Simple Policies for Managing Flexible Capacity

Ganesh Janakiraman[*], University of Texas, Dallas
Mahesh Nagarajan[†], University of British Columbia
Senthil Veeraraghavan[‡], University of Pennsylvania.

September 2010

Abstract

In many production scenarios, a fixed capacity is shared flexibly between multiple products. To manage such multi-product systems, firms need to make two sets of decisions. The first one requires setting an inventory target for each product and the second decision requires dynamically allocating the scarce capacity among the products. It is not known how to make these decisions optimally. In this paper, we propose easily implementable policies that have both theoretical and practical appeal. We first suggest simple and intuitive allocation rules that determine how such scarce capacity is shared. Given such a rule, we calculate the optimal inventory target for each product. We demonstrate analytically that our policies are optimal under certain assumptions in two asymptotic regimes represented by high service levels (i.e. high shortage costs) and heavy traffic (i.e. tight capacity). We also demonstrate numerically that they significantly outperform policies suggested previously over a wide range of problem parameters. In particular, the cost savings from our policies become more significant as the capacity gets more restrictive.

Keywords: Flexible Capacity, Allocation Rules, Multiple Product Problem, Asymptotic Optimality.

1 Introduction

An increasingly common trend in industries where capacity investments are capital-intensive is to invest in flexibility (i.e., capacity that can be utilized to serve demand across several products). Examples include auto-manufacturing, and high tech industries such as semiconductors, consumer electronics and pharmaceuticals. One important reason for this trend is the proliferation in product variety, which causes lower average demand volumes and greater variability in demand for the individual products. As a result, investment in *dedicated capacity* (capacity which can be used only for one product) is no longer economical.

In these industries, a firm's ability to carefully manage flexible capacity is often a significant factor for its long-term success (Limbaugh (2008)). Hence, it is important for firms to manage flexible capacity well. Firms that are able to do so can operate with smaller capacities and still

[*]School of Management, Univ. of Texas, Dallas, TX. ganesh@utdallas.edu
[†]Sauder School of Business, Univ. of British Columbia, Vancouver, Canada. mahesh.nagarajan@sauder.ubc.ca
[‡]The Wharton School, University of Pennsylvania, Philadelphia, PA 19104. senthilv@wharton.upenn.edu

manage to satisfy market demand across several products. The focus of this paper is to provide simple decision rules for managing flexibility efficiently - more specifically, rules for determining how limited capacity can be dynamically allocated across several products.

To achieve our goal, we study a firm that produces multiple products in every period, using a shared resource with limited capacity. We represent the firm's decisions using a periodically reviewed stochastic inventory model. Production occurs at the beginning of each period. A random demand (for each product) occurs during the period. For all products, the unsatisfied demand at the end of any period is backordered. Linear holding and shortage costs are assessed for all products at the end of every period.

We explore the objective of minimizing the long-run average cost per period as the performance measure. This optimization problem comprises of two sets of related decisions in every period. The first one involves setting the target level for each product, and the second requires an allocation rule that determines how the *scarce* capacity is shared among the products. It is well known that performing these two tasks optimally is difficult (more details on this difficulty in the next section). Therefore, in this work, we propose implementable heuristic policies that have both theoretical and practical appeal.

Given the mathematical difficulty of analyzing our problem, we take an approach similar in spirit to papers that study limiting regimes of such stochastic control problems. Limiting regimes yield insights on the structure of optimal policies. Policies constructed using this structural insight are then empirically shown to perform well in non-limiting regimes. In the same spirit, we first suggest an intuitive *class* of allocation rules called *weighted balancing rules*. These rules are parametrized by a *weight* for each product, and they determine how the scarce capacity in any period is shared amongst multiple products. For every rule in this class, the optimal target level for each product is obtained directly from an application of the newsvendor formula – we refer to the combination of weighted balancing rules with these target levels as *weighted balancing policies*. To provide theoretical validity to this class of policies, we study two different asymptotic regimes represented by (*i.*) service levels approaching one (i.e., when shortages are prohibitively expensive), and (*ii.*) utilization approaching one (i.e., when there is little slack between capacity and expected aggregate demand). For each of these two regimes, we identify an allocation rule (i.e. a vector of weights) within our class which is asymptotically optimal under certain assumptions. To investigate how our class of policies performs, in general, we study a set of problems that span a wide range of costs, demand variabilities and capacity utilizations; the majority of these problems do not "belong" to the asymptotic regimes which our analytical results required. For each of these problems, we perform a quick search to identify weights which lead to a policy with good cost performance within our class; then, we compare the cost of this policy with the costs of existing heuristics and with the optimal cost *over all feasible policies* (or a lower bound on the optimal cost). Our computational investigation reveals that our policy outperforms existing policies even at moderate service levels and moderate capacity utilizations.

The organization of the rest of the paper is as follows: We briefly describe the logic of our approach and the construction of our policies in Section 2. In Section 3, we explain our relative contribution to the literature. We present the mathematical formulation of the model and the policies we propose in Section 4. In Section 5, we state a result from Shaoxiang (2004) on the structure of the optimal policy for the two-product problem in the finite and infinite horizon discounted cost case. We also show that our weighted balancing policies have these structural properties. We follow this with some preliminary results on these policies in Section 6. In Section 7, we show, under certain assumptions, that our proposed policies are asymptotically optimal along a sequence of problems in which the backorder cost parameter increases to ∞ (i.e. as the service level approaches 1). In Section 8, we provide similar asymptotic results in heavy traffic (i.e., as the utilization approaches one). In Section 9, we construct a theoretical lower bound on the optimal cost and use them in computational experiments that verify the performance of our policies relative to the optimal policy. Finally, we conclude in Section 10. All relevant proofs are relegated to the appendix.

2 Our Approach

In this paper, we focus on a class of policies called *stationary base-stock policies*, which we define below. There is a target or base-stock level corresponding to each product. This target is constant, i.e., stationary across periods. At the beginning of a period, let us assume that the inventory level of each product will be at most equal to that product's base-stock level. (It will be easy to see that our definition of a stationary base-stock policy is such that this assumption is satisfied in every period if it is satisfied in the first period). The difference between the inventory level and the base-stock level is called the "opening shortfall". If the aggregate shortfall of all products is smaller than the capacity limit, we produce enough of each product to raise its inventory to its base-stock level. The resulting shortfall ("ending shortfall") is thus zero for every product. If the aggregate shortfall exceeds the capacity, then the entire production capacity is used in such a way that the inventory level of each product does not exceed its base-stock level. This concludes our definition of a stationary base-stock policy. It is known from Evans (1967) and Shaoxiang (2004) that this class of policies contains an optimal policy for the two-product case (albeit, under the finite horizon and infinite horizon discounted cost criteria).

We note that, in our discussion of stationary base-stock policies, we have not described how the capacity is allocated to the different products in any period in which capacity is insufficient (scarce) for all products to reach their base-stock levels. We refer to such a description as an *allocation rule*. Clearly, even if we restrict our attention to the class of stationary base-stock policies, calculating the exact base-stock levels in an optimal policy within this class entails an understanding of the optimal allocation rule in periods when capacity is scarce. Thus, *even* within this class, the optimal policy involves two sets of decisions that affect the choice of each other, namely, base-stock levels (production policy) and an allocation rule. The lack of knowledge of the structure of the optimal allocation rule is thus the main stumbling block. To resolve this difficulty, we suggest and work with a set of policies that *decouple* the base-stock and allocation decisions. We will now explain our approach in detail.

The class of policies we advocate is the following. For any given vector of base-stock levels, we raise all inventory levels to the base-stock levels whenever it is feasible to do so. This is possible in periods in which the aggregate shortfall does not exceed the capacity. In periods in which the capacity is insufficient, all the capacity is used. Only in such periods, the allocation rule becomes relevant. An important aspect of the class of allocation rules proposed by us is that, in any period, the *only* state information these decisions require is the opening shortfall of each product (equivalently, the allocation decisions depend on the inventory levels only through the shortfalls). In other words, for a given vector of opening shortfalls, the allocation decisions remain the same for any choice of base-stock levels. For any such allocation rule, the stationary distribution of the vector of ending shortfalls in a period is independent of the base stock levels. This finding has an important implication – the optimal base-stock level for each product can be computed using the newsvendor formula applied to the convolution of that product's demand and its ending shortfall. (This application of the newsvendor formula in this multi-product context has appeared earlier in continuous review models - see, for example, Rubio and Wein (1996).)

Our approach is the following: We restrict attention to simple choices of allocation rules within the aforementioned class, and we choose the base-stock vector corresponding to any particular allocation rule optimally. In the following paragraph, we describe the allocation rules which we propose in detail. We subsequently explain the benefits of our approach.

We use a family of allocation rules which we refer to as *weighted balancing rules*. These rules work as follows. Each product is assigned a strictly positive weight which is constant through time. Next, at the beginning of each period, we rank order the products based on their weighted shortfalls (i.e. the shortfall divided by the weight). We then take the highest ranked product (i.e., the one with the largest weighted shortfall), and use the capacity to bring its weighted shortfall to be equal to the weighted shortfall of the second highest product. Next, we allocate capacity to both these products simultaneously until their weighted shortfalls coincide with the third highest product. We continue this procedure with subsequent products until the entire capacity is exhausted. As mentioned earlier, for any vector of weights, the base-stock level for each product is chosen optimally. This completes the description of a *weighted balancing policy*, given the vector of weights.

We now discuss the issue of choosing the weight vector. One special choice is that all weights are equal to 1 - we call the resulting allocation rule as the *symmetric rule*. At the other extreme are choices of the following type: There is some permutation $\{(1), (2), \ldots, (N)\}$ of the N products such that the weight for (1) << the weight for (2) << ... << the weight for (N) (here, we use << to mean "much smaller than"). Intuitively, such a choice mimics the *priority rule*, i.e. the rule which devotes all the available capacity to (1) until its shortfall is zero and then devotes all the remaining capacity to (2) until its shortfall is zero and so on. Later, we will prove that this priority rule, can be approximated by a suitable weighted balancing policy, for every beginning shortfall. We will show under certain assumptions that the symmetric rule is asymptotically optimal in high service level regimes while the priority rule is asymptotically optimal in heavy traffic. But, in general, the heuristic we propose searches over the space (more precisely, a grid) of weight vectors and picks the best vector – we will refer to the policy of using this weight vector along with the corresponding optimal base-stock levels as the *search policy*. Thus, for every problem instance, the search policy is at least as good as the two asymptotically optimal rules mentioned above; therefore, this policy also has the desired optimality property in both the asymptotic regimes.

We conclude this section by summarizing the benefits of the class of weighted balancing policies.

1. In the single product case, this policy (when there is only one product, there is only one policy in this class) is optimal. This is a trivial consequence of the way the base-stock level is chosen in this policy.

2. In high service level regimes, the policy with symmetric weights is asymptotically optimal under certain assumptions.

3. In heavy traffic (i.e. when utilization approaches one), the policy with weights chosen to mimic a priority policy is asymptotically optimal under certain assumptions.

4. When all products are symmetric (i.e. they have the same cost parameters and the same demand distributions), the policy with symmetric weights is optimal. (We will show this result formally in Section 6.)

This concludes our introductory discussion of the policies we propose in this paper. Next, we discuss our contribution to the literature.

3 Related Literature

We note that the single product capacitated inventory problem is a special case of our problem. This special case has been analyzed in detail. It is well known that a modified base stock policy is optimal for the single product problem as noted in Federgruen and Zipkin (1986). In short, the optimal policy is to produce until the inventory level reaches a threshold (base-stock) if there is enough capacity available; if capacity is not available, the policy is to produce until the entire capacity is exhausted. Furthermore, if the inventory level is above the threshold, no production is done. Clearly, there is no allocation decision involved in this single product problem. While Federgruen and Zipkin establish the optimality of a policy of this type, Tayur (1992) derives the steady state distribution of the shortfall process, and shows that the optimal base stock level can be obtained using the newsvendor formula applied to the convolution of the demand distribution and the shortfall distribution. In this paper, we use the aforementioned result in our model with multiple products and limited capacity.

As mentioned earlier, not much is known about the multi-product problem. There are two sets of difficulties, computational and theoretical, which we will outline below.

From a computational perspective, a dynamic programming approach to solve this problem becomes intractable due to the curse of dimensionality. Furthermore, we believe that policies which are easy to describe and understand dominate complex but theoretically optimal policies from the perspective of implementation. Thus, providing simple and cost-effective heuristics which scale well to problems with many products is valuable - we will see that the policies we propose have these desirable attributes.

The theoretical difficulty is as follows. In the finite horizon dynamic program for the single product problem, the cost-to-go function is convex, and that guarantees the optimality of base-stock policies. The cost-to-go function can be shown to be convex even for the multi-product problem; however, this only guarantees the existence of a minimizer (interpreted as the vector of optimal *after-order* inventory levels) but it does not guarantee the optimality of *base-stock policies*.[1] (We will show in Theorem 6 that base-stock policies are optimal for the special case in which all products have identical demand distributions and costs.)

Moreover, even temptingly simple and intuitive statements do not follow from convexity. For instance, one would imagine that the optimal policy possesses the following property which base-stock policies satisfy: If the inventory levels of all products at the beginning of a period are smaller than their optimal after-order inventory levels, then the inventory level of every product after ordering will be no larger than the optimal after-order inventory level. The careful reader will note that this property does not follow from convexity. In fact, a satisfactory description of the optimal policy has so far been provided only for the two product case (that too, only for the finite or infinite horizon discounted cost problems, not the average cost problem) by Shaoxiang (2004) who expands on the early work by Evans (1967). For this case, Shaoxiang shows that the optimal policy is a base-stock policy. Moreover, he shows that there is a "monotone switching curve" in the space of

[1] By a base-stock policy, we refer to a policy in which there is a base-stock vector every period such that (a) if any product's inventory exceeds its base-stock level, that product is not ordered, (b) when all products' inventory levels are below the respective base-stock levels, then they are all raised to these levels if capacity is sufficient; otherwise, the entire capacity is used and no product's inventory level is allowed to exceed its base-stock level.

inventory levels such that when the starting inventory levels lie on this curve, capacity is used for both products in such a way that the inventory levels remain on this curve; otherwise, capacity is initially used exclusively on one of the two products such that the inventory levels approach this switching curve and once that is done, capacity is used for both products to raise the inventories along the curve. In the two product case, our weighted balancing rules can be viewed as linear approximations of Shaoxiang's switching curves.

While the papers by Shaoxiang and Evans, study the structure of the optimal policies for the two product case, the papers by DeCroix and Arreola-Risa (1998) and Aviv and Federgruen (2001) propose heuristic policies for our problem. We discuss these two papers in the next two paragraphs.

DeCroix and Arreola-Risa (1998) study the periodic review version of the multi-product problem under both the finite horizon and the infinite horizon discounted cost criteria. They prove the optimality (for the finite horizon) of base-stock policies for the special case, where all products are identical both in terms of cost parameters and demand distributions, and when the inventory level of each product in the first period is below its target level. For the general case, they propose the following heuristic: Find the single period newsvendor level for each product. Then, consider any multiple of this vector of newsvendor levels as a potential base-stock vector; perform a one-dimensional search for the best multiple (at each step of this search, the cost of the current policy is evaluated by simulation). It is easy to see that this policy is optimal for the single product problem and for the multi-product problem with symmetric products. However, there are no results on the asymptotic performance of these policies.

Aviv and Federgruen (2001) study a multi-product inventory system in which production occurs in two stages in every period. In the first stage, a common product ("blank") is produced, and, in the second stage, blanks are converted into finished products. There is a production capacity limit in the first stage and lead times could be present in both stages. In the main part of their study, they assume that blanks cannot be held in inventory. They present a heuristic and a lower bound for this problem, and use those results to develop insights on the business practice of *postponement* or *delayed differentiation*. When the lead times at both stages are zero, their problem is the same as ours. Thus, we are interested in studying how their heuristic performs in this special case. In this case, their heuristic works as follows: The problem is first relaxed in such a way that the inventories at the beginning of a period can be freely redistributed among the products without penalty – this relaxed problem has the same structure as a single product, capacitated inventory problem – the optimal base-stock level for this problem is used as the aggregate base-stock level over all products for the original problem – individual base-stock levels and the allocation decisions in every period in the original problem are computed by solving the myopic, i.e. single period, problem. While this policy is optimal for the single product problem, it can be verified that it is not optimal even for the multi-product problem with symmetric products – this is because the base-stock levels used are obtained by solving a relaxed problem. As with DeCroix and Arreola-Risa's heuristic, there are no results on the asymptotic performance of this policy.

While our focus in this paper is on periodically reviewed systems (i.e. discrete time), there are obvious counterparts to our systems in continuous time - we briefly review this literature. Ha (1997) studies a special case with two products and shows several structural properties of the optimal policy – these results are similar to Shaoxiang's results in the discrete time case. The other papers in continuous time (Zheng and Zipkin, 1990; Zipkin, 1995; Veatch and Wein, 1996; Rubio and Wein, 1996; Pena-Perez and Zipkin, 1997) study multi-product inventory systems in the framework of multi-class make-to-stock queues – that is, the entire attention is on the class of base-stock policies and on finding good policies within this class. A common theme of these papers is an answer to the following questions:

⋄ When, i.e. at what states of the system, should the production resource be switched on? – This question is equivalent to finding the base-stock vector.

⋄ At any given state in which the production resource is on, which product should be produced? – In the language of our paper, this question is equivalent to finding the allocation rule.

This body of work uses a combination of heavy traffic analysis and computational tests to motivate and evaluate various choices of base-stock levels and allocation rules. Most works in this literature stream assume *Poisson* demand processes. Among these papers, Pena-Perez and Zipkin (1997) and Veatch and Wein (1996) are closely related to our paper, as explained below.

Pena-Perez and Zipkin (1997) argue that for systems in "heavy traffic", i.e. systems where the aggregate demand rate is close to the production capacity rate, a specific *priority rule* is asymptotically optimal under certain assumptions. Their asymptotic analysis is based on the results of Wein (1992) and uses diffusion approximations. In this paper, we show a similar result in discrete time with two main differences: (a) our notion of asymptotic optimality is *strong* (i.e. the difference between the cost of the priority policy and the optimal cost is bounded, while the optimal cost itself approaches infinity in heavy traffic) whereas their notion is *weak* (i.e. the ratio between the cost of the priority policy and the optimal cost approaches one) and (b) our proof is from first principles and does not rely on diffusion approximations.

Veatch and Wein (1996) propose and evaluate *index rules* - these rules suggest that when production occurs, it should be devoted to the product with the lowest index at that time. Moreover, the index rules they propose are such that the index of a product is independent of the indices of the other products. Our weighted balancing rules are analogous to "linear" index rules.

We conclude this section with the following summary on how our work contributes to the literature on multi-product inventory systems. The continuous time literature has identified multi-class queues as the natural framework for studying multi-product, base-stock systems. The papers in this literature propose different rules for assigning priorities (or indices) dynamically to the products and use the resulting steady state queue length distributions to find the base-stock levels. In this literature, there has also been some interest in studying the heavy traffic asymptotic regime for motivating and validating policies. The discrete time literature has progressed somewhat independently thus far. Our paper bridges this divide – our weighted balancing policies resemble index policies used in the continuous time literature – we also prove asymptotic optimality in heavy traffic. Furthermore, we prove a second asymptotic optimality result in high service levels. Both these results are new to the discrete time literature. Also, our policies are optimal for the two special cases in which the optimal policy is known, namely, the single product case and the symmetric, multi-product case. Finally, we will see in Section 9 that our policies consistently perform better than the existing policies in this literature.

4 Model Description

We index the products by n, $1 \leq n \leq N$. The holding and backorder cost associated with product n in \$/unit/period are h^n and b^n, respectively. Periods are indexed by $t \geq 1$. In period t, the net-inventory, x_t^n (inventory on hand minus backorders) for each product n is observed and the production quantity, q_t^n, for each product is decided. The total production quantity $q_t = \sum_{n=1}^N q_t^n$ is constrained above by a capacity limit κ. Next, the demand, D_t^n for each product n is observed. Finally, the cost C_t incurred for this period is computed based on the inventory levels and backorder levels at the end of the period as follows:

$$C_t \;=\; \sum_{n=1}^{N} \left(h^n \cdot (x_t^n + q_t^n - D_t^n)^+ \;+\; b^n \cdot (D_t^n - x_t^n - q_t^n)^+ \right) \; .$$

The optimization problem that we are interested in solving is that of minimizing the long run average cost when the set of admissible (or feasible) policies is the set of all non-anticipatory policies. A formal definition of this problem follows. A non-anticipatory policy π is described by a set of vector-valued functions $\{\pi_t : t = 1, 2, \ldots, \}$ where $q_t^n = \pi_t^n(\mathbf{x}_t)$; here, $\mathbf{x}_t$ is the state vector $(x_t^1, \ldots, x_t^N)$ in period t and π_t is a function from $\Re^n \to \Re^{n,+}$. Let Π denote the set of all non-anticipatory policies π such that the capacity constraint

$$\sum_{n=1}^{N} \pi_t^n(\mathbf{x}) \leq \kappa \text{ for all } \mathbf{x} \in \Re^n \text{ and for all } t \in \{1, 2, \ldots\}$$

is satisfied. If C_t^π denotes the random cost incurred by the system in period t when the system follows the policy π, our performance measure is

$$C^\pi = \lim_{T\to\infty} \sup E[\sum_{t=1}^{T} C_t^\pi]/T \ .$$

The optimal long run average cost is defined as

$$C^* = \inf_{\pi\in\Pi} C^\pi \ .$$

Throughout the paper, we assume that the sequence of random vectors $\{\mathbf{D}_t\}$ is independent and identically distributed across time periods, where $\mathbf{D}_t = (D_t^1, \ldots, D_t^N)$. Note that we allow for the demands of the products to be correlated. We use D^n to denote a random variable with the same distribution as the single period demand for product n and D to denote a random variable with the same distribution as the aggregate single period demand. Let $\mu^n = E[D^n]$. We also assume that capacity exceeds aggregate expected demand, i.e.,

$$\mu := \sum_{n=1}^{N} \mu^n < \kappa \ ,$$

which is a necessary condition for the existence of a policy with a finite long-run average cost. Moreover, we assume that there is a positive probability that aggregate demand in a period exceeds capacity, i.e.,

$$P\left(\sum_{n=1}^{N} D^n > \kappa\right) > 0 \ .$$

(Without this assumption, the system under study is nothing but a set of N newsvendor problems and is, therefore, trivial.)

Let Π denote the class of all non-anticipatory policies, and let π denote a generic element in this class. Let Π_{BS} denote the subset of stationary base-stock policies described at the beginning of Section 2. We now introduce some notation specific to Π_{BS}.

Let S^n denote the target or base-stock level for product n, and $\mathbf{S}$ denote the vector of base-stock levels. In our analysis of stationary base-stock policies, we assume that $x_1^n \leq S^n$ for all n. Let $W_t^n = S^n - x_t^n$; we refer to W_t^n as the opening shortfall for n in period t. Let V_t^n denote the ending shortfall, i.e. shortfall after ordering. So, $V_t^n = W_t^n - q_t^n$. By definition of a base-stock policy, the following condition holds:

$$\text{if } \sum_{n=1}^{N} W_t^n \leq \kappa \ , \quad \text{then } q_t^n \ = \ W_t^n \quad \text{for all } n \ .$$

That is, all inventory levels are raised to their respective targets, if that is feasible. Otherwise, the entire capacity is used for production without the inventory level of any product exceeding its target, i.e.,

$$\text{if } \sum_{n=1}^{N} W_t^n > \kappa \ , \quad \text{then } \sum_{n=1}^{N} q_t^n \ = \ \kappa \ \text{ and } q_t^n \leq W_t^n \text{ for all } n \ .$$

Notice that the exact manner in which the capacity is allocated among products in such periods has not been completely specified yet. We will specify these allocation rules shortly.

Let Π_{BS-B} denote the set of stationary base-stock policies in which the weighted balancing allocation rule is followed. We will refer to these as weighted balancing policies. A verbal description of these allocation rules was given in Section 2. Clearly, Π_{BS-B} is a subset of Π_{BS}. A mathematical description of a policy in this class follows.

Weighted Balancing Allocation: Rank the products according to the 'weighted' shortfalls $\{W_t^j/\alpha^j\}$, where α^j is the weight corresponding to product j. Let $\boldsymbol{\alpha} = (\alpha^1, \ldots, \alpha^N)$. The symmetric rule chooses $\boldsymbol{\alpha} = \mathbf{1}$, where $\mathbf{1} = (1, 1, \ldots, 1)$. Let $\tilde{n}$ denote the product with the n^{th} largest value of the weighted shortfall, W_t^j/α^j, breaking ties arbitrarily. Allocate production to product $\tilde{n} = 1$ until its weighted shortfall equals that of $\tilde{n} = 2$, or until the capacity is exhausted. Using the remaining capacity, allocate production to products $\tilde{1}$ and $\tilde{2}$ (proportionally based on their weights so that their weighted shortfalls are always equal) until their weighted shortfalls equal that of $\tilde{3}$, or until the capacity is exhausted. This process is continued until the entire capacity available in the period is exhausted. While the description above applies when inventory and production quantities are real-valued, a simple uniform randomization scheme can be used to define the policy when these quantities are integer-valued. Note that any policy $\pi \in \Pi_{BS-B}$ is completely specified by a pair $(\mathbf{S}, \boldsymbol{\alpha})$ where $\mathbf{S}$ is a vector of base-stock levels $\mathbf{S}$ and $\boldsymbol{\alpha}$ is a vector of weights.

Priority Allocation: Let $\{(1), (2), \ldots, (N)\}$ denote any permutation of $\{1, 2, \ldots, N\}$. Then, a priority rule defined by this permutation works as follows: In every period, allocate production to (1) until its shortfall is zero or the entire capacity is consumed; then, proceed to (2) and do the same until all product shortfalls are zero or the capacity is consumed. Intuitively, this rule can be imitated by a weighted balancing rule which assigns the weights α^js in an extremely disparate fashion. We show this formally in Section 6.

5 Optimal Policy Structure

In this section, we present Shaoxiang's results on the optimal policy for the two product case of our problem under the finite horizon criterion and discuss how these results translate to our multi-product case.

Theorem 1 (Shaoxiang (2004)). *Consider the two-product problem under the infinite horizon discounted cost criterion, i.e.* $\min \lim_{T\to\infty} \sum_{t=1}^{T} \eta^t \cdot E[C_t]$ *for some discount factor* $\eta \in (0,1)$. *For every period* t, *there exists a decreasing function* $x^{2,*}(x^1)$ *and an increasing function* $\overline{x}^2(x^1)$ *which meet at a point* (S^1, S^2) *such that the policy defined by (a)-(d) below is an optimal policy:*

(a) If $(x^1, x^2) \geq (S^1, S^2)$, *then* $q^1 = 0$ *and* $q^2 = 0$*; i.e., if inventories exceed the base-stock vector, no order is placed.*

(b) If $(x^1, x^2) \leq (S^1, S^2)$ *and* $(S^1 - x^1) + (S^2 - x^2) \leq \kappa$, *then* $q^1 = S^1 - x^1$ *and* $q^2 = S^2 - x^2$*; i.e., order up to the base-stock vector if it is feasible to do so.*

(c) If $x^1 \geq S^1$ *and* $x^2 \leq S^2$, *then* $q^1 = 0$ *and* $q^2 = \max(0, \min(\kappa, x^{2,*}(x^1) - x^2))$*; i.e., if product* 1*'s inventory exceeds its base-stock level, no order is placed for that product and product 2 orders up to* $x^{2,*}(x^1)$ *if it is feasible to do so. The same statement holds when the superscripts* 1 *and* 2 *are interchanged; then, the function* $x^{1,*}(x^2)$ *denotes the inverse of the function* $x^{2,*}(x^1)$.

(d) If $(x^1, x^2) \leq (S^1, S^2)$ *and* $(S^1 - x^1) + (S^2 - x^2) > \kappa$, *then the following conditions apply:*

(i) If $x^2 = \overline{x}^2(x^1)$, *then* $q^1 + q^2 = \kappa$ *and* $x^1 + q^1 = y^1$ *and* $x^2 + q^2 = y^2$, *where* (y^1, y^2) *solves* $y^2 = \overline{x}^2(y^1)$.

(ii) If $x^2 < \overline{x}^2(x^1)$, *then* $q^1 + q^2 = \kappa$. *Moreover, if* $\overline{x}^2 - x^2 \leq \kappa$, *then* $x^1 + q^1 = y^1$ *and* $x^2 + q^2 = y^2$, *where* (y^1, y^2) *solves* $y^2 = \overline{x}^2(y^1)$. *Otherwise,* $q^2 = \kappa$ *and* $q^1 = 0$. *The same statement holds when the superscripts* 1 *and* 2 *are interchanged.*

Next, we explain how our weighted balancing policies relate to the optimal policy structure described above for the two-product case. It is easy to show that the optimal policy of Theorem 1 is such that once the inventory vector reaches a point which is componentwise smaller than (S^1, S^2), then the inventory vector in every subsequent period is also smaller than (S^1, S^2). Thus, the effective state space (i.e. possible inventory vectors) is $(-\infty, S^1] \times (-\infty, S^2]$; so, it is sufficient to study the optimal policy within this "rectangle". We know that, within this region, the optimal policy is completely described by the function $\overline{x}^2(x^1)$. Our weighted balancing policies work exactly like the optimal policy except that they replace $\overline{x}^2(x^1)$ with the function $(\frac{\alpha_2}{\alpha_1}) \cdot x^1$. In other words, computing the optimal policy involves finding or searching for the function $\overline{x}^2(x^1)$, i.e. the optimal "switching curve" within the space of all increasing functions, whereas, the best weighted balancing policy is found by searching for the best ratio $\frac{\alpha_2}{\alpha_1}$.

6 Weighted Balancing Policies: Preliminaries

Consider a stationary base-stock policy with base-stock levels $S^1, \ldots, S^N$ for the N products. Let $V_t = \sum_{n=1}^{N} (S^n - x_t^n - q_t^n)^+$ denote the aggregate ending shortfall in period t. Let $D_t = \sum_{n=1}^{N} D_t^n$ similarly denote the aggregate demand the system faces in period t.

We begin by making a simple observation about the aggregate shortfall process.

Lemma 2. *Consider any policy in* Π_{BS} *with some base-stock vector* $\mathbf{S}$*. Assume* $x_1^n = S^n$ *for all* n*. The following recursive equation describes the evolution of the aggregate shortfall process,* $\{V_t\}$*.*

$$V_{t+1} \;=\; (V_t + D_t - \kappa)^+ \;.$$

Moreover, (i) the distribution of V_t *is independent of* $\mathbf{S}$ *for all* t*, (ii)* V_t *converges in distribution to a limiting random variable* V_∞ *as* $t \to \infty$ *and (iii) the distribution of* V_∞ *is also independent of* $\mathbf{S}$*.*

We now make a few observations about the vector of individual shortfalls of the products.

Lemma 3. *Consider any policy in* Π_{BS-B} *defined by a base-stock vector* $\mathbf{S}$ *and a weight vector* $\boldsymbol{\alpha}$*. Assume* $x_1^n = S^n$ *for all* n*. Then, (i) the distribution of the vector of ending shortfalls is independent of* $\mathbf{S}$ *for all* t*, (ii) the sequence of distributions of this vector converges to a limiting distribution as* $t \to \infty$ *and (iii) this limiting distribution is also independent of* $\mathbf{S}$*.*

For any policy $\pi \in \Pi_{BS-B}$ defined by the pair $(\mathbf{S}, \boldsymbol{\alpha})$, we use $\mathbf{V}_t^{\boldsymbol{\alpha}}$ to denote the vector of shortfalls in period t. (Note that it is not necessary to include the argument $\mathbf{S}$ in the notation for the shortfall vector since its distribution does not depend on $\mathbf{S}$.) Let $\mathbf{V}^{\boldsymbol{\alpha}}$ denote the limiting distribution of $\mathbf{V}_t^{\boldsymbol{\alpha}}$.

Let $V^{\boldsymbol{\alpha},n}$ denote the limiting random variable representing product n's ending shortfall under the weighted balancing allocation rule with the weight vector $\boldsymbol{\alpha}$. Let $\Phi^{\boldsymbol{\alpha},n}$ denote the distribution of the convolution of $V^{\boldsymbol{\alpha},n}$ and D^n. Let

$$\mathbf{S}^{\boldsymbol{\alpha}*} = (S^{\boldsymbol{\alpha}*,1}, \ldots, S^{\boldsymbol{\alpha}*,N}), \text{ where } S^{\boldsymbol{\alpha}*,n} = (\Phi^{\boldsymbol{\alpha},n})^{-1}\left(\frac{b^n}{b^n + h^n}\right).$$

We will now show that the base-stock vector $\mathbf{S}^{\boldsymbol{\alpha}*}$ is the optimal choice of $\mathbf{S}$ within the subset of those policies in Π_{BS-B} that use the weight vector $\boldsymbol{\alpha}$.

Lemma 4. *Consider the class of weighted balancing policies,* Π_{BS-B}*. Within the subclass of policies which use the weight vector* $\boldsymbol{\alpha}$*, the policy with the base-stock vector* $\mathbf{S}^{\boldsymbol{\alpha}*}$ *is optimal.*

Next, we discuss the special case in which all products are "symmetric", i.e. identical in terms of cost parameters and demand distributions. We are able to make stronger statements about the optimal policy for this special case. We first formally state our assumption.

Assumption 1. *The following conditions hold.*
(a) $h^n = h$ *and* $b^n = b$ *for all* n*.*
(b) $(D^1, \ldots, D^N)$ *has a symmetric distribution, that is, the joint distribution of* $(D^1, \ldots, D^N)$ *is identical to the joint distribution of* $(D^{\theta(1)}, \ldots, D^{\theta(N)})$ *for any permutation* $(\theta(1), \ldots, \theta(N))$ *of* $(1, \ldots, N)$*.*

Theorem 5. *Consider the policy in* Π_{BS-B} *defined by a base-stock vector* $\mathbf{S}$ *and the weight vector* $\mathbf{1}$*. Assume that* $x_1^n = S^n$*. Under Assumption 1 (b), the following statements hold.*

(i) The distribution of $\mathbf{V}_t^{\mathbf{1}}$ *is symmetric for all* t.
(ii) The distribution of $\mathbf{V}_\infty^{\mathbf{1}}$*, the limiting random vector mentioned in Lemma 3, is symmetric.*

Next, we show that the policy in Π_{BS-B} that uses the symmetric allocation rule and the corresponding optimal base-stock vector (as defined in Lemma 4) is optimal over all policies, not just base-stock policies, when all products are identical. This result is the average cost version of Theorem 3 of DeCroix and Arreola-Risa (1998)[2], which pertains to the finite horizon and infinite horizon discounted cost problems.

Theorem 6. *Consider the policy in* Π_{BS-B} *with the weight vector* $\mathbf{1}$ *and the base-stock vector* $\mathbf{S}^{\mathbf{1}*}$*. Under Assumption 1, this policy minimizes the long run average cost per period* $\lim_{T\to\infty}\sup E[\sum_{t=1}^{T} C_t^{\pi}]/T$ *over* Π*, the class of all non-anticipatory policies.*

In the next result, we show that shortfalls under the priority policy can be approximated by policies in Π_{BS-B}.

Lemma 7. *Let* $((1),(2),\ldots,(N))$ *denote any permutation of* $\{1,2,\ldots,N\}$*. Consider any given shortfall vector* $\mathbf{W}$ *(before ordering) in any period. Let* $\boldsymbol{\alpha}_m$ *be defined by* $\alpha_m^{(1)} = 1$ *and* $\alpha_m^{(j)} = m\cdot\alpha_m^{(j-1)}$ *for* $j \in \{2,\ldots,N\}$*. Let* $\mathbf{V}^P$ *and* $\mathbf{V}^{\boldsymbol{\alpha}_m}$ *denote the shortfall vectors after ordering under the priority rule (with priorities* $(1) > (2) > \ldots > (N)$*) and the weighted balancing rule (with weight vector* $\boldsymbol{\alpha}_m$*), respectively. Then, for every* $\epsilon > 0$*, there exists a sufficiently large* M *such that* $|\mathbf{V}^{\boldsymbol{\alpha}_m} - \mathbf{V}^P| < \epsilon$ *for all* $m > M$*, where* $|(u_1,u_2,\ldots,u_n)| = max\{u_1,u_2,\ldots,u_n\}$*.*

7 High Service Level Asymptotics

We show that if the joint distribution of demands for all the products is symmetric and the holding costs for all products are identical, then the best base-stock policy under the symmetric allocation rule is asymptotically optimal along a sequence of problems in which the backorder costs are scaled by a factor β that approaches ∞. In more practical terms, when the cost parameters are such that service levels for all products are high (in any reasonable policy), the best base-stock policy under the symmetric allocation rule is close to being optimal. We note that we do *not* restrict the backorder cost parameters for the products to be identical in this analysis.

We proceed to state our assumptions formally, and then present our analysis.

[2]We note that the conclusion of Theorem 3 of DeCroix and Arreola-Risa (1998) is not completely correct. For example, they claim the following: if, in some period, some products have inventory levels which exceed their optimal base-stock levels and if it is feasible to raise the inventory levels of the other products to their optimal base-stock levels, then the optimal policy is to not produce any of the products in the former category while bringing the other products' inventories to their optimal base-stock levels. This claim is incorrect because the optimal inventory level for a product after ordering depends non-trivially on the inventory levels of the other products since the cost-to-go function is *not* separable even though the single period cost function is separable with respect to the inventory levels of the products. However, we note that their claims are correct for every inventory vector in which every component is below its corresponding optimal base-stock level. The above comments also apply to Theorem 1 of their paper.

Assumption 2. *The following conditions hold.*
(a) All products have identical holding costs, that is, $h^n = h$ *for all* $n \in \{1, \ldots, N\}$.
(b) $(D^1, \ldots, D^N)$ *has a symmetric distribution.*

When the demand vector has a symmetric distribution, let us employ $C^*(h, b)$ to denote the optimal long run average cost of a system in which *all* the products have the same holding cost parameter h and the same backorder cost parameter b. When the backorder costs are not identical (which might generally be the case), we use $C^*(h, \mathbf{b})$ to denote the same except that $\mathbf{b}$ represents the vector of backorder costs over all the products. We denote the long-run average cost of the policy in Π_{BS-B} with parameters $(\mathbf{S^{1*}}, \mathbf{1})$ as $C^{\mathbf{1}*}(h, \mathbf{b})$. Finally, we denote the lowest backorder cost parameter in $\mathbf{b}$ by $min(\mathbf{b})$ and the average of all the individual itemwise backorder costs by $avg(\mathbf{b})$.

In what follows, we note that $C^*(h, b)$ can be evaluated using Theorem 5. In our analysis, we use the cost $C^*(h, b)$ as a basis for cost comparisons across various policies because we know how it can be computed. In fact, we know that

$$C^*(h, b) \;=\; N \cdot L(h, b, V_\infty^{\mathbf{1},1} + D^1) \;, \tag{1}$$

where (i) $V_\infty^{\mathbf{1},1}$ is the marginal distribution of any component of the vector $\mathbf{V}_\infty^{\mathbf{1}}$ (recall that the distribution of $\mathbf{V}_\infty^{\mathbf{1}}$ is symmetric when the distribution of $(D^1, \ldots, D^N)$ is symmetric) and (ii) $L(h, b, X)$ is the optimal cost of a single product newsvendor problem with holding and penalty cost parameters h and b respectively, and facing a demand distribution of X, i.e.,

$$L(h, b, X) \;=\; \min_y \; h \cdot E[(y - X)^+] \;+\; b \cdot E[(X - y)^+] \;.$$

Before proceeding to the details of the analysis leading to the asymptotic optimality result of Theorem 9, we outline the main steps. In Lemma 8, we derive an upper bound on $C^{\mathbf{1}*}(h, \mathbf{b})$ - the long run average cost of the optimal symmetric policy - and a lower bound on the optimal long run average cost $C^*(h, \mathbf{b})$. More specifically, $C^*(h, avg(\mathbf{b}))$ is shown to be an upper bound on $C^{\mathbf{1}*}(h, \mathbf{b})$ while $C^*(h, \min(\mathbf{b}))$ is shown to be a lower bound on $C^*(h, \mathbf{b})$. Notice that both the bounds are optimal costs of systems in which the products are symmetric in costs. (Recall that throughout this section we assume that the product demands are symmetric.) Thus, we can express these bounds as the optimal costs of certain newsvendor problems involving the convolution of demands and shortfalls as explained in the previous paragraph. Our goal is to show that the ratio $\frac{C^{\mathbf{1}*}(h, \beta \cdot \mathbf{b})}{C^*(h, \beta \cdot \mathbf{b})}$ approaches 1 as β approaches ∞. Thus, it is sufficient to show that the ratio of the optimal costs of the two newsvendor problems alluded to above converges to 1 since one of these optimal costs is an upper bound on the numerator of the ratio of interest and the other is a lower bound on its denominator. We establish this convergence in the proof of Theorem 9 by making use of a result in Huh et al. (2009) (presented in our appendix as Lemma 13) for the standard newsvendor problem under a mild distributional assumption on demand. Since the newsvendor problems we are interested in involve the convolution of a product's demand and its shortfall, we have to demonstrate that this convolution also satisfies their assumption - this step is done in Lemma 14 of our appendix.

Lemma 8. *Under Assumption 2, the following inequalities hold:*

$$C^*(h, min(\mathbf{b})) \;\leq\; C^*(h, \mathbf{b}) \;\leq\; C^{\mathbf{1}*}(h, \mathbf{b}) \;\leq\; C^*(h, avg(\mathbf{b})) \;. \tag{2}$$

Lemma 8 bounds the cost of the symmetric policy by the optimal cost of two comparable systems. The upper (lower) bound is provided by a system in which the backorder cost parameter for every product is replaced by the average (minimum) value of the backorder cost parameters.

We are now ready to derive an upper bound on the ratio $\frac{C^{\mathbf{1}*}(h,\mathbf{b})}{C^*(h,\mathbf{b})}$ and show that this ratio approaches 1 as $\mathbf{b}$ is scaled by a factor β which approaches ∞. This is the asymptotic optimality result that we have been referring to all along - for this result, we assume that the demand distribution for every product is IFR, i.e. has an increasing failure rate which is a condition satisfied by several common distributions.

Theorem 9. *Under Assumption 2, the increase in cost due to using the symmetric allocation rule and its corresponding optimal base-stock vector relative to the optimal cost can be bounded as follows:*

$$\left(\frac{C^{\mathbf{1}*}(h, \mathbf{b})}{C^*(h, \mathbf{b})}\right) \;\leq\; \left(\frac{C^*(h, avg(\mathbf{b}))}{C^*(h, min(\mathbf{b}))}\right) \;.$$

Moreover, if the common marginal distribution of the random variables $\{D^j\}$ *is an IFR distribution, this ratio converges to* 1 *as the backorder cost parameters grow, in the following sense:*

$$\lim_{\beta \to \infty} \left(\frac{C^{\mathbf{1}*}(h, \beta\mathbf{b})}{C^*(h, \beta\mathbf{b})}\right) \;=\; 1 \;.$$

8 Heavy Traffic Asymptotics

In this section, we assume without loss of generality that the products are numbered in such a way that $h^1 \geq h^2 \geq \ldots \geq h^N$. We show that when $b^N = \min\{b^j\}$, the priority rule which assigns priorities based on the order $(1, 2, \ldots, N)$ is asymptotically optimal in heavy traffic, i.e. as the capacity κ approaches the expected aggregate demand $E[\sum_1^N D^j]$. As mentioned earlier, Pena-Perez and Zipkin (1997) argued that such a result holds in continuous time systems by appealing to diffusion approximations based on Wein (1992). Our proof, as we will see, is from first principles and does not use such approximations. Moreover, our result is that the asymptotic optimality discussed above holds in the *strong* sense whereas Pena-Perez and Zipkin use it in the *weak* sense. We say that a policy π is asymptotically optimal in the weak sense along a sequence of systems indexed by n if the optimal cost approaches ∞ as n approaches ∞ and the ratio between the cost of π and the optimal cost approaches one. Furthermore, if the absolute difference between the cost of π and the optimal cost is bounded, we say π is strongly asymptotically optimal.

To proceed with our asymptotic analysis, we first introduce some notation. Let $C^*(\mathbf{h}, \mathbf{b}, \kappa)$ be the optimal long run average cost of our inventory system when the holding cost vector is $\mathbf{h}$, the backorder cost vector is $\mathbf{b}$ and the capacity is $\kappa \in (\mu, \infty)$. Let $C^*(h, b, \kappa)$ be the same as $C^*(\mathbf{h}, \mathbf{b}, \kappa)$ when $\mathbf{h} = (h, h, \ldots, h)$ and $\mathbf{b} = (b, b, \ldots, b)$. Let $C^P(\mathbf{h}, \mathbf{b}, \kappa)$ denote the long run average cost of the priority policy, P, which assigns priority based on the order $(1, 2, \ldots, N)$ and uses the corresponding optimal base-stock levels according to Lemma 4.

We present a preliminary lemma on the asymptotic behavior of the optimal cost $C^*(\mathbf{h}, \mathbf{b}, \kappa)$ using a well known result due to Kingman (1962) that a suitably scaled distribution of the waiting time in a single server queue converges to an exponential distribution in heavy traffic.

Lemma 10. *As the capacity κ approaches the expected aggregate demand μ, the optimal cost approaches ∞, i.e.*

$$\lim_{\kappa \downarrow \mu} C^*(\mathbf{h}, \mathbf{b}, \kappa) = \infty \ .$$

Next, we present our assumption on the cost parameters formally before stating and proving our asymptotic result in Theorem 11.

Assumption 3. *The cost parameters satisfy the following conditions: $h^1 \geq h^2 \geq \ldots \geq h^N$ and $b^N = \min\{b^j : 1 \leq j \leq N\}$.*

Theorem 11. *Under Assumption 3, the following statement holds:*

$$C^P(\mathbf{h}, \mathbf{b}, \kappa) - C^*(\mathbf{h}, \mathbf{b}, \kappa) < \infty \text{ for all } \kappa > \mu \ .$$

Therefore, $\lim_{\kappa \downarrow \mu} \frac{C^P(\mathbf{h}, \mathbf{b}, \kappa)}{C^*(\mathbf{h}, \mathbf{b}, \kappa)} = 1$.

9 Policy Performance and Results

Theorem 9 establishes that, as the backorder costs grow (or required service level increases), the optimal cost under the symmetric allocation rule asymptotically approaches the optimal cost when the holding costs and demand distributions of all products are identical. While this result is of theoretical interest, it is also important to know how large the optimality gap of our policy is for reasonable values of the backorder cost parameters and when product demands are not symmetric. We conduct a numerical investigation of this issue in this section.

Lower Bound for Benchmarking the Heuristics: Since the optimal cost is virtually impossible to calculate for a large set of problem instances due to the usual problems associated with dynamic programming, we require an easily computable lower bound on the optimal cost. Although we already have such a lower bound in Lemma 8 for the case of symmetric demands, we require a more generally applicable lower bound because the case of asymmetric demands is also included in our numerical investigation. We state such a lower bound in Lemma 12.

Let $G^n(x) = h^n \cdot E[(x - D^n)^+] + b^n \cdot E[(D^n - x)^+]$ be the expected single period newsvendor cost function for product n. We now develop a lower bound on the optimal long run average cost by using the free balancing relaxation (see, for example, Eppen and Schrage (1981) or Aviv and Federgruen (2001)). Let $F_1(y)$ be defined as follows:

$$F_1(y) = \min_{y} \sum_{n=1}^{N} G^n(y^n) \text{ s.t. } \sum_{n=1}^{N} y^n = y. \tag{3}$$

Note that the computation of $F_1(y)$ can be done quite efficiently using a greedy algorithm to solve the optimization problem above. We can now construct a lower bound on the optimal long run average cost using the function $F_1(\cdot)$.

Recall that V_∞ is the limiting random variable of the stochastic process representing aggregate ending shortfalls, i.e. $\{V_t\}$. We employ this limiting distribution to derive a lower bound on the optimal cost.

Lemma 12. *Let $LB_1 = \min_S E[F_1(S - V_\infty)]$. Then, LB_1 is a lower bound on the optimal long run average cost over Π, the class of all non-anticipatory policies.*

9.1 Existing Heuristics

In this section, we describe the heuristics of DeCroix and Arreola-Risa (1998) and Aviv and Federgruen (2001), and compare their heuristics with our weighted balancing heuristics.

The heuristic of DeCroix and Arreola-Risa (1998) is a stationary base-stock policy. We now explain how their base-stock levels are chosen and what their allocation rule is. Let $(S^1, \ldots, S^N)$ denote the vector of base-stock levels for the N products. For any such vector, the allocation rule they use in every period in which capacity is insufficient to raise the inventory levels of all products to their base-stock levels is that of "balancing" the ratios $\{x^n/S^n\}$, where x^n is the net-inventory of product n at the beginning of the period. That is, allocate capacity to the product with the lowest value of this ratio until this ratio equals the next highest ratio; from then, allocate capacity to these two products until their ratios equal the next highest ratio and so on, until the capacity is exhausted. It still remains to specify how the base-stock vector is chosen. This is done as follows. For every $n \in \{1, \ldots, N\}$, let z^n denote the newsvendor level for product n. That is, $z^n = \max\{\arg\min_y G^n(y^n)\}$. For products $n \in \{2, \ldots, N\}$, let $\gamma^n = z^n/z^1$. Let $f(S^1)$ denote the long run average cost of using the policy with the base-stock vector $(S^1, S^1 \cdot \gamma^2, S^1 \cdot \gamma^3, \ldots, S^1 \cdot \gamma^N)$ and the allocation rule described above. The prescribed value of S^1 is that which minimizes $f(\cdot)$ and the prescribed value of S^n for any $n \neq 1$ is $S^1 \cdot \gamma^n$. Note that the evaluation of $f(S^1)$ for a given value of S^1 requires the computation of the steady state distribution of the shortfall vector. The computational effort for any of our weighted balancing heuristics is essentially the effort required to obtain this distribution. However, the heuristic above requires evaluating $f(S^1)$ over an entire search set for S^1, whereas we compute the steady state distribution of the shortfall vector *only once*. Finally, as evidenced by our numerical experiments, as the cost and demand parameters become more asymmetric or when capacity becomes tighter, our heuristic significantly outperforms the policy described above.

The heuristic of Aviv and Federgruen (2001) is also a stationary base-stock policy. Let $(S^1, \ldots, S^N)$ denote the vector of base-stock levels for the N products, the computation of which we discuss after discussing their allocation rule, which is a myopic allocation rule. That is, in every period in which there is not enough capacity for the inventory levels of all the products to attain their base-stock levels, the vector of inventory levels after ordering, say $(y^1, \ldots, y^N)$ is chosen to be a solution to the following optimization problem:

$$\min_{\mathbf{y}} \sum_{n=1}^{N} G^n(y^n) \text{ s.t } y^n \geq x^n \ \forall\, n \ \text{ and } \sum_{n=1}^{N} (y^n - x^n) = \kappa \ ,$$

where x^n is the net inventory of product n at the beginning of the period. The base-stock vector $(S^1, \ldots, S^N)$ is chosen as the solution to the optimization problem

$$\min \sum_{n=1}^{N} G^n(S^n) \text{ s.t. } \sum_{n=1}^{N} S^n = s \ ,$$

where $s = \arg\min_S E[F_1(S - V_\infty)]$. Recall that

$$F_1(y) \ = \ \min_{y} \sum_{n=1}^{N} G^n(y^n) \ \text{ s.t. } \sum_{n=1}^{N} y^n \ = \ y.$$

In terms of computational effort, this heuristic also requires the computation of the steady state distribution of the aggregate shortfall, in order to obtain the function F_1; thus, this method is comparable to our weighted balancing policies since the main component of the effort required for these policies is the computation of the steady state distribution of the shortfall vector. However, we note that when the distribution of the vector of demands for all the products is symmetric and when all products have identical costs, this heuristic is *not* guaranteed to be optimal (this is because the base-stock levels are not chosen optimally); on the other hand, any weighted balancing policy which uses identical weights and the corresponding optimal base-stock levels is optimal and so is the policy of DeCroix and Arreola-Risa.

9.2 Search Policy under Weighted Balancing

Recall that we have established the optimality of the symmetric policy (a weighted balancing policy with weights of 1) and of the priority policy (a weighted balancing policy with extremely different weights) in the asymptotic regimes of high service levels and heavy traffic, respectively. Motivated by the fact that these two policies are very different in terms of their weight vectors, we propose searching over the space of weight vectors[3] and finding the best weighted balancing policy (i.e. the policy with the lowest cost) - we refer to this as the *search policy*.

We conducted several computational experiments and compared the performance of our search policy (simply represented as "Search" in the tables), with those of the heuristics of DeCroix and Arreola-Risa (1998) (which we refer to as the "DA heuristic" in the tables) and Aviv and Federgruen (2001) (which we refer to as the "AF heuristic" in the tables). We also compared our heuristic policy against the priority policy (represented as "Pri" in the tables).

In addition, we tested the efficiency of our heuristic with respect to the proposed lower bound in Section §9, and for a limited number of "small" problems, against the optimal policy in Section 9.10 captured in Tables 12 through 15.

For the purpose of brevity, we report a summary of our computational results run for a large range of Erlang (k, λ) demands. The choice of the Erlang distribution as the demand distribution allows us to work with any variance to mean ratio. In fact, by using Erlang demand with appropriate parameters k and λ, one can arbitrarily approximate any continuous demand distribution up to the first two moments.

9.3 Computational Design

In the computational study, we report problems with three products, i.e. $N = 3$. Nevertheless, there are several parameters in the multi-product problem, viz. the capacity available, the mean and the standard deviation of the demands of individual products, the holding and penalty costs of each product. We seek to demonstrate the strong performance of our policy, by comparing it to the two existing heuristics in the literature, and to the lower bound in all different possible scenarios. However, to carefully calibrate the performance of our policies against the existing ones, we have to gradually demonstrate the performance with respect to each parameter in the model. To achieve this end, we build our computational experiments by gradually adding more complexity to the problem in each subsection.

First, we test all heuristics while gradually increasing the asymmetry of problem parameters, and demonstrate that our heuristic outperforms all extant heuristics. We also compare our heuristic costs (and a lower bound) against the optimal policy costs, and demonstrate that in fact our policy gets better as the products become more asymmetric.

[3] While an exhaustive search for the weights would involve searching over the $N-1$ dimensional space of positive reals, we perform a one dimensional search using a weight vector which is prescribed by one number m similar to Lemma 7.

⋄ Then, in Section 9.4, we compute the costs of our policy under asymmetric demands to make the following observations: (i) The performance improvement of our policy, as the capacity gets tighter, is more pronounced for asymmetric demands, and (ii) given a fixed capacity, our policy performs better as the demands across different products get more asymmetric.

⋄ In Section 9.5, we explore the impact of asymmetric penalty costs. To do this, we revert back to the symmetric demand case and vary only the backorder costs. To make costs asymmetric, we first assume that the holding cost parameters $\{h^n\}$ are all identical to one initially, and use different backorder cost parameters $\{b^n\}$ for constructing problems with different service levels, thus also depicting cost asymmetry among the products. Again, we notice the performance of our policy improves (i) as the capacity gets scarcer, and (ii) as the penalty costs become more asymmetric.

⋄ In Section 9.6, we combine the effect of asymmetric demands and asymmetric penalty costs. Again, we fix the holding costs to be identical across products, and test different asymmetric possibilities in costs and demands. We allow for the products with highest demand variability to have the highest penalty cost, or lowest penalty cost. Thus, the penalty costs and demand variability can be positively or negatively correlated across the product portfolio. We notice that our policy outperforms the existing heuristics in all these asymmetric scenarios.

⋄ In Section 9.7, we allow for all the parameters, i.e., the holding costs, the penalty costs and the demand distributions to be asymmetric, and find that our policy outperforms the existing policies consistently.

⋄ In Section 9.8, we show that for the cases with high asymmetric (both in demand and costs), that our policy performs well when capacity is ample, and then demonstrate that as the total capacity becomes scarce (i.e, κ decreases), our policy does increasingly better. So, the critical advantage of our approach is that our policy performance improves as managing flexibility becomes more critical.

⋄ In Section 9.9, we show an illustrative table, that shows how our policy continues to outperform all policies, as the capacity becomes tighter, in the case where the product costs are heterogeneous. Furthermore, the demands distributions are heterogenous, and are (negatively) correlated between a pair of products.

⋄ Finally, in Section 9.10, we compare our derived lower bound and the cost from our proposed heuristic against the optimal policy cost over a variety of scenarios, while gradually increasing their complexity. We note that our performance improves (i.e., optimality gap with DP reduces), as the products become more asymmetric and demands become more asymmetric. In fact, the proposed heuristic is much closer to the optimal cost than the lower bound itself is to the optimal cost.

To summarize, in every problem instance we tested, our search policy has a better cost performance than the priority heuristic, the DA and the AF heuristics. Recall from our discussion in Section 2 that the DA heuristic also performs a one dimensional search for the base-stock level of product 1 and that the computational effort required for each step of this search is comparable to the computational effort required for our heuristic at every step of the search. Thus, our policy dominates the DA heuristic when we couple computational effort and cost performance.

In addition, our heuristic performs favorably to the AF heuristic because as the capacity gets tighter, and with more asymmetry in costs and demand distributions of the products - the cases where flexibility is most valuable - the performance of our heuristic becomes increasingly better than the AF heuristic. Finally, in all the asymmetric instances our policy performs better then the priority heuristic.

To summarize, our heuristic consistently outperforms the extant heuristics in the literature. We report the performance in detail in the tables that follow, and finally compare our solution to the optimal solution calculated using the Dynamic Program for small problem instances.

In all computational tables that follow, the % reported under the heading "% gap of Search" and the sub-heading LB refers to the quantity

$$100 * (\text{Cost of Search - Lower Bound})/(\text{Lower Bound}).$$

Similarly, the % reported under the sub-heading DA refers to the quantity

$$100 * (\text{Cost of DA - Cost of Search})/(\text{Cost of Search}),$$

the % reported under the sub-heading Pri refers to the quantity

$$100 * (\text{Cost of Pri - Cost of Search})/(\text{Cost of Search})$$

and finally, the % reported under the sub-heading AF refers to the quantity,

$$100 * (\text{Cost of AF - Cost of Search})/(\text{Cost of Search}).$$

9.4 Effect of Asymmetric Demands

We begin with a base case (the first instance in Table 1), where the capacity is held at 48, and perturb only the demand distributions. (The products continue to be symmetric, except for their demand distributions).

To create systematic demand asymmetry effects, we hold the k-parameter of the Erlang distribution identical across all products and vary λ. Since the mean demand of the Erlang(k, λ) distribution is $\frac{k}{\lambda}$ and the variance is $\frac{k}{\lambda^2}$, by increasing the λ parameter, we decrease the mean demand and the demand variance of a product. For the experiments shown in Table 1, we increase λ for product 1 and decrease λ for product 3, in every successive simulation experiment, while holding all other parameters at the same value. Therefore, as we progress down Table 1, the VTMR (variance to mean ratio) for product 1's demand decreases while the VTMR for product 3's demand increases. Whenever we vary the demand, we will repeat this scheme for all computational experiments that follow in the paper.

$\mathbf{b} = (3,3,3)$			Costs					% gap of heuristic			
λ^1	λ^2	λ^3	LB	Pri	DA	AF	Search	LB	Pri	DA	AF
1	1	1	49.5	56.9	53.1	53.8	53.1	7.2%	7.3%	0.0%	1.4%
1.1	1	0.9	49.7	57.3	53.5	54.3	53.5	7.6%	7.2%	0.0%	1.6%
1.2	1	0.8	50.5	58.7	54.7	55.8	54.7	8.3%	7.3%	0.1%	2.0%
1.3	1	0.7	52.0	61.3	56.9	58.5	56.9	9.4%	7.7%	0.0%	2.8%
1.4	1	0.6	54.2	65.5	60.5	62.9	60.5	11.6%	8.4%	0.0%	4.0%
1.5	1	0.5	57.9	72.9	66.7	70.5	66.6	15.0%	9.4%	0.1%	5.8%
$\mathbf{b} = (10,10,10)$											
1	1	1	84.2	97.2	88.5	89.8	88.5	5.1%	9.8%	0.0%	1.5%
1.1	1	0.9	84.5	97.3	89.0	90.5	89.0	5.3%	9.4%	0.0%	1.8%
1.2	1	0.8	85.4	98.8	90.3	92.2	90.3	5.7%	9.3%	0.0%	2.1%
1.3	1	0.7	86.7	101.5	92.6	95.1	92.6	6.8%	9.6%	0.0%	2.7%
1.4	1	0.6	88.1	108.2	95.8	99.3	95.8	8.7%	13.0%	0.0%	3.8%
1.5	1	0.5	89.5	123.0	105.0	107.5	104.0	16.2%	18.2%	0.9%	3.3%
$\mathbf{b} = (15,15,15)$											
1	1	1	96.5	111.4	100.9	102.5	100.9	4.6%	10.4%	0.0%	1.5%
1.1	1	0.9	96.7	111.3	101.3	102.8	101.3	4.8%	9.8%	0.0%	1.5%
1.2	1	0.8	97.3	112.4	102.5	104.1	102.5	5.3%	9.7%	0.1%	1.6%
1.3	1	0.7	98.0	116.0	104.2	106.6	104.2	6.3%	11.3%	0.1%	2.3%
1.4	1	0.6	98.0	124.4	107.3	110.0	106.7	8.9%	16.7%	0.6%	3.1%
1.5	1	0.5	98.3	141.5	118.2	123.0	117.8	19.8%	20.1%	0.3%	4.3%

Table 1: Cost behavior of the proposed heuristic as the demands become asymmetric. For the displayed table: All products have the same holding cost $h = 1$ and backorder cost (as mentioned). However, the demands are asymmetric. We have $k^1 = k^2 = k^3 = 11$, $\lambda^j, j = 1, 2, 3$ are indicated as above. Product 1 (Product 3) has a lower (higher) mean demand and lower (higher) variance, when compared to product 2. The total capacity is 48.

We note three successive tabulations in Table 1. In each sub-table, we progressively increase the value of the symmetric penalty costs, beginning from 3 to 15 in the bottom of the table. Within each table, the demand is made more and more asymmetric. We note that the typical deviation from the lower bound is between 4.9% to 16.1%.

Our heuristic performs significantly better than the priority policy in all the cases. In fact, as the penalty costs increase, the relative superior performance of our policy is more pronounced. Also observe that our heuristic performs better than the AF heuristic. Note that the performance improves as the demands become more asymmetric. We also note that the cost performance of our heuristic is comparable to the DA heuristic under asymmetric demands and symmetric costs. (In most cases, the deviation is negligible and our heuristic performs better within a 1.0%).

Nevertheless this comparison raises a natural question. What are the effects of products having asymmetric costs? Does our heuristic improve its performance (relative to the extant heuristics in the literature)? We explore this question in Section 9.5.

9.5 Effect of Asymmetric Penalty Costs

To explore the effect of asymmetric backorder costs, we revert to our base case. Beginning with our base case (with ample capacity, symmetric costs, symmetric demands), we examine the performance of our heuristic under the scenario where the demand distributions for all products are identical, but the backorder costs are asymmetric. Just as in the base case, the mean demand and variance are fixed as in the first line of Table 1.

In the set of simulations in Tables 2 and 3, we gradually increase the backorder cost of product 1 successively, while decreasing the backorder cost of product 3, thus making the newsvendor fractiles of the products more asymmetric. Again, throughout the computations in Table 2, the demand distributions of the products are kept symmetric, and the capacity is kept fixed at the same level. In Table 3, we follow the scheme as in Table 2, except that the available overall capacity is reduced.

$K = 48$			Costs of Heuristics					% gap of Search			
b^1	b^2	b^3	LB	Pri	DA	AF	Search	LB	Pri	DA	AF
4.5	3	2	50.4	59.9	56.7	57.0	56.2	11.5%	6.5%	0.9%	1.4%
6	3	1.5	50.9	59.0	58.2	57.0	56.3	10.6%	4.8%	3.4%	1.2%
8	3	1	51.0	57.4	59.9	55.8	55.7	9.2%	3.0%	7.6%	0.3%
9	3	1	52.1	58.5	61.4	57.1	56.9	9.2%	2.9%	8.0%	0.4%
12	3	0.75	53.3	58.8	64.2	57.7	57.5	7.9%	2.2%	11.7%	0.3%
15	3	0.6	54.4	59.2	66.7	58.4	58.2	7.0%	1.8%	14.6%	0.4%

Table 2: In the above instance, we consider three products with symmetric demands distributed Erlang(k, λ) with $k = 12$ and $\lambda = 1$. The backorder costs are asymmetric. While the holding costs are identical, the backorder cost of product 1, b^1 is progressively increased down the table, and the backorder cost of product 3, b^3 is decreased. The total capacity is $K = 48$.

From Table 2, our performance is significantly better than the priority policy. As one would expect, as the costs become asymmetric, the priority policy improves, but our heuristic continues to outperform the priority policy.

For asymmetric costs, we observe that our Search policy performs significantly better than the DA heuristic. As the cost difference between the highest and the lowest backorder costs of all products $\{\max\{b^j\} - \min\{b^j\}\}$ increases, we note that our heuristic performs much superior to the DA heuristic, providing as much as a 14.6% cost difference.

When examined against the AF heuristic, our performance is comparable. Nevertheless, our cost performance is slightly better (by about 1.4%), when the penalty costs are high and less asymmetric, while our heuristic performs comparably (about 0.5% better) when the costs are more asymmetric.

$K = 44$			Costs of Heuristics					% gap of Search			
b^1	b^2	b^3	LB	Pri	DA	AF	Search	LB	Pri	DA	AF
4.5	3	2	53.1	68.3	63.9	64.4	62.5	17.7%	9.3%	2.2%	3.1%
6	3	1.5	53.2	66.0	66.1	63.2	62.0	16.5%	6.5%	6.7%	2.0%
8	3	1	52.7	62.8	69.3	61.1	60.4	14.6%	4.1%	14.8%	1.3%
9	3	1	53.9	64.0	71.4	62.1	61.5	14.1%	4.0%	16.0%	0.9%
12	3	0.75	54.8	63.3	75.9	61.9	61.5	12.2%	2.9%	23.4%	0.7%
15	3	0.6	55.6	63.1	79.8	62.2	61.6	10.8%	2.4%	29.5%	0.9%

Table 3: Cost behavior of the proposed heuristic as the backorder costs become asymmetric. The computational experiments are structured similar to those in Table 2, except that the total capacity is reduced to 44.

Suppose the aggregate capacity becomes tighter. How does our heuristic perform when products have asymmetric costs? To address this question, we decrease the total capacity from $K = 48$ (in Table 2) to $K = 44$ in Table 3. This table shows that as the capacity becomes tighter (as the capacity utilization increases from 75% to 82%), the performance of our heuristic (over other heuristics) improves on *every* instance.

As the capacity got tighter, the performance of our heuristic improves over the priority heuristic significantly. For instance, comparing the first rows of Table 2 and 3, as capacity got tighter the relative performance of our policy over the priority heuristic improves from about 8% to 11.6%.

Again, the performance of the Search policy is equal or better than the AF heuristic in all instances. In almost all cases, the relative performance nearly *doubled* from the values in Table 2. The Search policy significantly outperforms the DA heuristic. In fact, in the last case reported in Table 3, which is the most asymmetric case, the our heuristic outperforms the DA heuristic by 29.5% (improving from 14.6%, in the corresponding case with more capacity seen in Table 2).

It is also worth noting that our performance weakens slightly with respect to the lower bound, as products become more asymmetric. But the relative improvement in comparison to other heuristics indicate that all the heuristics are weak, and our heuristic is less so.

Given that the lower bound is not very tight in this instance, one might surmise that the optimality gap of our heuristic is low (less than 10% in the most asymmetric case). The strong performance of our policy relative to the lower bound is significant, since the optimal allocation rule structure remains unknown for the multi-product capacitated problem.

9.6 Effect of Asymmetric Demand and Penalty Costs

While the previous two subsections focus on problems with 1) asymmetric demands but symmetric costs, and 2) asymmetric costs but symmetric demands, we now study problems where *both* demands and costs are asymmetric.

To begin the analysis, we report a set of experiments in Table 4. In Table 4, we construct a set of experiments in which both the demand distributions and penalty costs are different for the three products. In each sub-table within the Table 4, we sequentially increase the demand asymmetry. From one sub-table to another sub-table, we test gradually increasing penalty costs beginning from the set of costs $(1, 5, 10)$ to $(3, 6, 15)$.

We find that as the demands becomes more asymmetric, our heuristic performs better than both the DA heuristic and the AF heuristic.

Backorder Costs $(1, 5, 10)$			Costs					% gap of Search			
λ_1	λ_2	λ_3	LB	Pri	DA	AF	Search	LB	Pri	DA	AF
1	1	1	59.7	98.2	78.3	68.3	67.6	13.2%	44.7%	15.6%	0.9%
1.1	1	0.9	58.5	99.9	77.9	67.3	66.7	14.0%	49.7%	16.7%	0.9%
1.2	1	0.8	58.5	107.5	81.1	68.0	67.3	15.0%	59.6%	20.5%	1.1%
1.3	1	0.7	60.4	125.8	90.0	71.3	70.6	16.9%	78.2%	27.5%	1.1%
1.4	1	0.6	67.9	170.2	114.7	81.4	80.1	18.0%	112.3%	43.1%	1.6%
$(2, 5, 12)$											
1	1	1	67.6	108.0	84.1	79.4	77.6	14.8%	37.9%	8.3%	2.4%
1.1	1	0.9	66.9	110.2	83.9	79.1	77.8	16.3%	41.7%	7.9%	1.7%
1.2	1	0.8	67.8	118.8	87.5	81.3	79.6	17.4%	49.3%	9.9%	2.1%
1.3	1	0.7	71.7	138.9	96.6	87.3	85.3	19.0%	62.8%	13.2%	2.3%
1.4	1	0.6	84.4	186.9	121.6	103.2	100.7	19.3%	85.5%	20.7%	2.4%
$(3, 6, 15)$											
1	1	1	75.8	119.6	92.1	89.6	87.0	14.8%	35.9%	4.6%	2.1%
1.1	1	0.9	75.4	122.3	92.3	89.7	87.8	16.4%	39.4%	5.1%	2.2%
1.2	1	0.8	76.9	131.9	95.9	92.7	90.3	17.4%	46.1%	6.2%	2.6%
1.3	1	0.7	82.1	154.2	105.9	100.7	97.5	18.8%	58.1%	8.6%	3.3%
1.4	1	0.6	98.3	206.5	132.9	120.3	117.1	19.1%	76.3%	13.5%	2.7%

Table 4: Cost behavior of the proposed heuristic as both costs and demands become asymmetric. The computational experiments are structured similar to those in Table 1 except that the product backorder costs are also asymmetric and $k^1 = k^2 = k^3 = 12$, $\lambda^j, j = 1, 2, 3$ are indicated as above. Available aggregate capacity is 44.

The performance of our policy in relation to the lower bound, is steady about 13-19% across all sub-tables. The gap with lower bound increases slightly as the demand becomes asymmetric, and the penalty costs increases.

Our policy performs significantly better than the priority policy in all cases (on average about 40% or better). Compared to the DA heuristic, our policy performs better (i) as the demand becomes more asymmetric for given penalty costs, and (ii) as the costs decreases given the same demand characteristics. Finally, our policy performance is better than the AF heuristic. Furthermore, as the penalty costs increase, the relative performance of our policy improves.

In Table 4, the items that had higher backorder costs had lower demand variability. In Table 5, the backorder costs are reversed for the three products such that the product with higher backorder cost also faces demand with less variability. In general, our policy performs strongly compared to the existing heuristics in these asymmetric cases. In fact, when compared to the symmetric demand and symmetric cost cases discussed before, the performance of our heuristic is now consistently better than the DA and the AF heuristic.

Backorder Costs $(10,5,1)$			Costs					% gap of Search			
λ_1	λ_2	λ_3	LB	Pri	DA	AF	Search	LB	Pri	DA	AF
1	1	1	59.7	70.4	78.2	68.26	67.6	13.2%	4.2%	15.6%	0.9%
1.1	1	0.9	61.7	73.5	80.8	70.98	70.3	13.9%	4.7%	15.0%	1.0%
1.2	1	0.8	64.7	78.8	87.3	75.80	74.6	15.3%	5.5%	16.9%	1.6%
1.3	1	0.7	69.2	87.3	101.4	83.53	81.7	18.1%	6.9%	24.1%	2.2%
1.4	1	0.6	77.8	103.4	134.9	98.14	95.0	22.1%	8.8%	42.0%	3.3%
$(12,5,2)$											
1	1	1	67.5	83.4	84.1	79.39	77.6	15.0%	7.4%	8.3%	2.3%
1.1	1	0.9	69.4	86.4	86.8	81.80	80.1	15.4%	7.9%	8.3%	2.1%
1.2	1	0.8	72.6	92.5	93.0	86.83	85.0	17.1%	8.8%	9.4%	2.2%
1.3	1	0.7	78.1	103.3	107.6	95.69	93.6	19.8%	10.4%	15.0%	2.3%
1.4	1	0.6	90.6	125.4	143.0	115.59	112.4	24.1%	11.6%	27.2%	2.9%
$(15,6,3)$											
1	1	1	75.7	95.6	91.9	89.50	87.0	14.9%	14.8%	5.6%	2.9%
1.1	1	0.9	77.5	98.8	94.7	92.33	89.4	15.4%	10.6%	6.0%	3.3%
1.2	1	0.8	80.7	105.6	101.4	97.35	94.4	17.0%	11.9%	7.4%	3.2%
1.3	1	0.7	86.8	118.2	116.8	107.19	103.9	19.7%	13.8%	12.5%	3.2%
1.4	1	0.6	102.6	144.8	154.9	132.67	127.7	24.5%	13.5%	21.4%	3.9%
$(20,10,5)$											
1	1	1	89.4	115.4	105.6	105.73	102.6	14.8%	14.1%	2.9%	3.1%
1.1	1	0.9	91.0	119.1	108.1	108.59	104.8	15.2%	13.6%	3.1%	3.6%
1.2	1	0.8	94.2	127.2	114.9	113.65	109.9	16.7%	15.7%	4.6%	3.4%
1.3	1	0.7	100.9	142.6	131.8	125.18	121.1	20.0%	17.8%	8.8%	3.4%
1.4	1	0.6	122.9	176.3	173.4	160.46	151.9	23.6%	16.1%	14.2%	5.7%

Table 5: Cost behavior of the proposed heuristic as both costs and demands become asymmetric. Available aggregate capacity is 44.

Table 5 indicates that as the backorder costs diverge further, the performance of the Search policy improves relative to the Priority, the DA and the AF heuristics, when compared with the same asymmetric demand case (see corresponding rows in Table 4). As products' demand variances diverge, our heuristic provides a cost performance of about, 4% or better than the Priority heuristic, 9% or better than the DA heuristic, and roughly 1% or better than the AF heuristic.

9.7 Asymmetric Demands, Holding, and Penalty costs

In this section, we examine general asymmetric problems with varying holding costs, penalty costs and demand distributions. We imposed asymmetric holding costs on a problem with significant asymmetry in penalty costs *and* demand distributions, and measured the performance of the Search policy.

Table 6 shows how our policy compares to the Priority, the DA and the AF heuristics. Our policies generally continue to perform better than the extant heuristics. Compared to the Priority policy, our heuristic provides about 9% or more cost savings. In fact, as the costs increase (in the lower sub-table of Table 6), we have the cost savings of our heuristics increase to 12% higher over the priority heuristic. The performance increases as the product demands become more asymmetric.

Comparing our policy to the DA policy, we find that our policy improves as the demand becomes more asymmetric and for lower penalty costs. Nevertheless, our policy significantly outperforms the DA policy with a cost benefit ranging from from 3.5% to 23.5%.

Costs $\mathbf{b} = (15, 6, 3)$ $\mathbf{h} = (1.1, 1, 0.9)$			Costs					% gap of Search			
λ_1	λ_2	λ_3	LB	Pri	DA	AF	Search	LB	Pri	DA	AF
1	1	1	77.1	96.0	93.3	90.8	88.0	14.1%	13.9%	6.0%	3.1%
1.1	1	0.9	79.1	99.6	96.6	93.9	90.8	14.8%	9.7%	6.4%	3.4%
1.2	1	0.8	82.8	106.7	103.2	99.8	96.2	16.2%	10.9%	7.3%	3.8%
1.3	1	0.7	89.4	119.4	119.5	110.2	106.2	18.8%	12.4%	12.5%	3.7%
1.4	1	0.6	105.3	145.8	159.7	134.6	129.3	22.8%	12.8%	23.5%	4.1%
$\mathbf{b} = (20, 10, 5)$ $\mathbf{h} = (1.1, 1, 0.9)$			Costs					% gap of Search			
1	1	1	90.7	115.3	106.9	107.0	103.3	13.9%	11.6%	3.4%	3.5%
1.1	1	0.9	92.6	119.3	109.8	110.1	106.0	14.5%	12.6%	3.7%	3.9%
1.2	1	0.8	96.3	127.7	117.1	116.3	111.6	15.9%	14.4%	5.0%	4.3%
1.3	1	0.7	103.6	143.1	135.0	128.4	123.0	18.7%	16.4%	9.8%	4.4%
1.4	1	0.6	125.1	175.9	178.1	161.9	152.9	22.2%	15.0%	16.5%	5.8%

Table 6: Cost behavior of the proposed heuristic all parameters become asymmetric. i.e, backorder costs, holding costs and the demand distributions are all asymmetric. Available aggregate capacity is 44.

In Table 7, we increase the holding costs to $(1.2, 1.0, 0.8)$, thus the problems become more asymmetric, in both penalty and holding costs. Our policy continues to outperform other heuristics. This observation persists at higher penalty costs. Again, our policy cost performance improves as the demands become more asymmetric.

Costs $\mathbf{b} = (15, 6, 3)$ $\mathbf{h} = (1.2, 1, 0.8)$			Costs					% gap of Search			
λ_1	λ_2	λ_3	LB	Pri	DA	AF	Search	LB	Pri	DA	AF
1	1	1	78.2	96.2	94.5	91.9	88.8	13.6%	8.3%	6.3%	3.5%
1.1	1	0.9	80.6	100.1	98.2	95.3	92.0	14.1%	8.9%	6.8%	3.7%
1.2	1	0.8	84.6	107.4	105.4	101.6	97.8	15.6%	9.9%	7.8%	3.9%
1.3	1	0.7	91.7	120.3	123.1	112.9	108.0	17.8%	11.3%	13.9%	4.5%
1.4	1	0.6	107.9	146.2	164.0	136.3	130.4	20.8%	12.1%	25.7%	4.5%
$\mathbf{b} = (20, 10, 5)$ $\mathbf{h} = (1.2, 1, 0.8)$			Costs					% gap of Search			
1	1	1	91.7	114.9	107.8	108.0	103.8	13.3%	10.7%	3.8%	4.1%
1.1	1	0.9	94.1	119.3	111.2	111.5	106.9	13.6%	11.6%	4.0%	4.3%
1.2	1	0.8	98.2	127.8	118.8	118.8	112.9	15.0%	13.2%	5.2%	5.2%
1.3	1	0.7	106.0	143.2	137.0	130.6	124.6	17.5%	15.0%	10.0%	4.9%
1.4	1	0.6	127.2	175.0	182.5	162.4	153.5	20.7%	14.0%	18.9%	5.8%

Table 7: Cost behavior of the proposed heuristic all parameters become asymmetric. i.e, backorder costs, holding costs and the demand distributions are all asymmetric. The parameters are identical to the Table 6 except for holding cost parameters that are more asymmetric.

9.8 Effect of Capacity

In this section, we explore the effect of capacity K by considering the set of asymmetric problems with different backorder costs, holding costs, and different demands for all the products.

In the first table we analyze the effect of capacity, for a problem, with following demands (Erlang distribution with $k = 12$ and $\lambda^1 = 1.5, \lambda^2 = 1.0$ and $\lambda^3 = 0.5$). The penalty costs are set at $(15, 6, 3)$ for each product respectively and holding costs are held at $(1.1, 1.0, 0.9)$ respectively. In Table 8 the results are summarized for an Erlang distribution for each product with parameters $k = 12$ and $\lambda^1 = 1.5, \lambda^2 = 1.0$ and $\lambda^3 = 0.5$). Thus product 1 has the lowest variability and product 3 has the highest variability.

In Tables 8 and 9, we decrease the capacity such that the utilization increases from 73.3% ($K = 60$) to 97.78% (for $K = 45$). The observation made earlier continues to hold: the performance of our policy is especially improved in the cases of limited capacity availability, which indicates that our policy is more efficient in allocating scarce resources amongst the products.

In Table 8, as the capacity falls from 60 to 45, the relative performance of our heuristic over all other heuristics improves consistently. Against the priority heuristic, the relative cost advantage of our heuristic improves from 6% to 9%. Against the DA heuristic, as the capacity gets tighter, the relative advantage increases from a low 6% to as high as 46%. Similarly, against the AF heuristic, the performance improves from 0.3% to about $3 - 5\%$.

In Table 8, we note that as the capacity gets tighter, the relative difference between the WB heuristic and the lower bound increases. This is due to the weakened nature of the lower bound under high utilization (as can be seen in the comparison against the dynamic program in the next subsection). This is intuitive because when capacity is unlimited the multi-product problem de-

$\lambda = (1.5, 1, 0.5)$ $\mathbf{b} = (15, 6, 3)$ $\mathbf{h} = (1.1, 1, 0.9)$	Costs					% gap of Search			
Capacity	LB	Pri	DA	AF	Search	LB	Pri	DA	AF
60	78.7	88.7	86.3	83.8	83.6	6.2%	6.1%	3.3%	0.3%
59	78.8	89.8	87.8	84.7	84.3	7.0%	6.5%	4.1%	0.4%
58	78.9	91.0	89.5	85.8	85.2	8.0%	6.8%	5.0%	0.6%
57	79.1	92.5	91.6	87.0	86.3	9.1%	7.1%	6.2%	0.8%
56	79.4	94.1	94.4	88.6	87.7	10.5%	7.4%	7.7%	1.1%
55	79.7	96.1	96.3	90.5	89.3	12.0%	7.7%	7.9%	1.4%
54	80.3	98.5	99.9	92.9	91.2	13.6%	8.0%	9.6%	1.9%
53	81.1	101.3	103.5	95.5	93.5	15.3%	8.3%	10.7%	2.1%
52	82.1	104.7	108.4	98.8	96.4	17.4%	8.6%	12.4%	2.4%
51	83.8	109.0	114.3	103.1	100.2	19.6%	8.8%	14.1%	2.9%
50	86.1	114.3	121.5	109.0	104.8	21.7%	9.1%	15.9%	3.9%
49	89.7	121.2	131.4	115.6	110.9	23.6%	9.3%	18.5%	4.2%
48	95.2	130.1	144.9	125.4	118.9	24.9%	9.4%	21.9%	5.5%
47	104.2	142.5	163.5	138.1	130.1	24.8%	9.5%	25.7%	6.1%
46	120.2	161.6	197.3	156.1	147.9	23.0%	9.3%	33.4%	5.5%
45	150.0	194.0	262.5	185.7	179.3	19.5%	8.2%	46.4%	3.5%

Table 8: Cost behavior of the proposed heuristic as the total capacity becomes tighter for highly asymmetric demand and cost cases. For the displayed table: we have three products with asymmetric demands distributed Erlang(k, λ) with $k = 12$ and $\lambda_1 = 1.5, \lambda_2 = 1.0, \lambda_3 = 0.5$.

composes into N individual newsvendor problems. In this case, the balancing heuristic is optimal and the lower bound coincides with the optimal cost. As the capacity gets tighter, the issue of allocating capacity becomes important and the lower bound benefits from the fact that it allows for costless redistribution of inventories in each period. In any case, the relative performance of our heuristic continues to improve as the capacity becomes tighter.

In Table 9, we follow the same schematic as in Table 8, except that the pattern of asymmetric demands are now reversed. The results are summarized for an Erlang distribution for each product with parameters $k = 12$ and $\lambda^1 = 0.5, \lambda^2 = 1.0$ and $\lambda^3 = 1.5$. In other words, for the same holding and penalty costs as in Table 8, product 1 has the highest variability and product 3 has the lowest variance to mean ratio. Similar to Table 8, we display the costs and relative performance of the policies as the capacity gets progressively scarcer.

As the capacity falls from 60 to 45, the relative performance of our heuristic over all other heuristics improves consistently. Against the priority heuristic, the relative cost advantage of our heuristic improves from 6.4% to 10.1%. Against the DA heuristic, as the capacity gets tighter, the relative advantage increases from a low 3.2% to as high as 34%. Similarly, against the AF heuristic, the performance of our heuristic improves from 0.7% to about 8.8%, as the capacity becomes tighter. In addition, when the the highest penalty cost product also has the highest variance to mean ratio

Parameters $\lambda = (0.5, 1, 1.5)$ $\mathbf{b} = (20, 10, 5)$ $\mathbf{h} = (1.2, 1, 0.8)$	Costs					% gap of Search			
Capacity	LB	Pri	DA	AF	Search	LB	Pri	DA	AF
60	86.6	99.3	96.3	93.9	93.3	7.7%	6.4%	3.2%	0.7%
59	86.7	100.7	98.0	95.2	94.3	8.8%	6.7%	3.9%	0.9%
58	86.9	102.4	100.1	96.7	95.6	10.0%	7.1%	4.8%	1.2%
57	87.2	104.3	101.9	98.6	97.1	11.4%	7.5%	4.9%	1.5%
56	87.5	106.6	104.3	100.3	98.9	13.0%	7.8%	5.5%	1.4%
55	88.1	109.2	106.9	102.5	101.0	14.6%	8.2%	5.9%	1.5%
54	88.8	112.4	110.4	105.9	103.5	16.5%	8.6%	6.7%	2.2%
53	89.9	116.1	114.3	109.4	106.6	18.6%	9.0%	7.3%	2.7%
52	91.5	120.7	119.7	114.2	110.3	20.6%	9.4%	8.5%	3.6%
51	93.8	126.3	125.4	120.4	114.9	22.5%	9.9%	9.1%	4.8%
50	97.2	133.4	133.5	127.3	120.8	24.3%	10.4%	10.5%	5.4%
49	102.3	142.3	144.3	137.0	128.3	25.4%	10.9%	12.5%	6.8%
48	110.1	154.0	159.0	149.5	138.3	25.6%	11.4%	14.9%	8.1%
47	122.5	170.2	179.7	164.7	152.5	24.5%	11.6%	17.9%	8.0%
46	144.5	195.5	216.7	190.6	175.5	21.5%	11.4%	23.5%	8.6%
45	184.3	238.6	290.1	234.9	216.8	17.6%	10.1%	33.8%	8.8%

Table 9: Cost behavior of the proposed heuristic as the total capacity becomes tighter for highly asymmetric demand and cost cases. For the displayed table: asymmetric demands are *reversed* from the previous table, i.e., Erlang(k, λ) with $k = 12$ and $\lambda_1 = 0.5, \lambda_2 = 1.0, \lambda_3 = 1.5$.

(VTMR), we note that our heuristic significantly outperforms the AF heuristic (compared to the case when the product with the highest penalty cost has the lowest variance to mean ratio). To summarize, under scarce capacity our approach does 9% or better cost-wise against *every* extant heuristic.

In Table 10 show the base-stock levels for the scenarios reported in Table 9. In general, it appears that the priority policy assigns a significantly higher base-stock for the product 3 (which is cheapest to hold). On the other hand, the DA heuristic chooses inventories such that a significantly higher base stock is assigned to Product 1. Our Search policy and the AF heuristic both choose base-stock levels that are in between the those chosen under the Priority and the DA heuristics. It appears that the AF heuristic chooses weakly lower base-stocks for the products, compared to our policy. These differences are more pronounced as the capacity becomes tighter. It also appears that our policy outperforms significantly better than other policies when the capacity is scarce, by setting up the base-stock parameters appropriately. The difference in the base-stock levels in our policy and those in the other heuristics may be possibly due to the better allocation approach used in our policy.

For instance, examine the case when $K = 45$ (the last line in Table 9). Although, it is hard to characterize the structure of the optimal policy decisions, we consider a simple scenario which illustrates the different decisions made under the policies which we study in this paper.

Capacity	Priority			DA			AF			Search		
60	60	35	24	60	29	16	60	29	16	60	32	20
59	60	35	25	60	29	16	60	29	16	60	33	20
58	60	36	26	60	29	16	60	29	16	60	33	21
57	60	36	27	62	30	17	60	29	16	60	33	21
56	60	37	28	64	31	17	60	30	16	60	34	22
55	60	38	30	66	32	18	60	30	17	60	34	25
54	60	39	32	68	33	18	60	30	17	60	35	26
53	60	39	35	70	34	19	60	31	17	62	36	28
52	60	41	38	72	35	19	60	31	18	63	37	29
51	60	42	42	78	38	21	60	32	18	64	38	32
50	60	43	47	82	40	22	60	33	20	66	40	34
49	60	45	55	86	42	23	60	35	21	67	41	41
48	60	47	64	93	45	25	60	37	24	70	44	47
47	60	49	78	105	51	28	60	41	30	72	46	59
46	60	52	101	122	59	33	60	46	44	75	49	77
45	60	56	139	150	73	40	60	48	79	77	52	113

Table 10: Base-stock levels for the proposed heuristics as the total capacity becomes tighter for highly asymmetric demand and cost cases for the parameters and costs shown in Table 9.

Let the beginning inventory levels in some period be $(70, 70, 80)$ for the three products. Under the priority policy, we have to produce 59 items for product 3, none for products 1 and 2. Due to limited capacity, shortfalls continue to exist (for product 3). Under the DA policy, we have to produce 80 units of product 1 and 3 items for product 2 and none for product 3. Even in this case, shortfalls continue to exist, and surely for product 1, since capacity available is 45 but 80 items have to be produced. Under the AF heuristic, all items are above their base stock level under the heuristic, so the entire capacity goes unused. However, under our policy, 33 items of product 3 are produced. There is no shortfall. In this scenario, the DA and priority heuristics allow for too much shortfall for different products, and under the AF heuristic, the capacity may go unused (compared to the Search policy). Our policy tries to find a balance between excessive shortfalls (due to high base-stock levels) and low utilization (due to low base-stock levels).

9.9 Correlated Demands

In the correlated demand case shown in Table 11, we follow the same schematic (i.e. the tables display the costs and relative performance of the policies as the capacity gets progressively scarcer) as in Table 9, except that we use demand distributions that are correlated as described in the caption of Table 11.

As the capacity falls from 60 to 45, the relative performance of our heuristic over all other heuristics improves consistently. Against the priority heuristic, the relative cost advantage of our heuristic improves from 6.1% to 9.4%. Against the DA heuristic, as the capacity gets tighter, the relative advantage increases from a low 3.2% to as high as 46%. Similarly, against the AF heuristic, the performance of our heuristic improves from 0.3% to 13%, as the capacity becomes tighter.

Parameters $\mathbf{b} = (20, 10, 5)$ $\mathbf{h} = (1.2, 1, 0.8)$	Costs					% gap of Search			
Capacity	LB	Pri	DA	AF	Search	LB	Pri	DA	AF
60	69.33	86.03	79.14	79.24	77.11	11.2%	11.6%	2.6%	2.8%
59	69.49	87.68	80.57	80.94	78.25	12.6%	12.0%	3.0%	3.4%
58	69.72	89.56	82.23	82.98	79.59	14.2%	12.5%	3.3%	4.3%
57	70.04	91.75	84.27	85.45	81.17	15.9%	13.0%	3.8%	5.3%
56	70.51	94.30	86.73	87.69	83.06	17.8%	13.5%	4.4%	5.6%
55	71.14	97.25	89.63	90.61	85.32	19.9%	14.0%	5.0%	6.2%
54	72.05	100.71	93.24	94.98	88.04	22.2%	14.4%	5.9%	7.9%
53	73.34	104.80	97.85	99.50	91.20	24.3%	14.9%	7.3%	9.1%
52	75.13	109.72	103.61	104.42	95.06	26.5%	15.4%	9.0%	9.8%
51	77.74	115.78	110.77	111.04	99.98	28.5%	15.9%	10.9%	11.1%
50	81.55	123.30	119.75	118.91	106.15	30.2%	16.2%	12.8%	12.0%
49	87.10	132.73	131.44	129.05	114.11	31.0%	16.3%	15.2%	13.1%
48	95.46	144.95	147.14	141.69	124.81	30.8%	16.1%	17.9%	13.5%
47	108.50	161.64	169.98	156.73	139.75	28.8%	15.7%	21.6%	12.2%
46	131.02	187.48	209.13	179.30	163.74	25.0%	14.5%	27.7%	9.5%
45	171.24	231.13	283.69	225.54	205.96	20.3%	12.2%	37.7%	9.6%

Table 11: Cost behavior of the proposed heuristic as the total capacity becomes tighter for highly asymmetric demand and cost cases. The means of the distributions are 19.61, 9.89 and 7.97 respectively, and the corresponding standard deviations are 13.18, 6.95 and 5.97 respectively. Demands are correlated for Product 1 and 2 are positively correlated (correlation factor $\simeq 0.3$), and they are both negatively correlated with demand for Product 3.

In the previous sections, we have established the relative performance of our heuristic, over the lower bound and the priority, the DA and the AF heuristics. In general, our heuristic performs consistently better than all the competing heuristics, and even more so when the capacity is very limited, and product costs/demands are asymmetric.

Nevertheless, it is important to explore how our policy performs against the optimal policy. The optimal cost can be computed using dynamic program, but the procedure suffers from the curse of dimensionality. Hence, we compare our heuristic and the lower bound against the optimal cost derived from the Dynamic Program formulation on a set of "small" problems. We report the costs of all other policies, but do not report their optimality gaps which can be calculated quickly from the information available in the tables. This also helps us focus on the relative performance of our heuristic and the lower bounds against the dynamic program.

9.10 Comparison with Dynamic Program

In this section, through Tables 12 through 15, we explore the performance of our lower bound against the optimal cost from the dynamic program over a set of computational experiments with limited state space. While the problems are large enough to be realistic comparisons, the state space is small enough for the dynamic programming optimal solution to be achieved within few hours in each instance.[4]

In all computational tables that follow, the costs are reported for the lower bound (LB), the optimal policy (DP), the priority policy (Pri), the DA heuristic, the AF heuristic and our Search policy. Then, we also report the gaps with optimal DP for the lower bound and our heuristic. The % reported under the heading "% Optimality Gap" and the sub-heading DP-LB refers to the quantity

$$100 * (\text{Dynamic Program Cost - Cost from Lower Bound})/(\text{Cost from the Lower Bound}).$$

Similarly, the % reported under the sub-heading Search-DP refers to the quantity

$$100 * (\text{Cost of Search - Optimal DP cost})/(\text{Cost of Search Policy}).$$

In all the computational experiments, the demands for product were kept according to an Erlang distribution with asymmetric parameters $k^1 = 3, k^2 = 4, k^3 = 3$. (Recall that the superscripts represent the product indices).

$(k^1, k^2, k^3) = (3, 4, 3)$ $(\lambda^1, \lambda^2, \lambda^3) = (1, 1, 1.5)$ $\mathbf{h} = (1, 1, 1)$			Costs						% Optimality Gap	
b^1	b^2	b^3	LB	DP	Pri	DA	AF	Search	DP-LB	Search-DP
2	2	2	8.7	10.3	12.1	10.8	10.9	10.9	15.4%	5.2%
3	3	3	10.8	12.6	15.2	13.2	13.4	13.4	14.7%	5.6%
4	4	4	12.2	14.2	17.5	14.9	15.2	15.1	14.2%	6.2%
6	6	6	14.0	16.3	20.6	17.3	17.5	17.5	14.1%	6.9%
10	10	10	15.7	18.8	24.8	20.5	20.6	20.7	16.4%	8.8%
12	12	12	16.2	19.8	26.4	21.5	21.6	21.8	17.9%	9.4%
15	15	15	16.8	20.9	28.4	22.7	23.1	23.0	19.7%	9.0%

Table 12: Cost behavior of the proposed heuristic and the lower bound against the optimal solution. In this problem, the holding costs and the demands are asymmetric. The symmetric penalty costs increase progressively down the column.

In Table 12, we find that the Lower Bound is much weaker (thus strengthening the results of the large numerical experiments reported in the previous sections). In general, both our heuristic and the lower bound grow weaker as the overall backorder costs of all products increase. However, we find that that our optimal policy is better than the other heuristics in the literature, and our costs are closer to the optimal cost solution from the Dynamic Program, deviating only between

[4]To be more precise, we solve the finite horizon un-discounted dynamic program with a horizon length of 50 periods and take the minimum (over all starting states) average cost as a proxy for the optimal infinite horizon, average cost. This is a legitimate proxy in the sense that it can easily be shown that the cost we get is a lower bound on the optimal infinite horizon, average cost. Thus the optimality gap that we report is a numerically close but a conservative estimate of the true optimality gap. In other words, our approach does not under-report the gap. The choice of fifty periods as the horizon length was based on our observation that our computed cost converges in the first two decimals, for a horizon of this length in our examples.

$5\% - 9\%$. So far, we studied the effect of symmetric penalty costs as they grew in cost. In Table 13, we study the effect of varying demand asymmetry for fixed symmetric backorder costs.

$(k^1, k^2, k^3) = (3, 4, 3)$ $\mathbf{b} = (10, 10, 10)$ $\mathbf{h} = (1, 1, 1)$			Costs						% Optimality Gap	
λ^1	λ^2	λ^3	LB	DP	Pri	DA	AF	Search	DP-LB	Search-DP
1	1	1.1	16.7	23.3	36.9	29.3	31.9	29.9	28.2%	22.2%
1	1	1.2	16.5	21.7	32.0	25.5	27.2	26.0	24.3%	16.4%
1	1	1.3	16.3	20.5	28.7	23.1	23.7	23.5	20.8%	12.7%
1	1	1.4	16.0	19.6	26.5	21.7	21.8	21.9	18.5%	10.4%
1	1	1.5	15.7	18.8	24.8	20.5	20.6	20.7	16.4%	8.8%

Table 13: Cost behavior of the proposed heuristic and the lower bound against the optimal solution. In this problem, the holding costs and the demands are asymmetric. The symmetric penalty costs are held at 10 and the the asymmetry of the demand (through variability) is increased progressively.

In Table 13, we fixed the backorder costs at 10 for all items and varied the asymmetric demand, holding the variance of product 1 at the smallest value, and product 3 at the highest value. We notice in this case, both the lower bound and our policy perform *better* as the demand asymmetry increases. This underlines that the performance of our heuristic is particularly beneficial in the multi-product cases that are most likely to occur. In particular the performance of our policy improves from 22% to 8.8% as the demand becomes more asymmetric and the effective capacity tightens.

$(k^1, k^2, k^3) = (3, 4, 3)$ $\mathbf{b} = (15, 6, 3)$ $\mathbf{h} = (1, 1, 1)$			Costs						% Optimality Gap	
λ^1	λ^2	λ^3	LB	DP	Pri	DA	AF	Search	DP-LB	Search-DP
1	1	1.1	14.9	19.0	25.2	27.0	23.2	22.9	21.9%	16.8%
1	1	1.2	14.7	18.1	22.6	23.6	20.4	20.5	18.7%	11.9%
1	1	1.3	14.4	17.3	20.8	21.4	18.8	19.0	16.4%	8.8%
1	1	1.4	14.3	16.6	19.6	19.6	17.6	17.9	14.4%	6.9%
1	1	1.5	14.1	16.2	18.7	18.5	17.0	17.1	12.7%	5.4%

Table 14: Cost behavior of the proposed heuristic and the lower bound against the optimal solution. In this problem, the holding costs and the demands are asymmetric. The asymmetry of the demands (through variability) is increased progressively.

In Table 14, we repeat the same scheme in Table 13, except that the penalty costs are also made asymmetric. Thus backorder costs, mean demands, and standard deviations of the demand distributions are all asymmetric in this set of experiments. We note that the relative performance of our heuristics and the lower bound have improved. Comparing the corresponding rows in Tables 13 and 14, we find that the introduction of asymmetric penalty costs has improved the performance of the heuristic has improved the performance of the heuristic and the lower bound. In fact, the cost performance of the Search improves by about 5% to 7% for the same set of demand distributions. In addition, the performance of our heuristic progressively improves as the demands become more asymmetric for the same set of backorder costs of the products. This is observed from decreasing performance gap (from 16.8% to 5.4%) of our Search policy compared to the optimal cost derived from the Dynamic Program.

$(k^1,k^2,k^3)=(3,4,3)$ $\mathbf{b}=(6,3,1)$ $\mathbf{h}=(1,1,1)$			Costs						% Optimality Gap	
λ^1	λ^2	λ^3	LB	DP	Pri	DA	AF	Search	DP-LB	Search-DP
1	1	1.1	11.1	13.2	15.7	19.1	14.9	15.1	16.4%	12.5%
1	1	1.2	10.9	12.8	14.5	16.5	13.7	14.0	14.2%	8.8%
1	1	1.3	10.9	12.4	13.7	15.1	13.1	13.2	12.3%	6.5%
1	1	1.4	10.8	12.1	13.1	14.3	12.7	12.7	10.7%	5.0%
1	1	1.5	10.7	11.8	12.7	13.7	12.2	12.3	9.7%	4.3%

Table 15: Cost behavior of the proposed heuristic and the lower bound against the optimal solution. In this problem, the holding costs are symmetric. The penalty costs and the demands are asymmetric. The asymmetry of the demands (through variability) is increased progressively.

In Table 15, we tested the performance of our heuristics for another set of backorder costs, increasing the ratio of the highest backorder cost to the lowest backorder cost. The introduction of more asymmetry in the backorder costs, further improves the performance of our policy. In fact, there is a performance gain ranging from 1% to 4.3%.

To summarize, the performance of our Search policy improves with respect to the optimal cost, (i) when the capacity constraints become tighter (ii) as demand variability increases, (iii) as demands become more asymmetric, and finally (iv) as the holding and penalty costs become more asymmetric.

9.11 Inference from the Computational Study

- When all product attributes (i.e., demand and costs) are symmetric, our policy and the DA heuristic are optimal. The AF heuristic is sub-optimal.
- When there is ample capacity (utilization is low), all heuristics perform comparably.
- In virtually all of the problem instances we computed, our policy *significantly outperforms* all the extant heuristics.

  The only instances in which the heuristics have comparable performances are those instances with low capacity utilization. Even in such instances, the cost of our Search policy is within 0.5% of the cost of the DA heuristic.
- As the capacity gets tighter, our heuristic consistently outperforms other heuristics.

- Our Search policy performs significantly better than priority heuristic in all cases.

- Our policy consistently outperforms the AF heuristic. The search policy particularly outperforms the AF heuristic by as much as 13% when the products are asymmetric and capacities are tight.

- Our performance is significantly better than DA heuristic, except for few symmetric cases with low utilization when the two heuristics are comparable. When the products and demands are asymmetric, it is possible that our heuristic gains more than 40% in costs.

- In general, the more asymmetric the costs are, or the more asymmetric the demands are, or the tighter the capacity is, the better the Search policy is (compared to all extant heuristics). In fact, the optimality gap between the DP and our heuristic is reduced when the costs and demands are highly asymmetric.

To wit, our policies generally outperform the existing policies in literature, especially as the available capacity becomes scarce, and the product demand and costs become asymmetric. As discussed in several Operations Management textbooks, firms find flexibility is most valuable when the available capacity is scarce and expensive, and when utilization is high. In many such cases, products also vary in their demand distributions and costs. Given that the optimal policies for allocating capacities remain unknown, the need for managers to have simple implementable schemes is paramount. To implement our allocation policies, firms just need to examine their current shortfalls to determine the allocation of capacity amongst different products. In sum, our policies perform well in cases in which flexible capacity is most valuable, have the theoretical appeal of being asymptotically optimal at high service levels and also at high utilization levels, and the practical appeal of being simple and intuitive to implement.

10 Concluding Remarks

We have provided an intuitive and easily implementable policy for a firm to manage finite flexible capacity under periodic review. We note that our policy outperforms the existing approaches in the literature. Moreover, our policy has some theoretical underpinning in that our policy class is asymptotically optimal under the two different regimes represented by high service levels (i.e., high penalty costs) and heavy traffic (i.e., tight capacity).

Nevertheless, there are several theoretically challenging questions that are left unanswered. Perhaps the most important among them involves understanding the structure of the optimal policy when products are asymmetric. Very little is known in this respect. These issues are generally present when one deals with multiple products in a dynamic setting. Recent progress on multi-product inventory problems with substitution by Song and Xue (2007) may be useful in advancing our knowledge on this issue. We leave such musings to future research.

Appendix

Proof of Lemma 2

Observe that $V_1^n = 0$ for all n by assumption. Therefore, $V_1 = 0$. We now proceed to the analysis of the aggregate shortfall in period $t+1$ for any $t \geq 0$.

The aggregate opening shortfall in period $t+1$ is $V_t + D_t$. If this is smaller than κ, all inventory levels are raised to the respective base-stock levels and thus V_{t+1}^n is zero for each n. Thus, $V_{t+1} = 0$ in this case. If $V_t + D_t > \kappa$, the entire production capacity is used, i.e. $\sum_{n=1}^{N} q_{t+1}^n = \kappa$ and none of the inventory levels exceeds the corresponding base-stock level. So, the new aggregate shortfall, V_{t+1} equals $\sum_{n=1}^{N}(V_t^n + D_t^n - q_{t+1}^n)$ which can be written as $V_t + D_t - \kappa$. Combining the two cases, we get

$$V_{t+1} = (V_t + D_t - \kappa)^+ \ .$$

Notice that this is the same as the recursion representing the evolution of the waiting time in a G/D/1 queue with a deterministic inter-arrival time of κ and a service time distribution represented by the random variable D. Statements (i) - (iii) follow jointly from the stability condition $E[D] < \kappa$ and induction arguments.

Proof of Lemma 3

Note that $V_1^n = 0$ for all n by assumption. This establishes statement (i) for $t = 1$. Notice that in Π_{BS-B}, for any $\boldsymbol{\alpha}$, the allocation rule makes production decisions which can be computed by knowing the opening shortfalls in a period. So, the ending shortfalls in period t, for any $t > 1$, depend exclusively on the ending shortfalls in period $t-1$, the demands in period $t-1$ and κ. This proves statement (i) for any t. Note that any weighted balancing allocation rule, i.e. any allocation rule in Π_{BS-B}, is monotone with respect to the shortfalls in the sense that if the shortfall of a product is perturbed non-negatively in period t, the shortfall of that product is perturbed non-negatively in period $t+1$. Then, (ii) follows from Lemma 1 of Loynes (1962) and the remark following that lemma. Statement (iii) follows directly from (i) and (ii).

Proof of Lemma 4

Consider the subclass of policies mentioned in the statement of the lemma. Once the base-stock vector $\mathbf{S}$ is chosen for a policy within this class, the policy is entirely specified. The long run average cost of this policy is

$$\sum_{n=1}^{N} E[h^n \cdot (S^n - V^{\boldsymbol{\alpha},n} - D^n)^+ \ + \ b^n \cdot (V^{\boldsymbol{\alpha},n} + D^n - S^n)^+] \ .$$

Since the distribution of $\mathbf{V}^{\boldsymbol{\alpha}}$ does not depend on the base-stock vector, the expression above is separable in $(S^1, \ldots, S^N)$; thus, the optimal value of S^n is simply the minimizer of the "newsvendor-type" expression within the summation above. The desired result is immediate.

Proof of Theorem 5

By assumption $V_1^{\mathbf{1},n} = 0$ for all n. This establishes statement (i) for $t = 1$. Notice that the symmetric allocation rule is a symmetric allocation rule. This implies that if the distribution of $\mathbf{V}_t^{\mathbf{1}}$ is symmetric across n for some t and the distribution of D_t^n is also symmetric across n, then the distribution of $\mathbf{V}_{t+1}^{\mathbf{1}}$ will also be symmetric. Statement (i) follows for all t by induction. Statement (ii) is a direct consequence of statement (i).

Proof of Theorem 6

Lemma 4 establishes the optimality of the base-stock vector $\mathbf{S}^{\mathbf{1}*}$ for policies in Π_{BS-B} that use the weight vector $\mathbf{1}$. It remains to show that the policy in Π_{BS-B} defined by the base-stock vector $\mathbf{S}^{\mathbf{1}*}$ and the weight vector $\mathbf{1}$ is an optimal policy when all policies in Π are considered.

Let us first consider the finite horizon discounted cost problem with a discount factor $\gamma \in (0, 1]$ and a planning horizon of T periods, that is, the problem of minimizing $E[\sum_{t=1}^{T} \gamma^t \cdot C_t]$ over Π. This finite horizon dynamic program can be represented through the cost-to-go functions $\{f_{t,T}^{\gamma} : t = 1, \ldots, T\}$ as follows:

$$
\begin{aligned}
f_{t,T}^{\gamma}(\mathbf{x}) \quad = \quad & \min_{\mathbf{y}} \; \sum_{n=1}^{N} \left(h^n \cdot E[(y^n - D^n)^+] \; + \; b^n \cdot E[(D^n - y^n)^+] \right) + \gamma \cdot E[f_{t+1,T}^{\gamma}(\mathbf{y} - \mathbf{D})] \\
& \text{s.t.} \quad \mathbf{y} \; \geq \; \mathbf{x} \quad \text{and} \quad \sum_{n=1}^{N} y^n \leq \sum_{n=1}^{N} x^n + \kappa \; ,
\end{aligned}
$$

where $f_{T+1,T}^{\gamma}(\mathbf{x}) := 0$ for all $\mathbf{x}$.

It is fairly easy to show using induction that under Assumption 1, the function $f_{t,T}^{\gamma}$ is convex and symmetric. Using standard dynamic programming arguments, we can establish the pointwise convergence of the finite horizon cost-to-go functions $\{f_{1,T}^{\gamma}(\mathbf{x})\}$ to $\{f^{\gamma}(\mathbf{x})\}$ the cost-to-go function of the infinite horizon, discounted cost dynamic program (defined for $\gamma \in (0, 1)$) represented below:

$$
f^{\gamma}(\mathbf{x}) \quad = \quad \min_{\mathbf{y}} \; g^{\gamma}(\mathbf{y}) \tag{4}
$$

$$
\text{s.t.} \quad \mathbf{y} \; \geq \; \mathbf{x} \quad \text{and} \quad \sum_{n=1}^{N} y^n \leq \sum_{n=1}^{N} x^n + \kappa \; , \tag{5}
$$

where

$$g^{\gamma}(\mathbf{y}) = \sum_{n=1}^{N} \left(h^n \cdot E[(y^n - D^n)^+] \ + \ b^n \cdot E[(D^n - y^n)^+] \right) + \gamma \cdot E[f^{\gamma}(\mathbf{y} - \mathbf{D})] \ .$$

The infinite horizon discounted cost optimal policy is defined by a selector $\mathbf{y}^{\gamma *}(\mathbf{x})$ such that for every $\mathbf{x}$, the vector $\mathbf{y}^{\gamma *}(\mathbf{x})$ is a solution to the above minimization problem. The convergence of $\{f^{\gamma}_{1,T}(\mathbf{x})\}$ to $\{f^{\gamma}(\mathbf{x})\}$ ensures that g^{γ} is also convex and symmetric. The convexity and symmetry of g^{γ} implies the existence of a vector $\mathbf{S}^{\gamma *}$ such that (a) it minimizes $g^{\gamma}(\mathbf{y})$ and (b) all its components are identical; let us denote this identical base-stock value for all components as $S^{\gamma *}$.

Next, we claim that the symmetric allocation rule applied in combination with the base-stock vector $\mathbf{S}^{\gamma *}$ is an optimal policy for the infinite horizon, discounted cost problem defined in (4)-(5) when $\mathbf{x} \leq \mathbf{S}^{\gamma *}$. There are two cases to study. The first case is the following: $\mathbf{x}$ is such that $\mathbf{x} \leq \mathbf{S}^{\gamma *}$ and $\sum_{n=1}^{N} x^n + \kappa \geq \sum_{n=1}^{N} S^{\gamma *}$. From such an inventory state $\mathbf{x}$, before ordering, we know that it is feasible to reach $\mathbf{S}^{\gamma *}$ after ordering. Moreover, this action is optimal since $\mathbf{S}^{\gamma *}$ minimizes $g^{\gamma}(\mathbf{y})$. The second case is the following: $\mathbf{x}$ is such that $\mathbf{x} \leq \mathbf{S}^{\gamma *}$ and $\sum_{n=1}^{N} x^n + \kappa < \sum_{n=1}^{N} S^{\gamma *}$. Convexity of the function $g^{\gamma}(\mathbf{y})$ ensures that there is an optimal solution such that $\sum_{n=1}^{N} y^n = \sum_{n=1}^{N} x^n + \kappa$. Moreover, it is easy to show the following inequality by using the convexity and symmetry of $g^{\gamma}(\cdot)$: Let $\mathbf{y}$ and $\tilde{\mathbf{y}}$ be two vectors such that $\sum_{n=1}^{N} y^n = \sum_{n=1}^{N} \tilde{y}^n = N \cdot \overline{y}$ for some constant $\overline{y}$, and such that, for every n, $y^n \in [\tilde{y}^n, \overline{y}]$ if $\tilde{y}^n \leq \overline{y}$ and $y^n \in [\overline{y}, \tilde{y}^n]$ if $\tilde{y}^n \geq \overline{y}$. Then, $g^{\gamma}(\tilde{\mathbf{y}}) \geq g^{\gamma}(\mathbf{y})$. In words, this inequality asserts that, given two inventory vectors $\tilde{\mathbf{y}}$ and $\mathbf{y}$ which are equal in terms of their aggregate inventories, the vector $\mathbf{y}$ which is more symmetric (i.e. whose components are closer to the average $\overline{y}$) is preferable with respect to the cost function $g^{\gamma}(\cdot)$. Since all products have the same mean demands per period, notice that the policy of using the symmetric allocation rule and the base-stock vector $\mathbf{S}^{\gamma *}$ leads to the "most symmetric" vector $\mathbf{y}$ among all vectors such that (5) holds. Thus, the inventory vector $\mathbf{y}$ chosen by this policy is optimal over all $\mathbf{y}$ satisfying the constraints above. This completes the proof of the claim.

Let us now return to the infinite horizon average cost problem. Schäl (1993) shows that, under certain conditions, the sequence of infinite horizon discounted cost optimal policies converges to an infinite horizon average cost optimal policy as the discount factor γ approaches 1. Huh et al. (2010) refer to this convergence as the preservation property and verify Schäl's conditions for the single product capacitated problem. A straightforward extension of their analysis can be used to verify Schäl's conditions for our multi-product problem. The proof follows below.

Proof of the Fact that Our Problem Satisfies the Conditions Required by Schäl (1993)

We only verify Condition (B) of Schäl (1993) on the relative discounted function. (The other conditions are trivial to verify.) Equivalently, we need only to verify the following condition. Let $f_T(\pi, \mathbf{x})$ denote the expected cost incurred under policy π over a finite horizon consisting of T periods when the starting state is x. Then, given any pair of initial states $\mathbf{x}$ and $\mathbf{x}'$, and any

feasible policy π, we can produce another feasible policy π', such that for all
$\forall\ T \geq 1$,

$$f_T(\pi', \mathbf{x}') \leq f_T(\pi, \mathbf{x}) + \eta(\mathbf{x}', \mathbf{x})$$

holds for some function η independent of T. (The reader is referred to Schäl (1993) and Huh et al. (2010).)

Following Huh et al. (2010), we need to consider only $\mathbf{x}$ and $\mathbf{x}'$ that differ only in one component, say $x_1' \neq x_1$ where $\mathbf{x} = (x_1, x_2, \ldots, x_N)$ and $\mathbf{x}' = (x_1', x_2, \ldots, x_N)$ at the beginning of period 1.

To produce the required η, we look at the following two complementary cases.

Case 1: $x_1' \geq x_1$.
The description of the π' policy is exactly the same policy described in the Huh et al. (2010) for the single product capacitated production problem with zero lead times. Note that this policy is feasible, and produced the desired η.

Case 2: $x_1' < x_1$.
In this case, the π' policy is similar to Huh (2010) in that π' for products $2, \ldots, N$ imitates π for product 1 and uses all the residual capacity until the inventory levels of product 1 in both the systems are the same.

Now, noting that the residual capacity is a random variable that depends on the policy (π and π') and the demand realizations, we observed that this cases reduced to Huh et al. (2010) with Markov modulated capacities. This produces the desired η.

Thus the above mentioned convergence of discounted cost optimal policies to an average cost optimal policy holds in our case. This implies that there exists a vector $\mathbf{S}^*$, in which all components are identical, such that the symmetric allocation rule applied in combination with the base-stock vector $\mathbf{S}^*$ is an average cost optimal policy. Finally, we know from Lemma 4 that, within Π_{BS-B}, the optimal base-stock vector corresponding to the weight vector $\mathbf{1}$ is $\mathbf{S}^{\mathbf{1}*}$. Thus, $\mathbf{S}^{\mathbf{1}*}$ is a valid choice for $\mathbf{S}^*$; this completes the proof.

Proof of Lemma 7

Without loss of generality, we assume that the priority order $(1), (2), \ldots, (N)$ is $1, 2, \ldots, N$.
When the capacity is not binding (i.e. $\sum_j^N (W^j) \leq \kappa$), the shortfalls after ordering are zero under both rules (for any m). Thus, the statement holds for any m.

Similarly, if the shortfall before ordering, W^j, is zero for any j, then the shortfalls after ordering $V^{P,j}$ and $V^{\boldsymbol{\alpha}_m,j}$ are both zero. Thus, it is sufficient to consider the case where $\mathbf{W}$ is strictly positive in every component.

When capacity is binding, there exists some $k, 1 \leq k < N$ such that $\sum_1^k W^j \leq \kappa$ and $\sum_1^{k+1}(W^j) > \kappa$. Then, under the priority policy $V^{P,j} = 0\ \forall\ j = 0, \ldots, k$, $V^{P,k+1} = W^{k+1} + \kappa - \sum_1^k W^j$, and $V^{P,j} = W^{P,j}\ \forall\ j = k+2, \ldots, N$. Let us define $\beta = W^{k+1} + \kappa - \sum_1^k W^j$, i.e.

$\beta = V^{P,k+1}$.

Let M be large enough that $W^{k+2}/M^{k+1} \geq W^{k+3}/M^{k+2} \geq \ldots \geq W^N/M^{N-1}$. That is, $k+2$ is the product with the largest weighted shortfall before ordering among products $\{k+2,\ldots,N\}$.

Let $\tilde{\epsilon} \in (0, \epsilon/k)$ and let $\tilde{\epsilon} \leq \min\{W^1,\ldots,W^k,\beta/k\}$. Moreover, let M be large enough that $\tilde{\epsilon} \geq W^{k+1}/M^k$, $\tilde{\epsilon}/M \geq W^{k+1}/M^k$, ..., $\tilde{\epsilon}/M^{k-1} \geq W^{k+1}/M^k$ and $(\beta - k\cdot\tilde{\epsilon})/M^k \geq W^{k+1}/M^k$. (All the inequalities above except the first and the last are redundant - but we present them here for ease of verification of our next claim). These inequalities ensure that even if the first $k+1$ components of the shortfall vector before ordering were reduced to $(\tilde{\epsilon},\ldots,\tilde{\epsilon},\beta - k\cdot\tilde{\epsilon})$, the weighted balancing rule defined by the vector $\boldsymbol{\alpha}_m$ prefers to allocate the next incremental amount of capacity to the first $k+1$ products and not the products in $\{k+2,\ldots,N\}$.

It is now easy to verify that $\mathbf{V}^{\boldsymbol{\alpha}_m}$ satisfies the following inequalities for all $m \geq M$:
$V^{\boldsymbol{\alpha}_m,j} = W^j = V^{P,j}$ for all $j \in \{k+2, k+3, \ldots, N\}$,
$V^{\boldsymbol{\alpha}_m,j} \in [0,\tilde{\epsilon}] = [V^{P,j}, V^{P,j} + \tilde{\epsilon}]$ for all $j \in \{1,2,\ldots,k\}$, and
$V^{\boldsymbol{\alpha}_m,k+1} \in [\beta - k\cdot\tilde{\epsilon},\beta] = [V^{P,k+1} - k\cdot\tilde{\epsilon}, V^{P,k+1}]$. The proof of the lemma is complete from the fact that $\tilde{\epsilon} \leq \epsilon/k$.

Proof of Lemma 8

The first inequality is trivial to establish because the cost incurred by any policy in any period when the backorder costs are given by $\mathbf{b}$ exceed the corresponding quantity when all backorder costs are $min(\mathbf{b})$. The second inequality follows from the definition of $C^*(h,\mathbf{b})$ and $C^{\mathbf{1}*}(h,\mathbf{b})$ as the optimal cost over all policies and the cost of the optimal weighted balancing policy, respectively.

We now show the third inequality. From Theorem 6, we know that

$$C^{\mathbf{1}*}(h, avg(\mathbf{b})) = C^*(h, avg(\mathbf{b})) \ .$$

Observe that $C^{\mathbf{1}*}(h,\mathbf{b})$ is a constant with respect to permutations to $\mathbf{b}$ due to the symmetric demand assumption and the symmetric nature of the symmetric allocation rule. The average of all possible permutations of $\mathbf{b}$ is

$$avg(\mathbf{b}) \cdot (1,1,\ldots,1) \ .$$

Since the single period function is linear with respect to $\mathbf{b}$ for any given state and action, it is easy to show that, for any policy π, $C^{\pi}(h,\mathbf{b})$ is concave with respect to $\mathbf{b}$. (See Janakiraman and Seshadri (2008) for a formal statement and proof of a more general result.) This implies that

$$C^{\mathbf{1}*}(h,\mathbf{b}) \ \leq \ C^{\mathbf{1}*}(h, avg(\mathbf{b})) \ .$$

Recalling that $C^{\mathbf{1}*}(h, avg(\mathbf{b})) = C^*(h, avg(\mathbf{b}))$, we have

$$C^{\mathbf{1}*}(h, \mathbf{b}) \leq C^*(h, avg(\mathbf{b})) ,$$

which is the desired result.

Proof of Theorem 9

The first statement follows directly from Lemma 8. We now show the asymptotic result.

We begin the proof by stating a result from Huh et al. (2009) which we will use.

Lemma 13 (Huh et al. (2009))**.** *Let X be a random variable such that $\overline{M} = \sup\{x : P(X \leq x) < 1\}$ and $\lim_{x \uparrow \overline{M}} \frac{E[X-x \mid X>x]}{x} = 0$, where $\overline{M} \in \Re^+ \cup \{\infty\}$. Then,*

$$\lim_{\beta \to \infty} \left(\frac{L(h, \beta \cdot b', X)}{L(h, \beta \cdot b, X)} \right) = 1 \quad \textit{for all} \quad (h, b', b) \ .$$

Note that any probability distribution with an increasing failure rate (IFR) satisfies the assumption made in Lemma 13.

Next, we show that, for every product, the convolution of the steady state shortfall distribution and the demand distribution satisfies the assumption made in Lemma 13.

Lemma 14. *Under Assumption 2, the following statement holds for all $j \in \{1, \ldots, N\}$: The convolution of $V_\infty^{\mathbf{1},j}$ and D^j is unbounded, i.e. $P(V_\infty^{\mathbf{1},j} + D^j < x) < 1$ for all x. Moreover, if the common marginal distribution of the random variables $\{D^j\}$ is an IFR distribution, then,*

$$\lim_{x \to \infty} \frac{E[(V_\infty^{\mathbf{1},j} + D^j) - x \mid (V_\infty^{\mathbf{1},j} + D^j) > x]}{x} = 0 \quad \textit{for all } j \ .$$

Proof: Recall that we assume throughout the paper that $P(D > \kappa) > 0$. Since the steady state distribution of the aggregate shortfall V_∞ is the same as that of the waiting time in a G/D/1 queue, it is easy to verify that the random variable V_∞ is unbounded under the above assumption. Since demands for all the products have been assumed to be symmetric, we know from Theorem 5 that the distributions of all the random variables $V_\infty^{\mathbf{1},j}$ are identical. By definition, V_∞ has the same distribution as the convolution $\sum_{j=1}^N V_\infty^{\mathbf{1},j}$. Since V_∞ is unbounded, it follows that the random variables $V_\infty^{\mathbf{1},j}$ are also unbounded; this implies that the random variables $V_\infty^{\mathbf{1},j} + D^j$ are also unbounded. This completes the proof of the first part of the lemma. We now proceed to prove the second part.

It is well known (see Bryson and Siddiqui (1969)) that if a non-negative random variable X has a finite mean and an IFR distribution, then its mean residual life $E[X - x|x > x]$ is decreasing and therefore bounded by the mean, $E[X]$. Thus, for every j, we know that the mean residual life of D^j is bounded.

Next, we use conditional expectations and observe that

$$E[(V_\infty^{\mathbf{1},j} + D^j) - x \mid (V_\infty^{\mathbf{1},j} + D^j) > x] = E_{V_\infty^{\mathbf{1},j}}\left[E_{D^j}[v + D^j - x \mid D^j > x - v] | V_\infty^{\mathbf{1},j} = v\right] .$$

Dividing both sides by x and taking the limit as $x \to \infty$, we obtain

$$\begin{aligned}
&\lim_{x\to\infty} \frac{E[(V_\infty^{\mathbf{1},j} + D^j) - x \mid (V_\infty^{\mathbf{1},j} + D^j) > x]}{x} \\
&= \lim_{x\to\infty} \frac{E_{V_\infty^{\mathbf{1},j}}\left[E_{D^j}[v + D^j - x \mid D^j > x - v] | V_\infty^{\mathbf{1},j} = v\right]}{x} \\
&\le \lim_{x\to\infty} \frac{E_{V_\infty^{\mathbf{1},j}}\left[\left(1(x \ge v) \cdot E[D^j] + 1(x < v) \cdot (E[D^j] + v - x)\right) | V_\infty^{\mathbf{1},j} = v\right]}{x} ,
\end{aligned}$$

where $1(\cdot)$ is the indicator operator; the inequality follows from the fact that the mean residual life of D^j is bounded by its unconditional mean. The expression on the right side of the inequality can be bounded above by

$$\lim_{x\to\infty} \frac{E[D^j + V_\infty^{\mathbf{1},j}]}{x} = 0 \text{ because } E[D^j] < \infty \text{ and } E[V_\infty^{\mathbf{1},j}] < \infty \text{ (since } E[D] < \kappa \text{ by assumption).}$$

This proves the desired result. □

We know from (1) that

$$\begin{aligned}
C^*(h, avg(\mathbf{b})) &= N \cdot L(h, avg(\mathbf{b}), V_\infty^{\mathbf{1},1} + D^1) \text{ and} \\
C^*(h, min(\mathbf{b})) &= N \cdot L(h, min(\mathbf{b}), V_\infty^{\mathbf{1},1} + D^1) .
\end{aligned}$$

Therefore,

$$\left(\frac{C^{\mathbf{1}*}(h, \mathbf{b})}{C^*(h, \mathbf{b})}\right) \le \left(\frac{L(h, avg(\mathbf{b}), V_\infty^{\mathbf{1},1} + D^1)}{L(h, min(\mathbf{b}), V_\infty^{\mathbf{1},1} + D^1)}\right) .$$

The desired asymptotic result now follows directly from Lemma 13 and Lemma 14. This completes the proof of Theorem 9.

Proof of Lemma 12

Consider any non-anticipatory policy π. Let y_t^π denote the aggregate inventory level after ordering in period t, when this policy is followed. Similarly let $\mathbf{y}_t^\pi$ ($\mathbf{x}_t^\pi$) denote the vector of inventory levels after (before) ordering in period t and let C_t^π be the cost incurred in that period. Thus,

$E[C_t^\pi] = \sum_{n=1}^N G^n(y_t^{\pi,n})$. Therefore, we know from the definition of F_1 that

$$E[C_t^\pi] \geq F_1(y_t^\pi) \ .$$

This implies that

$$\inf_{\pi\in\Pi} \lim_{T\to\infty} \sup \frac{E\left[\sum_{t=1}^T C_t^\pi\right]}{T} \geq \inf_{\pi\in\Pi} \lim_{T\to\infty} \sup \frac{E\left[\sum_{t=1}^T F_1(y_t^\pi)\right]}{T} \ .$$

Note that Π is the class of non-anticipatory policies satisfying the constraints $\mathbf{y_t^\pi} \geq \mathbf{x_t^\pi}$ and $\sum_{n=1}^N y_t^{\pi,n} \leq \sum_{n=1}^N x_t^{\pi,n} + \kappa$, in every period. Let Π' denote the larger class of policies which are non-anticipatory and require that only the second constraint, i.e. the capacity constraint, is satisfied in every period. This implies that

$$\inf_{\pi\in\Pi} \lim_{T\to\infty} \sup \frac{E\left[\sum_{t=1}^T F_1(y_t^\pi)\right]}{T} \geq \inf_{\pi\in\Pi'} \lim_{T\to\infty} \sup \frac{E\left[\sum_{t=1}^T F_1(y_t^\pi)\right]}{T} \ .$$

The quantity on the right side of the above inequality is nothing but the long run average optimal cost for a single product inventory problem with a capacity limit of κ and an expected single period cost $F_1(\cdot)$, which is a convex function. We know from Federgruen and Zipkin (1986) (for the case of countable state spaces, i.e. integer demands and inventories for example) and from Huh et al. (2010) (for the case of real-valued demands and inventories) that this optimal cost is achieved by a base-stock policy. Thus, we obtain

$$\inf_{\pi\in\Pi'} \lim_{T\to\infty} \sup \frac{E\left[\sum_{t=1}^T F_1(y_t^\pi)\right]}{T} = \min_S E[F_1(S - V_\infty)]$$

using the strong law of large numbers for Markov Chains (see Resnick (1992) for details). This leads to the desired result that

$$\inf_{\pi\in\Pi} \lim_{T\to\infty} \sup \frac{E\left[\sum_{t=1}^T C_t^\pi\right]}{T} \geq \min_S E[F_1(S - V_\infty)] \ .$$

Proof of Lemma 10

Proof. First, we observe that

$$C^*(\mathbf{h}, \mathbf{b}, \kappa) \geq C^*(h, b, \kappa) \ , \text{ if } \ 0 \ < \ h \ \leq \ h^j \ \forall \ j \ \text{ and } \ 0 \ < \ b \ < \ b^j \ \forall \ j \ , \tag{6}$$

where $C^*(h, b, \kappa)$ is the optimal cost of a system in which all products have the same holding cost h and the same backorder cost b. Thus, it suffices to show that, for any $h > 0$ and $b > 0$,

$$\lim_{\kappa\downarrow\mu} C^*(h, b, \kappa) = \infty \ .$$

Next, let us define $V_\infty(\kappa)$ as the steady state version of the aggregate shortfall process $\{V_t(\kappa)\}$ defined by the recursion $V_{t+1}(\kappa) = (V_t(\kappa) + D - \kappa)^+$ (recall that $D = \sum_{j=1}^N D^j$). We now claim that

$$C^*(h, b, \kappa) \geq \min_S h \cdot E[(S - V_\infty(\kappa) - D)^+] + b \cdot E[(D + V_\infty(\kappa) - S)^+] \; . \tag{7}$$

The proof of the claim is the following: Consider any feasible policy in the multi-product system. We can use this policy to construct a feasible policy in the "aggregate system" whose optimal long run average cost is represented on the right side of (7) such that the cost in the latter system (and therefore, the long run average cost) is smaller than that in the former system every period. This is done by ordering, in the latter system, the sum of the quantities ordered for all the products in the former system – the fact that the cost in the latter system is smaller in every period follows from the inequalities

$$\sum_{j=1}^N (x^j - d^j)^+ \geq \left(\sum_{j=1}^N (x^j - d^j)\right)^+ \quad \text{and} \quad \sum_{j=1}^N (d^j - x^j)^+ \geq \left(\sum_{j=1}^N (d^j - x^j)\right)^+ \; .$$

This proves the claim.

Thus, it only remains to show that

$$\lim_{\kappa \downarrow \mu} \min_S h \cdot E[(S - V_\infty(\kappa) - D)^+] + b \cdot E[(D + V_\infty(\kappa) - S)^+] = \infty.$$

To show this, we first note that replacing D by its expectation, μ, in the expression within the limit above we obtain a lower bound on that expression (this is a consequence of Jensen's inequality and the convexity of the function $(x)^+$). Letting $\tilde{S} = S - \mu$, it is sufficient to show that

$$\lim_{\kappa \downarrow \mu} \min_{\tilde{S}} h \cdot E[(\tilde{S} - V_\infty(\kappa))^+] + b \cdot E[(V_\infty(\kappa) - \tilde{S})^+] = \infty. \tag{8}$$

Next, observe that the recursion for $\{V_t(\kappa)\}$ is the same as that for the waiting time process for a $G/G/1$ queue in which the inter-arrival times are deterministic and equal to κ and the service time for the t^{th} customer is D_t. We know from Kingman (1962) that the distribution of the random variable $[\frac{(\kappa-\mu)}{\sigma^2}] \cdot V_\infty(\kappa)$ converges to an exponential distribution with mean 1/2, i.e.,

$$\lim_{\kappa \downarrow \mu} P\left(\frac{(\kappa - \mu)}{\sigma^2} \cdot V_\infty(\kappa) \geq z\right) = e^{-2z} \; , \quad \text{for all} \;\; z \; \geq \; 0 \; ,$$

where σ^2 is the variance of the aggregate demand D. We can verify using straight forward calculus that this implies that

$$\begin{aligned} & \lim_{\kappa \downarrow \mu} \min_{S'} h \cdot E\left[(S' - \frac{(\kappa-\mu)}{\sigma^2} \cdot V_\infty(\kappa))^+\right] + b \cdot E\left[(\frac{(\kappa-\mu)}{\sigma^2} \cdot V_\infty(\kappa) - S')^+\right] \\ = \; & (h/2) \cdot \ln\left(\frac{b+h}{h}\right). \end{aligned} \tag{9}$$

It is easy to verify that the desired inequality in (8) follows directly from (9). □

Proof of Theorem 11

Proof. The second statement follows directly from the first statement and Lemma 10. We proceed to show the first statement. Our plan is to find an upper bound on $C^P(\mathbf{h}, \mathbf{b}, \kappa)$ and a lower bound on $C^*(\mathbf{h}, \mathbf{b}, \kappa)$ and show that the difference between these bounds is finite for all κ.

Let $S(\kappa)$ be defined as $\arg\min_S h^N \cdot E[(S - D - V_\infty(\kappa))^+] + b^N \cdot E[(D + V_\infty(\kappa) - S)^+]$. Now, consider a policy π which uses the same priority rule as P but uses the following non-optimal base-stock levels:

$$S^j = 0 \text{ for all } j < N \text{ and } S^N = S(\kappa) .$$

Let $C^\pi(\mathbf{h}, \mathbf{b}, \kappa)$ ($C^{\pi,N}(h^N, b^N, \kappa)$) denote the long run average cost for the system (product N) under π given the respective parameters. Since P uses the optimal base-stock levels under the given priority allocation rule and π does not, we obtain the following relations:

$$\begin{aligned} C^P(\mathbf{h}, \mathbf{b}, \kappa) & \leq C^\pi(\mathbf{h}, \mathbf{b}, \kappa) \\ & = \sum_{j=1}^{N-1} b^j \cdot E[D^j + V_\infty^{P,j}] + C^{\pi,N}(h^N, b^N, \kappa) . \end{aligned} \tag{10}$$

The equality above follows from the fact that under π, there is never any inventory of products 1 through $N-1$ on hand and from the fact that the shortfall process under π is the same as that under P. From (10) and Assumption 3, it follows that

$$\begin{aligned} C^P(\mathbf{h}, \mathbf{b}, \kappa) \leq \; & b^1 \cdot \sum_{j=1}^{N-1} \left(E[D^j + V_\infty^{P,j}]\right) + C^{\pi,N}(h^N, b^N, \kappa) \\ = \; & b^1 \cdot \sum_{j=1}^{N-1} \left(E[D^j + V_\infty^{P,j}]\right) \\ & + h^N \cdot E[(S(\kappa) - D^N - V_\infty^{P,N}(\kappa))^+] + b^N \cdot E[(D^N + V_\infty^{P,N}(\kappa) - S(\kappa))^+] . \end{aligned} \tag{11}$$

The inequality above provides an upper bound on $C^P(\mathbf{h}, \mathbf{b}, \kappa)$.

Next, we proceed to identify a lower bound on $C^*(\mathbf{h}, \mathbf{b}, \kappa)$. By Assumption 3, we have

$$C^*(\mathbf{h}, \mathbf{b}, \kappa) \geq C^*(h^N, b^N, \kappa) \ . \tag{12}$$

Now, observe that $C^*(h^N, b^N, \kappa)$ is the optimal cost of a multi-product inventory system in which all products have identical holding and shortage costs. We have shown in the proof of Lemma 10 that this quantity exceeds the optimal cost of a single product inventory system with a holding cost h^N, backorder cost b^N, capacity κ and demand distribution D. That is,

$$\begin{aligned} C^*(h^N, b^N, \kappa) &\geq \min_S h^N \cdot E[(S - D - V_\infty(\kappa))^+] + b^N \cdot E[(D + V_\infty(\kappa) - S)^+] \ , \\ &= h^N \cdot E[(S(\kappa) - D - V_\infty(\kappa))^+] + b^N \cdot E[(D + V_\infty(\kappa) - S(\kappa))^+] \ . \end{aligned} \tag{13}$$

Let us define $V_\infty^{P,[1,N-1]}$ as $\sum_{j=1}^{N-1} V_\infty^{P,j}$ and $D^{[1,N-1]}$ as $\sum_{j=1}^{N-1} D^j$. Now, comparing (11) and (13) and using (12), we can write

$$\begin{aligned} & C^P(\mathbf{h}, \mathbf{b}, \kappa) - C^*(\mathbf{h}, \mathbf{b}, \kappa) \\ \leq \ & b^1 \cdot \sum_{j=1}^{N-1} \left(E[D^j + V_\infty^{P,j}] \right) + h^N \cdot E[D^{[1,N-1]} + V_\infty^{P,[1,N-1]}(\kappa)] \\ = \ & (b^1 + h^N) \cdot E[D^{[1,N-1]} + V_\infty^{P,[1,N-1]}(\kappa)] \ . \end{aligned} \tag{14}$$

Notice that $V_\infty^{P,[1,N-1]}(\kappa)$ is the steady state distribution of the stochastic process $\{V_t^{P,[1,N-1]}(\kappa)\}$ which evolves according to the recursion

$$V_{t+1}^{P,[1,N-1]}(\kappa) = \left(V_t^{P,[1,N-1]}(\kappa) + D_t^{[1,N-1]} - \kappa \right)^+ \ .$$

Since $\mu > E[D^{[1,N-1]}]$, it is easy to see that $\overline{V} := \lim_{\kappa \downarrow \mu} E[V_\infty^{P,[1,N-1]}(\kappa)]$ exists and is finite. Thus, we obtain

$$C^P(\mathbf{h}, \mathbf{b}, \kappa) - C^*(\mathbf{h}, \mathbf{b}, \kappa) \leq (b^1 + h^N) \cdot E[D^{[1,N-1]} + \overline{V}] \ < \ \infty \ \text{ for all } \kappa \ > \ \mu \ .$$

This completes the proof of the theorem. □

References

Aviv, Y., and A. Federgruen. 2001. Capacitated Multi-Item Inventory Systems with Random and Seasonally Fluctuating Demands: Implications for Postponement Strategies. *Management Science* 47 (4): 512.

Bryson, M., and M. M. Siddiqui. 1969. Some Criteria for Aging. *Journal of the American Statistical Association* 64 (328): 1472–1483.

DeCroix, G., and A. Arreola-Risa. 1998. Optimal Production and Inventory Policy for Multiple Products under Resource Constraints. *Management Science* 44 (7): 950–961.

Eppen, G., and L. Schrage. 1981. Centralized Ordering Policies in a Multi-Warehouse System with Lead Times and Random Demand. *Multi-Level Production / Inventory Systems: Theory and Practice, ed. L. B. Schwarz*:51–67.

Evans, R. 1967. Inventory Control of a Multiproduct System with a Limited Production Resource. *Naval Research Logistics Quarterly* 14:173–184.

Federgruen, A., and P. Zipkin. 1986. An Inventory Model with Limited Production Capacity and Uncertain Demands I. The Average-Cost Criterion. *Mathematics of Operations Research* 11 (2): 193–207.

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

Huh, W., G. Janakiraman, J. Muckstadt, and P. Rusmevichientong. 2009. Asymptotic Optimality of Order-up-to Policies in Lost Sales Inventory Systems. *Management Science* 55 (3): 404–420.

Huh, W. T., G. Janakiraman, and M. Nagarajan. 2010. Average Cost Inventory Models: An Analysis Using a Vanishing Discount Approach. *Operations Research (forthcoming)*.

Janakiraman, G., and S. Seshadri. 2008. Parametric Concavity in Stochastic Dynamic Programs. Working Paper.

Kingman, J. 1962. On Queues in Heavy Traffic. *Journal of the Royal Statistical Society. Series B (Methodological)* 24 (2): 383–392.

Limbaugh, K. 2008. Honda's Flexible Plants Provide Edge. *Wall Street Journal* Sept 23.

Loynes, R. 1962. The stability of a queue with nonindependent inter-arrival and service. *Proc. Camb. Philos. Soc.* 58:497–520.

Pena-Perez, A., and P. Zipkin. 1997. Dynamic Scheduling Rules for a Multiproduct Make-To-Stock Queue. *Operations Research* 45 (6): 919–930.

Resnick, S. 1992. *Adventures in Stochastic Processes*. Birkhauser, Boston, MA.

Rubio, R., and L. Wein. 1996. Setting Base Stock Levels Using Product-form Queueing Networks. *Management Science* 42 (2): 259–268.

Schäl, M. 1993. Average Optimality in Dynamic Programming with General State Space. *Mathematics of Operations Research* 18 (1): 163–172.

Shaoxiang, C. 2004. The Optimality of Hedging Point Policies for Stochastic Two-Product Flexible Manufacturing Systems. *Operations Research* 52 (2): 312–322.

Song, J., and Z. Xue. 2007. Demand Management and Inventory Control for Substitutable Products. Working Paper (Duke University).

Tayur, S. 1992. Computing the Optimal Policy for Capacitated Inventory Models. *Communications in Statistics: Stochastic Models* 9:585–598.

Veatch, M. H., and L. M. Wein. 1996. Scheduling a Make-To-Stock Queue: Index Policies and Hedging Points. *Operations Research* 44 (4): 634–647.

Wein, L. 1992. Dynamic Scheduling of a Multiclass Make-To-Stock Queue. *Operations Research* 40 (4): 724–735.

Zheng, Y., and P. Zipkin. 1990. A queueing model to analyze the value of centralized inventory information. *Oper. Res* 38 (2): 296–307.

Zipkin, P. 1995. Performance Analysis of a Multi-Item Production-Inventory System under Alternative Policies. *Management Science* 41 (4): 690–703.