Infinitesimal Perturbation Analysis for the Capacitated Finite-Horizon Multi-Period Multiproduct Newsvendor Problem

Brigham B. Wilson

A thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Master of Science

Sean Warnick, Chair
Robin Roundy
Jay McCarthy

Department of Computer Science

Brigham Young University

April 2012

ABSTRACT

Infinitesimal Perturbation Analysis for the Capacitated Finite-Horizon Multi-Period Multiproduct Newsvendor Problem

Brigham B. Wilson
Department of Computer Science, BYU
Master of Science

An optimal ordering scheme for the capacitated, finite-horizon, multi-period, multiproduct newsvendor problem was proposed by Shaoxiang [2006] using a hedging point policy. This solution requires the calculation of a central curve that divides the different ordering regions and a vector that defines the target inventory levels. The central curve is a nonlinear curve that determines the optimal order quantities as a function of the initial inventory levels. In this paper we propose a method for calculating this curve and vector using spline functions, infinitesimal perturbation analysis (IPA), and convex optimization. Using IPA the derivatives of the cost with respect to the variables that determine the spline function are efficiently calculated. A convex optimization algorithm is used to optimize the spline function, resulting in a optimal policy. We present the mathematical derivations and simulation results validating this solution.

Keywords: Infinitesimal Perturbation Analysis, Spline Function, Capacitated Newsvendor Problem

ACKNOWLEDGMENTS

I am indebted to Sean Warnick and Robin Roundy for letting me work with them and for all of their support throughout the past years. I owe them a lot. Thanks to my parents for their strong emotional support and encouragement. Thanks to my fellow IDeA Lab students for their friendship - it has been great to be a part of the group. BYU has been a glorious place for me and I should be grateful forever for the opportunity to have been here and do so much.

Contents

List of Figures

Chapter 1

Introduction

The classic inventory management problem is deciding how many products to order in every time period over a specific time horizon in order to minimize costs. The newsvendor problem is a special case where there are no fixed costs, and there is uncertainty in the demand for the products. If at the end of the time horizon all products have not been sold, then they must be salvaged at a cost. There is a tradeoff between having sufficient inventory to meet demand while minimizing the holding and backorder costs as well as the cost of any leftover inventory. As a result of the uncertainty in the demand, any solution involves an understanding of both the current situation and what is likely to happen in the future.

1.1 Literature Review

Inventory management has been studied for decades. For the most part our discussion of the inventory theory literature is limited to multiproduct, stochastic, single-stage systems with a capacity constraint. Optimal ordering policies for one-product systems have been extensively studied (Bertsekas [2005], Glasserman [1994], Tayur [1993]), but multiproduct systems have not received as much attention. Most of the following papers propose simple heuristics with no theoretical performance guarantees.

Evans [1967] was one of the first to analyze the multiproduct inventory management problem. While Evans establishes certain results pertaining to convexity and some key characteristics of the optimal policy, he does not completely characterize the optimal policy for a capacitated model that can start from any possible inventory level. DeCroix and

Arreola-Risa [1998] analyze the problem, but assume that all products are homogeneous. The results by Wein [1992] propose certain qualitative characteristics of the optimal policy obtained by approximating the production of goods as a heavy traffic system. Gershwin [1994] surveys previous work and analyzes inventory management systems with constant demands.

Ha [1997] proves the optimality of the hedging point policy for two identical products. Using sample path comparisons and dynamic programming, de Vericourt et al. [2000] partially characterize the optimal hedging point policy. Pena-Perez and Zipkin [1997] develop a heuristic for a multiproduct, capacitated, stochastic demand system and test their method using numerical methods. Graves [1980] and Gallego [1990] analyze the problem as a cyclic scheduling system. Srivatsan and Dallery [1998] partially characterize the optimal hedging point policy for a two-product system.

Shaoxiang [2006] establishes the optimality of hedging point policies for stochastic two-product systems. This type of policy uses a central curve and a vector of inventory levels to define five regions (see Figure 1.1). The policy is most easily described by imagining the vector of inventory levels increasing continuously from its pre-order level to its post-order level and describing the path that the vector traverses. In Region I, product 2 is not ordered. The inventory level for product 1 increases until either the inventory vector reaches the diagonal line that separates Regions I and V or until all of the capacity has been used. In Region II the inventory vector moves right until reaching the central curve. Then it follows the central curve until the target inventory levels are reached. Movement is halted before that if the available capacity has all been used. In Region III the inventory vector moves up until it reaches the central curve. Then it follows the central curve until the target inventory levels are reached, but it stops before that if the capacity constraint becomes tight. In Region IV the inventory vector moves up until the diagonal line adjacent to Region V has been reached, or the capacity has all been used. In Region V nothing is ordered. Although this policy is shown to be optimal, Shaoxiang does not show how to compute the central curve that separates Regions II and III, nor does he say how to calculate the vector that defines the target inventory levels or the lines adjacent to Region V.

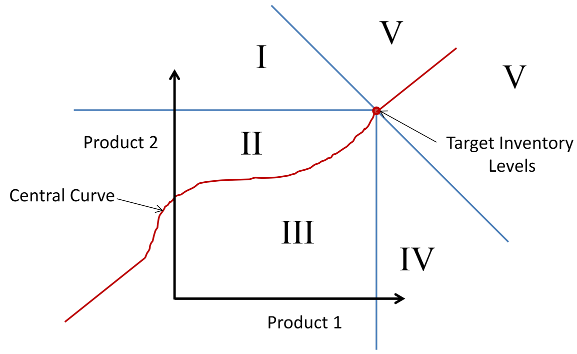

Figure 1.1: Five Regions

Infinitesimal Perturbation Analysis (IPA) is explained in Glasserman [1991] and Ho and Cao [1991]. Our work builds on the single product IPA work done by Kapuscinski and Tayur [2003]. They analyze a capacitated serial system for one product, where the product is transported between multiple stages and demand is a random variable. Although many of the assumptions of our model match theirs, Kapuscinski and Tayur do not account for multiple, heterogeneous products nor a finite time horizon.

The goal of this thesis is to compute the central curve, which determines the optimal ordering amounts given any starting inventory level. We represent the central curve using a spline function. Given an initial estimate of the spline parameters that determine the central curve, we simulate the system and use Infinitesimal Perturbation Analysis (IPA) to efficiently calculate the derivatives of the expected cost with respect to the spline parameters. These parameters are then adjusted by a convex minimization algorithm and the process is repeated, until the optimization is complete. In summary, we calculate an ordering scheme for the capacitated finite-horizon multi-period multiproduct newsvendor problem using spline functions, IPA, and convex optimization.

This paper is organized as follows. In Chapter III, the model assumptions are explained along with the inventory dynamics and model costs. Chapter IV presents our approach, which uses spline functions, IPA and convex optimization. Chapter V discusses the numerical experiment we performed. Chapter VI gives our conclusions. Chapter VII is the appendix, which includes a number of derivations and proofs.

Chapter 2

Model Assumptions

The inventory model we use is the traditional model presented in the literature (Bertsekas [2005]). In each time period the order quantity is calculated, demand occurs, and the inventory levels are updated.

The version of the newsvendor problem that we analyze has a finite time horizon, two different products, and a fixed capacity. In every time period a decision must be made as to how many of each type of product to order. The costs are proportional to the inventory and backorder levels. Demand has a stationary distribution where the demand in each period is independent of demand in every other period. We assume that the demand in each time period has density. Excess demand is backordered; excess inventory is stored. The goal is to manage the inventory system in the near future, say, the next 5-10 time periods.

The inventory dynamics are as follows: $\mathbf{x}_{t+1} = \bar{X}_{t+1}(\mathbf{x}_t) = \mathbf{x}_t + \mathbf{q}_t - \mathbf{d}_t$ for $t = 0, 1, ..., T-1$, where t is the time period, $\mathbf{x}_t \in \mathbb{R}^n$ is a vector of the inventory levels of the products at the beginning of time period t, $\mathbf{q}_t \in \mathbb{R}^n$ is a vector of the order quantities in time period t, and $\mathbf{d}_t \in \mathbb{R}^n$ is the demand vector in time period t. We must choose $\mathbf{q}_t$ before $\mathbf{d}_t$ has been observed. The capacity is $c \in \mathbb{R}^1$. $\mathbf{x}_0$ is the initial inventory level and the first order is in time period $t = 1$. We use $n = 2$. Hence each vector is 2-dimensional and there are 2 different products.

The system is represented graphically in Figure 2.1. The block K is a static controller, a function that receives the initial inventory levels and uses the central curve and the target inventory levels to determine the order quantities of each product. The block P keeps track of the inventory levels of the previous period and receives as input the changes in inventory due to the order quantity and the demand quantity, and increments the time period t by one. The loop that takes the plant output (the end-of-period inventory levels) and becomes input to the controller (the starting point for the next time period) is called state feedback. The block C evaluates the holding or backorder costs due to the end-of-period inventory levels.

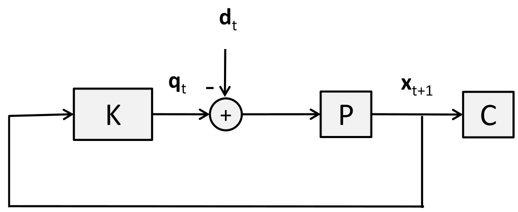

Figure 2.1: System Diagram

Qualitatively, certain features of the solution are consistent with our understanding of this type of problem. For example, with state feedback a static controller should be optimal. Moreover, with a nonlinear cost function the controller should also be nonlinear. This is the case.

The objective is to minimize the expected costs. The inventory must be maintained until the product is shipped. Holding costs are incurred when end-of-period inventory levels are positive. Backorder costs are incurred when demand exceeds supply. (We assume that when there is not inventory arriving customers do not simply go to another supplier, but instead wait until the next time period to receive their product.) Leftover inventory is disposed of at a set salvage value at the end of the horizon. Thus, the total costs are given by

$$\begin{aligned} C(\vec{\omega_t}) &= \sum_{t=1}^{T} C_t(\vec{\omega_t}), \\ C_t(\vec{\omega_t}) &= \mathbf{h}_t^T \max(\mathbf{x}_t, 0) + \mathbf{p}_t^T \max(-\mathbf{x}_t, 0), \text{ for } 1 \le t < T, \text{ and} \\ C_T(\vec{\omega_t}) &= \mathbf{1}_{t=T} \boldsymbol{\sigma}^T \mathbf{x}_t \end{aligned} \tag{2.1}$$

where $C_t(\vec{\omega_t})$ is the period-t cost, $\mathbf{h}_t \in \mathbb{R}^2$ is the holding cost vector in time period t, $\mathbf{p}_t \in \mathbb{R}^2$ is the backorder cost vector in time period t, and $\boldsymbol{\sigma} \in \mathbb{R}^2$ is the salvage value vector. These parameters are given at the beginning of the simulation. The vector $\vec{\omega}_t$ of decision variables is defined in Section 3.1 below.

Chapter 3

Approach

Although Shaoxiang [2006] shows that a hedging point policy is optimal, he does not discuss how one can calculate the central curve or the vector of target inventory levels that determine the different regions. To date no one has succeeded in defining these curves analytically.

We propose a solution with three elements: spline functions, IPA, and convex optimization. We use a spline function to define the central curve. We define a fixed number of knots and then create a function that is composed of low order polynomials between each pair of consecutive knots. The vector of variables $\vec{\omega_t}$ determines the low order polynomials. After guessing $\vec{\omega_t}$ we run a large number of simulations in order to evaluate the expected costs. The simulations use IPA in order to calculate the gradient of the expected costs with respect to $\vec{\omega_t}$. The convex optimization algorithm then adjusts $\vec{\omega_t}$ in order to minimize the costs. The process is repeated until it converges.

The control policy (the central curve and the target inventory levels) differ in each time period. Even if demands and costs are stationary, the proximity of the current time period to the final time period can cause the optimal control policy to be different in different time periods.

Dynamic programming can solve this problem for a single product; however, this method does not scale effectively as the number of products grows. In our simulation we only use two different products, but our methods should scale well as the number of products grows.

The computations required to evaluate the expected cost and its gradient via simulation with IPA grow linearly in n.

There is a shortcoming in our work: according to Shaoxiang [2006] the lines that separate Region V from Regions I and IV in Figure 1.1 are not necessarily straight. The methods that we are using for the central curve can clearly be applied to these curves as well. We approximate them with straight lines, with slopes that make intuitive sense.

3.1 Spline Functions

We use spline functions to determine the central curve that separates regions II and III of Figure 1.1. The spline function is defined by a set of fixed points $\{(a_i, 0) : 1 \leq i \leq m\}$ located along the horizontal axis of Figure 1.1 and a corresponding set of variable lengths θ_{it}. The distance from $(a_i, 0)$ to the central curve, along the line $x_{1t}+x_{2t} = a_i$, is θ_{it} (see Figure 3.1). We define $\omega_{it} = \frac{\theta_{it}}{\sqrt{2}}$. From here on we will use ω_{it}, not θ_{it}. The spline function $S_t(a)$ interpolates the points $(a_1, \omega_{1t}), (a_2, \omega_{2t}), ..., (a_m, \omega_{mt})$. Thus the point $(x_{1t}, x_{2t}) = (a_i - S(a_i), S(a_i))$ is on the central curve. The spline function is a piece-wise cubic polynomial, defined as follows for $a_i \leq a \leq a_{i+1}$: $S_t(a) = E(a - a_i)^3 + F(a - a_i)^2 + G(a - a_i) + H$. We constrain $S_t'(a_i)$ so that the slope of the central curve at each point $(a_i - S_t(a_i), S_t(a_i))$ is equal to the slope of the line segment connecting the two adjoining points, $(a_{i-1} - S_t(a_{i-1}), S_t(a_{i-1}))$ and $(a_{i+1} - S_t(a_{i+1}), S_t(a_{i+1}))$. Because of this ω_{it} has an impact on $S_t(a)$ only for $a \in (a_{i-2}, a_{i+2})$. This property is very useful in IPA because it means that we only need to calculate a limited number of derivatives for each time period. Many commonly used spline functions do not have this property. In addition, this ensures that the spline function will be smooth at a_i. In period t, for each i, given $\omega_{i-1,t}, \omega_{it}, \omega_{i+1,t}$ and $\omega_{i+2,t}$, we determine $S_t(a_i), S_t(a_{i+1}), S_t'(a_i)$ and $S_t'(a_{i+1})$, and we solve for E, F, G and H. In this way we fully define the spline function between a_i and a_{i+1}.

The vector (R_{1t}, R_{2t}) is the vector of target inventory levels. It is determined by the variable a_{Rt} and the spline function. The point $(a_{Rt}, 0)$ is located along the horizontal

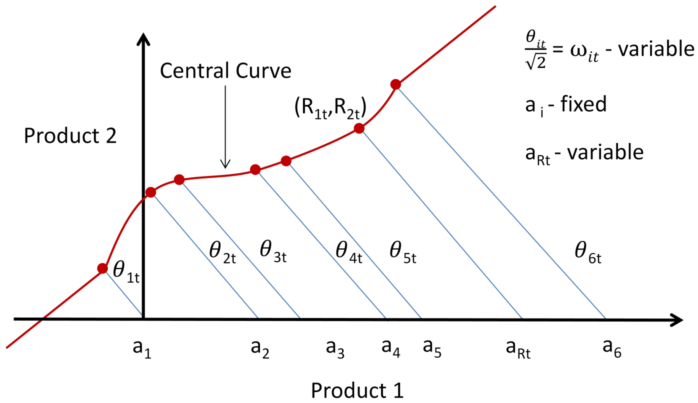

Figure 3.1: Spline Function

axis of Figure 3.1 and (R_{1t}, R_{2t}) is the intersection of the line with a slope of -1 passing through $(a_{Rt}, 0)$, and the central curve. Note that both the central curve and (R_{1t}, R_{1t}) change from period to period. The vector of all variables whose values we optimize is $\vec{\omega_t} = (\omega_{it} : 1 \leq i \leq m, 0 \leq t \leq T; a_{Rt} : 1 \leq t \leq T)$.

As the inventory levels go to $-\infty$ or $+\infty$, the central curve asymptotically approaches a straight line with known slope. We pre-specify the fixed points a_i with the goal of covering the areas where the inventory will most likely be at any given time period. We append straight lines to the far left and right of the central curve. On the far left the slope of the line depends on the cost parameters and is either 0, ∞, or 1. On the far right the slope of the line does not matter because no orders are placed in Region V (see Figure 1.1). We give the central curve a slope of 1 there.

3.2 Simulation with IPA

IPA is combined with our simulation so that a single simulation gives us the gradient as well as the cost for the entire finite time horizon. We simulate the T-period system many times; the number of replications is K. According to IPA, as a simulation is executed we observe and record what the impact of a small perturbation to each of the variables would have been on the costs incurred in the replication. In this way we compute the gradient of the cost with respect to $\vec{\omega_t}$, as a function of the specific random numbers used in that replication. We then average those gradients over the different replications. This converges, as $K \to \infty$, to the expectation of the gradient of the cost. We want to minimize the expected cost, so we need the gradient of the expectation of the cost, not the expectation of the gradient. Expectation and differentiation are not always invertible, but we will prove that they are in this case, i.e. $E\left[C_t'(\vec{\omega_t})\right] \cong E\left[C_t(\vec{\omega_t})\right]'$ (see Section 6.2).

3.3 Optimization

A convex optimization algorithm passes an initial vector $\vec{\omega_t}$ of variables to the simulation. It receives from the simulation estimates of the expected cost and the gradient of the expected cost with respect to $\vec{\omega_t}$. Then the algorithm adjusts $\vec{\omega_t}$ and passes the new vector to the simulation. This continues until the optimization algorithm converges.

The optimization problem has constraints that guarantee that the central curve will be monotone increasing. This optimization problem is a convex minimization with convex inequality constraints,

$$\begin{aligned} \underset{\vec{\omega_t}}{\text{minimize}} \quad & C(\vec{\omega_t}) \\ \text{subject to} \quad & \Psi^*_{it}(\vec{\omega_t}) \leq 0, \text{for } i = 2, \ldots, m-2 \\ & \Phi^*_{it}(\vec{\omega_t}) \leq 0, \text{for } i = 2, \ldots, m-2 \end{aligned}$$

The two constraints guarantee that the central curve is monotone increasing between a_i and a_{i+1}, in other words, that the vector $(a - S_t(a), S_t(a))$ is a nondecreasing function of a for $a \in [a_i, a_{i+1}]$. In the appendix we derive expressions for Ψ^*_{it} and Φ^*_{it}, and show that they are convex functions of $\vec{\omega_t}$ that only depend on $\omega_{it}, k-1 \leq i \leq k+2$.

Chapter 4

Computational Procedure

The simulation and optimization were programmed in MATLAB. For the convex optimization we use the MATALB function $fmincon$, which satisfies our need for a convex minimization with inequality constraints and a previously calculated gradient. The code for determining the order quantity using the spline function and calculating the gradient using IPA are integrated into the simulation.

The gradient is calculated differently depending on which section of the central curve is used to determine the order quantity. In order to obtain more accurate approximations of the derivatives, the inventory position needs to fall between every pair of knots multiple times. This requires locating the a_i to cover the area where the inventory levels will most likely occur. It also helps to use different starting inventories. Consequently, the starting inventories were determined for every replication as random variables from a uniform distribution of each product between $[0, 10]$ or $[-3, 7]$ depending on the case. To mitigate the impact of the randomness in the demand on the optimization we use common random number seeds. In other words, each vector of variables $\vec{\omega_t}$ is tested against the same set of random demands.

4.1 Computational Results

We define a Base Case that has stationary demands, the number of replication $K = 500$, the number of time periods $T = 10$, the holding cost $h_{jt} = 1$, the backorder cost $p_{jt} = 10$, the salvage value $\boldsymbol{\sigma} = .1$, $\sigma_{d_{jt}} = 1$, $\mu_{d_{jt}} = 3$, capacity $c = (3/4) * \sqrt{(\sigma^2_{d_{1t}} + \sigma^2_{d_{2t}})} + \mu_{d_{1t}} + \mu_{d_{2t}}$, where d_{jt} has mean $\mu_{d_{jt}}$ and standard deviation $\sigma_{d_{jt}}$.

The simulation is coded in MATLAB. There are 10 knots located between $a_1 = -4$ and $a_{10} = 14$. The optimal central curve and target inventory levels were reached in between 50-150 optimization steps, depending on the number of replications and starting central curve and target inventory levels.

A dashed red line runs through the origin with slope of 1 simply as a reference. The three green dotted lines have slope -1 and indicate the $0.1, 0.5$, and 0.9 quantiles of W_t. Since the derivative is only calculated near W_t, the central curve is really only accurate between the two extreme green lines. The greatest accuracy is near the middle green line. Outside of this region, a few simulations can cause large shifts in the curve that would not be corrected later simply because the inventory levels never return there. Also, since beyond the point (R_{1t}, R_{2t}) only affects the derivative of the cost with respect to a_{Rt}, the central curve beyond that point is not accurate.

The Base Case is shown in Figure 4.1. We can see that it follows a 45-degree line (as expected) so that the ideal inventory amounts are identical for both product.

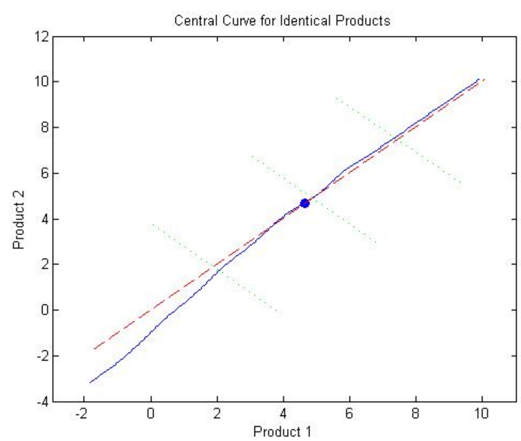

Figure 4.1: Simulation Results for Base Case

When all parameters are equal except that $p_{1t} = 0$, see Figure 4.2. Since there is no punishment for not having enough of product 1, a horizontal central curve makes sure that product 2 reaches its target inventory level before any of product 1 is ordered.

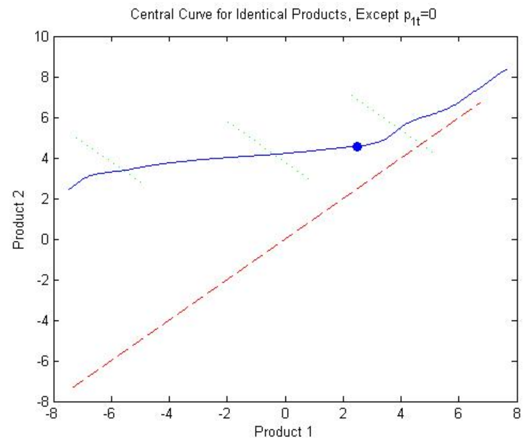

Figure 4.2: Simulation Results for Identical Products, Except $p_{1t} = 0$

When all parameters are equal except that the variance of the demand for product 1 is zero, see Figure 4.3. When $\mathbf{x}_t + \mathbf{q}_t \leq (3, 0)^T$ backorders are guaranteed in the next time period for both products and it does not matter how they are allocated between the products. Therefore any non-decreasing central curve that leads to $(3, 0)^T$ is optimal. Above and to the right of $(3, 0)^T$ the central curve should be vertical. But because very large inventory levels do not occur in the simulation, the optimization algorithm simply ignores points beyond (R_{1t}, R_{2t}) since they have no effect on the cost. Thus the trend of the line beyond that point eventually dies. Since the demand for product 1 is always 3, the vertical line segment at $x_{1t} = 3$ matches what would be expected.

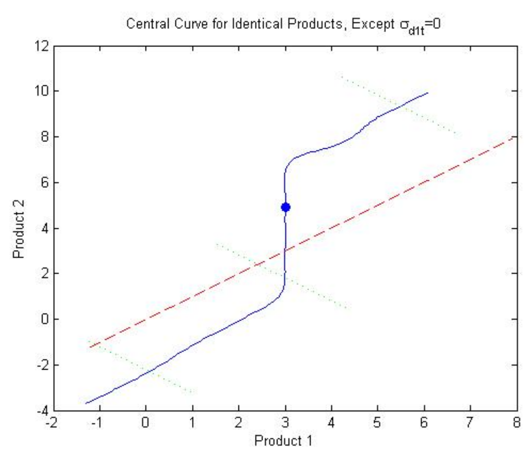

Figure 4.3: Simulation Results for Identical Products, Except $\sigma_{d_{1t}} = 0$

These simulation results of central curves and target inventory levels match the theoretical results that are appropriate for each case. In addition to the abstract tracking done with the code and the derivations in the appendix, these results provide greater confidence in the accuracy of our solution.

Chapter 5

Conclusions

We have presented a method for calculating the optimal ordering policy for the capacitated, finite-horizon, multi-period, multiproduct newsvendor problem proposed by Shaoxiang [2006] using a hedging point policy. We used spline functions, IPA, and convex optimization in order to calculate the central curve and target inventory levels. The simulation results agreeing with their expected cases and the mathematical derivations and proofs validate our solution.

Chapter 6

Appendix

6.1 Derivation of Derivatives

This section explains how we calculated the gradient of the cost with respect to the control variables, i.e. $C_t'(\vec{\omega}_t)$. We will define $C_t(\vec{\omega}_t) = \bar{C}_t(\mathbf{x}_{t+1})$ and $\mathbf{x}_{t+1} = X_{t+1}(\vec{\omega}_t)$. The future evolution of the system after period t is a function of $\mathbf{x}_{t+1}$. The period-t cost is also a function of $\mathbf{x}_{t+1}$ (see Equation 2.1). When we view $C_t(\vec{\omega}_t)$ as a function of $\vec{\omega}_s$ for some $s \leq t$, with all other elements of $\vec{\omega}_t$ temporarily fixed, we have

$$C_t(\vec{\omega}_t) = \bar{C}_t(\mathbf{x}_{t+1}(\mathbf{x}_t(\cdots(\mathbf{x}_{s+1}(\vec{\omega}_s))\cdots))). \tag{6.1}$$

If we use $\bar{X}_{t+1}'(\mathbf{x}_t)$ to refer to the 2×2 matrix whose $j-l$ entry is $\frac{\partial \mathbf{x}_{l,t+1}}{\partial \mathbf{x}_{j,t+1}}$, and define $\bar{C}_t'(\mathbf{x}_{t+1})$ and $X_{s+1}(\vec{\omega}_s)$ in like manner as 1×2 and $2 \times (m+1)$ matrices, then the Chain Rule for vector functions implies that

$$(C_t'(\vec{\omega}_t)) = \bar{C}_t'(\mathbf{x}_{t+1})\bar{X}_{t+1}'(\mathbf{x}_t)\cdots\bar{X}_{t+1}'(\mathbf{x}_t)X_{s+1}'(\vec{\omega}_s) \tag{6.2}$$

The computation of the first term, the derivatives of the period-t cost with respect to the inventory levels at the end of period t, is straightforward. In the remainder of this section we focus on the other terms.

The derivatives that explain the impact of the inventory level of one product at one time period on the inventory level of another product at the next time period (each of the terms of the form $X_{t+1}'(\mathbf{x}_t)$ above), are called the intertemporal derivatives. The intertemporal derivatives alone are not sufficient to solve the inventory management problem since they only relate how the inventory in one time period affects another.

The single-period derivatives are $X'_{s+1}(\vec{\omega}_s)$. The time periods are different because the inventory level is measured at the beginning of the time period but the effect of the control variables is felt and costs are incurred at the end of a time period. Consequently, it is the period-s control variables that have the direct effect on the period-$s+1$ inventory levels and the period-s costs.

The following subsection describes the derivation of the spline parameters, which are necessary for the following subsection that derives the single-period derivatives. Finally, the derivation of the intertemporal derivatives is given.

6.1.1 Derivation of Spline Parameters

We first need to express $S_t(a)$ as a function of $\vec{\omega_t}$. More specifically, we need E, F, G, and H as functions of $\vec{\omega_t}$ which are easy to differentiate. Let $a = x_{1t} + x_{2t} + q_{1t} + q_{2t}$. In period t, for $a_i \leq a \leq a_{i+1}$ we have

$$S_t(a) = E(a - a_i)^3 + F(a - a_i)^2 + G(a - a_i) + H. \tag{6.3}$$

There are five cases to consider.

Case A, $a_i \leq a \leq a_{i+1}$ where $2 \leq i \leq m - 2$: We see that

$$H = S_t(a_i) = \omega_{it} \text{ and } G = S'_t(a_i) = \frac{\omega_{i+1,t} - \omega_{i-1,t}}{a_{i+1} - a_{i-1}} \tag{6.4}$$

(see Section 3.1). For convenience, let $\delta = a_{i+1} - a_i$. Then $\omega_{it} = S_t(a_{i+1}) = E\delta^3 + F\delta^2 + G\delta + H$ and $\frac{\omega_{i+2,t} - \omega_{i,t}}{a_{i+2} - a_i} = S'_t(a_{i+1}) = 3E\delta^2 + 2F\delta + G$. We obtain

$$F = \left[-2\delta G - 3H + 3\omega_{i+1,t} - \delta \frac{(\omega_{i+2,t} - \omega_{it})}{(a_{i+2,t} - a_i)} \right] / \delta^2 \tag{6.5}$$

$$E = \left[\delta G + 2H - 2\omega_{i+1,t} + \delta \frac{(\omega_{i+2,t} - \omega_{it})}{(a_{i+2,t} - a_i)}\right] / \delta^3 \tag{6.6}$$

From Equations 6.3-6.6, for a given a, $a_i \leq a \leq a_{i+1}$, it is straightforward to compute $\frac{\partial S_t(a)}{\partial \omega_{kt}}$ for $k \in \{i-1, i, i+1, i+2\}$. These are the only elements of $\vec{\omega_t}$ that have an impact on $S_t(a)$.

Case B, $a < a_1$: In Equation 6.3, $i = 1$. E, F, G, and H become

$$E = 0, F = 0, G = \left\{\frac{1}{2} \text{ if } \Pi_{1t} = \Pi_{2t}, 1 \text{ if } \Pi_{1t} < \Pi_{2t}, 0 \text{ if } \Pi_{1t} > \Pi_{2t}\right\}, H = \omega_{1t} \tag{6.7}$$

where $\Pi_{jt} = \sum_{\tau=t}^{T-1} -p_{j,\tau}$. Note that since the central curve is $(a - S_t(a), S_t(a))$, if $S_t'(a)$ is equal to 0, $\frac{1}{2}$, or 1 then the slope of the central curve at $(a - S_t(a), S_t(a))$ is, respectively, equal to 0, 1, or ∞.

Case C, $a > a_m$: In Equation 6.3, $i = m$. E, F, G, and H become

$$E = 0, F = 0, G = \frac{1}{2}, H = \omega_{mt}. \tag{6.8}$$

Case D, $a_i \leq a \leq a_{i+1}$ where $i = 1$: In this case, the value of $S_t'(a_i)$ is given, and $S_t'(a_i) \in \{0, \frac{1}{2}, 1\}$. This leads to

$$\begin{aligned} E &= \left[\delta G + 2H - 2\omega_{2t} + \delta \frac{(\omega_{3t} - \omega_{1t})}{(a_{3t} - a_1)}\right] / \delta^3, \\ F &= \left[-2\delta G - 3H + 3\omega_{2t} - \delta \frac{(\omega_{3t} - \omega_{1t})}{(a_{3t} - a_1)}\right] / \delta^2, \\ G &= S_t'(a_1), H = \omega_{1t} \end{aligned} \tag{6.9}$$

Case E, $a_i \leq a \leq a_{i+1}$ where $i = m-1$: In this case, $S_t'(a_{i+1}) = \frac{1}{2}$. This leads to

$$\begin{aligned} E &= \left[\delta G + 2H - 2\omega_{mt} + \delta \frac{1}{2})\right] / \delta^3, \\ F &= \left[-2\delta G - 3H + 3\omega_{mt} - \delta \frac{1}{2})\right] / \delta^2, \\ G &= \frac{\omega_{mt} - \omega_{m-2,t}}{a_m - a_{m-2}}, H = \omega_{m-1,t} \end{aligned} \tag{6.10}$$

6.1.2 Single-Period Derivatives With Respect to ω_{it}

We now establish the derivatives of $\mathbf{x}_{t+1} = (x_{1,t+1}, x_{2,t+1})^T$ with respect to $\omega_{it}, 1 \leq i \leq m$. Note that these derivatives are computed as the simulation is executed. Consequently, in this context we think of the manner in which the system evolves for one specific set of randomly generated demands. These derivatives are all equal to zero if $\mathbf{x}_t + \mathbf{q}_t$ is not on the central curve. If $\mathbf{x}_t + \mathbf{q}_t$ is on the central curve then we assume that $a_i \leq a < a_{i+1}$ where, for convenience, we define

$$a_0 = -\infty, \text{ and } a_{m+1} = \infty$$

and

$$W_t = x_{1t} + x_{2t}. \tag{6.11}$$

Then $x_{2t} + q_{2t} = S_t(W_{t+1})$ and $x_{1t} + q_{1t} = W_{t+1} - S_t(W_{t+1})$. Since $\mathbf{x}_{t+1} = \mathbf{x}_t + \mathbf{q}_t - \mathbf{d}_t$ and the demands $\mathbf{d}_t$ are exogenous, and since $S_t'(a) = 3E(a-a_i)^2 + 2F(a-a_i) + G$, we have

$$\begin{aligned}
\frac{\partial x_{2,t+1}}{\partial \omega_{kt}} &= \frac{\partial}{\partial \omega_{kt}}(x_{2t} + q_{2t}) \\
&= \frac{\partial}{\partial \omega_{kt}}[S_t(W_{t+1})] \\
&= \left.\frac{\partial S_t(a)}{\partial \omega_{kt}}\right|_{a=W_{t+1}} + S_t'(W_{t+1})\frac{\partial W_{t+1}}{\partial \omega_{kt}} \\
&= (W_{t+1} - a_i)^3 \frac{\partial E}{\partial \omega_{kt}} + (W_{t+1} - a_i)^2 \frac{\partial F}{\partial \omega_{kt}} + (W_{t+1} - a_i)\frac{\partial G}{\partial \omega_{kt}} + \frac{\partial H}{\partial \omega_{kt}} \\
&\quad + (3E(W_{t+1} - a_i)^2 + 2F(W_{t+1} - a_i) + G)\frac{\partial W_{t+1}}{\partial \omega_{kt}}.
\end{aligned} \tag{6.12}$$

If all of the capacity is being used, $q_{1t} + q_{2t} = c$ and the central curve in period t has no effect on W_{t+1}. If all of the capacity is not used, then $W_{t+1} = a_{Rt}$ and the central curve still has no effect. Therefore $\frac{\partial W_{t+1}}{\partial \omega_{jt}} = 0$. We are left with

$$\frac{\partial x_{2,t+1}}{\partial \omega_{kt}} = (W_{t+1} - a_i)^3 \frac{\partial E}{\partial \omega_{kt}} + (W_{t+1} - a_i)^2 \frac{\partial F}{\partial \omega_{kt}} + (W_{t+1} - a_i)\frac{\partial G}{\partial \omega_{kt}} + \frac{\partial H}{\partial \omega_{kt}} \tag{6.13}$$

These derivatives are all 0 if either $k < \max(1, i-1)$ or $k > \min(m, i+2)$. Since $W_{t+1} = x_{1t} + x_{2t} + q_{1t} + q_{2t} - d_{1t} - d_{2t}$, and $\frac{\partial W_{t+1}}{\partial \omega_{kt}} = 0$,

$$\frac{\partial x_{1,t+1}}{\partial \omega_{kt}} = -\frac{\partial x_{2,t+1}}{\partial \omega_{kt}} \tag{6.14}$$

Equations 6.3-6.14 and 2.1 can now be used to obtain the partial derivatives of the period-t cost and of $\mathbf{x}_{t+1}$, with respect to ω_{kt}, in a straightforward manner.

6.1.3 Single-Period Derivatives With Respect to a_{Rt}

We now consider the partial derivatives of $x_{1,t+1}$ and $x_{2,t+1}$ with respect to a_{Rt}. This derivative is equal to zero unless $x_{1t} + x_{2t} = a_{Rt}$, i.e unless $\mathbf{x}_{t+1}$ lies on the boundary of Region V in Figure 1.1. Note that $W_{t+1} = a_{Rt}$. There are two cases to consider.

Case A, $\mathbf{x}_{t+1} = (R_{1t}, R_{2t})$: Equation 6.12 applies if we replace ω_{kt} with a_{Rt}. But the derivatives of E, F, G, and H with respect to a_{Rt} are all zero. Since $W_{t+1} = a_{Rt}$, Equation 6.12 becomes

$$\frac{\partial x_{2,t+1}}{\partial a_{Rt}} = 3E(a_{Rt} - a_i)^2 + 2F(a_{Rt} - a_i) + G \tag{6.15}$$

where if $a < a_1$ we use Equation 6.7 and if $a_m < a$ we use Equation 6.8.

Case B, $\mathbf{x}_{t+1}$ on the boundary of Region V but not at (R_{1t}, R_{2t}): The derivative depends on whether the post-order inventory level is on the boundary of Regions I and V or on the boundary of Regions IV and Vl.

$$\begin{aligned} \frac{\partial x_{2,t+1}}{\partial a_{Rt}} &= 1 \text{ if } x_{1,t+1} > R_{1t} \text{ and} \\ &= 0 \text{ if } x_{2,t+1} > R_{2t} \end{aligned} \tag{6.16}$$

In both Case A and Case B,

$$\frac{\partial x_{1,t+1}}{\partial a_{Rt}} = \frac{\partial}{\partial a_{Rt}}(a_{Rt} - S(a_{Rt})) = 1 - \frac{\partial x_{2,t+1}}{\partial a_{Rt}} \tag{6.17}$$

These equations are tedious to derive, but once programmed, they execute quickly.

6.1.4 Intertemporal Derivatives

In the previous section we computed the derivative of the period-t cost and of $\mathbf{x}_{t+1}$ with respect to the period-t parameters ω_{it} and a_{Rt}. In other words, we computed the last term of Equation 6.2. The first term is easy using Equation 2.1. To complete the derivation of $C_t'(\vec{\omega}_s)$ for $s < t$ we need to compute the intertemporal derivatives $\bar{X}_{t+1}'(\mathbf{x}_t)$, which are the derivatives of $\mathbf{x}_{t+1}$ with respect to $\mathbf{x}_t$.

The dependence of $\mathbf{x}_{t+1} = \mathbf{x}_t + \mathbf{q}_t - \mathbf{d}_t$ on $\mathbf{x}_t$ takes the form of a 2×2 matrix called the Jacobian. The Jacobian J_t is defined by $(J_t)_{kl} = \frac{\partial x_{k,t+1}}{\partial x_{lt}}$. Since the dynamics of the system are stationary over time, J_t has the same structure for all t.

Recall that $W_t = x_{1t} + x_{2t}$. Thus $X'(W_t)$ is a 2-element column vector. Often only W_t effects $\mathbf{x}_{t+1}$. When that is the case, by the Chain Rule,

$$J_t = X'(W_t)(\nabla_{\mathbf{x}_t} W_t) = X'(W_t)(1, 1) \tag{6.18}$$

$$\frac{dW_{t+1}}{dW_t} = (\nabla_{\mathbf{x}_{t+1}} W_{t+1})X'(W_t) = (1, 1)X'(W_t), \text{ and} \tag{6.19}$$

$$(\nabla_{\mathbf{x}_t} W_{t+1}) = (\nabla_{\mathbf{x}_{t+1}} W_{t+1})\bar{X}_{t+1}'(\mathbf{x}_t) = (1, 1)J_t \tag{6.20}$$

We proceed by cases (see Figure 6.1)

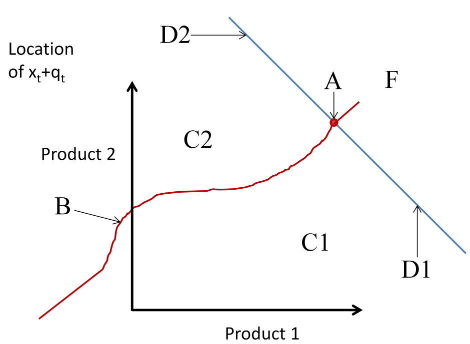

Figure 6.1: Intertemporal Derivative Cases

Case A, $\mathbf{x}_t + \mathbf{q}_t = (R_{1t}, R_{2t})$: In this case a small perturbation in $\mathbf{x}_t$ has no effect on $\mathbf{x}_{t+1} - \mathbf{d}_t = (R_{1t}.R_{2t})$. Therefore $J_t = \vec{0}$, where $\vec{0}$ is a 2×2 matrix of zeros. With respect to derivatives, all dependence of the future on the past is wiped out.

Case B, $\mathbf{x}_t + \mathbf{q}_t \neq (R_{1t}, R_{2t})$; $\mathbf{x}_t + \mathbf{q}_t$ on the central curve: With respect to small perturbations in $\mathbf{x}_t$, only W_t affects $\mathbf{x}_{t+1}$. $\mathbf{x}_{t+1} = \left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right)(W_t + c) + \left(\begin{smallmatrix} -1 \\ 1 \end{smallmatrix}\right) S_t(W_t + c) - \mathbf{d}_t$. Thus $X'(W_t) = \frac{d\mathbf{x}_{t+1}}{W_t} = \left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right) + \left(\begin{smallmatrix} -1 \\ 1 \end{smallmatrix}\right) S'(W_t + c)$. By Equation 6.18,

$$J_t = \left[\left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right) + \left(\begin{smallmatrix} -1 \\ 1 \end{smallmatrix}\right) S'(W_t + c)\right](1, 1) \tag{6.21}$$

Note that by Equation 6.20, $(\nabla_{\mathbf{x}_t} W_{t+1}) = (1, 1)$. By Equation 6.19, $\frac{dW_{t+1}}{dW_t} = 1$.

Case C2, use all capacity; produce only product 1: $\mathbf{x}_{t+1} = \mathbf{x}_t + \left(\begin{smallmatrix} c \\ 0 \end{smallmatrix}\right) - \mathbf{d}_t$ and $J_t = I$. Since $W_{t+1} = W_t + c - d_{1t} - d_{2t}$, $\nabla_{\mathbf{x}_t} W_{t+1} = (1, 1)$ and $\frac{dW_{t+1}}{dW_t} = 1$.

Case C1, use all capacity; produce only product 2: Clearly the conclusions of Case C2 regarding J_t, $\nabla_{\mathbf{x}_t} W_{t+1}$, and $\frac{dW_{t+1}}{dW_t}$ apply.

Case D2, $\mathbf{x}_t + \mathbf{q}_t$ is on the boundary between Regions I and V: By definition, $\mathbf{x}_t + \mathbf{q}_t = \left(\begin{smallmatrix} R_{1t} + R_{2t} - x_{2t} \\ x_{2t} \end{smallmatrix}\right)$. (R_{1t}, R_{2t}) is not impacted by $\mathbf{x}_t$. $\mathbf{x}_{t+1} = \mathbf{x}_t + \mathbf{q}_t - \mathbf{d}_t$ is not impacted by x_{1t}, so $J_{11} = J_{21} = 0$. $J_{*2} = X'(x_{2t}) = \left(\begin{smallmatrix} -1 \\ 1 \end{smallmatrix}\right)$. Thus, $J = \left(\begin{smallmatrix} 0 & -1 \\ 0 & 1 \end{smallmatrix}\right)$. Since $W_{t+1} = a_{Rt}$, $\nabla_{\mathbf{x}_t} W_{t+1} = \vec{0}$ and $\frac{dW_{t+1}}{dW_t} = 0$.

Case D1, $\mathbf{x}_t + \mathbf{q}_t$ on the boundary of Regions IV and V: Clearly, $J = \left(\begin{smallmatrix} 1 & 0 \\ -1 & 0 \end{smallmatrix}\right)$, $\nabla_{\mathbf{x}_t} W_{t+1} = \vec{0}$ and $\frac{dW_{t+1}}{dW_t} = 0$.

Case F, $\mathbf{x}_t + \mathbf{q}_t$ in Region V: $\mathbf{q}_t = \vec{0}$. $\mathbf{x}_{t+1} = \mathbf{x}_t - \mathbf{d}_t$. $J = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$, $\nabla_{\mathbf{x}_t} W_{t+1} = (1, 1)$ and $\frac{dW_{t+1}}{dW_t} = 1$.

6.2 Proof for IPA Validity

In order for the IPA results to be valid, the following must hold: $E\left[C_t'(\vec{\omega_t})\right] = E\left[C_t(\vec{\omega_t})\right]'$. This proof follows the methodology, and some of the notation, presented in Glasserman [1994].

All of the random demands $\mathbf{d}_t = (d_{jt} : 1 \leq j \leq 2, 1 \leq t \leq T)$ are defined on a probability space $(\omega_t, \mathcal{F}, P)$. Let ξ be any one of the elements of the vector $\vec{\omega}_t$. We view the cost function $C_t(\vec{\omega}_t)$ as an $\mathcal{R}$-valued, measurable random function of ξ. By selecting a specific $\mathbf{d}_t \in \Omega$, we obtain a random sample from the function $C_t(\vec{\omega}_t)$. An event is true almost surely (a.s.) if it is true with probability 1.

For a given $\xi \in \vec{\omega}_t$ we define

$$\mathcal{D}_\xi = \{f(\xi) : f(\xi) \text{ is differentiable at } \xi \text{ a.s.}\}$$

Lemma 1. *$C_t(\vec{\omega}_t)$, viewed as a function of ξ, belongs to $\mathcal{D}_\xi$ for all ξ, i.e. $\frac{\partial C_t(\vec{\omega}_t)}{\partial \xi}$ exists a.s for each fixed value of ξ.*

Proof. As defined in Equation 6.1, $C_t(\vec{\omega}_t)$ is a function of $\bar{C}_t(\mathbf{x}_{t+1})$, $\bar{X}_{t+1}(\mathbf{x}_t)$, and $X(\vec{\omega}_t)$. It can be shown that for any fixed value of ξ, every one of the functions given is differentiable in ξ almost surely. $\bar{C}_t(\mathbf{x}_{t+1})$ is differentiable everywhere except when $\mathbf{d}_t = \mathbf{x}_t$. Since each $\mathbf{d}_t$ has density, $P(\mathbf{d}_t = \mathbf{x}_t) = 0$. Therefore $\bar{C}_t(\mathbf{x}_{t+1})$ is differentiable almost surely. $\bar{X}_{t+1}(\mathbf{x}_t)$ and $X(\vec{\omega}_t)$ are differentiable everywhere except when the inventory levels reach Cases A, B, D1, and D2 (see Figure 6.1) and $q_{1t} + q_{2t} = c$. Since demand is stochastic, this is a probability zero event. Therefore $\bar{X}_{t+1}(\mathbf{x}_t)$ and $X(\vec{\omega}_t)$ are differentiable almost surely. Since the component functions of $C_t(\vec{\omega}_t)$ are differentiable almost surely, $C_t'(\vec{\omega}_t)$ exists with probability one. $\square$

We now define

$$\mathcal{D} = \bigcap_\xi \mathcal{D}_\xi = \{f(\xi) : \forall \xi \in \mathbb{R}, f(\xi) \text{ is differentiable at } \xi \text{ a.s.}\}$$

A $\mathbb{R} \to \mathbb{R}$ function $f(\xi)$ is Lipschitz if $|f(\xi_2) - f(\xi_1)| \leq k_f |\xi_2 - \xi_1|$. The constant k_f cannot depend on ξ_1 or ξ_2, but it can be a function of the random demands. Define

$$\text{Lip} = \{f(\xi) : f(\xi) \text{ is a Lipschitz function of } \xi \text{ a.s.}\}.$$

A subclass can be defined as

$$\text{Lip}^1 = \{f(\xi) \in \text{Lip} : \text{a modulus } k_f \text{ of } f(\xi)\text{has a finite mean}\}$$

Lemma 2. *Our cost function $C_t(\vec{\omega_t})$, viewed as a function of ξ, belongs to Lip^1.*

Proof. Because $C_t(\vec{\omega_t})$ is a continuous function of ξ and is differentiable almost everywhere,

$$C_t(\xi_2) - C_t(\xi_1) = \int_{\xi_1}^{\xi_2} C_t'(\vec{\omega_t}) d\xi$$

It suffices to show that $|C_t'(\vec{\omega_t})| < k < \infty$, i.e. that if ξ is perturbed by ϵ then $C_t(\vec{\omega_t})$ will only change by $\kappa\epsilon$. From Equation 6.2 we obtain $C_t'(\vec{\omega_t})$ by selecting the appropriate column from $\nabla_{\vec{\omega}_s}(C_t(\vec{\omega}))$ on the left and from $X_{s+1}'(\vec{\omega}_s)$ on the right. Therefore it suffices to show that there is a finite upper bound on every entry of every matrix in Equation 6.2.

By Equation 2.1 the entries of $\bar{C}_t'(\mathbf{x}_{t+1})$ are bounded by the maximum of the elements of the vectors $\mathbf{h}_t$, $\mathbf{p}_t$, and $\boldsymbol{\sigma}$, which is finite. By Section 6.1.4, $\bar{X}_{u+1}'(\mathbf{x}_u)$ is equal to either $[\left(\begin{smallmatrix}1\\0\end{smallmatrix}\right) + \left(\begin{smallmatrix}-1\\1\end{smallmatrix}\right) S'(W_t + c)]\,(1,1)$, $\vec{0}$, I, $\left(\begin{smallmatrix}0 & -1\\0 & 1\end{smallmatrix}\right)$, or $\left(\begin{smallmatrix}1 & 0\\-1 & 0\end{smallmatrix}\right)$. By Section 6.1.1 $S'(a)$ is continuous on $[a_1, a_m]$ and hence bounded there. By Equations 6.7 and 6.8, if $a \notin [a_1, a_m]$ then $S'(a) \in \{0, \frac{1}{2}, 1\}$. Hence the entries of $X_{u+1}'(\mathbf{x}_u)$ are bounded. Since each component is bounded, $C_t'(\vec{\omega_t})$ is bounded.

□

Theorem 1. *For each $\xi \in \vec{\omega}_t$, $E\left[C_t'(\vec{\omega}_t)\right] = E\left[C_t(\vec{\omega}_t)\right]'$*

Proof. For each $\xi \in \vec{\omega}_t$, $C_t(\vec{\omega}_t) \in \text{Lip}^1 \cap \mathcal{D}$ by Lemmas 1 and 2. Lemma 6.3.1 in Glasserman [1994] completes the proof. □

6.3 Derivation of Constraint Functions

We need the central curve to be monotonically non-decreasing. In this section we construct the constraint functions $\Phi_{it}^*(\vec{\omega_t})$ and $\Psi_{it}^*(\vec{\omega_t})$ that guarantee this. Since the central curve is $\{(a - S_t(a), S_t(a)) : -\infty < a < \infty\}$, the central curve is non-decreasing if both $S_t(a)$ and $a - S_t(a)$ are non decreasing. Thus we seek constraints on $\vec{\omega}_t$ that are equivalent to $0 \leq S_t'(a) \leq 1$ for all a.

We begin by calculating $S_t'(a)$. For $a \in [a_i, a_{i+1}]$, $S_t(a) = E(a - a_i)^3 + F(a - a_i)^2 + G(a - a_i) + H$. By design, $S_t'(a_i) = \frac{\omega_{i+1} - \omega_{i-1}}{a_{i+1} - a_{i-1}} = m_1$ and similarly $S_t'(a_{i+1}) = \frac{\omega_{i+2} - \omega_i}{a_{i+2} - a_i} = m_2$.

Next we normalize the parameters of the spline function so that the spline function is defined over a domain of $[0, 1]$ and has range $[0, 1]$. This makes some of the algebra easier later on. Define

$$\begin{aligned}
&\Delta = a_{i+1} - a_i, && \bar{\omega}_t = \omega_{i+1} - \omega_i > 0, \\
&\delta = (a - a_i)/\Delta, && \alpha = E\Delta^3/\bar{\omega}_t, \\
&\beta = F\Delta^2/\bar{\omega}_t, && \gamma = G\Delta/\bar{\omega}_t, \\
&\bar{\Delta} = \Delta/\bar{\omega}_t, && M_1 = m_1\bar{\Delta}, \text{ and} \\
&M_2 = m_2\bar{\Delta}.
\end{aligned} \tag{6.22}$$

Then $\frac{1}{\bar{\omega}_t} S_t(a) = g(\delta) = \alpha\delta^3 + \beta\delta^2 + \gamma\delta + H/\bar{\omega}_t$, $0 \leq \delta \leq 1$. Also, $g(0) = 0$, $g(1) = 1$, $g'(0) = M_1$ and $g'(1) = M_2$. Using these equalities we can express $g(\delta)$ and $g'(\delta)$ as follows: $g(\delta) = (M_1 + M_2 - 2)\delta^3 + (3 - 2M_1 - M_2)\delta^2 + M_1\delta + \omega_i/\bar{\omega}_t$ and $g'(\delta) = 3(M_1 + M_2 - 2)\delta^2 + 2(3 - 2M_1 - M_2)\delta + M_1$.

The central curve is non-decreasing if $0 \leq g'(\delta) \leq \bar{\Delta}$ for $0 \leq \delta \leq 1$. The constraint $\Phi_{it}^*(\vec{\omega}_t) \leq 0$ is equivalent to $g'(\delta) \geq 0$ for $0 \leq \delta \leq 1$, and $\Psi_{it}^*(\vec{\omega}_t) \leq 0$ matches $g'(\delta) \leq \bar{\Delta}$. We begin with $\Phi_{it}^*(\vec{\omega}_t) \leq 0$ by seeking a constraint on M_1 and M_2 that is equivalent to $0 \leq g'(\delta) = 3(M_1 + M_2 - 2)\delta^2 + 2(3 - 2M_1 - M_2)\delta + M_1$, for $0 \leq \delta \leq 1$. Substituting in $\delta = 0$ and $\delta = 1$ we see that $M_1 \geq 0$ and $M_2 \geq 0$. We know that $0 \leq g'(\delta)$ on $[0, 1]$ if $g'(\delta)$ is concave, i.e. if $M_1 + M_2 \leq 2$. So we assume that $g'(\delta)$ is convex, i.e. $M_1 + M_2 > 2$. Because $0 \leq g'(0)$ and $0 \leq g'(1)$, the only way $g'(\delta)$ can fail to be non-negative on $[0, 1]$ is if the minimizer of $g'(\delta)$ lies between 0 and 1, and the minimum value is negative.

Let $M_n = \min(M_1, M_2)$ and $M_x = \max(M_1, M_2)$. Algebraic manipulations show that the minimizer lies between 0 and 1 if $3 \leq 2M_n + M_x$, and that the minimum value of $g'(\delta)$ is negative if $(M_1-2)^2+(M_1-2)(M_2-2)+(M_2-2)^2 > 6$. We choose to combine these results in the following manner: The set of vectors (M_1, M_2) such that $0 \leq g'(\delta)$ for $0 \leq \delta \leq 1$, is $\{(M_1, M_2) : 0 \leq M_1,\ 0 \leq M_2$, and either $6 \geq 2M_x + M_n$ or $(M_1-2)^2+(M_1-2)(M_2-2)+(M_2-2)^2 \leq 3\}$ (see Figure 6.2). We express this condition as a convex constraint by defining $\Psi(M_1, M_2) = \{(M_1-2)^2+(M_1-2)(M_2-2)+(M_2-2)^2$, if $6 \leq 2M_x + M_n$; $\frac{3}{4}(2-M_n)^2$, otherwise$\}$. Thus $g'(\delta) \geq 0$ for $0 \leq \delta \leq 1$ if and only if $\Psi(M_1, M_2) \leq 3$. Noticing that the points $(M_1, M_2) = (3,0), (0,3)$, and $(2,2)$ are all on the boundary between the cases, the set $\{(M_1, M_2) : \Psi(M_1, M_2) \leq 3\}$ is easily seen to be a convex set.

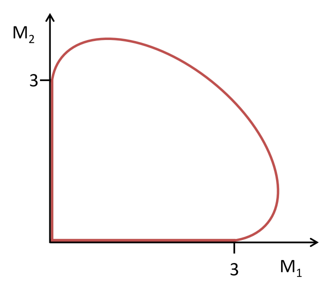

Figure 6.2: Convex Region Defined by $\Psi(M_1, M_2)$

We need to cast this constraint in terms of $\omega_{i-1,t}$, ω_{it}, $\omega_{i+1,t}$, and $\omega_{i+2,t}$. This can be done with the following substitutions:

$$\begin{aligned}
\lambda_0 &= \omega_{it} - \omega_{i-1,t}, & \lambda_1 &= \omega_{i+1,t} - \omega_{it}, \\
\lambda_2 &= \omega_{i+2,t} - \omega_{i+1,t}, & h_0 &= \frac{(a_{i+1} - a_i)}{(a_{i+1} - a_{i-1})}, \\
h_2 &= \frac{(a_{i+1} - a_i)}{(a_{i+2} - a_i)}, & h_x &= \max(h_0, h_2, \\
h_n &= \min(h_0, h_2), & \lambda_x &= \max(\lambda_0, \lambda_2), \text{ and} \\
\lambda_n &= \min(\lambda_0, \lambda_2),
\end{aligned}$$

By Equation 6.22 $M_1 = h_0(\lambda_1+\lambda_0)/\lambda_1$ and $M_2 = h_2(\lambda_1+\lambda_2)/\lambda_1$. Let $\lambda = (\lambda_0, \lambda_1, \lambda_2)^T$. We need the constraint to be convex in $\vec{\omega}_t$. It suffices to make it convex in λ. To this end we restate $\Psi(M_1, M_2) \leq 3$ as $0 \geq \lambda_1 \left[\sqrt{\Psi(M_1, M_2)} - \sqrt{3}\right] = \Psi^*_{it}(\lambda_0, \lambda_1, \lambda_2)$. Thus

$$\Psi^*_{it}(\lambda_0,\lambda_1,\lambda_2) = \begin{cases} \infty & \text{if } \min\{\lambda_0,\lambda_1,\lambda_2\} < 0, \\ \\ ((h_0(\lambda_1+\lambda_0)-2\lambda_1)^2 + & \text{if } 6\lambda_1 \leq 2h_x(\lambda_1+\lambda_x)+ \\ (h_0(\lambda_1+\lambda_0)-2\lambda_1)(h_2(\lambda_1+\lambda_2)-2\lambda_1)+ & h_n(\lambda_1+\lambda_n) \text{ and } \lambda \geq 0 \\ (h_2(\lambda_1+\lambda_2)-2\lambda_1)^2))^{\frac{1}{2}} - \sqrt{3}\lambda_1 & \\ \\ \frac{-\sqrt{3}}{2}(\lambda_1 - h_n(\lambda_1+\lambda_n)) & \text{otherwise} \end{cases} \tag{6.23}$$

The final form of the constraint is $\Psi^*_{it}(\lambda_0, \lambda_1, \lambda_2) \leq 0$. We note that $\Psi^*_{it}(\eta\lambda) = \eta\Psi^*_{it}(\lambda)$ for all $\eta \geq 0$. Also note that $\{(\lambda_0, \lambda_2) : \Psi^*_{it}(\lambda_0, 1, \lambda_2) \leq 0\}$ is convex because it is a linear transformation of the set $\{(M_1, M_2) : \Psi(M_1, M_2) \leq 3\}$, which is a convex set. These facts, together with line 1 of Equation 6.23, imply that Ψ^*_{it} is a convex function. Since λ is a linear function of $\vec{\omega}_t$, we apparently have what we need, an inequality constraint based on a convex function of $\vec{\omega}_t$ that ensures that $0 \leq S'(a)$ for all a. However there is one more thing to check. In deriving Ψ^*_{it} we multiplied Ψ by λ_1^2. To see that we did not introduce spurious solutions so doing note that if $\lambda_1 = 0$ and $\Psi^*_{it}(\lambda_0, \lambda_1, \lambda_2) \leq 0$ then $\lambda_0 = \lambda_2 = 0$. Hence $\omega_{i-1,t} = \omega_{it} = \omega_{i+1,t} = \omega_{i+2,t}$, so by Equation 6.22, $M_1 = M_2 = 0$ and $\Psi(M_1, M_2) = 3$.

We can now obtain the second constraint, $S'_t(a) \leq 1$, by applying the same logic to $a - S_t(a)$ that we applied to $S_t(a)$. The resulting constraint derivation is of the same structure, with identical parameters, except that the λ_i must be translated as follows:

$$
\begin{aligned}
&\lambda_0^n = \delta\left(\frac{1}{h_0} - \lambda_0\right), \quad \lambda_1^n = \delta - \lambda_1, \\
&\lambda_2^n = \delta\left(\frac{1}{h_2} - \lambda_2\right), \quad \lambda_x^n = \max(\lambda_0^n, \lambda_2^n), \text{ and} \\
&\lambda_n^n = \min(\lambda_0^n, \lambda_2^n).
\end{aligned}
$$

The end result is

$$
\Phi_{it}^*(\lambda_0^n, \lambda_1^n, \lambda_2^n) = \begin{cases}
\infty & \text{if } \min\{\lambda_0^n, \lambda_1^n, \lambda_2^n\} < 0, \\
\\
((h_0(\lambda_1^n + \lambda_0^n) - 2\lambda_1^n)^2 + & \text{if } 6\lambda_1^n \leq 2h_x(\lambda_1^n + \lambda_x^n) + \\
(h_0(\lambda_1^n + \lambda_0^n) - 2\lambda_1^n)(h_2(\lambda_1^n + \lambda_2^n) - 2\lambda_1^n) + & h_n(\lambda_1^n + \lambda_n^n) \text{ and } \lambda \geq 0 \\
(h_2(\lambda_1^n + \lambda_2^n) - 2\lambda_1^n)^2))^{\frac{1}{2}} - \sqrt{3}\lambda_1^n & \\
\\
\frac{-\sqrt{3}}{2}(\lambda_1^n - h_n(\lambda_1^n + \lambda_n^n)) & \text{otherwise}
\end{cases} \tag{6.24}
$$

References

R. Anupindi and S. Tayur. Managing stochastic multiproduct systems: Model, measures, and analysis. *Operations Research*, 46:S98–S111, 1998.

D. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 2005.

F. de Vericourt, F. Karaesmen, and Y. Dallery. Dynamic scheduling in a make-to-stock system: A partial characterization of optimal policies. *Operations Research*, 48(5):811–819, 2000.

G. DeCroix and A. Arreola-Risa. Optimal production and inventory policy for multiple products under resource constraints. *Operations Research*, 48:950–961, 1998.

R. Evans. Inventory control of a multiproduct system with a limited production resource. *Naval Research Logistics Quarterly*, 14:173–184, 1967.

A. Federgruen and Z. Katalan. Determining production schedules under base-stock policies in single facility multi-item production systems. *Operations Research*, 46:883–898, 1998.

A. Federgruen and P. Zipkin. An inventory model with limited production capacity and uncertain demands i: The average-cost criterion. *Mathematics of Operations Research*, 11: 193–207, 1986.

G. Gallego. Scheduling the production of several items with random demands in a single facility. *Management Science*, 36:1579–1592, 1990.

S. Gershwin. *Manufacturing Systems Engineering*. Prentice Hall, 1994.

P. Glasserman. *Gradient Estimation via Perturbation Analysis*. Kluwer Academic Publishers, 1991.

P. Glasserman. *Perturbation Analysis of Production Networks*. Springer-Verlag, 1994.

P. Glasserman and S. Tayur. The stability of a capacitated, multi-echelon production-inventory system under a base-stock policy. *Operations Research*, 42:913–925, 1994.

P. Glasserman and S. Tayur. A simple approximation for a multistage capacitated production-inventory system. *Naval Research Logistics*, 43:41–58, 1996.

S. Graves. The multi-product production cycling problem. *AIIE Transactions*, 12:233–240, 1980.

A. Ha. Optimal dynamic scheduling policy for a make-to-stock production system. *Operations Research*, 45(1):42–53, 1997.

Y. Ho and X. Cao. Perturbation analysis and optimization of queueing networks. *Journal of Optimization Theory and Applications*, 40(4):559–582, 1983.

Y. Ho and X. Cao. *Perturbation Analysis of Discrete Event Dynamic Systems*. Kluwer Academic Publishers, 1991.

R. Kapuscinski and S. Tayur. *Optimal Policies and Simulation-based Optimization for Capacitated Inventory Systems*. Kluwer Academic Publishers, 2003.

S. Karlin. Dynamic inventory policy with varying stochastic demands. *Management Science*, 6:231–258, 1960.

R. Parker and R. Kapuscinski. Optimal policies for a capacitated two-echelon inventory system. *Operations Research*, 52:739–755, 2004.

A. Pena-Perez and P. Zipkin. Dynamic scheduling rules for a multiproduct make-to-stock queue. *Operations Research*, 45(6):919–930, 1997.

K. Rosling. Optimal inventory policies for assembly systems under random demands. *Operations Research*, 37:565–579, 1989.

C. Shaoxiang. The optimality of hedging point policies for stochastic two-product flexible manufacturing systems. *INFORMS: Operations Research*, 52(2):312–322, 2006.

N. Srivatsan and Y. Dallery. Partial characterization of optimal hedging point policies in unreliable two-part-type manufacturing systems. *Operations Research*, 46(1):36–45, 1998.

S. Tayur. Computing the optimal policy in capacitated inventory models. *Stochastic Models*, 9:585–598, 1993.

L. Wein. Dynamic scheduling of a multiclass make-to-stock queue. *Operations Research*, 40 (4):724–735, 1992.