INVENTORY CONTROL IN A BUILD-TO-ORDER ENVIRONMENT

A Thesis
Presented to
The Academic Faculty

by

Melda Ormeci

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

School of Industrial and Systems Engineering
Georgia Institute of Technology
August 2006

INVENTORY CONTROL IN A BUILD-TO-ORDER ENVIRONMENT

Approved by:

Professor John Vande Vate, Advisor
School of Industrial and Systems Engineering
Georgia Institute of Technology

Professor Jim Dai, Co-Advisor
School of Industrial and Systems Engineering
Georgia Institute of Technology

Professor Pinar Keskinocak
School of Industrial and Systems Engineering
Georgia Institute of Technology

Professor Leon McGinnis
School of Industrial and Systems Engineering
Georgia Institute of Technology

Professor Mark E. Ferguson
College of Management
Georgia Institute of Technology

Date Approved: 16 June 2006

To My Family...

ACKNOWLEDGEMENTS

First of all I would like to thank my advisor Professor John Vande Vate for all his guidance and support during my graduate studies. His knowledge, experience and insights have been very influential in my studies. He has been a great mentor and tremendous source of motivation. I would also like to express my special thanks to Professor Jim Dai, my co-advisor, without his guidance the accomplishment of this thesis would not be possible.

I would like to thank Professor Mark E. Ferguson, Professor Pinar Keskinocak, and Professor Leon McGinnis. I appreciate their time and effort in serving on my Ph.D. committee.

I would also like to thank all my friends. Your friendship has been a big support during my studies. I would especially like to mention Demet Batur, Tiravat Assavapokee, Yun Fong Lim, Yetkin Ileri, Marcos Goycoolea, Secil Savasaneril and Feryal Kuran. I am grateful to Betsy R. Thompson for her friendship and guidance, as well as Sylvia Wrobel and Robert P. Nicholls. They have made my life in Atlanta feel more like home.

Finally, I would like to thank my parents, Lale and Cankut Ormeci, and my sister Ceyda Ormeci for their love, support, guidance, and encouragement. I would like to thank my husband Erdem Matoglu for his love, friendship and support. Thanks for being there for me!

TABLE OF CONTENTS

LIST OF FIGURES

SUMMARY

This dissertation addresses inventory control in a build-to-order environment. It consists of three independent sections: challenges of managing supply in the auto industry from a build-to-order (BTO) perspective and some solution approaches, a stochastic control problem dealing with managing supply, and exploring Lagrangian methods in stochastic models.

In the first part, focusing on the auto industry we look at the challenges and solution strategies of employing BTO with global supply. We consider some familiar tools for managing domestic supply and exploit them for managing international supply, and propose some alternative new methods. We study frequency of supply, a tool widely used in local supply, as a way to improve performance in global supply operations. We study the impact of forecast accuracy, and conclude that improvements there alone may not be sufficient to obtain desired savings. In fact, our analysis strongly suggests that reducing the level of demand detail communicated to distant suppliers can simultaneously improve their quality of service and reduce their cost in providing it. Within this perspective we look at a new shipping policy, "Ship-to-Average", which prescribes sending a fixed quantity, based on the long term average forecast, with each shipment and making adjustments only if the inventory strays outside a prescribed range.

In the second part we look at a stochastic control problem, which also provides theoretical evidence in support of the Ship-to-Average policies. When a manufacturer places repeated orders with a supplier to meet changing production requirements, he faces the challenge of finding the right balance between holding costs and the operational costs involved in adjusting the shipment sizes. We consider an inventory whose content fluctuates as a Brownian motion in the absence of control. At any moment a controller can adjust the inventory level by any positive or negative quantity, but incurs both a fixed cost and a cost proportional to the magnitude of the adjustment. The inventory level must be nonnegative

at all times and continuously incurs a linear holding cost. The objective is to minimize long-run average cost. We show that control band policies are optimal for the average cost Brownian control problem and explicitly calculate the parameters of the optimal control band policy. This form of policy is described by three parameters $\{q, Q, S\}$, $0 < q \leq Q < S$. When the inventory falls to 0 (rises to S), the controller expedites (curtails) shipments to return it to q (Q). Employing apparently new techniques based on methods of Lagrangian relaxation, we show that this type of policy is optimal even with constraints on the size of adjustments and on the maximum inventory level. We also extend these results to the discounted cost problem.

The Brownian Control problem can be viewed as an idealization – without delivery delays, of the problem of supplying build-to-order operations. The conclusion that Control Band Policies are optimal in this idealized setting provides some theoretical explanation for the observed performance of Ship-to-Average policies. In fact, Ship-to-Average policies are a practical implementation of Control Band policies in the setting with delivery delays.

In the final part of the thesis we further study the Lagrangian approach developed in the second part. We explore the power and applicability in more details, demonstrate specific technical characteristics and describe additional applications. In this process, we look at several Brownian control problems with different forms of constraints. We solve these constrained Brownian control problems using Lagrange methods and note that the modelling aspect plays a critical role in the solution.

CHAPTER I

INTRODUCTION

This thesis consists of three independent sections. The first part serves as the motivation and highlights the challenges the auto industry faces. We address these challenges from a Build-to-Order perspective and offer some solution approaches. The second part addresses a Stochastic Control Problem and its extensions, where the task of managing supply is modelled through Brownian Motion. The Stochastic Control Problem is motivated by the discussions in the first part of the thesis. The third and final part of the thesis, further explores the methods developed to address extensions of the Stochastic Control Problem.

Build-to-Order Meets Global Sourcing

In Chapter 2, focusing on the auto industry in general we look at the challenges and solution strategies of employing BTO in the presence of off-shore suppliers. We identify the general trends and problems in the auto industry, look at some familiar tools for managing domestic supply and exploit them for managing international supply, and propose some alternative new methods. In particular, we propose frequency of supply, which manufacturers commonly rely on for local supply, as a way to manage the variability in the international supply with long and variable lead times. We study the impact of forecast accuracy, and conclude that improvements there alone may not be sufficient to obtain desired savings. Our studies show that as things stand in the global auto industry even the Herculean feat of doubling forecast accuracy would reduce inventory and expediting costs by less than 10 %. In fact, our analysis strongly suggests that reducing the level of demand detail communicated to distant suppliers can simultaneously improve their quality of service and reduce their cost in providing it. Within this perspective we look at a new shipping policy, "Ship-to-Average", which prescribes sending a fixed quantity, based on the long term average forecast, with each shipment and making adjustments only if the inventory strays outside a prescribed range.

Optimal Control of Brownian Inventory

In Chapter 3 we focus on a problem related to the one identified in Chapter 2 and provide theoretical evidence in support of the Ship-to-Average policies proposed in Chapter 2. We consider a manufacturer that places repeated orders for a part with a supplier to meet changing production requirements. Because of the costs involved in idling manufacturing capacity, backordering or stockouts of the part are not acceptable. Thus, if inventory of this part falls to precariously low levels, the manufacturer may expedite shipments or take other actions to increase it. On the other hand, space and capital constraints limit the inventory the manufacturer is willing to hold. When inventory grows too large, the manufacturer may take actions to reduce it. The manufacturer's challenge is to minimize the space and capital costs associated with holding inventory and the operational costs involved in adjusting supply.

This problem is common in the automobile industry, where the costs of idling production at an assembly plant can exceed \$1,000 per minute. Although an assembly plant typically produces vehicles at a remarkably constant rate, the composition of those vehicles can vary widely either in terms of the options they require or – as manufacturers move to more flexible lines, in terms of the mix of models produced. It is not unusual in the industry to see usage of a part vary by over 70% from one day to the next. The increasing number of models and options have increased the variability in usage while growing reliance on suppliers in lower cost countries like Mexico and China have compounded the complexity of supply. Thus, we focus on the balance between the capital and space costs of carrying inventory and the unplanned costs like expediting and curtailing shipments incurred in controlling inventory levels.

We model this problem as a Brownian control problem and seek a policy that minimizes the long-run average cost. We model the netput process or the difference between supply and demand for the part in the absence of any control, as a Brownian Motion with drift μ and variance σ^2. Inventory incurs linear holding costs continuously and must remain non-negative at all times. The manufacturer may, at any time, adjust the inventory level by, for example, expediting or curtailing shipments, but incurs both a fixed cost for making the adjustment and a variable cost that is proportional to the magnitude of the adjustment. The fixed cost and the unit variable cost depend on whether the adjustment increases or decreases inventory as these involve different kinds of interventions.

We show that a simple form of policy, called Control Band Policies, is optimal for the Brownian Control Problem. Control Band Policies can be interpreted as shipping the same quantity each time and adjusting this quantity only when inventory strays outside prescribed limits. To reflect the realities of the problem we introduce several constraints to the Brownian Control Problem and prove that control band policies are still optimal in the presence of these constraints. We solve the constrained problems by developing apparently new techniques based on Lagrangian relaxation.

A Lagrangian Approach to Stochastic Control

A major contribution of the third chapter is the method developed based on Lagrangian relaxation to solve constrained stochastic problems. Lagrangian relaxation methods have been widely used in deterministic optimization problems, both to solve constrained problems optimally and to obtain lower bounds on the optimal solution. Chapter 3 shows that Lagrangian relaxation techniques can be adapted to solve stochastic control problem. In the final part of the thesis, in Chapter 4, the Lagrangian approach, its power and applicability are explored in more details, specific technical characteristics are demonstrated and additional applications are described. In this process, we look at several Brownian control problems with different forms of constraints. The Brownian control problems are based on the ones solved in Chapter 3, but usually involve multiple parts and more general constraints. We solve these constrained Brownian control problems using Lagrange methods and note that the modelling aspect plays a critical role in the solution.

CHAPTER II

BUILD-TO-ORDER MEETS GLOBAL SOURCING: PLANNING CHALLENGE FOR THE AUTO INDUSTRY

2.1 Introduction

Auto manufacturers today face many challenges: The industry is plagued with excess capacity that drives down prices, international competitors are seizing share at both ends of the market and today's more savvy consumers are better informed about options and prices. All heighten competitive pressures, squeeze margins, and leave manufacturers struggling to increase revenues and market share.

One nearly universal strategy in the ensuing battle for market share and survival has been to increase product offerings both in terms of models and options. Long gone are the days when a black Model T, or a black Hongqi for that matter, was the only choice. Today's increased product variety, however, complicates operations and confounds demand forecasting. Correcting the inevitable forecast errors with discounts and rebates not only erodes manufacturers' margins, but also damages brand image. Consumers now expect to get less than they wanted and to pay less for it.

As the competitive pressures increase, companies look overseas for new markets and low cost suppliers: elbowing their way into high-growth developing markets with new international assembly operations and sourcing more and more components from distant low-cost suppliers.

Offshore suppliers and international assembly operations bring long and variable lead times that complicate demand forecasting, production planning and supply. Pressured to keep inventories lean, manufacturers often find needed supplies are still at sea and end up expediting parts to keep production lines running. It's no wonder that Womack and Jones [39] said, "Oceans and lean production are not compatible". Incompatible or not, both lean production and transoceanic supply are here to stay and manufacturers are left with the daunting task of managing the ”inherent incompatibility”. This is today’s planning challenge for the auto industry.

We address this challenge from the perspective of a global build-to-order auto manufacturer. Build-to-Order (BTO) is an attractive strategy for dealing with increased product variety and reducing finished goods inventories. BTO is “the practice of building customized or standard products as they are ordered and shipping them directly to customers, instead of building-to-forecast and shipping from inventory”. Today, many auto manufacturers, such as BMW, Toyota Scion, Renault, Mercedes-Benz etc. employ BTO, but to different extents.

As BTO meets global sourcing we observe that some familiar tools for managing domestic supply can improve international supply, but have yet to be effectively exploited. In other cases, traditional approaches come up short or are fatally flawed. For example, we find that increasing the frequency of shipments, a fundamental strategy of lean production long employed in improving local supply, also reduces inventory and risk for international supply. On the other hand, improving the accuracy and detail of demand signals, while still important, loses much of its impact in the face of long and variable lead times. Our studies show that as things stand in the global auto industry even the Herculean feat of doubling forecast accuracy would reduce inventory and expediting costs by less than 10 %. In fact, our analysis strongly suggests that reducing the level of demand detail communicated to distant suppliers can simultaneously improve their quality of service and reduce their cost in providing it.

We achieve this counterintuitive result through a new shipping policy, called “Ship-to-Average”, which ships the same quantity each time and adjusts this quantity only when inventory drifts out of prescribed ranges or the average rate of demand changes. Ship-to-Average is much easier to implement than currently accepted “Ship-to-Forecast” policies that slavishly follow detailed demand signals and, in the process, unnecessarily amplify the bullwhip effect creating wild swings in capacity requirements on both the suppliers and the transportation providers. In Section 2.2 we elaborate on the challenges auto manufacturers face today. We focus on the resulting product proliferation as manufacturers target smaller and smaller segments of the market in an attempt to maintain and grow market share. We discuss the shortcomings of traditional push systems when faced with a wide variety of product offerings and observe how BTO helps manage this variety. But BTO simply shifts the challenge from forecasting finished goods demand to forecasting individual component demands. We look at how variability in part usage and the trend toward global outsourcing in the auto industry impact forecast accuracy and the value of improving it. Finally in Section 2.3 we propose first step strategies for managing the inherent incompatibilities in automotive supply chains.

2.2 Refining the challenge definition

Auto manufacturing is a capital-intensive industry. Developing a new vehicle can cost $1 billion and a new assembly plant to produce the vehicle typically costs another $1-3 billion. To be price competitive, manufacturers must spread these capital investments over large volumes, but vehicle life cycles are shortening and the number of different models on the market is growing.

Manufacturers often look for new sales volumes in overseas markets. North American manufacturers like Ford and GM, for example, have long had operations in Europe and are rapidly ramping up operations in Asia, especially China and India. European manufacturers are also highly involved in foreign operations. Manufacturers from Japan and Korea have made deep in-roads in the US and are taking a second run at Europe and Chinese Manufacturers are poised to follow.

Globalization has been both a blessing and a curse. While it has opened up new markets, it has also brought new competitors. The big three in North America (GM, Ford, and the Chrysler unit of DaimlerChrysler) have lost over 20% of their share in the US market primarily to Japanese and Korean competitors in the past two decades. In 1965 the Big Three accounted for 95 % of all vehicles sold in North America. Today, that figure has fallen to only 58.5% and will continue to decline. In fact, some industry analysts forecast the Big Three's market share will fall to 50% by 2008. Japanese brands alone now account for 30.6 % and the more recent entrants from Korean already account for 4.1 %. Now, Chinese manufacturers are poised to enter the market with lower costs and prices as Chinese auto assembly workers typically work for as little as $2/hour including wages and benefits compared with $22/hour in Korea and nearly $60/hour in the U.S. [3] For example, the Chinese manufacturer, Chery, is aiming at the premium end of the market but with prices 30 per cent below those of its rivals [9].

EU manufacturers face a similar threat. While initial efforts by the Japanese manufacturers in Europe were not as successful as in the U.S., today, Japanese brands are the largest external players in the European markets. In the passenger car market, Japanese and Korean brands' market share climbed from 11% to more than 17% between 1990 and 2005. American manufacturers also compete in the European markets, but mostly through their European branches and hence brands.

As auto manufacturers increase their global presence, open new plants in new or emerging markets they contribute to significant overcapacity already present in the industry. Globally it is estimated that in 2005, the industry had enough idle capacity to produce an additional 18 million cars - equivalent to nearly 33 times the annual production of the largest assembly plant in North America (Smyrna, TN Nissan plant makes around 550,000 vehicles per year). There is clearly a mismatch between capacity and demand.

Product proliferation

In the race for market share, nearly every manufacturer has pursued smaller and smaller segments of the market with more and more models and options. Even a mass-market auto manufacturer like Ford Motor Company offers a dizzying array of products. Consumers can choose from among 23 models of Ford vehicles and for any given model, there are several million possible configurations to choose from. For example, among the 5 different Ford Escape models (XLS manual, XLS automatic, XLT automatic, XLT sport, Limited automatic) consumers can choose:

- o either front-wheel drive or four-wheel drive;
- o a 2.3L or 3.0L engine;

- o 4-speed or 5-speed transmission;
- o from nine exterior color options, three interior colors, four wheel options, two choices of tires, four options of electronics and four options of seats;
- o various combinations of five special option packages representing 32 different possibilities; and
- o various combinations of four different upgrades representing a further 16 options.

All told these options lead to something like 70 million different configurations of the Ford Escape.

A more extreme example, BMW, offers its customers an essentially infinite number of products. By its own estimates, BMW offers its customers 10^{32} different vehicle configurations to choose from. Just the 7-Series with over 350 model variants, 175 interior trims, 500 options and 90 standard colors, represents 10^{17} possible configurations. That's nearly 17 million different configurations for each man, woman and child on the planet. To put these astronomical figures into context, consider this: the Spartanburg, SC plant produced a quarter million Z3's (the predecessor of Z4) before it produced two that were identical.

From "Push" to Build-to-Order

The traditional "push" systems in which manufacturers build to forecast and meet customer demand from available finished goods inventory, are struggling to keep up with the new challenges accompanying product proliferation. Can Ford Motor Company, for example, really expect to accurately forecast demands for each of its 70 million variants of the Escape? The company only sold about 183,000 Escapes in 2004. In such a crowded market, even forecasting total annual sales for the model is a challenge.

When forecasts are wrong, manufacturers are forced to offer significant incentives to sell remaining inventories. In the US alone, automakers are estimated to have spent $60 billion in rebates in 2004[1], with over 90 percent of all cars sold having some form of incentive. The Big Three spent over $4,500 per vehicle on incentives that year and even popular Japanese

[1] The US Bureau of Economic Analysis estimates "Final sales of motor vehicles to domestic purchasers" at $518 billion in 2004.

brands that long shunned the practice succumbed to the inevitable pressure: "Toyota's incentives in all forms were up 31.6 % in 2004, reaching a level of over $3,100 per vehicle. Nissan's incentives were up 26.0 % to almost $2,000 a vehicle and Honda increased incentives by 79.5 % to almost $2,000 a vehicle"[24]. Despite these big incentives, however, no one is really satisfied with the end result. Manufacturers give up margin and consumers have to settle for what's available, not necessarily what they really wanted. It's no wonder so many manufacturers are moving towards build-to-order (BTO) systems that convert orders to products without holding any finished goods inventory.

For build-to-order strategies to be successful it is essential to have short order-to-delivery lead times. Typically lead times for a custom-built vehicle range from six weeks to ten weeks while customers expect their vehicles within two to three weeks. As a result a minority of vehicles in the market gets custom built, and manufacturers end up holding up to 100 days of sales in the market place [26]. Especially, in the U.S. customers rely mainly (more than 90% of sales [26]) on the build-to-stock model and expect vehicles (close to what they really wanted) in two or three weeks (i.e., the time to transfer it to their dealer). On the other hand, in Europe customers rely more heavily on BTO and accept lead times of several months, but expect exactly what they ordered.

To get consumers and dealers to accept lead times of 2-3 months some manufacturers allow them to change their choices within this period. For example, BMW has implemented a Customer-Oriented Sales and Production Process (KOVP) that allows customers to change their order up to 6 days before the vehicle is produced. In fact a customer may change major specifications like the engine, transmission, color or optional equipment within days before the vehicle is assembled, without affecting the agreed upon delivery date. And customers exercise this flexibility: BMW responds to more than 120,000 change requests every month. This flexibility also allows BMW dealers to meet individual customer requirements more quickly. Typically dealers place an order for a basic vehicle in advance, and make changes to that order as customer demand takes shape. In many cases the dealer can offer a customer exactly the vehicle he or she wants within a two or three-week window.

A side benefit of this flexibility is that, more often than not, customers tend to upgrade to more expensive options like navigation systems, xenon lights, and electronically adjustable comfort seats etc., as the delivery date approaches. Thus, by allowing its customers flexibility, BMW is not only able to get those orders earlier - generally months in advance, it also enjoys enhanced revenues from the resulting upgrades.

Capacity and Variability

In the auto industry production capacity represents a major capital expense and production labor is skilled, highly organized and expensive. Consequently, companies rely on a variety of strategies to smooth demand. In fact, an assembly plant's daily production is set to a takt time, e.g., a vehicle every 50 seconds. Changes in production are accomplished by speeding or slowing the takt time, adding or reducing shifts or shutting down the facility for a period of time during the year. These latter two adjustments are very crude indeed and are planned far in advance. Reductions to the takt time translate immediately into increases in the labor cost per vehicle produced and so are again only made reluctantly and as part of a broader plan to manage capacity.

Thus over significant time periods the production rate in terms of vehicles produced per day, per shift and even per hour is very constant. Significant deviations from this rate are rare and usually the result of quality or supply problems.

Although an assembly plant typically produces vehicles at a remarkably constant rate, the composition of those vehicles can vary widely either in terms of the options they require or - as manufacturers move to more flexible lines, in terms of the mix of models produced. For example Hyundai's plan in Montgomery, AL builds Sonata sedans and Santa Fe SUVs on the same line. That plant is designed to accommodate as many as 4 different models simultaneously. Similarly Honda's plant in East Liberty, OH produces cars and light trucks on the same assembly line, while Ford's flexible plant in Chicago is capable of building eight models off two platforms and the Dearborn plant is capable of assembling nine vehicles off three platforms. This flexibility helps spread capital costs and risk.

As a consequence, even though automobile assembly plants make the same number of vehicles every day, demand for the components that go into those vehicles is increasingly

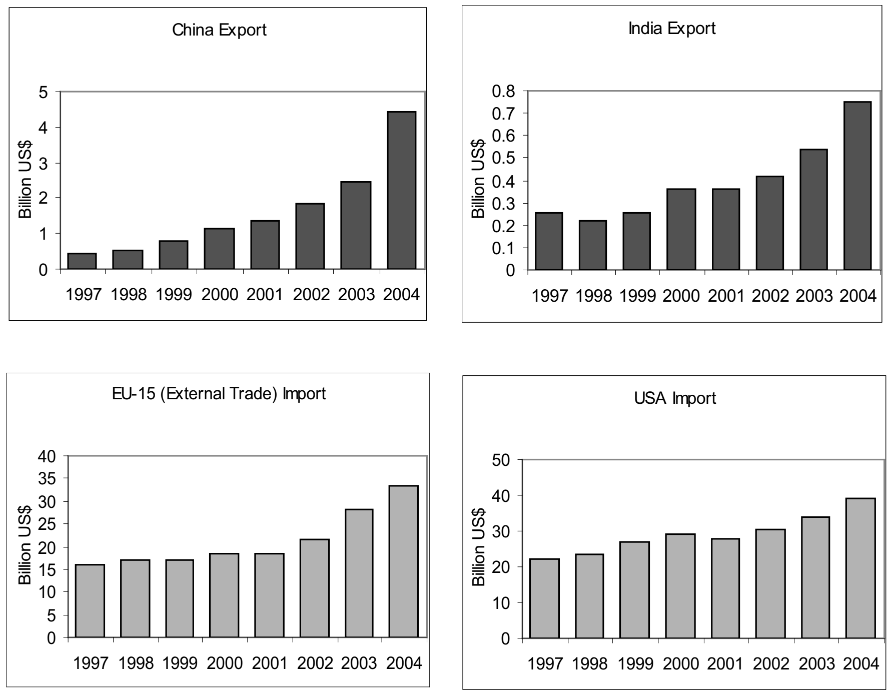

Figure 1: Import and Export of auto components. Source: OECD ITCS - International Trade By Commodity Statistics, Rev.3, India Department of Commerce

unpredictable. In fact, it is not unusual in the industry to see usage of a part vary by over 70% from one day to the next. BTO, which was intended to free manufacturers from the tyranny of poor forecasts, simply shifts the problem from forecasting finished vehicle demands to forecasting demands for components.

Global Sourcing

One of the major trends in the auto industry is globalization. Besides from producing global products at global production sites auto manufacturers increasingly rely on low-cost overseas suppliers. For example, US manufacturers import from Mexico, Brazil and now China and India. Sourcing from offshore suppliers is not limited the original equipment manufacturers. This trend is visible even among first tier suppliers like Cummins International (engine parts), Delphi, Visteon etc.[16] As a result, in the past decade, China's export of auto components have increased more than eight folds and US and European imports of auto components have doubled (See Figure 1).

On the other hand, many companies supply international operations from domestic markets for quality reasons, economic reasons and sometimes out of social obligations. EU

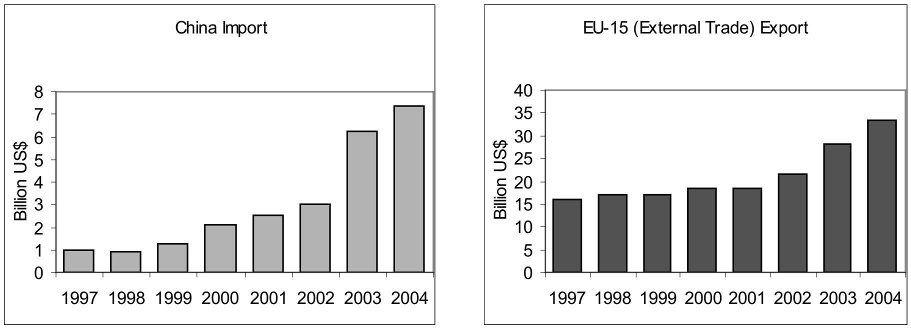

Figure 2: Import and Export of auto components increase due to global sourcing. Source: OECD ITCS - International Trade By Commodity Statistics, Rev.3

manufacturers, for example BMW, Volvo, Daimler Chrysler and PSA, supply operations in US, Brazil and increasingly China with parts from European suppliers. The decision to use domestic markets for supplying international operations is based on several factors. Having already invested in the tooling in Europe, manufacturers figure out that, in some cases, the volumes in this hemisphere are not sufficient to replicate those costs. Furthermore the quest for quality sometimes prevents the use of untried, new suppliers in the international locations. Likewise, existing relationships with suppliers and social responsibility issues all play a part in the decision to use domestic markets in supplying international operations. Hence, the imports of auto components (not vehicles) into countries like China and Brazil are on the increase along with the export volumes of EU countries where these parts are being sourced from (See Figure 2).

2.3 Rethinking familiar tools and new models

The wide variety of product offerings commitments to rapid order fulfillment and near-zero inventories create demand for individual parts that vary wildly from day to day. Combined with global sourcing- hence, long and variable lead times- supplying individual parts becomes a very complex process often forcing manufacturers to expedite parts to keep production lines running. Managing this "inherent incompatibility" between lean production and transoceanic supply is a daunting task for manufacturers.

In this section we observe that a traditional approach - increasing the frequency of shipments - that is a fundamental strategy of lean production long employed in improving local supply, also reduces inventory and risk for international supply. In other cases, traditional approaches come up short or are fatally flawed. For example, improving the accuracy and detail of demand signals, while still important, loses much of its impact in the face of long and variable lead times. Finally we look at a new shipping strategy, Ship-to-Average, which relies less heavily on the inevitably erroneous forecasts that accompany international supply and instead focuses on longer-term trends in demand. In fact, our analysis strongly suggests that reducing the level of demand detail communicated to distant suppliers can simultaneously improve their quality of service and reduce their cost in providing it.

Increasing Frequency

Auto manufacturers have long recognized the value of frequent deliveries from local suppliers: more frequent shipments mean smaller shipments and smaller shipments mean less inventory. More frequent shipments also mean less time between shipments and so less risk of interrupted supply. If something goes wrong with one delivery, another is not far behind.

Frequent deliveries are a cornerstone of lean production that Toyota has exploited to the fullest. The carmaker brings parts into its plants more than once an hour. Can these same ideas improve international supply? Hourly shipments may not be feasible, but increasing the frequency of shipments and reducing the risk is still a viable approach. Today most manufacturers make weekly shipments for internationally sourced parts. This means that on average they must carry not only a large safety stock to protect against delays in delivery or sudden increases in demand, but also half a week's supply as cycle stock. More frequent shipments reduce cycle inventory both at the plant and at the supplier and have a secondary benefit of reducing safety inventories. Shipping frequency's impact on cycle inventory is straightforward; its impact on safety stock or safety lead time and expediting are not.

The results of our analysis, suggest that increasing the frequency of international shipments can offer significant savings both in inventory and in expediting costs. Our preliminary simulations, modelled on the characteristics of a European manufacturer supplying assembly operations in the US, show that doubling the frequency of shipments from once-per-week to twice-per-week simultaneously reduce in-plant inventories and expediting costs. Moving to three shipments per week reduces in-plant inventories and expediting costs even further.

The ocean container lines' sailing schedules force manufacturers that source parts from overseas to work under a periodic review system since, regardless of when orders are placed, they can only be delivered when vessels are scheduled to arrive. Clark and Scarf [6] established the optimality in this setting of (s, S) policies, which each time either place an order to bring the inventory level up to level S or, if the order quantity is too small and inventory levels are already greater than s, place no order at all. As a consequence, although the time between orders is relatively constant, the order quantities can vary widely.

The risk of stocking out in an order cycle depends on the variability of demand in a period determined by both the time between orders and the lead time [35]. Increasing the frequency of orders reduces the time between orders and so reduces the risk of stocking out in an order cycle (assuming we hold safety stocks constant). On the other hand, increasing the frequency of orders increases the number of orders in a year and so means we face this reduced risk of stocking out more often. Increased frequency's final impact on safety stock depends on the balance between these two factors. In our experience the benefits of increased frequency outweigh the costs well beyond the frequencies achievable with ocean shipping schedules.

Increasing frequency only reduces inventory and expediting costs if the shipments are mixed. Simply increasing the number of vessels used has little or no impact beyond spreading the risks if individual components are still shipped once per week. In other words, shipping containers on three different vessels each week is little different from shipping them all on a single vessel if each part number is only on one vessel. To realize the savings, each component has to ride on each vessel and achieving this may require more packaging flexibility and increase handling costs. In our experience, the inventory and expediting savings more than compensate for any extra handling.

Increasing the frequency of shipments from local suppliers typically increases transportation costs because it reduces capacity utilization or requires a larger number of smaller vehicles to make the deliveries. This is not generally the case for international shipments or at least the impact on transportation costs is less pronounced. This is because, with the exception of a few high-volume suppliers that ship direct, most internationally sourced components are already consolidated for international packaging. Thus, while increasing the frequency of these shipments will generally increase transportation costs between the supplier and the consolidation center, it has little effect on the international transportation costs. Ocean shipping costs are generally incurred per container shipped so shipping 100 containers on three different vessels costs essentially the same as sending 300 containers on a single vessel each week.

Unfortunately, sailing schedules make it difficult for companies to increase international shipment frequencies. The carriers all try to set sail at the end of the week so their vessels aren't idled at port over the weekend. The impact: 90% of the fastest 30% of services from Hamburg to Charleston and 80% of the fastest 30% of services between Hong Kong and Long Beach are scheduled to arrive between Friday and Sunday. To achieve higher frequencies shippers are forced to use services out of and into alternative ports. While this certainly complicates the logistics, it has the added benefit of reducing the risks of disruptions in the case of a port strike or a hurricane.

Improving Forecasts

As manufacturers struggle to manage the "inherent incompatibility" between lean inventories and long lead times their first reaction is typically directed at improving the accuracy of demand forecast through investments in IT technology. They develop advanced forecasting models in order to capture the nature of demand and invest in new sales and operations planning software to get a better handle on their supply chain. While in general there is a consensus about the positive impacts of improved forecast accuracy, there really is no clear single method to achieve it or to identify the extent to which it can be achieved. In any case, everyone agrees it is impossible to eliminate forecast errors all together.

While forecast accuracy is important and efforts to improve it should not be abandoned, this avenue offers little prospect for resolving a significant portion of the "inherent incompatibility". Demand variability for option parts is an inherent component of the very flexibility manufacturers must allow consumers (and dealers) in order to get orders far enough in advance. As a result a certain amount of forecast inaccuracy is inherent in the Build-to-Order.

There are four ways to reduce the demand variability engendering forecast errors:

1. **Reduce the number of options:** Low volume parts are the most difficult to forecast and managing them takes the same effort if not more than high volume parts. Trimming the least popular options improves forecast accuracy for the more popular ones and reduces the complexity of managing supply.

2. **Increase the "frozen horizon":** Freezing orders earlier fixes the production schedule further in advance allowing the manufacturers to work with shorter forecast horizon i.e. closer to the actual demand.

3. **Source from local suppliers:** Reducing lead times significantly improves forecast accuracy.

4. **Exploit Postponement Strategies:** By delaying the point of product differentiation as close as possible to actual order information and waiting until actual order signals are received to complete the products, postponement offers manufacturers the flexibility they need to efficiently produce customized products.

Toyota effectively exploits postponement to provide U.S. customers with almost unlimited customization of the Scion even though that model is built in Japan. Scion customers in the U.S. can chose from over 40 different options, leading to 2^{40} different versions of the vehicle - more than 3,000 different versions for each person in the U.S. But the vehicles produced in Japan are standardized; distinguished only by transmission (automatic or manual) and color (there are 6 choices). These standardized vehicles are customized to order in the U.S and delivered within 5 to 7 days, if not sooner.

Toyota's postponement strategy for the Scion would be difficult to implement for higher end vehicles with fewer standard features and more complicated option offerings that can only be added during production. But that doesn't mean postponement isn't a viable strategy for complex subassemblies like wiring harnesses and cockpits.

Unfortunately, none of these options is particularly attractive or available to manufacturers. Furthermore, with the exception of the Scion, these "solutions" are contrary to the trends in the industry.

Our studies strongly suggest that even if we could improve forecast accuracy significantly, it would have limited impact. In fact, our studies indicate that the Herculean accomplishment of halving forecast errors would only reduce inventory and expediting costs by less than 10 %.

Using a history of orders and rolling forecasts similar to those of a European auto manufacturer, we artificially improved the forecast accuracy and evaluated the impact on inventory and expediting.

Our simulations suggest that even significant improvements in forecast accuracy yield only relatively small improvements in inventory and expediting costs. In fact, our studies indicated that cutting the forecast errors in half – from 75% to about 37% – reduced inventory and expediting costs by less than 10%. The reason: Poor forecasts are the scapegoat for all the excess inventories, stock outs, and premium freight charges. But they simply don't deserve the blame, at least not by themselves. Inaccurate forecasting is just one factor contributing to the problem. The other culprit is lead time variability. Accurate demand forecasts can tell you how much you'll need, but if lead times are unreliable you are still left with the question of when to ship so it arrives when you need it.

And lead time variability has been on the rise driven by increasingly violent and unpredictable weather (The 2005 Atlantic hurricane season, with 28 nameable storms is the most active season on record, surpassing the 1933 season's 21 [28]) and the general strain on the global transportation system.

North America has a growing port capacity problem, and the resulting congestion is affecting global supply chains negatively. Over the last twenty years, container volumes in North American ports have grown at an average annual rate of 7 %, but port capacity has not kept pace with volume growth [23]. A study by the National Chamber Foundation of the U.S. Chamber of Commerce, pointed out that most major North American ports are already operating at or near full capacity and will have significant capacity deficits by 2010 [27]. Port congestion leads to unpredictable delays, causing manufacturers to increase inventory levels and adjust supply networks to minimize the risk of stock outs and shutdowns. A 2004 survey by Logistics Management revealed that respondents were experiencing average

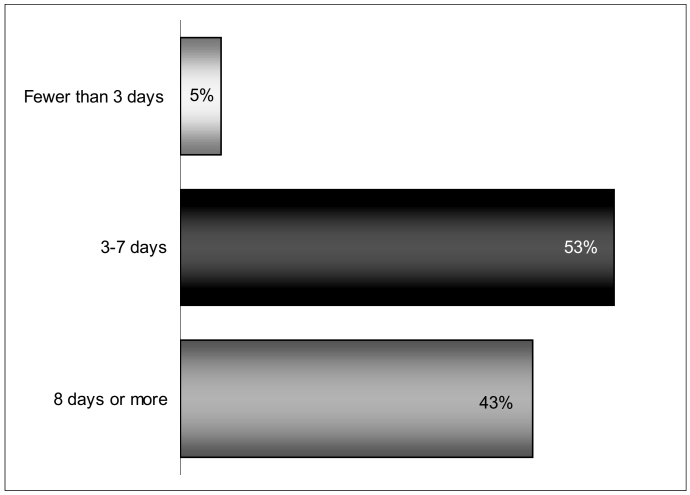

Figure 3: "How much delay are you experiencing in delivering your products due to West Coast backup?" [19]

delivery delay of 6.5 days. 43 % of the respondents reported 8 or more days of delay while 53% reported delays between 3 and 7 days (See Figure 3) [19]. Even if manufacturers' ability to predict what they need improves, their ability to predict when it will get there is declining.

The principle impact of improved forecasts is to reduce the need for expedited shipments. The auto industry with its heavier, lower value parts, typically reserves airfreight for emergencies. So increasing forecast accuracy principally reduces the already small airfreight costs and so has a relatively small impact over all. A parallel study with a telecommunications equipment manufacturer, on the other hand, indicated that improved forecast accuracy can have significant impact on companies that rely heavily on airfreight and expedited shipments to meet customer orders.

Ship to Average

Traditional methods for managing supply rely on detailed forecasts. Each order quantity is based on forecasted demand over the period the order is intended to cover: a week's forecasted production for weekly shipments, a few days' forecasted production with more frequent shipments. These methods suffer from two obvious difficulties: First, as we have seen, actual production quantities over periods as short as a week or a few days vary widely; second, forecasts of these quantities are remarkably inaccurate. The result: Traditional methods magnify the bullwhip effect and force international suppliers and logistics service providers on a wild goose chase after phantom peaks in demand.

We propose a simpler, and it turns out, more effective strategy for managing supply, which we call Ship-to-Average. The idea is to keep order quantities constant, adjusting them only when inventory drifts out of prescribed ranges or the average rate of demand changes. Ship-to-Average offers several advantages: First, since changes in the order quantity are the exception rather than the rule, Ship-to-Average policies reduce the effort involved in managing supplies. Second, since order quantities are consistent and reliable, suppliers and service providers can more efficiently plan production and manage labor requirements. Finally, the manufacturer can count on more consistent and reliable shipment quantities and no longer needs complicated calculations to determine whether it is necessary to expedite parts.

Ship-to-Average policies do use forecasts to calculate fixed order quantities, but ignore the details in these forecasts and instead focus on longer term trends. Figure 4 shows the relationship between forecast errors and the period of demand forecasted in the case of a European manufacturer forecasting demands for an option-driven component at a North American plant 30 days in advance. Figure 5 compares the stability of the forecasts averaged over different time periods. As expected, forecast accuracy and stability improves significantly with the length of the period forecasted. While daily forecasts are inaccurate and swing wildly, weekly, monthly and quarterly forecasts are increasingly accurate and stable. Ship-to-Average policies ship to the more stable and accurate average of forecasts covering longer periods.

Although there is no theoretical proof that Ship-to-Average is a best strategy, Ormeci et al.[30] proved that, in the case of zero lead times, policies of this form are optimal. In addition, initial studies based on the characteristics of a European auto manufacturer suggest that Ship-to-Average policies are at least as effective as the current Ship-to-Forecast strategies in terms of average inventory and expediting costs. In fact, in many cases Ship-to-Average simultaneously reduces total avoidable cost, expediting costs and the variability in order quantities. In many cases the reductions in expediting costs and order variability

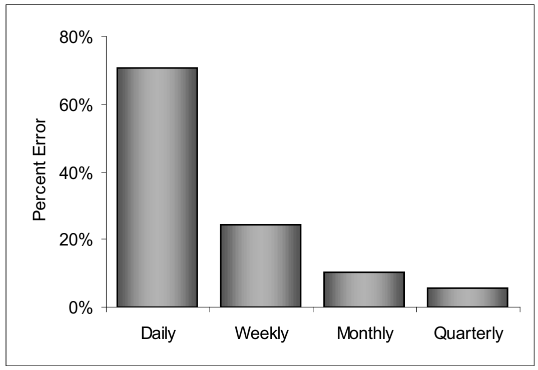

Figure 4: Forecast errors covering different periods

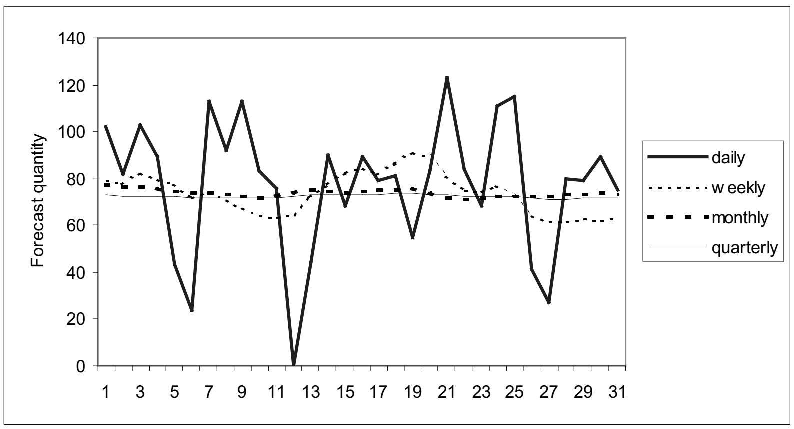

Figure 5: Forecasts over different time periods

are significant (60% and higher reductions for expediting; 50% and higher reductions for order variability) with no increase in total avoidable costs. So far, we have always been able to identify a Ship-to-Average policy that significantly improves order stability without increasing total avoidable cost when compared to the best Ship-to-Forecast policy.

Ship-to-Average policies produce significantly more stable order patterns, which simplifies the suppliers' task of managing labor and capacity. Consider the example of a large, high-value option-driven part shipped from Europe to a North American plant. Adopting the standard that a change of more than 10% in successive order quantities creates planning and scheduling challenges for the supplier, we found that while the best Ship-to-Forecast policy exceeded this limit more than 60% of the time, our Ship-to-Average policy exceed it less than 15% of the time with the same total avoidable cost. That additional constancy and predictability helps suppliers better manage their resources and realize savings that should eventually be reflected in piece prices.

2.4 Conclusions

One of the most important challenges facing build-to-order auto manufacturers is the wide variety of product offerings, commitments to rapid order fulfillment and near-zero inventories creating demand for individual parts that vary wildly from day to day, while supplying these parts globally with long and variable lead times. Increasing frequency is a simple and efficient tool to improve performance in terms of inventory costs and the risk of stocking out. However, given current carrier schedules increasing frequency is far from a simple task. Besides, since traditional Ship-to-Forecast policies base each order quantity on forecasted demand over the period the order is intended to cover, more frequent shipments rely on more detailed forecasts which are both less accurate and more variable. As a consequence, more frequent shipments can magnify the bullwhip effect unnecessarily. In the auto industry, attempts to redress this by improving forecast accuracy face daunt-ing challenges and even if they succeed beyond all expectations, will have little impact on the problem. We introduce a new type of policy as an alternative way to deal with variable lead times and inaccurate forecasts. Our Ship-to-Average policy, not only performs at least as well as the more complicated Ship-to-Forecast policies, but also significantly smoothes the order quantities sup-pliers must manage. Theoretical studies [30] indicate this form of policy is optimal in the case of zero lead times, simulation studies based on the parameters of actual component demands at the North American operations of a European manufacturer provide further support and suggest Ship-to-Average policies are at least as good as traditional Ship-to-Forecast policies in terms of inventory and expediting costs. These results are still new and much work is left to do. We are currently working to identify conditions that especially favor Ship-to-Average policies and develop analytic tools for quickly finding optimal policy parameters and assessing the long term impacts of implementing them.

CHAPTER III

IMPULSE CONTROL OF BROWNIAN MOTION: THE CONSTRAINED AVERAGE COST CASE

3.1 Introduction and Motivation

Consider a manufacturer that places repeated orders for a part with a supplier to meet changing production requirements. Because of the costs involved in idling manufacturing capacity, backordering or stockouts of the part are not acceptable. Thus, if inventory of this part falls to precariously low levels, the manufacturer may expedite shipments or take other actions to increase it. On the other hand, space and capital constraints limit the inventory the manufacturer is willing to hold. When inventory grows too large, the manufacturer may take actions to reduce it. The manufacturer's challenge is to minimize the space and capital costs associated with holding inventory and the operational costs involved in adjusting supply.

This problem is common in the automobile industry, where the costs of idling production at an assembly plant can exceed \$1,000 per minute. Although an assembly plant typically produces vehicles at a remarkably constant rate, the composition of those vehicles can vary widely either in terms of the options they require or – as manufacturers move to more flexible lines, in terms of the mix of models produced. It is not unusual in the industry to see usage of a part vary by over 70% from one day to the next. Electronically transmitted releases against a standing purchase order have essentially eliminated ordering costs and careful packaging, loading and transportation planning have squeezed planned transportation costs to the last penny. But the increasing number of models and options have increased the variability in usage while growing reliance on suppliers in lower cost countries like Mexico and China have compounded the complexity of supply. Thus, we focus on the balance between the capital and space costs of carrying inventory and the unplanned costs like expediting and curtailing shipments incurred in controlling inventory levels.

We model this problem as a Brownian control problem and seek a policy that minimizes the long-run average cost. We model the netput process or the difference between supply and demand for the part in the absence of any control, as a Brownian Motion with drift μ and variance σ^2. Inventory incurs linear holding costs continuously and must remain non-negative at all times. The manufacturer may, at any time, adjust the inventory level by, for example, expediting or curtailing shipments, but incurs both a fixed cost for making the adjustment and a variable cost that is proportional to the magnitude of the adjustment. The fixed cost and the unit variable cost depend on whether the adjustment increases or decreases inventory as these involve different kinds of interventions.

We address the average cost problem rather than the more traditional discounted cost problem for several reasons: First, although the notion of a discount factor may be natural and intuitive in many applications including finance, it is generally alien to material planners and the challenges of motivating it and eliciting a value for it outweigh the potential benefits. Second, researchers have traditionally pursued the discounted problem because that version of the dynamic programming operator exhibits favorable contraction properties that facilitate analysis. While the discounted cost version of our Brownian control problem has been studied in the literature [12], the average cost version has not.

A common method for solving long-run average cost problems is to utilize the discounted cost problem and take the limit as the discount rate goes to zero; see, for example, Feinberg and Kella [10], Hordijk and Van Der Duyn Schouten [15], Robin[33], and Sulem[36]. In this paper we address the average cost problem directly without the traditional reliance on a limit of the discounted cost version. This approach is more direct, more elegant and opens the possibility of tackling average cost problems with more general holding cost structures as it does not require explicit solutions to the discounted cost problem. We show that control band policies are optimal for the average cost Brownian control problem. This form of policy is described by three parameters $\{q, Q, S\}$ with $0 < q \leq Q < S$. When the inventory falls to 0, the manufacturer expedites a shipment to return it to q. When the inventory rises to S, the maximum allowed, the manufacturer curtails shipments, reducing the balance to

Q. The simplicity of this policy greatly facilitates its application in industrial settings, like automobile assembly, with thousands of parts to manage.

We extend the Brownian control problem by introducing constraints that reflect more of the realities of the inventory management problem. First, we impose an upper bound on the inventory level to reflect physical limits on total available inventory space or financial limits on the inventory budget. In addition, noting that the magnitude of adjustments may be limited, for example, by the nominal shipment quantity or the capacity of the transportation mode, we introduce bounds on the magnitude of each control. We prove that control band policies are optimal for the average cost Brownian control problem even with these constraints on the maximum inventory level and on the magnitudes of adjustments to the inventory. In each case, we provide optimality conditions that allow explicit calculation of the optimal control parameters. In fact, employing apparently new applications of Lagrangian relaxation techniques [11], we show how to reduce the constrained problem to a version of the original unconstrained problem and, in the process, provide methods for computing the optimal control band policy in the presence of the constraints. This approach extends to constrained versions of the discounted cost problem and we state the analogous results in that setting.

In their paper, Harrison et al.[12] addressed a problem in finance called the Stochastic Cash Management Problem in which a certain amount of income or revenue is automatically channelled into a cash fund from which operating disbursements are paid. If the balance in the cash fund grows too large, the controller may invest the excess. If it becomes too small, he may sell off investments to replenish it. The challenge is to minimize the *discounted* opportunity costs associated with holding cash in the fund and the transaction costs involved in buying and selling investments. They showed that a control band policy is optimal for the Stochastic Cash Management Problem and provided methods for computing the optimal control band policy. While they studied exactly our problem but with discounted costs rather than average costs, Taksar [38] studied the average cost problem but with a different cost structure on the controls: He minimized the average holding and control costs when there are *no* fixed costs for control and singular control is employed. He showed that the optimal policy, characterized by two constants $a < b$, keeps the process inside $[a, b]$ with minimal effort. Constantinides [7] studied a similar cash management problem that allowed both positive and negative cash balance. Although he looked at the average cost problem, he assumed the optimal policy to be of a simple form and proceeded to find the optimal parameters of this policy. Richard [32] on the other hand looked at a diffusion process with fixed plus proportional adjustments costs and general holding costs and showed that the optimal policy is one of impulse control in both finite and infinite horizon discounted case, without addressing the existence of such a control.

Our model differs from the ones in these works in many ways. While Harrison et al. [12] and Richard [32] allow adjustments of any magnitude, we introduce constraints on the magnitude of adjustments to the inventory level. While Constantinides [7] did not have any constraints on the inventory level, even allowing negative values, Harrison et al. [12] required that inventory remain non-negative at all times, and Taksar [38] allowed the holding cost to be infinite outside a range, which amounts to constraints on the inventory level. We, on the other hand, introduce a constraint on the maximum inventory level while keeping inventory non-negative. We develop a method built on ideas from Lagrangian relaxation techniques to handle these additional constraints. The method is quite general and we extend it to analogously constrained versions of the discounted cost problem. Finally, Harrison et al. [12] considered the Brownian control problem in a financial setting where discounting costs over time is natural and appropriate. We consider the problem in an industrial setting where the long-run average cost is more natural and accepted.

The rest of the paper is organized as follows: In Section 3.2, we describe the average cost Brownian control problem and its policy space. The main results of this paper, the optimality of control band policies and optimality conditions that permit ready computation of the optimal policy parameters, are stated here. Sections 3.3 and 3.4 set up the preliminaries for the solution of the problem. In Section 3.3 we introduce a lower bound for the optimal cost and in Section 3.4 we define a relative value function for control band policies with average cost criteria and show that the average cost can be calculated through this function. In Section 3.5 we first consider the Bounded Inventory Average Cost Brownian Control Problem in which M, the maximum inventory level allowed, is finite. We prove that a control band policy is optimal for the Bounded Inventory Average Cost Brownian Control Problem and derive explicit equations used in calculating the optimal control parameters. As a special case, we characterize an optimal solution for the unconstrained average cost problem with no bounds on the maximum inventory level. We also demonstrate the optimal policy for the discounted cost setting with finite M. In Section 3.6 we introduce constraints on the magnitude of the adjustments. In particular, employing Lagrangian relaxation techniques, we reduce the constrained problem to a version of the original unconstrained problem. Once again we solve the constrained problem in the average cost setting and characterize the optimal policies for the discounted cost setting as well. We conclude the paper in Section 3.7 by looking at possible extensions and by stating the optimal policy for the Bounded Inventory, Constrained Average Cost Brownian Control Problem, which simultaneously imposes bounds on the controls and a finite upper limit on inventory.

3.2 *Impulse control of Brownian motion*

In this paper we use the following notation and assumptions. Let Ω be the space of all continuous functions $\omega : [0, \infty) \to \mathbb{R}$, the real line. For $t \geq 0$ let $X_t : \Omega \to \mathbb{R}$ be the coordinate projection map $X_t(\omega) = \omega(t)$. Then $X = (X_t,\, t \geq 0)$ is the canonical process on Ω. Let $\mathcal{F} = \sigma(X_t,\, t \geq 0)$ denote the smallest σ-field such that X_t is $\mathcal{F}$-measurable for each $t \geq 0$, and similarly let $\mathcal{F}_t = \sigma(X_s,\, 0 \leq s \leq t)$ for $t \geq 0$. When we mention adapted processes and stopping times hereafter, the underlying filtration is understood to be $\{\mathcal{F}_t,\, t \geq 0\}$. Finally, for each $x \in \mathbb{R}$ let $\mathbb{P}_x$ be the unique probability measure on $(\Omega, \mathcal{F})$ such that X is a Brownian motion with drift μ, variance σ^2 and starting state x under $\mathbb{P}_x$. Let $\mathbb{E}_x$ be the associated expectation operator.

We are to control a Brownian motion $X = \{X_t,\, t \geq 0\}$ with mean μ, variance σ^2 and starting state x. Upward or downward adjustments, ξ_n, are exerted at discrete times, T_n, so that the resulting inventory process represented by $Z = \{Z_t,\, t \geq 0\}$ remains within $[0, M]$, where M is a possibly infinite bound on inventory. We adopt the convention that the sample path of Z is right continuous on $[0, \infty)$ having left limits in $(0, \infty)$. (The time parameter of a process may be written either as a subscript or as an argument, depending on which is more convenient.)

A policy $\varphi = \{(T_n, \xi_n),\ n \geq 0\}$ consists of stopping times $\{T_0, T_1, \ldots\}$ at which control is exerted and random variables $\{\xi_0, \xi_1, \ldots\}$ representing the magnitude and direction of each control. So T_n is the time at which we make the $(n+1)$st adjustment to inventory and ξ_n describes the magnitude and direction of that adjustment. We only consider policies that are non-anticipating, i.e., each adjustment ξ_n must be $\mathcal{F}_{T_n^-}$ measurable, where, for a stopping time τ, $\mathcal{F}_{\tau^-}$ is defined as in Definition I.1.11 of Jacod and Shiryaev [17].

When a policy increases inventory by $\xi > 0$, it incurs cost $K + k\xi$ representing the fixed costs $K > 0$ of changing the inventory and the variable costs $k\xi \geq 0$ that grow in proportion to the size of the adjustment. When a policy reduces inventory, i.e., when it adjusts inventory by $\xi < 0$, it incurs cost $L - \ell\xi$ where $L > 0$ is the fixed cost for reducing inventory and $-\ell\xi \geq 0$ is the variable cost. Finally, we assume that inventory incurs a positive holding cost of $h > 0$ per unit per unit of time.

We consider the Average Cost Brownian Control Problem, which is to find a non-anticipating policy $\varphi = \{(T_n, \xi_n),\ n \geq 0\}$ that minimizes:

$$\mathrm{AC}(x, \varphi) = \limsup_{n \to \infty} \mathbb{E}_x \left[\frac{1}{T_n} \left(\int_0^{T_n} hZ_t\, dt + \sum_{i=1}^{n} \left((K + k\xi_i) 1_{\{\xi_i > 0\}} + (L + \ell|\xi_i|) 1_{\{\xi_i < 0\}} \right) \right) \right], \tag{1}$$

the expected long-run average cost starting at a given initial point $x \in \mathbb{R}_+ = [0, \infty)$. Setting

$$\phi(\xi) = \begin{cases} K + k\xi & \text{if } \xi > 0, \\ 0 & \text{if } \xi = 0, \\ L - \ell\xi & \text{if } \xi < 0, \end{cases} \tag{2}$$

the control problem (1) can be written compactly as

$$\mathrm{AC}(x, \varphi) = \limsup_{T \to \infty} \frac{1}{T} \mathbb{E}_x \left[\int_0^T hZ_t\, dt + \sum_{i=1}^{N(T)} \phi(\xi_i) \right], \tag{3}$$

where, for each time $t \geq 0$, $N(t) = \sup\{n \geq 0 : T_n \leq t\}$ denotes the number of jumps by time t.

We introduce a possibly infinite upper bound M on the inventory level and restrict our attention to the the policy space $\mathcal{P}$, which is the set of all non-anticipating policies satisfying

$$\mathbb{P}_x(0 \leq Z_t \leq M \text{ for all } t > 0) = 1 \text{ for all } x \in \mathbb{R}_+. \tag{4}$$

When the upper bound M on inventory is finite, we refer to the problem as the Bounded Inventory Average Cost Brownian Control Problem. When M is infinite, adding the constraints

$$-d \leq \xi_i \leq u \qquad \text{for each } i = 1, 2, \ldots$$

to the Average Cost Brownian Control Problem gives rise to the Constrained Average Cost Brownian Control Problem.

Harrison, Sellke and Taylor [12] proved that a simple form of policy, called a control band policy, is optimal for the discounted cost problem. A control band policy is defined by three parameters $\{q, Q, S\}$ with $0 < q \leq Q < S$ (they define control band policies with strict inequalities between the parameters, $0 < q < Q < S$, however, we allow $0 < q \leq Q < S$). When inventory falls to 0, the policy exerts a control to bring it up to level q. When the inventory rises to S, the maximum allowed, the policy exerts a control to reduce it by $s = S - Q$ and bring it to level Q. If the initial inventory level lies outside the range $[0, S]$, the policy exerts a one-time control ξ_0 to bring it to the closer of level q or level Q. Thus, the control band policy $\{q, Q, S\}$ is defined by $\{(T_n, \xi_n),\, n \geq 0\}$, where $T_0 = 0$,

$$\xi_0 = \begin{cases} q - X_0, & \text{if } X_0 \leq 0, \\ 0, & \text{if } 0 < X_0 < S, \\ Q - X_0, & \text{if } X_0 \geq S \end{cases}$$

thereafter, $\{T_n,\, n > 0\}$ are the hitting times for $\{0, S\}$, i.e., $T_n = \{t > T_{n-1} : Z_t = 0 \text{ or } Z_t = S\}$ for all $n = 1, 2, \ldots$ and

$$\xi_n = \begin{cases} q, & \text{if } Z(T_n^-) = 0, \\ -s, & \text{if } Z(T_n^-) = S. \end{cases}$$

When φ is a control band policy with parameters $\{q, Q, S\}$ the average cost does not depend on the initial state x and hence we also use $\mathrm{AC}(\varphi)$ or $\mathrm{AC}(q, Q, S)$ to denote its average cost.

Now we state the main results, Theorems 3.2.1 and 3.2.2, of this paper. Theorem 3.2.1 says that a control band policy is optimal for the Bounded Inventory Average Cost Brownian Control Problem. Theorem 3.2.2 says that a control band policy is also optimal for the Constrained Average Cost Brownian Control Problem. Both Theorems 3.2.1 and 3.2.2 provide explicit formulas for calculating the optimal control parameters.

Theorem 3.2.1. *The Bounded Inventory Average Cost Brownian Control Problem admits an optimal policy that is a control band policy. Furthermore, If $\mu \neq 0$, the parameters $\{q^*, Q^*, S^*\}$ of the optimal control band policy φ^* are defined by the unique non-negative values of λ, s, Δ and $Q \geq \Delta$ satisfying:*

$$\begin{aligned}
L &= -h\left(\frac{s^2(1+e^{\beta s})}{2\mu(1-e^{\beta s})} + \frac{s}{\beta\mu}\right) + \lambda\left(\frac{s}{1-e^{-\beta s}} - \frac{1}{\beta}\right), && (5)\\
k+\ell &= -\frac{h\Delta}{\mu} - \frac{hs(1-e^{\beta\Delta})}{\mu(1-e^{-\beta s})} + \lambda\left(\frac{e^{\beta\Delta}-1}{1-e^{-\beta s}}\right), && (6)\\
K &= \frac{h(Q-\Delta)se^{\beta s}}{\mu(1-e^{\beta s})} + \frac{h(\Delta^2-Q^2)}{2\mu} + \frac{hs(e^{\beta Q}-e^{\beta\Delta})}{\mu\beta(1-e^{-\beta s})} - (\ell+k)(Q-\Delta)\\
&\quad +\lambda\left(\frac{e^{\beta Q}-e^{\beta\Delta}}{\beta(1-e^{-\beta s})} - \frac{(Q-\Delta)}{1-e^{-\beta s}}\right), && (7)
\end{aligned}$$

such that

$$\begin{aligned}
\lambda(S-M) &= 0, && (8)\\
S &\leq M, && (9)
\end{aligned}$$

where $\beta = \frac{2\mu}{\sigma^2}$, $q \equiv Q - \Delta$ and $S \equiv Q + s$.

If $\mu = 0$, the parameters of the optimal control band policy φ^* are defined by the unique non-negative values of λ, s, Δ and Q satisfying (8), (9), *and*

$$\begin{aligned}
L &= \frac{hs^3}{6\sigma^2} + \frac{\lambda s}{2}, && (10)\\
\ell+k &= \frac{h(\Delta^2+\Delta s)}{\sigma^2} + \lambda\frac{\Delta}{s}, && (11)\\
K &= \frac{Q^3 h}{3\sigma^2} + \frac{Q^2 hs}{2\sigma^2} - \frac{\Delta^2 hs}{2\sigma^2} - \frac{\Delta^3 h}{3\sigma^2} + (\Delta - Q)(\ell+k) + \lambda\left(\frac{Q^2-\Delta^2}{2s}\right), && (12)
\end{aligned}$$

where $q \equiv Q - \Delta$ and $S \equiv Q + s$.

For each fixed $\lambda \geq 0$, (5) alone determines $s \equiv S - Q$, after which (6) determines $\Delta \equiv Q - q$, and then (7) determines q. The value of λ that also satisfies (8) and (9) gives the optimal control policy.

Theorem 3.2.2. *The Constrained Average Cost Brownian Control Problem admits an optimal policy that is a control band policy. Furthermore, if $\mu \neq 0$, the parameters $\{q^*(d,u), Q^*(d,u), S^*(d,u)\}$ of the optimal control band policy φ^* are defined by the unique non-negative values of λ, η, s, Δ and $Q \geq \Delta$ satisfying:*

$$\begin{aligned} L - \lambda d &= -h\left(\frac{s^2(1+e^{\beta s})}{2\mu(1-e^{\beta s})} + \frac{s}{\beta\mu}\right), && (13)\\ k + \eta + \ell + \lambda &= -\frac{h\Delta}{\mu} - \frac{hs(1-e^{\beta\Delta})}{\mu(1-e^{-\beta s})}, && (14)\\ K - \eta u &= \frac{h(Q-\Delta)se^{\beta s}}{\mu(1-e^{\beta s})} + \frac{h(\Delta^2 - Q^2)}{2\mu} + \frac{hs(e^{\beta Q} - e^{\beta\Delta})}{\mu\beta(1-e^{-\beta s})} \\ &\quad -(\ell + \lambda + k + \eta)(Q - \Delta) && (15) \end{aligned}$$

such that

$$\begin{aligned} \lambda(d - s) &= 0, && (16)\\ \eta(u - Q + \Delta) &= 0, && (17)\\ s &\leq d, && (18)\\ Q - \Delta &\leq u, && (19) \end{aligned}$$

where $\beta = \frac{2\mu}{\sigma^2}$, $q \equiv Q - \Delta$ and $S \equiv Q + s$. If $\mu = 0$, the parameters of the optimal control band policy φ^* are defined by the unique non-negative values of λ, η, s, Δ and $Q \geq \Delta$ satisfying (16)-(19) and:

$$\begin{aligned} L - \lambda d &= \frac{hs^3}{6\sigma^2}, && (20)\\ \ell + \lambda + k + \eta &= \frac{h(\Delta^2 + \Delta s)}{\sigma^2}, && (21)\\ K - \eta u &= \frac{Q^3 h}{3\sigma^2} + \frac{Q^2 hs}{2\sigma^2} - \frac{\Delta^2 hs}{2\sigma^2} - \frac{\Delta^3 h}{3\sigma^2} + (\Delta - Q)(\ell + \lambda + k + \eta), && (22) \end{aligned}$$

where $q \equiv Q - \Delta$ and $S \equiv Q + s$.

Note that when $d = \infty$, we take $\lambda = 0$ as the unique solution to (16) and when $u = \infty$, we take $\eta = 0$ as the unique solution to (17).

Control band policies are also optimal for the discounted, Brownian control problem, with constraints on the inventory space and magnitude of adjustments; see Theorem 3.5.1 in Section 3.5 and Theorem 3.6.1 in Section 3.6.

The rest of this paper is devoted to the proofs of Theorems 3.2.1 and 3.2.2. In Section 3.3, we establish a lower bound on the average cost over all feasible policies. In Section 3.4 we define a relative value function for each control band policy with average cost criteria and show that the average cost can be calculated through this function. In Section 3.5 we complete the proof of Theorem 3.2.1 by showing that the average cost of a control band policy with a particular choice of control parameters achieves the lower bound. We prove Theorem 3.2.2, in Section 3.6, by employing Lagrangian relaxation techniques. We reduce the constrained problem to a version of the original unconstrained problem.

3.3 Lower Bound

In this section, we show how to construct a lower bound on the average cost over all feasible policies. Then in Section 3.4 we define the relative value functions for control band policies and show how to compute their average costs. In Section 3.5 we construct a particular control band policy $\{q^*, Q^*, S^*\}$ and show that its value function provides a lower bound, and thus establishing the optimality of the control band policy.

Proposition 3.3.1. *Suppose that $f : [0, M] \to \mathbb{R}$ is continuously differentiable, has a bounded derivative, and has a continuous second derivative at all but a finite number of points. Then for each time $T > 0$, initial state $x \in \mathbb{R}_+$ and policy $\{(T_n, \xi_n),\, n \geq 0\} \in \mathcal{P}$*

$$\mathbb{E}_x[f(Z_T)] = \mathbb{E}_x[f(Z_0)] + \mathbb{E}_x\left[\int_0^T \Gamma f(Z_t)\,dt\right] + \mathbb{E}_x\left[\sum_{n=1}^{N(T)} \theta_n\right], \tag{23}$$

where

$$\theta_n = f(Z(T_n)) - f(Z(T_n-)), \quad \text{for } n = 1, 2, \ldots \text{ and } \Gamma f = \frac{1}{2}\sigma^2 f'' + \mu f'. \tag{24}$$

Proof. The proof follows from an application of Ito's formula and is similar to the proof of (2.16) in Harrison, Sellke and Taylor [12]. (Ito's formula for semimartingales can be found, for example, in Theorem I.4.57 of Jacod and Shiryaev [17].) □

The following proposition shows that each function satisfying certain conditions provides a lower bound on the optimal average cost.

Proposition 3.3.2. *Suppose that $f : [0, M] \to \mathbb{R}$ satisfies all the hypotheses of Proposition 3.3.1 plus*

$$\Gamma f(x) - hx - \Gamma f(0) \leq 0 \quad \text{for almost all } 0 \leq x \leq M, \tag{25}$$

$$f(x) - f(y) \leq K + k(x-y) \quad \text{for } 0 \leq y < x \leq M, \tag{26}$$

$$f(x) - f(y) \leq L - \ell(x-y) \quad \text{for } 0 \leq x < y \leq M. \tag{27}$$

Then $\mathrm{AC}(x, \varphi) \geq -\Gamma f(0)$ for each policy $\varphi \in \mathcal{P}$ and each initial state $x \in \mathbb{R}_+$.

Proof. Recall the definition of θ_n in (24) and $\phi(\xi_n)$ in (2). Note that for each n, $\theta_n \leq \phi(\xi_n)$ by conditions (26)-(27). It follows from (23) and (25) that

$$\mathbb{E}_x[f(Z_T)] \leq \mathbb{E}_x[f(Z_0)] + \mathbb{E}_x\left[\int_0^T hZ_t\,dt\right] + \Gamma f(0)T + \mathbb{E}_x\left[\sum_{n=1}^{N(T)} \phi(\xi_n)\right]. \tag{28}$$

Dividing both sides of (28) by T and letting $T \to \infty$ gives

$$-\Gamma f(0) + \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}_x[f(Z_T)] \quad \leq \quad \mathrm{AC}(x, \varphi). \tag{29}$$

If $\limsup_{T\to\infty} \frac{1}{T}\mathbb{E}_x[f(Z_T)] \geq 0$, (29) yields $-\Gamma f(0) \leq \mathrm{AC}(x, \varphi)$, proving the proposition. Now suppose that $\limsup_{T\to\infty} \frac{1}{T}\mathbb{E}_x[f(Z_T)] < 0$. We show that this implies

$$\mathrm{AC}(x, \varphi) = \infty, \tag{30}$$

which again yields $-\Gamma f(0) \leq \mathrm{AC}(x, \varphi)$, as desired. To prove (30), set $a = \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}_x[f(Z_T)]$. Since $a < 0$ by assumption, it follows that there exists a constant $t^* > 0$ such that

$$\frac{1}{T}\mathbb{E}_x[f(Z_T)] < a/2 \quad \text{for } T > t^*$$

and so $\mathbb{E}_x[f(Z_T)] < Ta/2$ for $T > t^*$. Since f has bounded derivatives, it is Lipschitz continuous. Thus, there exists a constant $c > 0$ such that

$$f(Z_0) - f(Z_T) \leq |f(Z_T) - f(Z_0)| \leq c|Z_T - Z_0| \leq c(Z_T + Z_0) \tag{31}$$

for $T \geq 0$. Taking expectations on both sides of (31), we see that

$$f(x) - \mathbb{E}_x[f(Z_T)] \leq c(\mathbb{E}_x[Z_T] + x)$$

for all $T \geq 0$. Therefore,

$$\mathbb{E}_x[Z_t] \geq \frac{1}{c}\left[f(x) + t|a|/2\right] - x = c_1 t + c_2 \text{ for all } t \geq t^*,$$

where $c_1 = |a|/(2c)$ and $c_2 = f(x)/c - x$. It follows that

$$\begin{aligned} AC(x, \varphi) &\geq \limsup_{T\to\infty} \mathbb{E}_x\left[\frac{1}{T}\int_0^T hZ_t\, dt\right] = h \limsup_{T\to\infty} \frac{1}{T}\int_0^T \mathbb{E}_x[Z_t]\, dt \\ &\geq h \limsup_{T\to\infty}\left[\frac{1}{T}\int_{t^*}^T c_1 t\, dt\right] + hc_2 \geq h \limsup_{T\to\infty}\left[\frac{c_1(T^2 - t^{*2})}{2T}\right] + hc_2 = \infty, \end{aligned}$$

proving (30). □

Remark. In the discounted cost problem, Harrison, Sellke and Taylor [12] obtained a bound similar to the one in Proposition 3.3.2. In proving their bound, they require their policies to satisfy conditions

$$\mathbb{P}_x(Z_t \geq 0 \text{ for all } t > 0) = 1 \text{ for all } x \in \mathbb{R}_+ \text{ and} \tag{32}$$

$$\mathbb{E}_x \sum_{i=0}^{\infty} e^{-\gamma T_i}(1 + |\xi_i|) < \infty \text{ for all } x \in \mathbb{R}_+, \tag{33}$$

where γ is the discount rate, and, as before, $\mathbb{R}_+ := [0, \infty)$. They used condition (33) to ensure that when f has bounded derivative $\mathbb{E}_x[e^{-\gamma T} f(Z_T)] \to 0$ as $T \to \infty$.

A natural analog of (33) for the average cost problem is

$$\limsup_{n\to\infty} \mathbb{E}_x \frac{1}{T_n}\sum_{i=0}^{n}(1 + |\xi_i|) \quad < \quad \infty \text{ for all } x \in \mathbb{R}_+. \tag{34}$$

One suspects that condition (34) should analogously lead to

$$\mathbb{E}_x[f(Z_T)/T] \to 0 \quad \text{as} \quad T \to \infty \tag{35}$$

as long as the corresponding average cost is finite. Surprisingly, one can construct counterexamples such that (35) does not hold even though condition (34) holds and the corresponding average cost is finite; see Appendix A.1 for a counterexample. We are able to obtain the lower bound in Proposition 3.3.2 without condition (34) on the policies.

3.4 Control Band Policies

In this section we show that, for a given control band policy $\varphi = \{q, Q, S\}$, the associated long-run average cost $\mathrm{AC}(x, \varphi)$ is independent of the initial state x. Furthermore, the average cost can be computed through a relative value function, which we define below.

Proposition 3.4.1 below shows that there is a constant g and a function V that satisfy the ordinary differential equation (ODE), known as the Poisson equation,

$$\Gamma V(x) - hx + g = 0, \quad 0 \leq x \leq S, \tag{36}$$

and the boundary conditions

$$V(0) = V(q) - K - kq, \tag{37}$$

$$V(S) = V(Q) - L - \ell s. \tag{38}$$

The constant g is unique and the function V is unique up to an additive constant. With a slight abuse of terminology, any such function V is called the *relative value function* associated with the control band policy φ. The significance of the relative value function and the constant g is that they provide the long-run average cost $\mathrm{AC}(x, \varphi)$ of the control band policy through the formula $\mathrm{AC}(x, \varphi) = g = -\Gamma V(0)$.

Proposition 3.4.1. *Let the parameters of the control band policy $\varphi = \{q, Q, S\}$ be fixed.*

(a) There is a function $V : [0, S] \to \mathbb{R}$ that is twice continuously differentiable on $[0, S]$ and satisfies (36)–(38).

(b) Such a function is unique up to a constant.

(c) The constant g is unique. The average cost of the control band policy $\{q, Q, S\}$ is independent of the starting point and is given by $g = -\Gamma V(0)$.

Proof. We address only the case $\mu \neq 0$. The case $\mu = 0$ is analogous. The general solution to the ODE (36) is

$$V(x) = Ax + Be^{-\frac{2\mu}{\sigma^2}x} + \frac{h}{2\mu}x^2 + E$$

for some constants A, B, and E. The boundary conditions (37) and (38) determine the values of A and B uniquely. Thus, we have proved both (a) and (b). Since g is a constant, it follows from (36) that $g = -\Gamma V(0)$, and thus g is unique.

To complete the proof of (c), consider the control band policy $\varphi = \{q, Q, S\} = \{(T_n, \xi_n), n \geq 0\}$. Since V is twice continuously differentiable and has bounded derivative on $[0, S]$, we have by Proposition 3.3.1 that

$$\mathbb{E}_x[V(Z_T)] = \mathbb{E}_x[V(Z_0)] + \mathbb{E}_x\left[\int_0^T \Gamma V(Z_t)\,dt\right] + \mathbb{E}_x\left[\sum_{n=1}^{N(T)} \theta_n\right],$$

where $\theta_n = V(Z(T_n)) - V(Z(T_n-))$ for $n = 1, 2, \ldots$. Since V satisfies (37)–(38), $\theta_n = V(Z(T_n)) - V(Z(T_n-)) = \phi(\xi_n)$ for $n = 1, 2 \ldots$. Therefore, since V and g satisfy (36),

$$\begin{aligned}
\mathbb{E}_x[V(Z_T)] - \mathbb{E}_x[V(Z_0)] + gT &= \mathbb{E}_x\left[\int_0^T \Gamma V(Z_t) + g\,dt\right] + \mathbb{E}_x\left[\sum_{n=1}^{N(T)} \theta_n\right] \\
&= \mathbb{E}_x\left[\int_0^T hZ_t\,dt\right] + \mathbb{E}_x\left[\sum_{n=1}^{N(T)} \phi(\xi_i)\right].
\end{aligned}$$

Finally, dividing both sides by T, taking the limit as $T \to \infty$ and observing that

$$\lim_{T\to\infty} \frac{1}{T}\mathbb{E}_x[V(Z_T)] = 0,$$

we have that $\mathrm{AC}(x, \varphi) = g$. $\square$

Remark. From the proof of Proposition 3.4.1, the average cost of the control band policy φ has the formula $\mathrm{AC}(x, \varphi) = g = -\Gamma V(0) = -(\sigma^2 h/(2\mu) + \mu A)$. A detailed computation in terms of A shows that the latter expression is equal to

$$\mathrm{AC}(x,\varphi) = \begin{cases} h\left(\frac{(S^2-Q^2)C}{2(sC-qD)} - \frac{q^2 D}{2(sC-qD)} - \frac{\sigma^2}{2\mu}\right) + \frac{((L+\ell s)C + (K+kq)D)\mu}{sC-qD}, & \mu \neq 0 \\ h\frac{3S^2-3Ss+s^2-q^2}{3(2S-s-q)} + \frac{(\frac{K}{q}+k)\sigma_2}{2S-s-q} + \frac{(\frac{L}{s}+\ell)\sigma_2}{2S-s-q}, & \mu = 0, \end{cases} \tag{39}$$

where

$$\begin{aligned}
C &= e^{-\beta q} - 1, \\
D &= e^{-\beta S} - e^{-\beta Q}.
\end{aligned}$$

Appendix A.2 provides an alternative derivation of (39) for the average cost $\mathrm{AC}(x, \varphi)$ of a control band policy φ. The derivation is based on a basic adjoint relationship, which may be of independent interest.

Note that the relative value function of any control band policy satisfies conditions (36)–(38), which are related to the conditions (25)–(27) used in Proposition 3.3.2 to construct a bound. In Section 3.5 we construct a control band policy, $\{q^*, Q^*, S^*\}$, whose relative value function can be extended to $[0, M]$. In order for the extended function to satisfy conditions (25)–(27) for all $x \in [0, M]$, the parameters $\{q^*, Q^*, S^*\}$ must be the unique values specified in Theorem 3.2.1.

3.5 Optimal Policy Parameters

One of the main results of this paper, stated in Theorem 3.2.1, is to prove that a control band policy is optimal for the Bounded Inventory Average Cost Brownian Control Problem and to provide optimality conditions that permit ready computation of the optimal policy parameters. In the remainder of this section we prove Theorem 3.2.1. The outline of the proof is as follows.

In Proposition 3.4.1, we have proved that for any control band policy $\varphi = \{q, Q, S\}$, its average cost is given by $-\Gamma V(0)$, where V, defined on $[0, S]$, is the relative value function. We are to find a particular choice of parameter set $\{q^*, Q^*, S^*\}$ such that the corresponding relative value function V can be extended on $[0, M]$ and the extended function satisfies the conditions of Proposition 3.3.2. Thus, $-\Gamma V(0)$, in addition to being the average cost of the control band policy $\{q^*, Q^*, S^*\}$, is also a lower bound on the average cost of the Bounded Brownian control problem. Therefore, the control band policy $\{q^*, Q^*, S^*\}$ is optimal.

Recall that for a given set of parameters $\{q, Q, S\}$, the corresponding relative value function satisfies (36)–(38). To search for the optimal parameter set $\{q^*, Q^*, S^*\}$, we impose the following conditions on $\{q, Q, S\}$ and V:

$$V'(q) = k, \tag{40}$$

$$V'(Q) = -\ell, \tag{41}$$

$$V'(S) = -\ell - \lambda, \tag{42}$$

$$\lambda(S - M) = 0 \text{ and} \tag{43}$$

$$S \leq M, \tag{44}$$

where $\lambda \geq 0$. Lemma 3.5.1 below shows that the parameter set $\{q^*, Q^*, S^*\}$ satisfying (40)-(44) exists. In the proof of Theorem 3.2.1, to be presented immediately following Lemma 3.5.2, the corresponding relative value function will be extended to $[0, M]$. Condition (40) is to ensure that the extended function satisfies inequality (26) of Proposition 3.3.2, and condition (41) is to ensure that the extended function satisfies inequality (27) of Proposition 3.3.2. Condition (42) is to ensure that the derivative of the extended function is also continuous at S^*.

Lemma 3.5.1. *(a) There exists a unique non-negative solution $s^*, \Delta^*, Q^*, \lambda^*$ satisfying* (5)–(9)*. (b) For the parameter set $\{s^*, \Delta^*, Q^*, \lambda^*\}$, the corresponding relative value function satisfies* (40)–(44)*.*

Proof. We demonstrate the proof for $\mu \neq 0$. When $\mu = 0$ the arguments are analogous.

Proof of Part (a) is given in Appendix A.3.

Now we prove Part (b). Set $S^* = Q^* + s^*$ and $q^* = Q^* - \Delta^*$. We now show that the relative value function V corresponding to the control band policy $\{q^*, Q^*, S^*\}$ satisfies (40)-(44).

First note that the function

$$\begin{aligned} f(x) &= -\frac{hs^*(1 - e^{-\frac{2\mu}{\sigma^2}x})}{\mu(1 - e^{-\frac{2\mu}{\sigma^2}s^*})} + \frac{hx}{\mu} - \ell - \lambda^* \frac{(1 - e^{-\frac{2\mu}{\sigma^2}x})}{1 - e^{-\frac{2\mu}{\sigma^2}s^*}} \\ &= -\left[\frac{hs^*}{\mu} + \lambda^*\right] \frac{(1 - e^{-\frac{2\mu}{\sigma^2}x})}{(1 - e^{-\frac{2\mu}{\sigma^2}s^*})} + \frac{hx}{\mu} - \ell \end{aligned} \tag{45}$$

is the unique solution to the ODE: $\Gamma f(x) - h = 0$ for $-Q^* \leq x \leq s^*$, and satisfies boundary conditions $f(0) = -\ell$ and $f(s^*) = -\ell - \lambda^*$. Let

$$\pi(x) = f(x - Q^*), \quad 0 \leq x \leq S^*.$$

It follows that π is the unique solution to the ODE

$$\Gamma\pi(x) - h = 0, \quad 0 \leq x \leq S^*, \tag{46}$$

satisfying boundary conditions

$$\pi(Q^*) = -\ell, \quad \text{and} \tag{47}$$

$$\pi(S^*) = -\ell - \lambda. \tag{48}$$

(When $\mu = 0$, $\pi(x) = \frac{hx^2}{\sigma^2} + Bx + C$ for some constants B and C).

For $x \in [0, S^*]$, let

$$V(x) = \int_0^x \pi(u)\, du. \tag{49}$$

We claim that $V(x)$, defined on $[0, S^*]$, is a relative value function of the control band policy $\{q^*, Q^*, S^*\}$. To see this, we first note that $(\Gamma V(x) - hx)' = 0$ on $[0, S^*]$, and thus $\Gamma V(x) - hx$ is a constant on $[0, S^*]$. Denoting the constant by $-g$, we have that $V(x)$, together with the constant g, satisfies the Poisson equation (36). We next prove that $V(x)$ satisfies boundary conditions (37) and (38). The boundary condition (38) can be written in terms of π as

$$\begin{aligned} L &= V(Q^*) - V(S^*) - \ell(S^* - Q^*) = -\int_{Q^*}^{S^*} (\pi(x) + \ell)\, dx = -\int_0^{s^*} (f(x) + \ell)\, dx \\ &= -h\left(\frac{s^{*2}(1 + e^{\beta s^*})}{2\mu(1 - e^{\beta s^*})} + \frac{s^*}{\beta\mu}\right) + \lambda^*\left(\frac{s^*}{1 - e^{-\beta s^*}} - \frac{1}{\beta}\right), \end{aligned} \tag{50}$$

which holds because of (5). Similarly boundary condition (37) can be written in terms of π as

$$\begin{aligned} K &= V(q^*) - V(0) - kq^* = \int_0^{q^*} (\pi(x) - k)\, dx = \int_{-Q^*}^{-\Delta^*} (f(x) - k)\, dx \\ &= \frac{h(Q^* - \Delta^*)s^* e^{\beta s^*}}{\mu(1 - e^{\beta s^*})} + \frac{h(\Delta^{*2} - Q^{*2})}{2\mu} + \frac{hs^*(e^{\beta Q^*} - e^{\beta \Delta^*})}{\mu\beta(1 - e^{-\beta s^*})} - (\ell + k)(Q^* - \Delta^*) \\ &\quad + \lambda^*\left(\frac{e^{\beta Q^*} - e^{\beta\Delta^*}}{\beta(1 - e^{-\beta s^*})} - \frac{(Q^* - \Delta^*)}{1 - e^{-\beta s^*}}\right), \end{aligned} \tag{51}$$

which holds because of (7). Therefore, $V(x)$ is the relative value function of the control band policy $\{q^*, Q^*, S^*\}$.

Clearly, $V(x)$ satisfies conditions (41) and (42). To complete the proof of the lemma, it remains to prove that $V(x)$ satisfies condition (40). To see this, condition (40) requires

$$\begin{aligned} k &= \pi(q^*) = f(q^* - Q^*) = f(-\Delta^*) \\ &= -\frac{hs^*(1 - e^{\frac{2\mu}{\sigma^2}\Delta^*})}{\mu(1 - e^{-\frac{2\mu}{\sigma^2}s^*})} - \frac{h\Delta^*}{\mu} - \ell - \lambda^*\frac{(1 - e^{\frac{2\mu}{\sigma^2}\Delta^*})}{1 - e^{-\frac{2\mu}{\sigma^2}s^*}}, \end{aligned} \tag{52}$$

which is equivalent to (6). Thus, function $V(x)$ is the relative value function satisfying (40)–(44). $\square$

The following properties of π are useful in the proof of Theorem 3.2.1.

Lemma 3.5.2. *Let $\pi : [0, S^*] \to \mathbb{R}$ be the unique solution to the ODE* (46) *satisfying* (47) *and* (48) *for the optimal parameters $\{s^*, \Delta^*, Q^*, \lambda^*\}$ satisfying* (5)–(9). *Extend $\pi(x)$ to $[S^*, M)$ via $\pi(x) = -\ell$ for $x \in [S^*, M)$. Then,*

(a) For $x \in [0, q^*]$, $\pi(x) \geq k$ and for $x \in [q^*, M]$, $\pi(x) < k$.

(b) For $x \in [0, Q^*]$, $\pi(x) \geq -\ell$ and for $x \in [Q^*, M]$, $\pi(x) \leq -\ell$.

Proof. Recall that $\pi(x) = f(x - Q^*)$ for all $x \in [0, S^*]$, where f is defined in (45). When $\mu \geq 0$, it is clear from (45) that π is strictly convex. The result follows from the convexity of π and the conditions (47)-(48) and (52).

When $\mu < 0$, we have two cases to consider. If $\frac{hs^*}{\mu} + \lambda^* < 0$, π is again strictly convex, hence same arguments apply. If $\frac{hs^*}{\mu} + \lambda^* > 0$, π is strictly concave and decreasing. To see this note that

$$f'(x) = -\left[\frac{hs^*}{\mu} + \lambda^*\right] \frac{2\mu}{\sigma^2} \frac{e^{-\frac{2\mu}{\sigma^2}x}}{1 - e^{-\frac{2\mu}{\sigma^2}s^*}} + \frac{h}{\mu} < 0.$$

In this case, the result follows from (47)-(48) and (52). $\square$

Proof of Theorem 3.2.1. Let $s^* > 0, \Delta^* \geq 0$, and $Q^* > \Delta^*$ and $\lambda^* \geq 0$ be the unique solution in Lemma 3.5.1. We now show that the control band policy $\{q^*, Q^*, S^*\}$ is optimal for the bounded inventory problem, where $q^* = Q(\lambda^*) - \Delta(\lambda^*)$, $Q^* = Q(\lambda^*)$ and $S^* = Q(\lambda^*) + s(\lambda^*) \leq M$.

Recall that, in the proof of Lemma 3.5.1, the relative relative value function of the control band policy $\{q^*, Q^*, S^*\}$ can be expressed as $V(x) = \int_0^x \pi(y)\, dy$ for $x \in [0, S^*]$, where $\pi(x) = f(x - Q)$ is given by (45).

For $S^* < x \leq M$, we extend V and π as

$$V(x) = V(Q^*) - L - \ell(x - Q^*), \qquad \pi(x) = -\ell. \tag{53}$$

Thus, the extended function V, still denoted by V, satisfies

$$V(x) = \int_0^x \pi(y) dy, \quad x \in [0, M). \tag{54}$$

We now show that the extended function satisfies all the conditions of Proposition 3.3.2.

First, condition (38) implies that V is continuous at S^*, thus continuous on $[0, M]$. Next, we show that V has continuous derivatives in $(0, M)$. If $S^* = M$, then $V'(x) = \pi(x)$ for $x \in (0, M)$, and thus $V'(x)$ is continuous in $(0, M)$. Now, assume that $S^* < M$. By (43), $\lambda^* = 0$. Therefore, condition (42) implies that the left side derivative of V at S^* is $-\ell$, which is equal to the right side derivative obtained from (53). Clearly, the second derivative of V is continuous on $[0, S^*)$ and on (S^*, M).

We now check that V satisfies condition (25) of Proposition 3.3.2. By construction

$$\Gamma V(x) - hx - \Gamma V(0) = 0 \text{ for } x \in [0, S^*]. \tag{55}$$

We show that V satisfies

$$\Gamma V - hx - \Gamma V(0) \leq 0 \text{ for } S^* < x \leq M. \tag{56}$$

It is enough to consider the case when $S^* < M$ and hence $\lambda^* = 0$.

From (55)

$$\begin{aligned} 0 &= V''(S^*-) + \mu V'(S^*-) - hS^* - \Gamma V(0) \\ &= V''(S^*-) - \mu\ell - hS^* - \Gamma V(0). \end{aligned}$$

Also $V''(S^*-) = \pi'(S^*-) = f'(s^*) \geq 0$. The latter inequality follows from the fact that $f(x)$ is strictly convex with its minimum in $(0, s^*)$ when $\lambda^* = 0$.

It follows that

$$-\mu\ell - hx - \Gamma V(0) \leq -\mu\ell - hS^* - \Gamma V(0) \leq 0$$

for all $S^* < x \leq M$, which proves (56).

We demonstrate that (26) holds. The arguments for (27) are analogous and we leave them to the reader.

Recalling that when $0 \leq y < x \leq M$, $V(x) - V(y) = \int_y^x \pi(z)dz$, we apply the observations of Lemma 3.5.2 to the following cases:

Case 1: $q^* \leq y < x \leq M$**.** In this case,

$$V(x) - V(y) = \int_y^x \pi(z)dz \leq k(x-y) \leq K + k(x-y).$$

Case 2: $0 \le y \le q^* < x \le M$**.** In this case,

$$\begin{aligned} V(y) - V(0) &= \int_0^y \pi(z)dz \ge ky \text{ and} \qquad (57) \\ V(x) - V(0) &= V(q^*) - V(0) + V(x) - V(q^*) = K + kq^* + \int_{q^*}^x \pi(z)dz \le K + kx \end{aligned}$$

from which it follows that $V(x) - V(y) \le K + k(x - y)$ as desired.

Case 3: $0 \le y < x \le q^*$**.** In this case, we still have (57) and

$$V(x) - V(0) = V(q^*) - V(0) - \int_x^{q^*} \pi(z)dz = K + kq^* - \int_x^{q^*} \pi(z)dz \le K + kx$$

from which (26) again follows.

Thus, by Proposition 3.3.2, $\mathrm{AC}(x, \varphi) \ge -\Gamma V(0)$ for each policy $\varphi \in \mathcal{P}$ and each initial state $x \in \mathbb{R}_+$, where $-\Gamma V(0)$ is the average cost of policy $\{q^*, Q^*, S^*\}$. It remains only to show that this same inequality holds fir $x < 0$, which we leave as an exercise. $\square$

Note that when M is infinite the Bounded Inventory Average Cost Brownian Control Problem becomes the Unconstrained Average Cost Brownian Control Problem. In this case it is clear that an optimal policy is the control band policy with parameters $\{q, Q, S\}$ determined by the solution to (5)-(7) (or (10)-(12)) with $\lambda = 0$.

Corollary 3.5.1. *A control band policy is optimal for the Unconstrained Average Cost Brownian Control Problem. The parameters of this optimal policy are the unique solution* $\{q^*, Q^*, S^*\}$ *to equations* (5)-(7) *(or* (10)-(12)*) with* $\lambda = 0$.

These results extend to the Bounded Inventory Discounted Cost Problem that imposes a bound on the maximum inventory level in the discounted problem described by Harrison, Sellke and Taylor [12]. We state the result without proof. Here, $\gamma > 0$ is the discount rate.

Theorem 3.5.1. *The Bounded Inventory Discounted Cost Brownian Control Problem admits an optimal policy that is a control band policy. Furthermore, the parameters* $\{q^*, Q^*, S^*\}$ *of the optimal control band policy* φ^* *are defined by the unique non-negative values of* λ, s,

Δ *and* $Q \geq \Delta$ *satisfying:*

$$L = rs + \frac{r(1-e^{-\rho s})(1-e^{\alpha s})}{(e^{\alpha s}-e^{-\rho s})}\left(\frac{1}{\alpha}+\frac{1}{\rho}\right) + \lambda\left(\frac{r}{e^{\alpha s}-e^{-\rho s}}\left(\frac{e^{\alpha s}-1}{\alpha}+\frac{e^{-\rho s}-1}{\rho}\right)\right) \tag{58}$$

$$c = r\left(\frac{1-e^{-\rho s}}{e^{\alpha s}-e^{-\rho s}}e^{-\alpha\Delta} + \frac{e^{\alpha s}-1}{e^{\alpha s}-e^{-\rho s}}e^{\rho\Delta}\right) + \lambda\left(\frac{r}{e^{\alpha s}-e^{-\rho s}}\left(e^{\rho\Delta}-e^{-\alpha\Delta}\right)\right), \tag{59}$$

$$\begin{aligned} K = {} & r\left[\frac{(1-e^{-\rho s})(e^{-\alpha\Delta}-e^{-\alpha Q})}{\alpha(e^{\alpha s}-e^{-\rho s})} + \frac{(e^{\alpha s}-1)(e^{\rho Q}-e^{\rho\Delta})}{\rho(e^{\alpha s}-e^{-\rho s})}\right] - c(Q-\Delta) \\ & +\lambda\left[\frac{r(e^{-\alpha Q}-e^{-\alpha\Delta})}{\alpha(e^{\alpha s}-e^{-\rho s})} + \frac{r(e^{\rho Q}-e^{\rho\Delta})}{\beta(e^{\alpha s}-e^{-\rho s})}\right], \end{aligned} \tag{60}$$

$$0 = \lambda(S-M), \tag{61}$$

$$S \leq M, \tag{62}$$

where,

$$r = h/\gamma - \ell, \tag{63}$$

$$c = h/\gamma + k > r, \tag{64}$$

$$\alpha = \left[(\mu^2+2\gamma\sigma^2)^{1/2}-\mu\right]/\sigma^2 > 0, \tag{65}$$

$$\rho = \left[(\mu^2+2\gamma\sigma^2)^{1/2}+\mu\right]/\sigma^2 > 0, \tag{66}$$

$S \equiv Q + s$ *and* $q \equiv Q - \Delta$.

3.6 Constrained Policies

In this section we add the constraints $-d \leq \xi_i \leq u$ on the magnitude of adjustments to the inventory to the Average Cost Brownian Control Problem.

One of the main contributions of this paper is a new technique based on methods of Lagrangian relaxation that reduce the constrained problem to a version of the original unconstrained problem and, in the process, provide methods for computing the optimal control band policy in the presence of the constraints.

Theorem 3.2.2 shows that a control band policy is optimal for the Constrained Average Cost Brownian Control Problem and provides optimality conditions that permit ready computation of the optimal policy parameters.

Proof. We prove Theorem 3.2.2 for the special case where $u = \infty$ and leave the rest as an exercise to the reader. In this case the constrained problem becomes

$$\mathrm{AC}^*(d,\infty) = \inf_{\varphi} \limsup_{n\to\infty} \mathbb{E}_x \left[\frac{1}{T_n} \left(\int_0^{T_n} hZ_t dt + \sum_{i=1}^{n} \left((K + k\xi_i) 1_{\{\xi_i > 0\}} + (L - \ell\xi_i) 1_{\{\xi_i < 0\}} \right) \right) \right] \tag{67}$$

$$\text{s.t. } -d \le \xi_i$$

$$\varphi \in \mathcal{P}.$$

Note that as in the unconstrained problem, $\mathrm{AC}^*(d,\infty)$ does not depend on the initial state x so we omit the initial state from our notation. The notation $\mathrm{AC}^*(d,\infty)$ indicates that $u = \infty$ and we are only imposing a bound on the magnitude of reductions. The notation $AC^*(\infty,\infty)$ refers to the unconstrained problem.

We proved in Corollary 3.5.1 that a control band policy is an optimal policy for the unconstrained problem:

$$\begin{aligned} \mathrm{AC}^* = \mathrm{AC}^*(\infty,\infty) \quad &= \quad \min_{\varphi} \mathrm{AC}(\varphi) \\ &\quad\ \ \text{s.t. } \varphi \in \mathcal{P}. \end{aligned} \tag{68}$$

We show that a control band policy solves the constrained problem (67) and, in fact, reduce this constrained problem to a version of the unconstrained problem (68).

Let $\{q^*, Q^*, S^*\}$ be the optimal solution to the unconstrained problem. There are two possible cases:

Case 1 $s^* = S^* - Q^* \le d$: In this case the control band policy that is optimal for the unconstrained problem is also optimal for the constrained problem.

Case 2 $s^* = S^* - Q^* > d$.

We prove that also in Case 2 a control band policy is an optimal policy for the constrained problem. To prove this result we use Lagrangian Relaxation, a classic method for constrained optimization that moves the constraint to the objective and assigns it a price λ. The resulting unconstrained problem provides a bound on the objective of the original constrained problem. We find a control band policy that achieves this bound thereby proving its optimality.

We introduce a Lagrange multiplier $\lambda \geq 0$ and move the constraint $-d \leq \xi_i$ to the cost function. For each scalar $\lambda \geq 0$ and policy $\varphi \in \mathcal{P}$, we define the Lagrangian function

$$\begin{aligned}\mathcal{L}(\varphi;\lambda) \quad = \quad & \limsup_{n\to\infty} \mathbb{E}_x\Bigg[\frac{1}{T_n}\Bigg(\int_0^{T_n} hZ_t dt \\ & + \sum_{i=1}^{n}\Big((K+k\xi_i)1_{\{\xi_i>0\}} + (L-\ell\xi_i)1_{\{\xi_i<0\}} - \lambda(d+\xi_i)1_{\{\xi_i<0\}}\Big)\Bigg)\Bigg].\end{aligned}$$

Since we show that the optimal policy does not depend on the initial state x we omit the initial state from the notation of the Lagrangian function. For fixed $\lambda \geq 0$, the Lagrangian primal is

$$\begin{aligned}\mathcal{L}(\lambda) \quad = \quad & \min \mathcal{L}(\varphi;\lambda) \\ & \text{s.t. } \varphi \in \mathcal{P}.\end{aligned} \tag{69}$$

Note that for each $\lambda \geq 0$, $\mathcal{L}(\varphi;\lambda) \leq \mathrm{AC}(\varphi)$ for each feasible control policy φ, and so $\mathcal{L}(\lambda) \leq \mathrm{AC}^*(d,\infty)$.

The Lagrangian problem (69) can be expressed as a version of the unconstrained problem with modified costs for reducing inventory. In particular,

$$\begin{aligned}\mathcal{L}(\lambda) \quad = \quad & \inf_{\varphi} \limsup_{n\to\infty} \mathbb{E}_x\Bigg[\frac{1}{T_n}\Bigg(\int_0^{T_n} hZ_t dt \\ & + \sum_{i=1}^{n}\Big((K+k\xi_i)1_{\{\xi_i>0\}} + (L-\ell\xi_i)1_{\{\xi_i<0\}} - \lambda(d+\xi_i)1_{\{\xi_i<0\}}\Big)\Bigg)\Bigg] \\ & \text{s.t. } \varphi \in \mathcal{P} \\ = \quad & \inf_{\varphi} \limsup_{n\to\infty} \mathbb{E}_x\Bigg[\frac{1}{T_n}\Bigg(\int_0^{T} hZ_t dt \\ & + \sum_{i=1}^{n}\big((K+k\xi_i)1_{\{\xi_i>0\}} + (L-\lambda d-(\ell+\lambda)\xi_i)1_{\{\xi_i<0\}}\big)\Bigg)\Bigg] \\ & \text{s.t. } \varphi \in \mathcal{P}.\end{aligned}$$

Hence, for each $0 \leq \lambda < L/d$, the Lagrangian problem admits a control band policy $\varphi_\lambda = \{q_\lambda, Q_\lambda, S_\lambda\} \in \mathcal{P}$ as an optimal policy and so $\mathcal{L}(\varphi_\lambda;\lambda) = \mathcal{L}(\lambda)$ for each $0 \leq \lambda < L/d$. Now consider the dual problem

$$\begin{aligned}\mathcal{L} \quad = \quad & \max_{\lambda} \mathcal{L}(\lambda) \\ & \text{s.t. } \lambda \geq 0.\end{aligned} \tag{70}$$

Note that $\mathcal{L} \leq \mathrm{AC}^*(d, \infty)$. If we can find a multiplier $0 \leq \lambda^* \leq L/d$ such that the control band policy $\varphi_{\lambda^*} = \{q_{\lambda^*}, Q_{\lambda^*}, S_{\lambda^*}\}$ satisfies

$$s_{\lambda^*} = S_{\lambda^*} - Q_{\lambda^*} \leq d, \text{ i.e., } \xi_i = -s_{\lambda^*} \geq -d \text{ whenever } \xi_i < 0 \tag{71}$$

and

$$\lambda^*(d - s_{\lambda^*}) = 0, \tag{72}$$

then

$$\mathrm{AC}(\varphi_{\lambda^*}) = \mathcal{L}(\varphi_{\lambda^*}; \lambda^*) = \mathcal{L}(\lambda^*) \leq \mathcal{L} \leq \mathrm{AC}(d, \infty) \leq \mathrm{AC}(\varphi_{\lambda^*}),$$

proving that λ^* is an optimal solution to the dual problem (70) and φ_{λ^*} is an optimal policy for the constrained problem. When the unconstrained problem yields $s^* = S^* - Q^* \leq d$, $\lambda^* = 0$ and the proof is complete. When this is not the case, existence of a Lagrange multiplier $0 \leq \lambda^* < L/d$ satisfying (71) and (72) for each $d > 0$ can also be shown. See Appendix A.4 for details of the proof. Thus we have proved Theorem 3.2.2 for the special case where $u = \infty$ $\square$

In the discounted cost problem when the magnitude of adjustments is bounded, using similar arguments it is possible to prove the optimality of control band policies. We briefly describe the mechanism behind the proof for the special case $u = \infty$, but omit the details. Following the notation of Harrison, Sellke and Taylor [12] in which the problem is formulated as finding a policy that maximizes the reward, we let $R(s) = -L + rs$ denote the return achieved each time the upper boundary is hit, where r is defined in (63). We introduce a Lagrange multiplier $\lambda \geq 0$ and move the constraint $s \leq d$ to the cost function. Thus whenever the upper bound is hit the system incurs a reward $R(s, \lambda) = -L + rs + \lambda(d - s)$. Note that since $\lambda \geq 0$ and $d - s \geq 0$, $R(s, \lambda) \geq R(s)$ for all feasible s and so the Lagrangian problem provides an upper bound on the objective.

We may rewrite $R(s, \lambda)$ as

$$R(s, \lambda) = -(L - \lambda d) + (r - \lambda)s.$$

So $R(s, \lambda)$ is equivalent to the original discounted cost problem with parameters $L - \lambda d$ and $r - \lambda$. Hence the optimal solution is again a control band policy. It is easy to show the existence of a $\lambda < L/d$ so that the solution of the Lagrangian relaxation yields $s \leq d$ and the optimal solution is given by Theorem 3.6.1.

Theorem 3.6.1. *The Constrained Discounted Cost Brownian Control Problem with $\xi_i \geq -d$ admits an optimal policy that is a control band policy. Furthermore, the parameters $\{q^*, Q^*, S^*\}$ of the optimal control band policy φ^* are defined by the unique non-negative values of $\lambda \leq L/d$, s, Δ and $Q \geq \Delta$ satisfying:*

$$
\begin{aligned}
L - \lambda d &= (r-\lambda)s + \frac{(r-\lambda)(1-e^{-\rho s})(1-e^{\alpha s})}{(e^{\alpha s}-e^{-\rho s})}\left(\frac{1}{\alpha}+\frac{1}{\rho}\right), && (73)\\
c &= (r-\lambda)\left(\frac{1-e^{-\rho s}}{e^{\alpha s}-e^{-\rho s}}e^{-\alpha\Delta} + \frac{e^{\alpha s}-1}{e^{\alpha s}-e^{-\rho s}}e^{\rho\Delta}\right), && (74)\\
K &= (r-\lambda)\left[\frac{(1-e^{-\rho s})(e^{-\alpha\Delta}-e^{-\alpha Q})}{\alpha(e^{\alpha s}-e^{-\rho s})} + \frac{(e^{\alpha s}-1)(e^{\rho Q}-e^{\rho\Delta})}{\rho(e^{\alpha s}-e^{-\rho s})}\right] - c(Q-\Delta), && (75)\\
0 &= \lambda(s-d), && (76)\\
s &\leq d, && (77)
\end{aligned}
$$

where, r, c, α, and ρ are given in (63)–(66)*, $S \equiv Q + s$ and $q \equiv Q - \Delta$.*

3.7 Concluding Remarks

In this paper, we showed that an optimal policy for the Average Cost Brownian Control Problem is a control band policy and demonstrated how to calculate the optimal control parameters explicitly. We also considered the Bounded Inventory Average Cost Brownian Control Problem in which the total available inventory space is bounded, and the Constrained Average Cost Brownian Control Problem in which the magnitude of each adjustment is bounded. The Bounded Inventory, Constrained Average Cost Brownian Control Problem combines these two problems and simultaneously imposes an upper bound d and a lower bound u on the controls, and an upper limit M on inventory. In this general setting, one can show that a control band policy is still optimal and its parameters, $\{q, Q, S\}$, when $\mu \neq 0$, can be determined from the unique non-negative values of λ, η, τ, s, Δ and $Q \geq \Delta$

satisfying:

$$L - \tau d = -h\left(\frac{s^2(1+e^{\beta s})}{2\mu(1-e^{\beta s})} + \frac{s}{\beta\mu}\right) + \lambda\left(\frac{s}{1-e^{-\beta s}} - \frac{1}{\beta}\right), \tag{78}$$

$$k + \eta + \ell + \tau = -\frac{h\Delta}{\mu} - \frac{hs(1-e^{\beta\Delta})}{\mu(1-e^{-\beta s})} + \lambda\left(\frac{e^{\beta\Delta}-1}{1-e^{-\beta s}}\right), \tag{79}$$

$$\begin{aligned} K - \eta u = {} & \frac{h(Q-\Delta)se^{\beta s}}{\mu(1-e^{\beta s})} + \frac{h(\Delta^2 - Q^2)}{2\mu} + \frac{hs(e^{\beta Q} - e^{\beta\Delta})}{\mu\beta(1-e^{-\beta s})} \\ & -(\ell + \tau + k + \eta)(Q-\Delta) + \lambda\left(\frac{e^{\beta Q} - e^{\beta\Delta}}{\beta(1-e^{-\beta s})} - \frac{(Q-\Delta)}{1-e^{-\beta s}}\right) \end{aligned} \tag{80}$$

such that

$$\lambda(S - M) = 0, \qquad S \leq M,$$

$$\tau(d - s) = 0, \qquad s \leq d,$$

$$\eta(u - Q + \Delta) = 0, \qquad Q - \Delta \leq u,$$

where $\beta = \frac{2\mu}{\sigma^2}$, $q \equiv Q - \Delta$ and $S \equiv Q + s$. (When $\mu = 0$, (78)-(80) are modified accordingly.)

This paper focused on an inventory control problem whose netput process follows a Brownian motion that has continuous sample paths. However, in most applications the netput process is a pure jump process; for example, a downward jump signifies the fulfillment of a customer order. One hopes to identify a class of inventory systems whose netput processes can be discontinuous such that our optimal policy to the average cost Brownian control problem provides some key insights to these systems. In the manufacturing setting, the justification of such procedure is often carried out through *heavy traffic limit theorems*; see, for example, Krichagina et al. [18]. Plambeck [31] proved a heavy traffic limit theorem for an assemble-to-order system with capacitated component production and fixed transport costs. The limit theorem enables her to find an asymptotically optimal control for the system.

An important contribution of this paper is to develop a method based on Lagrangian relaxation to solve constrained stochastic problems. Lagrangian relaxation methods have been widely used in deterministic optimization problems, both to solve constrained problems optimally and to obtain lower bounds on the optimal solution where it can not be solved to optimality. In this paper we showed that Lagrangian relaxation techniques can be adapted to solve stochastic control problems as well. This approach makes it possible to study a whole new venue of problems. The next chapter explores this approach in more detail and describes additional applications of it.

CHAPTER IV

LAGRANGIAN RELAXATION APPROACHES FOR CONSTRAINED STOCHASTIC CONTROL PROBLEMS

4.1 Introduction

This chapter has three kinds of motivations: To illustrate and explore the power and applicability of the Lagrange technique, to demonstrate specific technical characteristics of the problem or approach, and to capture additional features encountered in industrial applications in the Brownian control model.

Lagrangian relaxation methods have been widely used in deterministic optimization problems, both to solve constrained problems optimally and to obtain lower bounds on the optimal solution. Lagrangian relaxation is a classic method for constrained optimization that assigns a price λ to the constraint and moves it to the objective. The resulting unconstrained problem provides a bound on the objective of the original constrained problem. Finding a policy that achieves this bound proves its optimality. For more detailed explanation of Lagrangian methods in deterministic problems see for example Bertsekas [2].

This chapter further illustrates and explores a Lagrangian solution approach initially used in Ormeci et al. [30] (which is also Chapter 3 of this thesis) to solve Brownian control problems with simple constraints. In this chapter we build on this method and extend its use to more general constraints.

Lagrangian methods, have been applied to stochastic optimization and control problems in the past in different capacities. For example in the works by Dentcheva and Römisch [8], Märkert and Schultz [25], Nowak and Römisch [29], Sen, Higle and Birge [34] and Takriti and Birge [37] the stochastic problem is converted to a deterministic problem by assigning probabilities to a finite number of scenarios. Then Lagrangian methods are used as in deterministic problems to decompose the problem or obtain bounds or as part of a heuristic in proposed algorithms. On the other hand Chen, Dubrawski and Meyn [4] develop a fluid network model of their problem and apply Lagrangian relaxation methods to this deterministic fluid model. Luh and Feng[20], Luh et al. [21, 22] develop a solution methodology that combines Lagrangian relaxation, stochastic dynamic programming and heuristics. They replace the constraint on a random variable like $Ax \leq d$ with its expectation, i.e., $\mathbb{E}[Ax] \leq d$, and then apply Lagrangian relaxation to this deterministic version of the constraint. Decomposing the problem into sublevel problems with this kind of Lagrangian relaxation method they then use stochastic dynamic programming and some heuristic search methods to solve it.

Chow [5] describes Lagrange multipliers method as an alternative to dynamic programming methods. He maximizes the objective function of a multi period problem subject to a dynamic equation that defines the evolution of the state. In this case the constraint defines how the state variables x_t change based on the control and some random shocks at each time period. Chow solves this multiple period problem by assigning a Lagrange multiplier at each period to this constraint and then solving the Lagrangian relaxation using backward induction.

Bäuerle [1] solves a mean-variance problem, where the solution approach resembles the methodology developed in Chapter 3. The aim is to minimize the risk of the terminal reserve measured by the variance over all admissible policies which yield the same terminal reserve, i.e. the variance is minimized subject to the condition that the mean equals a certain value $(\min\{\mathbb{E}_{x_0}[(X_T - b)^2] : \mathbb{E}_{x_0}[X_T] = b\})$. A Lagrange multiplier is assigned to the constraint and moved to the objective so that the problem reduces to its unconstrained version. This approach differs from the one introduced in Chapter 3 in that the constraint is already within an expectation operator, and the process is not modelled by Brownian Motion.

In this chapter we consider several stochastic control problems to investigate the conditions under which the type of solution approach developed in Chapter 3 is applicable, and to capture additional features encountered in industrial applications. In the process, we solve several Brownian Motion Control Problems with constraints. The unconstrained problem on which our variations are based was initially solved by Harrison et al. [12] in the discounted cost setting. Ormeci et al. [30] solve a similar problem in the average cost setting. In this chapter we extend the one dimensional Brownian control problem to multiple dimensions, and then introduce constraints on total inventory and adjustment quantities. We look at several different problems that reflect these constraints. In Section 4.3 we introduce a bound on the total available average inventory. We introduce the constraint $\sum_{i=1}^{m} \mathbb{E}Z_i \leq M$, where $\mathbb{E}Z_i$ is the average inventory for part i. While the warehouse space is fixed, quite often manufacturers are able to extend it by keeping parts in the yard in trailers. As long as the average inventory is below a certain quantity manufacturers are able to accommodate these extra parts. The model in this section approximates this situation.

One of the principles of lean manufacturing is the notion of footprinting or clearly identifying the floor space allocated to each part or function. One consequence of this philosophy is that floor space is dedicated to a single use. In Section 4.4 we look at a problem that reflects this feature. We allocate a maximum amount of space to each part and require the sum of the space allocated to each part to be less than an upper bound M.

To guarantee that the total expedites do not exceed the capacity of the mode of transportation, in Section 4.5 we introduce bounds on the maximum adjustment quantity so that the sum is less than a given quantity $p < \infty$.

To demonstrate specific technical characteristics of the approach in Section 4.6, we add constraints on a problem solved by Taksar [38], which is a similar one dimensional Brownian control problem, but has singular control and no fixed costs of adjustments. While the general idea is similar in solving these constrained problems, each problem requires some adaptation of the approach.

4.2 Impulse control of Brownian motion

In this chapter, unless otherwise stated, we use the following notation and assumptions. Let Ω be the space of all continuous functions $\omega : [0, \infty) \to \mathbb{R}^m$. For $t \geq 0$ let $X(t) : \Omega \to \mathbb{R}^m$ be the coordinate projection map $X(t)(\omega) = \omega(t)$. Then $X = (X(t),\, t \geq 0)$ is the canonical process on Ω. Let $\mathcal{F} = \sigma(X(t),\, t \geq 0)$ denote the smallest σ-field such that $X(t)$ is $\mathcal{F}$-measurable for each $t \geq 0$, and similarly let $\mathcal{F}_t = \sigma(X(s),\, 0 \leq s \leq t)$ for

$t \geq 0$. When we mention adapted processes and stopping times hereafter, the underlying filtration is understood to be $\{\mathcal{F}_t,\ t \geq 0\}$. Finally, for each $x \in \mathbb{R}^m$ let $\mathbb{P}_x$ be the unique probability measure on $(\Omega, \mathcal{F})$ such that each component of X, X_j, is an independent Brownian motion with drift μ_j, variance σ_j^2 and starting state x_j under $\mathbb{P}_x$. Let $\mathbb{E}_x$ be the associated expectation operator.

We are to control m independent Brownian motions, where each $X_j = \{X_j(t),\ t \geq 0\}, j = 1, \ldots, m$, has mean μ_j, variance σ_j^2 and a starting state x_j. For each $j = 1, \ldots, m$ upward or downward adjustments, ξ_{jn}, are exerted at discrete times, T_{jn}, so that the resulting inventory processes represented by $Z_j = \{Z_j(t), t \geq 0\}$ remains such that $|Z| = Z_1 + \ldots + Z_m$ is within $[0, M]$. The case in which M is infinite and so Z is simply restricted to be nonnegative, is a special case. We adopt the convention that the sample path of Z is right continuous on $[0, \infty)$ having left limits in $(0, \infty)$. (The subscript j identifying a part may be omitted at times when it is clear from the context and will not be confused with the vector. Also the time parameter of a given process may be written either as a subscript or as functional argument for convenience.)

A policy $\{(T_{jn}, \xi_{jn}),\ n \geq 0, j = 1, \ldots, m\}$ consists of stopping times $\{T_{j0}, T_{j1}, \ldots\}$ at which control is exerted and random variables $\{\xi_{j0}, \xi_{j1}, \ldots\}$ representing the magnitude and direction of each control. So T_{jn} is the time at which we make the $(n+1)$st adjustment to inventory of part j and ξ_{jn} describes the magnitude and direction of that adjustment. Each adjustment ξ_{jn} must be $\mathcal{F}_{T_{jn}^-}$ measurable, where, for a stopping time τ, $\mathcal{F}_{\tau^-}$ is defined as in Definition I.1.11 of Jacod and Shiryaev [17].

When a policy increases inventory by $\xi > 0$, it incurs cost $K_j + k_j\xi$ representing the fixed costs $K_j > 0$ of changing the inventory and the variable costs $k_j\xi \geq 0$ that grow with the size of the adjustment for part j . When a policy reduces inventory, i.e., when it adjusts inventory by $\xi < 0$, it incurs cost $L_j - \ell_j\xi$ where $L_j > 0$ is the fixed cost for reducing inventory and $-\ell\xi_j \geq 0$ is the variable cost. Finally, we assume that inventory incurs a positive holding cost of $h_j > 0$ per unit per unit of time for part j.

Harrison et al. [12] proved that a simple form of policy, called a control band policy, is optimal for the discounted cost problem in the case of a single part. Ormeci et al. [30] proved a similar result in the average cost setting, furthermore they included some constraints on the magnitude of the control and maximum available space. They showed that control band policy is optimal for each problem. A control band policy is defined by three parameters $\{q, Q, S\}$ with $0 < q \leq Q < S$. When inventory falls to 0, the policy exerts a control to bring it up to level q. When the inventory rises to S, the maximum allowed, the policy exerts a control to reduce it by $s = S - Q$ and bring it to level Q. If the initial inventory level lies outside the range $[0, S]$, the policy exerts a one-time control ξ_0 to bring it to the closer of level q or level Q. Thus, the control band policy $\{q, Q, S\}$ is defined by $\{(T_n, \xi_n),$ $n \geq 0\}$, where $T_0 = 0$,

$$\xi_0 = \begin{cases} q - X_0, & \text{if } 0 \geq X_0, \\ 0, & \text{if } 0 < X_0 < S, \\ Q - X_0, & \text{if } X_0 \geq S \end{cases} \tag{81}$$

thereafter, $\{T_n,\ n > 0\}$ are the hitting times for $\{0, S\}$, i.e., $T_n = \{t > T_{n-1} : Z(t) = 0 \text{ or } Z(t) = S\}$ for all $n = 1, 2, ...$ and

$$\xi_n = \begin{cases} q, & \text{if } Z(T_n^-) = 0, \\ -s, & \text{if } Z(T_n^-) = S. \end{cases} \tag{82}$$

We consider the Average Cost Brownian Control Problem with multiple parts, $X_j, j = 1, \ldots, m$, where each part behaves as an independent Brownian Motion with drift μ_j and variance σ_j^2 in the absence of control. Control of size ξ_{jn} is exerted at time T_{jn} on part j. Z_j is the resulting net inventory process as before, and $N_j(t) = \max\{n \geq 0 : T_{jn} \leq t\}$. The aim is to find a non-anticipating policy $\varphi = \{(T_{jn}, \xi_{jn}),\ n \geq 0,\ j = 1, \ldots, m\}$ that minimizes:

$$\begin{aligned} \mathrm{AC}(x, \varphi) \quad &= \quad \limsup_{T \to \infty} \sum_{j=1,\ldots,m} \mathbb{E}_x \Bigg[\frac{1}{T} \Bigg(\int_0^T h Z_j(t)\, dt \\ &\qquad + \sum_{i=1}^{N_j(T)} \left((K + k\xi_{ji}) 1_{\{\xi_{ji} > 0\}} + (L + \ell|\xi_{ji}|) 1_{\{\xi_{ji} < 0\}} \right) \Bigg) \Bigg] \end{aligned} \tag{83}$$

the expected long-run average cost starting at a given initial point x.

We restrict our attention to $\mathcal{P}$, the set of policies satisfying only

$$\mathbb{P}_x(\infty > Z_j(t) \geq 0 \text{ for all } t > 0) \quad = \quad 1, \text{ for all } j \text{ and } x \in \mathbb{R}_+^m. \tag{84}$$

We impose additional constraints in each section and find the policy that minimizes the expected long-run average cost subject to these constraints.

We denote the average cost of part j under policy φ as

$$\begin{aligned}\mathrm{AC}_j(x,\varphi) \quad = \quad \limsup_{T\to\infty}\mathbb{E}_x\Bigg[\frac{1}{T}\Bigg(\int_0^T hZ_j(t)\,dt \\ + \sum_{i=1}^{N_j(T)}\Big((K+k\xi_{ji})1_{\{\xi_{ji}>0\}}+(L+\ell|\xi_{ji}|)1_{\{\xi_{ji}<0\}}\Big)\Bigg)\Bigg].\end{aligned}$$

4.3 Bounding The Average Total Inventory Space

Although available space in a manufacturer's warehouse is a fixed quantity, quite often, manufacturers are able to expand this space through the use of trailers in the yard. Thus, as long as the average inventory does not exceed a certain quantity, i.e. we don't rely on the trailers all the time, manufacturers are able to exceed available warehouse space. We model this by introducing a bound, $M<\infty$, on the average space occupied by all parts in the warehouse. We assume there are m parts, where $Z_j(t)$ denotes the number of part j available in inventory at time t, and $\mathbb{E}Z_j$ is the average inventory for part j. Letting a_j denote the amount of space each unit of part j requires we minimize the average cost:

$$\begin{aligned}\min \quad & \sum_j \mathrm{AC}_j \\ s.t. \quad & \sum_j a_j\mathbb{E}Z_j \le M \\ & \varphi\in\mathcal{P}.\end{aligned} \tag{85}$$

We call this the Bounded Average Total Inventory Average Cost Brownian Control Problem.

Theorem 4.3.1, states that a control band policy is optimal for the Bounded Average Total Inventory Average Cost Brownian Control Problem and provides optimality conditions that permit ready computation of the optimal policy parameters.

Theorem 4.3.1. *The Bounded Average Total Inventory Average Cost Brownian Control Problem admits an optimal policy that is a control band policy. Furthermore, if $\mu_j\neq 0$, the parameters $\{q_j^*,Q_j^*,S_j^*\}$ of the optimal control band policy φ^* for each part j are defined by* *the unique non-negative values of* λ, s, Δ *and* Q *satisfying:*

$$\begin{aligned} L &= -(h+a_j\lambda)\left(\frac{s^2(1+e^{\beta s})}{2\mu(1-e^{\beta s})}+\frac{s}{\beta\mu}\right), & (86)\\ k+\ell &= -\frac{(h+a_j\lambda)\Delta}{\mu}-\frac{(h+a_j\lambda)s(1-e^{\beta\Delta})}{\mu(1-e^{-\beta s})}, & (87)\\ K &= \frac{(h+a_j\lambda)(Q-\Delta)se^{\beta s}}{\mu(1-e^{\beta s})}+\frac{(h+a_j\lambda)(\Delta^2-Q^2)}{2\mu}+\frac{(h+a_j\lambda)s(e^{\beta Q}-e^{\beta\Delta})}{\mu\beta(1-e^{-\beta s})}\\ &\quad -(\ell+k)(Q-\Delta), & (88) \end{aligned}$$

such that

$$\begin{aligned} \lambda(\sum_j a_j\mathbb{E}Z_j-M) &= 0 & (89)\\ \sum_j a_j\mathbb{E}Z_j &\le M & (90) \end{aligned}$$

where $\beta=2\mu/\sigma^2$, $q\equiv Q-\Delta$ *and* $S\equiv Q+s$.

If $\mu_j=0$, *the parameters of the optimal control band policy* φ^* *for part* j *are defined by the unique non-negative values of* λ, s, Δ *and* Q *satisfying* (89) *and* (90)

$$\begin{aligned} L &= \frac{(h+a_j\lambda)s^3}{6\sigma^2}, & (91)\\ \ell+k &= \frac{(h+a_j\lambda)(\Delta^2+\Delta s)}{\sigma^2}, & (92)\\ K &= \frac{Q^3(h+a_j\lambda)}{3\sigma^2}+\frac{Q^2(h+a_j\lambda))s}{2\sigma^2}-\frac{\Delta^2(h+a_j\lambda)s}{2\sigma^2}-\frac{\Delta^3(h+a_j\lambda)}{3\sigma^2}\\ &\quad +(\Delta-Q)(\ell+k), & (93) \end{aligned}$$

where $q\equiv Q-\Delta$ *and* $S\equiv Q+s$.

Here for ease of notation we omitted the subscript j from the cost and Brownian Motion parameters.

In [30], Theorem 4.3.1 is proved for a single part. Here we extend the proof to m parts my showing that there exists a $\lambda\ge 0$ such that $\sum_j a_j\mathbb{E}Z_j(t)\le M$.

Proof. Let $\{q_j^*, Q_j^*, S_j^*\}$ be the optimal solution to the unconstrained problem for each part j. Letting $\mathbb{E}Z_j$ denote the average inventory for part j under this policy, there are two possible cases:

Case 1 $\sum_j a_j \mathbb{E}Z_j \leq M$: In this case the control band policy that is optimal for the unconstrained problem is also optimal for the constrained problem.

Case 2 $\sum_j a_j \mathbb{E}Z_j > M$.

We prove that even when Case 1 does not apply, a control band policy for each part is still an optimal policy for the constrained problem. To prove this result we use Lagrangian Relaxation, a classic method for constrained optimization that moves the constraint to the objective and assigns it a price λ. The resulting unconstrained problem provides a bound on the objective of the original constrained problem. We find a control band policy that achieves this bound thereby proving its optimality. We introduce a Lagrange multiplier $\lambda \geq 0$ and move the space constraint to the cost function. For each scalar $\lambda \geq 0$ and policy $\varphi \in \mathcal{P}$, we define the Lagrangian function

$$\mathcal{L}(\varphi; \lambda) = \sum_j \mathrm{AC}_j + \lambda(\sum_j a_j \mathbb{E}_x Z_j - M). \tag{94}$$

For fixed $\lambda \geq 0$, the Lagrangian primal is

$$\begin{aligned} \mathcal{L}(\lambda) \quad &= \quad \min \mathcal{L}(\varphi; \lambda) \\ &\quad\quad \text{s.t. } \varphi \in \mathcal{P}. \end{aligned} \tag{95}$$

Note that for each $\lambda \geq 0$, $\mathcal{L}(\varphi; \lambda) \leq \mathrm{AC}(\varphi)$ for each feasible control policy φ, and so $\mathcal{L}(\lambda) \leq \min_\varphi \mathrm{AC}(\varphi) = \mathrm{AC}(\varphi^*)$.

We show that there exists a multiplier $\lambda \geq 0$ and a feasible policy φ^* under which each part is governed by a control band policy $\{q_j, Q_j, S_j\}$ and $\lambda(\sum_j a_j \mathbb{E}_x Z_j - M) = 0$. Hence, $\mathcal{L}(\varphi^*; \lambda) = \mathrm{AC}(\varphi^*)$ and φ^* is optimal.

When minimizing $\mathcal{L}(\varphi; \lambda)$ over $\varphi \in \mathcal{P}$ we can ignore the last term, λM, since it is a constant. Then the Lagrangian problem separates. Define:

$$\begin{aligned} \mathcal{L}_j(\lambda) = \min_\varphi \mathcal{L}_j(\varphi; \lambda) = \min_\varphi \quad & \mathrm{AC}_j + \lambda a_j \mathbb{E}_x Z_j \\ & \varphi \in \mathcal{P}. \end{aligned}$$

We rewrite $\mathcal{L}_j(\varphi;\lambda)$ as

$$\begin{aligned}\mathcal{L}_j(\varphi;\lambda) &= \limsup_{T\to\infty}\frac{1}{T}\mathbb{E}_x\left[\int_0^T h(Z_j(t))\,dt + \sum_{i=1}^{N_j(T)}\phi(\xi_{ji})\right] + \lambda(a_j\mathbb{E}_x Z_j)\\ &= \limsup_{T\to\infty}\frac{1}{T}\mathbb{E}_x\left[\int_0^T (h+a_j\lambda)(Z_j(t))\,dt + \sum_{i=1}^{N_j(T)}\phi(\xi_{ji})\right].\end{aligned}$$

The problem of minimizing $\mathcal{L}(\varphi;\lambda)$ over $\varphi\in\mathcal{P}$ is equivalent to minimizing the average cost Brownian control problem with holding cost $h+a_j\lambda$.

From the Basic Adjoint Relationship (BAR) of a single part (see Appendix A.2), we know that for a $\{q,Q,S\}$ policy, the average inventory, $\mathbb{E}Z$, is

$$\mathbb{E}Z = \frac{(S^2-Q^2)C}{2(sC-qD)} - \frac{q^2 D}{2(sC-qD)} - \frac{\sigma^2}{2\mu}$$

where,

$$\begin{aligned}C &= e^{-\frac{2\mu}{\sigma^2}q}-1\\ D &= e^{-\frac{2\mu}{\sigma^2}S}-e^{-\frac{2\mu}{\sigma^2}Q}.\end{aligned}$$

It can be shown that for the s,Δ,Q satisfying (86)-(88), $\frac{ds(\lambda)}{d\lambda}<0$, $\frac{d\Delta(\lambda)}{d\lambda}<0$, $\frac{dQ(\lambda)}{d\lambda}<0$ and that $s(\lambda),\Delta(\lambda),Q(\lambda)\to 0$ as $\lambda\to\infty$.

What remains is to show that there is a $\lambda\geq 0$ such that $\sum_j a_j\mathbb{E}Z_j = M$. To see this note that when $\lambda=0$, $\sum_j a_j\mathbb{E}Z_j > M$, and when $\lambda\to\infty$, $\sum_j \mathbb{E}Z_j \to 0$ since each $s_j,\Delta_j,Q_j\to 0$. Hence, by the intermediate value theorem there exist a $\lambda\geq 0$ such that $\sum_j a_j\mathbb{E}Z_j = M$. □

4.4 Bounding The Sum Of Maximum Allocation For Each Part

One of the principles of lean manufacturing is the notion of footprinting or clearly identifying the floor space allocated to each part or function. In one extreme example of this, DaimlerChrysler's Smart Car facility in France has footprints for the potted plants decorating the assembly line. One consequence of this philosophy is that floor space is dedicated to a single use. In this setting we solve the following optimization problem:

$$\begin{aligned}
\min \quad & \sum_j \mathrm{AC}_j && (96)\\
s.t. \quad & a_j Z_j(t) \le M_j && (97)\\
& \sum_j M_j \le M, && (98)
\end{aligned}$$

where M_j is the maximum amount of space that may be dedicated to part j. The space M_j dedicated to each part j is itself a decision variable and the problem is to allocate the available space among the parts so as to minimize the total cost of inventory and control. For ease of exposition and without loss of generality, we assume each $a_j = 1$.

Once again there are two possible cases. Let $\{q_j^*, Q_j^*, S_j^*\}$ be the optimal solution to the unconstrained problem for each part j.

Case 1 $\sum_j S_j^* \le M$: In this case the control band policy that is optimal for the unconstrained problem is also optimal for the constrained problem.

Case 2 $\sum_j S_j^* > M$.

When the space constraint is not satisfied we turn to Lagrangian Relaxation. We assign two sets of multipliers: λ_j associated with constraint (97) for each part j, and υ associated with constraint (98). Moving the constraints, with their associated multipliers to the cost function we obtain an unconstrained problem that provides a bound on the objective of the original constrained problem. We find a control band policy that achieves this bound thereby proving its optimality.

First, consider the Lagrangian primal where only constraint (98) is moved to the cost function:

$$\begin{aligned}
\mathcal{L}(\upsilon) \quad &= \quad \min_{\varphi} \mathcal{L}(\varphi;\upsilon) = \min_{\varphi} \sum_j \mathrm{AC}_j(\varphi) + \upsilon(\sum_j M_j - M) && (99)\\
s.t. \quad & \quad\quad Z_j(t) \le M_j.
\end{aligned}$$

It is clear that for each fixed $\upsilon \ge 0$, $\mathcal{L}(\varphi;\upsilon) \le \mathrm{AC}(\varphi)$ for each feasible policy φ. Since the last term, $-\upsilon M$, in the objective is a constant it can be ignored for the moment. Now, observe that for each fixed $\upsilon \ge 0$, the problem (99) separates for each part.

In [30] the optimal solution to the bounded inventory average cost Brownian control problem for a single part is proved to be a control band policy. Let $\mathrm{AC}_B^j(M_j)$ denote the cost of the solution of the single part problem when part j is bounded by M_j, i.e., $\mathrm{AC}_B^j(M_j) = \min_\varphi\{\mathrm{AC}_j(\varphi) : 0 \leq Z_j(t) \leq M_j \text{ for } t \geq 0\}$.

To extend the solution to multiple parts, consider the following problem

$$\mathcal{L}_j(\upsilon) = \min_{M_j} \mathrm{AC}_B^j(M_j) + \upsilon M_j. \tag{100}$$

Observe $\mathrm{AC}_B^j(M_j)$ is a nonincreasing function of M_j. To see this note that for a given value of M_j if the optimal control band policy $\varphi_j = \{q_j, Q_j, S_j\}$ will be so that $S_j \leq M_j$ is not tight then it will still be optimal when M_j is replaced with $M_j' \geq M_j$, and $\mathrm{AC}_B^j(M_j) = \mathrm{AC}_B^j(M_j')$. Now suppose $M_j = S_j$, where S_j is the optimal solution to the unconstrained problem for part j. Then for any $M_j' < M_j$ the feasible policy space will be smaller hence $\mathrm{AC}_B^j(M_j) \leq \mathrm{AC}_B^j(M_j')$. Combining the arguments for these two cases it is clear that $\mathrm{AC}_B^j(M_j) \leq \mathrm{AC}_B^j(M_j')$ for $M_j' \leq M_j$ in general, and $AC_B^j(M_j)$ is a decreasing (not strictly) function of M_j.

Then for each fixed υ, solving $\mathcal{L}_j(\upsilon)$ is equivalent to minimizing a convex function (sum of a decreasing then constant function and linearly increasing function). Finding the optimal M_j for a given υ is just a matter of simple search.

To prove optimality we find a υ^* such that $\mathcal{L}(\upsilon^*) = AC(\varphi)$. Note that when $\upsilon = 0$, setting $M_j = S_j$, where S_j is the optimal upper control parameter of the unconstrained problem for part j, is an optimal solution to $\mathcal{L}(0)$. If $\sum_j S_j \leq M$, case 1 holds and $\upsilon^* = 0$ and this is the optimal solution. Otherwise observe that the optimal solution to (100) will yield decreasing M_j with increasing υ values. Thus, picking υ^* so that $\sum_j M_j^*(\upsilon^*) = M$, will ensure $\mathcal{L}(\upsilon^*) = AC(\varphi)$, hence proving optimality.

4.5 Bounding The Total Adjustment Quantities

In this section we further explore the capabilities of the Lagrange methods and introduce a constraint on the magnitude of upward adjustments in multiple parts setting. This constraint also reflects the desire by manufacturers to ensure that expedites do not exceed the capacity of the mode of transportation We let M, the bound on total inventory, be infinite and consider the problem in which the sum of upward adjustments is bounded. We model this problem as

$$\begin{aligned} \min \quad & \sum_j \mathrm{AC}_j && (101) \\ s.t. \quad & a_j \xi_j(t) \leq m_j, \text{ if } \xi_j > 0 && (102) \\ & \sum_j m_j \leq p, && (103) \end{aligned}$$

where m_j is the maximum amount of upward adjustment allowed for part j. The bound m_j assigned to each part j is itself a decision variable. The parameter p represents the maximum amount of total upward adjustments allowed. For ease of exposition and without loss of generality, we assume each $a_j = 1$.

Once again there are 2 possible cases. Let $\{q_j^*, Q_j^*, S_j^*\}$ be the optimal solution to the unconstrained problem for each part j.

Case 1 $\sum_j q_j^* \leq p$: In this case the control band policy that is optimal for the unconstrained problem is also optimal for the constrained problem.

Case 2 $\sum_j q_j^* > p$.

As in Section 4.4 we use Lagrangian Relaxation, and assign two sets of multipliers: λ_j associated with (102) constraint for each part j, and v associated with constraint (103). Moving the constraints, with their associated multipliers to the cost function we obtain an unconstrained problem that provides a bound on the objective of the original constrained problem. We find a control band policy that achieves this bound thereby proving its optimality.

First, consider the Lagrangian primal in which only constraint (103) is moved to the cost function:

$$\begin{aligned} \mathcal{L}(v) \quad &= \quad \min_{\varphi} \mathcal{L}(\varphi; v) = \min_{\varphi} \sum_j \mathrm{AC}_j(\varphi) + v(\sum_j m_j - p) && (104) \\ s.t. \quad & \quad \xi_j \leq m_j. \end{aligned}$$

It is clear that for each fixed $v \geq 0$, $\mathcal{L}(\varphi; v) \leq \mathrm{AC}(\varphi)$ for each feasible policy φ. Since the last term, $-vp$, in the objective is a constant it can be ignored for the moment. Now, observe that for each fixed $v \geq 0$, the problem (104) separates.

From [30] the optimal solution to the constrained average cost Brownian control problem for a single part is a control band policy. As before, we let $\mathrm{AC}_q^j(m_j)$ denote the cost of the solution of the single part problem when the upward adjustment on part j is bounded by m_j, i.e., $\mathrm{AC}_q^j(m_j) = \min_\varphi\{\mathrm{AC}_j(\varphi) : \xi_j \leq m_j, Z_t \geq 0 \text{ for } t \geq 0\}$.

To extend the solution to multiple parts, we consider the following problem

$$\mathcal{L}_j(\upsilon) = \min_{m_j} \mathrm{AC}_B^j(m_j) + \upsilon m_j.$$

Arguments similar to those in Section 4.4 yield that $\mathrm{AC}_q^j(m_j)$ is a decreasing function of m_j, and for each fixed υ, solving $\mathcal{L}_j(\upsilon)$ is equivalent to minimizing a convex function (sum of a decreasing then constant function and linearly increasing function).

To prove optimality we find an υ^* such that $\mathcal{L}(\upsilon^*) = AC(\varphi)$. Following arguments similar to those in Section 4.4 we conclude that picking υ^* that yields $\sum_j m_j^*(\upsilon) = p$ will result with the optimal solution where each part is governed by a control band policy.

4.6 Bounding The Available Inventory Space: Single Part, Singular Control

In this section we demonstrate a subtlety of the Lagrangian relaxation method. This also serves to show that the Lagrangian method requires careful consideration while relaxing the constraints. In this process, we study a special case of the problem described in Section 4.2 where we consider only one part, i.e. $m = 1$. Furthermore we assume that the fixed costs are zero, $K = L = 0$, and that singular control is employed. This is similar to the problem studied by Taksar [38]. We introduce an upper bound on the available inventory space, but instead of general holding costs we adopt a linear holding cost, $h > 0$. We consider the Average Cost Brownian Control Problem, which is to find a non-anticipating policy φ that minimizes:

$$\mathrm{AC}(x, \varphi) = \limsup_{T\to\infty} \mathbb{E}_x \left[\frac{1}{T} \left(\int_0^T hZ(t)\,dt + \int_0^T (kdA_t + \ell dR_t) \right) \right] \tag{105}$$

the expected long-run average cost starting at a given initial point x. Here $A(t)$ ($R(t)$) is the cumulative upward (downward) adjustment up to time t. We show that the optimal policy is a control limit policy. A control limit policy with parameters $\{0, S\}$, prescribes exerting minimal effort to keep the process between 0 and S. When φ is a control limit policy with parameters $\{0, S\}$, the average cost does not depend on the initial state x and hence we also use $AC(\varphi)$ or $AC(0,S)$ to denote its average cost. We minimize the average cost subject to the constraint that the available inventory space is bounded by M.

$$\min_{\varphi} \mathrm{AC}(x,\varphi) = \min_{\varphi} \limsup_{T\to\infty} \mathbb{E}_x \left[\frac{1}{T}\left(\int_0^T hZ(t)\,dt + \int_0^T k dA_t + \int_0^T \ell dR_t\right)\right] \tag{106}$$

$$\text{s.t.} \quad 0 \le Z_t \le M, \ \text{for } t \ge 0. \tag{107}$$

We observe that for any feasible policy

$$\limsup_{T\to\infty} \frac{1}{T}\int_0^T M dR_t \le \limsup_{T\to\infty} \frac{1}{T}\int_0^T Z_t dR_t.$$

We replace constraint (107) with this inequality and associate a Lagrangian multiplier λ with it.

Remark: Note that when using the Lagrangian relaxation on (106) the natural approach is to assign a multiplier to the constraint and move it to the objective. However, this approach does not work because when $(Z_t - M)\lambda$ is used the relaxation will only achieve $\mathbb{E}Z \le M$.

We define

$$\begin{aligned} \mathcal{L}(\varphi;\lambda) &= \limsup_{T\to\infty} \mathbb{E}_x \left[\frac{1}{T}\left(\int_0^T hZ(t)\,dt + \int_0^T k dA_t + \int_0^T \ell dR_t + \lambda(\int_0^T (Z_t - M) dR_t\right)\right] \\ &\le \mathrm{AC}(\varphi), \end{aligned}$$

which we rewrite as

$$\mathcal{L}(\varphi;\lambda) = \limsup_{T\to\infty} \mathbb{E}_x \left[\frac{1}{T}\left(\int_0^T hZ(t)\,dt + \int_0^T k dA_t + \int_0^T (\ell_\lambda(Z_t)) dR_t\right)\right],$$

where

$$\ell_\lambda(x) = \ell + \lambda(x - M).$$

Define

$$\mathcal{L}(\lambda) = \min_{\varphi} \mathcal{L}(\varphi;\lambda). \tag{108}$$

As before we consider the dual problem

$$\begin{aligned} \mathcal{L} &= \max_{\lambda} \mathcal{L}(\lambda) \\ &\quad \text{s.t. } \lambda \ge 0 \end{aligned}$$

and observe that if for a multiplier $\lambda^* \geq 0$ the control limit policy $\varphi_\lambda^* = \{0, S^*\}$ satisfying $S^* \leq M$ and $\lambda^*(S^* - M) = 0$ is optimal for the Lagrangian problem, then

$$\mathrm{AC}(\varphi_{\lambda^*}) = \mathcal{L}(\varphi_{\lambda^*}; \lambda^*) = \mathcal{L}(\lambda^*) \leq \mathcal{L} \leq \mathrm{AC}(\varphi_{\lambda^*}).$$

Using arguments similar to those in Section 3.3 it is a straightforward exercise to show that a lower bound, η, is obtained for $\mathcal{L}(\lambda)$, through a function f, called the relative value function, satisfying the following conditions:

$$\frac{\sigma^2}{2} f''(x) + \mu f'(x) - hx + \eta \quad \leq \quad 0, \text{ for almost all } x \geq 0 \tag{109}$$

$$-\ell_\lambda(x) \leq f'(x) \leq k \qquad \text{for all } x \geq 0. \tag{110}$$

Next we show that optimality is attained in problem (108) if (111)–(114) hold as well:

$$\frac{\sigma^2}{2} f''(x) + \mu f'(x) - hx + \eta \quad = \quad 0, \text{ for } 0 \leq x \leq S \tag{111}$$

$$f'(S) \quad = \quad -\ell_\lambda(S) \tag{112}$$

$$f'(0) \quad = \quad k \tag{113}$$

$$f''(S) \quad = \quad -\lambda. \tag{114}$$

Choosing λ so that complementary slackness is achieved, i.e. $\lambda(\int_0^T (Z_t - M) dR_t) = 0$ will ensure that the objective of the Lagrangian dual problem is equal to the average cost, hence, proving the optimality of the control limit policy $\{0, S\}$. In the remainder of this section we construct the result for $\mu \neq 0$, the proof for $\mu = 0$ is analogous.

Observe that

$$\pi(x) = f'(x) = \frac{hx}{\mu} + \frac{e^{-\beta x} - 1}{1 - e^{-\beta S}} \frac{hS}{\mu} + \frac{(e^{-\beta x} - 1)(\ell + k)}{1 - e^{-\beta S}} + k, \qquad \text{for } 0 \leq x \leq S$$

is a solution to $\sigma^2/2\ \pi''(x) + \mu \pi'(x) - h = 0$. Letting

$$f(x) = \int_0^x \pi(y) dy \text{ for } 0 \leq x \leq S$$

and

$$f(x) = f(S) - \int_S^x \ell_\lambda(y) dy \qquad \text{for } x > S, \tag{115}$$

$f(x)$ satisfies (111), (112) and (113). Furthermore it is easy to show that, this $f(x)$ satisfies (110).

It only remains to show that (114) is a necessary condition for (109) to hold when $x > S$. By (115), $f''(x) = -\lambda$ for $x \geq S$, and $f'(x) = -\ell_\lambda(x) = -(\ell + \lambda(x - M)) < -(\ell + \lambda(y - M)) = f'(y)$ for $S \leq y \leq x$. It follows that

$$-\frac{\sigma^2}{2}\lambda - \mu(\ell + \lambda(x - M)) - hx + \eta \leq -\frac{\sigma^2}{2}\lambda - \mu(\ell + \lambda(S - M)) - hS + \eta = 0$$

for all $x \geq S$. Choosing λ^* such that $\lambda^*(S^* - M) = 0$ ensures that the Lagrangian problem and the Brownian control problem yield the same objective, proving the optimality of the control limit policies.

4.7 Conclusion

The aim of this chapter was to illustrate and explore the power and applicability of the Lagrange technique, to demonstrate specific technical characteristics of the problem or approach, and to capture additional features encountered in industrial applications.

In this process, we considered several constrained stochastic control problems. We were able to solve these problems using Lagrangean relaxation methods, however, while solving these problems the approach had to be adapted for each problem. This was especially crucial in the example depicted in section 4.6. Thus, we conclude that while utilizing Lagrangian methods, a strong tool to solve stochastic constrained problems, the modelling aspect plays a very important role in the solution.

APPENDIX A

PROOFS AND ADDITIONAL NOTES

A.1 Policy Space Counterexample

In the discounted cost problem, Harrison et al. [12] required their policies to satisfy condition (33) to ensure that when f has bounded derivative $\mathbb{E}_x[e^{-\gamma T} f(Z_T)] \to 0$ as $T \to \infty$. One suspects that condition (34), a natural analog of condition (33) in our average cost setting, should analogously lead to

$$\mathbb{E}_x[f(Z_T)/T] \to 0 \quad \text{as} \quad T \to \infty \tag{116}$$

as long as the corresponding average cost is finite. In this section we construct a control policy φ with finite average cost under which condition (34) does not guarantee (116) when $f(x) = x$ for $x \in \mathbb{R}_+$.

Consider the policy $\hat{\varphi}$ defined as follows: Normally, we follow the control band policy $\varphi_0 = \{1, 2, 3\}$ that keeps inventory between 0 and 3. At times $T = 2^m$, for each integer m, however, we exert a control $\xi = 2^m$, and for the subsequent one unit of time shift the control band by 2^m. Thus in the time interval $(2^m, 2^m + 1]$ we follow a modified control band policy $\varphi_m = \{2^m, 2^m+1, 2^m+2, 2^m+3\}$ that keeps inventory between 2^m and 2^m+3, and brings the inventory up to $2^m + 1$ whenever $Z_t = 2^m$, and down to $2^m + 2$ whenever it hits $2^m + 3$. At time $2^m + 1$, we exert another control $\xi = -2^m$, move inventory back to a level between 0 and 3, and revert to the original control band policy $\varphi_0 = \{1, 2, 3\}$. We show that this policy has finite average cost and satisfies condition (34). Clearly,

$$\frac{Z(2^m)}{2^m} \geq 1 \qquad \text{for each } m \geq 1,$$

and thus $\mathbb{E}_x[Z_T/T]$ does not go to zero as $T \to \infty$.

To see that condition (34) is satisfied, let $\hat{\varphi} = \{(T_i, \xi_i) : i \geq 0\}$ and define

$$U = \limsup_{n\to\infty} \frac{1}{T_n} \sum_{i=1}^{n} (1 + |\xi_i|).$$

Let φ_1 be the set of controls exerted to shift the control band, i.e., $\varphi_1 = \{(T_m^\wedge, \xi_m^\wedge), m \geq 1\} \cup \{(T_m^\vee, \xi_m^\vee), m \geq 1\}$, where $T_m^\wedge = \xi_m^\wedge = 2^m$, $T_m^\vee = 2^m + 1$ and $\xi_m^\vee = -2^m$. We partition the set of positive integers into two index sets $\mathcal{I}_0$ and $\mathcal{I}_1$ such that $\varphi_1 = \{(T_i, \xi_i) : i \in \mathcal{I}_1\}$ and $\varphi_0 = \{(T_i, \xi_i) : i \in \mathcal{I}_0\}$. Then,

$$\begin{aligned} U &= \limsup_{n\to\infty} \left[\frac{1}{T_n} \sum_{i\in\mathcal{I}_0}^{n} (1 + |\xi_i|) + \frac{1}{T_n} \sum_{i\in\mathcal{I}_1}^{n} (1 + |\xi_i|) \right] \\ &\leq \limsup_{n\to\infty} \left[\frac{1}{T_n} \sum_{i\in\mathcal{I}_0}^{n} (1 + |\xi_i|) + \frac{2}{2^m} \sum_{i=1}^{m=\log T_n} (1 + |2^i|) \right] \\ &\leq \limsup_{n\to\infty} \frac{1}{T_n} \sum_{i\in\mathcal{I}_0}^{n} (1 + |\xi_i|) + 6 \\ &\leq 6 + \limsup_{n\in\mathcal{I}_0} \frac{1}{T_n} \sum_{i\in\mathcal{I}_0}^{n} (1 + |\xi_i|) = 6 + F \end{aligned}$$

Since $\varphi_0 = \{(T_i, \xi_i), i \in \mathcal{I}_0\}$ is exactly the controls of the usual control band policy $\{1, 2, 3\}$, F is simply the average cost of φ_0 when $h = 0$, $K = L = k = \ell = 1$ and so is finite.

To see that $AC(\hat{\varphi})$ is finite observe that an analogous argument shows that the average cost of the controls under $\hat{\varphi}$ is bounded by the average cost of the controls under φ_0 plus a finite constant that depends on K, L, k and ℓ. Moreover, the additional inventory cost incurred because of each shifted control band is $h2^m$ and so these shifts add at most

$$\sum_{i=1}^{m} \frac{h2^i}{2^m + 1} \leq \sum_{i=0}^{m-1} \frac{h}{2^i} < 2h$$

to the average inventory holding cost. Thus, the policy $\hat{\varphi}$ has finite average cost and satisfies condition (34), but since $\mathbb{E}[\frac{Z_{2^m}}{2^m}] \geq 1$ for $m \geq 1$, $\frac{\mathbb{E}[Z_T]}{T}$ does not converge to 0.

A.2 An Alternative Derivation of the Average Cost

Consider a control band policy φ with parameters $\{q, Q, S\}$. In this section, we give an alternative derivation of its average cost $\mathrm{AC}(\varphi)$ given in (39).

Under the control band policy we can express the inventory at time t as

$$Z_t = X_t - (S - Q)R_t + qA_t, \quad t \geq 0,$$

where $A_t = \sum_{s<t} 1\{Z_{s^-} = 0\}$ is the cumulative number of upwards adjustments by time t, and $R_t = \sum_{s<t} 1\{Z_{s^-} = S\}$ is the cumulative number of downward adjustments.

Under such a policy, the inventory process $Z = \{Z_t, t \geq 0\}$ is a strong Markov process, and since the state space is bounded it has a stationary distribution.

Suppose π is a stationary distribution for Z. Note that we use π in this section with a different meaning than in previous sections, and it should not to be confused with the derivative of the relative value function. Let $\mathbb{E}_\pi$ denote the expectation with respect to $\mathbb{P}_\pi$, where, for a Borel set A, $\mathbb{P}_\pi(A) \equiv \int_0^S \mathbb{P}_x(A)\pi(dx)$. For a continuous function $f : [0, S] \to \mathbb{R}$, since $\{f(Z_t), t \geq 0\}$ is stationary under probability $\mathbb{P}_\pi$,

$$\mathbb{E}_\pi [f(Z_t)] = \mathbb{E}_\pi [f(Z_0)] = \int_0^S f(x)\pi(dx) \quad \text{for all } t \geq 0. \tag{117}$$

We define

$$\begin{aligned} \overline{Z} &\equiv \mathbb{E}_\pi[Z_0], \\ \delta_S &= \mathbb{E}_\pi \left[\int_0^1 dR(s) \right], \\ \delta_0 &= \mathbb{E}_\pi \left[\int_0^1 dA(s) \right]. \end{aligned}$$

Here, $\overline{Z}$ is interpreted as the long-run average inventory level, δ_0 can be interpreted as the frequency with which inventory hits the lower control point 0 and δ_S as the frequency with which it hits the upper control point S.

For any function $f : [0, S] \to \mathbb{R}$ that is twice continuously differentiable, applying the Ito's formula as in Jacod and Shiryaev [17], using the stationarity condition (117), and following the derivation of the basic adjoint relationship in Harrison and Williams [13, 14], one has

$$\int_0^S \left(\frac{\sigma^2}{2} f^{''}(x) + \mu f^{'}(x) \right) \pi(dx) + (f(q) - f(0))\delta_0 + (f(Q) - f(S))\delta_S = 0. \tag{118}$$

Equation (118) is known as the basic adjoint relationship (BAR) and it holds for all twice continuously differentiable functions $f : [0, S] \to \mathbb{R}$. We now use BAR to obtain the long-run average cost for the policy φ.

We demonstrate the case when $\mu \neq 0$. When $\mu = 0$ the arguments are analogous. We evaluate the BAR for the three functions $f(x) = x$, $f(x) = x^2$ and $f(x) = e^{-\frac{2\mu}{\sigma^2}x}$ (we replace this last choice with $f(x) = x^3$ when $\mu = 0$) giving rise to the equations (119), (120) and

(121), respectively:

$$\mu + q\delta_0 - s\delta_S = 0, \tag{119}$$

$$\sigma^2 + 2\mu\overline{Z} + q^2\delta_0 + (Q^2 - S^2)\delta_S = 0, \tag{120}$$

$$(e^{-\frac{2\mu}{\sigma^2}q} - 1)\delta_0 + (e^{-\frac{2\mu}{\sigma^2}Q} - e^{-\frac{2\mu}{\sigma^2}S})\delta_S = 0. \tag{121}$$

From (119) – (121) we see that δ_0, δ_S and $\overline{Z}$ are given by:

$$\begin{aligned}
\delta_0 &= \frac{\mu D}{sC - qD}, \\
\delta_S &= \frac{\mu C}{sC - qD}, \\
\overline{Z} &= \frac{(S^2 - Q^2)C}{2(sC - qD)} - \frac{q^2 D}{2(sC - qD)} - \frac{\sigma^2}{2\mu},
\end{aligned}$$

where,

$$\begin{aligned}
C &= e^{-\frac{2\mu}{\sigma^2}q} - 1 \\
D &= e^{-\frac{2\mu}{\sigma^2}S} - e^{-\frac{2\mu}{\sigma^2}Q}.
\end{aligned}$$

Observe that when $\mu \neq 0$, $sC - qD \neq 0$ and so the average cost under policy φ is given by

$$\begin{aligned}
\mathrm{AC}(\varphi) &= h\overline{Z} + (K + kq)\delta_0 + (L + \ell s)\delta_S \\
&= h\left(\frac{(S^2 - Q^2)C}{2(sC - qD)} - \frac{q^2 D}{2(sC - qD)} - \frac{\sigma^2}{2\mu}\right) + \frac{(K + kq)\mu D}{sC - qD} + \frac{(L + \ell s)\mu C}{sC - qD},
\end{aligned}$$

which is identical to (39). Note that the average cost $\mathrm{AC}(\varphi)$ of a control band policy φ is independent of the initial state of the system.

A.3 Proof of Lemma 3.5.1, Part (a)

To facilitate the proof of part (a) of Lemma 3.5.1, we re-write (5)–(7) as:

$$\begin{aligned}
L &= u_1(s) + \lambda v_1(s), \\
k + \ell &= u_2(s, \Delta) + \lambda v_2(s, \Delta), \\
K &= u_3(s, \Delta, Q) + \lambda v_3(s, \Delta, Q),
\end{aligned}$$

and first prove the following technical lemmas. Lemma A.3.1 shows that the right-hand-sides of (5) – (7) are increasing functions of the relevant arguments and so, for each $\lambda \geq 0$,

(5) – (7) admits a unique solution $\{s(\lambda), \Delta(\lambda), Q(\lambda)\}$. Lemma A.3.2 shows that $s(\lambda)$, $\Delta(\lambda)$ and $Q(\lambda)$ are decreasing functions of λ and so there is a unique $\lambda \geq 0$ satisfying (8) and (9).

Lemma A.3.1. *The functions*

$$
\begin{aligned}
u_1(s) &\equiv -h\left(\frac{s^2(1+e^{\beta s})}{2\mu(1-e^{\beta s})} + \frac{s}{\beta\mu}\right) \text{ and} \\
v_1(s) &\equiv \left(\frac{s}{1-e^{-\beta s}} - \frac{1}{\beta}\right)
\end{aligned}
$$

are strictly increasing functions of $s \geq 0$ satisfying $u_1(0) = 0$, $\lim_{s\to\infty} u_1(s) = \infty$ and $v_1(0) = 0$. For each fixed $s > 0$, the functions

$$
\begin{aligned}
u_2(s,\Delta) &\equiv -\frac{h\Delta}{\mu} - \frac{hs(1-e^{\beta\Delta})}{\mu(1-e^{-\beta s})} \text{ and} \\
v_2(s,\Delta) &= \left(\frac{e^{\beta\Delta}-1}{1-e^{-\beta s}}\right)
\end{aligned}
$$

are strictly increasing functions of Δ satisfying $u_2(s,0) = v_2(s,0) = 0$. Further, $u_2(s,\Delta)$ is a strictly convex function of $\Delta \geq 0$ and so $\lim_{\Delta\to\infty} u_2(s,\Delta) = \infty$.

Finally, define

$$
\begin{aligned}
u_3(s,\Delta,Q) &\equiv \frac{h(Q-\Delta)se^{\beta s}}{\mu(1-e^{\beta s})} + \frac{h(\Delta^2-Q^2)}{2\mu} + \frac{hs(e^{\beta Q}-e^{\beta\Delta})}{\mu\beta(1-e^{-\beta s})} - (\ell+k)(Q-\Delta) \text{ and} \\
v_3(s,\Delta,Q) &\equiv \left(\frac{e^{\beta Q}-e^{\beta\Delta}}{\beta(1-e^{-\beta s})} - \frac{(Q-\Delta)}{1-e^{-\beta s}}\right).
\end{aligned}
$$

For each fixed $\lambda \geq 0$, s and Δ satisfying $u_2(s,\Delta) + \lambda v_2(s,\Delta) = k + \ell$,

$$
w(Q) \equiv u_3(s,\Delta,Q) + \lambda v_3(s,\Delta,Q)
$$

is a non-negative, strictly increasing, convex function of Q for $Q \geq \Delta$ such that $w(\Delta) = 0$.

Proof of Lemma A.3.1. We start with $u_1(s)$. First, consider the case in which $\mu \neq 0$, with $\beta = 2\mu/\sigma^2$. Observe that since $e^x > 1 + x$ for $x \neq 0$, $e^x(e^x) = e^{2x} > e^x + xe^x$ and hence $e^{2x} - e^x - xe^x > 0$ for $x \neq 0$. Hence, $(\beta e^{2x} - \beta e^x - \beta xe^x)/\mu > 0$ for $x \neq 0$ and so, replacing x with βs, we see that $(\beta e^{2\beta s} - \beta e^{\beta s} - \beta^2 se^{\beta s})/\mu > 0$ for $s > 0$.

Now, consider the function $\phi(s) = (e^{2\beta s} - 1 - 2s\beta e^{\beta s})/\mu$. Since $\phi(0) = 0$ and $\phi'(s) = 2\left(\beta e^{2\beta s} - \beta e^{\beta s} - \beta^2 se^{\beta s}\right)/\mu > 0$ for $s > 0$, it follows that $\phi(s)$ is a strictly positive, increasing function for $s > 0$.

Now,

$$\begin{aligned}
\phi(s) &= (e^{\beta s}-1)(e^{\beta s}+1) + s\beta e^{\beta s}(e^{\beta s}-1) - s(1+e^{\beta s})\beta e^{\beta s}/\mu \\
&= (e^{\beta s}-1)^2\left[\frac{1+e^{\beta s}}{e^{\beta s}-1} + \frac{s\beta e^{\beta s}}{e^{\beta s}-1} - \frac{s(1+e^{\beta s})\beta e^{\beta s}}{(e^{\beta s}-1)^2}\right]/\mu \\
&= (e^{\beta s}-1)^2 y(s)
\end{aligned}$$

where

$$y(s) = \left[\frac{1+e^{\beta s}}{e^{\beta s}-1} + \frac{s\beta e^{\beta s}}{e^{\beta s}-1} - \frac{s(1+e^{\beta s})\beta e^{\beta s}}{(e^{\beta s}-1)^2}\right]/\mu.$$

It follows that $y(s)$ is strictly positive for $s > 0$.

But $y(s)$ is the derivative of

$$Y(s) = \frac{s(1+e^{\beta s})}{\mu(e^{\beta s}-1)}$$

and so, $Y(s)$ is an increasing function for $s > 0$. We can rewrite $u_1(s)$ as

$$u_1(s) = \frac{sh}{\mu}\left(\frac{s(1+e^{\beta s})}{2(e^{\beta s}-1)} - \frac{1}{\beta}\right) = sh\left(\frac{Y(s)}{2} - \frac{1}{\mu\beta}\right).$$

Since $\lim_{s\to 0} Y(s) = \frac{2}{\mu\beta}$, $Y(s)$ is increasing with s and $\lim_{s\to 0} Y(s) = \infty$, $u_1(s)$ increases continuously from 0 to ∞ as s increases from 0 to ∞.

Consider the partial derivative

$$\begin{aligned}
\frac{\partial v_1(s)}{\partial s} &= \frac{1}{1-e^{-\beta s}} - \frac{\beta s e^{-\beta s}}{(1-e^{-\beta s})^2} \\
&= \frac{1-e^{-\beta s}(1+\beta s)}{(1-e^{-\beta s})^2}.
\end{aligned}$$

Note the the denominator is always positive. So it suffices to show that the numerator is strictly positive. Since $e^x > 1+x$ for $x \neq 0$, it follows that $\frac{(1+\beta s)}{e^{\beta s}} < 1$. We conclude that since $\frac{\partial v_1(s)}{\partial s} > 0$, $\psi(s)$ is an increasing function.

We next demonstrate that for each fixed value of $s > 0$, $u_2(\Delta)$ is a strictly increasing convex function.

$$u_2(s,\Delta) = -\frac{h\Delta}{\mu} - \frac{hs(1-e^{\beta\Delta})}{\mu(1-e^{-\beta s})},$$

and

$$\frac{\partial u_2(s,\Delta)}{\partial \Delta} = \frac{h}{\mu}\left(-1 + \frac{s\beta}{1-e^{-\beta s}}e^{\beta\Delta}\right).$$

When $\mu > 0$, $e^{\beta\Delta} > 1$ for $\Delta > 0$. Further, since $e^{-x} > 1 - x$, $x > 1 - e^{-x}$ for $x > 0$, replacing x with βs, we see that $\frac{s\beta}{1-e^{-\beta s}} > 1$ for $s > 0$. Thus, $\frac{\partial u_2(s,\Delta)}{\partial \Delta} > 0$ and $u_2(\Delta)$ is an increasing function for $\Delta > 0$. When $\mu < 0$, analogous arguments show that $\frac{\partial u_2(s,\Delta)}{\partial \Delta} > 0$ and hence $u_2(s, \Delta)$ is an increasing function for $\Delta > 0$. Furthermore $\frac{\partial^2 u_2(s,\Delta)}{\partial \Delta^2} > 0$, so that $u_2(s, \Delta)$ is convex for $\mu \neq 0$.

We next show that

$$v_2(s, \Delta) \quad = \quad \left(\frac{e^{\beta\Delta} - 1}{1 - e^{-\beta s}} \right),$$

is a strictly increasing function. To show that $v_2(s, \Delta)$ is non-negative and increasing, first note that $v_2(s, 0) = 0$ and consider the partial derivative

$$\frac{\partial v_2(s, \Delta)}{\partial \Delta} \quad = \quad \frac{\beta e^{\beta\Delta}}{1 - e^{-\beta s}}$$

When $\mu > 0$, β is positive as well and so both the numerator and the denominator will be strictly positive for $\Delta > 0$. Likewise, when $\mu < 0$, β is negative and so both the numerator and denominator are negative. Thus the derivative will be strictly positive.

Finally, we show that for each fixed $\lambda \geq 0$ and $s > 0$ and $\Delta \geq 0$ that satisfy $u_2(s, \Delta) + \lambda v_2(s, \Delta) = k + \ell$, $w(Q) = u_3(s, \Delta, Q) + \lambda v_3(s, \Delta, Q)$ is a strictly increasing function for $Q \geq \Delta$. Note that

$$\frac{\partial w(Q)}{\partial Q} \quad = \quad u_2(s, Q) + \lambda v_2(s, Q) - (\ell + k).$$

By hypothesis $u_2(s, \Delta) + \lambda v_2(s, \Delta) - (\ell + k) = 0$, and we have shown that $u_2(s, \Delta) + \lambda v_2(s, \Delta)$ is an increasing function. Hence $w(Q)$ is an increasing convex function for $Q \geq \Delta$. $\square$

Lemma A.3.2. *For each $\lambda \geq 0$, define $\{s(\lambda), \Delta(\lambda), Q(\lambda)\}$ to be the unique solution to* (5)–(7)*. For some open neighborhood about $\{s(\lambda), \Delta(\lambda), Q(\lambda)\}$,*

$$\frac{ds(\lambda)}{d\lambda} < 0, \quad \frac{d\Delta(\lambda)}{d\lambda} < 0 \text{ and } \quad \frac{dQ(\lambda)}{d\lambda} < 0.$$

Proof of Lemma A.3.2. Recall that $\{s(\lambda), \Delta(\lambda), Q(\lambda)\}$ is the unique solution to (5)–(7). We show that for each $\lambda \geq 0$

$$\begin{aligned}
\frac{ds(\lambda)}{d\lambda} &= \frac{-v_1(s(\lambda))}{\frac{du_1(s(\lambda))}{ds(\lambda)} + \lambda \frac{dv_1(s(\lambda))}{ds(\lambda)}} < 0, \\
\frac{d\Delta(\lambda)}{d\lambda} &= \frac{-v_2(s(\lambda),\Delta(\lambda))}{\frac{\partial u_2(s(\lambda),\Delta(\lambda))}{\partial \Delta(\lambda)} + \lambda \frac{\partial v_2(s(\lambda),\Delta(\lambda))}{\partial \Delta(\lambda)}} - \frac{\frac{\partial u_2(s(\lambda),\Delta(\lambda))}{\partial s(\lambda)} + \lambda \frac{\partial v_2(s(\lambda),\Delta(\lambda))}{\partial s(\lambda)}}{\frac{\partial u_2(s(\lambda),\Delta(\lambda))}{\partial \Delta(\lambda)} + \lambda \frac{\partial v_2(s(\lambda),\Delta(\lambda))}{\partial \Delta(\lambda)}} \frac{ds(\lambda)}{d\lambda} < 0 \text{ and} \\
\frac{dQ(\lambda)}{d\lambda} &= \frac{-v_3(s(\lambda),\Delta(\lambda),Q(\lambda))}{\frac{\partial u_3(s(\lambda),\Delta(\lambda),Q(\lambda))}{\partial Q(\lambda)} + \lambda \frac{\partial v_3(s(\lambda),\Delta(\lambda),Q(\lambda))}{\partial Q(\lambda)}} - \\
&\quad \frac{\frac{\partial u_3(s(\lambda),\Delta(\lambda),Q(\lambda))}{\partial s(\lambda)} + \lambda \frac{\partial v_3(s(\lambda),\Delta(\lambda),Q(\lambda))}{\partial s(\lambda)}}{\frac{\partial u_3(s(\lambda),\Delta(\lambda),Q(\lambda))}{\partial Q(\lambda)} + \lambda \frac{\partial v_3(s(\lambda),\Delta(\lambda),Q(\lambda))}{\partial Q(\lambda)}} \frac{ds(\lambda)}{d\lambda} - \\
&\quad \frac{\frac{\partial u_3(s(\lambda),\Delta(\lambda),Q(\lambda))}{\partial \Delta(\lambda)} + \lambda \frac{\partial v_3(s(\lambda),\Delta(\lambda),Q(\lambda))}{\partial \Delta(\lambda)}}{\frac{\partial u_3(s(\lambda),\Delta(\lambda),Q(\lambda))}{\partial Q(\lambda)} + \lambda \frac{\partial v_3(s(\lambda),\Delta(\lambda),Q(\lambda))}{\partial Q(\lambda)}} \frac{d\Delta(\lambda)}{d\lambda} < 0
\end{aligned}$$

for some open neighborhood about $\{s(\lambda), \Delta(\lambda), Q(\lambda)\}$.

Note that by Lemma A.3.1 all the terms in the denominator are positive, so it is enough to show that

$$-v_1(s(\lambda)) < 0, \tag{122}$$

$$-v_2(s(\lambda),\Delta(\lambda)) - \left(\frac{\partial u_2(s(\lambda),\Delta(\lambda))}{\partial s(\lambda)} + \lambda \frac{\partial v_2(s(\lambda),\Delta(\lambda))}{\partial s(\lambda)} \right) \frac{ds(\lambda)}{d\lambda} < 0, \tag{123}$$

$$\begin{aligned}
-v_3(s(\lambda),\Delta(\lambda),Q(\lambda)) - \left(\frac{\partial u_3(s(\lambda),\Delta(\lambda),Q(\lambda))}{\partial s(\lambda)} + \lambda \frac{\partial v_3(s(\lambda),\Delta(\lambda),Q(\lambda))}{\partial s(\lambda)} \right) \frac{ds(\lambda)}{d\lambda} - \\
\left(\frac{\partial u_3(s(\lambda),\Delta(\lambda),Q(\lambda))}{\partial \Delta(\lambda)} + \lambda \frac{\partial v_3(s(\lambda),\Delta(\lambda),Q(\lambda))}{\partial \Delta(\lambda)} \right) \frac{d\Delta(\lambda)}{d\lambda} < 0
\end{aligned} \tag{124}$$

The first inequality is a direct result of Lemma A.3.1. To prove the second inequality we first write $\frac{ds(\lambda)}{d\lambda}$ explicitly, and rearrange terms to see that it is equivalent to showing that

$$\begin{aligned}
& v_2(s(\lambda),\Delta(\lambda)) \left(\frac{\partial u_1(s(\lambda),\Delta(\lambda))}{\partial s(\lambda)} + \lambda \frac{\partial v_1(s(\lambda),\Delta(\lambda))}{\partial s(\lambda)} \right) \\
& \quad - v_1(s(\lambda)) \left(\frac{\partial u_2(s(\lambda),\Delta(\lambda))}{\partial s(\lambda)} + \lambda \frac{\partial v_2(s(\lambda),\Delta(\lambda))}{\partial s(\lambda)} \right) \\
&= \frac{\lambda e^{\beta s}(e^{\beta \Delta} - 1)}{e^{\beta s} - 1}
\end{aligned}$$

is positive, which is easy to see. To prove (124) we expand $\frac{ds(\lambda)}{d\lambda}$ and $\frac{d\Delta(\lambda)}{d\lambda}$ and rearrange terms to see that it is equivalent to proving

$$\begin{aligned}
&v_3(s(\lambda), \Delta(\lambda), Q(\lambda)) \left(\frac{\partial u_1(s(\lambda))}{\partial s(\lambda)} + \lambda \frac{\partial v_1(s(\lambda))}{\partial s(\lambda)} \right) \left(\frac{\partial u_2(s(\lambda), \Delta(\lambda))}{\partial \Delta(\lambda)} + \lambda \frac{\partial v_2(s(\lambda), \Delta(\lambda))}{\partial \Delta(\lambda)} \right) - \\
&v_1(s(\lambda)) \left(\frac{\partial u_3(s(\lambda), \Delta(\lambda), Q(\lambda))}{\partial s(\lambda)} + \lambda \frac{\partial v_3(s(\lambda), \Delta(\lambda), Q(\lambda))}{\partial s(\lambda)} \right) \\
&\quad \left(\frac{\partial u_2(s(\lambda), \Delta(\lambda))}{\partial \Delta(\lambda)} + \lambda \frac{\partial v_2(s(\lambda), \Delta(\lambda))}{\partial \Delta(\lambda)} \right) - \\
&v_2(s(\lambda), \Delta(\lambda)) \left(\frac{\partial u_3(s(\lambda), \Delta(\lambda), Q(\lambda))}{\partial s(\lambda)} + \lambda \frac{\partial v_3(s(\lambda), \Delta(\lambda), Q(\lambda))}{\partial \Delta(\lambda)} \right) \\
&\quad \left(\frac{\partial u_1(s(\lambda))}{\partial s(\lambda)} + \lambda \frac{\partial v_1(s(\lambda))}{\partial s(\lambda)} \right) + \\
&v_1(s(\lambda)) \left(\frac{\partial u_3(s(\lambda), \Delta(\lambda), Q(\lambda))}{\partial s(\lambda)} + \lambda \frac{\partial v_3(s(\lambda), \Delta(\lambda), Q(\lambda))}{\partial \Delta(\lambda)} \right) \\
&\quad \left(\frac{\partial u_2(s(\lambda), \Delta(\lambda))}{\partial \Delta(\lambda)} + \lambda \frac{\partial v_2(s(\lambda), \Delta(\lambda))}{\partial s(\lambda)} \right) > 0
\end{aligned} \tag{125}$$

Substituting that $k + \ell = u_2(s, \Delta) + \lambda v_2(s, \Delta)$, (125) simplifies to

$$\frac{e^{\beta s} \lambda \left[e^{\beta Q} - e^{\beta \Delta} - \beta(Q - \Delta) \right] \left[h(1 - e^{\beta s} + \beta s e^{\beta(s+\Delta)}) + \beta \lambda \mu e^{\beta(s+\Delta)} \right]}{\beta \mu (e^{\beta s} - 1)^2}.$$

Note that since $\beta = 2\mu/\sigma^2$ the denominator is always positive. Similarly it is easy to verify that $\left[e^{\beta Q} - e^{\beta \Delta} - \beta(Q - \Delta) \right]$ is strictly positive. Hence, it remains to show that $Y(s, \Delta) = \left[h(1 - e^{\beta s} + \beta s e^{\beta(s+\Delta)}) + \beta \lambda \mu e^{\beta(s+\Delta)} \right]$ is positive. When $s = 0$, $Y(0, \Delta) = \lambda \mu \beta > 0$. Since $\frac{\partial (Y(s, \Delta))}{\partial s} > 0$ it is clear that $Y(s, \Delta) > 0$, hence the desired result is obtained. □

At this point we are ready to prove part (a) of Lemma 3.5.1.

Proof of Lemma 3.5.1. From Lemma A.3.1 it is clear that for each $\lambda \geq 0$ there exists a unique solution $\{s(\lambda), \Delta(\lambda), Q(\lambda)\}$ that satisfies (5)–(7).

Let $S(\lambda) = s(\lambda) + Q(\lambda)$. If $S(0) \leq M$ the uniqueness of $\{s(\lambda), \Delta(\lambda), Q(\lambda)\}$ and λ is immediate. Next we show that when $S(0) > M$, there is a unique $\lambda^* > 0$ such that $S(\lambda^*) = s(\lambda^*) + Q(\lambda^*) = M$.

Lemma A.3.2 shows that $S(\lambda)$ is a decreasing function of λ and so, if there is a value of $\lambda > 0$ for which $S(\lambda) = M$, then it is unique. To show that such a value exists, we show that $s(\lambda)$ and $Q(\lambda)$ go to 0 as λ goes to infinity.

Observe that since

$$\frac{L}{\lambda} \geq \frac{L - u_1(s(\lambda))}{\lambda} = v_1(s(\lambda)) \geq 0,$$

it follows that $v_1(s(\lambda)) \to 0$ and hence $s(\lambda) \to 0$ as $\lambda \to \infty$. Similarly, since

$$\begin{aligned} \frac{k+\ell}{\lambda} &\geq \frac{k+\ell-u_1(s(\lambda),\Delta(\lambda))}{\lambda} = \frac{e^{\beta\Delta}-1}{1-e^{-\beta s}} \geq 0, \\ \frac{(k+\ell)|1-e^{-\beta s}|}{\lambda} &\geq |e^{\beta\Delta}-1| \geq 0 \end{aligned}$$

and so $\Delta(\lambda) \to 0$ as $\lambda \to \infty$.

Finally, we show that for all $\lambda \geq 0$, $u_3(s(\lambda),\Delta(\lambda),Q(\lambda)) \geq -Q(0)(k+\ell)$ from which it follows that

$$\frac{K+(k+\ell)Q(0)}{\lambda} \geq \frac{K-u_3(s(\lambda),\Delta(\lambda),Q(\lambda))}{\lambda} = v_3(s(\lambda),\Delta(\lambda),Q(\lambda))$$

and so $v_3(s(\lambda),\Delta(\lambda),Q(\lambda))$ goes to 0 as λ goes to infinity. It is easy to verify that if $v_3(s(\lambda),\Delta(\lambda),Q(\lambda))$ goes to 0, then $Q(\lambda)$ must go to $\Delta(\lambda)$, which we already showed goes to zero with increasing λ. It remains then only to show that $u_3(s(\lambda),\Delta(\lambda),Q(\lambda)) \geq -Q(0)(k+\ell)$.

Note that for all s, $\Delta > 0$, $u_3(s,\Delta,\Delta) = 0$. Furthermore, $\frac{\partial u_3(s,\Delta,Q)}{\partial Q} = u_2(s,Q)-(k+\ell)$. Since u_2 is non-negative, $\frac{\partial u_3(s,\Delta,Q)}{\partial Q} \geq -(k+\ell)$ and $u_3(s(\lambda),\Delta(\lambda),Q(\lambda)) \geq -(Q(\lambda)-\Delta(\lambda))(k+\ell) \geq -Q(\lambda)(k+\ell)$. Finally, since $Q(\lambda)$ is decreasing in λ, $Q(0) \geq Q(\lambda)$ and so $u_3(s(\lambda),\Delta(\lambda),Q(\lambda)) \geq -Q(0)(k+\ell)$. □

A.4 Proof of Existence of Lagrange Multiplier in Theorem 3.2.2

In this section we prove the following lemma.

Lemma A.4.1. *For each $d > 0$ there is a Lagrange multiplier $0 \leq \lambda^* < L/d$ satisfying* (71) *and* (72).

Proof. To find $\lambda^* \geq 0$ satisfying (71) and (72), we consider two cases:

Case 1 The optimal reduction quantity in the unconstrained problem, $s^* = S^* - Q^* \leq d$ and $\lambda^* = 0$ or

Case 2 $0 \leq \lambda^* < L/d$ and $s_{\lambda^*} = S_{\lambda^*} - Q_{\lambda^*} = d$

We prove that if Case 1 does not apply, i.e., if $s^* = S^* - Q^* > d$, then Case 2 does, i.e., there is $0 \le \lambda^* < L/d$ such that $s_{\lambda^*} = d$.

From (13) we have that when $\mu \neq 0$, s_λ is defined by the solution to

$$z(\lambda) \equiv L - \lambda d = h\left(\frac{s^2(1+e^{\beta s})}{2\mu(e^{\beta s}-1)} - \frac{s}{\beta\mu}\right) \equiv u(s) \tag{126}$$

where $\beta = \frac{2\mu}{\sigma^2}$. Note that $\lim_{s\to 0} u(s) = 0$ and, by Lemma A.3.1, $u(s)$ is a continuous and increasing function of s.

Thus, if $u(d) \ge L$ then $s_0 = s^* \le d$ and Case 1 applies. On the other hand if $u(d) < L$ then, since $z(\lambda)$ is a continuous function of λ with $z(0) = L > u(d)$ and $z(L/d) = 0 < u(d)$, there exists $0 < \lambda^* < L/d$ such that $z(\lambda^*) = u(d)$ and so $s_{\lambda^*} = d$.

When $\mu = 0$, (126) reduces to

$$z(\lambda) \equiv L - \lambda d = \frac{hs^3}{6\sigma^2} \equiv u(s)$$

and similar arguments apply. $\square$

REFERENCES

[1] Bäuerle, N., "Benchmark and mean-variance problems for insurers," *Math. Meth. Oper. Res.*, vol. 62, pp. 159–165, 2005.

[2] Bertsekas, D. P., *Constrained Optimization and Lagrange Multiplier Methods*. Belmont: Athena Scientific, 1996.

[3] Bremner, B. and Kerwin, K., "Here come Chinese cars," *Business Week*, May 26, 2005.

[4] Chen, M., Dubrawski, R., and Meyn, S. P., "Management of demand-driven production systems," *IEEE Transactions On Automatic Control*, vol. 49, pp. 686–698, 2004.

[5] Chow, G. C., *Dynamic Econimics: Optimization by the Lagrange method*. New York: Oxford University Press, 1997.

[6] Clark, A. and Scarf, H., "Optimal policies for a multi-echelon inventory problem," *Management Science*, p. 475, 1960.

[7] Constantinides, G. M., "Stochastic cash management with fixed and proportional transaction costs," *Management Science*, vol. 22, pp. 1320–1331, 1976.

[8] Dentcheva, D. and Römisch, W., "Duality gaps in nonconvex stochastic optimization," *Mathematical Programming, Series A*, vol. 101, pp. 515–535, 2004.

[9] Dyer, G. and Mackintosh, J., "Next for the west are cars "Made in China"," *The Financial Times*, 1 June 2005.

[10] Feinberg, E. A. and Kella, O., "Optimality of D-policies for an $M/G/1$ queue with a removable server," *Queueing Systems: Theory and Applications*, vol. 42, no. 4, pp. 355–376, 2002.

[11] Fisher, M. L., "The Lagrangean relaxation method for solving integer programming problems," *Management Science*, vol. 27, pp. 1–18, 1981.

[12] Harrison, J. M., Sellke, T. M., and Taylor, A. J., "Impulse control of Brownian motion," *Mathematics of Operations Research*, vol. 8, pp. 454–466, 1983.

[13] Harrison, J. M. and Williams, R. J., "Brownian models of open queueing networks with homogeneous customer populations," *Stochastics*, vol. 22, pp. 77–115, 1987.

[14] Harrison, J. M. and Williams, R. J., "Multidimensional reflected Brownian motions having exponential stationary distributions," *Annals of Probability*, vol. 15, pp. 115–137, 1987.

[15] Hordijk, A. and Van Der Duyn Schouten, F. A., "On the optimality of (s, S) policies in continuous review inventory models," *SIAM Journal of Applied Mathematics*, vol. 46, pp. 912–928, 1986.

[16] IBEF. India Brand Equity Foundation, http://www.ibef.in/industry/autocomponents.aspx, (Accessed February 2006).

[17] Jacod, J. and Shiryaev, A. N., *Limit Theorems for stochastic processes.* Berlin: Springer, second ed., 2003.

[18] Krichagina, E. V., Lou, S. X. C., and Taksar, M. I., "Double band policy for stochastic manufacturing systems in heavy traffic," *Mathematics of Operations Research*, vol. 19, pp. 560–595, 1994.

[19] Levans, M. A., "Tight squeeze," *Logistics Management*, vol. 44, p. 45, January 2005.

[20] Luh, P. B. and Feng, W., "From manufacturing scheduling to supply chain coordination: The control of complexity and uncertainty," *The Fourth International Conference onControl and Automation, ICCA 2003.*, 2003.

[21] Luh, P. B., Liu, F., and Moser, B., "Scheduling of design projects with uncertain number of iterations," *European Journal of Operational Research*, vol. 113, pp. 575–592, 1999.

[22] Luh, P. B., Yu, D., Soorapanth, S., Khibnik, A. I., and Rajamani, R., "A Lagrangian relaxation based approach to schedule asset overhaul and repair services," *IEEE Transactions On Automation Science and Engineering*, vol. 2, pp. 145–157, 2005.

[23] Maloni, M. and Jackson, E. C., "North American container port capacity: An exploratory analysis," *Transportation Journal*, p. 1, Summer 2005.

[24] Manufacturing and Services International Trade Administration U. S. Department of Commerce, "The road ahead 2005 for the U. S. auto industry," June 2005. http://ita.doc.gov/td/auto/domestic/RA2005.pdf, (Accessed March 2006).

[25] Märkert, A. and Schultz, R., "On deviation measures in stochastic integer programming," *Operations Research Letters*, vol. 33, pp. 441–449, 2005.

[26] Miemczyk, J. and Holweg, M., "Building cars to order: What does it mean for inbound logistics operations?," *J. Bus. Logistics*, vol. 25, no. 2, pp. 171–197, 2004.

[27] National Chamber Foundation of the U. S. Chamber of Commerce, "Trade and transportation, a study of North American port and intermodal systems," 2003.

[28] National Oceanic & Atmospheric Administration U S Department of Commerce, "Noteworthy records of the 2005 Atlantic hurricane season," http://www.noaanews.noaa.gov/stories2005/s2540b.htm, (Accessed March 2006).

[29] Nowak, M. P. and Römisch, W., "Stochastic Lagrangian relaxation applied to power scheduling in a hydro-thermal system under uncertainty," *Annals of Operations Research*, vol. 100, pp. 251–272, 2000.

[30] Ormeci, M., Dai, J., and Vande Vate, J., "Impulse control of brownian motion: The constrained average cost case," *Operations Research*, forthcoming.

[31] Plambeck, E. L., "Asymptotically optimal control for an assemble-to-order system with capacitated component production and fixed transport costs," tech. rep., Graduate School of Business, Stanford University, 2005.

[32] Richard, S. F., "Optimal impulse control of a diffusion process with both fixed and proportional costs of control," *SIAM Journal of Control and Optimization*, vol. 15, pp. 79–91, 1977.

[33] Robin, M., "On some impulse control problems with long run average cost," *SIAM Journal of Control and Optimization*, vol. 19, pp. 333–358, 1981.

[34] Sen, S., Higle, J. L., and Birge, J. R., "Duality gaps in stochastic integer programming," *Journal of Global Optimization*, vol. 18, pp. 189–194, 2000.

[35] Silver, E. A., Pyke, D. F., and Peterson, R., *Inventory Management and Production Planning and Scheduling.* New York: John Wiley & Sons, 3rd ed., 1998. Section 7.10.

[36] Sulem, A., "A solvable one- dimensional model of a diffusion inventory system," *Mathematics of Operations Research*, vol. 11, pp. 125–133, 1986.

[37] Takriti, S. and Birge, J. R., "Lagrangian solution techniques and bounds for loosely coupled mixed-integer stochastic programs," *Operations Research*, vol. 48, pp. 91–98, 2000.

[38] Taksar, M. I., "Average optimal singular control and a related stopping problem," *Mathematics of Operations Research*, vol. 10, pp. 63–81, 1985.

[39] Womack, J. and Jones, D., *Lean Thinking. Banish waste and create wealth in your corporation.* New York: Simon & Schuster, 1996. p.224.

VITA

Melda Ormeci was born in Istanbul, Turkey. She received her B.S. degree in Industrial Engineering from Bogazici University in Istanbul, Turkey in 1999. Awarded with a Fulbright Fellowship, she joined Georgia Institute of Technology for her M.S. degree in Industrial and Systems Engineering. After receiving her M.S. Degree in 2001, she started to work with Professor John H. Vande Vate and Professor Jim Dai towards her Ph.D. in manufacturing and logistics field at Georgia Institute of Technology in the School of Industrial and Systems Engineering. Her present research interests are in the field of supply chain and logistics.