiary goals. We consider  $K = 100$  for both tables and choose vendor 1 and vendor 4 for illustration. The service rate for vendor 1 is maximal and increasing as the profile number goes from 1 to 6 while the service rate for vendor 4 is minimal and decreasing in the profile number. The results in Table 3.1 are computed based on cost model 1 with goodwill cost parameter  $d^1 = 10$  while Table 3.2 is for cost model 2 with goodwill cost parameter  $d^2 = 1000$ . These choices of cost parameters  $d^1$  and  $d^2$  aim to provide a broadly similar level in terms of the resulting optimal costs rate obtained, as indicated from Table 3.3 to Table 3.8. We further take the repair cost parameters  $c = 0$  and  $c = 2$  for both tables.

Vendor	Profile	$B^2(c=0)$	$OA^2(c=0)$	$B^2(c=2)$	$OA^2(c=2)$
1	1	100 <sup>+</sup>	25	100 <sup>+</sup>	25
1	2	100 <sup>+</sup>	32	100 <sup>+</sup>	32
1	3	100 <sup>+</sup>	40	100 <sup>+</sup>	40
1	4	100 <sup>+</sup>	50	100 <sup>+</sup>	50
1	5	100 <sup>+</sup>	61	100 <sup>+</sup>	61
1	6	100 <sup>+</sup>	72	100 <sup>+</sup>	72
4	1	100 <sup>+</sup>	25	100 <sup>+</sup>	25
4	2	100 <sup>+</sup>	19	100 <sup>+</sup>	19
4	3	100 <sup>+</sup>	12	100 <sup>+</sup>	12
4	4	100 <sup>+</sup>	6	100 <sup>+</sup>	6
4	5	100 <sup>+</sup>	1	100 <sup>+</sup>	1
4	6	100 <sup>+</sup>	0	100 <sup>+</sup>	0

Table 3.2: Convexity boundary points for cost model 2 when  $c = 0$ ,  $c = 2$  and  $d^2 = 1000$ 

The results from Table 3.1 show that there are strong positive correlations between the boundary points, the optimal allocations and the vendor's service rates. For example, the lowest convexity boundary point is 11 in Table 3.1. With a very low service capacity the vendor is likely to receive the smallest allocation of workload, which turns out to be 0 by using both DP and GA. As we expect from Proposition 3.4, the boundary points for  $c = 0$  are higher than those for  $c = 2$  and the optimal allocations are all below the boundary points. We record the boundary points beyond the value of the population size 100 as 100<sup>+</sup>, and so the evidence from Table 3.2 is not as clear as that from Table 3.1. However, the stronger convexity of cost model 2 (i.e. all boundary points are beyond 100 in all cases studied in Table 3.2) will make the optimality of GA more likely for these problems. Another observation from Tables 3.1 and 3.2 is that the optimal allocations are less sensitive to  $c$  than are the boundary points. This might imply that the optimal allocation to the problem sets is likely to be insensitive to the choice of parameters. The latter point will be verified further by the evidence obtained from subsequent tables.

Table 3.3 contains greedy allocations under cost model 1 with six profiles of service rate distributions among the four vendors, as given in (3.22), for  $K = 100$  and  $K = 500$ . We set  $c_v = 1, 1 \leq v \leq 4$  and  $d^1 = 10$ . The associated goodwill cost rates are headed  $g^1$  while the repair cost rates are under the heading  $r^1$ . Note that the goodwill cost rates are decreasing in the degree of dissimilarity among the four vendors' service capacities. On the contrary, the repair cost rate is increasing in the profile number, but less significantly. Overall, the total cost rate is decreasing in the profile number. Under the above parameter settings, we find that the dissimilarity among service vendors might benefit manufacturers in terms of cutting costs under the static allocation model. Table 3.4 presents an equivalent set of results for cases in which, for each row, repair costs  $c_v, 1 \leq v \leq 4$  have been sampled independently from a uniform  $U(0.90, 1.10)$  distribution. Note that the instances studied range from those in which goodwill costs dominate ( $K = 100$ ) to those in which repair costs do ( $K = 500$ ). The allocations for  $K = 500$  in Table 3.4 are somewhat sensitive to the variations in repair costs among the four vendors. However, comparing the two tables the greedy allocations remain identical throughout the six profiles when the goodwill costs dominate for  $K = 100$ .

Profile	$K$	Greedy Allocation				$r^1$	$g^1$
1	100	25	25	25	25	2.857	2.857
2	100	31	27	23	19	1.327	2.804
3	100	38	28	20	14	1.103	2.693
4	100	47	29	18	9	1.001	2.622
5	100	56	33	12	4	1.103	2.167
6	100	65	37	7	0	1.624	1.701
1	500	135	107	119	128	3.0471	0.0916
2	500	170	112	116	97	3.7754	0.0909
3	500	195	126	130	69	3.9490	0.0796
4	500	216	137	130	47	3.7531	0.0623
5	500	239	146	127	28	3.9530	0.0473
6	500	253	157	119	10	3.0075	0.0116

Table 3.3. Greedy allocations and associated cost rates for cost model 1 with six profiles of service rate distributions among the four vendors, as given in (3.22), for  $K = 100$  and  $K = 500$ . We set  $c_v = 1, 1 \leq v \leq 4$  and  $d^1 = 10$ . The associated goodwill cost rates are headed  $g^1$  while the repair cost rates are under the heading  $r^1$ .

Profile	$K$	Greedy Allocation				$r^1(\times 10^{-2})$	$g^1(\times 10^{-2})$
1	100	25	25	25	25	1.1607	2.8358
2	100	31	27	23	19	1.1610	2.8034
3	100	38	28	20	14	1.1619	2.6913
4	100	47	28	16	9	1.1636	2.4822
5	100	57	27	12	4	1.1664	2.1601
6	100	68	25	7	0	1.1704	1.7207
<hr/>							
1	500	125	125	125	125	5.9569	0.0714
2	500	164	136	111	89	5.9571	0.0698
3	500	210	143	93	54	5.9578	0.0639
4	500	263	146	69	22	5.9591	0.0519
5	500	322	139	39	0	5.9618	0.0302
6	500	383	113	4	0	5.9659	0.0113

Table 3.3: Greedy allocation and associated cost rates for cost model 1 when  $c_v = 1$ ,  $1 \leq v \leq 4$ , and  $d^1 = 10$ .

Profile	$K$	Greedy Allocation				$r^1(\times 10^{-2})$	$g^1(\times 10^{-2})$
1	100	25	25	25	25	1.1557	2.8358
2	100	31	27	23	19	1.1327	2.8034
3	100	38	28	20	14	1.1783	2.6913
4	100	47	28	16	9	1.1922	2.4822
5	100	56	28	12	4	1.1782	2.1617
6	100	68	25	7	0	1.1828	1.7207
<hr/>							
1	500	135	107	129	129	5.9132	0.0816
2	500	178	113	116	93	5.7754	0.0865
3	500	195	136	100	69	6.0440	0.0766
4	500	266	157	50	27	6.0831	0.0622
5	500	269	166	57	8	5.9880	0.0775
6	500	383	117	0	0	6.0075	0.0116

Table 3.4: Greedy allocation and associated cost rates for cost model 1 when repair costs  $c_v$ ,  $1 \leq v \leq 4$  are drawn independently from a  $U(0.90, 1.10)$  distribution, and  $d^1 = 10$ .

Tables 3.5 and 3.6 include the results obtained from the application of the greedy algorithm to cost model 2 for the equivalent sets of problems to those that have been covered by Tables 3.3 and 3.4 respectively. We now take the goodwill cost parameter to be  $d^2 = 1000$ . Clearly, most of the features of Tables 3.3 and 3.4 are reflected in Tables 3.5 and 3.6. We confirm that the greedy allocations are insensitive to the modest variations of repair costs  $c_v^2, 1 \leq v \leq 4$  among vendors given that the goodwill cost rate dominates under cost model 2.

We are certainly keen to find out the sensitivity of greedy allocations to the setting of the goodwill costs parameter  $d^I, I = 1, 2$  due to the practical difficulty of assigning values to such nonphysical parameters. For this purpose, we increase  $d^2$  from 1000 to 10000 in the problem set of Table 3.5. See Table 3.7. Again, there are no changes of greedy allocations for  $K = 100$  but there are some modest changes when  $K = 500$ .

We proceed to carry out more experiments to validate the claim that the greedy allocations are insensitive to the choice of  $d^I, I = 1, 2$  if the goodwill cost rate predominates. From the relative size of the repair cost rate and the goodwill cost rate for  $K = 500$ , we infer from the tables that the current setting of overall service capacities of four vendors is more than is needed. Therefore, we reduce the overall service capacities by half but still use (3.22) to get six profiles of service rates. The results obtained subsequently for  $K = 500$  under cost model 1 are presented in Table 3.8. The greedy allocations when  $d^1 = 10$  are virtually identical to those when  $d^1 = 100$ . Consequently, the corresponding goodwill cost rates are increased by a factor of 10 with the repair cost rates unchanged. Similar outcomes are obtained under cost model 2. The experiments conducted show clearly that the optimal solutions are insensitive to the choice of key parameters, namely the repair costs  $c_v, 1 \leq v \leq 4$  and the goodwill costs parameter  $d^I, I = 1, 2$  in most cases studied. Hence, we believe that the outstanding performance of the greedy allocations (which are optimal) is sustainable beyond the problems set constructed. This is an encouraging finding.



Profile	$K$	Greedy Allocation				$r^2(\times 10^{-2})$	$g^2(\times 10^{-2})$
1	100	25	25	25	25	1.1607	7.7542
2	100	32	27	22	19	1.1610	7.6458
3	100	40	29	19	12	1.1618	7.2602
4	100	50	29	15	6	1.1635	6.5058
5	100	61	28	10	1	1.1663	5.2891
6	100	72	24	4	0	1.1703	3.7177
<hr/>							
1	500	125	125	125	125	5.9569	0.0414
2	500	165	136	111	88	5.9570	0.0405
3	500	211	144	92	53	5.9577	0.0369
4	500	266	146	68	20	5.9589	0.0297
5	500	325	138	37	0	5.9615	0.0160
6	500	389	111	0	0	5.9652	0.0056

Table 3.5: Greedy allocation and associated cost rates for cost model 2 when  $c_v = 1$ ,  $1 \leq v \leq 4$ , and  $d^2 = 1000$ .

Profile	$K$	Greedy Allocation				$r^2(\times 10^{-2})$	$g^2(\times 10^{-2})$
1	100	25	25	25	25	1.1557	7.7542
2	100	32	27	22	19	1.1321	7.6458
3	100	40	29	19	12	1.1821	7.2602
4	100	50	29	15	6	1.1918	6.5058
5	100	61	28	10	1	1.1848	5.2891
6	100	72	24	4	0	1.1809	3.7177
<hr/>							
1	500	138	98	132	132	5.9045	0.0564
2	500	182	102	120	96	5.7621	0.0638
3	500	191	135	103	71	6.0382	0.0541
4	500	273	162	37	28	5.9923	0.0464
5	500	258	172	61	9	5.9970	0.0708
6	500	384	116	0	0	6.0074	0.0053

Table 3.6: Greedy allocation and associated cost rates for cost model 2 when repair costs  $c_v$ ,  $1 \leq v \leq 4$  are drawn independently from a  $U(0.90, 1.10)$  distribution, and  $d^2 = 1000$ .

Profile	$K$	Greedy Allocation				$r^2(\times 10^{-2})$	$g^2(\times 10^{-3})$
1	100	25	25	25	25	1.1607	7.7542
2	100	32	27	22	19	1.1610	7.6458
3	100	40	29	19	12	1.1618	7.2602
4	100	50	29	15	6	1.1635	6.5058
5	100	61	28	10	1	1.1663	5.2891
6	100	72	24	4	0	1.1703	3.7177
<hr/>							
1	500	125	125	125	125	5.9569	0.0414
2	500	164	136	111	89	5.9571	0.0404
3	500	211	144	92	53	5.9577	0.0369
4	500	264	146	69	21	5.9591	0.0296
5	500	322	139	39	0	5.9618	0.0159
6	500	381	115	4	0	5.9660	0.0052

Table 3.7: Greedy allocation and associated cost rates for cost model 2 when  $c_v = 1, 1 \leq v \leq 4$ , and  $d^2 = 10,000$ .

Profile	$d^1$	Greedy Allocation				$r^1(\times 10^{-2})$	$g^1(\times 10^{-3})$
1	10	125	125	125	125	5.6678	2.7718
2	10	143	129	118	110	5.6652	2.7518
3	10	165	134	111	90	5.6647	2.6899
4	10	193	137	100	70	5.6685	2.5919
5	10	225	138	86	51	5.6761	2.4553
6	10	262	135	69	34	5.6880	2.2813
<hr/>							
1	100	125	125	125	125	5.6678	27.7181
2	100	143	129	118	110	5.6652	27.5178
3	100	165	134	111	90	5.6647	26.8986
4	100	193	137	100	70	5.6685	25.9193
5	100	225	138	86	51	5.6761	24.5530
6	100	262	135	69	34	5.6880	22.8126

Table 3.8: Greedy allocation and associated cost rates for cost model 1 when  $c_v = 1, 1 \leq v \leq 4$ ,  $K = 500$ , and the total service rate is reduced by half.

Most impressive, however, is the fact that of more than 500 problem instances we have studied, in only one case was the greedy allocation **not** optimal. This instance was under cost model 1 with  $K = 500$  and  $d = 10$  for which the repair costs were drawn from a  $U(0.50, 1.50)$  distribution. The greedy allocation for this case was (302, 175, 21, 2) while the optimal allocation was (303, 175, 20, 2). The overall cost rates for these allocations differed by just 0.02%.

Opp et al. (2003) mentioned five other heuristics to guide the allocation. We quote as follows:

- H1** : Equal allocation to all vendors. This may be appropriate if the manufacturer has little information about each vendor.
- H2** : Allocation is proportional to the reciprocal of the unit repair cost. This favours low-cost vendors. This allocation strategy might be used when the manufacturer has cost information about the vendors, but does not have good estimates of vendor's service rates.
- H3** : Allocation is proportional to the service rate of each vendor divided by their unit repair cost. This favors vendors with low repair cost and high system service rate. This allocation policy might be used when the manufacturer has good estimates of both cost and service rate.
- H4** : All items are allocated to the vendor with the smallest unit repair cost. This is an "all-or-nothing" version of H2.
- H5** : All items are allocated to the vendors with the largest value of the system service rate divided by the unit repair cost. This heuristic is an "all-or-nothing" version of H3.

We conducted some trials to compare greedy solutions with these simple heuristics to show how much the manufacturer might be losing by employing a suboptimal approach to workload allocation. The comparison is based on the cases from Table 3.6 for cost model 2. Table 3.9 presents the difference of overall cost rate in percentage between

GA and other heuristics respectively in three columns. We did not include H4 and H5 as these two "all-or-nothing" versions are even worse than the other three under cost model 2.

Profile	$K$	DIF(H1)%	DIF(H2)%	DIF(H3)%
1	100	0.0000	0.2360	0.2360
2	100	12.9876	12.0617	2.1453
3	100	67.2862	91.1349	14.0093
4	100	236.3158	230.4512	26.7825
5	100	802.2178	806.6620	66.9917
6	100	2478.7482	2233.2419	108.2621

Table 3.9: the difference between GA and H1, H2 and H3 in terms of overall cost rates for cost model 2 when repair costs  $c_v, 1 \leq v \leq 4$  are drawn independently from a  $U(0.90, 1.10)$  distribution, and  $d^2 = 1000$ .

### 3.6 Conclusions to the Chapter

We have proposed static allocation models for the optimal distribution of warranty repair work among a collection of service vendors. In Chapter 2 we have argued the importance of approaches to the modelling of goodwill costs which take explicit account of the delays experienced by customers. While the cost rates which arise from these approaches are such that no claim to the optimality of simple heuristics can be made in any generality, nevertheless a range of evidence (both theoretical and numerical) is adduced in support of the strong performance of greedy approaches to work allocation. Through numerical investigations, we also find that greedy/optimal allocations are not very sensitive to the choice of key parameters in our proposed models. Hence, we are confident of the greedy algorithm's effectiveness beyond our problem sets. Further, we show the manufacturer could save substantially by using greedy heuristics instead of some *ad hoc* ones.

the same allocation strategy as in Chapter 3 in terms of decision time is applied (i.e. items are allocated to alternative vendors upon purchase). An arrival stream of new equipment purchases are assumed to follow a compound Poisson process. Under this assumption, bulk orders placed by business purchasers are included in our studies. All the items within a single order are assumed to be allocated to one vendor. We discuss how this assumption may be relaxed in Section 4.3 after main analysis. All the possible repair work within the order will be carried out by this assigned vendor until those items leave the system when their warranty period has expired. Allocation decisions at every decision epoch are based upon the number of items currently allocated to each vendor along with the amount of their remaining time under warranty. The unexpired warranty period of each item could be obtained by simply logging its purchase date. Plainly, the difference between the purchase date and the current decision epoch is the elapsed time of the warranty period for each item. Above all, the allocation procedure designed for the model in this chapter will *not* consider the repair queue length at the vendor which is the locus at which costs are actually incurred. We believe that the additional information (e.g., the queue length) would make our decisions better though it involves a substantial administrative overhead for continuous observation of all vendors. Decisions based on the low-dimensional indices computed for the system where the repair queue length at each vendor is known will be our focus in the next chapter.

In outline, the rest of the chapter has the following structure. Section 4.2 will give a detailed description of a stochastic and dynamic optimization problem whose system state is only partially observable. Our non-standard setup makes the direct use of dynamic programming to analyze the model unrealistic. Alternatively, in Section 4.3 we develop an effective heuristic by deploying the idea of dynamic programming policy improvement in two steps along with its interpretation and relevant discussions. Followed by a brief reflection on our earlier static allocation model in Chapter 3, we derive some other simple heuristics, namely the dynamic greedy heuristic and tracking heuristic plus the empirical smallest workload heuristic. Then, the implementation of this DP policy improvement heuristic within a simulation framework is illustrated in Section 4.4. We report and examine the results obtained by alternative heuristics and

explore the underlying implications. This chapter ends with some concluding remarks.

## 4.2 The System

Purchases are made according to a compound Poisson process  $(\eta, F)$ , where  $\eta \in \mathbb{R}^+$  is the rate at which orders occur and  $F$  is the distribution function (d.f.) to describe each order size. We denote  $X \sim F$  as a generic order size, which is a positive integer valued random variable with mean  $\beta$  and finite second moment  $\beta_2$ . Upon receipt of an order, a decision is made regarding which one of  $V$  vendors should be chosen to take care of the items in that order during the ensuing warranty period ( $\Omega$  years). Once an allocation decision is made for that order, the data including the vendor chosen, the order size and the purchase date are all logged and are available to inform future decisions. We believe that a good allocation policy will certainly take account of the amount of the workload already committed to each vendor in relation to their service capacities. The workload in this context considers not only the number of items already at each vendor but also the durations of their unexpired warranties. In summary, the following information will be available to form the basis for allocation decisions every time an order is placed.

- (i)  $x$ : the size of the incoming order;
- (ii)  $N_v$ : the number of items currently at vendor  $v$  along with the durations  $(t_n^v, 1 \leq n \leq N_v)$  of their unexpired warranties,  $1 \leq v \leq V$ ;
- (iii)  $M_v$ : the number of orders currently at vendor  $v$ . Given that items within the same order have identical unexpired warranties, the information in (ii) for vendor  $v$  may be alternatively presented as

$$(\mathbf{x}^v, \mathbf{t}^v) \equiv \{(x_1^v, t_1^v), (x_2^v, t_2^v), \dots, (x_{M_v}^v, t_{M_v}^v)\} \quad (4.1)$$

where the  $x_m^v, 1 \leq m \leq M_v$ , are order sizes and the  $t_m^v, 1 \leq m \leq M_v$ , are the corresponding durations of unexpired warranties, numbered such that

$$0 \leq t_1^v < t_2^v < \dots < t_{M_v-1}^v < t_{M_v}^v \leq \Omega, \quad 1 \leq v \leq V.$$

Plainly we have

$$N_v = \sum_{m=1}^{M_v} x_m^v, \quad 1 \leq v \leq V, \quad \text{with} \quad N = \sum_{v=1}^V N_v,$$

the total number of items under warranty (i.e., at any vendor). The following quantities

$$\{N_v, (x^v, t^v), 1 \leq v \leq V\} \tag{4.2}$$

which include all the information needed (and which we take as the system state) evolve through time driven by the stochastic dynamics of the order process  $(\eta, F)$  and the chosen policy for making allocations. Standard properties of the compound Poisson process indicate that, once the above system has been in operation for (at least) time  $\Omega$ , the mean and variance of the total number of items under warranty are given by

$$E(N) = \eta\Omega\beta \quad \text{and} \quad \text{var}(N) = \eta\Omega\beta_2. \tag{4.3}$$

We proceed to describe the dynamics of the breakdown/repair process at each vendor where the costs for repairing broken items under warranty actually occur. Though this breakdown/repair process for items at each vendor is assumed *not* observable to inform the above decisions concerning allocations a succinct description of the dynamics of this process is necessary for a better understanding of the overall system and the actual costs calculations in the simulation study described later in this chapter. The breakdown/repair process itself is assumed to be Markovian in character. If at some time  $t$ ,  $N_v(t)$  items are under the care of vendor  $v$  we write  $D_v(t)$  for the number of those items which are awaiting or undergoing repair (i.e. down) and  $U_v(t) = N_v(t) - D_v(t)$  for those (up) items which are functioning satisfactorily. Let  $\delta t > 0$  and consider system evolution during time interval  $[t, t + \delta t)$ . In the absence of any new items arriving at vendor  $v$  or of any departures due to warranty expiration then during  $[t, t + \delta t)$  the

following describe the dynamics of the  $\{D_v(t), U_v(t)\}$  process:

$$\begin{aligned}
 P\{\text{single breakdown requiring repair during } [t, t + \delta t)\} &= \lambda U_v(t)\delta t + o(\delta t); \\
 P\{\text{no breakdown requiring repair during } [t, t + \delta t)\} &= 1 - \lambda U_v(t)\delta t + o(\delta t); \\
 P\{\text{single repair completed during } [t, t + \delta t)\} &= \mu_v \delta t + o(\delta t); \\
 P\{\text{no repair completed during } [t, t + \delta t)\} &= 1 - \mu_v \delta t + o(\delta t),
 \end{aligned}
 \tag{4.4}$$

where  $o(\delta t)$  denotes any quantity satisfying  $o(\delta t)/\delta t \rightarrow 0, \delta t \rightarrow 0$ . Please refer to Section 2.3 for a detailed discussion of the model assumptions for the above breakdown/repair process.

As the above process evolves through time, the costs associated with all repairs are aggregated and are averaged over time to obtain an average cost rate. Chapter 2 contains several different approaches to formulating the expected goodwill costs which predominate. In the simulation study conducted, a response time  $r$  is obtainable. Hence the actual cost  $c_v(r)$  incurred for every repair under the models we consider where  $I$  is the cost model index and  $I = 1, 2, 3$  will be calculated respectively by

$$c_v^1(r) = c_v + d^1 I(r > \tau), \tag{4.5}$$

$$c_v^2(r) = c_v + d^2 (r - \tau)^+, \tag{4.6}$$

$$c_v^3(r) = c_v + h(r - \tau)^- + d^3 (r - \tau)^+, \tag{4.7}$$

where  $I(\cdot)$  is the indicator function. By using the policies that are developed in the next sections we aim to minimize the resulting average cost rate over an infinite horizon or come close to doing so. The policies consist of sets of decisions that determine to which vendor each incoming order is sent on the basis of the partially observed system state described in (4.2). Any attempt to make conventional use of dynamic programming in seeking optimal allocation policies for the problem concerned encounters the following particular difficulties for the study of this non-standard stochastic problem:

- (a) the breakdown/repair process for each vendor which is the direct generator of the costs incurred is assumed not to be observable. Costs for vendor  $v$  may only be *inferred* from its current state  $\{N_v, (x^v, t^v), 1 \leq v \leq V\}$ ;

- (b) the observed system state is itself complex, being both continuous and of high (and variable) dimension.

Despite these formidable difficulties, we shall develop an effective allocation heuristic in the upcoming sections. It will then be subject to numerical evaluation against other simpler heuristics for the problem concerned. Note that we drop the cost model index  $I$  in the rest of this chapter for simplicity of notation.

### 4.3 DP Policy Improvement Approach by Using Approximate Costs Model

In this section, we employ a DP policy improvement approach to develop a close-to-optimal allocation heuristic for our problem. The basic idea is to derive a vendor-specific calibrating index which is a function of the system's current state information given in (4.2). The proposed policy will route an incoming order to the vendor with minimal index. The origin of the policy improvement approach may be traced back to 1960s. At that time Howard pioneered a policy-iteration algorithm<sup>1</sup> for solving probabilistic sequential decision processes over an infinite planning horizon by using basic principles from Markov Chain theory and dynamic programming (which was popularized by Bellman (1957)). Following their pioneering work, the theory of Markov decision processes has been developed considerably over the next two decades. Additionally, various application domains like queueing control, scheduling and maintenance have been explored subsequently and utilize the policy improvement approach. For examples, see Glazebrook et al. (2004), Tijms (1994) and Krishnan (1987). Given the complexity of our problem, executing several iterations of policy improvement would be numerically intractable. We then follow Krishnan (1987) in developing a "one policy-iteration step" approach by applying a single DP policy improvement step to an initial optimal static (state independent) allocation policy.

---

<sup>1</sup>a sequence of monotonically improving policies and objective values until it achieves optimality.

Before proceeding to the design of an initial static policy, we need to confront the difficulties highlighted in Section 4.2 and labeled (a) to ease the development of this allocation heuristic. We adopt an approximate costs model for this purpose to give an observable cost rate approximation to the unobservable costs incurred at the vendors. Explicitly, the approximate costs model shares the the system description up to (4.3), where orders of known size arrive according to a compound Poisson process  $(\eta, F)$  and must be dispatched to one of  $V$  vendors based on the state information given in (4.2). The actual total unobservable cost rate at any vendor  $v, 1 \leq v \leq V$  is now approximated by  $\sum_{v=1}^V f_v(N_v)$  when the system is in state  $\{N_v, (x^v, t^v), 1 \leq v \leq V\}$ , where the  $f_v$ s are appropriate vendor-specific cost rates as given in (3.5). Hence, we take the vendor-specific average cost rate  $f_v(N_v)$ , computed on the basis that vendor  $v$  has a *fixed* warranty population of  $N_v$  as an approximation to the instantaneous cost rate at vendor  $v$  wherever it has responsibility for  $N_v$  items under warranty in the following development of our DP policy improvement approach. Throughout the thesis, we confine ourselves to the optimality criterion of the average cost per unit time over an infinite horizon. It is believed that this criterion is more appropriate than other alternatives for most application of Markov decision processes according to Tijms (1994). Hence we expect that in the long run these observable cost rates given by the approximate costs model should be in good agreement with those unobservable ones actually incurred in the full system described in Section 4.2.

This powerful *approximate* DP approach is developed in two stages by the deployment of an assumption that all decisions beyond the current one are made according to a strongly performing static allocation policy. The first stage of the approach concerns the development of such a static policy which indicates the proportion of overall workload which should be allocated to each vendor. An effective dynamic heuristic (which utilizes the system information in (4.2)) is then developed at the second stage by the application of a single DP policy improvement step.

### Stage 1. Initial static policy

At the first stage in developing our DP policy improvement approach, we design an

initial static policy  $\mathbf{p} = (p_1, p_2, \dots, p_V)$  for the problem which allocates each incoming order to vendor  $v$  with probability  $p_v$ , where

$$p_v \geq 0, 1 \leq v \leq V, \text{ and } \sum_{v=1}^V p_v = 1.$$

The decisions under this initial policy are state independent. The stochastic system under policy  $\mathbf{p}$  is such that the steady state is achieved once time  $\Omega$  has passed from the beginning of the process. Under policy  $\mathbf{p}$ , the number of orders at vendor  $v$  in steady state whose warranties have yet to expire is  $\bar{M}_v \sim \text{Poisson}(\eta\Omega p_v), 1 \leq v \leq V$ . Moreover, the order size  $X_i$ s are independent random variables, we write

$$S_{M_v} = X_1 + X_2 + \dots + X_{M_v}, \quad M_v \in \mathbb{N}, \quad 1 \leq v \leq V$$

with  $X_i \sim F$ , the generic order size distribution,  $1 \leq i \leq M_v$ . We further write

$$\gamma_v(M_v) = E\{f_v(S_{M_v})\}, \quad 1 \leq v \leq V, \quad M_v \in \mathbb{N} \tag{4.8}$$

where the function  $\gamma_v(M_v)$  is the cost rate for vendor  $v$  when responsible for items from  $M_v$  orders. The overall expected cost rate in steady state when applying policy  $\mathbf{p}$  to the problem is as follows:

$$\sum_{v=1}^V E_{\eta\Omega p_v}\{\gamma_v(\bar{M}_v)\} \equiv G(\mathbf{p}) \tag{4.9}$$

where the subscript  $\eta\Omega p_v$  in (4.9) denotes taking of an expectation of  $\bar{M}_v$  with respect to the Poisson distribution with this mean. An intermediate goal of analysis is the search for an optimal static policy  $\mathbf{p}^*$  such that

$$G(\mathbf{p}^*) = \min_{\mathbf{p}} G(\mathbf{p}) \tag{4.10}$$

Following the discussion of the  $f_v, 1 \leq v \leq V$ , in Section 3.2 it is important to point out (and straightforward to show) that if the  $f_v$  are all increasing convex then so are the  $\gamma_v, 1 \leq v \leq V$ . In this case, the optimization problem in (4.10) is convex and separable and simple efficient algorithms exist for its solution. Specifically, if there exists an optimizing  $\mathbf{p}^*$  which is an interior point ( $0 < p_v^* < 1, 1 \leq v \leq V$ ) then it will be true that the quantities

$$E_{\eta\Omega p_v^*}\{\gamma_v(\bar{M}_v + 1) - \gamma_v(\bar{M}_v)\}$$

are equal for all vendors. Plainly,  $G(p^*)$  is an accessible upper bound on the cost rate incurred under an optimal dynamic policy when adopting the above approach to the approximation of costs.

**Stage 2. DP policy improvement step**

The basic idea for improving the optimal static policy  $p^*$  obtained at the first stage is to seek an optimal initial allocation at time zero when in the system state  $N \equiv \{N_v, (x^v, t^v), 1 \leq v \leq V\}$  on the basis that all subsequent ones are made according to  $p^*$ . The optimal initial allocation only has an effect on the expected (future) costs for the duration  $[0, \Omega]$ . This is because any decision made at time zero does not impact the system state after the warranty duration  $\Omega$  has elapsed. Specifically, we shall use  $C(v, p^*, T|N, x)$  to denote the expected cost incurred by using the approximate costs model up to time  $T, T \geq \Omega$  when the initial allocation of order size  $x$  is made to vendor  $v$  and all subsequent allocations are according to  $p^*$ . Our PI heuristic will, in any state  $N$ , choose to allocate an order of size  $x$  to vendor  $v^*$  satisfying

$$C\{v^*, p^*, T|N, x\} = \min_{1 \leq v \leq V} C(v, p^*, T|N, x), \quad T \geq \Omega. \tag{4.11}$$

A more explicit form of  $C(v, p^*, T|N, x)$  will be derived in the following analysis.

Should the incoming order at time zero be allocated to vendor  $v$  then, from (4.1), the vendor  $v$  state undergoes a transition

$$(x^v, t^v) \rightarrow (x^v, t^v)^x,$$

where

$$(x^v, t^v)^x \equiv \{(x_1^v, t_1^v), (x_2^v, t_2^v), \dots, (x_{M_v}^v, t_{M_v}^v), (x, \Omega)\}. \tag{4.12}$$

After the first allocation, subsequent evolution under policy  $p^*$  is independent for distinct vendors. Further, at any time  $T > \Omega$  the time 0 allocation of the size  $x$  order does not impact the system state then. Hence the total expected cost under  $p^*$  up to time  $T$  when the initial allocation at time zero is made to vendor  $v$  is divided into three components. In a natural notation, we have

$$C(v, p^*, T|N, x) = C_v\{p_v^*, \Omega|(x^v, t^v)^x\} + \sum_{w \neq v} C_w\{p_w^*, \Omega|(x^w, t^w)\}$$

$$+ (T - \Omega) \sum_{w=1}^V E_{\eta\Omega p_w^*} \{\gamma_w(\overline{M}_w)\}, 1 \leq v \leq V, T \geq \Omega. \quad (4.13)$$

Expression (4.13) partitions the expected system cost during  $[0, T]$  between costs incurred during  $[0, \Omega]$ , and those incurred during  $[\Omega, T]$  (the final term in (4.13)). Costs incurred during  $[0, \Omega]$  are partitioned between those incurred at vendor  $v$  under  $p_v^*$  (the first term in the right-hand side of (4.13)) and those at other vendors (the second term in the right-hand side of (4.13)). Now rewrite (4.13) as

$$C(v, p^*, T|N, x) = C_v\{p_v^*, \Omega|(x^v, t^v)^x\} - C_v\{p_v^*, \Omega|(x^v, t^v)\} + \sum_{w=1}^V C_w\{p_w^*, \Omega|(x^w, t^w)\} + (T - \Omega) \sum_{w=1}^V E_{\eta\Omega p_w^*} \{\gamma_w(\overline{M}_w)\}, 1 \leq v \leq V, T \geq \Omega. \quad (4.14)$$

It is straightforward to infer the following result from (4.14).

**Theorem 4.1 (Characterization of PI heuristic)** *The allocation policy obtained when a single DP policy improvement step is applied to optimal static policy  $p^*$  operates as follows: If an order of size  $x$  arrives when the system state is*

$$N \equiv \{N_v, (x^v, t^v), 1 \leq v \leq V\}$$

*it should be allocated to any vendor  $v$  which has a minimal value of the calibrating index*

$$I_{vx}(x^v, t^v) \equiv C_v\{p_v^*, \Omega|(x^v, t^v)^x\} - C_v\{p_v^*, \Omega|(x^v, t^v)\} \quad (4.15)$$

We proceed to develop an appropriate formula to compute the calibrating indices in (4.15), namely the difference of the expected costs at vendor  $v$  incurred during time  $[0, \Omega]$  with and without the first allocation made to vendor  $v$ . To simplify the notation, the following computation for indices will be based on an individual vendor, so we drop the vendor identifier  $v$  and write

$$I_x(x, t) = C\{p^*, \Omega|(x, t)^x\} - C\{p^*, \Omega|(x, t)\} \quad (4.16)$$

where

$$(x, t) \equiv \{(x_1, t_1), (x_2, t_2), \dots, (x_M, t_M)\}$$

By using the approximate costs model, we also expand the notation in (4.8) such that for any  $y \in \mathbb{N}, M \in \mathbb{N}$  we have that

$$\gamma(y, M) = E\{f(y + S_M)\} \tag{4.17}$$

**Lemma 4.2 (Vendor-specific calibrating indices)** *The vendor-specific indices which determine the PI allocation heuristic are given by*

$$I_x(\mathbf{x}, \mathbf{t}) = \sum_{m=0}^M \int_{t_m}^{t_{m+1}} E_{\eta t p^*} \left\{ \gamma\left(x + \sum_{r=m+1}^M x_r, \overline{M}^t\right) - \gamma\left(\sum_{r=m+1}^M x_r, \overline{M}^t\right) \right\} dt \tag{4.18}$$

for all values of the arguments concerned, where  $t_0 = 0, t_{M+1} = \Omega$ .

*Proof.* Consider the computation of the first term on the r.h.s. of (4.16) for the time being. The basic idea is to accumulate the future costs incurred among  $M + 1$  intervals from time zero till the end of the warranty period of  $\Omega$  to obtain

$$C\{p^*, \Omega | (\mathbf{x}, \mathbf{t})^x\}$$

We assume an arbitrary time point  $t$  falling within the interval of  $[t_m, t_{m+1}]$ ,  $0 \leq m \leq M$ .  $M$  is the number of live orders that have been previously allocated to the vendor concerned when the order of size  $x$  arrives. It is obvious that by time  $t$  the items with the logged information

$$(x_1, t_1), \dots, (x_m, t_m)$$

will have left the system with expired warranties. Of those items which had been under warranty at time 0 a total of  $\sum_{r=m+1}^M x_r$  will remain under warranty at the vendor. These items will then be joined by the order of size  $x$  which is allocated to the vendor at time 0 and a Poisson distributed number of orders, denoted  $\overline{M}^t$  with mean  $\eta t p^*$  which will be allocated to the vendor during  $[0, t]$ .  $S_{\overline{M}^t}$  denotes the associated number of items from  $\overline{M}^t$  orders where  $S_{\overline{M}^t}$  is Poisson distributed with mean  $\eta t p^* \beta$  (using (4.3)). Hence in the expression of (4.18) under the approximate costs model the expected cost rate at  $t$  is given by

$$E_{\eta t p^*} \left\{ f\left(x + \sum_{r=m+1}^M x_r + S_{\overline{M}^t}\right) \right\} \equiv E_{\eta t p^*} \left\{ \gamma\left(x + \sum_{r=m+1}^M x_r, \overline{M}^t\right) \right\}. \tag{4.19}$$

We now consider the computation of the second term on the r.h.s. of (4.16), namely

$$C\{p^*, \Omega | (x, t)\},$$

Equivalent expected costs to those above can be obtained by inserting  $x = 0$  into the expressions in (4.19). The expression in (4.18) is then derived from (4.16) by integrating the appropriate cost rates over the  $M + 1$  time intervals. This concludes the proof.  $\square$

To give the reader a feel of the likely form and the interpretation of the indices in (4.16), we consider a simple example in which all orders are of size 1 ( $x_m^v = 1$  for all  $m, v$ ) and all vendor-specific cost rates are quadratic such that

$$f_v(k) = a_v + b_v k + c_v k^2, \quad 1 \leq v \leq V, \quad k \in \mathbb{N}, \quad (4.20)$$

where in (4.20) and what follows we restore the vendor identifier  $v$ . It is not difficult to show that the vendor-specific calibrating indices derived above are then given by

$$I_v(t^v) = c_v \left( 2 \sum_1^{M_v} t_m^v - \eta p_v^* \Omega^2 \right) + const, \quad 1 \leq v \leq V, \quad (4.21)$$

where in (4.21) *const* denotes a constant common to all indices for all vendors. The expression in (4.21) has a simple interpretation:  $\sum_1^{M_v} t_m^v$  is the total of all unexpired warranties at vendor  $v$  and  $\eta p_v^* \Omega^2 / 2$  is the mean value of this quantity under the static policy  $p_v^*$ . Hence an allocation policy based on the indices in (4.21) (see Theorem 4.1) will favour vendors whose current commitments (as measured by  $\sum_1^{M_v} t_m^v$ ) are most below those indicated by the optimal static policy, the difference being factored by the (positive) cost constant  $c_v$ . In this example the derived indices yield an allocation policy which dynamically tracks the optimal static solution. This provides a straightforward illustration of this PI heuristic policy developed in Theorem 4.1.

Further, the PI heuristic developed in this section is readily adjusted to encompass some major extensions. For example, the assumption that each order should be allocated to a single vendor is intended to reflect the interests of customers, but other possibilities

exist. The items within a single order of size  $x$  at time zero could have been distributed among all available vendors as expressed by

$$x \rightarrow (z_1, z_2, \dots, z_V), \sum_{v=1}^V z_v = x. \quad (4.22)$$

where  $z_v$  denotes the corresponding allocation to vendor  $v$  from the order of size  $x$ .

As a result, the total expected cost under  $p^*$  up to time  $T$  when the initial allocation of the items within the order of size  $x$  is made to all the vendors is as follows:

$$\sum_{v=1}^V C(v, p^*, T | N, z_v) = \sum_{v=1}^V C_v \{p_v^*, \Omega | (x^v, t^v)^{z_v}\} + (T - \Omega) \sum_{v=1}^V E_{\eta W p_v^*} \{\gamma_v(\overline{M}_v)\},$$

$$1 \leq v \leq V, T \geq \Omega. \quad (4.23)$$

Hence, cost minimization is achieved by the minimization of

$$\sum_{v=1}^V [C_v \{p_v^*, \Omega | (x^v, t^v)^{z_v}\} - C_v \{p_v^*, \Omega | (x^v, t^v)\}] \equiv \sum_{v=1}^V I_{vz_v}(x^v, t^v) \quad (4.24)$$

Hence, the allocation policy under a single DP policy improvement step in this adjusted scenario operates to seek the minimal value of the index

$$\min_{z_v} \sum_{v=1}^V I_{vz_v}(x^v, t^v) \quad (4.25)$$

$$\text{s.t.} \quad \sum_{v=1}^V z_v = x, z_v \in \mathbb{N}$$

(4.25) is a resource allocation problem with discrete variables. In Chapter 3, we have indicated that such a problem is solvable by a greedy heuristic given that each additive component in the objective function is increasing and convex. See Gross (1956). In this case, if we could prove that each of  $I_{vz_v}(x^v, t^v)$  in (4.25) is increasing and convex in  $z_v$ , then a simple greedy heuristic can be developed to solve the problem (4.25). For ease of notation, we drop the vendor subscript  $v$  for the following Theorem 4.3 and its associated proof.

It is straightforward from Lemma 4.2 that  $I_z(x, t)$  has the following form

$$I_z(x, t) = \sum_{m=0}^M \int_{t_m}^{t_{m+1}} E_{\eta t p^*} \left\{ \gamma(z + \sum_{r=m+1}^M x_r, \overline{M}^t) - \gamma(\sum_{r=m+1}^M x_r, \overline{M}^t) \right\} dt \quad (4.26)$$

**Theorem 4.3 (Properties of the vendor-specific calibrating indices)** *If the vendor-specific cost rate  $f$  is increasing convex then the calibrating index  $I_z(\mathbf{x}, \mathbf{t})$  is increasing and convex in  $z$  for all  $(\mathbf{x}, \mathbf{t})$ .*

*Proof.* Since  $f$  is increasing convex it follows that, for any fixed and positive  $s$ ,  $f(z+s) - f(s)$  is increasing in  $z$ . We readily conclude that any quantity of the form

$$E_{\eta|p^*}\{\gamma(z+y, \overline{M}^t) - \gamma(y, \overline{M}^t)\} \equiv H(z)$$

will be increasing in  $z$  for any fixed  $y$ . The increasing property of the calibrating index  $I_z(\mathbf{x}, \mathbf{t})$  is now immediate followed from (4.26).

To prove the convexity of the index derived, write  $h(z) = f(z+s) - f(s)$ . We observe that the increment

$$h(z+1) - h(z) = f(z+1+s) - f(z+s)$$

is always positive and increasing in  $z, z \in \mathbb{N}$  for any positive constant  $s$  because of the increasing convex property of  $f$ . Similarly, we conclude that the increment

$$H(z+1) - H(z) = E_{\eta|p^*}\{\gamma(z+1+y, \overline{M}^t) - \gamma(z+y, \overline{M}^t)\}$$

is always positive and increasing in  $z$  for any fixed  $y$ . The convexity of the calibrating index  $I_z(\mathbf{x}, \mathbf{t})$  follows.  $\square$

With the result in Theorem 4.3, an efficient heuristic may be developed to solve (4.25) and is described as follows. We restore the vendor subscript  $v$  for this purpose.

#### Greedy Algorithm for (4.25)

Step 0: Set  $z_v = 0, 1 \leq v \leq V$ ;

Step 1: Choose any  $j \in \operatorname{argmin}_{1 \leq v \leq V} [I_{v(z_v+1)}(\mathbf{x}^v, \mathbf{t}^v) - I_{vz_v}(\mathbf{x}^v, \mathbf{t}^v)]$ ;

Step 2: Set  $z_j = z_j + 1$ ;

Step 3: If  $\sum_{v=1}^V z_v < x$ , go to Step 1; otherwise stop.

where  $I_{vz_v}(\mathbf{x}^v, \mathbf{t}^v) = 0$  when  $z_v = 0$ .

In this scenario, we expect the resulting cost rates obtained under the *one-item-at-a-time* allocation strategy implemented by utilizing the above greedy heuristic would be certainly no more than those incurred when all the items in a single order are allocated to a single vendor.

Before we proceed to evaluate the indices derived in this section numerically, the following might support the upcoming simulation study and further clarify matters.

### Lower and Upper bounds

We now establish upper and lower bounds for the cost rates obtained under the DP policy improvement allocation policy. It has already been pointed out that  $G(\mathbf{p}^*)$  is an accessible upper bound on the average cost rate incurred by an optimal dynamic policy for the approximate costs model. The following are theoretical attempts to obtain a lower bound for the cost rate concerned. We use  $\tilde{G}(K)$  for the optimal value of the optimization problem in (3.7) but without the integrality constraints,

$$\begin{aligned} \tilde{G}(K) = \min_{k_v} \quad & \sum_{v=1}^V f_v(k_v) \\ \text{s.t.} \quad & \sum_{v=1}^V k_v = K \geq 0, \\ & k_v \geq 0, \quad v = 1, \dots, V. \end{aligned} \tag{4.27}$$

**Lemma 4.4** *When the vendor-specific cost rates  $f_v, 1 \leq v \leq V$ , are convex,  $\tilde{G}(\eta\Omega\beta)$  is a lower bound for the average cost rate achieved by any allocation policy for the approximate costs model.*

*Proof.* By standard theory we may restrict to stationary allocation policies. Consider one such policy  $\pi$ , say, and use  $P_v(\pi, n)$  for the proportion of time (over an infinite horizon) during which vendor  $v$  has  $n$  items under warranty. Under the approximate

costs model, the average cost rate under policy  $\pi$  may be written

$$\sum_{v=1}^V \sum_{n \in \mathbb{N}} f_v(n) P_v(\pi, n) \geq \sum_{v=1}^V f_v \left\{ \sum_{n \in \mathbb{N}} n P_v(\pi, n) \right\} \quad (4.28)$$

$$\geq \tilde{G} \left\{ \sum_{v=1}^V \sum_{n \in \mathbb{N}} n P_v(\pi, n) \right\} \quad (4.29)$$

$$= \tilde{G}(\eta \Omega \beta). \quad (4.30)$$

Inequality (4.28) uses the convexity of cost rate function  $f_v$  and is a consequence of Jensen's inequality while inequality (4.29) follows from the definition of  $\tilde{G}$  in (4.27). Equality (4.30) uses the fact that the average number of items under warranty (i.e. with any vendor) is  $\eta \Omega \beta$ . See (4.3). This concludes the proof.  $\square$

### Other heuristics

For comparison purposes in the upcoming simulation study, other simpler heuristics are described as follows.

**Dynamic greedy heuristic** We recall the greedy algorithm designed for the static allocation model in Chapter 3. This yields a natural allocation policy for the problem described in Section 4.2. We take  $f_v(N_v)$  as an approximation to the instantaneous costs rate incurred when  $N_v$  items are under warranty at vendor  $v$ . Suppose at some arbitrary time that an order of size  $x$  is awaiting to be allocated and that the current system state is given by (4.2). This heuristic policy is determined by the expected costs rate escalations

$$f_v(N_v + x) - f_v(N_v) \quad (4.31)$$

when a warranty population at vendor  $v$  is increased from  $N_v$  to  $N_v + x$ . The dynamic greedy allocation heuristic will assign the incoming order of size  $x$  to the vendor with smallest associated value of (4.31). As will become clear, numerical evidence suggests that this very simple heuristic performs extraordinarily well..

**Tracking heuristic** An alternative way of using static models to develop allocation heuristics is as follows: Suppose that the mean size of the warranty population  $\eta\Omega\beta$  is an integer (and otherwise take the nearest integer to it). Now consider the static optimization problem in (3.7) with  $K = \eta\Omega\beta$ . Use

$$K(\eta\Omega\beta) = \{k_1(\eta\Omega\beta), k_2(\eta\Omega\beta), \dots, k_V(\eta\Omega\beta)\}$$

for an optimal set of vendor allocations. We propose an allocation heuristic for the problem described in Section 4.2 which dynamically tracks this static solution by allocating an incoming order to any vendor for which the difference

$$N_v - k_v(\eta\Omega\beta) \quad (4.32)$$

is minimal. Note that this heuristic takes no account of order size.

**Smallest workload heuristic** This is a version of the *join-the-shortest-queue* heuristic for conventional dynamic routing problems. The incoming order  $x$  now joins any vendor with the smallest current committed workload as measured by the sum of unexpired warranty periods

$$\sum_{m=1}^{M_v} (t_m^v x_m^v) \quad (4.33)$$

where  $M_v$  is the current number of orders at vendor  $v$  whose warranties have not expired yet.

Experience of the utilization of DP policy improvement in the manner advocated above in a variety of application contexts (see Glazebrook et al. (2004) and Ansell et al. (2001)) suggests that the index heuristic described in Theorem 4.1 and Lemma 4.2 will perform well for the problem described in Section 4.2. In the simulation study conducted in the following section this heuristic will be used both as a benchmark by which other (simpler) allocation policies may be judged and also as a policy of great interest in its own right.

## 4.4 Simulation Study

An extensive simulation study has been undertaken to address the following issues of interest:

1. Is the PI heuristic more effective than other simpler heuristics derived from the static allocation model in Chapter 3 given that it makes more extensive use of system state information?
2. If simple heuristics, namely the dynamic greedy heuristic and the tracking heuristic perform just as well as the PI heuristic, what does that imply?
3. Are the cost rates obtained by the PI heuristic actually constrained within the upper and lower bounds discussed theoretically in the previous section?
4. As the variability of the population size increases, will those costs rates obtained from alternative heuristics depart further away from the lower bounds based on fixed allocations?
5. When bulk orders arrive, do the 'one-item-at-a-time' allocations among all available vendors significantly reduce the cost rates obtained from the cases in which the assumption is made that all the items within a single order are simply allocated to a single vendor?
6. What cost penalties might be incurred by using the smallest workload heuristic which takes no account of different service capacities across all the vendors?

Before we present numerical evidence to shed light on such matters, the system dynamics coded in the simulation are illustrated by a flow chart in Figure 4.1. It demonstrates the event-driven Monte Carlo simulation that we conducted for the system in which the length of the life time is set to be long enough to ensure stable states. We actually start to calculate the costs when the system is in steady state after a burn-in period. This is not shown on the chart for simplicity of presentation. While the vendor with smallest index evaluated under our alternative heuristics is chosen to take care of the arriving order, the underlying system evolutions at individual vendors are independent

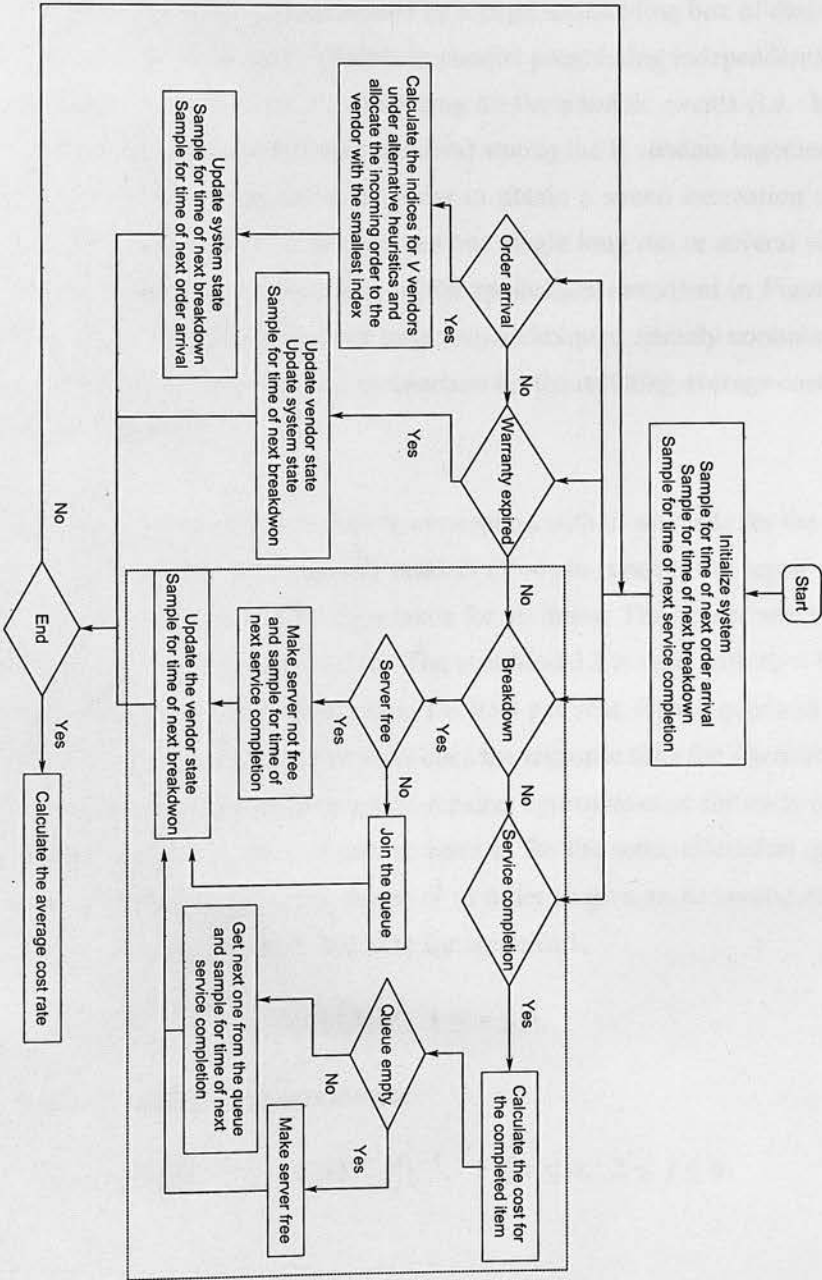


Figure 4.1: A flow chart of the simulation for the full system

from each other. Hence, for simplicity, we illustrate the dynamics evolving over time within a single vendor, which are highlighted by a large surrounding box of dash lines on the chart. Given that there are  $V$  vendors in parallel progressing independently, the simulation chooses the next event by comparing all the possible events (i.e. breakdowns, service completions and warranty expiries) among the  $V$  vendors together with the event of the next incoming order. In order to obtain a sound estimation of the cost rate with small standard error, one can use one single long run or several shorter ones. Technically, hundreds of iterations of the application described in Figure 4.1 have been conducted along with variance reduction techniques, namely common random numbers (CRN<sup>2</sup>), to ensure a good comparison for the resulting average cost rates under alternative heuristics.

The system parameters are chosen mostly to correspond with those made for the static model. In all cases studied, we have four vendors ( $V=4$ ) to conduct the repair work. A warranty period of two years ( $\Omega = 2$ ) is taken for all items. The breakdown rate of an individual item is 1.2 per year ( $\lambda = 1.2$ ). The cost Model 2 is used with  $c_v = 0, 1 \leq v \leq 4, d = 1$  and the service time threshold is  $\tau = 0.04$  per year. Hence goodwill costs grow linearly (at the rate of one unit per year) once the response time for a broken item exceeds 10 working days. We adopt a single-repairer approximation for each vendor and we have the same six profiles of service rates as for the static allocation model, except that we introduce an adjusting factor  $H$  in order to give an economic service capacity for the problem concerned. We have for scenario 1,

$$\mu_{v1} = (62.5)H, \quad 1 \leq v \leq 4,$$

while in scenario  $j$ , vendor  $v$  has service rate

$$\mu_{vj} = 250Hy_j^{v-1}(1-y_j)(1-y_j^4)^{-1}, \quad 1 \leq v \leq 4, \quad 2 \leq j \leq 6,$$

where

$$y_j = 1 - 0.1(j-1), \quad 2 \leq j \leq 6.$$

---

<sup>2</sup>CRN requires synchronization of the random number streams for the system in parallel, which has also been called correlated sampling, matched streams or matched pairs.

$H$  takes the values 0.7 for Tables 4.2(a) and 4.3(a), 1.5 for Table 4.3(b), 2.5 for Table 4.2(b), 3.0 for Table 4.3(c), 3.8 for Table 4.3(d) and 4.7 for Table 4.3(e). Table 4.2 contains cases in which items are purchased as singletons at rates of  $\eta = 50$  per year (Table 4.2(a)) and  $\eta = 250$  per year (Table 4.2(b)) respectively. Table 4.3 contains cases in which orders are placed in bulk at a rate of  $\eta = 25$  per year and these are required to be allocated to a single vendor. The positive order size  $X$  is such that  $X - 1$  has a Poisson distribution with mean  $\beta - 1$ , which takes values 1 (Table 4.3(a)), 5 (Table 4.3(b)), 11 (Table 4.3(c)), 15 (Table 4.3(d)) and 19 (Table 4.3(e)). Using (4.3), the means and standard deviations of the total population of items under warranty in steady state for the cases presented in Table 4.2 and Table 4.3 are given in Table 4.1 below.

	2(a)	2(b)	3(a)	3(b)	3(c)	3(d)	3(e)
$E(N)$	100	500	100	300	600	800	1000
$\sqrt{\text{var}(N)}$	10.00	22.36	15.81	45.28	88.03	116.40	144.74

Table 4.1: Means and standard deviations of  $N$ , the total number of items under warranty, for examples in subsequent tables.

Note that with the introduction of random order sizes into the instances in Table 4.3, the variability of the entire population size  $N$  is increased. The standard deviation of  $N$  grows approximately in a linear fashion with its mean.

Rows in Table 4.2 and Table 4.3 correspond to a given scenario of service rates while column heads in the two tables are explained as follows:

$\bar{G}(\eta\Omega\beta)$ : These columns contain the best cost rates obtained from the optimization problem in (3.7) for the relevant mean population sizes. Since the vendor-specific cost rates for our examples are close to convex, then these values are close to the lower bound,  $\tilde{G}(\eta\Omega\beta)$ , given by Lemma 4.3 on achievable cost performance under the approximate costs model;

- $G(p^*)$ : These columns contain values of the average cost rates incurred when an optimal static policy  $p^*$  in steady state is applied to the problem under the approximate costs model;
- PIH: These columns contain simulation-based estimates of the average cost rates incurred when the index-based PI heuristic of Theorem 4.1 and Lemma 4.2 is applied to the full system as described in Figure 4.1. Bracketed figures are the standard errors of the corresponding cost rate estimates;
- GRE: These columns contain simulation-based estimates of the average cost rate incurred when the dynamic greedy allocation heuristic described in Section 4.3 around (4.31) is applied to the full system as described in Figure 4.1. Bracketed figures are the standard errors of the corresponding estimates of the *difference* between the cost rates incurred by the PI and the dynamic greedy heuristics;
- TRA: These columns contain simulation-based estimates of the average cost rates incurred when the tracking heuristic is applied to the full system as described in Figure 4.1. This heuristic is described around (4.32) at the final part of Section 4.3. Bracketed figures are the standard errors of the corresponding estimates of the *difference* between the cost rates incurred by the PI and the tracking heuristics;
- SMA: These columns contain simulation-based estimates of the average cost rate incurred in the full system when each incoming order is allocated to any vendor whose current committed workload is smallest, as described at the end of Section 4.3. Bracketed figures are the standard errors of the corresponding estimates of the *difference* between the cost rates incurred by the PI and this *equal shares* heuristic.

SCENARIO NUMBER	$\bar{G}(100)$	$G(p^*)$	PIH	GRE	TRA	SMA
1	2.976	3.620	3.174 [0.015]	3.171 (0.010)	3.171 (0.010)	3.174 (0.0)
2	2.957	3.594	3.155 [0.014]	3.135 (0.009)	3.131 (0.010)	3.522 (0.012)
3	2.886	3.506	3.064 [0.014]	3.053 (0.011)	3.069 (0.010)	4.903 (0.020)
4	2.745	3.343	2.924 [0.015]	2.909 (0.010)	2.915 (0.010)	8.146 (0.032)
5	2.517	3.076	2.704 [0.014]	2.705 (0.011)	2.801 (0.010)	13.964 (0.044)
6	2.183	2.685	2.382 [0.013]	2.369 (0.009)	2.402 (0.010)	21.562 (0.062)

Table 4.2(a)

SCENARIO NUMBER	$\bar{G}(500)$	$G(p^*)$	PIH	GRE	TRA	SMA
1	11.224	15.098	12.174 [0.076]	12.183 (0.050)	12.183 (0.050)	12.174 (0.0)
2	11.186	15.045	12.146 [0.069]	12.155 (0.045)	12.150 (0.054)	19.626 (0.093)
3	11.042	14.868	12.007 [0.080]	12.017 (0.052)	12.043 (0.056)	42.255 (0.114)
4	10.765	14.524	11.777 [0.079]	11.675 (0.052)	11.811 (0.048)	72.637 (0.127)
5	10.315	13.963	11.364 [0.074]	11.290 (0.053)	11.449 (0.054)	107.330 (0.138)
6	9.643	13.124	10.672 [0.069]	10.728 (0.050)	10.984 (0.052)	135.736 (0.220)

Table 4.2(b)

Table 4.2: Results of a simulation study of the comparative performance of four allocation heuristics when orders are singletons. See above text for further details.

SCENARIO NUMBER	$\bar{G}(100)$	$G(p^*)$	PIH	GRE	TRA	SMA
1	2.976	4.855	3.416 [0.025]	3.427 (0.012)	3.427 (0.012)	3.420 (0.010)
2	2.957	4.823	3.405 [0.025]	3.408 (0.010)	3.412 (0.011)	3.820 (0.014)
3	2.886	4.714	3.333 [0.025]	3.340 (0.011)	3.349 (0.011)	5.291 (0.023)
4	2.745	4.505	3.195 [0.024]	3.205 (0.010)	3.219 (0.010)	8.670 (0.038)
5	2.517	4.166	2.955 [0.023]	2.963 (0.011)	3.026 (0.011)	14.364 (0.065)
6	2.183	3.667	2.665 [0.022]	2.639 (0.010)	2.802 (0.006)	22.025 (0.037)

Table 4.3(a)

SCENARIO NUMBER	$\bar{G}(300)$	$G(p^*)$	PIH	GRE	TRA	SMA
1	11.548	30.401	16.371 [0.196]	16.400 (0.048)	16.400 (0.048)	16.383 (0.047)
2	11.512	30.322	16.381 [0.206]	16.390 (0.058)	16.386 (0.052)	19.921 (0.073)
3	11.379	30.051	16.228 [0.198]	16.266 (0.049)	16.372 (0.055)	30.610 (0.104)
4	11.116	29.526	16.034 [0.194]	16.098 (0.049)	16.228 (0.050)	47.078 (0.121)
5	10.694	28.677	15.681 [0.191]	15.729 (0.052)	16.195 (0.057)	66.734 (0.143)
6	10.066	27.398	15.255 [0.190]	15.290 (0.048)	16.232 (0.058)	88.561 (0.174)

Table 4.3(b)

SCENARIO NUMBER	$\bar{G}(600)$	$G(p^*)$	PIH	GRE	TRA	SMA
1	10.793	52.511	21.591 [0.384]	21.661 (0.064)	21.661 (0.064)	21.607 (0.061)
2	10.753	52.375	21.515 [0.391]	21.646 (0.068)	21.749 (0.073)	30.147 (0.125)
3	10.611	51.907	21.537 [0.386]	21.639 (0.069)	21.817 (0.084)	54.946 (0.123)
4	10.336	51.003	21.435 [0.383]	21.555 (0.072)	21.938 (0.065)	89.695 (0.247)
5	9.887	49.533	20.977 [0.370]	21.057 (0.067)	21.940 (0.067)	130.094 (0.303)
6	9.214	47.339	20.705 [0.375]	20.793 (0.073)	22.610 (0.076)	173.376 (0.321)

Table 4.3(c)

SCENARIO NUMBER	$\bar{G}(800)$	$G(p^*)$	PIH	GRE	TRA	SMA
1	19.056	84.417	38.578 [0.683]	38.649 (0.099)	38.649 (0.099)	38.649 (0.100)
2	19.007	84.243	38.403 [0.674]	38.770 (0.093)	38.672 (0.102)	51.672 (0.150)
3	18.830	83.645	38.451 [0.662]	38.600 (0.100)	38.814 (0.096)	86.559 (0.236)
4	18.483	82.490	38.381 [0.703]	38.632 (0.102)	39.048 (0.097)	132.994 (0.267)
5	17.916	80.611	38.028 [0.647]	38.188 (0.097)	39.078 (0.107)	185.422 (0.320)
6	17.060	77.816	37.761 [0.634]	37.925 (0.083)	39.955 (0.101)	240.898 (0.309)

Table 4.3(d)

SCENARIO NUMBER	$\bar{G}(1000)$	$G(p^*)$	PIH	GRE	TRA	SMA
1	22.406	108.173	48.982 [0.652]	49.249 (0.079)	49.249 (0.079)	49.110 (0.078)
2	22.356	107.961	48.937 [0.646]	49.250 (0.086)	49.331 (0.072)	66.162 (0.138)
3	22.166	107.233	48.935 [0.658]	49.098 (0.082)	49.458 (0.088)	110.842 (0.203)
4	21.798	105.826	48.915 [0.642]	49.207 (0.082)	49.753 (0.088)	168.940 (0.224)
5	21.194	103.541	48.583 [0.626]	48.715 (0.089)	49.839 (0.094)	233.745 (0.278)
6	20.278	100.135	48.455 [0.639]	48.528 (0.086)	50.898 (0.091)	302.330 (0.268)

Table 4.3(e)

Table 4.3: Results of a simulation study of the comparative performance of four allocation heuristics when order sizes are random. See above text for further details.

We further present some example results in Tables 4.4(a) and 4.4(b) following the *one-item-at-a time* rule when orders are of random sizes. We use the modified index under the PIH in (4.24) and discussions around (4.24) to estimate the associated cost rate. The other heuristics are also adjusted accordingly toward this *one-item-at-a time* version of our problems. Tables 4.4(a) and 4.4(b) contain the equivalent sets of problems to those in Tables 4.3(a) and 4.3(b) respectively.

SCENARIO NUMBER	$\bar{G}(100)$	$G(p^*)$	PIH	GRE	TRA	SMA
1	2.976	4.855	3.394 [0.026]	3.385 (0.011)	3.385 (0.011)	3.394 (0.0)
2	2.957	4.823	3.391 [0.025]	3.374 (0.010)	3.389 (0.012)	3.789 (0.011)
3	2.886	4.714	3.317 [0.023]	3.301 (0.011)	3.317 (0.011)	5.240 (0.024)
4	2.745	4.505	3.181 [0.024]	3.171 (0.009)	3.182 (0.011)	8.624 (0.044)
5	2.517	4.166	2.962 [0.023]	2.952 (0.013)	2.999 (0.011)	14.311 (0.061)
6	2.183	3.667	2.666 [0.025]	2.617 (0.014)	2.761 (0.012)	21.879 (0.075)

Table 4.4(a)

SCENARIO NUMBER	$\bar{G}(300)$	$G(p^*)$	PIH	GRE	TRA	SMA
1	11.548	30.401	16.315 [0.218]	16.337 (0.052)	16.337 (0.052)	16.315 (0.0)
2	11.512	30.322	16.216 [0.204]	16.305 (0.055)	16.230 (0.048)	19.704 (0.063)
3	11.379	30.051	16.211 [0.195]	16.169 (0.053)	16.141 (0.057)	30.586 (0.106)
4	11.116	29.526	15.991 [0.196]	15.931 (0.052)	16.096 (0.056)	46.893 (0.121)
5	10.694	28.672	15.586 [0.185]	15.499 (0.056)	15.825 (0.049)	66.525 (0.142)
6	10.066	27.395	15.166 [0.192]	15.063 (0.052)	15.847 (0.056)	88.084 (0.153)

Table 4.4(b)

Table 4.4: Equivalent Results of a simulation study for Table 4.3(a) and Table 4.3(b) but under *one-item-at-a time* rule. See the following text for further details.

SCENARIO 1	Sample Mean	Sample Standard Deviation	Sample Size
PIH	0.0786	0.0399	1017
GRE	0.0774	0.0668	1030
p*	0.0826	0.0808	967
SCENARIO 2	Sample Mean	Sample Standard Deviation	Sample Size
PIH	0.0678	0.0372	1180
GRE	0.0675	0.0597	1184
p*	0.0641	0.0638	1247
SCENARIO 3	Sample Mean	Sample Standard Deviation	Sample Size
PIH	0.0543	0.0307	1474
GRE	0.0541	0.0517	1477
p*	0.0558	0.0572	1433
SCENARIO 4	Sample Mean	Sample Standard Deviation	Sample Size
PIH	0.0446	0.0264	1794
GRE	0.0445	0.0424	1798
p*	0.0454	0.0454	1760
SCENARIO 5	Sample Mean	Sample Standard Deviation	Sample Size
PIH	0.0391	0.0267	2044
GRE	0.0389	0.0348	2054
p*	0.0398	0.0402	2010
SCENARIO 6	Sample Mean	Sample Standard Deviation	Sample Size
PIH	0.0322	0.0233	2485
GRE	0.0320	0.0306	2501
p*	0.0313	0.0306	2557

Table 4.5: Estimated means and standard deviations of inter-allocation times at vendor 1 generated by three heuristics for the scenarios in Table 4.3(a)

Observations on the results presented in the above tables help us to answer the questions listed at the beginning of the section. The findings in regards to those matters are summarized sequentially below:

**Strong performance of dynamic greedy and tracking heuristics** For the instances of singleton-order problems in Table 4.2, the dynamic greedy and tracking heuris-

tics based on solutions to the optimization problem (3.7) perform comparably to the PI heuristic policies. Where the introduction of random order sizes in Table 4.3 results in increased variability of the warranty population, the PI heuristic outperforms the others (i.e. has the lowest estimated cost rate) in 29 of the 30 problem configurations. However, the margins are often small and in many individual cases fail to be statistically significant. In the worst case for the dynamic greedy heuristic its average cost rate exceeds that of the PI heuristic by just 1%; see scenario 2 in Table 4.3(d). We conclude that in most environments in which the warranty population is subject to moderate temporal variability, the dynamic greedy heuristic will perform well. The equivalent worst case figure for the tracking heuristic is 9.20%; see scenario 6 in Table 4.3(c). With regard to the latter allocation procedure there is some evidence of deteriorating performance as the differences between the vendor service rates increase. It appears that the tracking heuristic's failure to take account of the order size in making its allocations may lead it on occasion to overload vendors with small service capacity;

**Implications of the above findings** An interesting inference from the above reported strong performance of the greedy heuristic is that in the current models, once the number of items at each vendor is known, further information regarding unexpired warranty times adds *relatively* little to effective decision-making. To understand why, see Table 4.5 which records the means and standard deviations of the times between successive allocations of newly purchased machines to vendor 1 for the six scenarios of Table 4.2(a) under the heuristics PIH, GRE and the optimal static policy  $p^*$ . Note that, while the *average* rates at which the three heuristics send machines to the vendor are equal (within sampling error), the dynamic heuristics impose greater regularity on the allocations as reflected in smaller standard deviations for the inter-allocation times. This is not wholly surprising since it goes along with the associated costs moving from the upper bound  $G(p^*)$  toward the approximate lower bound  $\bar{G}(\eta\Omega\beta)$  from the static allocation model (3.7). We conclude that, if past allocations have been made effectively it will be rare to encounter a situation in which, from the perspective

of making an optimal allocation decision, one vendor dominates another with regard to machine numbers ( $N_v$ ) but is dominated with regard to (any reasonable measure of) the unexpired warranties ( $t^v$ ). Note that all of this relates to the Poisson assumption (with uniform rate) concerning arriving orders. Should the arrival process be, for example, non-homogeneous Poisson with substantial fluctuations in the arrival rate, then we would expect the patterns of unexpired warranties to be *necessarily* much more irregular and potentially more informative for (good) allocation decisions. Further consideration of this issue will be the subject of future research;

**The cost rates under the PIH are between the values of  $G(p^*)$  and  $\bar{G}(\eta\Omega\beta)$**  This claim is valid throughout the tables. The difference between the upper bound  $G(p^*)$  and the approximate lower bound  $\bar{G}(\eta\Omega\beta)$  in practice bounds the biggest possible reduction in the cost rate that the PI heuristic could achieve over the optimal static policy  $p^*$ ;

**$\bar{G}(\eta\Omega\beta)$  is not achievable by the PI heuristic if the population has great variability .**

Note that the more variability the warranty population size has, the greater is the difference between the values of  $G(p^*)$  and  $\bar{G}(\eta\Omega\beta)$ . See Table 4.1, Table 4.2(a) and Table 4.3(a) for examples. Further, through Table 4.3(a) to Table 4.3(e), the cost rates obtained by the PI heuristic depart further away from the corresponding values of  $\bar{G}(\eta\Omega\beta)$ . Hence, we come to the conclusion that the lower bound does not give a realistic indication of achievable cost rates especially when the warranty population is subject to substantial temporal variability;

**The further reduction in the cost rate under the *one-item-at-a-time* rule is not significant.**

As expected before conducting the simulation study, the cost rates obtained under the 'one-item-by-another' rule shown in Table 4.4 are lower than those of Table 4.3(a) and Table 4.3(b). However, the difference is not significant enough to persuade us that adopting this more sophisticated rule when orders are placed in bulk will achieve any significant cost reduction. Further, from the practical point of view it is less confusing for both the manufacturer and customers if all the items within a bulk order are allocated to a single vendor;

**Poor performance of the SMA heuristic** The smallest workload heuristic is a version of the *join-the-shortest-queue* rule for conventional dynamic routing problems. It works well only when the service rates are close to equal for all the vendors. It is observed in Table 4.2 to 4.4 that the PI heuristic and the SMA heuristic operate identically when the vendors have equal service rates. Nevertheless, as the distribution of service rates among vendors becomes more diverse from scenario 2 to scenario 6 in Tables 4.2 and 4.3, the SMA heuristic performs increasingly poorly. Even when the difference in vendor service rates are fairly small, for example, in scenario 2 in Table 4.2(b), the cost rate obtained by SMA heuristic exceeds that from the PI heuristic policy by over 60%.

## 4.5 Conclusions to the chapter

This chapter presented non-standard allocation models for a dynamic warranty population with partially observable system state. As it is not amenable to conduct the analysis via conventional dynamic programming, we instead developed a DP policy improvement heuristic by utilizing an approximating costs model in order to minimize the resulting cost rate (which penalizes long response times in the breakdown/repair process) or come close to doing so. This heuristic makes thorough use of the information regarding previous allocations at each vendor and operates by attaching derived indices to all available vendors. The incoming order (whether a singleton or in bulk) is assigned to the vendor with smallest index. Three other simpler heuristics were proposed to compare against this sophisticated allocation procedure in an extensive simulation study. In contrast to our expectation, the deployment of additional information concerning unexpired warranties of previous allocations at each vendor does not result in the DP policy improvement heuristic significantly outperforming the simpler dynamic greedy heuristic in most cases studied. The potential reasons behind this have been discussed in the above text. We would expect the data of unexpired warranties to be more informative for a non-homogeneous arrival process.

## Chapter 5

# Dynamic Allocations upon Real-time Breakdowns

### 5.1 Introduction to the chapter

This chapter is concerned with the dynamic allocation of the incoming breakdowns to one of the service vendors for repairs. Each breakdown will be sent to a vendor, chosen by the central decision maker on the basis of the current system state. Therefore, unlike the static allocation model for which every time an item fails it is sent to its pre-assigned vendor, a same item under the dynamic allocation setting might be repaired at different service vendors if it fails more than once during its warranty. With the same objective of minimizing the overall cost rate in the long run, we are to expect a lower cost rate obtained under the dynamic allocation model in comparison to the static solution, because the decision epoch now is delayed until the real-time breakdown when the information regarding the queue length at each vendor becomes available. There is already a moderately large quantity of literature devoted to versions of this relatively conventional model concerning dynamic routing of items to a collection of service stations. See, for example, Hordijk and Koole (1990), Weber (1978) and Winston (1977).

In this multi-vendor scenario, the complexity of the dynamic routing problem together with the particular cost structure proposed in Chapter 2 make direct application of

stochastic dynamic programming to produce the optimal policies numerically unrealistic. Hence, we focus on the development of effective routing heuristics which are state dependent in utilizing the data of queue length at each vendor. There is a large literature dedicated to elucidating simple structure in optimal policies for such problems involving homogeneous queues. The simple *join-the-shortest-queue* (in Section 5.4 which is referred to as JSQ) heuristic has been shown to be optimal under a range of general settings. See, for example, Weber(1978), Johri (1989), Hordijk and Koole (1990) and Menich and Serfozo (1991). However, there exist counterexamples showing that JSQ may not be optimal even for some simple cases. See, Whitt (1986), Houck (1987) and Foley and McDonald (2001). In this chapter, we focus on problems with a more general setting for which we do not assume service vendor homogeneity. Consequently, the performance of JSQ can be extremely poor (see the numerical evidence in the computation study of Section 5.4). This chapter primarily deploys Whittle's proposal for restless bandits in deriving near optimal index policies for the problem concerned. Whittle's index has appeared very promising in a range of application areas. See, for example, Ansell et al. (2003), Glazebrook, Niño-Mora, and Ansell (2002), Glazebrook, Mitchell and Ansell (2005), Glazebrook, Kirkbride, and Oueniche (2005). Its superiority in terms of computational efficiency is verified once again in this outsourcing warranty repairs problem. Above all, the numerical results demonstrate that the restless bandit heuristic is not only effective but very close to optimal. It also outperforms other simpler heuristics and generates huge savings over static allocation.

The rest of the chapter has the following structure. We formulate the dynamic allocation problem as a Markov decision process in Section 5.2. The derived optimality equation is solvable by the value-iteration algorithm for small instances (i.e. small populations and limited numbers of vendors) in Section 5.4. However, direct use of stochastic dynamic programming in dynamic routing problems of realistic size is numerically intractable. We therefore derive a near-optimal index policy by deploying Whittle's restless bandits approach in Section 5.3. Followed by an extensive numerical investigation in Section 5.4, we compare cost rates obtained under the restless ban-

dit heuristic to those of the optimal dynamic solution and the optimal static solution by employing value-iteration algorithms. We further examine the performance of the restless bandit heuristic developed against two other simpler heuristics in a simulation study where a collection of more general settings are studied. This chapter ends with some concluding remarks in Section 5.5.

## 5.2 Model Formulation

In the dynamic allocation, decisions regarding the choice of vendor are delayed until an item is actually broken. The data of queue length at each vendor is available to inform such decisions in routing each incoming failure to alternative service vendors. As before, we consider  $V$  vendors in parallel with each modelled as a single-server queue with exponentially distributed service at rate  $\mu_v, 1 \leq v \leq V$ . The generic queueing network of this dynamic allocation problem is depicted in Figure 5.1 to accommodate the stochastic nature of the warranty population. All the items that are working properly reside in the UP pool. When one of them fails, its owner will phone the after-sale service center to transfer the broken item into the DOWN pool. Meanwhile, the central decision-maker denoted as a question mark decides to which vendor the item should be sent. If service vendor  $v$  is chosen to take care of this breakdown, a fixed cost  $c_v$  and the expected goodwill cost  $a_v^I(x), I = 1, 2, 3$  which penalizes large response time based on the information of the current queue length  $x$  would be deemed to be incurred immediately. See the discussion of various goodwill cost models proposed in Chapter 2. We drop the costs model index  $I, I = 1, 2, 3$  to simplify notation in this chapter. Last but not least, items that are just out of their specified warranty periods in both UP and DOWN pools will be discharged and transferred to the DEAD pool (i.e. those items are assumed to have no impact on the system considered any longer).

However, when the warranty population size is either steady or subject to moderate variability, it is not inappropriate to assume a fixed warranty population of size  $K$ . This enables the formulation of the dynamic allocation problem as a routing control prob-

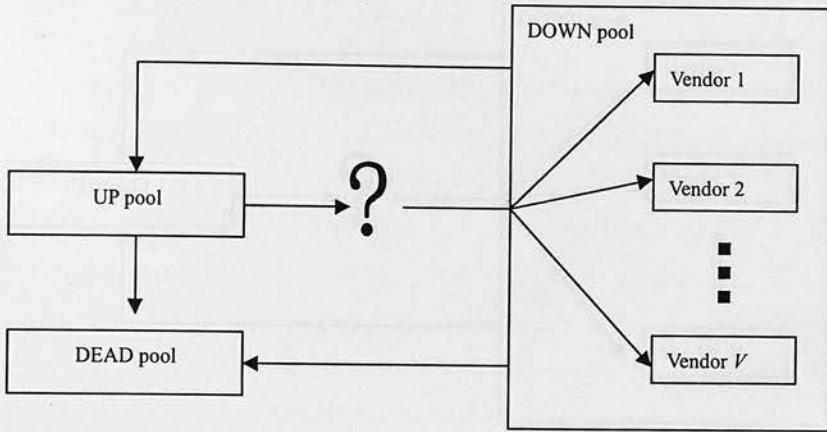


Figure 5.1: Dynamic routing system with a variable population

lem in a finite queueing network as depicted in Figure 5.2. It simplifies matters if the population size is held constant. Therefore, from the standard approach we formulate this controlled dynamic system as a Markov decision model as follows.

At each decision epoch concerning the allocation of an incoming failure, the system is observed to be in one of its possible states. The set of finite states is denoted by  $S$  that includes all the possible combinations of queue lengths (where queue lengths are defined to include any items in service at the  $V$  service vendors). This is given by  $S \equiv \{\mathbf{x} = (x_1, \dots, x_V) \in \mathbb{Z}^+ : x_v \geq 0, \sum_{v=1}^V x_v \leq K\}$ , where  $x_v$  is the number of broken items that have been routed to vendor  $v$ . For each state  $\mathbf{x} \in S$ , we choose from a set of actions  $A(\mathbf{x}) = \{1, \dots, V\}$ , where action  $v \in A(\mathbf{x})$  indicates that the vendor  $v$  is chosen to undertake repair work for the incoming failure when the system is in state  $\mathbf{x}$ .

Next we enumerate all the transition rates out of state  $\mathbf{x}$ . A new failure arrives at rate  $\lambda(K - \sum_{v=1}^V x_v)$ . When it is to be routed to vendor  $v$ , the fixed repair cost  $c_v$  and the goodwill cost are incurred, and the state changes from  $\mathbf{x}$  to  $\mathbf{x} + \hat{e}_v$ , where  $\hat{e}_v$  denotes the  $v^{\text{th}}$  unit vector. A repair completion occurs at rate  $\sum_{v=1}^V \mu_v$ . When the repair service is done at vendor  $v$ , the state changes from  $\mathbf{x}$  to  $\mathbf{x} - \hat{e}_v$ . Hence, a total transition rate out




Figure 5.2: Dynamic routing system with a closed population

of state  $\mathbf{x}$  is given by

$$\lambda(K - \sum_{v=1}^V x_v) + \sum_{v=1}^V \mu_v$$

We follow Opp, Glazebrook and Kulkarni (2005) for the standard uniformization of a time scale such that  $\lambda K + \sum_{v=1}^V \mu_v = 1$ . By doing so, the total transition rate out of state  $\mathbf{x}$  is 1 and we have 'fictitious' transitions which result in no change of the state with the following rate.

$$1 - \lambda(K - \sum_{v=1}^V x_v) - \sum_{v=1}^V \mu_v = \lambda \sum_{v=1}^V x_v. \tag{5.1}$$

Let  $V_t(\mathbf{x})$  denote the optimal expected cost incurred over the first  $t$  time epochs when the initial state is  $\mathbf{x}$ . The state space  $S$  and the action space  $A$  are both finite. Using the transition rates enumerated above, the dynamic programming equation to obtain the optimal costs for this sequential decision problem is given by

$$\begin{aligned} V_t(\mathbf{x}) &= \lambda \sum_{v=1}^V x_v V_{t-1}(\mathbf{x}) + \sum_{v=1}^V \mu_v V_{t-1}(\mathbf{x} - \hat{e}_v) \\ &+ \lambda [K - \sum_{v=1}^V x_v] \min_{j=1, \dots, V} [c_j + a_j(x_j) + V_{t-1}(\mathbf{x} + \hat{e}_j)], t \geq 1 \end{aligned} \tag{5.2}$$

$V_t$  is computed from the  $V_{t-1}$  terms on the right-hand side of (5.2).  $V_{t-1}(\mathbf{x})$  denotes the optimal expected cost incurred over the remaining  $t - 1$  decision epochs from state  $\mathbf{x}$ .

We set  $V_0 \equiv 0$

The optimality equation (5.2) is recursively solvable for small populations and limited numbers of vendors (for instance,  $K \leq 200$  and  $V = 2$ ) by the implementation of value-iteration algorithms, which will be described in Section 5.4.1. However, seeking an optimal solution to the DP equation (5.2) for large values of  $K$  or  $V$  (e.g.,  $K \geq 500$  and  $V = 4$ ) often results in huge computational costs. To facilitate the determination of good decisions for our dynamic allocation problem in a cost-effective manner, we thereby develop a near-optimal index policy derived from a restless bandit approach. This is explained in the following section.

### 5.3 The Restless Bandit Approach

We describe the essential ideas underlying the restless bandit approach proposed by Whittle (1988) and which has been introduced in Chapter 1. We will show how it may be applied to our outsourcing warranty problem. A detailed description for deriving the index policy from the restless bandit heuristic is given in the rest of this section.

The original dynamic routing problem with the objective to minimize the overall expected average cost rate for the system can be written as

$$C^{opt} = \min_{u \in \hat{U}} \sum_{v=1}^V \tilde{C}_v(u) \quad (5.3)$$

where  $\tilde{C}_v(u)$  denotes the time average cost rate incurred at vendor  $v$  under policy  $u \in \hat{U}$ , while  $\hat{U}$  denotes a general class of stationary policies by which each incoming arrival can be routed to any service vendor available. Generally,  $\tilde{C}_v(u)$  is obtained by the aggregation of the actual costs associated with all individual repairs averaged over time. As explained in Section 4.2 of Chapter 4, the actual cost  $c_v(r)$  in terms of a response time  $r$  is incurred for every repair after it is admitted at vendor  $v$ .

We consider an infinite population approximation to our finite population problems with a fixed arrival rate of  $\bar{\lambda} = K\lambda$ , which satisfies all of the sufficient conditions in the development of the vendor indices using Whittle's idea (see Opp, Glazebrook and Kulkarni (2005)). We further assume that each breakdown should be admitted to the system. Hence,  $\bar{\lambda}$  is also the overall admission rate.

We then relax the class of policies to those which route incoming arrivals to one single service vendor on average. Expressed differently, we now suppose a single vendor is facing an incoming Poisson stream of failures, which has a breakdown rate of  $\bar{\lambda}$ . Two actions are available for this single vendor  $v$  to choose from, i.e., to *accept* or *reject* each incoming repair. Accordingly, the arriving repair will be routed to vendor  $v$  if it is accepted or to other vendors in the full multi-vendor scenario if vendor  $v$  declines to provide the service. Suppose  $U$  to be a set of stationary policies for routing each incoming breakdown to any number of chosen vendors and let  $u \in U$ .  $\tilde{A}_v(u)$  denotes the admission rate at vendor  $v$  under policy  $u \in U$  while  $\tilde{R}_v(u)$  denotes the rejection rate where

$$\tilde{R}_v(u) = \bar{\lambda} - \tilde{A}_v(u). \quad (5.4)$$

The original dynamic routing problem (5.3) is therefore relaxed and can be expressed by

$$\begin{aligned} \min_{u \in U} \quad & \sum_{v=1}^V \tilde{C}_v(u) \\ \text{s.t.} \quad & \sum_{v=1}^V \tilde{R}_v(u) = (V-1)\bar{\lambda}. \end{aligned} \quad (5.5)$$

The constraint in (5.5) represents the idea underlying the set of stationary policies  $U$  in which newly arriving failures are routed to one single vendor on average.

In addition, a Lagrange relaxation to further relax the problem in (5.5) is given by

$$\hat{C}(W) = \min_{u \in U} \sum_{v=1}^V \{ \tilde{C}_v(u) + W\tilde{R}_v(u) \} - W(V-1)\bar{\lambda} \quad (5.6)$$

Where in (5.6),  $W$  is a Lagrange multiplier and can be consider as a *charge of passivity* (i.e. the penalty for rejecting an incoming breakdown). Plainly,  $\hat{C}(W) \leq C^{opt}$  for

all  $W \in \mathbb{R}$ . But the nature of the policy class  $U$  and of the objective in (5.6) admits a decomposition of the Lagrange problem by vendor. Hence, we rewrite (5.6) as follows:

$$\sum_{v=1}^V \bar{C}_v(W) - W(V-1)\bar{\lambda}$$

where

$$\bar{C}_v(W) = \min_{u \in U_v} \{ \bar{C}_v(u) + W\bar{R}_v(u) \}, \quad v = 1, 2, \dots, V \tag{5.7}$$

The Lagrange relaxation in (5.6) is thus decomposed into  $V$  optimization problems, each of which involves a single vendor as shown in (5.7).  $U_v$  denotes the set of policies for accepting or rejecting each incoming failure at vendor  $v$ . In discussing the single-vendor problem we drop the vendor identifier  $v$  to simplify the notation until further notice. The single-vendor optimization problem with rejection charge  $W$  in (5.7) aims to design a policy for choosing actions (admission or rejection) to minimize an aggregate of the average admission cost rate  $\bar{C}(u)$  and the rejection charge rate  $W\bar{R}(u)$ . Note that whether it is optimal to reject or accept an arrival is dependent on the vendor state  $x$ .

We shall next express the notion of *indexability*, which was developed by Whittle (1988). Basically, it means the set of states in which it is optimal to reject an incoming failure will decrease as we increase the value of  $W$ . Unfortunately, there are no simple sufficient conditions to guarantee indexability and it cannot be simply assumed. See, for example, Whittle (1988) and Niño-Mora (2002). Nevertheless, the following theoretical discussions will show that the system in our context is indeed indexable. Once we have indexability, we may derive a state-dependent index  $W(x)$  as a minimum charge for passivity which causes the admission action to be optimal in state  $x$ . We write

$$W(x) = \inf\{W; \text{it is optimal to accept an arrival in state } x\}.$$

On the basis of the above, it is trivial to show that the single vendor problem in (5.7) has the following structure: for  $x$ , it is optimal to accept an incoming failure when  $W \geq W(x)$  and to reject an incoming failure when  $W \leq W(x)$ . In consequence, we

have  $W = W(x)$  when both actions (admission and rejection) are optimal for a single vendor of queue length  $x$ , where  $W(x)$  might be thought of as a *fair charge* for rejecting an arrival when the vendor is in state  $x$ .

We shall assume that all optimal policies belong to the family of threshold policies (see Ziya et al. (forthcoming)). It follows from discussions above that under indexability either accepting or rejecting an incoming breakdown is optimal for the vendor in state  $x$  when the rejection penalty  $W$  equals the fair charge  $W(x)$ . Accordingly, the policy which accepts incoming breakdowns only in states  $\{0, 1, \dots, x-1\}$  under the fair charge  $W(x)$  is optimal. The policy which accepts incoming breakdowns only in states  $\{0, 1, \dots, x\}$  under the fair charge  $W = W(x)$  is also optimal. Therefore, it follows that the fair charge  $W(x)$  makes both of the following policies optimal for the vendor:

- 1 .Policy  $u(x)$  : Accept an incoming breakdown at the vendor in states  $\{0, 1, \dots, x-1\}$ , and reject an incoming item to the vendor in states  $\{x, x+1, \dots\}$ .
- 2 .Policy  $u(x+1)$  : Accept an incoming item at the vendor in states  $\{0, 1, \dots, x\}$ , and reject an incoming item to the vendor in states  $\{x+1, x+2, \dots\}$ .

**Lemma 5.1 (Policy-specific costs)** *The expected cost associated with policy  $u(x)$  is given by*

$$\begin{aligned} \text{Cost}_{u(x)}(W) &= \Pi_x(x)\bar{\lambda}W + \sum_{n=0}^{x-1} \Pi_n(x)\bar{\lambda}(a(n)+c), \quad x \geq 1, \\ \text{Cost}_{u(0)}(W) &= \bar{\lambda}W. \end{aligned} \quad (5.8)$$

where  $\Pi_n(x)$  is the stationary distribution of the repair queue length under policy  $u(x)$ .

*Proof.* Under policy  $u(x)$ , the repair process at the vendor is modelled as a birth-death process on the states  $\{0, 1, \dots, x\}$ , where the birth rate is the total breakdown rate  $\bar{\lambda} = \lambda K$  for states  $0, 1, \dots, x-1$ . The birth rate in state  $x$  is 0. The death rate of this birth-death process is given by the service rate  $\mu$  for states  $1, 2, \dots, x$ . The death rate for state 0 is 0.

Consider the first term in (5.8). Under policy  $u(x)$  an incoming breakdown is rejected by the vendor in state  $x$ . Hence in  $x$  a rejection charge is incurred at rate  $\bar{\lambda}W$ .  $\Pi_x(x)$

denotes the steady state probability that the queue length at the vendor is  $x$  under policy  $u(x)$ . For the second term in (5.8), admission costs are incurred at rate  $\bar{\lambda}(a(n) + c)$  in states  $n = 0, 1, \dots, x - 1$ . The quantity  $\Pi_n(x)$  stands for the steady state probability that  $n$  breakdowns are present at the vendor under policy  $u(x)$ . The rejection charge is incurred at rate  $\bar{\lambda}W$  under policy  $u(0)$ . Note that  $\Pi_0(0) = 1$ . This concludes the proof.  $\square$

We proceed to present the formula for  $\Pi_n(x), 0 \leq n \leq x$ . In this  $M/M/1$  queue network with infinite population and finite buffer space, the steady state probability under policy  $u(x)$  is given by

$$\Pi_n(x) = \rho^n \Pi_0(x), \quad 0 \leq n \leq x \quad (5.9)$$

where  $\rho = \frac{\bar{\lambda}}{\mu}$  and

$$\Pi_0(x) = \frac{1 - \rho}{1 - \rho^{x+1}}, \quad \rho \neq 1; \quad (5.10)$$

and

$$\Pi_n(x) = \frac{1}{x+1}, \quad 0 \leq n \leq x, \quad \rho = 1. \quad (5.11)$$

**Lemma 5.2** *The sequence of the steady state probabilities  $\{\Pi_x(x), x \in \mathbb{N}\}$  is decreasing in  $x$  for  $\rho > 0$ .*

*Proof.* Using equations (5.9) and (5.10), we have

$$\Pi_x(x) = \frac{\rho^x}{1 + \rho + \rho^2 + \dots + \rho^x}, \quad \rho \neq 1, \quad (5.12)$$

$$\Pi_{x+1}(x+1) = \frac{\rho^{x+1}}{1 + \rho + \rho^2 + \dots + \rho^{x+1}}, \quad \rho \neq 1, \quad (5.13)$$

When  $\rho \neq 1$ , both numerator and denominator of (5.12) are divided by  $\rho^x$  while the equivalent parts of (5.13) are divided by  $\rho^{x+1}$ . We then obtain the following inequality

$$\begin{aligned} \frac{1}{1 + \frac{1}{\rho} + \frac{1}{\rho^2} + \dots + \frac{1}{\rho^x}} &> \frac{1}{1 + \frac{1}{\rho} + \frac{1}{\rho^2} + \dots + \frac{1}{\rho^{x+1}}} \\ \Leftrightarrow \Pi_x(x) &> \Pi_{x+1}(x+1) \end{aligned}$$

When  $\rho = 1$ , using (5.11) it is trivial that

$$\Pi_x(x) = \frac{1}{x+1} > \frac{1}{x+2} = \Pi_{x+1}(x+1).$$

The result follows. □

With the discussions above, we now proceed to the development of appropriate indices for an individual vendor. The fair charge  $W(x)$  is the value of the rejection penalty  $W$  for which  $Cost_{u(x)}(W) = Cost_{u(x+1)}(W)$ . The solution to this equation yields a vendor-specific index for the restless bandit heuristic. It is given by the following theorem.

**Theorem 5.3 (Vendor-specific indices for the restless bandit)** *The fair charge  $W(x)$  for rejection in state  $x$  is given by*

$$\begin{aligned} W(x) &= \sum_{n=0}^x \rho^n b(x) - \sum_{n=1}^x \rho^n b(n-1), \quad x \in \mathbb{N} \\ W(0) &= b(0). \end{aligned} \quad (5.14)$$

where  $b(\bullet) = c + a(\bullet)$  and  $\rho = \frac{\bar{\lambda}}{\mu} > 0$ .

*Proof.* The equilibrium distribution under policy  $u(x+1)$  is given by

$$\Pi_n(x+1) = \rho^n \Pi_0(x+1), \quad 0 \leq n \leq x+1 \quad (5.15)$$

with

$$\Pi_0(x+1) = \frac{1-\rho}{1-\rho^{x+2}}, \quad \rho \neq 1. \quad (5.16)$$

and

$$\Pi_n(x+1) = \frac{1}{x+2}, \quad 0 \leq n \leq x+1, \quad \rho = 1. \quad (5.17)$$

When  $\rho \neq 1$ , it follows from Lemma 5.1, Lemma 5.2 and expressions (5.15) and (5.16) by using equilibrium equations presented above that we have that  $W = W(x)$  is the solution to

$$\begin{aligned} Cost_{u(x)}(W) &\equiv \left\{ \rho^x \bar{\lambda} W + \sum_{n=0}^{x-1} \rho^n \bar{\lambda} b(n) \right\} \frac{1-\rho}{1-\rho^{x+1}} \\ &= \left\{ \rho^{x+1} \bar{\lambda} W + \sum_{n=0}^x \rho^n \bar{\lambda} b(n) \right\} \frac{1-\rho}{1-\rho^{x+2}} \\ &\equiv Cost_{u(x+1)}(W), \quad x \in \mathbb{N} \end{aligned} \quad (5.18)$$

For  $x = 0$ , we have that  $W = W(0)$  is the solution to

$$\begin{aligned} \text{Cost}_{u(0)}(W) &\equiv \bar{\lambda}W \\ &= \frac{\rho}{1+\rho}\bar{\lambda}W + \frac{1}{1+\rho}\bar{\lambda}b(0) \\ &\equiv \text{Cost}_{u(1)}(W) \end{aligned} \quad (5.19)$$

Upon solving (5.18) and (5.19) for  $W$  it is straightforward to show that when  $\rho \neq 1$  the vendor-specific index  $W(x)$  in state  $x$  is expressed by

$$\begin{aligned} W(x) &= b(x) + \sum_{n=1}^x \rho^n b(x) - \sum_{n=1}^x \rho^n b(n-1), \quad x \in \mathbb{N} \\ W(0) &= b(0). \end{aligned} \quad (5.20)$$

When  $\rho = 1$ , an equivalent analysis yields

$$\begin{aligned} W(x) &= (x+1)b(x) - \sum_{n=1}^x b(n-1), \quad x \in \mathbb{N} \\ W(0) &= b(0). \end{aligned}$$

This completes the proof. □

Note that the first term on the right-hand side of (5.20) is the expected cost incurred when an incoming breakdown is sent to the vendor in state  $x$ . This could be thought of as the *direct* cost for admitting the breakdown and is sensibly an element of the fair charge for rejection. However, the breakdown also has impact on other subsequent incoming breakdowns which will be routed to the same vendor. Hence, the second and third term on the right-hand side of (5.20) may be thought of as estimating this impact. This may be interpreted as an *indirect* cost of admitting a breakdown to the vendor.

**Lemma 5.4** *The sequence  $\{W(x), x \in \mathbb{N}\}$  is increasing in  $x$ .*

*Proof.* Rewrite (5.20) as

$$W(x) = b(x) + \sum_{n=0}^{x-1} \rho^{n+1} \{b(x) - b(n)\} \quad (5.21)$$

From (5.21), it is not difficult to show the increment in the fair charge is given by

$$W(x) - W(x - 1) = [b(x) - b(x - 1)]\{1 + \rho + \rho^2 + \dots + \rho^x\}$$

which is positive for all  $x$  since the cost  $b(x)$  is increasing. This latter property of the sequence  $\{b(x), x \in \mathbb{N}\}$  has been ensured by Lemma 2.1. The conclusion then follows easily.  $\square$

It follows from Theorem 5.3 and Lemma 5.4 that the vendor-specific indices  $\{W(x), x \in \mathbb{N}\}$  form an increasing sequence. The following key result in Theorem 5.6 describes a strictly optimal policy for the problem (5.7) for any  $W \in (W(x - 1), W(x))$ . The proof of Theorem 5.6 will utilize the results in Lemmas 5.4 and 5.5.

**Lemma 5.5** (i) If  $W > W(x)$ , then  $Cost_{u(x)}(W) > Cost_{u(x+1)}(W)$ ;  
(ii) If  $W < W(x)$ , then  $Cost_{u(x)}(W) < Cost_{u(x+1)}(W)$ .

*Proof.* It has been shown that  $W = W(x)$  renders  $Cost_{u(x)}(W) = Cost_{u(x+1)}(W)$ . Note that only the first term on the right-hand side of (5.8) involves  $W$ . When  $W > W(x)$ , policy  $u(x)$  is no longer optimal for the vendor with the rejection penalty  $W$ . Because  $W - W(x)$  is now positive and  $\Pi_x(x) > \Pi_{x+1}(x + 1)$  (which is the result from Lemma 5.2), the conclusion in (i) follows. When  $W < W(x)$ , policy  $u(x + 1)$  is certainly no longer optimal for the vendor with the rejection penalty  $W$ . Because  $W - W(x)$  is now negative and  $\Pi_x(x) > \Pi_{x+1}(x + 1)$ , the result in (ii) follows too.  $\square$

**Theorem 5.6** If  $W(x - 1) < W < W(x)$ ,  $x \geq 1$ , policy  $u(x)$  is strictly optimal for the problem (5.7),  $x \geq 1$ .

*Proof.* Consider the first inequality  $W > W(x - 1)$ . The direct utilization of the result (i) in Lemma 5.5 together with Lemma 5.4 yields that:

$$\begin{aligned} W > W(x - 1) &\rightarrow Cost_{u(x-1)}(W) > Cost_{u(x)}(W) \\ W > W(x - 1) > W(x - 2) &\rightarrow Cost_{u(x-2)}(W) > Cost_{u(x-1)}(W) \\ &\vdots \\ W > W(0) &\rightarrow Cost_{u(0)}(W) > Cost_{u(1)}(W) \end{aligned} \tag{5.22}$$

It follows that  $Cost_{u(x)}(W) < Cost_{u(y)}(W)$  for any  $0 \leq y \leq x - 1$ .

Next consider the other inequality  $W < W(x)$ . By using the result (ii) in Lemma 5.5 and Lemma 5.4, we have that

$$\begin{aligned}
 W < W(x) &\rightarrow Cost_{u(x)}(W) < Cost_{u(x+1)}(W) \\
 W < W(x) < W(x+1) &\rightarrow Cost_{u(x+1)}(W) < Cost_{u(x+2)}(W) \\
 W < W(x+2) &\rightarrow Cost_{u(x+2)}(W) < Cost_{u(x+3)}(W) \\
 &\vdots \\
 & \hspace{15em} (5.23)
 \end{aligned}$$

From the set of inequalities presented in (5.23), we conclude that  $Cost_{u(x)}(W) < Cost_{u(y)}(W)$  for any  $y \geq x + 1$ .

Hence, it is straightforward to infer that policy  $u(x)$  must be strictly optimal for the problem (5.7) when  $W \in (W(x-1), W(x))$ . The proof is completed.  $\square$

The theoretical discussions above show that the system concerned is indeed *indexable* following the terminology given by Whittle (1988). The vendor-specific index  $W(x)$  does exist and is given by the expressions in the statement of Theorem 5.3.

Following Whittle (1988) the development of the indices in (5.14) using the restless bandit approach based on a single-vendor problem now yields an allocation policy for a full multi-vendor scenario in which each incoming breakdown is routed to any vendor with the smallest current fair charge. We restore the vendor subscript  $\nu$  from now on. Formally, suppose a broken item arrives for repair when the system is in state  $\mathbf{x} = \{(x_1, \dots, x_\nu) \in \mathbb{Z}^+ : x_\nu \geq 0, \sum_{\nu=1}^V x_\nu \leq K\}$ . The breakdown should then be routed to any vendor  $\nu^*$  which satisfies

$$W_{\nu^*}(x_{\nu^*}) = \min_{1 \leq \nu \leq V} W_\nu(x_\nu) \quad (5.24)$$

In the upcoming numerical study, the fixed breakdown rate  $\bar{\lambda}$  under the assumption of an infinite population approximation will be modified to  $\lambda(K - \sum_{\nu=1}^V x_\nu)$  in computing

the derived indices for our finite population problem.

Though no optimality is guaranteed by applying the derived indices to the dynamic equation (5.2), the following computational studies will show the effectiveness of the restless bandit heuristic developed for the dynamic allocation model.

## 5.4 Computational Study

The performance of the restless bandit heuristic developed in the previous section will be assessed comparatively in relation to other simpler routing heuristics as well as compared to the optimal solutions for small instances. Computational experiments will be conducted using both the value-iteration algorithm of dynamic programming and simulation models. The former algorithm generates optimal solutions to Markov decision problems, including the dynamic routing problem, in principle. However, the convergence sought by the value-iteration algorithm for problem instances where the state spaces are of high dimensions becomes computationally very expensive. Hence, we develop simulation models to estimate the cost rates resulting from the implementation of different heuristic policies for large values of  $K$  and  $V$ . Above all, the computational results reported in this section suggest the following principle findings:

1. Restless bandits index policies outperform optimal static solutions in all cases studied;
2. Expected costs incurred under the restless bandit heuristic are very close to optimal;
3. The comparative performance of the restless bandit heuristic declines marginally relative to that of two simpler heuristics, namely the greedy heuristic and the join-the-shortest-queue heuristic, as the system becomes increasingly congested;
4. The performance of the restless bandit heuristic improves compared with that of the two other heuristics as service rates across the vendors become more unequal;

5. Under our system scaled-up proposal, it does not matter hugely (in terms of system performance) whether it is achieved via an increase in the committed service capacities of existing vendors or by contracting with an additional vendor.

### 5.4.1 The application of the value-iteration algorithm

Firstly, we consider the value-iteration algorithm to determine solutions to the optimality equation (5.2). At each recursive step, this simple iterative method computes the value function  $V_t$  in terms of  $V_{t-1}$  using (5.2). For large  $t$ , the difference  $V_t(x) - V_{t-1}(x)$  gives an estimate of the minimal cost rate per unit time achievable by the system. The iteration will stop until the difference between the upper and lower bounds of these value function differences are less than a given tolerance error  $\epsilon$  for all states  $x$ . The general form of the value-iteration algorithm (see Tijms(1994)) adapted to our specific problem scenario is described as follows:

#### Value-iteration algorithm for computing the optimal cost rate

Step 0: Set  $V_0(x) = 0$  for all  $x \in S$ . Let  $t:=1$ ;

Step 1: Compute the value function  $V_t(x)$  by using  $V_{t-1}(x), x \in S$ , in (5.2);

Step 2: Compute the bounds

$$M_t = \max_{x \in S} \{V_t(x) - V_{t-1}(x)\} \text{ and } m_t = \min_{x \in S} \{V_t(x) - V_{t-1}(x)\}$$

The algorithm is stopped when  $0 \leq M_t - m_t \leq \epsilon m_t$  where  $\epsilon$  is a specified tolerance (e.g., 0.0001). The optimal cost rate is then estimated by  $0.5(M_t + m_t)$ .

Otherwise go to Step 3;

Step 3:  $t:=t+1$  and go to Step 1.

The number of iterations needed for the above algorithm to converge depends on the problem structure as well as the given parameters. In all the cases studied using the value-iteration algorithm, we have two service vendors ( $V = 2$ ) available to carry out

the repair work. The individual breakdown rate for a purchased item is 1.2 per year as before ( $\lambda = 1.2$ ). The cost model 3 is adopted in this context with  $c_v = 0$ ,  $1 \leq v \leq 4$ ,  $h = 1$ ,  $d = 10$  and the service time threshold is taken to be  $\tau = 0.04$  per year throughout the numerical studies of this chapter. Hence the goodwill costs initially grow linearly at the rate of one unit per year until the response time for a repair exceeds 10 working days. After that point, the goodwill costs still increase linearly but with a larger slope of ten units per year. This could be understood as a service delay penalty. The two vendors operate as single servers working at rates  $\mu_1$  and  $\mu_2$ . One of the service rates,  $\mu_1$  say, is randomly generated from the uniform interval  $[0.2TS, 0.8TS]$ , where  $TS$  stands for the total service rate. We then choose  $\mu_2 = TS - \mu_1$ .  $TS$  takes the values 140, 170, 200, 230, 260, 290, 320 and 350 for  $K = 100$ , and 170, 200, 230, 260, 290, 320, 350 and 380 for  $K = 200$ . Because these values chosen cover the full range of the system load (i.e., from low to high), the status of the restless bandit heuristic can be evaluated across a range of traffic levels.

Before we proceed to present the results for these parameter settings, it is important to show that the above value-iteration algorithm can also be used to estimate the cost rate associated with a specified policy. The value-iteration algorithm for a specified policy  $u$  is given by:

#### Value-iteration algorithm for computing the cost rate from a specified policy $u$

Step 0: Set  $V_0^u(\mathbf{x}) = 0$  for all  $\mathbf{x} \in S$ . Let  $t:=1$ ;

Step 1: Compute the value function  $V_t^u(\mathbf{x})$ ,  $\mathbf{x} \in S$  in terms of  $V_{t-1}^u(\mathbf{x})$  from

$$\begin{aligned}
 V_t^u(\mathbf{x}) &= \lambda \sum_{v=1}^V x_v V_{t-1}^u(\mathbf{x}) + \sum_{v=1}^V \mu_v V_{t-1}^u(\mathbf{x} - \hat{e}_v) \\
 &+ \lambda [K - \sum_{v=1}^V x_v] (c_{u(\mathbf{x})} + a_{u(\mathbf{x})}(\mathbf{x}) + V_{t-1}^u(\mathbf{x} + \hat{e}_{u(\mathbf{x})})) \quad (5.25)
 \end{aligned}$$

where  $V_t^u(\mathbf{x})$  is the expected cost incurred over the first  $t$  time epochs from time zero when the initial state is  $\mathbf{x}$  under policy  $u$ , and  $u(\mathbf{x})$  indicates the vendor chosen by policy  $u$  in state  $\mathbf{x}$ .

Step 2: Compute the bounds

$$M_t^u = \max_{x \in S} \{V_t^u(x) - V_{t-1}^u(x)\} \text{ and } m_t^u = \min_{x \in S} \{V_t^u(x) - V_{t-1}^u(x)\}$$

The algorithm is stopped when  $0 \leq M_t^u - m_t^u \leq \epsilon m_t^u$  where  $\epsilon$  is a prespecified tolerance (e.g., 0.0001), and hence the expected cost rate associated with policy  $u$  is estimated by  $0.5(M_t^u + m_t^u)$ . Otherwise go to Step 3;

Step 3:  $t:=t+1$  and go to Step 1.

Each row in Tables 5.1 and 5.2 corresponds to a given value of the total service rate  $TS$ . Each summarizes the average cost rates from 20 randomly generated combinations of  $\mu_1$  and  $\mu_2$ . For each value of  $TS$  across the two tables, we randomly generate 20 instances. For each of those, the cost rates are computed under OPI DYN, RB and OPI STA. The column heads in the two tables are explained as follows:

**OPI DYN** : These columns contain the averages of the optimal cost rates obtained from 20 randomly generated instances using the value-iteration algorithm described above;

**99% C. I.** : These columns contain half widths of 99% confidence intervals for the corresponding mean cost rates. This is so throughout the tables;

**RB** : These columns contain the averages of the cost rates obtained from 20 randomly generated instances using the second value-iteration algorithm described above under the restless bandit index policy derived in the previous section;

**OPI STA** : These columns contain the averages of the optimal static cost rates obtained from 20 randomly generated instances. See Chapter 3 for details;

**Max DIF%** : These columns contain the maximum percentage *difference* between the cost rate incurred by the optimal dynamic policy and that of the restless bandit heuristic among the 20 cases;

**Min DIF%** : These columns contain the minimum percentage *difference* between the cost rate incurred by the optimal dynamic policy and that of the restless bandit heuristic among the 20 cases.

<i>TS</i>	OPI DYN	99% C. I.	RB	99% C. I.	OPI STA	99% C. I.	Max DIF%	Min DIF%
140	18.394	±0.032	18.661	±0.121	32.564	±0.876	3.809	0.000
170	6.319	±0.018	6.409	±0.045	12.547	±0.548	4.342	0.000
200	3.075	±0.012	3.103	±0.022	5.864	±0.340	2.779	0.035
230	1.896	±0.010	1.904	±0.013	3.235	±0.220	2.042	0.010
260	1.360	±0.009	1.363	±0.009	2.060	±0.152	0.921	0.004
290	1.071	±0.009	1.072	±0.009	1.473	±0.111	0.696	0.001
320	0.892	±0.010	0.893	±0.010	1.148	±0.086	0.775	0.000
350	0.769	±0.012	0.769	±0.012	0.948	±0.071	0.026	0.000

Table 5.1: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 100$  and cost model 3. See above text for further details.

<i>TS</i>	OPI DYN	99% C. I.	RB	99% C. I.	OPI STA	99% C. I.	Max DIF%	Min DIF%
170	522.341	±0.019	522.342	±0.019	522.411	±0.028	0.001	0.000
200	264.958	±0.004	265.050	±0.049	277.015	±0.595	0.076	0.000
230	82.259	±0.021	82.775	±0.228	114.752	±1.301	1.124	0.000
260	23.539	±0.028	23.918	±0.190	46.120	±1.075	4.331	0.000
290	9.021	±0.029	9.174	±0.088	20.603	±0.683	4.428	0.000
320	4.713	±0.022	4.761	±0.037	10.541	±0.418	2.086	0.000
350	3.088	±0.016	3.104	±0.019	6.177	±0.264	0.908	0.000
380	2.324	±0.011	2.330	±0.011	4.096	±0.177	0.820	0.000

Table 5.2: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 200$  and cost model 3. See above text for further details.

Some important observations from Tables 5.1 and 5.2 are summarized below to shed light on the key findings mentioned at the beginning of this section, among which the first two are justified in particular.

1. As long as the total service rate is held fixed, variations around the mean cost rates are very small when the optimal dynamic and the restless bandit index poli-

cies are applied (see the associated 99% C.I. for the 20 cases studied). However, this is certainly not the case for the optimal static solution, whose cost rate is more sensitive to the choice of the individual service rates. Specifically, when values of the service rates are close to each other the optimal static cost rate exceeds that when they are quite different. See detailed data in Appendix B. This is consistent with the findings of Opp, Glazebrook and Kulkarni (2005) in which a GINI coefficient is used as a measure to demonstrate this point.

2. The restless bandit heuristic performs extremely well in comparison to the optimal dynamic solution throughout the above tables. In the worst case of more than 300 problems studied, the cost rate incurred by the restless bandit heuristic exceeds that of the optimal dynamic policy by just under 4.5%. When the distribution of the total service rates between the two vendors is near equal, the cost rates obtained by the restless bandit heuristic are very close to optimal. See data in Appendix B.
3. The results above include some interesting features because of the wide range of the values of  $TS$  chosen. When the system is overloaded (i.e. the system load measure  $\tilde{\rho} = \frac{\lambda K}{\mu_1 + \mu_2}$  is greater than 1), the difference between the optimal static cost rate and the optimal dynamic one is extremely small. See the first two rows in Table 5.2 for example. Consequently, the differences between the cost rates obtained under the restless bandit heuristic and the optimal dynamic policy are likely to be even smaller.
4. In most cases, there are very substantial savings by applying the near-optimal restless bandit heuristic instead of an optimal static policy. Such savings may be very worthwhile even taking account of the substantial administrative overhead involved in observing the queue lengths at all times.

### 5.4.2 Simulation study

The simulation model developed for the purpose of evaluating the performance of the restless bandit heuristic against two other simpler heuristics has the flexibility and efficiency to study a general system with more than two service vendors and a large

population size.

In the first part of the simulation study, we are interested in investigating the impact of different levels of the total system load on the outsourcing warranty problem, in particular, the performance of the restless bandit heuristic developed. As suggested by Ziya et al. (forthcoming), we consider four service vendors ( $V = 4$ ) with service rates  $\mu_v, 1 \leq v \leq 4$  drawn independently from a uniform distribution on  $[0, K]$ , where  $K = 500$ . The individual failure rate  $\lambda$  is drawn from a  $U[0, 4]$  distribution. We discard instances for which the randomly generated values of the system load ( $\tilde{\rho} = \lambda K / \sum_{v=1}^4 \mu_v$ ) lies outside of the range  $[0, 1]$ . We consider a categorization of problem instances into three different traffic levels in terms of the system load  $\tilde{\rho}$ :

1. Light traffic:  $0 \leq \tilde{\rho} \leq 0.6$ ;
2. Medium traffic:  $0.6 < \tilde{\rho} \leq 0.8$ ;
3. Heavy traffic:  $0.8 < \tilde{\rho} \leq 1$ .

Under the above three different traffic levels, we assess the performance of the alternative dynamic heuristics, namely the restless bandit (RB) index policy which has been developed in Section 5.3, the individually optimal (IO) heuristic and the join-the-shortest-queue (JSQ) heuristic. The latter two heuristics are described as follows.

**IO** : The individually optimal policy is a standard proposal for dynamic routing problems. It is also called the greedy heuristic in some texts. See Ziya et al. (forthcoming). Under this proposal, an incoming failure is routed to any service vendor  $v$  such that the expected cost associated with this breakdown alone is smallest. For each vendor  $v$ , the expected cost of a particular item is computed on the basis of the queue length  $x_v, 1 \leq v \leq 4$ . Therefore, the IO policy in this context sends the failed item to the vendor with the minimal associated cost  $b_v(x_v) = c_v + a_v(x_v), 1 \leq v \leq 4$ .

**JSQ** : This is also a conventional approach adopted widely in the queueing literature. It routes an incoming failure to any vendor  $v$  with the shortest queue length  $x_v$ . Any ties occurring when the shortest queue length occurs at more than one

vendor will be broken by sending the item to the candidate vendor with highest service rate.

Cost model 3 is assumed with holding cost  $h = 1$ , long delay penalty  $d = 10$  and fixed repair cost  $c_v = 0, 1 \leq v \leq 4$ . The time threshold  $\tau$  which determines whether the holding cost or the long delay penalty is applied is set to be 10 working days. For each traffic level (light, medium and heavy), we randomly generate 500 instances for each of which the cost rates are computed under three alternative heuristics. We evaluate their performance and present results in the following Tables 5.3 to 5.5. We summarize the 500 outcomes in each case by the lower quartile, median and upper quartile. We show the percentage deviation of the cost rates obtained under the IO and JSQ heuristics from those of the RB heuristic via a 99% confidence interval. We also give the number of cases for which the corresponding heuristic produces the smallest cost rate among the three. Note that the numbers of instances in the last columns in the tables add up to a value more than 500 due to ties among the heuristics. Tables 5.3, 5.4 and 5.5 report these results for light, medium and heavy traffic levels respectively.

Heuristic	Lower quartile	Median	Upper quartile	99% C.I. for the mean deviation %	Best heuristics in 500 cases
RB	1.772	2.105	2.818	0.000	355
IO	1.771	2.126	2.895	$2.106 \pm 0.537$	162
JSQ	2.130	2.961	6.555	$120.745 \pm 26.943$	1

Table 5.3: Comparative performance of three alternative heuristics under light traffic for cost Model 3. See above text for further details.

Heuristic	Lower quartile	Median	Upper quartile	99% C.I. for the mean deviation %	Best heuristics in 500 cases
RB	1.840	2.316	3.330	0.000	317
IO	1.858	2.350	3.451	$2.164 \pm 0.655$	187
JSQ	2.267	3.499	8.745	$127.671 \pm 29.448$	6

Table 5.4: Comparative performance of three alternative heuristics under medium traffic for cost Model 3. See above text for further details.

Heuristic	Lower quartile	Median	Upper quartile	99% C.I. for the mean deviation %	Best heuristic in 500 cases
RB	1.919	2.366	3.476	0.000	297
IO	1.919	2.383	3.558	$2.174 \pm 0.894$	207
JSQ	2.308	3.569	10.056	$134.945 \pm 30.408$	1

Table 5.5: Comparative performance of three alternative heuristics under high traffic for cost Model 3. See above text for further details.

The results from Tables 5.3, 5.4 and 5.5 demonstrate that the overall performance of RB is statistically better than IO and JSQ though it deteriorates modestly as the traffic level increases relative to the performance of IO. But, even when the best heuristic is not RB in the 500 instances studied, it is usually the case that the cost rate of RB is very close to that of the best heuristic among the three. Clearly, the performance of JSQ is consistently weakest throughout the tables. Note that the mean deviation of the cost rate of IO from RB is not very sensitive to the traffic levels. Instead, it is rather stable around 2.2% throughout. The small advantage over the individual optimal heuristic stimulates the second part of the simulation study in which the superior status of the restless bandit heuristic will be more clearly articulated. Though the optimal dynamic solutions are not available for comparison with these three heuristics because of the computational difficulty of these large problems (e.g.,  $V = 4$  and  $K = 500$ ), we find that the worst scenarios in which the cost rates under the RB departed furthest from optimal are those for which the system load  $\tilde{\rho}$  ranges from 0.7 to 1. See Tables 5.1

and 5.2 for examples. We suspect that the similar conclusion would be obtained for scaled-up problem instances in this part of the simulation study.

We now proceed to present the second set of experimental results. In these we scrutinize the influence of the distribution of service rates across the vendors on the dynamic allocation model. To measure the degree of inequality, we introduce the *Gini coefficient*. This was used by Opp, Glazebrook and Kulkarni (2006) as a measure of inequality of optimal static allocations across vendors. It is also commonly used in economics as a measure of inequality in a population (see Glasser (1962) and Sen (1973)). In this context, the Gini coefficient, denoted  $G$  when used to measure inequality of the service rates among vendors is calculated by using the following formula:

$$G = \frac{\sum_{v=1}^V \sum_{j=v+1}^V |\mu_v - \mu_j|}{V \sum_{v=1}^V \mu_v} \quad (5.26)$$

Plainly, we have  $G = 0$  when all the service rates are equal and  $G = (V - 1)/V$  when all the service rates but one are zero. For  $V = 4$ , the value of  $G$  ranges from 0 to 0.75. We categorize the problems studied into three groups in terms of the degree of inequality determined by the value of  $G$  as follows:

1. Low inequality:  $0 \leq G < 0.25$ ;
2. Medium inequality:  $0.25 \leq G < 0.5$ ;
3. High inequality:  $0.5 \leq G \leq 0.75$ .

Again, the service rates  $\mu_v, 1 \leq v \leq 4$  are drawn from a uniform distribution  $[0, K]$ , where  $K = 500$  for this part of the simulation study. The individual failure rate  $\lambda$  is determined by the total service rate, a given system load ( $\tilde{\rho} = 0.7$ ) and the population parameter  $K$ . It is given by

$$\lambda = \frac{0.7 \sum_{v=1}^V \mu_v}{K}$$

The choice of other parameters are the same as in the first part of the simulation study. Similarly, for each category of heterogeneity level of service rates we randomly generate 500 cases. For each of those the cost rates are computed under the three heuristics.

We describe their performance in the following Tables 5.6 to 5.8 where the cost rates in each set of 500 problem instances is summarized by their lower quartile, median and upper quartile. We describe the average percentage deviation of the cost rates from that of RB via a 99% confidence interval. We also give the number of cases for which the corresponding heuristic produces the smallest cost rate among the three in the last columns in the tables, where the number of instances add up to a value more than 500 due to ties among the heuristics. Tables 5.6, 5.7 and 5.8 report results for low, medium and high inequality among the service rates respectively.

Heuristic	Lower quartile	Median	Upper quartile	99% C.I. for the mean deviation %	Best heuristic in 500 cases
RB	1.791	2.074	2.694	0.000	335
IO	1.807	2.096	2.756	$1.258 \pm 0.330$	172
JSQ	2.101	2.519	3.484	$29.241 \pm 6.022$	4

Table 5.6: Comparative performance of three alternative heuristics for low inequality of the service rates for cost Model 3. See above text for further details.

Heuristic	Lower quartile	Median	Upper quartile	99% C.I. for the mean deviation %	Best heuristic in 500 cases
RB	2.465	3.371	6.419	0.000	330
IO	2.475	3.477	6.996	$4.703 \pm 1.066$	161
JSQ	6.042	12.943	34.795	$269.183 \pm 37.425$	10

Table 5.7: Comparative performance of three alternative heuristics for medium inequality of the service rates for cost Model 3. See above text for further details.

Heuristic	Lower quartile	Median	Upper quartile	99% C.I. for the mean deviation %	Best heuristic in 500 cases
RB	7.140	15.696	196.963	0.000	329
IO	7.740	19.534	229.427	11.027±1.703	135
JSQ	74.666	112.634	315.661	611.046±76.531	36

Table 5.8: Comparative performance of three alternative heuristics for high inequality of the service rates for cost Model 3. See above text for further details.

From the results in Tables 5.6 to 5.8, it is of note that the more heterogeneous are the service rates across the four vendors, the better the RB policy performs against the IO and the JSQ policies. Specifically, the average percentage deviation between the cost rate under the IO policy and that of RB increases from 1.3 % to 11% as the degree of service rate inequality increases from low to high. The number of occasions that the IO yields the smallest cost estimates among the three also deteriorates in a heterogeneous scenario. The performance of the JSQ heuristic is consistently the weakest. The average percentage difference of its cost rate from that of RB increases more dramatically as the the degree of service rate inequality increases. The finding that the RB heuristic is increasingly superior over the other two simpler heuristics in a heterogeneous context is of great interest.

In the third part of the simulation study, we are concerned about an important question arising from business practice relating to this outsourcing warranty problem. Specifically, the question is addressed regarding how to scale up the current range of service rates of the outsourced contractors to meet increased demand. We consider three options, namely, increasing the service capacity of the most capable vendor only, increasing the service rate of the least capable vendor only or outsourcing to an additional service vendor to meet the increased demand for warranty repairs. On the basis of the numerical study that we have presented above, we are confident of the performance of the restless bandit heuristic for the dynamic allocation problem. Hence, the following study of the scaled-up process for the current system is guided by the cost rates obtained under the RB heuristic. We consider a current system in which the individual

breakdown rate  $\lambda$  is fixed at 1.2 per year. The total service rate for the four service vendor sums to 750, which makes the overall system load  $\tilde{\rho} = 0.8$ . The distribution of the service rates among the four service vendors is given by

$$\mu_1 = 296.092, \mu_2 = 207.264, \mu_3 = 145.085, \mu_4 = 101.650. \quad (5.27)$$

With the service rates provided, we can calculate the Gini coefficient for this problem setting by using (5.26). The resulting  $G$  is 0.2. The choices of other parameters are as in the first part of the simulation study. The options to scale up the current problem are described as follows:

**Option 1** : Starting with four service vendors and a warranty population of 500, we increase the population size by 100, 200, 500 and 1000. We search for the additional service rate needed from vendor 1 so that the average cost rate for individual customers remains unchanged;

**Option 2** : Starting with the same setting, we search for the additional service capacity required from vendor 4 when the population size is increased to 600, 700, 1000 and 1500 to obtain the same average cost rate for individual customers as was the case for  $K = 500$ ;

**Option 3** : Again starting with the same setting, now the manufacturer is allowed to contract with an additional vendor (say, vendor 5 with service rate  $\mu_5$ ) to provide repair work for an increased population size. We search for the exact service capacity required for  $K=600, 700, 1000$  and 1500 so that the average cost rate per customer is as in the  $K = 500$  case.

This experimental study aims to find the option for which the least additional service rate in total is required to cope with the expanded warranty population sizes while sustaining the same service level (i.e. the expected cost rate for an individual item remains unchanged). We compare the above three options to provide useful information for the decision-making of a manufacturer in a similar business situation.

For this experimental study, we apply a traditional *bisection search* method (see Burden and Faires (2000)) to look for the service rate of the vendor concerned to cope with

the increased warranty population . Technically, an interval wide enough to include the sought value of the service rate is set at the beginning of the simulation. At each step, we divide the interval or the subsequent subinterval into half to shrink the search range. We also allow a tolerance error relating to the cost rates targeted and the cost rates estimated for any scaled-up problems. The degree of the closeness between those two cost rates together with the width of the initial search interval of the service rate determine the number of steps needed to estimate the additional service rate required for each option. In the simulation, we set the tolerance error to be 0.001. The time length is 250 years including the burn-in period of 100 years. We also apply 100 independent runs to yield better estimates for the targeted cost rate.

The first row in Table 5.9 presents the current problem setting with  $K = 500$ . From the second row onwards, the result presented corresponds to a single option (which is numbered beside the population size) to cope with the increased population. The columns heads in Table 5.9 are explained as follows:

**RB** : the associated average cost rate obtained by the RB heuristic under the setting of the service rates in the same row. From the second row onwards, the cost rate of the RB heuristic is calculated on the basis of the current simulated cost rate 6.29 for  $K = 500$ , for which the average cost rate incurred by an individual item is  $6.29/500$ . Hence, for example, the cost rate 7.55 for  $K = 600$  is calculated as  $600 * 6.29/500$ .

**Percentage of long waits** : the proportion of failed items waiting longer than the time threshold  $\tau$  expressed as a percentage. We intend to control this measure regarding the system efficiency under 2% in the scaled-up problems.

**Total service rate** : the sum of the service rates from all the vendors available.

$\mu_1, \mu_4$  and  $\mu_5$  : The corresponding service rates estimated for vendor 1, vendor 4 and additional vendor 5. They are estimates obtained from the simulation.

$K$	RB	Percentage of long waits	Total service rate	$\mu_1$	$\mu_4$	$\mu_5$
500	6.29	1.13	750.00	296.09	101.56	0.00
600(1 <sup>st</sup> )	7.55	1.14	858.73	404.82	101.56	0.00
600(2 <sup>nd</sup> )	7.55	1.18	856.32	296.09	207.88	0.00
600(3 <sup>rd</sup> )	7.55	1.10	882.90	296.09	101.56	132.90
700(1 <sup>st</sup> )	8.81	1.13	969.62	515.71	101.56	0.00
700(2 <sup>nd</sup> )	8.81	1.17	966.85	296.09	318.41	0.00
700(3 <sup>rd</sup> )	8.81	1.10	989.02	296.09	101.56	239.02
1000(1 <sup>st</sup> )	12.58	1.16	1308.64	854.73	101.56	0.00
1000(2 <sup>nd</sup> )	12.58	1.16	1306.94	296.09	658.50	0.00
1000(3 <sup>rd</sup> )	12.58	1.21	1325.83	296.09	101.56	575.83
1500(1 <sup>st</sup> )	18.88	1.05	1885.79	1431.88	101.56	0.00
1500(2 <sup>nd</sup> )	18.88	1.08	1883.47	296.09	1235.03	0.00
1500(3 <sup>rd</sup> )	18.88	1.18	1898.25	296.09	101.56	1148.25

Table 5.9: Problem scaled-up from  $K = 500$  under the three options available for cost Model 3. See above text for further details.

The results in Table 5.9 show that the second option is consistently the best in terms of the least additional service rate required to meet the increased demand of warranty repairs to approach a given cost rate under the RB heuristic. The difference between the additional service rates required by the first and the second options is not large. Nevertheless, the experimental study delivers a clear message that outsourcing to an additional service vendor is the least preferred option in this scenario, though the difference between the three options becomes less significant as  $K$  increases from  $K = 600$  to 1500. This simplifies matters for the manufacturer in practice because the way by which we need to scale up the problem is not very sensitive to the options available. It implies that the manufacturer could feel free to choose any option which is the most convenient for them to scale up the current outsourcing problem without too much consideration of the difference between the amounts of additional service rates required at alternative vendors to cope with the enlarged warranty population.

Note that in Table 5.9 the additional service rate required for each additional 100 in the warranty population is steadily and slowly increasing in  $K$  (e.g. an additional 109 service rate is needed for the additional 100 warranty population when  $K$  is increased to 600, while an average additional 114 service rate is required per additional 100 in the warranty population increased when  $K$  is scaled up to 1500 from the same starting point  $K = 500$ ). However, these conclusions are limited by the number of instances studied. We are not sure, for example, whether the service rate required over each additional 100 warranty population will be increasing indefinitely or if it will reach a limit or even decrease at some point.

With the information obtained from the scaled-up experiments, we proceed to present six profiles of service rates distribution when the warranty population  $K$  is scaled up to 10,000 from 500. As before, we have four service vendors available to undertake the repairs. The total service rate is calculated such that  $750 + 118 * 95 = 11,960$ . We would like to see if the observations above are extendable beyond the results presented in Table 5.9. 118 is chosen to be the average additional service rate required for each additional 100 in the warranty population increased from  $K = 500$  based on some trial studies. Specifically, the service rate for vendor  $v$  in profile  $j$  is given by

$$\mu_{v1} = 2,990, \quad 1 \leq v \leq 4$$

and

$$\mu_{vj} = 11,960 y_j^{v-1} (1 - y_j) (1 - y_j^4)^{-1}, \quad 1 \leq v \leq 4, 2 \leq j \leq 6$$

where  $y_j = 1 - 0.1(j - 1)$ ,  $2 \leq j \leq 6$ . The choice of the other parameters remains unchanged for the results in Table 5.10.

	cost rate	$G$	$\mu_1$	$\mu_2$	$\mu_3$	$\mu_4$
RB	94.07					
IO	94.07	0.00	2990.00	2990.00	2990.00	2990.00
JSQ	94.07					
RB	94.34					
IO	95.20	0.07	3477.76	3129.98	2816.98	2535.28
JSQ	95.11					
RB	94.03					
IO	94.53	0.14	4051.49	3241.19	2592.95	2074.36
JSQ	95.07					
RB	93.99					
IO	94.79	0.22	4721.67	3305.17	2313.62	1619.53
JSQ	96.50					
RB	94.82					
IO	95.47	0.30	5496.32	3297.79	1978.68	1187.21
JSQ	106.16					
RB	94.79					
IO	95.87	0.38	6378.67	3189.33	1594.67	797.33
JSQ	134.94					

Table 5.10: Comparative performance of three alternative heuristics for cost Model 3 and  $K = 10,000$ . See above text for further details.

The results in Table 5.10 confirm that the overall performance of RB remains better than that of the other two heuristics when the size of the warranty population is scaled up to 10,000. The benefit of applying the restless bandit heuristic to the problem considered becomes more evident as the Gini coefficient  $G$  measuring the degree of inequality in the service rates increases from 0 to 0.38. One interesting finding is that the RB cost rate for an individual item in Table 5.10 is less than that for  $K = 500$  (e.g.  $0.0095 = 95/10000 < 0.0126 = 6.29/500$ ) under our scaled-up proposal. However, in contrast, we note that the overall system load  $\tilde{\rho}$  increases from 0.8 ( $=1.2*500/750$ ) to 1.003 ( $=1.2*10,000/11960$ ). This shows that the scaled-up process for this outsourc-

ing warranty problem is a subtle issue and further investigations are required for this aspect in our future research.

## 5.5 Conclusions to the Chapter

In this chapter, we have considered the problem of dynamically allocating real-time breakdowns to alternative service vendors such that the overall cost rate incurred is minimized. The system is modelled as a Markovian decision process. However, the resulting dynamic programming problem is not solvable for problems of practical size due to the large dimension of the state space. Therefore, an index-based heuristic is established by deploying Whittle's restless bandit approach. An important achievement of this chapter is the theoretical proof to show that our system is indeed *indexable*. Through an extensive computational study we further demonstrated that

- The RB heuristic is very close to optimal. For small problems, the cost rates of RB differ by just 0.65% on average from those obtained under the optimal dynamic policy.
- In comparison to the static allocation, the application of the RB heuristic secures significant cost savings except for cases in which the system is heavily loaded (i.e.  $\tilde{\rho} > 1$ ) where the difference between the cost rates under the optimal dynamic and the optimal static policies is small.
- The RB heuristic consistently outperforms the IO and the JSQ heuristics, particularly in cases where the Gini coefficient measuring the inequality between the service rates is greater than 0.25. Put another way, the more heterogeneity in the service rates of the vendors, the better is the comparative performance of the RB heuristic against the IO and the JSQ heuristics.
- Should an increase in the total service rate be needed, our simulation evidence suggest that it does not matter significantly (in terms of system performance) whether this is achieved via an increase in the committed service capacities of existing vendors or by contracting with an additional vendor.

## Chapter 6

# Conclusions and Future Research

In this chapter, we summarize the findings from our systematic study in response to a range of key research questions, which have been articulated in the introductory chapter. We further present implications of this study of academic interest and managerial relevance. Finally, we are aware that the investigation concerning the allocation problem of outsourcing warranty repairs in this thesis has not been fully accomplished, which provides space for further extensions.

### 6.1 Answers to the research questions of interest

*1. What level of service capacity among the contracted vendors needs to be available to meet the anticipated demand for the manufacturer's post-sales repair service effectively?*

This research question is closely related with the second one below. Plainly, the level of service capacity needed is dependent on the degree of system congestion which is often implicitly set by the manufacturer in its *Code of Practice*, given that the anticipated warranty repairs are optimally/effectively distributed among the service vendors. For instance, the Code of Practice may guarantee that no more than 2% of customers should experience a response time longer than two weeks.

Besides the factor of system load, our analyses within the static allocation framework suggest that the total service capacity for any average cost rate incurred is less needed as the distribution of the service capacity among a collection of vendors becomes more unequal. Put in another way, the average cost rate achievable for any total service capacity may be well estimated (from any of good allocation heuristics) by considering the case when the service rates are equal across all the vendors.

In contrast, the study from the dynamic allocation model indicates that the total service capacity for any average cost rate achievable is remarkably insensitive to how it is divided between the vendors. The results from the system scaled-up simulation using the dynamic allocation model further confirms the following: should an increase in the total service rate be needed, it does not matter hugely (in terms of system performance) whether this is attained via an increase in the service rates of existing vendors or by contracting with an additional vendor.

*2. Given that the manufacturer's currently contracted vendors do possess sufficient service capacity, how should the repair work be best distributed among them?*

The second research question is the main theme of this thesis. It has been shown that the special cost structure and the high dimensionality of the state space make the development of an optimal policy to the distribution of the workload very difficult. Consequently, the utilisation of effective and near-optimal methods appears to be the most realistic way to tackle such allocation problems. As a result, we develop three effective heuristics for our three allocation models, namely, the greedy heuristic for the simplest static allocation model with a fixed warranty population, the DP policy improvement heuristic for the non-standard allocation model with a dynamic warranty population, and the restless bandit heuristic for the dynamic allocation model. Both theoretical and numerical evidence supports the claim that the overall cost rates obtained from the three heuristics we developed are either optimal or near-optimal. The following paragraphs summarize the findings under the three modelling frameworks

respectively.

In the static allocation model with a fixed warranty population, the proposed greedy heuristic achieves a workload allocation between the vendors which minimizes the overall cost rate for almost every circumstance. We indicate that such a strong performance of the greedy heuristic is dependent on the increasing and convex properties of the goodwill costs (which are deemed to predominate). A range of sensitivity analyses with regards to the key parameters has been undertaken to demonstrate the followings:

- The optimal static allocations are sensitive to the distribution of service capacity among the vendors and suggest that the able vendor should do more work.
- The optimal static allocations are insensitive to the values of the goodwill cost, which are practically difficult to determine.
- The optimal static allocations are insensitive to modest variations of the physical repair costs among vendors.

Last but not least (this is also of relevance to the first question), the heterogeneity among the service vendors, given that the total service capacity is fixed, leads to cost reductions for the manufacturer within the static allocation framework.

In the non-standard allocation model for a variable warranty population, standard methods are not available to seek an optimal solution to the repair workload allocation among the multiple vendors. This is because the stochastic dynamic optimization problem formulated is merely partially observed and has a system state of high and varying dimensions. We develop an effective dynamic heuristic by the adoption of an approximate DP approach based on policy improvement. Although it emerged that the DP policy improvement heuristic we develop is not highly superior to a simpler dynamic greedy one in terms of the overall cost performance, the DP policy improvement heuristic does impose greater regularity on the static allocations. Put in another way, such a dynamic heuristic reduces variability and hence lessens the chance of excessive queue lengths for repair. The simply calculated values of certain quantities (denoted  $\bar{G}(\eta\Omega\beta)$  and  $G(p^*)$  as lower and upper bounds) give a good prior indication of the scope of the cost reduction achievable by strong performing dynamic heuristics.

In the dynamic allocation model, for similar reasons we do not pursue direct application of stochastic dynamic programming to produce optimal policies. We instead deploy the restless bandit approach in deriving policies which utilise the information on the repair queue lengths. It emerged that the index policies derived perform outstandingly well and are very close to optimal. By application of the value-iteration algorithm, the cost rates obtained by the restless bandit approach are just 0.65% on average above optimal in the 320 instances studied. Extensive simulation studies investigate large problem sets in further assessing the performance of the restless bandit heuristic across various scenarios. Computational evidence suggests that the index-based heuristic policies consistently outperform the empirical suboptimal routing heuristics, whose performances are particularly poor in the problem instances for which the service rates of the multiple vendors are highly heterogeneously distributed.

*3. How much might the manufacturer have saved by employing a central decision controller to determine to which vendor a warranty repair is sent every time an item fails rather than preassigning them to the service vendors when they are purchased?*

We have seen that the savings generated by the dynamic solutions over the static ones are substantial throughout the computational studies. An average 65%<sup>1</sup> reduction in overall costs under the dynamic allocation model may render the administrative overhead for continuously observing the repair queue lengths negligible. We further show that the employment of the restless bandit approach in the dynamic allocation model also yields a similar degree of cost savings over the corresponding static solutions. However, previously related research for the outsourcing warranty repairs problem did not produce such impressive cost improvement from the utilization of the information regarding the repair queues (see the work of Opp et al. (2005)). This might be a consequence of the questionable linear holding costs model used in their papers.

Note that the cost reduction between the optimal dynamic and the optimal static is most significant when the service capacity is equally distributed across the vendors

---

<sup>1</sup>based on the results from the 320 instances randomly generated which is included in Appendix B.

and the system load is moderate or low, while less noticeable when the system is either overloaded or of high heterogeneity.

### 6.2.2 Recommendations of managerial relevance

*4. How much might the manufacturer be losing (economically and in customer goodwill) by maintaining an existing suboptimal approach to workload distribution?*

In cutting costs, the existing suboptimal approaches (such as the JSQ/SMA) to workload distributions have been shown to be seriously inferior at various levels compared to the near-optimal heuristics developed throughout this thesis. Nevertheless, a good lesson is learned from the substantial additional costs incurred when applying those suboptimal heuristics to vendors with heterogeneous service rates. This powerfully illustrates the importance of developing near-optimal heuristics to facilitate decision-making for the manufacturer in this outsourcing warranty repairs problem.

## 6.2 Implications and recommendations of the study

### 6.2.1 Implications of academic interest

In summary, the problems of the outsourcing of warranty repairs studied in this thesis belong to the regimes of stochastic optimization and queueing control. We develop models and heuristics to support decisions concerning how responsibility for the repair workload should be distributed among a collection of external service vendors. The effectiveness of the heuristic policies developed from alternative near-optimal approaches (i.e. Whittle's restless bandit, the DP policy improvement and the greedy algorithm) has been demonstrated consistently in the theoretical discussions and evaluated extensively through the computational studies. It is true that the heuristic policies derived do not always have simple explicit forms for the allocation models we consider. Nevertheless, theoretical evidence has indicated that they do possess sensible properties, such as monotonicity and convexity. Above all, the unique formulations of the cost models which penalize the delays experienced by individual customers have

not been seen in previously published papers to date to our knowledge.

### 6.2.2 Recommendations of managerial relevance

Based on the analyses undertaken in this thesis, we are able to provide the manufacturer guidance on how the workload should be divided in a way which appropriately respects the (possibly diverse) capabilities and cost characteristics of the vendors. As mentioned in Chapter 1, there are business concerns primarily regarding cost reduction, service quality and customer satisfaction in the transition of such a function to the external service providers. To some extent, these concerns are reflected and resolved in our model formulations. We develop simple allocation models for the problems which accommodate different scenarios but share a goal of the minimization of the average cost (economically and in customer goodwill). For a manufacturer facing to a choice concerning the allocation models, we make the following practical recommendations:

1. Calculating the system load and the Gini coefficient for measuring the heterogeneity of the vendors;
2. If the system is over congested, one should explore the possibility of increasing the service capacity from the existing vendors first, or by contracting with a new vendor. Before the values of the system load are back in a reasonable range, the manufacturer should use the static allocation model, which is simple and cheap (no central call center is needed) and as good as the dynamic one in such cases.
3. If the heterogeneity of the vendors is very high (i.e. the Gini coefficient is approaching  $(V - 1)/V$ ), there is also little advantage to apply dynamic routing. Under such situation, we also recommend to use the static allocation model;
4. Other than the rare cases described in 2 and 3 above, using dynamic allocation model shall make substantial savings by implementing the effective restless bandit heuristic developed even at the cost of administrative overheads to make the information of the queue lengths available.

### 6.3 Future research

This thesis that studies allocation models regarding the outsourcing of warranty repairs is self-contained. Nevertheless, further extensions are plausible and suggested in the following directions:

First, the dynamic allocation model could be extended to allow a variable warranty population. We believe that the modifications are not substantial for the system concerned and the associated index-based heuristic derived for this extension. It will be of interest to conduct an equivalent set of computational studies as we did for a fixed warranty population in this thesis.

Second, customers may be offered different warranty lengths upon product purchase and expired warranties could be renewed upon requests from customers. To accommodate this aspect, more than one type of fixed-length warranties need be considered in future research to yield the flexibility and generality of the allocation models developed. One could begin with the modification to include just two different lengths of warranties in the problem formulation. Methodologically, this extension only impacts the variable allocation model.

Third, it has been indicated that the DP policy improvement approach which makes use of previous allocations at the vendors does not always outperform the simple dynamic greedy heuristic. The underlying reasons may primarily lie in the Poisson assumption with a constant rate concerning arriving orders. Hence, the arrivals process can be generalised to the non-homogeneous Poisson with substantial fluctuations in the arrival rate. Seasonal effects or product life cycles could be perfect justifications for this modification. Meanwhile, the service rates of the vendor can be allowed to vary by the volume of the committed workload to improve the system efficiency.

# Appendix A

## Detailed Results from the Value Iteration Algorithm

	OPI DYN	RB	dif%	IO	dif%	JSQ	dif%	OPI STA	$\nu_1$	$\nu_2$	$\mu_1$	$\mu_2$
1	18.420	18.441	0.115	18.422	0.012	18.422	0.014	34.279	45	55	63.982	76.018
2	18.408	18.422	0.072	18.414	0.030	18.417	0.048	34.222	44	56	62.884	77.116
3	18.385	18.601	1.175	20.361	10.746	23.951	30.275	30.973	79	21	106.751	33.249
4	18.388	18.527	0.757	18.593	1.113	18.880	2.677	33.695	63	37	86.478	53.522
5	18.372	18.635	1.432	20.646	12.379	24.126	31.317	30.681	20	80	32.867	107.133
6	18.376	18.451	0.406	18.517	0.766	18.760	2.087	33.785	38	62	54.825	85.175
7	18.400	18.668	1.452	19.516	6.062	23.287	26.559	31.251	22	78	34.795	105.205
8	18.303	19.000	3.809	20.059	9.597	25.576	39.735	30.053	82	18	110.019	29.981
9	18.392	18.774	2.078	19.841	7.877	21.363	16.155	32.250	74	26	99.673	40.327
10	18.380	18.856	2.593	19.217	4.553	20.792	13.123	32.452	28	72	42.425	97.575
11	18.576	18.576	0.000	18.576	0.000	18.576	0.000	34.310	49	51	68.702	71.298
12	18.394	18.679	1.553	20.557	11.761	23.580	28.197	30.973	21	79	34.096	105.904
13	18.410	18.776	1.990	19.607	6.502	22.858	24.165	31.517	77	23	104.124	35.876
14	18.389	18.523	0.728	18.588	1.077	18.888	2.713	33.695	63	37	86.555	53.445
15	18.405	18.750	1.879	19.722	7.160	22.625	22.932	31.517	23	77	36.495	103.505
16	18.348	18.947	3.267	18.824	2.593	20.482	11.633	32.633	29	71	43.695	96.305
17	18.383	18.846	2.515	19.201	4.448	20.819	13.249	32.452	72	28	97.681	42.319
18	18.390	18.419	0.158	18.612	1.203	18.580	1.033	33.948	40	60	57.327	82.673
19	18.356	18.906	2.996	18.821	2.533	20.577	12.101	32.633	71	29	96.706	43.294
20	18.395	18.427	0.176	18.629	1.272	18.565	0.926	33.956	40	60	57.600	82.400

Table A.1: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 100$ , cost model 3 and  $TS = 140$  in Table 5.1 See text in Chapter 5 for further details.

	OPI DYN	RB	dif%	IO	dif%	JSQ	dif%	OPI STA	$v_1$	$v_2$	$\mu_1$	$\mu_2$
1	6.324	6.330	0.085	6.324	0.001	6.325	0.011	13.607	45	55	77.692	92.308
2	6.318	6.331	0.208	6.318	0.000	6.321	0.046	13.589	44	56	76.359	93.641
3	6.294	6.372	1.232	7.189	14.208	9.687	53.896	11.551	83	17	129.626	40.374
4	6.311	6.342	0.503	6.347	0.574	6.584	4.329	13.246	65	35	105.008	64.992
5	6.295	6.568	4.342	7.157	13.703	9.798	55.652	11.365	16	84	39.910	130.090
6	6.315	6.348	0.533	6.332	0.283	6.514	3.152	13.307	36	64	66.573	103.427
7	6.304	6.415	1.774	7.258	15.134	9.266	46.991	11.731	18	82	42.251	127.749
8	6.302	6.459	2.489	6.915	9.728	10.732	70.307	10.970	87	13	133.594	36.406
9	6.342	6.431	1.402	6.934	9.333	8.070	27.240	12.337	77	23	121.031	48.969
10	6.325	6.467	2.244	7.121	12.591	7.722	22.090	12.474	25	75	51.516	118.484
11	6.424	6.424	0.000	6.424	0.000	6.424	0.000	13.652	49	51	83.423	86.577
12	6.297	6.399	1.619	7.174	13.917	9.451	50.077	11.551	17	83	41.402	128.598
13	6.316	6.367	0.821	6.741	6.734	8.996	42.444	11.901	80	20	126.436	43.564
14	6.310	6.374	1.016	6.345	0.554	6.588	4.417	13.246	65	35	105.102	64.898
15	6.326	6.386	0.952	6.772	7.050	8.850	39.899	11.901	20	80	44.315	125.685
16	6.329	6.519	3.004	6.458	2.028	7.535	19.050	12.588	26	74	53.059	116.941
17	6.325	6.466	2.226	7.111	12.422	7.738	22.347	12.474	75	25	118.613	51.387
18	6.302	6.332	0.485	6.385	1.319	6.410	1.721	13.421	38	62	69.612	100.388
19	6.327	6.505	2.821	6.620	4.635	7.592	20.004	12.588	74	26	117.429	52.571
20	6.302	6.336	0.533	6.393	1.444	6.402	1.575	13.435	39	61	69.943	100.057

Table A.2: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 100$ , cost model 3 and  $TS = 170$  in Table 5.1 See text in Chapter 5 for further details.

	OPI DYN	RB	dif%	IO	dif%	JSQ	dif%	OPI STA	$v_1$	$v_2$	$\mu_1$	$\mu_2$
1	3.072	3.073	0.035	3.072	0.027	3.072	0.008	6.521	44	56	91.403	108.597
2	3.067	3.070	0.121	3.070	0.090	3.068	0.050	6.506	43	57	89.834	110.166
3	3.075	3.155	2.625	3.381	9.960	5.078	65.157	5.249	87	13	152.501	47.499
4	3.053	3.067	0.464	3.059	0.214	3.210	5.144	6.297	67	33	123.539	76.461
5	3.076	3.144	2.181	3.358	9.141	5.148	67.346	5.132	12	88	46.953	153.047
6	3.060	3.062	0.094	3.060	0.000	3.171	3.636	6.337	34	66	78.321	121.679
7	3.070	3.084	0.437	3.439	12.012	4.813	56.759	5.361	15	85	49.708	150.292
8	3.067	3.087	0.656	3.479	13.437	5.749	87.442	4.878	91	9	157.170	42.830
9	3.094	3.118	0.755	3.337	7.843	4.078	31.790	5.736	80	20	142.389	57.611
10	3.095	3.127	1.027	3.433	10.893	3.870	25.020	5.819	21	79	60.607	139.393
11	3.133	3.133	0.000	3.133	0.000	3.133	0.000	6.550	49	51	98.145	101.855
12	3.073	3.158	2.779	3.431	11.653	4.929	60.416	5.249	13	87	48.708	151.292
13	3.067	3.090	0.751	3.507	14.318	4.645	51.437	5.461	84	16	148.748	51.252
14	3.053	3.066	0.442	3.059	0.197	3.212	5.234	6.297	67	33	123.650	76.350
15	3.068	3.098	0.967	3.231	5.310	4.555	48.464	5.461	16	84	52.136	147.864
16	3.083	3.137	1.768	3.177	3.065	3.759	21.933	5.893	23	77	62.422	137.578
17	3.097	3.126	0.931	3.427	10.673	3.880	25.281	5.819	79	21	139.545	60.455
18	3.058	3.065	0.224	3.085	0.870	3.114	1.830	6.405	37	63	81.896	118.104
19	3.086	3.130	1.429	3.166	2.586	3.793	22.902	5.893	77	23	138.151	61.849
20	3.058	3.065	0.242	3.088	1.012	3.109	1.690	6.411	37	63	82.286	117.714

Table A.3: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 100$ , cost model 3 and  $TS = 200$  in Table 5.1 See text in Chapter 5 for further details.

	OPI DYN	RB	dif %	IO	dif%	JSQ	dif%	OPI STA	$v_1$	$v_2$	$\mu_1$	$\mu_2$
1	1.894	1.894	0.010	1.894	0.007	1.894	0.005	3.660	44	56	105.113	124.887
2	1.890	1.891	0.034	1.891	0.023	1.891	0.033	3.649	42	58	103.309	126.691
3	1.904	1.908	0.246	2.009	5.521	3.107	63.231	2.837	91	9	175.376	54.624
4	1.877	1.878	0.091	1.877	0.000	1.963	4.604	3.515	69	31	142.070	87.930
5	1.899	1.903	0.187	1.996	5.127	3.153	66.016	2.759	8	92	53.996	176.004
6	1.879	1.885	0.315	1.882	0.187	1.942	3.340	3.542	32	68	90.069	139.931
7	1.896	1.934	2.042	2.061	8.721	2.937	54.924	2.908	11	89	57.164	172.836
8	1.872	1.875	0.168	2.021	7.973	3.548	89.520	2.591	95	5	180.745	49.255
9	1.900	1.904	0.187	2.030	6.812	2.476	30.268	3.154	83	17	163.748	66.252
10	1.914	1.922	0.419	2.082	8.794	2.349	22.709	3.207	18	82	69.698	160.302
11	1.936	1.936	0.000	1.936	0.000	1.936	0.000	3.676	48	52	112.867	117.133
12	1.899	1.922	1.226	2.036	7.253	3.011	58.600	2.837	9	91	56.015	173.985
13	1.893	1.895	0.099	2.088	10.310	2.830	49.454	2.976	88	12	171.060	58.940
14	1.877	1.878	0.089	1.877	0.000	1.965	4.696	3.515	69	31	142.197	87.803
15	1.893	1.896	0.142	1.956	3.318	2.772	46.456	2.976	12	88	59.956	170.044
16	1.910	1.927	0.898	1.935	1.312	2.282	19.492	3.255	20	80	71.785	158.215
17	1.913	1.921	0.387	2.079	8.674	2.355	23.058	3.207	82	18	160.476	69.524
18	1.884	1.893	0.509	1.893	0.480	1.911	1.477	3.585	35	65	94.180	135.820
19	1.914	1.933	1.013	1.929	0.813	2.302	20.303	3.255	80	20	158.874	71.126
20	1.883	1.885	0.087	1.894	0.572	1.909	1.376	3.589	36	64	94.628	135.372

Table A.4: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 100$ , cost model 3 and  $TS = 230$  in Table 5.1 See text in Chapter 5 for further details.

	OPI DYN	RB	dif%	IO	dif%	JSQ	dif%	OPI STA	$v_1$	$v_2$	$\mu_1$	$\mu_2$
1	1.366	1.366	0.051	1.366	0.000	1.366	0.002	2.350	43	57	118.824	141.176
2	1.362	1.362	0.004	1.362	0.000	1.362	0.022	2.344	41	59	116.784	143.216
3	1.356	1.358	0.157	1.407	3.785	2.109	55.573	1.786	94	6	198.251	61.749
4	1.347	1.347	0.031	1.347	0.000	1.395	3.538	2.253	71	29	160.601	99.399
5	1.352	1.354	0.138	1.399	3.486	2.139	58.204	1.732	4	96	61.039	198.961
6	1.349	1.350	0.097	1.350	0.051	1.383	2.516	2.271	31	69	101.818	158.182
7	1.359	1.372	0.921	1.439	5.891	1.996	46.878	1.836	7	93	64.620	195.380
8	1.334	1.335	0.078	1.393	4.475	2.405	80.329	1.614	99	1	204.321	55.679
9	1.362	1.362	0.049	1.448	6.316	1.701	24.878	2.006	86	14	185.106	74.894
10	1.373	1.373	0.054	1.373	0.000	1.622	18.137	2.042	15	85	78.789	181.211
11	1.396	1.396	0.000	1.396	0.000	1.396	0.000	2.362	48	52	127.589	132.411
12	1.362	1.364	0.149	1.424	4.567	2.045	50.161	1.786	6	94	63.321	196.679
13	1.357	1.359	0.182	1.464	7.919	1.927	42.011	1.882	91	9	193.373	66.627
14	1.347	1.347	0.026	1.347	0.000	1.395	3.598	2.253	71	29	160.745	99.255
15	1.356	1.356	0.030	1.390	2.479	1.889	39.342	1.882	9	91	67.776	192.224
16	1.375	1.383	0.629	1.382	0.550	1.581	14.996	2.075	17	83	81.149	178.851
17	1.372	1.373	0.055	1.372	0.000	1.625	18.466	2.042	85	15	181.408	78.592
18	1.354	1.357	0.245	1.357	0.234	1.368	1.007	2.300	34	66	106.465	153.535
19	1.377	1.380	0.190	1.379	0.119	1.593	15.671	2.075	83	17	179.597	80.403
20	1.354	1.359	0.367	1.359	0.355	1.367	0.942	2.303	34	66	106.971	153.029

Table A.5: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 100$ , cost model 3 and  $TS = 260$  in Table 5.1 See text in Chapter 5 for further details.

	OPI DYN	RB	dif%	IO	dif%	JSQ	dif%	OPI STA	$v_1$	$v_2$	$\mu_1$	$\mu_2$
1	1.084	1.084	0.021	1.084	0.000	1.084	0.002	1.686	42	58	132.534	157.466
2	1.080	1.080	0.001	1.080	0.000	1.080	0.018	1.681	40	60	130.259	159.741
3	1.060	1.060	0.046	1.091	2.983	1.543	45.625	1.271	97	3	221.126	68.874
4	1.064	1.064	0.005	1.064	0.000	1.090	2.424	1.615	73	27	179.132	110.868
5	1.057	1.058	0.039	1.086	2.722	1.564	47.947	1.231	1	99	68.082	221.918
6	1.066	1.066	0.018	1.066	0.002	1.084	1.680	1.628	29	71	113.566	176.434
7	1.065	1.072	0.696	1.114	4.619	1.468	37.818	1.308	4	96	72.076	217.924
8	1.041	1.041	0.041	1.070	2.804	1.747	67.875	1.147	100	0	227.896	62.104
9	1.071	1.073	0.127	1.131	5.607	1.274	18.925	1.434	89	11	206.465	83.535
10	1.080	1.081	0.016	1.080	0.000	1.224	13.295	1.461	13	87	87.880	202.120
11	1.108	1.108	0.000	1.108	0.000	1.108	0.000	1.694	48	52	142.310	147.690
12	1.066	1.067	0.070	1.103	3.471	1.501	40.717	1.271	3	97	70.627	219.373
13	1.064	1.064	0.050	1.082	1.748	1.421	33.608	1.343	94	6	215.685	74.315
14	1.064	1.064	0.009	1.064	0.000	1.090	2.469	1.615	73	27	179.292	110.708
15	1.064	1.064	0.072	1.088	2.335	1.397	31.313	1.343	6	94	75.597	214.403
16	1.082	1.088	0.545	1.088	0.518	1.199	10.752	1.486	14	86	90.512	199.488
17	1.080	1.080	0.015	1.080	0.000	1.226	13.566	1.461	87	13	202.340	87.660
18	1.071	1.073	0.187	1.073	0.184	1.077	0.627	1.650	33	67	118.749	171.251
19	1.085	1.086	0.095	1.085	0.072	1.206	11.221	1.486	86	14	200.319	89.681
20	1.071	1.071	0.005	1.073	0.273	1.077	0.586	1.651	33	67	119.314	170.686

Table A.6: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 100$ , cost model 3 and  $TS = 290$  in Table 5.1 See text in Chapter 5 for further details.

	OPI DYN	RB	dif%	IO	dif%	JSQ	dif%	OPI STA	$v_1$	$v_2$	$\mu_1$	$\mu_2$
1	0.911	0.911	0.009	0.911	0.000	0.911	0.002	1.314	42	58	146.245	173.755
2	0.908	0.908	0.000	0.908	0.000	0.908	0.012	1.310	40	60	143.734	176.266
3	0.877	0.877	0.012	0.898	2.419	1.196	36.498	0.987	100	0	244.001	75.999
4	0.890	0.890	0.002	0.890	0.000	0.904	1.565	1.259	74	26	197.663	122.337
5	0.874	0.874	0.008	0.893	2.208	1.211	38.518	0.958	0	100	75.125	244.875
6	0.892	0.892	0.006	0.892	0.000	0.902	1.043	1.269	28	72	125.314	194.686
7	0.880	0.880	0.023	0.888	0.878	1.144	30.074	1.017	2	98	79.532	240.468
8	0.857	0.857	0.015	0.876	2.196	1.338	56.100	0.905	100	0	251.472	68.528
9	0.891	0.891	0.049	0.891	0.045	1.015	13.965	1.117	91	9	227.823	92.177
10	0.900	0.900	0.010	0.900	0.000	0.983	9.252	1.138	10	90	96.972	223.028
11	0.932	0.932	0.000	0.932	0.000	0.932	0.000	1.320	48	52	157.032	162.968
12	0.880	0.883	0.275	0.908	3.119	1.167	32.581	0.987	0	100	77.934	242.066
13	0.880	0.880	0.010	0.896	1.813	1.113	26.430	1.045	97	3	237.997	82.003
14	0.890	0.890	0.002	0.890	0.001	0.904	1.597	1.259	74	26	197.840	122.160
15	0.881	0.881	0.024	0.901	2.339	1.096	24.469	1.045	3	97	83.417	236.583
16	0.900	0.907	0.775	0.907	0.766	0.967	7.485	1.157	12	88	99.875	220.125
17	0.899	0.899	0.010	0.899	0.003	0.985	9.477	1.138	90	10	223.272	96.728
18	0.896	0.898	0.208	0.898	0.207	0.900	0.369	1.286	31	69	131.034	188.966
19	0.901	0.905	0.363	0.904	0.355	0.972	7.844	1.157	88	12	221.042	98.958
20	0.897	0.897	0.001	0.899	0.274	0.900	0.344	1.287	32	68	131.657	188.343

Table A.7: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 100$ , cost model 3 and  $TS = 320$  in Table 5.1 See text in Chapter 5 for further details.

	OPI DYN	RB	dif%	IO	dif%	JSQ	dif%	OPI STA	$v_1$	$v_2$	$\mu_1$	$\mu_2$
1	0.793	0.794	0.004	0.793	0.000	0.793	0.001	1.086	41	59	159.955	190.045
2	0.790	0.790	0.000	0.790	0.000	0.790	0.007	1.082	39	61	157.210	192.790
3	0.750	0.750	0.001	0.767	2.212	0.970	29.294	0.815	100	0	266.876	83.124
4	0.771	0.771	0.000	0.771	0.000	0.779	0.982	1.039	76	24	216.194	133.806
5	0.748	0.748	0.009	0.763	2.052	0.980	31.057	0.795	0	100	82.167	267.833
6	0.774	0.774	0.002	0.774	0.000	0.778	0.611	1.048	26	74	137.062	212.938
7	0.752	0.752	0.008	0.761	1.115	0.933	24.075	0.837	0	100	86.988	263.012
8	0.732	0.732	0.003	0.746	1.933	1.070	46.278	0.757	100	0	275.047	74.953
9	0.767	0.767	0.026	0.767	0.024	0.846	10.287	0.920	94	6	249.181	100.819
10	0.776	0.776	0.003	0.776	0.001	0.825	6.289	0.938	8	92	106.063	243.937
11	0.813	0.813	0.000	0.813	0.000	0.813	0.000	1.091	48	52	171.754	178.246
12	0.752	0.752	0.011	0.756	0.535	0.949	26.253	0.815	0	100	85.240	264.760
13	0.754	0.754	0.001	0.769	1.974	0.911	20.933	0.860	99	1	260.309	89.691
14	0.771	0.771	0.000	0.771	0.000	0.779	1.005	1.039	76	24	216.387	133.613
15	0.755	0.755	0.006	0.773	2.441	0.900	19.249	0.860	1	99	91.237	258.763
16	0.774	0.774	0.019	0.782	1.054	0.815	5.199	0.954	10	90	109.239	240.761
17	0.775	0.775	0.003	0.775	0.000	0.826	6.479	0.938	92	8	244.203	105.797
18	0.778	0.778	0.000	0.780	0.237	0.779	0.217	1.062	30	70	143.318	206.682
19	0.775	0.775	0.002	0.780	0.644	0.818	5.451	0.954	90	10	241.765	108.235
20	0.778	0.778	0.001	0.780	0.288	0.780	0.203	1.063	31	69	144.000	206.000

Table A.8: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 100$ , cost model 3 and  $TS = 350$  in Table 5.1 See text in Chapter 5 for further details.

	OPI DYN	RB	dif%	IO	dif%	JSQ	dif%	OPI STA	$v_1$	$v_2$	$\mu_1$	$\mu_2$
1	522.299	522.299	0.000	522.299	0.000	522.299	0.000	522.467	93	107	77.692	92.308
2	522.301	522.301	0.000	522.301	0.000	522.301	0.000	522.468	91	109	76.359	93.641
3	522.375	522.376	0.000	522.376	0.000	522.384	0.002	522.360	145	55	129.626	40.374
4	522.320	522.320	0.000	522.320	0.000	522.321	0.000	522.448	121	79	105.008	64.992
5	522.375	522.377	0.000	522.377	0.000	522.386	0.002	522.353	53	147	39.910	130.090
6	522.318	522.317	0.000	522.318	0.000	522.318	0.000	522.449	81	119	66.573	103.427
7	522.370	522.372	0.000	522.372	0.000	522.379	0.002	522.368	57	143	42.251	127.749
8	522.382	522.387	0.001	522.384	0.000	522.398	0.003	522.336	150	50	133.594	36.406
9	522.355	522.356	0.000	522.357	0.000	522.359	0.001	522.397	136	64	121.031	48.969
10	522.349	522.349	0.000	522.351	0.000	522.353	0.001	522.405	66	134	51.516	118.484
11	522.291	522.291	0.000	522.291	0.000	522.291	0.000	522.470	98	102	83.423	86.577
12	522.372	522.374	0.000	522.374	0.000	522.381	0.002	522.360	55	145	41.402	128.598
13	522.367	522.368	0.000	522.369	0.000	522.375	0.001	522.378	141	59	126.436	43.564
14	522.321	522.320	0.000	522.321	0.000	522.322	0.000	522.448	121	79	105.102	64.898
15	522.366	522.367	0.000	522.367	0.000	522.373	0.001	522.378	59	141	44.315	125.685
16	522.346	522.346	0.000	522.346	0.000	522.348	0.000	522.411	67	133	53.059	116.941
17	522.350	522.350	0.000	522.350	0.000	522.353	0.001	522.405	134	66	118.613	51.387
18	522.312	522.312	0.000	522.312	0.000	522.312	0.000	522.458	84	116	69.612	100.388
19	522.347	522.347	0.000	522.348	0.000	522.350	0.001	522.411	133	67	117.429	52.571
20	522.311	522.311	0.000	522.312	0.000	522.312	0.000	522.456	85	115	69.943	100.057

Table A.9: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 200$ , cost model 3 and  $TS = 170$  in Table 5.2 See text in Chapter 5 for further details.

	OPI DYN	RB	dif%	IO	dif%	JSQ	dif%	OPI STA	$v_1$	$v_2$	$\mu_1$	$\mu_2$
1	264.951	264.955	0.001	264.959	0.003	264.958	0.002	278.201	93	107	91.403	108.597
2	264.951	264.956	0.002	264.965	0.005	264.965	0.005	278.148	91	109	89.834	110.166
3	264.966	265.151	0.070	265.544	0.218	266.862	0.716	275.959	148	52	152.501	47.499
4	264.952	264.995	0.016	265.018	0.025	265.155	0.076	277.763	122	78	123.539	76.461
5	264.965	265.152	0.070	265.519	0.209	266.926	0.740	275.745	50	150	46.953	153.047
6	264.953	264.969	0.006	264.992	0.015	265.113	0.060	277.859	80	120	78.321	121.679
7	264.965	265.124	0.060	265.445	0.181	266.624	0.626	276.160	53	147	49.708	150.292
8	264.966	265.168	0.076	265.712	0.281	267.472	0.946	275.303	154	46	157.170	42.830
9	264.959	265.078	0.045	265.252	0.111	265.965	0.380	276.753	139	61	142.389	57.611
10	264.955	265.074	0.045	265.271	0.119	265.777	0.310	276.899	63	137	60.607	139.393
11	264.962	264.962	0.000	264.962	0.000	264.962	0.000	278.214	98	102	98.145	101.855
12	264.965	265.143	0.067	265.491	0.199	266.728	0.665	275.959	52	148	48.708	151.292
13	264.963	265.106	0.054	265.449	0.183	266.473	0.570	276.311	145	55	148.748	51.252
14	264.952	264.994	0.016	265.016	0.024	265.157	0.077	277.763	122	78	123.650	76.350
15	264.962	265.109	0.056	265.347	0.145	266.392	0.540	276.311	55	145	52.136	147.864
16	264.956	265.025	0.026	265.218	0.099	265.676	0.272	277.046	65	135	62.422	137.578
17	264.955	265.071	0.044	265.267	0.118	265.786	0.314	276.899	137	63	139.545	60.455
18	264.951	264.969	0.007	264.996	0.017	265.047	0.036	277.953	83	117	81.896	118.104
19	264.956	265.021	0.025	265.204	0.094	265.707	0.284	277.046	135	65	138.151	61.849
20	264.951	264.970	0.007	264.999	0.018	265.041	0.034	278.015	84	116	82.286	117.714

Table A.10: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 200$ , cost model 3 and  $TS = 200$  in Table 5.2 See text in Chapter 5 for further details.

	OPI DYN	RB	dif %	IO	dif%	JSQ	dif%	OPI STA	$v_1$	$v_2$	$\mu_1$	$\mu_2$
1	82.276	82.296	0.024	82.305	0.035	82.306	0.036	117.272	92	108	105.113	124.887
2	82.267	82.297	0.036	82.325	0.071	82.337	0.085	117.212	90	110	103.309	126.691
3	82.261	83.076	0.992	85.179	3.548	91.513	11.247	112.401	152	48	175.376	54.624
4	82.242	82.540	0.363	82.548	0.372	83.337	1.331	116.406	124	76	142.070	87.930
5	82.266	83.149	1.073	85.058	3.395	91.782	11.568	111.959	46	154	53.996	176.004
6	82.239	82.359	0.145	82.592	0.428	83.114	1.064	116.564	78	122	90.069	139.931
7	82.264	83.152	1.079	85.289	3.677	90.480	9.987	112.818	50	150	57.164	172.836
8	82.296	83.313	1.235	86.127	4.656	94.020	14.246	110.993	158	42	180.745	49.255
9	82.233	83.046	0.988	84.046	2.205	87.451	6.346	114.260	142	58	163.748	66.252
10	82.236	82.853	0.750	84.051	2.206	86.536	5.228	114.568	60	140	69.698	160.302
11	82.369	82.369	0.000	82.369	0.000	82.369	0.000	117.371	98	102	112.867	117.133
12	82.259	83.083	1.002	85.080	3.429	90.936	10.548	112.401	48	152	56.015	173.985
13	82.274	83.199	1.124	84.938	3.237	89.810	9.159	113.211	148	52	171.060	58.940
14	82.243	82.533	0.353	82.541	0.363	83.351	1.347	116.406	124	76	142.197	87.803
15	82.272	83.133	1.046	84.123	2.250	89.445	8.718	113.211	52	148	59.956	170.044
16	82.221	82.804	0.709	83.550	1.616	86.036	4.639	114.858	62	138	71.785	158.215
17	82.236	82.838	0.732	84.030	2.182	86.580	5.282	114.568	140	60	160.476	69.524
18	82.247	82.349	0.124	82.572	0.395	82.765	0.630	116.832	82	118	94.180	135.820
19	82.224	82.756	0.648	83.475	1.522	86.190	4.824	114.858	138	62	158.874	71.126
20	82.246	82.364	0.144	82.515	0.327	82.733	0.592	116.863	82	118	94.628	135.372

Table A.11: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 200$ , cost model 3 and  $TS = 230$  in Table 5.2 See text in Chapter 5 for further details.

	OPI DYN	RB	dif%	IO	dif%	JSQ	dif%	OPI STA	$v_1$	$v_2$	$\mu_1$	$\mu_2$
1	23.526	23.547	0.091	23.539	0.056	23.552	0.112	48.213	91	109	118.824	141.176
2	23.516	23.532	0.069	23.546	0.128	23.582	0.278	48.151	89	111	116.784	143.216
3	23.610	24.201	2.506	25.635	8.579	32.449	37.440	44.177	156	44	198.251	61.749
4	23.489	23.693	0.866	23.672	0.779	24.540	4.472	47.483	126	74	160.601	99.399
5	23.610	24.237	2.658	25.544	8.192	32.713	38.557	43.804	42	158	61.039	198.961
6	23.495	23.740	1.043	23.752	1.092	24.326	3.539	47.624	76	124	101.818	158.182
7	23.579	24.123	2.308	25.808	9.454	31.441	33.346	44.534	46	154	64.620	195.380
8	23.592	24.614	4.331	26.445	12.091	34.904	47.945	43.000	162	38	204.321	55.679
9	23.541	24.063	2.218	25.232	7.184	28.499	21.063	45.706	145	55	185.106	74.894
10	23.503	23.911	1.733	24.221	3.053	27.615	17.493	45.961	57	143	78.789	181.211
11	23.614	23.614	0.000	23.614	0.000	23.614	0.000	48.278	98	102	127.589	132.411
12	23.590	24.278	2.917	25.769	9.240	31.886	35.168	44.177	44	156	63.321	196.679
13	23.557	24.096	2.288	26.045	10.562	30.789	30.701	44.868	151	49	193.373	66.627
14	23.489	23.688	0.846	23.667	0.758	24.553	4.529	47.483	126	74	160.745	99.255
15	23.550	24.167	2.620	24.795	5.284	30.433	29.226	44.868	49	151	67.776	192.224
16	23.509	23.874	1.554	24.388	3.738	27.132	15.414	46.205	59	141	81.149	178.851
17	23.504	23.900	1.686	24.207	2.995	27.657	17.671	45.961	143	57	181.408	78.592
18	23.503	23.596	0.398	23.724	0.942	23.992	2.082	47.847	80	120	106.465	153.535
19	23.505	23.898	1.672	24.330	3.510	27.281	16.061	46.205	141	59	179.597	80.403
20	23.505	23.586	0.345	23.742	1.009	23.962	1.941	47.847	81	119	106.971	153.029

Table A.12: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 200$ , cost model 3 and  $TS = 260$  in Table 5.2 See text in Chapter 5 for further details.

	OPI DYN	RB	dif%	IO	dif%	JSQ	dif%	OPI STA	$v_1$	$v_2$	$\mu_1$	$\mu_2$
1	8.989	9.000	0.120	8.992	0.034	9.004	0.174	21.926	91	109	132.534	157.466
2	8.982	9.005	0.260	8.991	0.099	9.022	0.445	21.899	88	112	130.259	159.741
3	9.080	9.291	2.327	10.220	12.560	14.910	64.217	19.367	160	40	221.126	68.874
4	8.974	9.087	1.251	9.034	0.667	9.615	7.139	21.470	128	72	179.132	110.868
5	9.080	9.273	2.131	10.166	11.966	15.097	66.272	19.132	38	162	68.082	221.918
6	8.966	9.073	1.188	9.078	1.242	9.482	5.748	21.563	75	125	113.566	176.434
7	9.074	9.326	2.776	10.231	12.750	14.203	56.521	19.595	42	158	72.076	217.924
8	9.075	9.477	4.428	10.673	17.614	16.669	83.686	18.636	166	34	227.896	62.104
9	9.013	9.365	3.903	9.875	9.562	12.188	35.224	20.353	148	52	206.465	83.535
10	9.013	9.145	1.467	9.305	3.245	11.599	28.696	20.506	54	146	87.880	202.120
11	9.046	9.047	0.000	9.046	0.000	9.046	0.000	21.980	98	102	142.310	147.690
12	9.083	9.288	2.260	10.157	11.823	14.514	59.797	19.367	40	160	70.627	219.373
13	9.066	9.246	1.984	9.722	7.236	13.749	51.658	19.795	155	45	215.685	74.315
14	8.975	9.085	1.223	9.033	0.650	9.624	7.225	21.470	128	72	179.292	110.708
15	9.060	9.259	2.192	9.769	7.817	13.504	49.038	19.795	45	155	75.597	214.403
16	8.991	9.189	2.204	9.386	4.395	11.281	25.473	20.657	56	144	90.512	199.488
17	9.015	9.141	1.392	9.303	3.188	11.627	28.967	20.506	146	54	202.340	87.660
18	8.967	9.015	0.541	9.076	1.220	9.274	3.426	21.688	79	121	118.749	171.251
19	8.996	9.169	1.929	9.358	4.028	11.378	26.485	20.657	144	56	200.319	89.681
20	8.968	9.002	0.376	9.085	1.302	9.255	3.199	21.707	79	121	119.314	170.686

Table A.13: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 200$ , cost model 3 and  $TS = 290$  in Table 5.2 See text in Chapter 5 for further details.

	OPI DYN	RB	dif%	IO	dif%	JSQ	dif%	OPI STA	$v_1$	$v_2$	$\mu_1$	$\mu_2$
1	4.684	4.687	0.050	4.684	0.000	4.692	0.170	11.348	90	110	146.245	173.755
2	4.680	4.686	0.125	4.681	0.031	4.700	0.440	11.333	88	112	143.734	176.266
3	4.753	4.841	1.850	5.333	12.206	8.320	75.045	9.784	164	36	244.001	75.999
4	4.679	4.729	1.073	4.695	0.360	5.025	7.409	11.071	130	70	197.663	122.337
5	4.753	4.830	1.634	5.320	11.942	8.444	77.670	9.644	34	166	75.125	244.875
6	4.670	4.718	1.015	4.715	0.952	4.950	5.992	11.124	73	127	125.314	194.686
7	4.757	4.826	1.442	5.021	5.542	7.852	65.060	9.922	38	162	79.532	240.468
8	4.761	4.828	1.419	5.551	16.596	9.510	99.750	9.328	171	29	251.472	68.528
9	4.712	4.783	1.516	4.884	3.652	6.563	39.281	10.383	152	48	227.823	92.177
10	4.701	4.745	0.943	4.859	3.349	6.198	31.847	10.487	50	150	96.972	223.028
11	4.724	4.724	0.000	4.724	0.000	4.724	0.000	11.384	98	102	157.032	162.968
12	4.757	4.811	1.144	5.325	11.940	8.057	69.374	9.784	36	164	77.934	242.066
13	4.756	4.831	1.577	5.076	6.713	7.556	58.870	10.046	159	41	237.997	82.003
14	4.679	4.727	1.020	4.705	0.542	5.030	7.495	11.071	130	70	197.840	122.160
15	4.748	4.847	2.086	5.091	7.218	7.397	55.796	10.046	41	159	83.417	236.583
16	4.702	4.744	0.891	4.886	3.916	6.004	27.708	10.575	53	147	99.875	220.125
17	4.701	4.744	0.922	4.875	3.708	6.215	32.212	10.487	150	50	223.272	96.728
18	4.668	4.683	0.322	4.723	1.188	4.835	3.576	11.207	77	123	131.034	188.966
19	4.703	4.755	1.106	4.873	3.614	6.064	28.942	10.575	147	53	221.042	98.958
20	4.668	4.683	0.320	4.728	1.289	4.824	3.357	11.213	78	122	131.657	188.343

Table A.14: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 200$ , cost model 3 and  $TS = 320$  in Table 5.2 See text in Chapter 5 for further details.

	OPI DYN	RB	dif%	IO	dif%	JSQ	dif%	OPI STA	$v_1$	$v_2$	$\mu_1$	$\mu_2$
1	3.070	3.071	0.028	3.070	0.000	3.074	0.139	6.688	89	111	159.955	190.045
2	3.066	3.068	0.061	3.066	0.001	3.077	0.336	6.676	87	113	157.210	192.790
3	3.115	3.134	0.614	3.474	11.511	5.284	69.613	5.700	168	32	266.876	83.124
4	3.064	3.080	0.513	3.072	0.237	3.245	5.889	6.513	132	68	216.194	133.806
5	3.116	3.136	0.649	3.458	10.969	5.366	72.214	5.609	29	171	82.167	267.833
6	3.058	3.083	0.804	3.079	0.694	3.204	4.768	6.545	71	129	137.062	212.938
7	3.116	3.131	0.482	3.274	5.052	4.976	59.676	5.786	34	166	86.988	263.012
8	3.127	3.142	0.478	3.584	14.618	6.086	94.652	5.410	176	24	275.047	74.953
9	3.092	3.120	0.908	3.178	2.792	4.154	34.340	6.081	155	45	249.181	100.819
10	3.078	3.100	0.712	3.180	3.331	3.930	27.690	6.142	47	153	106.063	243.937
11	3.102	3.102	0.000	3.102	0.000	3.102	0.000	6.707	97	103	171.754	178.246
12	3.115	3.143	0.895	3.280	5.295	5.110	64.044	5.700	32	168	85.240	264.760
13	3.120	3.142	0.711	3.303	5.890	4.784	53.334	5.866	163	37	260.309	89.691
14	3.065	3.079	0.474	3.071	0.202	3.247	5.964	6.513	132	68	216.387	133.613
15	3.120	3.142	0.704	3.317	6.309	4.681	50.029	5.866	37	163	91.237	258.763
16	3.076	3.091	0.507	3.193	3.825	3.813	23.956	6.201	50	150	109.239	240.761
17	3.078	3.099	0.688	3.178	3.248	3.940	28.011	6.142	153	47	244.203	105.797
18	3.056	3.065	0.304	3.089	1.111	3.142	2.838	6.597	76	124	143.318	206.682
19	3.076	3.090	0.461	3.186	3.580	3.848	25.115	6.201	150	50	241.765	108.235
20	3.056	3.066	0.353	3.092	1.197	3.137	2.655	6.602	76	124	144.000	206.000

Table A.15: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 200$ , cost model 3 and  $TS = 350$  in Table 5.2 See text in Chapter 5 for further details.

	OPI DYN	RB	dif%	IO	dif%	JSQ	dif%	OPI STA	$v_1$	$v_2$	$\mu_1$	$\mu_2$
1	2.318	2.318	0.031	2.318	0.000	2.320	0.093	4.438	89	111	173.666	206.334
2	2.314	2.316	0.065	2.314	0.000	2.319	0.223	4.430	86	114	170.685	209.315
3	2.340	2.345	0.209	2.437	4.151	3.682	57.335	3.775	172	28	289.752	90.248
4	2.307	2.313	0.253	2.311	0.145	2.401	4.073	4.322	134	66	234.725	145.275
5	2.340	2.344	0.184	2.583	10.386	3.738	59.734	3.714	25	175	89.210	290.790
6	2.304	2.323	0.820	2.317	0.596	2.379	3.283	4.343	69	131	148.810	231.190
7	2.341	2.351	0.450	2.455	4.882	3.477	48.543	3.834	31	169	94.444	285.556
8	2.349	2.356	0.306	2.511	6.892	4.230	80.085	3.577	180	20	298.623	81.377
9	2.327	2.337	0.412	2.383	2.407	2.945	26.564	4.032	159	41	270.540	109.460
10	2.316	2.322	0.247	2.389	3.147	2.806	21.127	4.074	44	156	115.154	264.846
11	2.345	2.345	0.000	2.345	0.000	2.345	0.000	4.451	97	103	186.476	193.524
12	2.340	2.347	0.308	2.453	4.861	3.566	52.409	3.775	28	172	92.546	287.454
13	2.344	2.349	0.214	2.476	5.627	3.350	42.924	3.888	167	33	282.622	97.378
14	2.308	2.313	0.223	2.310	0.120	2.403	4.129	4.322	134	66	234.935	145.065
15	2.346	2.353	0.317	2.489	6.112	3.283	39.965	3.888	33	167	99.058	280.942
16	2.313	2.321	0.325	2.403	3.876	2.734	18.174	4.113	47	153	118.602	261.398
17	2.317	2.322	0.234	2.388	3.066	2.812	21.383	4.074	156	44	265.135	114.865
18	2.303	2.305	0.125	2.328	1.090	2.347	1.943	4.378	74	126	155.602	224.398
19	2.314	2.322	0.349	2.398	3.648	2.755	19.088	4.113	153	47	262.488	117.512
20	2.303	2.306	0.139	2.303	0.000	2.345	1.806	4.381	75	125	156.343	223.657

Table A.16: The comparative performance of the restless bandit heuristic against the optimal dynamic solution and the optimal static solution for  $K = 200$ , cost model 3 and  $TS = 380$  in Table 5.2 See text in Chapter 5 for further details.

## Appendix B

### The published paper

I include the paper published in the *Journal of the Operational Research Society* with the permission from my coauthor, Professor K. D. Glazebrook.



# A static allocation model for the outsourcing of warranty repairs

L Ding\* and KD Glazebrook

Edinburgh University, UK

There has been a strong recent trend among original equipment manufacturers toward the outsourcing of work relating to the repair of items under warranty. In the typical cases where a large manufacturer uses several vendors to perform such work, we develop and analyse models to support decisions concerning how the work should be distributed among them. We depart from previous work in arguing the importance of an approach to the modelling of goodwill costs which takes explicit account of the delays experienced by customers. Theoretical considerations and numerical work both lend strong support to the contention that simple greedy approaches to workload allocation work well.

*Journal of the Operational Research Society* (2005) 56, 825–835. doi:10.1057/palgrave.jors.2601904

Published online 1 December 2004

**Keywords:** dynamic programming; greedy heuristic; finite-source queueing system; resource allocation; warranty repairs

## Introduction

Opp *et al*<sup>1</sup> discuss the development of policies for the optimal utilization of service vendors in the out-sourcing of warranty repairs by an equipment manufacturer. In this they were motivated by interaction with a leading manufacturer of PCs. Such work reflects a strong recent trend among producers of electronic equipment toward the outsourcing of warranty work in the interests of focusing on core business. For example, claims have been made that outsourcing by equipment manufacturers has grown at a compound annual growth rate of more than 30% over the past 5 years.<sup>2</sup> Opp *et al*<sup>1</sup> themselves quote a recent report by Merrill Lynch<sup>3</sup> to the effect that warranty services represent a 100 billion dollar opportunity for manufacturers and subcontractors. In light of such figures, serious consideration of how such outsourcing should be effectively managed is important and overdue.

Typically, a large manufacturer will have contracts with several service vendors to perform warranty work. While the manufacturer will be concerned to distribute the workload among these in a cost-effective manner, he will also be concerned to minimize the risks associated with long delays to repairs (and resulting customer dissatisfaction) when vendors become overloaded. Opp *et al*<sup>1</sup> develop a resource allocation model for making decisions about how warranty work should be distributed among a collection of vendors of known characteristics to minimize a combination of repair costs (parts and labour) and goodwill costs, while arguing that it is the concern for the latter which should predominate.

Plainly, in such a study the approach to the modelling of goodwill costs is absolutely critical. Opp *et al*<sup>1</sup> take an approach, conventional in such areas as inventory management and queueing control, whereby goodwill costs are modelled as linear holding costs in which a single unit of time spent by a single item awaiting repair incurs a fixed goodwill cost of  $h$ , say. This approach has the consequence that the optimization problem underlying the distribution of warranty work then becomes a separable convex resource allocation problem (under simply stated conditions) for which greedy solutions are known to be optimal, see Gross.<sup>4</sup> However, in several application areas, this approach to the modelling of goodwill costs has been brought into question. In queueing control, for example, see the discussions in van Meighem<sup>5</sup> and Ansell *et al*.<sup>6</sup> We argue that in the current context it is important to take explicit account of the delays experienced by individual customers when an item is at a vendor for repair. The manufacturer may, for example, have a code of service quality whereby every effort is made to keep customer delays below a specified amount,  $\tau$  say. Allocation to vendors should then be made in a way which makes delays greater than  $\tau$  improbable. We propose two cost models which reflect this reality. Sadly, these models result in much more challenging optimization problems for solution. Instead of the vendor-specific cost rates which are convex in the allocated work which result from the linear holding cost model, we now typically have convexity of vendor-specific cost rates over part of the range concerned only. Nonetheless, we give grounds in the upcoming sections for believing that simple greedy solution approaches should still perform well. This point of view is roundly endorsed in the concluding numerical study where, of 450 problems studied, a greedy solution is optimal in 449 cases.

\*Correspondence: L Ding, School of Management, University of Edinburgh, WRB George Square, EH8 9JY, UK.  
E-mail: kevin.glazebrook@ed.ac.uk

The essentials of our allocation models are presented in the next two sections and theoretical results are established which shed light on the subsequent discussion of partially convex cost rates and greedy heuristics. The subsequent section uses those results as a basis of the discussion of how vendor cost rates vary with allocated work under our models and how that in turn impacts upon the status of greedy heuristics. The paper concludes with details of a numerical study and some concluding remarks.

### A static allocation model

$K$  identical items under warranty are available for distribution among  $V$  service vendors. Decision variable  $k_i$  denotes the number of items allocated to vendor  $i$ ,  $1 \leq i \leq V$ . Upon allocation of  $k_i$  items to vendor  $i$  there is a consequential cost per unit time incurred, denoted  $g_i(k_i)$ ,  $1 \leq i \leq V$ . In the models discussed here this cost has two components. The first (and less important) component concerns the cost of repair of each item. We shall assume that for each repair performed by vendor  $i$ , the manufacturer must pay the vendor a fixed amount  $c_i$ ,  $1 \leq i \leq V$ . It is not unreasonable to assume that these fixed costs do not vary greatly across vendors. The second (and more important) component is the goodwill cost rate. We shall assume that such goodwill costs are incurred whenever an item requiring repair remains at the vendor for a length of time exceeding some given specified amount  $\tau$ . Explicit expressions for  $g_i(k_i)$ ,  $1 \leq i \leq V$ , will be developed in the next section via a finite population queueing model.

The problem of optimally allocating items to vendors (ie to achieve minimal total cost rate) is a resource allocation problem with integer variables, expressed by

$$\begin{aligned} \min \quad & \sum_{i=1}^V g_i(k_i) \\ \text{s.t.} \quad & \sum_{i=1}^V k_i = K \\ & k_i \in \mathbb{Z}^+, \quad i = 1, \dots, V \end{aligned} \quad (1)$$

See, for example, Gross,<sup>4</sup> Fox<sup>7</sup> and Ibaraki and Katoh.<sup>8</sup> It is very easy to show by means of a pairwise interchange argument that when each of the  $g_i$ 's is increasing convex then the problem in (1) may be solved by a greedy algorithm. This solution approach was first proposed by Gross,<sup>4</sup> see also Fox.<sup>7</sup> The algorithm may be expressed as follows:

#### Greedy Algorithm for (1)

- Step 0:* Set  $k_i = 0$ ,  $1 \leq i \leq V$ ;  
*Step 1:* Choose any  $j \in \arg \min_{1 \leq i \leq V} [g_i(k_i + 1) - g_i(k_i)]$ ;  
*Step 2:* Set  $k_j = k_j + 1$ ;  
*Step 3:* If  $\sum_{i=1}^V k_i < K$ , go to Step 1; otherwise stop.

Dynamic programming is available as a solution method for (1) even when the above convexity does not hold. However,

this is not a viable option computationally for problems of realistic size. We shall find that the cost rate functions  $g_i$  for our models are not fully increasing convex in general. Rather, they are increasing functions which are convex in some range of the form  $0 \leq k_i \leq B_i$ ,  $1 \leq i \leq V$ . If  $B_i$  is maximal in this respect, we shall refer to it as a *convexity boundary point* for the function  $g_i$ . A formal definition is given later in the paper. In light of the above result of Gross,<sup>4</sup> and the difficulty of utilizing dynamic programming as a solution method, it is natural to explore the performance of greedy heuristics in such cases. We shall conclude this section with a motivating theoretical result.

Let  $B_i$  be the convexity boundary point for  $g_i$ ,  $1 \leq i \leq V$ . Suppose that  $\sum_{i=1}^V B_i \geq K$  and consider the following variant of optimization problem (1):

$$\begin{aligned} \min \quad & \sum_{i=1}^V g_i(k_i) \\ \text{s.t.} \quad & \sum_{i=1}^V k_i = K \\ & k_i \leq B_i \\ & k_i \in \mathbb{Z}^+, \quad i = 1, \dots, V \end{aligned} \quad (2)$$

A greedy solution to problem (2) proceeds as follows:

#### Greedy Algorithm for (2)

- Step 0:* Set  $k_i = 0$ ,  $1 \leq i \leq V$ ;  
*Step 1:* Choose any  $j \in \arg \min_{1 \leq i \leq V} [g_i(k_i + 1) - g_i(k_i)]$  when the *argmin* is taken over those  $i$  for which  $k_i < B_i$ ;  
*Step 2:* Set  $k_j = k_j + 1$ ;  
*Step 3:* If  $\sum_{i=1}^V k_i < K$ , go to Step 1; otherwise stop.

**Lemma 1** *The above greedy algorithm solves optimization problem (2).*

**Proof** Define the constant  $G$  by

$$G = \max_{1 \leq i \leq V} [g_i(B_i) - g_i(B_i - 1)]$$

and the functions  $\bar{g}_i$ ,  $1 \leq i \leq V$ , by

$$\bar{g}_i(k_i) = \begin{cases} g_i(k_i), & 0 \leq k_i \leq B_i \\ g_i(B_i) + (k_i - B_i)G, & k_i \geq B_i + 1 \end{cases} \quad (3)$$

Now consider the optimization problem

$$\begin{aligned} \min \quad & \sum_{i=1}^V \bar{g}_i(k_i) \\ \text{s.t.} \quad & \sum_{i=1}^V k_i = K \\ & k_i \in \mathbb{Z}^+, \quad i = 1, \dots, V \end{aligned} \quad (4)$$

Since  $\bar{g}_i$  coincides with  $g_i$  over convex ranges  $0 \leq k_i \leq B_i$  it must be true that any solution to (4) for which  $k_i \leq B_i$ ,  $1 \leq i \leq V$ , must necessarily solve (2). However, by

construction,  $\bar{g}_i$  is increasing convex for all  $i, 1 \leq i \leq V$ . Hence (4) is solved by a greedy algorithm, as with (1) above. However, note from (3) that it follows from the convexity of  $g_i$  in the range  $0 \leq k_i \leq B_i$ , that

$$\begin{aligned} k_i \geq B_i &\Rightarrow \bar{g}_i(k_i + 1) - \bar{g}_i(k_i) = G \\ &= \max_{1 \leq j \leq V} \{g_j(B_j) - g_j(B_j - 1)\} \\ &= \max_{1 \leq j \leq V} \{ \max_{k_j \leq B_j} [g_j(k_j) - g_j(k_j - 1)] \} \\ &= \max_{1 \leq j \leq V} \{ \max_{k_j \leq B_j} [\bar{g}_j(k_j) - \bar{g}_j(k_j - 1)] \} \end{aligned} \quad (5)$$

It follows from (5) that the greedy algorithm designed for (4) needs never allocate further items to any vendor  $i$  for which  $k_i = B_i$ . Hence this algorithm exactly coincides with the above form for (2). This plainly yields a solution to (4) for which  $k_i \leq B_i, 1 \leq i \leq V$ , and hence solves (2) as well as (4). This concludes the proof.  $\square$

**Comment** The *practical* impact of Lemma 1 is as follows: should cost rate function  $g_i(k_i)$  be convex in some range  $0 \leq k_i \leq B_i, 1 \leq i \leq V$ , and should Cartesian product  $X_{i=1}^V \{0 \leq k_i \leq B_i\}$  contain the solution to problems which are realistic in industrial terms then we will expect greedy heuristics to provide effective solutions to practical problems.

**A finite population queueing model for a single vendor**

Throughout this section, we shall focus on a single service vendor and, in the interests of notational simplicity, shall drop the vendor suffix  $i$ . Our ultimate goal is the elucidation of suitable forms of cost rate  $g(k)$ . In order to achieve that we develop a finite population queueing model for the items requiring repair, awaiting (or receiving) service at the vendor when the latter is assigned  $k$  items.

For fixed population size  $k$ , the repair process at the vendor will be modelled as an  $M/M/1/\infty/k$  queueing, system with arrival rate  $\lambda$ , service rate  $\mu$ , a single server and infinite buffer space. This is as in Opp *et al.*,<sup>1</sup> except that we shall exploit from the outset the fact that, for our purposes, a model with single server and service rate  $\mu$  will yield excellent approximations to more general models with  $s$  servers working in parallel, each with service rate  $\mu/s$ . As it is standard in this area, we shall make Markovian assumptions in which the up time for any item has an exponential distribution with rate  $\lambda$ , while repair times are assumed to be exponentially distributed with rate  $\mu$ . All up times and repair times are taken to be independent. Gross and Harris<sup>9</sup> confirm machine repair as the archetypical application of finite population queueing models.

Consider now a situation in which  $x$  items requiring repair are queued at the vendor, while the remaining  $k-x$  items are up. Standard results indicate that the time until one of the latter breaks down (and requires repair) has an exponential

distribution with rate  $\lambda(k-x)$ . It follows that the repair process at the vendor may be modelled as a birth-death process, where the sets of arrival rates  $\{\lambda_x, 0 \leq x \leq k\}$  and system service rates  $\{\mu_x, 1 \leq x \leq k\}$  are given by

$$\lambda_x = \lambda(k-x), \quad 0 \leq x \leq k$$

and

$$\mu_x = \mu, \quad 1 \leq x \leq k \quad (6)$$

By standard theory, this process is ergodic and has an equilibrium distribution  $\Pi(k) \equiv \{\Pi_x(k), 0 \leq x \leq k\}$  given by

$$\Pi_x(k) = \rho^x \left\{ \prod_{r=0}^{x-1} (k-r) \right\} \Pi_0(k), \quad 0 \leq x \leq k \quad (7)$$

where  $\rho = \lambda/\mu$  in (7) and

$$\Pi_0(k) = \left\{ \sum_{x=0}^k \rho^x \left[ \prod_{r=0}^{x-1} (k-r) \right] \right\}^{-1} \quad (8)$$

Now let  $L(k)$  (queue length) be a random variable such that  $L(k) \sim \Pi(k)$ . The following result is the key to understanding the above repair process. In Lemma 2 we use the following terminology: let  $X$  be a Poisson random variable with parameter (mean)  $\rho^{-1}$  and let  $Y$  be a random variable whose distribution is that of  $X$  conditioned on the event  $X \leq k$ . We say that  $Y$  has the Poisson distribution with parameter  $\rho^{-1}$ , truncated at  $k$ .

**Lemma 2** *The number of up items,  $k-L(k)$ , has the Poisson distribution with parameter  $\rho^{-1}$ , truncated at  $k$ .*

**Proof** Fix  $x$  in the range  $0 \leq x \leq k$ . From (7) and (8) we have that

$$\begin{aligned} P(k-L(k) = x) &= P(L(k) = k-x) = \Pi_{k-x}(k) \\ &= \rho^{k-x} \left\{ \prod_{r=0}^{k-x-1} (k-r) \right\} \\ &\quad \times \left\{ \sum_{y=0}^k \rho^y \left[ \prod_{r=0}^{y-1} (k-r) \right] \right\}^{-1} \\ &= \rho^{-x} (x!)^{-1} \left[ \sum_{y=0}^k \rho^{-y} (y!)^{-1} \right]^{-1} \\ &= \rho^{-x} (x!)^{-1} \exp(-\rho^{-1}) \\ &\quad \times \left[ \sum_{y=0}^k \rho^{-y} (y!)^{-1} \exp(-\rho^{-1}) \right]^{-1} \end{aligned}$$

The result follows.  $\square$

**Comment** The above result implies that when the vendor population  $k$  is large, the number of up items has a distribution which is close to being  $k$ -independent. More specifically, it is close to having a Poisson distribution with

parameter  $\rho^{-1}$ . It follows that when  $k$  is large, the effect in steady state of increasing the vendor population size by a single item is very close to increasing the queue length at the vendor by one. It is this fact which lies behind some of the characterizations of goodwill costs discussed in the next section.

We can in fact say more. In order to do so we introduce some terminology which may be unfamiliar to some readers.

**Definitions** If  $X$  is a random variable with distribution function,  $F$  and  $Y$  a random variable with distribution function  $G$  we say that  $X$  is *stochastically larger* than  $Y$ , written  $X \geq_{st} Y$ , if  $F(x) \leq G(x)$ ,  $x \in \mathbb{R}$ . The sequence of random variables  $\{X_k, k \in \mathbb{N}\}$  is *stochastically increasing* if  $X_{k+1} \geq_{st} X_k$ ,  $k \in \mathbb{N}$ .

**Lemma 3** *The sequences of the number of up items  $\{k - L(k), k \in \mathbb{N}\}$  and the number of down items  $\{L(k), k \in \mathbb{N}\}$  are both stochastically increasing.*

**Proof** The claim in the statement of the lemma in relation to the sequence  $\{k - L(k), k \in \mathbb{N}\}$  is a trivial consequence of Lemma 2. In considering the number of down items, write  $F_k$  for the distribution function of  $L(k)$ . Using Lemma 2 we have that, for any  $x$  in the range  $0 \leq x \leq k$ ,

$$\begin{aligned} 1 - F_k(x) &= P\{k - L(k) \leq k - x\} \\ &= \left[ \sum_{y=0}^{k-x} \rho^{-y} (y!)^{-1} \exp(-\rho^{-1}) \right] \\ &\quad \times \left[ \sum_{y=0}^k \rho^{-y} (y!)^{-1} \exp(-\rho^{-1}) \right]^{-1} \\ &= A(k-x)[A(k)]^{-1} \end{aligned}$$

where

$$A(k) \equiv \sum_{y=0}^k \rho^{-y} (y!)^{-1} \exp(-\rho^{-1})$$

It is straightforward to show that

$$A(k-x)[A(k)]^{-1} \leq A(k+1-x)[A(k+1)]^{-1}$$

if and only if

$$\begin{aligned} \left[ \sum_{y=0}^{k-x} \rho^{-y} (y!)^{-1} \right] \rho^{-k-1} [(k+1)!]^{-1} &\leq \left[ \sum_{y=0}^k \rho^{-y} (y!)^{-1} \right] \\ \times \rho^{-k-1+x} [(k+1-x)!]^{-1} & \end{aligned}$$

if and only if

$$\begin{aligned} \sum_{y=0}^{k-x} \rho^{-k-1-y} \{x!(k+1)!\}^{-1} \\ \leq \sum_{y=0}^{k-x} \rho^{-k-1-y} \{(x+y)!(k+1-x)!\}^{-1} \end{aligned}$$

However, the latter inequality is a trivial consequence of the fact that

$$(x+y)!(k+1-x)! \leq (k+1)!y!, 0 \leq y \leq k-x$$

It follows from the above calculation that

$$1 - F_k(x) \leq 1 - F_{k+1}(x) \Rightarrow F_k(x) \geq F_{k+1}(x), 0 \leq x \leq k$$

and hence trivially that

$$F_k(x) \geq F_{k+1}(x), x \in \mathbb{R}$$

It now follows that  $L(k+1) \geq_{st} L(k)$ , as required.  $\square$

**Comment** It follows from Lemma 3 that, for our models, both of the contributing components to the overall cost rate will increase as the population size  $k$  grows. These components are (i) the repair cost rate, which is related to  $k - L(k)$ , the number of up items (hence vulnerable to breakdown), and (ii) the goodwill cost rate, which is related to  $L(k)$ , the number of down items.

It is imperative that we have efficient means of computing equilibrium distributions  $\Pi(k)$ ,  $k \in \mathbb{Z}^+$ , and quantities related to them.

Firstly note that the reciprocal of  $\Pi_0(k)$ , written  $\Pi_0^{-1}(k)$ , satisfies the recursion

$$\begin{aligned} \Pi_0^{-1}(0) &= 1 \\ \Pi_0^{-1}(k) &= 1 + \rho k \Pi_0^{-1}(k-1), k \in \mathbb{Z}^+ \end{aligned}$$

See (8). Further, we can re-express (7) by writing

$$\Pi_x(k) = B_x(k) \Pi_0(k), 0 \leq x \leq k$$

where the quantity

$$B_x(k) = \rho^x \left\{ \prod_{r=0}^{x-1} (k-r) \right\}$$

satisfies the recursion

$$B_0(k) = 1$$

$$B_x(k) = \rho(k-x+1)B_{x-1}(k), 0 < x \leq k$$

Finally, we consider the computation of the mean of distribution  $\Pi(k)$ , denoted  $\bar{L}(k)$ . We write

$$\bar{L}(k) = B(k) \Pi_0(k)$$

where

$$B(k) = \sum_{x=0}^k xB_x(k), k \in \mathbb{Z}^+$$

It is straightforward to show algebraically that  $B(k)$  satisfies the recursion

$$B(0) = 0$$

$$B(k) = \rho k \Pi_0^{-1}(k-1) + \rho k B(k-1), k \in \mathbb{Z}^+$$

With the above in place, we can now proceed to the development of cost rate functions  $g(k)$ . We shall consider two different approaches to the modelling of goodwill costs, labelled Models 1 and 2.

**Model 1** Under Model 1, a fixed cost  $c$  is incurred for every repair. In addition, a goodwill cost  $d$  is incurred whenever an item requiring repair spends more than time  $\tau$  at the vendor. It is natural to assume that all vendors share a common goodwill cost parameter  $d$ , although this is not required for our approach. Use  $W$  for the waiting time (including service) of items in steady state. When an item arriving at the vendor finds a queue of length  $x$ , its waiting time will have a gamma distribution with shape parameter  $x+1$  and scale parameter  $\mu$ . We express this by writing

$$W|L(k) = x \sim \Gamma(x+1, \mu)$$

It follows from the standard properties of the gamma distribution that

$$dP(W > \tau | L(k) = x) = d \sum_{y=0}^x (\mu\tau)^y (y!)^{-1} \exp(-\mu\tau) \equiv a(x), \text{ say} \tag{9}$$

It follows from (9) that the expected cost rate incurred in steady state from both repairs and goodwill costs when the vendor population is  $k$  may be expressed as

$$\sum_{x=0}^k \Pi_x(k) \lambda(k-x) \{c + a(x)\} \equiv g^1(k), \text{ say} \tag{10}$$

**Model 2** As in Model 1, a fixed cost  $c$  is incurred for every repair. Now, however, we shall suppose that a cost of  $d(w-\tau)^+$  is incurred when an item spends time  $w$  at the vendor. As in Model 1, it is natural though not mandatory to assume that all vendors share a common  $d$ . Hence, as before, no penalty is incurred provided the item spends less time than  $\tau$  at the vendor, but now we suppose that the goodwill cost is proportional to the excess over  $\tau$  of the waiting time. It follows from standard properties of the gamma distribution and of

expectation that

$$\begin{aligned} dE\{(W-\tau)^+ | L(k) = x\} &= d \int_{\tau}^{\infty} P(W > w | L(k) = x) dw \\ &= d(\mu)^{-1} \sum_{y=0}^x (x+1-y)(\mu\tau)^y (y!)^{-1} \exp(-\mu\tau) \\ &= b(x), \text{ say} \end{aligned} \tag{11}$$

It follows from (11) that the expected cost rate incurred in steady state from both repairs and goodwill costs when the vendor population is  $k$  may be expressed as

$$\sum_{x=0}^k \Pi_x(k) \lambda(k-x) \{c + b(x)\} \equiv g^2(k), \text{ say} \tag{12}$$

**The cost rate functions, monotonicity, convexity and optimal allocations**

Following the last two sections, our optimal allocation problem for Model 1 can be expressed by

$$\begin{aligned} \min \quad & \sum_{i=1}^V g_i^I(k_i) \\ \text{s.t.} \quad & \sum_{i=1}^V k_i = K \\ & k_i \in \mathbb{Z}^+, i = 1, \dots, V \end{aligned} \tag{13}$$

where  $I=1,2$ . From the earlier discussion of greedy algorithms it is plain that the status of greedy heuristics as solutions to (13) will relate to the degree of convexity (suitably defined) of the expected cost rate functions  $g_i^I$ ,  $1 \leq i \leq V$ ,  $I=1,2$ . This section will contain a discussion of this issue. To further this discussion we shall again drop the vendor subscript  $i$ .

Before proceeding to our Models 1 and 2, we pause to point out that the linear holding cost model of Opp *et al*<sup>1</sup> has an associated cost rate function  $g^L$  given by

$$\begin{aligned} g^L(k) &= \sum_{x=0}^k \Pi_x(k) \{\lambda(k-x)c + hx\} \\ &= \lambda kc + (h-\lambda c)\bar{L}(k) \end{aligned} \tag{14}$$

The known convexity of  $\bar{L}(k)$  for the finite population queueing model of the previous section guarantees the convexity of  $g^L$  when  $h > \lambda c$ . Under this condition (satisfied for all vendors) the standard greedy heuristic must indeed solve the optimal allocation problem for the Opp *et al*<sup>1</sup> model. Of direct relevance to our Models 1 and 2 is the related fact that the repair cost element of expected cost rates  $g^L$ ,  $g^1$  and  $g^2$ , namely

$$\sum_{x=0}^k \Pi_x(k) \lambda(k-x)c = \lambda kc - \lambda c \bar{L}(k) \tag{15}$$

is concave in  $k$ .

While this concavity is a concern, we would argue (along with Opp *et al*<sup>1</sup>) that the goodwill costs contribution to the expected cost rate functions  $g^1$  and  $g^2$  should predominate. These latter contributions will be strongly influenced by the nature, respectively, of the queue length-dependent cost sequences  $\{a(0), a(1), a(2), \dots\} \equiv \{a(x), x \in \mathbb{N}\}$  and  $\{b(0), b(1), b(2), \dots\} \equiv \{b(x), x \in \mathbb{N}\}$  introduced in (9) and (11). The next result enunciates some important properties of these sequences. Note that we say that some function  $f : \mathbb{N} \rightarrow \mathbb{R}$  is convex/concave if there is some  $B \in \mathbb{Z}^+$  such that  $f$  is convex over the range  $[0, 1, \dots, B]$  and concave over  $[B, B + 1, \dots]$ .

**Lemma 4** (i) *Model 1: The sequence  $\{a(x), x \in \mathbb{N}\}$  is increasing and convex/concave;*  
 (ii) *Model 2: The sequence  $\{b(x), x \in \mathbb{N}\}$  is increasing and convex.*

**Proof** For (i), see (9) and note that it is a well-known property of the Poisson distribution that the difference

$$a(x + 1) - a(x) = d(\mu\tau)^{x+1} \{(x + 1)!\}^{-1} \exp(-\mu\tau) \geq 0$$

is increasing/decreasing in  $x$ . The increasing and convex/concave nature of  $\{a(x), x \in \mathbb{N}\}$  follows.

For (ii), observe that

$$b(x + 1) - b(x) = d(\mu)^{-1} \sum_{y=0}^{x+1} (\mu\tau)^y (y!)^{-1} \exp(-\mu\tau)$$

is increasing in  $x$  and always positive. The conclusions in (ii) follow easily.  $\square$

We can now use Lemmas 3 and 4 to show that the expected cost rates for both Models 1 and 2 are increasing in the vendor population  $k$ . In order to accomplish this we shall need to invoke the standard result that if  $X \geq_{st} Y$  then  $E\{\Phi(X)\} \geq E\{\Phi(Y)\}$  for any increasing function  $\Phi$ .

**Proposition 5** *Expected cost rate functions  $g^I(k)$  are increasing in  $k, I=1, 2$ .*

**Proof** Take  $I=1$  for definiteness. The proof for  $I=2$  is along identical lines. First, note from (7) that

$$\Pi_x(k)\lambda(k-x) = \Pi_{x+1}(k)\mu, 0 \leq x \leq k-1 \quad (16)$$

Hence, from (10), the formula for  $g^1(k)$  may be rewritten

$$\begin{aligned} & \mu c \sum_{x=0}^{k-1} \Pi_{x+1}(k) + \mu \sum_{x=0}^{k-1} \Pi_{x+1}(k)a(x) \\ & = \mu E[\Phi_1\{L(k)\} + \Phi_2\{L(k)\}] \end{aligned}$$

where

$$\Phi_1(x) = \begin{cases} c, & x \geq 1 \\ 0, & x < 1 \end{cases}$$

and

$$\Phi_2(x) = \begin{cases} a(x-1), & x \geq 1 \\ 0, & x < 1 \end{cases}$$

However, from Lemma 4, both  $\Phi_1$  and  $\Phi_2$  are increasing. The result now follows from Lemma 3 and the property of stochastic ordering enunciated before the statement of the proposition.  $\square$

Before proceeding further, we pause to point out that both theoretical considerations and computational experience testify to the fact that the goodwill costs elements of  $g^1$  and  $g^2$ , namely

$$h^1(k) \equiv \sum_{x=0}^k \Pi_x(k)\lambda(k-x)a(x)$$

and

$$h^2(k) \equiv \sum_{x=0}^k \Pi_x(k)\lambda(k-x)b(x)$$

strongly reflect the convexity properties reported in Lemma 4. For example, utilizing (16) and Lemma 2 we have that

$$\begin{aligned} h^1(k) &= \mu \sum_{x=0}^{k-1} \Pi_{x+1}(k)a(x) \\ &\sim \mu \sum_{r=0}^{\infty} \rho^{-r}(r!)^{-1} \exp(-\rho^{-1})a(k-1-r) \\ &\equiv H^1(k) \end{aligned} \quad (17)$$

where we take  $a(x) = 0, x < 0$  in (17). Similarly we have

$$\begin{aligned} h^2(k) &\sim H^2(k) \\ &= \mu \sum_{r=0}^{\infty} \rho^{-r}(r!)^{-1} \exp(-\rho^{-1})b(k-1-r) \end{aligned}$$

From Lemma 2 the approximating goodwill cost rates  $H^I$  satisfy

$$\lim_{k \rightarrow \infty} \{h^I(k) - H^I(k)\} = 0, I = 1, 2$$

It is straightforward to infer from Lemma 4 that approximating goodwill cost rate  $H^2$  is convex. In the case of  $H^1$  it is possible to assert the existence of positive integer  $B$  such that  $H^1$  is convex over the range  $[0, 1, \dots, B]$ . However, our computational experience suggests that we can say more. In fact, in all our computational experiments,  $h^1(k)$  has been found to be convex/concave.

Further,  $h^2(k)$  has always been close to convex. In Figures 1 and 2 see typical plots of  $d^{-1}h^1$  and  $d^{-1}h^2$ , respectively, for a case with  $\lambda = 1.2$ ,  $\tau = 0.04$  and  $\mu = 62.5$ . These problem instances are discussed more extensively in the next section ( $K = 100$ ). Note that  $d^{-1}h^1$  is convex up to  $B = 52$  and concave beyond while  $d^{-1}h^2$  is convex throughout the range  $[0, 100]$ . These findings are replicated in the top lines of Tables 1 and 2 following. We now proceed to reflect on how the (concave) repair cost element identified in (15) impacts the convexity of the expected cost rate functions  $g^1$  and  $g^2$ .

To further the discussion, we introduce additional terms and notations. To explore dependence upon repair cost  $c$ , we shall expand the notation for the expected cost rate functions to  $g_c^I, I = 1, 2$ .




Figure 1 A plot of  $d^{-1} \times$  (goodwill cost rate) against allocated items  $k$  for Model 1.




Figure 2 A plot of  $d^{-1} \times$  (goodwill cost rate) against allocated items  $k$  for Model 2.

Table 1 Convexity boundary points for Model 1 when  $c = 0$  and 2

Vendor	Scenario	$B_0^1$	$B_2^1$
1	1	52	51
1	2	61	60
1	3	71	71
1	4	84	83
1	5	98	97
1	6	100 +	100 +
4	1	52	51
4	2	44	43
4	3	35	34
4	4	27	26
4	5	19	18
4	6	11	10

Table 2 Convexity boundary points for Model 2 when  $c = 0$  and 2

Vendor	Scenario	$B_0^2$	$B_2^2$
1	1	100 +	100 +
1	2	100 +	100 +
1	3	100 +	100 +
1	4	100 +	100 +
1	5	100 +	100 +
1	6	100 +	100 +
4	1	100 +	100 +
4	2	100 +	100 +
4	3	95	95
4	4	82	80
4	5	68	68
4	6	55	55

Definitions We say that expected cost rate function  $g_c^I$  is initially convex if

$$g_c^I(2) \geq 2g_c^I(1), I = 1, 2$$

If  $g_c^I$  is initially convex then its convexity boundary point  $B_c^I$  is given by

$$B_c^I = \max\{k; g_c^I(l) + g_c^I(l - 2) \geq 2g_c^I(l - 1), 2 \leq l \leq k\}$$

If  $g_c^I$  is not initially convex then we assign  $B_c^I$  the value 1.

Proposition 6 (i) Expected cost rate function  $g_c^I$  is initially convex if and only if  $c \leq c^I$ , where

$$c^I = d\{(I + \mu\tau)(1 + \rho) - (1 + 2\rho)\} \times \rho^{-1} \exp(-\mu\tau), I = 1, 2 \tag{18}$$

- (ii)  $c^1 \leq c^2$
- (iii) Convexity boundary point  $B_c^I$  is decreasing in  $c, I = 1, 2$ .

**Proof** (i) For definiteness consider Model 1. Expected cost rate function  $g_c^I$  is initially convex if and only if  $g_c^I(2) \geq 2g_c^I(1)$ , which becomes, upon evaluation of these quantities

$$\begin{aligned} & \{2\rho[c + a(0)] + 2\rho^2[c + a(1)]\}(1 + 2\rho + 2\rho^2)^{-1} \\ & \geq 2\rho[c + a(0)](1 + \rho)^{-1} \\ & \Leftrightarrow [c + a(1)][c + a(0)]^{-1} \\ & \geq (1 + 2\rho)(1 + \rho)^{-1} \end{aligned} \tag{19}$$

This becomes, upon utilization of the form of  $a(x)$  from (9),

$$\begin{aligned} & [d + d\mu\tau + c \exp(\mu\tau)][d + c \exp(\mu\tau)] \\ & \geq (1 + 2\rho)(1 + \rho)^{-1}. \end{aligned} \tag{20}$$

The conclusion in (i) for the  $I = 1$  case now follows trivially from (20). The case  $I = 2$  is dealt with similarly.

(ii) This follows immediately from (i).

(iii) Consider two distinct repair costs  $\bar{c} < \bar{c}'$ . Firstly, note that it follows from (i) that if  $\bar{c} > c^I$  then we must have  $B_{\bar{c}}^I = 1$  and  $B_{\bar{c}'}^I \geq B_{\bar{c}}^I$  trivially. Suppose now that  $\bar{c} \leq c^I$  and hence that  $B_{\bar{c}}^I \geq 2$ . Explicitly, we have that

$$g_{\bar{c}}^I(k) + g_{\bar{c}}^I(k - 2) \geq 2g_{\bar{c}}^I(k - 1), 2 \leq k \leq B_{\bar{c}}^I, I = 1, 2 \tag{21}$$

But, from the definition of the cost rate functions in the previous section, it follows that:

$$g_{\bar{c}'}^I(k) = g_{\bar{c}}^I(k) + (\bar{c}' - \bar{c}) \sum_{x=0}^k \Pi_x(k) \lambda(k - x), I = 1, 2 \tag{22}$$

where from (15), the second term the right-hand side of (22) is concave in  $k$ . It then follows from (21) and (22) that:

$$\begin{aligned} g_{\bar{c}'}^I(k) + g_{\bar{c}'}^I(k - 2) &= g_{\bar{c}}^I(k) + g_{\bar{c}}^I(k - 2) \\ &+ (\bar{c}' - \bar{c}) \sum_{x=0}^k \Pi_x(k) \lambda(k - x) \\ &+ (\bar{c}' - \bar{c}) \sum_{x=0}^{k-2} \Pi_x(k - 2) \lambda \\ &\times (k - 2 - x) \\ &\geq 2g_{\bar{c}}^I(k - 1) + 2(\bar{c}' - \bar{c}) \\ &\times \sum_{x=0}^{k-1} \Pi_x(k - 1) \lambda(k - 1 - x) \\ &= 2g_{\bar{c}'}^I(k - 1), 2 \leq k \leq B_{\bar{c}'}^I, I = 1, 2 \end{aligned} \tag{23}$$

It now follows from (23) that  $B_{\bar{c}'}^I \geq B_{\bar{c}}^I, I = 1, 2$  and the result follows.  $\square$

**Comments**

(1) Proposition 6 gives a formal mathematical expression to a range of ideas which have an important bearing on the convexity of our cost rate functions and consequential

likely status of greedy heuristics. First, the reader should note that the conditions expressed in (18) will be met comfortably for sensibly chosen model parameters and hence our cost rate functions can be safely assumed to be at least initially convex. Proposition 4(ii) is an indicator of the fact that the expected cost rate  $g^2$  enjoys stronger convexity properties than does  $g^1$  when a common cost parameter  $d$  is chosen for Models 1 and 2. Proposition 4(iii) points to the fact that the initial range in which the cost rate functions are increasing, convex will decrease as the repair cost element increases. For the linear holding cost model, Opp *et al*<sup>1</sup> were able to insist that the repair cost  $c$  should be sufficiently small that the resulting overall expected cost rate  $g^I$  remain increasing, convex and hence that greedy heuristics are guaranteed to provide optimal solutions to the allocation problem. This is not possible for our more complex models. Rather, we can comparably require that repair costs  $c$  are sufficiently small that the corresponding convexity boundary points for the cost rates are large (see Proposition 4(ii)) and hence, following Lemma 1, that greedy heuristics are likely to perform well for the allocation problem. Our numerical results indicate that this happens much of the time.

(2) Many of the essential insights of the above study will continue to hold for a range of departures from our stated model assumptions. For example, if up times and/or service times are not exponentially distributed then expected cost rates can be obtained numerically for a wide range of models and may be obtained via simulation more generally. Our experience is that when this is performed, expected cost rates continue to have properties very close to those described in Propositions 5 and 6, and the greedy heuristic for the distribution of warranty work continues to perform strongly. Expected cost rates may be obtained numerically for Markov models in which item up times and repair times both have phase-type distributions. This is a general class of which gamma distributions with integer shape parameters are special cases. The resulting finite population queueing system can then be modelled as a multi-dimensional continuous time Markov process and a wide range of numerical methods are available for the computation of the equilibrium distribution and expected cost rates. We have also found that allocations based on expected cost rates calculated using approximating models with exponential up times and repair times often perform extremely well. For example, the case reported as scenario 3,  $K = 100$  in Table 5 is based on a model which assumes exponential service times with rates  $\mu_i = 84.7(0.8)^{i-1}, 1 \leq i \leq 4$ , for the four vendors. Consider now a model in which service times at vendor  $i$  have a gamma distribution with shape parameter 2 and scale parameter  $2\mu_i, 1 \leq i \leq 4$ . These distributions have the same means as the above exponential distributions

but the variances are halved. Calculations based on estimates of cost rates for the gamma model obtained via simulation result in a greedy allocation of (41, 28, 19, 12) which is confirmed to be optimal by dynamic programming. The greedy allocation for the approximating model with exponential service times is (40, 29, 19, 12).

**Numerical investigation**

In this section, we present some results from our computational study of the static allocation problem. Our prime goal here is to assess the status of the greedy algorithm as a heuristic. To achieve this, in all cases studied, the overall cost rate achieved by the greedy algorithm has been compared with the optimum cost rate obtained by dynamic programming (DP). The scope of the numerical study is inevitably limited by the inability of DP to produce solutions to large problems; this indeed is a prime motivation of our study based on the greedy heuristic. However, the performance of the greedy heuristic is of such quality in the cases we have studied that we are confident of its effectiveness beyond our problem set. Subsidiary goals of the computational study include exploration of the sensitivity of optimal/greedy allocations to choice of key problem parameters. Of particular interest is the evidence that we have obtained that optimal/greedy allocations are insensitive to the choice of  $d$ , the key parameter for goodwill costs in both Models 1 and 2. Given the practical difficulty in assigning values to such parameters, this is an encouraging feature of our modelling approach.

We have followed Opp *et al*<sup>1</sup> in making some parameter choices on the basis of information acquired in an industrial context in which the items under warranty are PCs. In all cases studied we shall suppose that there are four vendors ( $n=4$ ) and that the fixed per item failure rate is  $\lambda=1.2$  breakdowns per year. We shall take the critical time  $\tau$  to be working days ( $\tau=0.04$  years). Further, we shall consider a range of (six) scenarios regarding the relative service rates of the four vendors. The service rate for vendor  $i$  in scenario  $j$  will be given by

$$0.625K, \quad 1 \leq i \leq 4, j = 1$$

$$2.5Kx_j^{i-1}(1-x_j)(1-x_j^4)^{-1}, \quad 1 \leq i \leq 4, 2 \leq j \leq 6 \quad (24)$$

where the  $x_j$ 's are chosen as follows:

$j$	2	3	4	5	6
$x_j$	0.9	0.8	0.7	0.6	0.5

Hence in scenario 1, all vendors have the same service rates while these rates become more unequal as the scenario number increases. It is always true that vendor 1 has the best service rate, followed by vendor 2, and so on. The

service rates always sum to  $2.5K$  where  $K$  is the size of the population under warranty, taken in our studies to be 100, 200, 300 or 500.

Tables 1 and 2 aim to give readers a feel for how the critical convexity boundary points  $B_c^l$  (see Lemma 1 and Proposition 6) vary with problem parameters. For both tables we have  $K=100$  and consider service rates for vendors 1 and 4 for all six scenarios, as given in (24). Table 1 concerns Model 1 and has goodwill cost parameter  $d=10$  while Table 2 is for Model 2 with  $d=1000$ . These choices of cost parameters give a broadly similar balance between repair costs and goodwill costs for the two models for the problems chosen. Please note that in our study only the relative sizes of  $c$  and  $d$  are of significance. However, it may help the reader to think of the basic cost unit being around \$100.

To obtain these boundary points we implemented in Matlab the computational schemes for  $g^1$  and  $g^2$ , given in the section on our finite population queueing model and identified the points directly. Any case for which the boundary point is beyond the value of  $K=100$  is recorded as  $100+$ . To interpret Tables 1 and 2, recall that the service rates for vendor 1 are maximal and increasing in the scenario number and that the service rates for vendor 4 are minimal and decreasing in the scenario number. A boundary point as low as 10 in Table 1 needs not be a concern for the greedy heuristic since in any event for this scenario, vendor 4 has a very low service rate capacity and is likely to receive a very small allocation of items. The evidence is clear from Tables 1 and 2 that the boundary points are increasing in the corresponding service rates, as (presumably) are the optimal item allocations. Further, the boundary points for Model 2 are all higher than for Model 1 and those for  $c=0$  higher than for  $c=2$ . All this is as we would expect from the previous sections. Finally, it is true for scenarios 1 and 2 in Model 2 that all vendors have boundary points of  $100+$  for both  $c=0$  and 2. From Gross,<sup>4</sup> this guarantees the optimality of the greedy heuristic for these cases.

Table 3 contains the item allocations obtained under Model 1 from the greedy heuristic for cases in which  $c_i=1$ ,  $1 \leq i \leq 4$ , namely that all vendors have an associated repair cost of 1. We take  $d=10$ . We also show the breakdown of the associated overall cost rate between repair costs and goodwill costs. Table 4 presents an equivalent set of results for problems in which, for each separate case, repair costs have been sampled independently from a uniform  $U(0.90, 1.10)$  distribution. Note that the instances studied range from those in which goodwill costs dominate ( $K=100$ ) to those in which repair costs do ( $K=500$ ). Comparing the allocations in the tables make it clear that the greedy allocations are somewhat sensitive to modest variations in the repair costs between vendors, especially for the  $K=500$  cases. We have, however, seen little evidence of sensitivity to the choice of goodwill cost parameter  $d$ . Running the cases

**Table 3** Greedy allocation and associated cost rates for Model 1 when  $c_i = 1, 1 \leq i \leq 4$ , and  $d = 10$

Scenario	K	Greedy allocation	Repair cost rate ( $\times 10^{-2}$ )	Goodwill cost rate ( $\times 10^{-2}$ )
1	100	25 25 25 25	1.1607	3.0491
2	100	31 27 23 19	1.1610	3.0218
3	100	38 28 20 14	1.1619	2.9187
4	100	47 29 16 8	1.1636	2.6993
5	100	58 28 12 2	1.1665	2.3238
6	100	69 25 6 0	1.1705	1.8116
1	500	125 125 125 125	5.9569	0.0415
2	500	162 135 112 91	5.9571	0.0400
3	500	207 142 94 57	5.9579	0.0351
4	500	259 144 70 27	5.9595	0.0264
5	500	315 136 43 6	5.9625	0.0147
6	500	375 111 14 0	5.9667	0.0049

**Table 4** Greedy allocation and associated cost rates for Model 1 when repair costs are drawn independently from a  $U(0.90, 1.10)$  distribution, and  $d = 10$

Scenario	K	Greedy allocation	Repair cost rate ( $\times 10^{-2}$ )	Goodwill cost rate ( $\times 10^{-2}$ )
1	100	25 25 25 25	1.1557	3.0491
2	100	31 27 23 19	1.1327	3.0218
3	100	38 28 20 14	1.1783	2.9187
4	100	47 29 16 8	1.1921	2.6993
5	100	58 28 12 2	1.1810	2.3238
6	100	70 25 5 0	1.1817	1.8121
1	500	136 103 130 131	5.9094	0.0536
2	500	179 106 119 96	5.7682	0.0612
3	500	190 135 103 72	6.0365	0.0493
4	500	265 158 44 33	6.0777	0.0406
5	500	253 169 63 15	5.9630	0.0656
6	500	380 120 0 0	6.0076	0.0068

in Table 3 for  $d=100$ , for example, leaves the greedy allocations virtually unchanged.

Tables 5 and 6 contain the details of the greedy allocations made under Model 2 for an equivalent set cases to those covered by Tables 3 and 4 for Model 1. However, now the goodwill cost parameter is set at  $d=1000$ . Again, the instances studied range from those in which the goodwill costs dominate ( $K=100$ ) to those in which repair costs do ( $K=500$ ). Many of the main features of Tables 3 and 4 are reflected in Tables 5 and 6. In substantiation of the claim that greedy allocations are insensitive to the setting of goodwill cost parameter  $d$ , see Table 7. The cases studied here are as in Table 5 but with  $d=10000$ . Note that, increasing  $d$  by a factor of 10 has resulted in modest changes to four of the greedy allocations. The corresponding goodwill cost rates are consequently approximately increased by a factor of 10 with the repair cost rates substantially unchanged.

**Table 5** Greedy allocation and associated cost rates for Model 2 when  $c_i = 1, 1 \leq i \leq 4$ , and  $d = 1000$

Scenario	K	Greedy allocation	Repair cost rate ( $\times 10^{-2}$ )	Goodwill cost rate ( $\times 10^{-2}$ )
1	100	25 25 25 25	1.1607	7.7542
2	100	32 27 22 19	1.1610	7.6458
3	100	40 29 19 12	1.1618	7.2602
4	100	50 29 15 6	1.1635	6.5058
5	100	61 28 10 1	1.1663	5.2891
6	100	72 24 4 0	1.1703	3.7177
1	500	125 125 125 125	5.9569	0.0414
2	500	165 136 111 88	5.9570	0.0405
3	500	211 144 92 53	5.9577	0.0369
4	500	266 146 68 20	5.9589	0.0297
5	500	325 138 37 0	5.9615	0.0160
6	500	389 111 0 0	5.9652	0.0056

**Table 6** Greedy allocation and associated cost rates for Model 2 when repair costs are drawn independently from a  $U(0.90, 1.10)$  distribution, and  $d = 1000$

Scenario	K	Greedy allocation	Repair cost rate ( $\times 10^{-2}$ )	Goodwill cost rate ( $\times 10^{-2}$ )
1	100	25 25 25 25	1.1557	7.7542
2	100	32 27 22 19	1.1321	7.6458
3	100	40 29 19 12	1.1821	7.2602
4	100	50 29 15 6	1.1918	6.5058
5	100	61 28 10 1	1.1848	5.2891
6	100	72 24 4 0	1.1809	3.7177
1	500	138 98 132 132	5.9045	0.0564
2	500	182 102 120 96	5.7621	0.0638
3	500	191 135 103 71	6.0382	0.0541
4	500	273 162 37 28	5.9923	0.0464
5	500	258 172 61 9	5.9970	0.0708
6	500	384 116 0 0	6.0074	0.0053

**Table 7** Greedy allocation and associated cost rates for Model 2 when  $c_i = 1, 1 \leq i \leq 4$ , and  $d = 10,000$

Scenario	K	Greedy allocation	Repair cost rate ( $\times 10^{-2}$ )	Goodwill cost rate ( $\times 10^{-2}$ )
1	100	25 25 25 25	1.1607	7.7542
2	100	32 27 22 19	1.1610	7.6458
3	100	40 29 19 12	1.1618	7.2602
4	100	50 9 15 6	1.1635	6.5058
5	100	61 28 10 1	1.1663	5.2891
6	100	72 24 4 0	1.1703	3.7177
1	500	125 125 125 125	5.9569	0.0414
2	500	164 136 111 89	5.9571	0.0404
3	500	211 144 92 53	5.9577	0.0369
4	500	264 146 69 21	5.9591	0.0296
5	500	322 139 39 0	5.9618	0.0159
6	500	381 115 4 0	5.9660	0.0052

Most impressive, however, is the fact that of the 450 problem instances we have studied, in only one was the greedy allocation *not* optimal. This was the case of Model 1 with  $K=500$  and  $d=10$  for which the repair costs were drawn from a  $U(0.50, 1.50)$  distribution. The greedy allocation for this case was (286, 177, 31, 6) while the optimal allocation was (287, 177, 30, 6). The overall cost rates for these allocations differed by just 0.02%.

### Concluding remarks

We have proposed static allocation models for the optimal distribution of warranty work among a collection of service vendors. In the development we have argued the importance of an approach to the modelling of goodwill costs which takes explicit account of the delays experienced by customers. While the expected cost rates which arise within our models are such that no claim to optimality of simple heuristics can be made in any generality, nevertheless a range of evidence (both theoretical and numerical) is adduced in support of the strong performance of greedy approaches to work allocation.

Among plans for future work is the development and analysis of dynamic allocation models where (in contrast to the present model) decisions are made about who should undertake a particular repair in light of the numbers of items already queued for repair at each vendor. This will result in a finite population queueing control problem involving the dynamic routing of repairs to vendors to minimize a cost rate related to delays in receiving service. This is a formidable research challenge. The question of interest which will arise will concern

whether substantial cost savings are available from the adoption of a dynamic approach to the assignment of work rather than a static one. If so, the additional administrative overhead necessitated by such an approach may be worthwhile.

*Acknowledgements*—The first author acknowledges support received from the University of Edinburgh through a research studentship.

### References

- 1 Opp M, Adan I, Kulkarni VG and Swaminathan JM (2003). *Outsourcing warranty repairs: static allocation*. Technical Report, University of North Carolina, Chapel Hill, NC.
- 2 *CNet news.com*, September 2002.
- 3 Serant C (2001). Solectron to Provide Xbox Support. *EBN*, October 22.
- 4 Gross O (1956). A class of discrete type minimization problems. Technical Report RM-1644, RAND Corp.
- 5 van Meighem J (1995). Dynamic scheduling with convex delay costs: the generalized  $c\mu$  rule. *Ann Appl Prob* 5: 809–833.
- 6 Ansell PS, Glazebrook KD, Nino-Mora J and O’Keeffe M (2003). Whittle’s index policy for a multiclass queueing system with convex holding costs. *Math Meth Oper Res* 57: 21–39.
- 7 Fox BL (1996). Discrete optimization via marginal analysis. *Mngt Sci* 13: 210–216.
- 8 Ibaraki T and Katoh N (1988). *Resource Allocation Problems: Algorithmic Approaches*. MIT Press: Cambridge, MA.
- 9 Gross D and Harris CM (1998). *Fundamentals of Queueing Theory*. Wiley: New York NY.

Received January 2004;  
accepted September 2004 after one revision

## Bibliography

- [1] P. S. Ansell, M. J. Dacre, K. D. Glazebrook, and C. Kirkbride. Optimal load balancing and scheduling in distributed multi-class service system. Technical Report, Newcastle University, 2001.
- [2] P. S. Ansell, K. D. Glazebrook, J. Nino-Mora, and M. O’Keeffe. Whittles’s index policy for a multi-class queueing system with convex holding costs. *Math. Meth. Oper. Res.*, 57:21–39, 2003.
- [3] R. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, 1957.
- [4] D. A. Berry and B. Fristedt. *Bandits Problems-Sequential Allocations of Experiments*. Chapman and Hall, London-New York, 1985.
- [5] D. Blackwell. Discrete dynamic programming. *Ann. Math. Statist.*, 33:719–726, 1962.
- [6] D. Blackwell. Discounted dynamic programming. *Ann. Math. Statist.*, 36:226–235, 1965.
- [7] S. Briskman. It’s all about right-sourcing. October. <http://services.silicon.com>, 2005.
- [8] P. S. Buczkowski, M. E. Hartmann, and V. G. Kulkarni. Outsourcing prioritized warranty repairs. *International Journal of Quality and Reliability Management*, 22:699–714, 2005.
- [9] R. L. Burden and J. D. Faires. *Numerical Analysis (7<sup>th</sup> ed)*. Books Cole, Belmont, CA., 2000.
- [10] Business Outsourcing Corporation. Case study #2. <http://www.businessoutsourcing.com>, 2006.
- [11] P. M. Byrne. Making warranty management manageable. *Logistics Management*, 43(i8):31–32, August 2004.

- [12] F. Chen and V. G. Kulkarni. Dynamic routing of prioritized warranty repairs. *Nav. Res. Log.*, forthcoming.
- [13] S. Chukova and M. R. Johnston. Two-dimensional warranty repair strategy based on minimal and complete repairs. *Mathematical and Computer Modelling*, 44:1133–1143, 2006.
- [14] Deloitte Consulting. Calling a change in the outsourcing market. <http://www.deloitte.com>, 2005.
- [15] L. Ding and K. D. Glazebrook. A static allocation model for the outsourcing of warranty repairs. *J. Oper. Res. Soc.*, 56:825–835, 2005.
- [16] R. D. Foley and D. R. McDonald. Join the shortest queue: stability and exact asymptotics. *Ann. Appl. Prob.*, 11:569–607, 2001.
- [17] B. L. Fox. Discrete optimization via marginal analysis. *Mgmt. Sci.*, 13:210–216, 1966.
- [18] E. Frostig and G. Weiss. Four proofs of Gittins' multi-armed bandit theorem. Technical Report, University of Haifa, Mount Carmel, Israel., Nov. 1999.
- [19] R. Garbe and K. D. Glazebrook. Reflections on a new approach to Gittins indexation. *J. Oper. Res. Soc.*, 47:1301–1309, 1996.
- [20] J. C. Gittins. Bandits processes and dynamic allocation indices. *J.R. Statist. Soc.B*, 41:148–177, 1979.
- [21] J. C. Gittins and D.M. Jones. A dynamic allocation index for the sequential design of experiments. *In Progress in Statistics Colloq. Math. Soc. Janos Bolyai*, 9:241–266, 1974.
- [22] G. J. Glasser. Variance formulas for the mean difference and coefficient of concentration. *J. Amer. Statist. Assoc.*, 57(299):648–654, 1962.
- [23] K. D. Glazebrook, P. S. Ansell, R. T. Dunn, and R. R. Lumley. On the optimal allocation of service to impatient tasks. *J. Appl. Prob.*, 41:51–72, 2004.
- [24] K. D. Glazebrook, H. M. Mitchell, and P. S. Ansell. Index policies for the maintenance of a collection of machines by a set of repairmen. *Eur. J. Oper. Res.*, 165:267–284, 2005.
- [25] K. D. Glazebrook, J. Niño-Mora, and P. S. Ansell. Index policies for a class of discounted restless bandit problem. *Adv. Appl. Prob.*, 34(4):754–774, 2002.
- [26] K.D Glazebrook. Optimal strategies for families of alternative bandit processes. *IEEE Trans, Automat. Control.*, 28:858–861, 1983.

- [27] D. Gross and C. M. Harris. *Fundamentals of Queueing Theory*. Wiley, New York, NY., 1998.
- [28] O. Gross. A class of discrete type minimization problems. Technical Report RM-1644, RAND Corp., 1956.
- [29] Gartner group. *Available from*. World Wide Web, <http://www.gartner.com>, 2004.
- [30] N. A. J Hastings. Bounds on the gain of a Markov decision process. *Oper. Res.*, 19:240–244, 1971.
- [31] W. J. Hopp and M. L. R. Sturgis. Quoting manufacturing due dates subject to a service level constraint. *IIE Trans.*, 32:771–784, 2000.
- [32] A. Hordijk and G. Koole. On the optimality of the generalised shortest queue policy. *Prob. Eng. Inf. Sci.*, 4:477–487, 1990.
- [33] D. J. Houck. Comparison of policies for routing customers to parallel queueing systems. *Oper. Res.*, 35:306–310, 1987.
- [34] R. A. Howard. *Dynamic Programming and Markov Process*. Wiley, New York, NY., 1960.
- [35] T. Ibaraki and N. Katoh. *Resource Allocation Problems: Algorithmic Approaches*. MIT Press, Cambridge, MA., 1988.
- [36] T. Ishikida and P. Varaiya. Multi-armed bandit problem revisited. *J. Opti. Theo. Appl.*, 83:113–154, 1994.
- [37] P. K. Johri. Optimality of the shortest line discipline with state-dependent service rates. *Eur. J. Oper. Res.*, 41:157–161, 1989.
- [38] H. Kaspi and A. Mandelbaum. Multi-armed bandits in discrete and continuous time. *Ann. Appl. Prob.*, 8:1270–1290, 1998.
- [39] M. N. Katehakis and A. F. Veinott. The multi-armed bandit problem: decomposition and computation. Technical Report, Dept. of Oper. Res. Stanford University, CA, 1985.
- [40] K. R. Krishnan. Joining the right queue: a state-dependent decision rule. *Auto. Con. IEEE Transactions*, 35:104–108, 1990.
- [41] K. R. Krishnan and T. J. Ott. State-dependent routing for telephone traffic: theory and results. *Proc. 25<sup>th</sup> IEEE Conf. Decision and Control*, pages 2124–2128, 1986.

- [42] K. R. Krishnan and T. J. Ott. Joining the right queue: A Markov decision rule. *Proc. 26<sup>th</sup> IEEE Conf. Decision and Control*, pages 1863–1868, 1987.
- [43] P. Kumar, M. Kalwani, and M. Dada. The impact of waiting time guarantees on customers' waiting experiences. *Marketing Science*, 16:295–314, 1997.
- [44] T. L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Adv. Appl. Math*, 6:4–22, 1985.
- [45] J. MacQueen. A modified dynamic programming method for Markov decision problems. *J. Math. Anal. Appl.*, 14:38–43, 1969.
- [46] A. Mahajan and D. Teneketzis. Multi-armed bandit problems. in *Foundations and Applications of Sensor Management*, eds. A Hero, D Castanon, D Cochran and K Kastella, 2007 (to appear).
- [47] D. Maister. Psychology of waiting lines. *Harvard Business School Cases*, April:71–78, 1984.
- [48] P. McDougall. Division of labour. <http://www.informationweek.com>, 2005.
- [49] B. Menich and R. Serfozo. Optimality of routing and servicing in dependent parallel processing systems. *Queueing Systems*, 9:403–418, 1991.
- [50] D. N. P. Murthy and I. Djameludin. New product warranty: A literature review. *International Journal of Production Economics*, 79:231–260, 2002.
- [51] P. Naor. The regulation of queue size by levying tolls. *Econometrica*, 37:15–24, 1969.
- [52] J. Niño-Mora. Restless bandits, partial conservation laws and indexability. *Adv. Appl. Prob.*, 33:76–98, 2001.
- [53] J. Niño-Mora. Dynamic allocation indices for restless projects and queueing admission control: a polyhedral approach. *Math. Program. Ser. A*, 93:361–413, 2002.
- [54] J. M. Norman. *Heuristic Procedures in Dynamic Programming*. Manchester University Press, Manchester, 1972.
- [55] A. Odoni. On finding the maximal gain for Markov decision processes. *Oper. Res.*, 17:857–860, 1969.
- [56] M. Opp, I. Adan, V. G. Kulkarni, and J. M. Swaminathan. Outsourcing warranty repairs: static allocation. Technical Report, University of North Carolina, Chapel Hill, NC., 2003.

- [57] M. Opp, K. D. Glazebrook, and V. Kulkarni. Outsourcing warranty repairs - dynamic allocation. *Nav. Res. Log.*, 52:381–398, 2005.
- [58] C. H. Papadimitriou and J. N. Tsitsiklis. The complexity of optimal queueing network control. *Math. Oper. Res.*, 24(2):293–305, 1999.
- [59] S. Patton. Outsourcing vendor management. <http://outsourcingmonitor.eu>, 2005.
- [60] M. L. Puterman. *Markov Decision Process: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, New York, 1994.
- [61] S. M. Ross. *Introduction to Stochastic Dynamic Programming*. Academic Press, New York, 1983.
- [62] A. Sen. *On Economic Inequality*. Clarendon, Oxford, 1973.
- [63] L. Sennott. *Stochastic Dynamic Programming and the Control of Queueing Systems*. Wiley, New York, 1999.
- [64] C. Serant. Solectron to Provide Xbox Support (merill lynch). *EBN*, October 2001.
- [65] S. Stidham. On the optimality of single-server queueing systems. *Oper. Res.*, 18:708–732, 1970.
- [66] S. Stidham. Optimality control of admission to a queueing system. *IEEE Trans. Auto. Control*, 30:705–713, 1985.
- [67] S. Taylor. Waiting for service: the relationship between delays and evaluations of service. *J. Marketing*, 58:56–69, 1994.
- [68] H. C. Tijms. *Stochastic models : An algorithmic approach*. Chichester : Wiley, 1994.
- [69] J. Tsitsiklis. A short proof of the Gittins index theorem. *Ann. Appl. Prob.*, 4:194–199, 1994.
- [70] J. van Meighem. Dynamic scheduling with convex delay costs: the generalized  $c\mu$  rule. *Ann. Appl. Prob.*, 5:809–833, 1995.
- [71] P. Varaiya, J. Walrand, and C. Buyukkoc. Extensions of the multiarmed bandit problem: the discounted case. *IEEE Trans, Automat. Control.*, 30:426–439, 1985.
- [72] B. Violino. What can logistics do for you? Global Services. June, 2006.
- [73] J. Walrand. Queueing networks. in *Handbook in Operations Research and Management Science: Stochastic Models*, 2:519–604, 1991.

- [74] R. R. Weber. On the optimal assignment of customers to parallel queues. *J. Appl. Prob.*, 15:406–413, 1978.
- [75] R. R. Weber. On the Gittins index for multi-armed bandits. *Ann. Appl. Prob.*, 2:1024–1033, 1992.
- [76] R. R. Weber and G. Weiss. On an index policy for restless bandits. *J. Appl. Prob.*, 27:637–648, 1990.
- [77] D. J. White. Dynamic programming, Markov chains and the method of successive approximations. *J. Math. Anal. Appl.*, 6:373–376, 1963.
- [78] W. Whitt. Deciding which queue to join: some counterexamples. *Oper. Res.*, 34:55–62, 1986.
- [79] P. Whittle. Multi-armed bandits and the Gittens index. *J.R. Statist. Soc.B*, 42:143–149, 1980.
- [80] P. Whittle. Arm-acquiring bandits. *Annals of Probability*, 9:284–292, 1981.
- [81] P. Whittle. *Optimization Over Time: Dynamic Programming And Stochastic Control*. John Wiley & Sons, New York, 1982.
- [82] P. Whittle. Restless bandits: activity allocation in a changing world. *J. Appl. Prob. special vol*, 25A:287–298, 1988.
- [83] P. Whittle. *Optimal Control: Basics and Beyond*. Wiley, New York, 1996.
- [84] J. Wijngaard. Decomposition for dynamic programming in production and inventory control. *Engineering and Process Economy*, 4:385–388, 1979.
- [85] W. Winston. Optimality of the shortest line discipline. *J. Appl. Prob.*, 14:181–189, 1977.
- [86] K. Wong. Top 10 challenges of outsourcing. <http://management.cadalyst.com>, April 2006.
- [87] S. Ziya, N. T. Argon, L. Ding, and K. D. Glazebrook. Dynamic routing of customers with general delay costs in a multi-server queueing system. *Prob. Eng. Inf. Sci.*, forthcoming.