STUDIES OF INVENTORY CONTROL AND CAPACITY PLANNING WITH MULTIPLE SOURCES

A Dissertation
Presented to
The Academic Faculty

By

Frederick Craig Zahrn

In Partial Fulfillment
Of the Requirements for the Degree
Doctor of Philosophy in Industrial Engineering

Georgia Institute of Technology

August 2009

STUDIES OF INVENTORY CONTROL AND CAPACITY PLANNING WITH MULTIPLE SOURCES

Approved by:

Dr. Shi-Jie Deng, Advisor
School of Industrial and Systems Engineering
Georgia Institute of Technology

Dr. John H. Vande Vate, Advisor
School of Industrial and Systems Engineering
Georgia Institute of Technology

Dr. Hayriye Ayhan
School of Industrial and Systems Engineering
Georgia Institute of Technology

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

Dr. Anton J. Kleywegt
School of Industrial and Systems Engineering
Georgia Institute of Technology

Date Approved: July 6, 2009

ACKNOWLEDGMENTS

I thank my advisors, Dr. Shi-Jie Deng and Dr. John H. Vande Vate, for their guidance of my research. I am particularly indebted to them—Prof. Vande Vate especially—for technical and expository advice in the portion of the dissertation on capacity planning. I also thank Dr. R. Gary Parker for his support of my graduate studies.

TABLE OF CONTENTS

SUMMARY

This dissertation consists of two self-contained studies.

The first study, in the domain of stochastic inventory theory, addresses the structure of optimal ordering policies in a periodic review setting. We take multiple sources of a single product to imply an ordering cost function that is nondecreasing, piecewise linear, and convex. Our main contribution is a proof of the optimality of a finite generalized base stock policy under an average cost criterion. Our inventory model is formulated as a Markov decision process with complete observations. Orders are delivered immediately. Excess demand is fully backlogged, and the function describing holding and backlogging costs is convex. All parameters are stationary, and the random demands are independent and identically distributed across periods. The (known) distribution function is subject to mild assumptions along with the holding and backlogging cost function. Our proof uses a vanishing discount approach. We extend our results from a continuous environment to the case where demands and order quantities are integral.

The second study is in the area of capacity planning. Our overarching contribution is a relatively simple and fast solution approach for the fleet composition problem faced by a retail distribution firm, focusing on the context of a major beverage distributor. Vehicles to be included in the fleet may be of multiple sizes; we assume that spot transportation capacity will be available to supplement the fleet as needed. We aim to balance the fixed costs of the fleet against exposure to high variable costs due to reliance on spot capacity.

We propose a two-stage stochastic linear programming model with fixed recourse. The demand on a particular day in the planning horizon is described by the total quantity to be delivered and the total number of customers to visit. Thus, daily demand throughout the entire planning period is captured by a bivariate probability

distribution. We present an algorithm that efficiently generates a "definitive" collection of bases of the recourse program, facilitating rapid computation of the expected cost of a prospective fleet and its gradient. The equivalent convex program may then be solved by a standard gradient projection algorithm.

The two investigations making up this dissertation are united by a concern with multiple sources within their respective contexts. The first study posits multiple sources of a single product (for example, multiple suppliers or production technologies) within an inventory control context. The second study considers multiple sources of transportation capacity, in the form of vehicles of differing sizes and the alternative of fleet versus spot capacity. For each of the two investigations, our focus on multiple sources entails a richer analysis than that arising from a single-source setting.

CHAPTER 1: AVERAGE OPTIMAL CONTROL IN AN INVENTORY MODEL WITH MULTIPLE SOURCES

1.1 Introduction

Consider a periodic review inventory control context in which there are multiple suppliers or production technologies for a single product, each source being subject to a fixed short-term capacity. On the premise that the cheapest source should be used first, a convex and piecewise linear ordering cost function—associating each source with a constant marginal ordering cost—may be an appropriate model element. In this chapter, we investigate the structure of optimal ordering policies for a class of models with this type of ordering cost function. In defining optimal policies, our focus is on an average cost criterion. While this type of cost criterion is relatively familiar in the field of stochastic processes, it is often regarded as more technically forbidding in comparison to finite-horizon or discounted cost criteria. Thus, our method of analysis may be of some interest as well as the structural result.

Our inventory model is formulated as a Markov decision process with complete observations. The ordering cost for any period is a nondecreasing, piecewise linear, and convex function of the order quantity for a single product. Here we assume that there is no upper limit on the order quantity in a given period; the most expensive source is modeled as uncapacitated. We are also assuming that there is no significant fixed cost to be incurred by ordering a positive quantity as against ordering nothing. Furthermore, we are assuming that all sources have nonnegative marginal cost. Orders are delivered immediately. Excess demand is fully backlogged, and the function describing holding and backlogging costs is convex. All parameters are stationary, and the random demands are independent and identically distributed across periods; the (known) distribution function is subject to mild assumptions along with the holding and backlogging cost function. We prove that a finite generalized base stock policy is optimal under a long-run average expected cost criterion. We focus on the case of a continuous state space in which demands and order quantities might take any nonnegative real value, and we extend our argument to the discrete case in which these quantities may take only nonnegative integral values.

Though convex ordering cost functions appeared early in the literature on stochastic inventory theory, and though there has been considerable effort directed at understanding optimal inventory control policies under average cost criteria, we are not aware of any prior work establishing our precise conclusions. Our result is not surprising, however, as finite generalized base stock policies have been claimed to be optimal in discounted cost settings when the ordering cost function is piecewise linear and convex. Furthermore, average cost results consistent with ours have been claimed for the special case when the convex cost function is composed of two linear pieces. Moreover, important work of Huh et al. (2008) aims to greatly facilitate proofs of optimal policy structures, under an average cost criterion, for a wide class of inventory control models very nearly encompassing ours (as well as models that are more complex than ours in fundamental ways).

A possibly unique element of our argument centers around an observation that, in a relaxed version of our model under a discounted cost criterion, the one-sided derivatives (with respect to the inventory level) of the optimal value functions are nonincreasing in the discount factor. This observation may be useful for further results, and it is not present in the paper by Huh et al. (2008). In its broad outlines, our method of proof is more familiar: it is based on the well-known vanishing discount strategy, and it incorporates a relaxation technique used before by Zheng (1991).

In Section 1.2, we review relevant literature. In Section 1.3, we precisely define our model and state our desired results, and we then summarize our proof. Section 1.4 contains our technical arguments. We conclude the chapter briefly in Section 1.5, and we give references in Section 1.6.

1.2 Literature review

In reviewing related literature, we focus on work featuring periodic review inventory models with a convex (and nonlinear) ordering cost function or with an average cost criterion. Broader coverage of the subject of inventory theory may be found in Veinott (1966), Porteus (1990), Zipkin (2000), and Porteus (2002). References on Markov decision process (MDP) models more generally, including models under average cost criteria, include Heyman and Sobel (1984), Puterman (1994), Arapostathis et al. (1993), Hernández-Lerma and Lasserre (1996), Sennott (1999), and Feinberg and Shwartz (2002).

1.2.1 Inventory control with a convex ordering cost function

Convex ordering cost functions appeared early in the literature on inventory control and production planning. Several studies in this area consider deterministic models, unlike our framework which features stochastic demand. For example, Veinott (1964) studies a production and inventory model with a convex ordering cost function (or rather a convex production cost function) in a finite-horizon setting with deterministic future demand; here, uncertainty is dealt with by sensitivity analysis. In a subsequent survey of inventory theory, Veinott (1966) discusses other work with convex ordering cost functions in a deterministic setting, much of which was published in the 1950s. A few additional references along these lines are given in Sethi et al. (2005, p. 11).

In a stochastic and dynamic setting, Karlin (1958) considers basic inventory control models featuring three types of ordering cost functions, associating each type of function with an optimal decision rule having a particular structure:

1. A *linear* ordering cost function is associated with what are now widely known as *base stock* rules, which have the form: order up to meet a target inventory level s^* when the current period's inventory level is below s^*; i.e., order the quantity $(s^* - I)$ if the current inventory level is $I < s^*$, and otherwise order

nothing.

2. An ordering cost function involving a fixed *set-up cost* incurred for all positive order quantities, in addition to a linear cost component, is associated with (s, S) rules: order up to a target level S^* when we see inventory below a critical level s^*, where $s^* \leq S^*$; i.e., order the quantity $(S^* - I)$ if the current inventory level is $I < s^*$, and otherwise order nothing.

3. A *convex* ordering cost function is associated with what have been called (in Porteus 1990) *generalized base stock* rules, which have the property that the order-up-to level is a nondecreasing function of the current inventory level, while the order quantity is a nonincreasing function of the current inventory level.

Here, Karlin's criterion for evaluating a given policy is the *discounted* expected cost incurred, which by nature diminishes the emphasis on the (perhaps very) long term. By contrast, we are concerned with the structure of optimal policies under an *average* cost criterion, which is intended to ignore the short-term, transient behavior of the system and focus on the steady state. Also notable is that Karlin assumes *strict* convexity of the ordering cost function, apparently making extensive further specification of optimal policies cumbersome in general. We instead assume a piecewise linear form that implies an intuitive and relatively simple optimal policy structure. Further results for stochastic inventory control with a strictly convex ordering cost function under a discounted cost criterion may be found in Bulinskaya (1967).

The case of a piecewise linear and convex ordering cost function is discussed in the survey of stochastic inventory theory by Porteus (1990). He describes the structure of a *finite generalized base stock* rule by reference to a hypothetical situation involving alternative production technologies. Each technology has a linear cost and a fixed per-period capacity—except the most expensive technology, which is uncapacitated.

This leads to a convex and piecewise linear ordering cost function, as in our setting, on the assumption that a particular technology is utilized only if all cheaper technologies are being used to capacity. The decision rule is then defined by a nonincreasing set of base stock levels corresponding to the technologies in increasing order of marginal cost. There may therefore be a range of inventory levels for which we do not utilize a given technology, though we utilize all cheaper technologies to capacity. Porteus asserts the optimality of a finite generalized base stock policy under a discounted expected cost criterion, but he offers no proof or reference for this proposition. The only optimality proof we have found that allows an ordering cost function with any number of linear pieces is in Bensoussan et al. (1983), under the finite-horizon total expected cost criterion. Unlike Bensoussan et al., we deal with an average cost criterion, and we also allow unbounded marginal holding and backlogging costs—as well as a marginal ordering cost equal to zero for the cheapest source.

Considerable attention has been given to stochastic models with piecewise linear and convex ordering cost functions for the special case with two linear pieces. Sobel (1970) studies such a model in which the location of the kink in the function is chosen at the outset and thereafter is fixed from period to period. He argues for the optimality of a finite generalized base stock policy when the location of the kink is given—under discounted cost criteria, and also under an average cost criterion for the case of discrete demand. His finite-horizon results are used in Kleindorfer and Kunreuther (1978). Henig et al. (1997) consider a similar model with an ordering cost function equal to zero for up to R units, with a cost of c per additional unit. (This is in the context of supply and transportation contracts, in which the available volume R per period may be specified by a long-term agreement; like Sobel, they aim to optimally choose R.) They argue for the optimality of a finite generalized base stock policy specified by two base stock levels (and the parameter R) with respect to the discounted expected cost. They conjecture that the same type of policy is optimal with respect to an average cost criterion. This conjecture is repeated in Geunes (1999) and in Serel et al. (2001). Yang et al. (2005) consider a model with capacitated "in-house" production and an uncapacitated "outsourcing" option. In their default setting, the capacity level fluctuates randomly and there is a fixed cost of outsourcing as well as a per-unit cost. These complexities aside, they offer in effect an argument for average optimality, in a discrete setting, of a finite generalized base stock policy when the ordering cost function is nondecreasing, convex, and piecewise linear with two linear pieces. Our argument accommodates a cost function with any (finite) number of linear pieces, and we also allow non-discrete demand distributions.

Stochastic production smoothing models such as that of Beckmann (1961) are also relevant here. Beckmann's model incorporates a linear cost of production along with per-unit costs of increasing and decreasing the production level relative to the level chosen for the preceding period. Thus we have in effect an inventory model in which the ordering cost function is convex and piecewise linear with two linear pieces, such that the location of the kink in the function may change from period to period. It is even allowed that the function may decrease up to the kink, signifying that it is very costly to reduce production. (This characteristic is also allowed in Sobel 1970.) Later work on production smoothing models in this vein includes Sobel (1969) and Sobel (1971). In our model, by contrast, the ordering cost function is nondecreasing and fixed across periods, while we allow any number of linear pieces.

Huh et al. (2008) is a study aimed at developing a framework under which the optimality of particular inventory control policy structures may be immediately extended from finite-horizon to infinite-horizon (including average cost) settings. Their framework includes the possibility of a nondecreasing, piecewise linear and convex ordering cost, though they do not discuss the specific structure of optimal policies for this situation. Their framework also requires bounds on the marginal holding and backlogging costs, whereas our arguments do not require such bounds. Furthermore, in the course of our technical argument we offer some structural insight for our problem that may be of wider use and that is not present in their paper.

1.2.2 Inventory control under average cost criteria

As observed in the general treatments of Markov decision processes in Heyman and Sobel (1984, p. 171) and Puterman (1994, p. 331), an average cost criterion may be appropriate for modeling systems in which decisions are made frequently. These authors also note the complexity of technical analysis under average cost criteria; problematic characteristics of inventory models in particular include the possibility of state spaces and feasible action sets that are unbounded (and perhaps continuous), as well as unbounded cost functions.

In Section 1.2.1 above, we have discussed work on inventory control under an average cost criterion when the ordering cost function is convex (and nonlinear). Here, we mention work on stochastic inventory control models with other types of ordering cost functions.

The case of an inventory model with a linear ordering cost function (as in the first case examined in Karlin 1958, mentioned above) is studied in Vega-Amaya and Montes-de-Oca (1998). They argue that the base stock structure remains optimal under an average cost criterion. Our model encompasses linear (nondecreasing) ordering cost functions, but we assume backlogging of unmet demand while they assume lost sales.

Literature on average optimality in models with a fixed cost of ordering in addition to a linear component (as in the second case examined in Karlin 1958) is briefly reviewed in Feinberg and Lewis (2006). They cite several papers arguing that the (s, S) structure remains optimal. One such paper of particular significance for us is Zheng (1991), which employs a relaxation technique that we use as well. Other studies cited include Iglehart (1963), Veinott and Wagner (1965), Beyer and Sethi

(1999), and Chen and Simchi-Levi (2004). The forthcoming volume of Beyer et al. (2009) also promises to discuss this case.

The work by Huh et al. (2008) mentioned above stands to establish average optimality of policy structures, for a wide class of cost functions and other model parameters, whenever the structures are known to be optimal in a finite-horizon setting.

1.3 Formal problem statement and proof strategy

We define our inventory model as a Markov decision process. Let $I \in \mathbb{R}$ represent the inventory level at a particular decision point. Excess demand is backlogged, so this quantity may be negative. Knowing I, we must choose the order quantity for the current period; we order so as to bring the inventory level up to some level $Z \geq I$, delivery being instantaneous. The cost of the order is $C(Z-I)$ according to a function $C : \mathbb{R}^+ \to \mathbb{R}^+$. After our decision is made, the demand level X for the current period is revealed. The demand level in each period is a nonnegative, real-valued random variable with known distribution; it is independent of the history of the process, which here consists of all past inventory levels and order-up-to levels, including those of the current period. In particular, the demand levels are independent and identically distributed across periods. The period inventory level net of demand is $(Z - X)$. An inventory holding cost (or backlogging penalty cost, if net inventory is negative) for the current period is incurred in the amount of $L(Z - X)$ according to a function $L : \mathbb{R} \to \mathbb{R}^+$. The quantity $(Z - X)$ will be observed as the inventory level at the next decision point.

In our setting, $C(Q)$ is a convex, nondecreasing, piecewise linear function of the order quantity $Q = (Z - I)$. The m sources of our product are numbered in order of marginal cost, with the cheapest source first and the most expensive source last. (We assume that m is finite.) Without loss of generality, we assume that no two sources have identical marginal cost, and that each source has nonzero capacity. The ordering cost function is defined by parameters as follows:

$$0 = c^0 \leq c^1 < c^2 < \ldots < c^{m-1} < c^m < +\infty$$

$$0 = r^0 < r^1 < r^2 < \ldots < r^{m-1} < r^m = +\infty$$

Here, r^i is the cumulative capacity of the first i sources. On the premise that the cheapest source(s) should be used first, the marginal cost of increasing the order quantity is c^i when the order quantity satisfies $r^{i-1} \leq Q < r^i$. For order quantity Q, then, we have:

$$C(Q) = \sum_{i=1}^{m} (c^i - c^{i-1})(Q - r^{i-1})^+$$

We impose the following additional conditions on the holding and backlogging cost function and the distribution of demand variables:

Assumption 1. $L : \mathbb{R} \to \mathbb{R}^+$ is a convex function.

Assumption 2. $L(I) \to +\infty$ as $|I| \to +\infty$.

Assumption 3. $\mathbb{E}[L(Z - X)] < +\infty$ for all $Z \in \mathbb{R}$.

Assumption 4. $P(X = 0) < 1$.

The property $\mathbb{E}[X] < +\infty$ is implicit in Assumptions 1–3, and we prove this later on. Our assumption of convexity of the holding and backlogging cost function is fairly typical—encompassing the straightforward case of two linear functions that are zero when the net inventory level is zero—though it is not universal. Given convexity, our additional assumption that this function approaches $+\infty$ for increasing or decreasing inventory levels is reasonable: the violation of this assumption would imply an incentive to pursue arbitrarily great inventory or backlog quantities. Furthermore, given Assumptions 1 and 2 and our optimality criterion, our assumption that $L \geq 0$ is without loss of generality. Similarly, it is without loss of generality that we assume $C(0) = 0$. Assumption 4 is imposed to eliminate a relatively trivial special case. Finally, note that our framework allows mixed demand distributions as well as distributions having a probability density or mass function.

Given the characteristics of our model as just described, we seek to understand the structure of optimal control policies under an average cost criterion. We now define the particular criterion we adopt.

Definition 1. For a general Markov decision process, let $\phi(y, a)$ be the cost incurred in any given period when the state is observed to be y and the action a is taken. An admissible policy π^* is *average optimal* if for all admissible policies π and all initial states y_0, we have

$$\limsup_{n \to +\infty} \frac{1}{n+1} \mathbb{E}_{y_0}^{\pi^*}\left[\sum_{t=0}^{n} \phi(y_t, a_t)\right] \leq \limsup_{n \to +\infty} \frac{1}{n+1} \mathbb{E}_{y_0}^{\pi}\left[\sum_{t=0}^{n} \phi(y_t, a_t)\right]$$

Above, a *policy* prescribes the decision rule to be used at all decision points; a *decision rule* specifies an action at a particular decision point. (Our usage of these terms follows Puterman 1994 and is not universal.) An *admissible* policy consists of rules that are allowed to depend on any information in the history of the process; in our case, the history includes (in sequence) all inventory levels I observed and all order-up-to levels Z chosen. Furthermore, randomized selection of actions may be employed. However, an admissible policy may not use rules depending on demand levels not yet realized. There is a further requirement that all policies under consideration satisfy certain measurability requirements. In general (and beyond the issue of policies) we take measurability for granted in the course of our arguments; to the best of our knowledge, our constructions do not run afoul of such technical subtleties at the foundations of fully rigorous treatments of probabilistic models.

Definitions of average optimality similar to ours are given Heyman and Sobel (1984) and Arapostathis et al. (1993). Our definition is equivalent to the "lim inf" average optimality of Puterman (1994, p. 129, in the context of maximizing rewards).

Our use of the limit superior instead of the plain limit is reasonable because it is possible in our framework to specify an inventory control model and policy for which the desired limit does not exist. We will, however, find that a plain limit is attained by the policy we establish as average optimal. (There is also a related "sample path average cost" criterion which we do not discuss; see Arapostathis et al. 1993.)

Our primary goal is to show that there exists a finite generalized base stock policy that is average optimal for our inventory model as described. In particular, we will show that there is such an optimal policy that is expressible in terms of the parameters $\{r^0, \ldots, r^m\}$ along with certain critical values $\{s^0, \ldots, s^m\}$. The following definition makes this goal precise.

Definition 2. Given our model parameters $\{r^0, \ldots, r^m\}$, and given critical values

$$s^m \leq s^{m-1} \leq \ldots \leq s^2 \leq s^1 < s^0 = +\infty$$

the corresponding *finite generalized base stock policy* ("FGB policy") prescribes an order-up-to level as a function of the inventory level at any decision point:

- If $s^i - r^i \leq I < s^i - r^{i-1}$ (equivalently, $r^{i-1} < s^i - I \leq r^i$) for some $i \in \{1, \ldots, m\}$, choose $Z = s^i$. Here we are able to reach the ith base stock level by utilizing the ith source, while using the first $(i-1)$ sources to their capacities.

- If $s^i - r^{i-1} \leq I < s^{i-1} - r^{i-1}$ (equivalently, $s^i \leq I + r^{i-1} < s^{i-1}$) for some $i \in \{1, \ldots, m\}$, choose $Z = I + r^{i-1}$. We are in a range where we use the first $(i-1)$ sources to their capacities, but we do not use source i. Our use of the first $(i-1)$ sources alone brings us above the ith base stock level.

Note that the cases above are mutually exclusive and exhaustive. Also note that if $s^i = s^{i-1}$ for some i, the second type of condition above becomes vacuous for that i. We allow for the possibility that $s^i = -\infty$ for some $i \neq 0$, which would imply that source i and any more expensive sources are never utilized. However, we will establish that s^i is finite for all $i \neq 0$ characterizing our particular average optimal FGB policy. (Our definition may be compared with Bensoussan et al. 1983, p. 332, Porteus 1990, p. 622, and Liu and Esogbue 1999, p. 43.)

Our proof strategy involves studying a relaxed version of the inventory model in which the inventory level may be reduced by an arbitrary amount at no cost. Note that by expanding the domain of the function C to be $\mathbb{R}$, our given formula for $C(Q)$ remains valid. This relaxed model will be associated with a *relaxed* FGB policy structure, which we now define.

Definition 3. A *relaxed finite generalized base stock policy* ("rFGB policy") is the same as an FGB policy, except that for $s^1 - r^0 \leq I < s^0 - r^0$ (that is, $I \geq s^1$), we choose $Z = \min\{I, s^*\}$ instead of $Z = I + r^0 = I$. Here s^* is an additional parameter satisfying $s^1 \leq s^* < +\infty$.

In addition to the average cost criterion defined above, we will in the course of our arguments make use of (rather standard) finite- and infinite-horizon discounted cost optimality criteria for Markov decision processes, which we define below.

Definition 4. For an MDP, let $\phi(y, a)$ be the cost incurred in period t when the state is y and action a is taken. Let $\phi_n(y)$ be the cost incurred for final state y in the terminal period $n \in \mathbb{Z}^+$. Let $\beta \in (0, 1)$ be the discount factor. An admissible policy π^* is *discount optimal* if for all admissible policies π and all initial states y_0,

$$\mathbb{E}_{y_0}^{\pi^*}\left[\sum_{t=0}^{n-1} \beta^t \phi(y_t, a_t) + \beta^n \phi_n(y_n)\right] \leq \mathbb{E}_{y_0}^{\pi}\left[\sum_{t=0}^{n-1} \beta^t \phi(y_t, a_t) + \beta^n \phi_n(y_n)\right]$$

If we take $\beta = 1$, the above becomes a total expected cost criterion. We will find it convenient to reverse the numbering of periods, so that subscript '0' is associated with the terminal period.

Definition 5. For an MDP, let $\phi(y, a)$ be the cost incurred when the state is y and action a is taken. Let $\beta \in (0, 1)$ be the discount factor. An admissible policy π^* is *discount optimal* if for all admissible policies π and all initial states y_0,

$$\mathbb{E}_{y_0}^{\pi^*}\left[\sum_{t=0}^{+\infty} \beta^t \phi(y_t, a_t)\right] \leq \mathbb{E}_{y_0}^{\pi}\left[\sum_{t=0}^{+\infty} \beta^t \phi(y_t, a_t)\right]$$

In our case, $\phi(y, a) \geq 0$. By the monotone convergence theorem, then, it is equivalent for us to take the limits outside of the expectations in the definition above.

Our technical arguments run as follows. In Section 1.4.1, we establish some notation and several analytical results that we will use later. In Section 1.4.2, we begin our study of the relaxed model. We find that rFGB decision rules are discount optimal in a finite-horizon setting with discount factor β. Specifically, given $n \geq 1$ decision points remaining and terminal cost equal to zero, critical values $\{s_n^*, s_n^1, \ldots, s_n^m\}$ defined in terms of the optimal cost function (also called the "value function") $f_{n-1}(I)$, which is convex, correspond to an optimal rFGB decision rule. Upon further examination, we find that the one-sided derivatives of $f_{n-1}(I)$ with respect to I are nonincreasing in n as well as in the discount factor β. As a consequence, the critical values s_n^i are nondecreasing in n and in β (provided that we choose the greatest possible critical values at each step, in cases where there are multiple optimizers). Finally, we show that the critical values $\{s_n^*, s_n^1, \ldots, s_n^m\}$ are finite for sufficiently large n and β.

In Section 1.4.3, we turn to the infinite-horizon setting, considering the situation as $n \to +\infty$. We observe for the finite-horizon model that we may restrict attention to a compact interval of actions in any given state, and we also establish $f_n \nearrow f$. After seeing its role in a solution of the discounted cost optimality equation, this function f turns out to be the optimal cost function for the infinite-horizon setting. Based on the convex structure of f, we find that rFGB policies are discount optimal. We also establish that the one-sided derivatives of $f(I)$ with respect to I are nonincreasing in β, and that the critical values s^i derived from f are nondecreasing in β (again provided that we choose the greatest possible critical values).

In Section 1.4.4, we use the above results to construct and validate average optimal policies. In the infinite-horizon discounted setting, for β close to 1 we find that we may restrict attention to a compact interval of actions that does not vary with β. We also establish for a (convex) relative value function $\bar{f}$ that $\bar{f} \nearrow \bar{f}_1$ as $\beta \to 1$. These facts are used to show the role of $\bar{f}_1$, along with a scalar ρ, in solving the average cost optimality equation. We then argue that an rFGB policy with critical values s_1^i defined by reference to the convex function $\bar{f}_1$ is average optimal for the relaxed problem, having cost ρ. Finally, we show that this implies that an FGB policy (with the same critical values, absent s_1^*) must be average optimal for the unrelaxed model.

In Section 1.4.5, we aim to attain our secondary goal: extension of our main result to the discrete case in which demands and order quantities may take only nonnegative integral values. We show how our main argument in the preceding sections may be used to establish that an average optimal FGB policy exists in this case as well.

Our arguments are fairly self-contained, making use of many ideas from prior work. Some notable examples of prior literature that inspired our argument in various ways include: Heyman and Sobel (1984, Theorems 8-14 and 8-15 and proofs), Zheng (1991, Section 4), Arapostathis et al. (1993, Theorem 5.1 and proof), Puterman (1994, Theorems 6.2.2 and 8.10.7 and proofs), Henig et al. (1997, Lemma 1 and Theorem 1 and proofs), and Yang et al. (2005, Lemmas 7 and 8 and proofs). It is also clear that our argument has elements in common with Sobel (1970) and Schäl (1993, Lemma 1.2 and Proposition 1.3).

1.4 Technical arguments

1.4.1 Preliminary analytical results

In the rest of the chapter, we will denote left and right derivatives with '$-$' and '$+$' superscripts, respectively. In Section 1.4.2 and after, these derivatives will always be taken with respect to I or Z. We define the left and right derivatives so that they are equal when the conventional derivative exists. We take the existence of a one-sided derivative to mean that the corresponding limit is finite.

The following lemma is essentially the same as Theorem 5.1.3 in Webster (1994):

Lemma 1. If a finite-valued, convex function g is defined on an open interval of $\mathbb{R}$, then $g^-(x)$ and $g^+(x)$ exist for all x in this interval. Given $x_1 < x_2$ in this interval, we also have

$$g^-(x_1) \leq g^+(x_1) \leq \frac{g(x_2) - g(x_1)}{x_2 - x_1} \leq g^-(x_2) \leq g^+(x_2)$$

In particular, if $g : \mathbb{R} \to \mathbb{R}$ is convex then it is continuous. Sums, positive scalar multiples, and pointwise limits of convex functions are themselves convex. A convex, finite-valued function g achieves a local (and therefore global) minimum over $\mathbb{R}$ at x^* if and only if $g^-(x^*) \leq 0 \leq g^+(x^*)$. For a convex, finite-valued function g, we also define

$$g^-(-\infty) := \lim_{x \to -\infty} g^-(x)$$

$$g^+(+\infty) := \lim_{x \to +\infty} g^+(x)$$

and we note that these are monotone limits (by convexity) that may be infinite.

Lemma 2. Let $g : \mathbb{R} \to \mathbb{R}$ be convex, and let Y be a real-valued random variable such that $\mathbb{E}[g(x - Y)]$ is finite for all $x \in \mathbb{R}$. $\mathbb{E}[g(x - Y)]$ is a convex, finite-valued function of x over $\mathbb{R}$.

Proof. Given reals $x_1 < x_2$ and $\lambda \in (0, 1)$, we have

$$g((1-\lambda)x_1 + \lambda x_2 - Y) \leq (1-\lambda)g(x_1 - Y) + \lambda g(x_2 - Y)$$

for any realized value of Y by the convexity of g. Since $\mathbb{E}[g(x - Y)]$ is finite for all $x \in \mathbb{R}$, we may apply linearity of expectation (Wheeden and Zygmund 1977, Theorem 10.23) as well as monotonicity of expectation (Wheeden and Zygmund 1977, p. 170) to this inequality, establishing convexity. □

The following lemma is essentially contained in Sobel (1970) and Heyman and Sobel (1984, p. 527):

Lemma 3. Given the premises of Lemma 2, the left and right derivatives with respect to x of $\mathbb{E}[g(x - Y)]$ are equal to the finite quantities $\mathbb{E}[g^-(x - Y)]$ and $\mathbb{E}[g^+(x - Y)]$ respectively for all $x \in \mathbb{R}$.

Proof. Let us consider the left derivative case. For any particular x and realized value of Y, convexity implies that $\delta^{-1}(g(x - Y) - g(x - \delta - Y)) \nearrow g^-(x - Y)$ as $\delta \to 0^+$. By our premises and linearity of expectation, $\mathbb{E}[\delta^{-1}(g(x-Y) - g(x-\delta-Y))]$ is finite for all real $\delta > 0$. Using the monotone convergence theorem (Wheeden and Zygmund 1977, Theorem 10.27), we conclude that

$$\delta^{-1}\left(\mathbb{E}[g(x - Y)] - \mathbb{E}[g(x - \delta - Y)]\right) \to \mathbb{E}[g^-(x - Y)]$$

By Lemmas 1 and 2, the left derivative of $\mathbb{E}[g(x-Y)]$ must exist. By the convergence above, then, $\mathbb{E}[g^-(x - Y)]$ is equal to this left derivative and is finite as desired. A similar argument may be applied for the case of right derivatives. □

Lemma 4. Given the premises of Lemma 2, we have:

- $\lim_{x\to-\infty} \mathbb{E}[g^-(x-Y)] =: \mathbb{E}[g^-(-\infty-Y)] = g^-(-\infty)$
- $\lim_{x\to+\infty} \mathbb{E}[g^+(x-Y)] =: \mathbb{E}[g^+(+\infty-Y)] = g^+(+\infty)$

Proof. For any realized value of Y, we have $g^-(x-Y) \searrow g^-(-\infty)$ as $x \to -\infty$ and $g^+(x-Y) \nearrow g^+(+\infty)$ as $x \to +\infty$. By Lemma 3, $\mathbb{E}[g^-(x-Y)]$ and $\mathbb{E}[g^+(x-Y)]$ are finite for every $x \in \mathbb{R}$. Using the monotone convergence theorem, we obtain $\lim_{x\to-\infty} \mathbb{E}[g^-(x-Y)] = g^-(-\infty)$ and $\lim_{x\to+\infty} \mathbb{E}[g^+(x-Y)] = g^+(+\infty)$. □

Rolle's theorem (Bartle 1976, p. 196) may be modified for one-sided derivatives:

Lemma 5. Let $g : \mathbb{R} \to \mathbb{R}$ be continuous on the closed interval $[x_1, x_2]$ with g^+ existing on the open interval (x_1, x_2), and suppose $g(x_1) = g(x_2) = 0$. There exist x' and x'' in (x_1, x_2) such that $g^+(x') \leq 0$ and $g^+(x'') \geq 0$.

Proof. If $g(x) = 0$ for all $x \in (x_1, x_2)$, we may choose $x' = x'' = \frac{1}{2}(x_1 + x_2)$. Suppose that $g(x) > 0$ for some $x \in (x_1, x_2)$. There must exist a maximizer x' of the continuous function g over the compact set $[x_1, x_2]$. By our supposition $x_1 \neq x' \neq x_2$, and since x' is a maximizer we must have $g^+(x') \leq 0$. Now there must be some $x'' \in (x_1, x')$ such that $g(x'') = \frac{1}{2}g(x')$ and $\frac{1}{2}g(x') < g(x)$ for all $x \in (x'', x')$. (By the intermediate value theorem, there exists $\hat{x}$ satisfying $g(\hat{x}) = \frac{1}{2}g(x')$ between x and x' whenever $g(x) < \frac{1}{2}g(x')$. If for every such $\hat{x}$ the function g goes below $g(\hat{x})$ inside $(\hat{x}, x')$, $\hat{x}$ may be taken infinitely close to x' and so g cannot be continuous on the left at x', a contradiction.) We conclude that $g^+(x'') \geq 0$. Finally, in the last remaining case we must have $g(x) < 0$ for some $x \in (x_1, x_2)$, and we may employ a similar argument. □

Lemma 5 may be used in a form analogous to the mean value theorem (Bartle 1976, p. 196) to prove, by contradiction, the following condition for convexity; cf. Theorem 5.3.1 in Hiriart-Urruty and Lemaréchal (1993, p. 34).

Lemma 6. If $g : \mathbb{R} \to \mathbb{R}$ is continuous with right derivatives for all $x \in \mathbb{R}$ and $g^+(x)$ is nondecreasing in x, then g is convex.

The following is implied by Theorem 24.1 in Rockafellar (1970):

Lemma 7. If $g : \mathbb{R} \to \mathbb{R}$ is convex and $\bar{x} \in \mathbb{R}$, then $\lim_{x \to \bar{x}^-} g^+(x) = g^-(\bar{x})$ and $\lim_{x \to \bar{x}^+} g^+(x) = g^+(\bar{x})$. Similarly,

The following is a slight modification of Lemma 8-5 in Heyman and Sobel (1984); cf. Lemma 3 of Sobel (1971).

Lemma 8. If on an open interval of $\mathbb{R}$ we have a sequence of finite-valued convex functions g_n such that $g_n \nearrow g$, the limit also being finite-valued, then for all x on this interval we have:

$$g^-(x) \leq \liminf_{n \to +\infty} g_n^-(x) \leq \limsup_{n \to +\infty} g_n^+(x) \leq g^+(x)$$

Proof. Let x be an element of the given interval. By Lemma 1, $g_n^-(x)$ and $g_n^+(x)$ exist for all n. The limiting function g is finite and inherits convexity, so Lemma 1 implies that $g^-(x)$ and $g^+(x)$ exist also.

Suppose that the first desired inequality does not hold. This means that there exists $\varepsilon > 0$ and a subsequence of elements $g_n^-(x)$ such that $g_n^-(x) < g^-(x) - \varepsilon$ for all n. But there also exists $\delta > 0$ such that $(x - \delta)$ lies in the given interval and also $\delta^{-1}(g(x) - g(x - \delta)) > g^-(x) - \varepsilon/2$. Furthermore, there exists n' in our subsequence of indices such that $g_{n'}(x) > g(x) - \delta\varepsilon/2$. We may now bring out a contradiction. By Lemma 1, $g_{n'}(x - \delta) \geq g_{n'}(x) - \delta g_{n'}^-(x)$. By our definitions of ε and n', we have $g_{n'}(x) - \delta g_{n'}^-(x) > g(x) - \delta g^-(x) + \delta\varepsilon/2$. By our definition of δ, $g(x) - \delta g^-(x) + \delta\varepsilon/2 > g(x - \delta)$. Putting these together, the implication is that $g_{n'}(x - \delta) > g(x - \delta)$, which contradicts the fact that $g_n \nearrow g$.

The second desired inequality follows from noting that $g_n^-(x) \leq g_n^+(x)$ for all n, which is a consequence of Lemma 1. To obtain the third inequality, we may use an argument symmetrical with that given for the first inequality. □

A consequence of Lemma 8 is that any sequence of minimizers of such a sequence of functions g_n will become infinitely close to the set of minimizers of g. (This is not to say that the sequence of minimizers necessarily converges to a point.) For observe that if a given x is less than the interval of minimizers of g, we have $g^+(x) < 0$, and Lemma 8 implies that $g_n^+(x) < 0$ for sufficiently large n, which implies that minimizers of g_n are eventually greater than x. We may argue similarly that, if a given x is greater than the interval of minimizers of g, then minimizers of g_n are eventually less than x.

The following is given as Dini's theorem in Bartle (1976, p. 173, without proof).

Lemma 9. If a monotone sequence of continuous functions f_n converges at each point of a compact set K in $\mathbb{R}^p$ to a function f which is continuous on K, then the convergence is uniform on K.

The next lemma is implied by the primary result 1.4.3 in Flett (1980, p. 22).

Lemma 10. Let $g : \mathbb{R} \to \mathbb{R}$ be continuous, and suppose that g has right derivatives for all $x \in \mathbb{R}$, such that $g^+(x) \geq 0$ for all $x \geq x'$. The function g is nondecreasing over $[x', +\infty)$.

Note that Lemma 10 may be "turned around" to imply that a function g having *left* derivatives is nonincreasing over $(-\infty, x']$ if $g^-(x) \leq 0$ for all $x \leq x'$.

Lemma 11. For all $n \geq 1$, let $g_n : \mathbb{R} \to \mathbb{R}$ be continuous with right derivatives for all $x \in \mathbb{R}$; let $g : \mathbb{R} \to \mathbb{R}$ have right derivatives for all $x \in \mathbb{R}$, and assume that $g_n \to g$. If for some $x' \in \mathbb{R}$ we have $g_n^+(x) \geq 0$ for all $x \geq x'$ and all $n \geq 1$, then $g^+(x') \geq 0$.

Proof. Suppose instead that we have $g^+(x') < 0$ for some x' as described. There exists $\delta > 0$ small enough so that $\delta^{-1}(g(x' + \delta) - g(x')) < 0$. Since $g_n \to g$, there exists n'

large enough so that $g_{n'}(x')$ and $g_{n'}(x'+\delta)$ are each strictly less than ε away from their limiting values as $n \to +\infty$, where we define $\varepsilon := (g(x') - g(x'+\delta))/2$, noting that by our supposition $\varepsilon > 0$. It follows that $g_{n'}(x') > g_{n'}(x'+\delta)$. Lemma 10, however, implies that $g_{n'}$ must be nondecreasing over $[x', +\infty)$, so we have a contradiction. $\square$

Turning around the preceding result yields $g^-(x') \leq 0$, if we are instead given $g_n^-(x) \leq 0$ for all $x \leq x'$ and the existence of *left* derivatives. These statements may also be applied to functions $-g_n$ with $-g$ to show the preservation of inequalities in the other direction, i.e., $g^+(x') \leq 0$ if $g_n^+(x) \leq 0$ for all $x \geq x'$, and $g^-(x') \geq 0$ if $g_n^-(x) \geq 0$ for all $x \leq x'$.

Lemma 12. Let $g : \mathbb{R} \to \mathbb{R}$ be a nondecreasing function, let Y be a real-valued random variable such that $E[g(Y)]$ is finite, and let y_1 and y_2 be real with $y_1 < y_2$. If $P(Y > y_2) > 0$, then $\mathbb{E}[g(Y) \,|\, Y > y_2] \geq \mathbb{E}[g(Y) \,|\, Y > y_1]$.

Proof. Define $\alpha := \mathbb{E}[g(Y) \,|\, Y > y_2]$. Define $g^*(y) := g(y)\mathbf{1}(y > y_2) + \alpha\mathbf{1}(y \leq y_2)$, where $\mathbf{1}$ is the indicator function. Using monotonicity of expectation and the fact that g is nondecreasing, for $y \leq y_2$ we obtain $g(y) \leq g(y_2) = \mathbb{E}[g(y_2) \,|\, Y > y_2] \leq \alpha$. Hence $g \leq g^*$, and by another application of monotonicity of expectation we find that $\mathbb{E}[g(Y) \,|\, Y > y_1] \leq \mathbb{E}[g^*(Y) \,|\, Y > y_1]$. Now we may observe

$$\begin{aligned}
\mathbb{E}[g^*(Y) \,|\, Y > y_1] &= \frac{\mathbb{E}[g(Y)\mathbf{1}(Y > y_2)] + \mathbb{E}[\alpha\mathbf{1}(y_1 < Y \leq y_2)]}{P(Y > y_1)} \\
&= \frac{\alpha P(Y > y_2) + \alpha P(y_1 < Y \leq y_2)}{P(Y > y_1)}
\end{aligned}$$

Simplifying, $\mathbb{E}[g^*(Y) \,|\, Y > y_1] = \alpha$, and the desired result follows. $\square$

For a *nonincreasing* function g, the above result may be applied to the function $-g$, from which we conclude instead $\mathbb{E}[g(Y) \,|\, Y > y_2] \leq \mathbb{E}[g(Y) \,|\, Y > y_1]$.

The following is a Tauberian theorem relating the infinite-horizon discounted and average cost criteria. A proof of an analogous statement for nonpositive sequences is offered after Lemma 8.10.6 in Puterman (1994).

Lemma 13. Suppose we are given a nonnegative sequence of elements a_t (e.g., expected costs in successive decision periods) with discount factor $\beta \in (0,1)$. We have

$$\limsup_{\beta \to 1}(1-\beta)\sum_{j=0}^{+\infty}\beta^j a_j \leq \limsup_{t \to +\infty}\frac{1}{t+1}\sum_{j=0}^{t} a_j$$

Lastly, we may bring out an implicit assumption of our framework which we will take for granted in the arguments ahead.

Lemma 14. Given our definition of X and Assumptions 1–3, $\mathbb{E}[X] < +\infty$ must hold.

Proof. By Assumption 1, Assumption 2, and Lemma 1, there exists $Z' \in \mathbb{R}$ such that $L^-(Z') < 0$. Using Lemma 1, we find that $L(Z') - L(Z' - X) \leq XL^-(Z')$ for all realized values of X. By linearity and monotonicity of expectation,

$$\mathbb{E}[X] \leq \frac{\mathbb{E}[L(Z' - X)] - L(Z')}{-L^-(Z')}$$

Using Assumption 3, we have a finite upper bound on $\mathbb{E}[X]$. □

1.4.2 Relaxed model, finite-horizon discounted setting

We begin our formal argument by stating the dynamic programming recursion for the *relaxed* model in a finite-horizon discounted setting. For $n \geq 1$:

$$f_n(I) := \min_{Z \in \mathbb{R}}\{C(Z - I) + \mathbb{E}[L(Z - X)] + \beta\mathbb{E}[f_{n-1}(Z - X)]\}$$

Here $\beta \in (0,1)$ is a discount factor; we define the terminal cost function $f_0(I) := 0$. $f_n(I)$ represents the minimum possible total discounted expected cost, to be incurred between the present (with n decision points remaining) and the terminal period, if the current inventory level is I. This fact is implied by the dynamic programming theorem in Hernández-Lerma and Lasserre (1996, p. 24, as generalized for discounted costs on p. 32). To be more precise, their theorem concludes that $f_n(I)$ is the minimal cost if we take the (undiscounted) one-stage cost function to be $C(Z-I)+\mathbb{E}[L(Z-X)]$. Using the tower property (Grimmett and Stirzaker 2001, p. 336), we find that this characterization of the one-stage cost is equivalent to $C(Z-I)+L(Z-X)$ as far as our optimality criteria are concerned.

Our first results in this section will show that the minimum in the recursion above is achieved by a decision rule of the rFGB type. In this effort, we will make use of the following function, defined for $n \geq 1$:

$$G_n(Z) := \mathbb{E}[L(Z-X)] + \beta\mathbb{E}[f_{n-1}(Z-X)]$$

Theorem 1. Let $n \geq 1$ be given. Suppose that $G_n(Z)$ is convex and finite-valued, with $G_n^-(-\infty) < 0$ and $G_n^+(+\infty) > 0$. Define $s_n^0 := +\infty$. For all $i \in \{*, 1, \ldots, m\}$, choose s_n^i minimizing $c^i Z + G_n(Z)$ over $\mathbb{R}$ (we define $c^* := 0$). If the minimum does not exist, define $s_n^i := -\infty$. Then for all $I \in \mathbb{R}$, the rFGB decision rule corresponding to parameters $\{r^0, \ldots, r^m\}$ and critical values $\{s_n^0, \ldots, s_n^m\}$ along with s_n^* achieves the minimum in the dynamic programming recursion for $f_n(I)$.

Proof. (In this proof, we drop all subscripts n from G_n and the critical values s_n^i.) By the given properties of $G(Z)$, we may make three comments immediately. First, to say that the minimum of $c^i Z + G(Z)$ does not exist is to imply that this function is nondecreasing everywhere (since $c^i \geq 0$); it is therefore natural in this case to define $s^i = -\infty$. Second, we may be assured that s^* is finite. Third, we have:

$$s^m \leq s^{m-1} \leq \ldots \leq s^2 \leq s^1 \leq s^* < s^0 = +\infty$$

For finite s^i, we know in particular that $G^-(s^i) \leq -c^i$ and $G^+(s^i) \geq -c^i$. Notice that for any particular $I \in \mathbb{R}$, the dynamic programming recursion is the problem of minimizing the convex, finite-valued objective function $C(Z-I)+G(Z)$ over $Z \in \mathbb{R}$. For a given I, an optimal Z^* is therefore one that satisfies

$$C^-(Z^*-I)+G^-(Z^*) \leq 0 \leq C^+(Z^*-I)+G^+(Z^*)$$

We proceed with an analysis of the possible cases for $I \in \mathbb{R}$:

- If $s^i - r^i \leq I < s^i - r^{i-1}$ for some $i \in \{1,\dots,m\}$, let $Z^* = s^i$. This case only applies if s^i is finite, so $G^-(Z^*)$ and $G^+(Z^*)$ exist. Moreover, the order quantity is in $(r^{i-1}, r^i]$. We have $C^-(Z^*-I) = c^i$ and $G^-(Z^*) \leq -c^i$, so decreasing Z^* does not improve the total cost. Similarly, $C^+(Z^*-I) \geq c^i$ and $G^+(Z^*) \geq -c^i$, so increasing Z^* does not improve the total cost, either.

- If $s^i - r^{i-1} \leq I < s^{i-1} - r^{i-1}$ for some $i \in \{2,\dots,m\}$, let $Z^* = I + r^{i-1}$. This case only applies if s^{i-1} is finite, so Z^* is some finite number in $[s^i, s^{i-1})$. $C^-(Z^*-I) = c^{i-1}$ and $G^-(Z^*) \leq -c^{i-1}$, so decreasing Z^* does not improve the total cost. Similarly, $C^+(Z^*-I) = c^i$ and $G^+(Z^*) \geq -c^i$, so increasing Z^* also fails to improve the total cost.

- If $s^1 \leq I < s^*$, let $Z^* = I$. $C^-(Z^*-I) = 0$ and $G^-(Z^*) \leq 0$, so decreasing Z^* does not improve the total cost; $C^+(Z^*-I) = c^1$ and $G^+(Z^*) \geq -c^1$, so increasing Z^* cannot improve the total cost.

- If $s^* \leq I$, let $Z^* = s^*$. (We know that s^* is finite.) $C^-(Z^*-I) = 0$ and $G^-(Z^*) \leq 0$, so decreasing Z^* will not improve the total cost; $C^+(Z^*-I) \geq 0$ and $G^+(Z^*) \geq 0$, so increasing Z^* does not improve the total cost.

In each case above, we see that the rFGB decision rule according to the given parameters achieves the minimum in the dynamic programming recursion. □

Theorem 2. For all $n \geq 1$,

- $G_n(Z)$ is convex and finite-valued with $G_n^-(-\infty) < 0$ and $G_n^+(+\infty) > 0$.
- $f_n(I)$ is convex and finite-valued with $f_n^-(-\infty) \geq -c^m$ and $f_n^+(I) = 0$ for I sufficiently large.

Proof. Our proof will be by induction on n. First, by our definition that $f_0(I) = 0$, we see that $f_0(I)$ satisfies the required conditions. Assume that these conditions are valid for $f_{n-1}(I)$ for some particular $n \geq 1$.

By the induction hypothesis, $f_{n-1}(Z)$ is finite for any $Z \in \mathbb{R}$ and $f_{n-1}(Z - X)$ is finite for any realized value of X. The conditions on the derivatives of f_{n-1} imply that $f_{n-1}(Z) \leq f_{n-1}(Z-X) \leq f_{n-1}(Z)+c^m X$. Since we know that $\mathbb{E}[X] < +\infty$, by taking expectations we find that $\mathbb{E}[f_{n-1}(Z - X)]$ is a finite-valued function of Z. Now by Lemmas 2 and 3 we may also conclude that $\mathbb{E}[f_{n-1}(Z-X)]$ is convex in Z and that its left and right derivatives are equal to $\mathbb{E}[f_{n-1}^-(Z-X)]$ and $\mathbb{E}[f_{n-1}^+(Z-X)]$ respectively. Lemmas 2 and 3 may be similarly applied to $\mathbb{E}[L(Z - X)]$ using Assumption 3. By Assumption 2 and Lemma 4, $\mathbb{E}[L^-(-\infty - X)] < 0$ and $\mathbb{E}[L^+(+\infty - X)] > 0$. Similarly, by the induction hypothesis we may argue that $\mathbb{E}[f_{n-1}^-(-\infty - X)] \leq 0$ and $\mathbb{E}[f_{n-1}^+(+\infty - X)] = 0$. We conclude that $G_n(Z)$ is a convex, finite-valued function with $G_n^-(-\infty) < 0$ and $G_n^+(+\infty) > 0$, and so we may apply Theorem 1 to obtain

$$f_n(I) = C(Z_n^*(I) - I) + G_n(Z_n^*(I))$$

where $Z_n^*(I)$ prescribes an optimal action for inventory level I according to an rFGB decision rule with critical values $\{s_n^0, \ldots, s_n^m\}$ and s_n^* satisfying the conditions of the theorem. $Z_n^*(I)$ is a continuous function of I, and the functions C and G_n are continuous, so $f_n(I)$ is continuous. Similarly, $f_n(I)$ is finite-valued. We now analyze the right derivatives of $f_n(I)$:

- If $s_n^i - r^i \le I < s_n^i - r^{i-1}$ for some $i \in \{1, \dots, m\}$, then $Z_n^*(I) = s_n^i$. Increasing I decreases the order quantity, but does not change the order-up-to level. The order quantity is in $(r^{i-1}, r^i]$, so $f_n^+(I) = -c^i$ here.

- If $s_n^i - r^{i-1} \le I < s_n^{i-1} - r^{i-1}$ for some $i \in \{2, \dots, m\}$, $Z_n^*(I) = I + r^{i-1}$. Increasing I does not change the order quantity, so $f_n^+(I) = G_n^+(I + r^{i-1})$ here. Over $[s_n^i, s_n^{i-1})$, $G_n^+(Z)$ is a nondecreasing function with values in $[-c^i, -c^{i-1}]$.

- If $s_n^1 \le I < s_n^*$, $Z_n^*(I) = I$. Increasing I does not change the order quantity, so $f_n^+(I) = G_n^+(I)$ here. Over $[s_n^1, s_n^*)$, $G_n^+(Z)$ is a nondecreasing function with values in $[-c^1, 0]$.

- If $s_n^* \le I$, $Z_n^*(I) = s_n^*$. Increasing I does not cause any change in ordering cost, and the order-down-to level is fixed, so $f_n^+(I) = 0$ here.

We observe that $f_n(I)$ has nondecreasing right derivatives; we may now invoke Lemma 6 to conclude that $f_n(I)$ is convex. Furthermore, by our analysis above we must have $f_n^-(I) \ge -c^m$ for all I, and $f_n(I)$ is constant over $[s_n^*, +\infty)$ where s_n^* has been shown to be finite. $\square$

By the dynamic programming theorem in Hernández-Lerma and Lasserre (1996), Theorems 1 and 2 establish that a policy consisting of rFGB decision rules as defined in Theorem 1 is discount optimal in our finite-horizon setting.

In preparation for our later arguments, we will study properties that depend on β. To make this dependence explicit, we may add a subscript 'β' to symbols already defined. The two theorems below have implications for limiting cases as $n \to +\infty$ and $\beta \to 1$.

Theorem 3. For all $n \ge 1$:

- $f_n^+(I) \le f_{n-1}^+(I)$ and $f_n^-(I) \le f_{n-1}^-(I)$ for all I. (These are derivatives with respect to I.)

- $s^i_{n+1} \geq s^i_n$ if these values are chosen to be the *greatest* minimizers defined in Theorem 1, for all $i \in \{*, 1, \ldots, m\}$.

- $f^+_{\beta_2,n}(I) \leq f^+_{\beta_1,n}(I)$ and $f^-_{\beta_2,n}(I) \leq f^-_{\beta_1,n}(I)$, for $\beta_1 < \beta_2$, for all I.

- $s^i_{\beta_2,n} \geq s^i_{\beta_1,n}$ for $\beta_1 < \beta_2$, if these values are chosen to be the *greatest* minimizers defined in Theorem 1, for all $i \in \{*, 1, \ldots, m\}$.

Proof. (We prove the first two statements first.) Our proof will be by induction on n. First, by our definition that $f_0(I) = 0$ and by Theorem 2, we have $f^+_1 \leq f^+_0$ and $f^-_1 \leq f^-_0$. Assume henceforth that $f^+_n \leq f^+_{n-1}$ and $f^-_n \leq f^-_{n-1}$ for some particular $n \geq 1$.

Using Lemmas 2 and 3 as in the proof of Theorem 2 and applying monotonicity of expectation with the induction hypothesis, we find that $G^+_{n+1} \leq G^+_n$ and $G^-_{n+1} \leq G^-_n$. From Theorem 2, we know that G_{n-1} and G_n are convex and finite-valued. For any $i \in \{*, 1, \ldots, m\}$, if s^i_n minimizes $c^i Z + G_n(Z)$ then we have $G^-_n(s^i_n) \leq -c^i$. We conclude that $G^-_{n+1}(s^i_n) \leq -c^i$ also, so s^i_{n+1} may be chosen to be greater than s^i_n as desired. (If $s^i_n = -\infty$ then we have nothing to prove.) Choosing the greatest minimizers at every step is therefore one way to make sure that $s^i_{n+1} \geq s^i_n$. (Convexity implies that the set of minimizers is a closed interval, and it follows in our situation that the greatest minimizer exists whenever a minimizer exists.)

Choosing the greatest minimizers as indicated, we may now use facts from the proof of Theorem 2 to verify that $f^+_{n+1}(I) \leq f^+_n(I)$ for all $I \in \mathbb{R}$:

- If $s^i_{n+1} - r^i \leq I < s^i_{n+1} - r^{i-1}$ for some $i \in \{1, \ldots, m\}$, then $f^+_{n+1}(I) = -c^i$. Meanwhile, since $s^i_n - r^i \leq s^i_{n+1} - r^i$, we have $f^+_n(I) \geq -c^i$.

- If $s^i_{n+1} - r^{i-1} \leq I < s^{i-1}_{n+1} - r^{i-1}$ for some $i \in \{2, \ldots, m\}$, then $f^+_{n+1}(I) = G^+_{n+1}(I + r^{i-1}) \leq -c^{i-1}$. At the same time, since $s^i_n - r^{i-1} \leq s^i_{n+1} - r^{i-1}$, we have $f^+_n(I) \geq \min\{G^+_n(I + r^{i-1}), -c^{i-1}\} \geq G^+_{n+1}(I + r^{i-1})$.

- If $s^1_{n+1} \leq I < s^*_{n+1}$, then $f^+_{n+1}(I) = G^+_{n+1}(I) \leq 0$. Since $s^1_n \leq s^1_{n+1}$, we also have $f^+_n(I) = \min\{G^+_n(I), 0\} \geq G^+_{n+1}(I)$.

- If $s^*_{n+1} \leq I$, then $f^+_{n+1}(I) = 0$. Since $s^*_{\beta_1,n} \leq s^*_{n+1}$, $f^+_n(I) = 0$ also.

In each case above, we see that the desired inequality holds. That $f^-_{n+1}(I) \leq f^-_n(I)$ for all I follows by Lemma 7. For if this inequality were invalid for some inventory level I', then any increasing sequence of inventory levels I_k approaching I' would see $f^+_{n+1}(I_k) > f^+_n(I_k)$ eventually.

(We now prove the last two statements of the theorem.) Let β_1 and β_2 be given such that $\beta_1 < \beta_2$. We will argue by induction on n that $f^+_{\beta_2,n} \leq f^+_{\beta_1,n}$ and $f^-_{\beta_2,n} \leq f^-_{\beta_1,n}$. In the base case, by our definition $f^+_{\beta,0} = f^-_{\beta,0} = 0$ for all β, and the desired inequalities are satisfied. Henceforth, let n be given and assume that the case for $n-1$ is settled. For any β and Z, we have, as suggested in the first half of this proof:

$$G^+_{\beta,n}(Z) = \mathbb{E}[L^+(Z - X)] + \beta\mathbb{E}[f^+_{\beta,n-1}(Z - X)]$$

The first term on the right side above is independent of the discount factor. Regarding the second term: the induction hypothesis and the fact (from Theorem 2) that $f^+_{\beta,n-1} \leq 0$ imply, via monotonicity of expectation, that $\beta_2\mathbb{E}[f^+_{\beta_2,n-1}(Z - X)] \leq \beta_1\mathbb{E}[f^+_{\beta_1,n-1}(Z - X)]$. Thus $G^+_{\beta_2,n} \leq G^+_{\beta_1,n}$, and essentially the same argument yields $G^-_{\beta_2,n} \leq G^-_{\beta_1,n}$. By reasoning as in the first half of this proof, it follows that choosing the greatest minimizers gives us $s^i_{\beta_2,n} \geq s^i_{\beta_1,n}$. Continuing with analogous reasoning, we find that $f^+_{\beta_2,n} \leq f^+_{\beta_1,n}$ and $f^-_{\beta_2,n} \leq f^-_{\beta_1,n}$. □

Theorem 4. If $L^-(-\infty)/(1-\beta) < -c^i$, then the critical values $\{s^*_n, s^1_n, \ldots, s^i_n\}$ defined according to Theorem 1 will be finite for n sufficiently large.

Proof. Since the critical values in question are in $[s^i_n, s^*_n]$ (see the proof of Theorem 1) and we are assured by Theorem 2 that $G^+_n(+\infty) > 0$, a sufficient condition for the desired finiteness is $G_n^-(-\infty) < -c^i$. Since $f_0(I) = 0$ we have $G_1(Z) = \mathbb{E}[L(Z - X)]$, and by Lemma 4 we see that $G_1^-(-\infty) = L^-(-\infty)$. Now if $s_{n-1}^i = -\infty$, through the definition of G_n and an analysis of the derivatives of f_{n-1} as in the proof of Theorem 2, we obtain $G_n^-(-\infty) = L^-(-\infty) + \beta G_{n-1}^-(-\infty)$. Summing this geometric series, we find that $s_{n'}^i$ will be finite for some n' if $L^-(-\infty)/(1-\beta) < -c^i$. By Theorem 3, s_n^i (and the other, greater critical values) will remain finite for all $n > n'$ as well. □

Since $L^-(-\infty) < 0$ by Assumptions 1–3, the condition in Theorem 4 is satisfied for any particular $i \in \{1, \ldots, m\}$ when β is sufficiently close to 1.

1.4.3 Relaxed model, infinite-horizon discounted setting

We now build a case for the optimality of an rFGB policy for the relaxed model with respect to the infinite-horizon discounted expected cost. In this effort, we will make use of results proved for the finite-horizon case. The first two theorems here will facilitate our solution of the discounted cost optimality equation.

Theorem 5. Let $I \in \mathbb{R}$ be given. There exist real numbers Z_I' and Z_I'' with $Z_I' < Z_I''$ such that the compact interval $[Z_I', Z_I'']$ contains the minimizer(s) of the dynamic programming recursion for $f_n(I)$, for all $n \geq 1$.

Proof. Consider the policy that always chooses $Z = 0$. The infinite-horizon discounted expected cost of this policy, given that the initial state is I, is

$$\chi(I) := C(-I) + \mathbb{E}[L(-X)] + \beta \frac{\mathbb{E}[C(X)] + \mathbb{E}[L(-X)]}{1-\beta}$$

which is finite for any given $\beta \in (0, 1)$ under our assumptions. Any finite-horizon problem with initial state I will have an optimal discounted expected cost no greater than this quantity.

Let $n \geq 1$ and $I \in \mathbb{R}$ be given. We define Z_I' such that $\mathbb{E}[L(Z - X)] > \chi(I)$ for all $Z \leq Z'$, and we define Z_I'' such that $\mathbb{E}[L(Z - X)] > \chi(I)$ for all $Z \geq Z''$.

By Assumptions 1–3 and Lemmas 2–4, Z_I' and Z_I'' may be chosen to be finite with $Z_I' < Z_I''$. $C \geq 0$ and $L \geq 0$ by definition, and by induction we may argue that $f_{n-1} \geq 0$, so we have $C(Z-I) + \beta\mathbb{E}[f_{n-1}(Z-X)] \geq 0$ for all $Z \in \mathbb{R}$. Therefore, for each $Z \notin [Z_I', Z_I'']$ it is the case that

$$C(Z-I) + \mathbb{E}[L(Z-X)] + \beta\mathbb{E}[f_{n-1}(Z-X)] > \chi(I)$$

But by the dynamic programming theorem and Theorems 1 and 2, $f_n(I)$ is the cost of a discount optimal policy, so we must have

$$C(Z-I) + \mathbb{E}[L(Z-X)] + \beta\mathbb{E}[f_{n-1}(Z-X)] \leq \chi(I)$$

for any minimizer of the recursive expression for $f_n(I)$. □

Theorem 6. As $n \to +\infty$, $0 \leq f_n \nearrow f$ where $f : \mathbb{R} \to \mathbb{R}$ is convex with $f^-(-\infty) \geq -c^m$ and $f^+(I) = 0$ for I sufficiently large.

Proof. By our definitions of the functions C and L and the terminal cost f_0, the cost incurred in any period in our finite-horizon setting is nonnegative. By Theorems 1 and 2 together with the dynamic programming theorem, then, for each $n \geq 0$ the optimal cost function f_n must be nonnegative. Furthermore, the optimal cost for our finite-horizon problem with $(n+1)$ decision points cannot be less than the optimal cost for the problem with n decision points, given a common initial state. The sequence $\{f_n(I) : n \geq 0\}$ is therefore nondecreasing for any given I. Additionally, continuing with notation introduced in the proof of Theorem 5, for any given I we have $f_n(I) \leq \chi(I)$ for all $n \geq 0$ where $\chi(I)$ is a finite quantity. Thus the sequence $\{f_n(I) : n \geq 0\}$ converges to a finite value as $n \to +\infty$.

As a pointwise limit, f inherits the convexity of the functions f_n established in Theorem 2. Since f is also finite-valued, its left and right derivatives exist. The property $f^-(-\infty) \geq -c^m$ is inherited from the analogous property established for the functions f_n in Theorem 2, as may be seen by using Lemma 11.

It remains to argue that $f^+(I) = 0$ for I sufficiently large. Consider again the proof of Theorem 5. If $I = 0$, we may define Z_0'' such that $\mathbb{E}[L(Z - X)] > \chi(0)$ for all $Z \geq Z_0''$. Now observe that $\chi(I) = \chi(0)$ for all $I \geq 0$. By Theorem 5, then, a minimizer of the dynamic programming recursion for $f_n(I)$ cannot be greater than Z_0'' for any $I \geq 0$ and $n \geq 1$. It follows that critical values s_n^* as defined in Theorem 1 cannot exceed Z_0''. This implies through the proof of Theorem 2 that $f_n^+(I) = 0$ for $I \geq Z_0''$, for all $n \geq 1$. Using Lemma 11, we conclude that $f^+(I) = 0$ for $I \geq Z_0''$. $\square$

Following the finite-horizon case, we define $G(Z) := \mathbb{E}[L(Z-X)] + \beta\mathbb{E}[f(Z-X)]$. We also define $J_n(I, Z) := C(Z - I) + G_n(Z)$ and $J(I, Z) := C(Z - I) + G(Z)$.

Theorem 7. For all $I \in \mathbb{R}$, f satisfies the discounted cost optimality equation:

$$f(I) = \min_{Z \in \mathbb{R}} \{C(Z - I) + \mathbb{E}[L(Z - X)] + \beta\mathbb{E}[f(Z - X)]\}$$

Proof. Theorem 6 implies that $0 \leq f_n(Z - X) \nearrow f(Z - X)$ for any given $Z \in \mathbb{R}$ and any realized value of X. The properties of the derivatives of f established there along with the fact that $\mathbb{E}[X] < +\infty$ yields that $\mathbb{E}[f(Z - X)]$ is finite for all Z, as argued in the proof of Theorem 2 for $\mathbb{E}[f_{n-1}(Z - X)]$. We may now invoke the monotone convergence theorem to establish that $G_n \nearrow G$ as $n \to +\infty$ (by monotonicity of expectation the convergence is monotone). We may also use Lemma 2 to see that $\mathbb{E}[f(Z - X)]$ is a finite-valued convex function of Z. It follows that, given $I \in \mathbb{R}$, $J(I, Z)$ as well as $J_n(I, Z)$ are finite-valued convex functions of Z. Moreover, that $G_n \nearrow G$ implies $J_n(I, Z) \nearrow J(I, Z)$ as $n \to +\infty$ for all real I and Z.

Let $I \in \mathbb{R}$ be given. By Theorem 5, there is a compact interval $\varphi(I) := [Z_I', Z_I'']$ such that any minimizer of $J_n(I, Z)$ over $Z \in \mathbb{R}$ must be in $\varphi(I)$. By the development above, Lemma 9 implies that $J_n(I, Z)$ converges uniformly to $J(I, Z)$ on the compact interval $\varphi(I)$. Letting $\varepsilon > 0$ be given, there exists N such that $n \geq N$ implies both $\min_{Z\in\varphi(I)} J_n(I,Z) \leq \min_{Z\in\varphi(I)} J(I,Z) + \varepsilon$ and $\min_{Z\in\varphi(I)} J_n(I,Z) \geq \min_{Z\in\varphi(I)} J(I,Z) - \varepsilon$. Since ε can be arbitrarily small and $f_n(I) = \min_{Z\in\varphi(I)} J_n(I,Z)$, we may take $n \to +\infty$ to conclude that $f(I) = \min_{Z\in\varphi(I)} J(I,Z)$.

We may now briefly argue that $\min_{Z\in\varphi(I)} J(I,Z) = \min_{Z\in\mathbb{R}} J(I,Z)$. Again, let $I \in \mathbb{R}$ be fixed. By the convexity of J_n and our selection of $\varphi(I)$, it follows that $J_n^-(I, Z_I') \leq 0$ and $J_n^+(I, Z_I'') \geq 0$, where these derivatives are taken with respect to Z, for all $n \geq 1$. By Lemma 11, we find that $J^-(I, Z_I') \leq 0$ and $J^+(I, Z_I'') \geq 0$, so a minimizer of $J(I,Z)$ over $Z \in \mathbb{R}$ will be found in $\varphi(I)$. This establishes that $f(I) = \min_{Z\in\mathbb{R}} J(I,Z)$ as desired. □

Our next result uses our solution to the discounted cost optimality equation to establish the optimality of an rFGB policy.

Theorem 8. Define $s^0 = +\infty$. For all $i \in \{*, 1, \ldots, m\}$, choose s^i minimizing $c^i Z + G(Z)$ over $\mathbb{R}$. (Again, we define $c^* = 0$.) If the minimum does not exist, define $s^i = -\infty$. Then for all $I \in \mathbb{R}$, the rFGB decision rule corresponding to parameters $\{r^0, \ldots, r^m\}$ and critical values $\{s^0, \ldots, s^m\}$ along with s^* achieves the minimum in the functional equation solved by $f(I)$. Moreover, the rFGB policy with these parameters is discount optimal in the infinite-horizon setting, and its cost is $f(I)$ given initial state I.

Proof. In the proof of Theorem 7, we argue that $\mathbb{E}[f(Z - X)]$ is finite-valued and convex over $Z \in \mathbb{R}$. In this context we may invoke Lemma 3 to the effect that the left and right derivatives of this funtion are $\mathbb{E}[f^-(Z - X)]$ and $\mathbb{E}[f^+(Z - X)]$ respectively. Furthermore, Lemma 4 and the properties of the derivatives of f imply that $\mathbb{E}[f^-(-\infty - X)] \leq 0$ and $\mathbb{E}[f^+(+\infty - X)] = 0$. Argument as in the proof of Theorem 2 shows that, in addition to G being finite-valued and convex, G satisfies $G^-(-\infty) < 0$ and $G^+(+\infty) > 0$. Theorem 1 (with G_n replaced by G) now implies that s^* is finite, that this and the other critical values specify a valid rFGB decision rule, and that this rule achieves the minimum in the functional equation. Adjusting the functional equation, we have:

$$f(I) \leq \min_{Z \in \mathbb{R}} \{\mathbb{E}[C(Z - I) + L(Z - X)] + \beta \mathbb{E}[f(Z - X)]\}$$

for any $I \in \mathbb{R}$. Let π be any admissible policy. Treating I_0 as an initial state implies by the above that

$$f(I_0) \leq \mathbb{E}^{\pi}_{I_0}[C(Z_0 - I_0) + L(Z_0 - X_0)] + \beta \mathbb{E}^{\pi}_{I_0}[f(I_1)]$$

where the next state is $I_1 = Z_0 - X_0$. The history of the process at decision point n, which we call H_n, consists of $(I_0, Z_0, \ldots, I_{n-1}, Z_{n-1}, I_n)$. Using the above inequality as a base for induction, suppose for given $n \geq 1$ that

$$f(I_0) \leq \sum_{t=0}^{n-1} \beta^t \mathbb{E}^{\pi}_{I_0}[C(Z_t - I_t) + L(Z_t - X_t)] + \beta^n \mathbb{E}^{\pi}_{I_0}[f(I_n)]$$

Now for any history H_n with present state I_n, we also have from our functional equation that

$$f(I_n) \leq \mathbb{E}^{\pi}_{H_n}[C(Z_n - I_n) + L(Z_n - X_n)] + \beta \mathbb{E}^{\pi}_{H_n}[f(I_{n+1})]$$

In the induction hypothesis we may therefore apply monotonicity of expectation and the tower property, obtaining:

$$f(I_0) \leq \sum_{t=0}^{n} \beta^t \mathbb{E}^{\pi}_{I_0}[C(Z_t - I_t) + L(Z_t - X_t)] + \beta^{n+1} \mathbb{E}^{\pi}_{I_0}[f(I_{n+1})]$$

The above holds for any $I_0 \in \mathbb{R}$ and any $n \geq 1$ by induction. A small adjustment to our argument shows that the above holds with equality for $\pi = \pi^*$, where π^*

represents the specified rFGB policy that achieves the minimum in the functional equation. In the remainder of this proof, we will argue that any policy that does not satisfy $\beta^{n+1}\mathbb{E}^{\pi}_{I_0}[f(I_{n+1})] \to 0$ cannot be optimal. Taking $n \to +\infty$, we may then obtain:

$$\begin{aligned} f(I_0) &= \lim_{n\to+\infty} \sum_{t=0}^{n} \beta^t \mathbb{E}^{\pi^*}_{I_0}[C(Z_t - I_t) + L(Z_t - X_t)] \\ &\leq \lim_{n\to+\infty} \sum_{t=0}^{n} \beta^t \mathbb{E}^{\pi}_{I_0}[C(Z_t - I_t) + L(Z_t - X_t)] \end{aligned}$$

for every policy π worth considering and for every $I_0 \in \mathbb{R}$. This will then establish the discount optimality of the rFGB policy π^* and validate the treatment of f as an optimal cost function.

For what kind of policy π would we see the condition $\beta^{n+1}\mathbb{E}^{\pi}_{I_0}[f(I_{n+1})] \to 0$ fail? Consider Theorem 6. Since f is nonnegative, such a policy would have to see $\limsup_{n\to+\infty} \mathbb{E}^{\pi}_{I_0}[f(I_{n+1})] = +\infty$. Since f is convex and finite-valued with $f^-(-\infty) \geq -c^m$ and $f^+(+\infty) = 0$, such a policy would have to see $\liminf_{n\to+\infty} \mathbb{E}^{\pi}_{I_0}[I_{n+1}] = -\infty$. Due to the geometric rate of decrease in β^{n+1}, a merely *linear* rate of decrease in a subsequence of $\{\mathbb{E}^{\pi}_{I_0}[I_{n+1}] : n \geq 0\}$ will not suffice to violate the condition. We will argue below that only a non-optimal policy could sustain a faster-than-linear rate of decrease in a subsequence of $\{\mathbb{E}^{\pi}_{I_0}[I_{n+1}] : n \geq 0\}$.

Recall from the proof of Theorem 5 that we need not consider policies with infinite-horizon expected discounted cost greater than the cost incurred by a policy of setting $Z = 0$ in every step. The cost of the $Z = 0$ policy with initial state I is, continuing with notation defined previously, $\chi(I) = \chi(0) + C(-I)$. As in Theorem 5, there exists some I' low enough so that $Z \leq I'$ implies $\mathbb{E}[L(Z - X)] > \chi(0)$. Assumptions 1–3 and Lemmas 2 and 3 imply that the left derivative of $\mathbb{E}[L(Z - X)]$ is $\mathbb{E}[L^-(Z - X)]$ and is no greater than $\delta := \mathbb{E}[L^-(I' - X)] < 0$ for $Z \leq I'$. For any given initial state

I, then,

$$\mathbb{E}\big[L\left(I' + \delta^{-1}c^m(-I)^+ - X\right)\big] > \chi(0) + c^m(-I)^+ \geq \chi(I)$$

Therefore, in state I we need not consider choosing Z below $I' + \delta^{-1}c^m(-I)^+$. When I decreases by one unit, this is to say that our lower bound on the order level Z decreases by (at most) the fixed quantity $|\delta^{-1}c^m|$. For any sequence of realized demands, the most rapid drop in inventory within our bounds would result from always ordering to this lower bound; call this policy π'. After our initial order dictated by π', this lower bound will decrease after each step by $|\delta^{-1}c^m|$ times the latest realized demand.

Suppose that the initial state is I_0, that the initial action according to π' is Z_0, and that for $n \geq 1$ we have $I_n = Z_{n-1} - X_{n-1}$. Given any finite sequence of realized demands, we have for $n \geq 1$ that $Z_n = Z_{n-1} + \delta^{-1}c^m X_{n-1}$, and so

$$I_{n+1} = Z_0 + \delta^{-1}c^m(X_0 + \ldots + X_{n-1}) - X_n$$

Given any initial state and a sequence of realized demands up to a given period, the above state calculated for π' may be taken as a lower bound on the state for any policy π we have not ruled out. Taking expectations,

$$\mathbb{E}^{\pi'}_{I_0}[I_{n+1}] = \mathbb{E}^{\pi'}_{I_0}[Z_0 + \delta^{-1}c^m(X_0 + \ldots + X_{n-1}) - X_n] = Z_0 + n\delta^{-1}c^m\mathbb{E}[X] - \mathbb{E}[X]$$

and so the expected inventory level under π' falls linearly in n. By monotonicity of expectation we conclude that, for those policies π we have not ruled out, the expected inventory level decreases *at most* at a linear rate (in any subsequence of the decision points). $\square$

To complete this section, we study the effect of β on our solution to the functional equation. The following result is an infinite-horizon counterpart to Theorem 3.

Theorem 9. $f^+_\beta(I)$ and $f^-_\beta(I)$ are nonincreasing in β for all I. Consequently: if s^i_β are chosen to be the *greatest* minimizers in Theorem 8, then these critical values are nondecreasing in β for all $i \in \{*, 1, \ldots, m\}$.

Proof. Let $\beta_1, \beta_2 \in (0, 1)$ be given with $\beta_1 < \beta_2$. Via Theorems 3 and 6, we may invoke Lemma 11 to conclude that $f^+_{\beta_1}(I) - f^+_{\beta_2}(I) \geq 0$ and $f^-_{\beta_1}(I) - f^-_{\beta_2}(I) \geq 0$ for all I. That the critical values are nondecreasing in β follows by an argument analogous to that used for Theorem 3. □

1.4.4 Average optimality in the relaxed and unrelaxed models

We will now develop an argument for the average optimality of an rFGB policy in the relaxed model. In this effort, we will make use of results proved for the finite-horizon and infinite-horizon discounted cases. As before, we use subscripts 'β' to explicitly indicate dependence on the discount factor. Theorems 10 and 11 function in this context similarly to Theorem 5 in its context.

Theorem 10. s^*_β, as defined in Theorem 8, is bounded above by a finite quantity independent of β.

Proof. We will show that there exists $I^* \in \mathbb{R}$ and $\beta' < 1$ such that $\beta \geq \beta'$ implies $s^*_\beta < I^*$. Since by Theorem 9 the greatest candidate for s^*_β is nondecreasing in β, we may then conclude that $s^*_\beta < I^*$ for all $\beta \in (0, 1)$.

Let π^* denote an optimal rFGB policy as specifed in Theorem 8. Suppose that the initial state is $I_0 = s^*_\beta$, and consider the operation of π^* on our model. Given a real parameter $\Delta \geq 0$, we define a history-dependent policy π^{**}_Δ by reference to its control of a second model, operating in parallel and subject to the same initial state I_0 and the same sequence of demand variables X_t (and the same cost functions C and L). At the initial decision point $t = 0$, the order quantity $(Z_0 - I_0)$ is chosen to be *one unit less* in the second model than the quantity chosen at $t = 0$ by π^* in the first model. (Given the initial state $I_0 = s^*_\beta$, this means that π^{**}_Δ chooses $Z_0 = I_0 - 1$ while π^* selects $Z_0 = I_0$.) In subsequent decision points $t \geq 1$, the order quantity $(Z_t - I_t)$ is chosen in the second model to be *the same* as the quantity chosen at t by π^* in the first model—with a single exception. In the first decision point t at which the inventory level is observed to be strictly below $s^*_\beta - \Delta$ in the first system (equivalently, strictly below $s^*_\beta - \Delta - 1$ in the second system), the order quantity $(Z_t - I_t)$ is chosen to be *one unit more* in the second system than in the first. We call this special decision point the "transitional step." To sum up from another perspective: after the initial step, the inventory level in the second system tracks exactly one unit below the level in the first system; then, after the inventory levels fall enough to trigger the transitional step, the inventory levels in the two systems match exactly. Note that the policy π^{**}_Δ controlling the second system is admissible, with the proviso that we have only defined its behavior when the initial state is $I_0 = s^*_\beta$.

Except for the transitional step, in which one more unit is ordered under π^{**}_Δ than under π^*, the ordering cost $C(Z_t - I_t)$ is the same for both policies at all decision points t. Meanwhile, if for all $I \geq s^*_\beta - \Delta - 1$ we have $L^+(I) > 0$, then the costs $L(Z_t - X_t)$ are *lower* for the second system in every step before the transitional step—excepting possibly the step immediately prior. In the transitional step and after, the costs $L(Z_t - X_t)$ are equal for the two systems. We will show that, if $s^*_\beta \geq I^*$ and $\beta \geq \beta'$, then for the initial state $I_0 = s^*_\beta$ the discounted expected cost of the policy π^{**}_Δ is less than that of π^*. Since we know that π^* is discount optimal, we may then conclude that $s^*_\beta \geq I^*$ is false for all $\beta \geq \beta'$.

By Assumptions 1 and 2, there exists $\varepsilon > 0$ and $I_{\min} \in \mathbb{R}$ such that $L^+(I) \geq \varepsilon$ for all $I \geq I_{\min} - 1$. We define $\Delta := 3k\mathbb{E}[X]$ and $\beta' := (3/4)^{1/(k-1)}$ for some integral value $k \geq 2$ to be specified. By Assumption 4, we may focus on the nontrivial case where $P(X = 0) < 1$, so we have $0 < E[X] < +\infty$. By Markov's inequality (Grimmett and

Stirzaker 2001, p. 311), we have

$$P(\sum_{t=0}^{k-1} X_t > \Delta) \leq \mathbb{E}[\sum_{t=0}^{k-1} X_t]/\Delta = k\mathbb{E}[X]/\Delta = \tfrac{1}{3}$$

Defining $T_\Delta := \min\{n : \sum_{t=0}^{n} X_t > \Delta\}$, we then have

$$P(T_\Delta \geq k) = P(\sum_{t=0}^{k-1} X_t \leq \Delta) \geq \tfrac{2}{3}$$

By the structure of rFGB policies and our choice of initial state $I_0 = s^*_\beta$, the time index of the transitional step must be (strictly) greater than T_Δ. If we posit $I_0 = s^*_\beta \geq I_{\min} + \Delta$ and $\beta \geq \beta'$, then the expected discounted savings under π^{**}_Δ (relative to π^*), before we first see $Z_t - X_t < I_0 - \Delta - 1$, is at least

$$\begin{aligned}
\mathbb{E}[\textstyle\sum_{t=0}^{T_\Delta - 1} \beta^t \varepsilon] &\geq \mathbb{E}[\textstyle\sum_{t=0}^{T_\Delta - 1} \beta^t \varepsilon \,|\, T_\Delta \geq k] P(T_\Delta \geq k) \\
&\geq \varepsilon(\textstyle\sum_{t=0}^{k-2} \beta^t)(\tfrac{2}{3}) \\
&\geq \varepsilon k(\beta')^{k-1}(\tfrac{2}{3}) \\
&= k\varepsilon/2
\end{aligned}$$

We may now fix k to be some integral value large enough so that $k\varepsilon/2 > c^m$. Here, c^m serves as an upper bound on the expected discounted excess in ordering cost incurred by π^{**}_Δ (relative to π^*) in the transitional step. (The bound remains valid in expectation even though we have not proved that the transitional step must occur.) It remains to account for the possible difference between the expected discounted costs incurred in the period immediately before the transitional step. We use T to denote the time index of the transitional step.

Suppose the random variable X is bounded above by some real quantity $X_{\max}$. By requiring $I^* \geq I_{\min} + X_{\max} + \Delta$, we ensure that $L^+(Z_{T-1} - X_{T-1}) > 0$ under π^{**}_Δ if the transitional step occurs. In this case of bounded demand, we may then infer that the expected discounted savings under π_Δ^{**} (relative to π^*) in the period immediately before the transitional step is nonnegative. In total, then, we have a positive net savings in expected discounted cost by using π_Δ^{**}.

Suppose the random variable X is not bounded above. If the transitional step occurs, then under π_Δ^{**} we have $I_0 - \Delta - 1 \leq Z_{T-1} \leq I_0 - 1$. Given $Z_{T-1} = Z$, the expected discounted savings under π_Δ^{**} (relative to π^*) in the period immediately before the transitional step is at least

$$\mathbb{E}[L^+(Z - X) \,|\, X > (Z - I_0 - \Delta - 1)]$$

Since L^+ is nondecreasing, the quantity above is at least

$$\mathbb{E}[L^+(I_0 - \Delta - 1 - X) \,|\, X > (Z - I_0 - \Delta - 1)]$$

By Assumption 3 and Lemma 12 (as applied to nonincreasing functions), the quantity above is at least

$$\mathbb{E}[L^+(I_0 - \Delta - 1 - X) \,|\, X > \Delta]$$

We are to take $I_0 \geq I^*$, and as $I^* \to +\infty$ we find $\mathbb{E}[L^+(I^* - \Delta - 1 - X) \,|\, X > \Delta] \to L^+(+\infty) > 0$. This may be seen by applying the monotone convergence theorem to the function $L^+(I^* - \Delta - 1 - X)\mathbf{1}(X > \Delta)$, where $\mathbf{1}$ is the indicator function, much as was done for Lemma 4. Thus we may fix I^* sufficiently high so that the expected discounted savings in the period of concern is nonnegative. In total, we again have a positive net savings by using π_Δ^{**}. □

In the following theorem we define the critical values that will play a role in the optimal solution to the unrelaxed problem. From this point on, we use a subscript '1' instead of 'β' to identify functions and critical values relevant to the undiscounted case. These functions and critical values are not to be confused with their finite-horizon counterparts when there is one decision period remaining.

Theorem 11. For all $i \in \{*, 1, \ldots, m\}$, choose s_β^i to be the greatest minimizer defined in Theorem 8. We have $s_\beta^i \nearrow s_1^i$ where s_1^i is some finite value, for all i as $\beta \to 1$ for $\beta \in (0,1)$. (We define $s_1^0 := +\infty$.) Furthermore, there exists a compact interval containing minimizers of the functional equation in Theorem 7 for all $\beta \in (0,1)$ sufficiently large.

Proof. By Theorem 4, there exists $\beta_0 < 1$ such that $s_{\beta_0,n'}^m$ is finite for some $n' \geq 1$. By the analysis of the derivatives of $f_{\beta,n}$ in Theorem 2, we have $f_{\beta_0,n}^-(I) = -c^m$ for all $I < s_{\beta_0,n}^m - r^{m-1}$. Since $s_{\beta_0,n'}^m \leq s_{\beta_0,n}^m$ for all $n > n'$ (by Theorem 3), the convergence of $f_{\beta_0,n}$ to f_{β_0} (by Theorem 6) implies via Lemma 11 that $f_{\beta_0}^-(I) = -c^m$ for all $I < s_{\beta_0,n'}^m - r^{m-1}$. It follows that $G_{\beta_0}^-(Z) < -c^m$ for sufficiently low Z, so $s_{\beta_0}^m$ as defined above is finite.

By Theorem 9, the values s_β^i defined above are all nondecreasing in β. By the structure of our rFGB policies and Theorem 10, there exists I^* such that, for all $\beta \geq \beta_0$

$$-\infty < s_{\beta_0}^m \leq s_\beta^m \leq s_\beta^{m-1} \leq \ldots \leq s_\beta^2 \leq s_\beta^1 \leq s_\beta^* < I^* < +\infty$$

Monotonicity and boundedness show the desired finite limits exist with $s_1^m \leq s_1^{m-1} \leq \ldots \leq s_1^2 \leq s_1^1 \leq s_1^*$. By the structure of our rFGB policies and Theorem 8, it also follows that, for each state I, a minimizer of the functional equation lies in $[s_{\beta_0}^m, s_1^*]$ for all $\beta \geq \beta_0$. □

For $\beta \in (0,1)$, we define the relative cost function $\bar{f}_\beta(I) := f_\beta(I) - f_\beta(s_1^*)$. The following theorem may be compared with Theorem 6 in the previous part.

Theorem 12. As $\beta \to 1$ for $\beta \in (0,1)$, $0 \leq \bar{f}_\beta \nearrow \bar{f}_1$ where $\bar{f}_1 : \mathbb{R} \to \mathbb{R}$ is convex with $\bar{f}_1^-(-\infty) \geq -c^m$ and $\bar{f}_1(I) = 0$ for $I \geq s_1^*$.

Proof. By Theorem 11, s_1^* is finite; by Theorem 6, so is $f_\beta(s_1^*)$ and hence $\bar{f}_\beta$. Therefore, we may directly obtain $\bar{f}_\beta^+(I) = f_\beta^+(I)$ and $\bar{f}_\beta^-(I) = f_\beta^-(I)$ for all I. From Theorem 6, then, $\bar{f}_\beta^-(-\infty) = f_\beta^-(-\infty) \geq -c^m$. Moreover, $\bar{f}_\beta$ inherits convexity from f_β. Using the properties of G_β shown in the proof of Theorem 8, an analysis of the derivatives of f_β as in the proof of Theorem 2 yields $f_\beta^+(I) = 0$ for $I \geq s_\beta^*$. Since (via Theorem 11) $s_\beta^* \leq s_1^*$ we have $\bar{f}_\beta(I) = 0$ for all $I \geq s_1^*$. We also have $f_\beta^-(I) \leq 0$ for $I \leq s_1^*$ by convexity, so Lemma 10 may be used to show that $f_\beta(I) \geq f_\beta(s_1^*)$ for all $I < s_1^*$. Putting the preceding two comments together, we conclude that $\bar{f}_\beta \geq 0$. All of the foregoing holds for all $\beta \in (0,1)$.

We may now argue by Theorem 9 that $\bar{f}_{\beta_1} \leq \bar{f}_{\beta_2}$ for $\beta_1 < \beta_2$. To see this, observe that $\bar{f}_{\beta_1} = \bar{f}_{\beta_2} = 0$ over $[s_1^*, +\infty)$ and $\bar{f}_{\beta_1}^- \geq \bar{f}_{\beta_2}^-$ over $(-\infty, s_1^*]$, each $\bar{f}_\beta$ being continuous, and apply Lemma 10 (in its "turned around" form) to $\bar{f}_{\beta_2}^- - \bar{f}_{\beta_1}^-$. $\bar{f}_\beta$ is also bounded above, since by the preceding paragraph we may write $\bar{f}(I) \leq c^m(s_1^* - I)^+$. This establishes the desired monotone convergence to the limiting function we call $\bar{f}_1$. Using Lemma 11, $\bar{f}_\beta^-(-\infty) \geq -c^m$ implies that $\bar{f}_1^-(-\infty) \geq -c^m$. Also, $\bar{f}_\beta(I) = 0$ implies $\bar{f}_1(I) = 0$ for $I \geq s_1^*$. □

Following the discounted case, we define $\bar{G}_\beta(Z) := \mathbb{E}[L(Z-X)] + \beta\mathbb{E}[\bar{f}_\beta(Z-X)]$ and $\bar{J}_\beta(I,Z) := C(Z-I) + \bar{G}_\beta(Z)$ for $\beta \in (0,1)$. Similarly, we define the functions $\bar{G}_1(Z) := \mathbb{E}[L(Z-X)] + \mathbb{E}[\bar{f}_1(Z-X)]$ and $\bar{J}_1(I,Z) := C(Z-I) + \bar{G}_1(Z)$.

By adding $f_\beta(s_1^*)$ to both sides of the functional equation in Theorem 7, we obtain for all $I \in \mathbb{R}$ and all $\beta \in (0,1)$ the modified functional equation

$$(1-\beta)f_\beta(s_1^*) + \bar{f}_\beta(I) = \min_{Z\in\mathbb{R}}\{C(Z-I) + \mathbb{E}[L(Z-X)] + \beta\mathbb{E}[\bar{f}_\beta(Z-X)]\}$$

which we will use to prove the following result.

Theorem 13. For all $I \in \mathbb{R}$:

- $(1-\beta)f_\beta(I) \to \rho$ as $\beta \to 1$ for some $\rho \in \mathbb{R}$.
- ρ and $\bar{f}_1$ satisfy the average cost optimality equation:

$$\rho + \bar{f}_1(I) = \min_{Z\in\mathbb{R}} \{C(Z-I) + \mathbb{E}[L(Z-X)] + \mathbb{E}[\bar{f}_1(Z-X)]\}$$

Proof. We will substantially follow the proof of Theorem 7. Using Theorem 12 as a basis instead of Theorem 6, we obtain $\bar{G}_\beta \nearrow \bar{G}_1$ and $\bar{J}_\beta \nearrow \bar{J}_1$ as $\beta \to 1$, where $\bar{G}_1(Z)$ and $\bar{J}_1(I,Z)$ are finite-valued and convex in Z for fixed I. Let $I \in \mathbb{R}$ be given. By Theorem 11, there is a compact interval φ such that a minimizer of $\bar{J}_\beta(I,Z)$ over $Z \in \mathbb{R}$ must exist in φ for β sufficiently large. By the convergences just noted, Lemma 9 implies that $\bar{J}_\beta(I,Z)$ converges uniformly to $\bar{J}_1(I,Z)$ on φ as $\beta \to 1$. Letting $\varepsilon > 0$ be given, there exists β' such that $\beta \geq \beta'$ implies both $\min_{Z\in\varphi} \bar{J}_\beta(I,Z) \leq \min_{Z\in\varphi} \bar{J}_1(I,Z) + \varepsilon$ and $\min_{Z\in\varphi} \bar{J}_\beta(I,Z) \geq \min_{Z\in\varphi} \bar{J}_1(I,Z) - \varepsilon$. Since ε can be arbitrarily small, we have $\lim_{\beta\to 1} \min_{Z\in\varphi} \bar{J}_\beta(I,Z) = \min_{Z\in\varphi} \bar{J}_1(I,Z)$, the latter term being finite, as it is the minimum of a continuous function over a compact interval. As we argued for Theorem 7, using the convexity of $\bar{J}_\beta$ and its minimization over φ along with Lemma 11 implies that $\min_{Z\in\varphi} \bar{J}_1(I,Z) = \min_{Z\in\mathbb{R}} \bar{J}_1(I,Z)$.

By our modified functional equation, $(1-\beta)f_\beta(s_1^*) = \min_{Z\in\mathbb{R}} \bar{J}_\beta(I,Z) - \bar{f}_\beta(I)$ for β sufficiently large. Taking $\beta \to 1$, we see $(1-\beta)f_\beta(s_1^*) \to \min_{Z\in\mathbb{R}} \bar{J}_1(I,Z) - \bar{f}_1(I) =: \rho$. This limit is finite, and it establishes the solution to the average cost optimality equation. Observe from Theorem 12 that $0 \leq (1-\beta)\bar{f}_\beta(I) \leq (1-\beta)c^m(s_1^* - I)^+$, and so $(1-\beta)\bar{f}_\beta(I) \to 0$ as $\beta \to 1$, which means that $((1-\beta)f_\beta(I) - (1-\beta)f_\beta(s_1^*)) \to 0$. Given our definition of ρ, this implies that $(1-\beta)f_\beta(I) \to \rho$ as desired. $\square$

The following theorem uses our solution of the average cost optimality equation to show the average optimality of an rFGB policy.

Theorem 14. The critical values $\{s_1^0, \ldots, s_1^m\}$ and s_1^* defined in Theorem 11, together with the given parameters $\{r^0, \ldots, r^m\}$, characterize an rFGB policy π^* that is average optimal for the relaxed model. Also, for any initial inventory level $I_0 \in \mathbb{R}$, we have

$$\lim_{n\to+\infty} \frac{1}{n+1} \sum_{t=0}^{n} \mathbb{E}_{I_0}^{\pi^*}[C(Z_t - I_t) + L(Z_t - X_t)] = \rho$$

Proof. We first argue that π^* satisfies the average cost optimality equation of Theorem 13 along with ρ and $\bar{f}_1$. Observe from the proof of Theorem 13 that $\bar{G}_1$ is finite-valued and convex. From Theorem 12, Lemmas 2–4, and Assumptions 1–3, we can see (similar to the proofs of Theorems 2 and 8) that $\bar{G}_1^-(Z) < 0$ and $\bar{G}_1^+(Z) > 0$. The proof of Theorem 1 still works if we take $\beta = 1$, so an rFGB decision rule with critical values minimizing $c^i Z + \bar{G}_1(Z)$ for $i \in \{*, 1, \ldots, m\}$ will achieve the minimum in the average cost optimality equation. Since $\bar{G}_\beta \nearrow \bar{G}_1$ as $\beta \to 1$ (from the proof of Theorem 13), Lemma 8 implies that minimizers of $c^i Z + \bar{G}_\beta(Z)$ will converge to minimizers of $c^i Z + \bar{G}_1(Z)$ if they converge at all. By Theorem 11, then, s_1^i minimizes $c^i Z + \bar{G}_1(Z)$ for all $i \in \{*, 1, \ldots, m\}$.

Consider any admissible policy π. Given initial state I_0, by Theorem 8 we have for all $\beta \in (0, 1)$ that:

$$\sum_{t=0}^{+\infty} \beta^t \mathbb{E}_{I_0}^{\pi}[C(Z_t - I_t) + L(Z_t - X_t)] \geq f_\beta(I_0)$$

Multiplying both sides of the above by $(1-\beta)$, it follows from Theorem 13 that

$$\limsup_{\beta\to1}(1-\beta) \sum_{t=0}^{+\infty} \beta^t \mathbb{E}_{I_0}^{\pi}[C(Z_t - I_t) + L(Z_t - X_t)] \geq \limsup_{\beta\to1}(1-\beta) f_\beta(I_0) = \rho$$

Applying Lemma 13,

$$\limsup_{t\to+\infty} \frac{1}{t+1} \sum_{t=0}^{+\infty} \mathbb{E}_{I_0}^{\pi}[C(Z_t - I_t) + L(Z_t - X_t)] \geq \rho$$

This means that ρ is a lower bound on the (lim sup) average expected cost. Now given our argument at the outset, we have for all t:

$$\rho + \bar{f}_1(I_t) = \mathbb{E}_{H_t}^{\pi^*}[C(Z_t - I_t) + L(Z_t - X_t)] + \mathbb{E}_{H_t}^{\pi^*}[\bar{f}_1(I_{t+1})]$$

where H_t represents the information available at the point of decision in period t (defined in the proof of Theorem 8). Given any initial state I_0, we may sum the equations corresponding to $t \in \{0 \dots n\}$, divide the resulting equation by $(n+1)$, and take expectations on both sides (based on π^* and I_0) to obtain

$$\frac{1}{n+1}\sum_{t=0}^{n} \mathbb{E}_{I_0}^{\pi^*}[C(Z_t - I_t) + L(Z_t - X_t)] = \rho + \frac{1}{n+1}\bar{f}_1(I_0) - \frac{1}{n+1}\mathbb{E}_{I_0}^{\pi^*}\big[\bar{f}_1(I_{n+1})\big]$$

Consider the right side of this equation as $n \to +\infty$. The first term is a constant. Since $\bar{f}_1$ is finite-valued, the second term converges to 0. Given the form of $\bar{f}_1$ (from Theorem 12) and π^*, we note that $\bar{f}_1(I_{n+1})$ is nonnegative and less than $\bar{f}_1(s_1^m - X_n)$ for all realized X_n. Hence the expectation in the third term is nonnegative and less than $\mathbb{E}[\bar{f}_1(s_1^m - X_n)]$, which is finite by Theorem 12 and the fact that $\mathbb{E}[X] < +\infty$. The third term therefore converges to 0. We conclude that the desired limit indeed holds, and this establishes the average optimality of π^*. □

We now relate our relaxed model with $Z \in \mathbb{R}$ to the unrelaxed model with the constraint $Z \geq I$, given what we have proved in Theorem 14. This accomplishes the primary mathematical goal of the chapter.

Theorem 15. The critical values $\{s_1^0, \dots, s_1^m\}$ defined in Theorem 11 (where all but $s_1^0 := +\infty$ were identified to be finite), together with the given parameters $\{r^0, \dots, r^m\}$, characterize an FGB policy $\bar{\pi}$ that is average optimal for the relaxed model, and is therefore average optimal for the unrelaxed model as well.

Proof. The FGB policy $\bar{\pi}$ specified differs from the average optimal rFGB policy of

Theorem 14 (which we again call π^*) only in that $Z = I$ is chosen when the inventory level is observed to be $I > s_1^*$, rather than choosing $Z = s_1^*$. Observe that for initial states $I_0 \leq s_1^*$, the two policies will behave equivalently at every decision point as the system operates. This is because an inventory level $I > s_1^*$ will never be seen given such an initial state, since $Z > s_1^*$ is never chosen by these policies for observed inventory levels $I \leq s_1^*$. The long-run average expected cost of $\bar{\pi}$ for initial states $I_0 \leq s_1^*$ is therefore ρ, the same as the long-run average expected cost of π^* identified in Theorem 14.

We now argue that, for any given initial state $I_0 > s_1^*$, the long-run average expected cost of $\bar{\pi}$ is also ρ. Define T to be a random variable indicating the first decision period for which $I_t < s_1^*$ under policy $\bar{\pi}$. Since $Z_t = I_t$ for all $t < T$ under this policy, T is determined by the properties $\sum_{t=0}^{T-2} X_t \leq (I_0 - s_1^*)$ and $\sum_{t=0}^{T-1} X_t > (I_0 - s_1^*)$. Because Assumption 4 tells us that $P(X = 0) < 1$, a standard theorem of renewal theory (Grimmett and Stirzaker 2001, p. 412) implies that T is finite with probability one. Furthermore, I_T will be finite with probability one. Given T and I_T, both finite, the cost incurred in periods 0 through $T - 1$ is bounded above by TM, where $M := \max\{L(I) : I \in [I_T, I_0]\}$. By Assumption 1, M is finite, and we also have a lower bound of 0 on the cost incurred through period $T - 1$. Given T and I_T, both finite, because the action of $\bar{\pi}$ depends only on the current inventory level, the cost incurred in any period $t \geq T$ is probabilistically the same as the cost incurred T steps earlier in a system with initial state I_T. Therefore, given T and I_T as noted, the sequence of averages $\frac{1}{n+1} \sum_{t=0}^{n} \mathbb{E}_{I_0}^{\bar{\pi}}[C(Z_t - I_t) + L(Z_t - X_t)]$ is bounded above for all $n \geq T$ by the sequence

$$\tfrac{1}{n+1} TM + \left(\tfrac{n-T+1}{n+1}\right)\left(\tfrac{1}{n-T+1}\right) \textstyle\sum_{t=0}^{n-T} \mathbb{E}_{I_T}^{\bar{\pi}}[C(Z_t - I_t) + L(Z_t - X_t)] \to \rho$$

and bounded below for all $n \geq T$ by the sequence

$$\left(\tfrac{n-T+1}{n+1}\right)\left(\tfrac{1}{n-T+1}\right)\sum_{t=0}^{n-T}\mathbb{E}_{I_T}^{\bar{\pi}}[C(Z_t - I_t) + L(Z_t - X_t)] \to \rho$$

where in taking limits we have used the fact that $I_T < s_1^*$. We conclude that the long-run average expected cost of the policy $\bar{\pi}$ must be ρ as desired.

We may now conclude that the FGB policy $\bar{\pi}$ is average optimal for the relaxed model: it performs as well as the average optimal rFGB policy π^* with respect to our optimality criterion. Furthermore, $\bar{\pi}$ is feasible for the unrelaxed model, as it never chooses a negative order quantity. Since any admissible policy that is feasible for the unrelaxed model is also feasible for the relaxed model, there cannot exist an admissible policy outperforming $\bar{\pi}$ in the unrelaxed model with respect to our optimality criterion. Such a policy would have to outperform π^* in the relaxed model, which by Theorem 14 is impossible. □

1.4.5 Extension to the discrete case

In this section, we extend our main result to the discrete case in which demands and order quantities may take only nonnegative integral values.

First, we may assume without loss of generality that the initial inventory level is an integer. Consequently, we may posit that the state of the process is always integral. This assumption is justified because we can shift our holding and backlogging cost function L by the same amount required to shift a fractional initial inventory level to reach an integral value. Policies in this "shifted" model then have the same cost characteristics as "unshifted" policies in the original model, and furthermore an "unshifted" FGB policy is still of the FGB type (featuring "unshifted" critical values).

We may also assume without loss of generality that the source capacity parameters $\{r^1, \ldots, r^{m-1}\}$ take integral values. This is because, if some parameters r^i

are fractional, we can replace the ordering cost function C with another function $\bar{C} : \mathbb{R}^+ \to \mathbb{R}^+$ which is equal to C for integral order quantities, but which is linear on the interval $[Q, Q+1]$ for each nonnegative $Q \in \mathbb{Z}$. By the nature of the discrete case, this substitution makes no difference to the cost incurred by a given policy. ($\bar{C}$ may, however, have a larger or smaller number of linear pieces than C.)

Given our assumption that the initial inventory level is integral, we may similarly replace the holding and backlogging cost function L by $\bar{L} : \mathbb{R} \to \mathbb{R}^+$, which is defined to be equal to L for integral inventory levels, but which is linear on the interval $[I, I+1]$ for each $I \in \mathbb{Z}$. By the nature of the discrete case, this substitution likewise makes no difference to the cost incurred by a given policy. Furthermore, $\bar{L}$ inherits the satisfaction of Assumptions 1–3 from L. In particular, Assumption 1 may be justified by Lemma 6: it is clear that $\bar{L}$ is continuous and that $\bar{L}^+$ is constant on $[I, I+1)$ for each $I \in \mathbb{Z}$; furthermore, by Lemma 1 we have $\bar{L}^+(I-1) \leq L^-(I) \leq L^+(I) \leq \bar{L}^+(I)$. Given this, Assumption 2 follows from noting that, as with L, there must exist reals I' and I'' such that $\bar{L}^-(I') < 0$ and $\bar{L}^+(I'') > 0$. It remains to deal with Assumption 3. Let $Z^* \in \mathbb{R}$ be given. Using Assumption 2 just shown, define Z' to be the greatest integer less than or equal to $\min\{Z : \bar{L}(Z) = \bar{L}(Z^*)\}$. Using the convexity of $\bar{L}$, we find for any (integral) realized demand X that $\bar{L}(Z^* - X) \leq \bar{L}(Z' - X) = L(Z' - X)$. That $\mathbb{E}[\bar{L}(Z^* - X)] < +\infty$ follows by monotonicity of expectation.

Let us designate by $\mathcal{F}$ the class of functions $\mathbb{R} \to \mathbb{R}$ exemplified by $\bar{L}$, defining the class to consist of those continuous functions that are linear on the interval $[x, x+1]$ for each $x \in \mathbb{Z}$. Observe that the ordering cost function $\bar{C}$ belongs to $\mathcal{F}$ when we extend its domain to the negative real line, where it is defined to be equal to zero in the context of our relaxed model (i.e., the model allowing negative order quantities).

Key facts about $\mathcal{F}$ are as follows. The class is closed with respect to addition, addition of real constants being a special case. Multiplication by real constants also preserves membership in $\mathcal{F}$, as does shifting of functions (left or right) by fixed integral amounts. Finite limits of sequences of functions in $\mathcal{F}$ are also in $\mathcal{F}$. Given $g, h \in \mathcal{F}$, we may conclude that $g(h(x)) \in \mathcal{F}$, if h maps integers to integers and the linear pieces of h have slopes among $\{-1, 0, 1\}$ only. If a function in $\mathcal{F}$ has a greatest minimizer, then the greatest minimizer is an integer.

Consider now the main argument of Sections 1.4.2–1.4.4, when we are assured $C, L \in \mathcal{F}$ as justified above, ignoring the requirement that order quantities be integral. In Theorem 1, we find if $G_n \in \mathcal{F}$ that $c^i Z + G_n(Z)$ is in $\mathcal{F}$ for each $i \in \{*, 1, \ldots, m\}$, and when any of these functions has a minimum its greatest minimizer is integral. In Theorem 2, we may argue that $G_n \in \mathcal{F}$ and $f_n \in \mathcal{F}$ for all $n \geq 1$. We first argue that $G_n \in \mathcal{F}$ if $f_{n-1} \in \mathcal{F}$. We find that $L(Z - X) \in \mathcal{F}$ for all realized X, and so $\mathbb{E}[L(Z-X)]$ belongs to $\mathcal{F}$ as a sum of functions in $\mathcal{F}$ or as a finite limit of such sums. Similarly, by hypothesis $\mathbb{E}[f_{n-1}(Z - X)]$ must be in $\mathcal{F}$, and now $G_n \in \mathcal{F}$ follows. We next argue that $f_n \in \mathcal{F}$ if $G_n \in \mathcal{F}$. Observe that

$$f_n(I) = C(Z_n^*(I) - I) + G_n(Z_n^*(I))$$

where $Z_n^*(I)$ is a function of I assigning order-up-to levels in accordance with the rFGB policy specified by the critical values s_n^i that are integral (whenever finite) as we have established. We find that $Z_n^*(I) \in \mathcal{F}$, that this function maps integers to integers, and that its linear pieces have slopes among $\{0, 1\}$ only. Additionally, $(Z_n^*(I) - I) \in \mathcal{F}$ maps integers to integers and its linear pieces have slopes among $\{-1, 0\}$ only. Since $C \in \mathcal{F}$, and since $G_n \in \mathcal{F}$ by hypothesis, we may conclude that $f_n \in \mathcal{F}$ as desired. Now, since our terminal value function $f_0 := 0$ is in $\mathcal{F}$, the desired memberships in $\mathcal{F}$ follow for all $n \geq 1$.

We turn to the infinite-horizon setting. By Theorem 6, $f_n \nearrow f$ where f is finite-valued, so we have $f \in \mathcal{F}$. Since G is defined in terms of f in the same way that G_n is defined in terms of f_{n-1}, the fact that $G \in \mathcal{F}$ follows by an argument analogous to that indicated for the finite-horizon setting. Also like the finite-horizon setting, we find in Theorem 8 that, by choosing the greatest minimizer of $c^i Z + G(Z)$ where possible, the critical values s^i are integral (whenever finite). From the perspective of the vanishing discount approach, these critical values s^i_β define a discount optimal rFGB policy when the discount factor is $\beta \in (0, 1)$. By Theorem 11, for each $i \in \{*, 1, \ldots, m\}$ we have $s^i_\beta \nearrow s^i_1$ as $\beta \to 1$, each limit being finite. By the integrality of each s^i_β for $i \in \{*, 1, \ldots, m\}$, each of these limiting values s^i_1 must be integral. In Theorem 15, then, the FGB policy that is average optimal for the unrelaxed model is specified by the integral critical values $\{s^1_1, \ldots, s^m_1\}$ as well as the integral parameters $\{r^0, \ldots, r^{m-1}\}$ (and, by definition, $s^0_1 = r^m = +\infty$). Since we are assured that the initial inventory level is integral (as justified above), it follows that the order quantity chosen by this FGB policy will be integral at each decision point. Since this policy is average optimal in the absence of the constraint mandating integral order quantities, this same policy is also average optimal when the constraint is imposed.

1.5 Conclusion

In this chapter, we have proved the existence of a finite generalized base stock policy that is optimal under an average cost criterion, for an inventory control model involving multiple sources of a single product. Our main argument concentrated on a continuous setting; we subsequently indicated how this argument can also be used to establish the result in a discrete case. Building on this work, one potentially interesting research direction would involve integration of our methods with other mathematical work on average optimal control in Markov decision processes (particularly including other inventory models). Such integration may enable streamlining of our main argument. Relatedly, methods we have employed (such as examination of the derivatives of discounted value functions within a vanishing discount approach) may also find wider use in proving new results or simplifying proofs of known results. Research establishing the degree to which our work is relevant for practical applications would also be most welcome. Such research may warrant investigation of methods for computing the critical values we have shown to exist.

1.6 References

Arapostathis, A., V. S. Borkar, E. Fernández-Gaucherand, M. K. Ghosh, S. I. Marcus. 1993. Discrete-time controlled Markov processes with average cost criterion: A survey. *SIAM Journal on Control and Optimization* **31**(2) 282–344.

Bartle, R. G. 1976. *The Elements of Real Analysis,* 2nd ed. Wiley, New York.

Beckmann, M. J. 1961. Production smoothing and inventory control. *Operations Research* **9**(4) 456–467.

Bensoussan, A., M. Crouhy, J.-M. Proth. 1983. *Mathematical Theory of Production Planning.* Elsevier Science, Amsterdam.

Beyer, D., S. P. Sethi. 1999. The classical average-cost inventory models of Iglehart and Veinott-Wagner revisited. *Journal of Optimization Theory and Applications* **101**(3) 523–555.

Beyer, D., F. Cheng, S. P. Sethi, M. Taksar. 2009. *Markovian Demand Inventory Models.* Springer, Secaucus, NJ. Forthcoming.

Bulinskaya, E. V. 1967. Optimum inventory policies with a convex ordering cost function. *Theory of Probability and its Applications* **12**(1) 9–21.

Chen, X., D. Simchi-Levi. 2004. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The infinite horizon case. *Mathematics of Operations Research* **29**(3) 698–723.

Feinberg, E. A., A. Shwartz, eds. 2002. *Handbook of Markov Decision Processes: Methods and Applications.* Kluwer, Boston, MA.

Feinberg, E. A., M. E. Lewis. 2006. Optimality inequalities for average cost Markov decision processes and the optimality of (s, S) policies. Technical Report TR1442, Cornell University, Ithaca, NY. `http://legacy.orie.cornell.edu/~melewis/pubs/mdp-ss-tech-report.pdf` Accessed May 23, 2009.

Flett, T. M. 1980. *Differential Analysis: Differentiation, Differential Equations, and Differential Inequalities.* Cambridge University Press, Cambridge, UK.

Geunes, J. 1999. *Models to Optimize Multi-Stage Linkages in Manufacturing and Distribution Operations.* Ph.D. dissertation, Smeal College of Business Administration, Pennsylvania State University, University Park, PA.

Grimmett, G. R., D. R. Stirzaker. 2001. *Probability and Random Processes,* 3rd ed. Oxford University Press, Oxford, UK.

Henig, M., Y. Gerchak, R. Ernst, D. F. Pyke. 1997. An inventory model embedded in designing a supply contract. *Management Science* **43**(2) 184–189.

Hernández-Lerma, O., J. B. Lasserre. 1996. *Discrete-Time Markov Control Processes: Basic Optimality Criteria.* Springer, New York.

Heyman, D. P., M. J. Sobel. 1984. *Stochastic Models in Operations Research, Volume II: Stochastic Optimization.* McGraw-Hill, New York.

Hiriart-Urruty, J.-B., C. Lemaréchal. 1993. *Convex Analysis and Minimization Algorithms I: Fundamentals.* Springer, Berlin.

Huh, W. T., G. Janakiraman, M. Nagarajan. 2008. Average cost inventory models: An analysis using a vanishing discount approach. Working paper dated September 15, 2008.

Iglehart, D. L. 1963. Dynamic programming and stationary analysis of inventory problems. H. E. Scarf, D. M. Gilford, M. W. Shelly, eds. *Multistage Inventory Models and Techniques,* Stanford University Press, Stanford, CA, 1–31.

Karlin, S. 1958. Optimal inventory policy for the Arrow-Harris-Marschak dynamic model. K. J. Arrow, S. Karlin, H. Scarf, eds. *Studies in the Mathematical Theory of Inventory and Production,* Stanford University Press, Stanford, CA, 135–154.

Kleindorfer, P., H. Kunreuther. 1978. Stochastic horizons for the aggregate planning problem. *Management Science* **24**(5) 485–497.

Liu, B., A. O. Esogbue. 1999. *Decision Criteria and Optimal Inventory Processes.* Kluwer Academic, Boston, MA.

Porteus, E. L. 1990. Stochastic inventory theory. D. P. Heyman, M. J. Sobel, eds. *Stochastic Models,* Elsevier Science, Amsterdam, 605–652.

Porteus, E. L. 2002. *Foundations of Stochastic Inventory Theory.* Stanford University Press, Stanford, CA.

Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming.* Wiley, New York.

Rockafellar, R. T. 1970. *Convex Analysis.* Princeton University Press, Princeton, NJ.

Schäl, M. 1993. Average optimality in dynamic programming with general state space. *Mathematics of Operations Research* **18**(1) 163–172.

Sennott, L. I. 1999. *Stochastic Dynamic Programming and the Control of Queueing Systems.* Wiley, New York.

Serel, D. A., M. Dada, H. Moskowitz. 2001. Sourcing decisions with capacity reservation contracts. *European Journal of Operational Research* **131** 635–648.

Sethi, S., H. Zhang, Q. Zhang. 2005. *Average-Cost Control of Stochastic Manufacturing Systems.* Springer, New York.

Sobel, M. J. 1969. Production smoothing with stochastic demand I: Finite horizon case. *Management Science* **16**(3) 195–207.

Sobel, M. J. 1970. Making short-run changes in production when the employment level is fixed. *Operations Research* **18**(1) 35–51.

Sobel, M. J. 1971. Production smoothing with stochastic demand II: Infinite horizon case. *Management Science* **17**(11) 724–735.

Vega-Amaya, O., R. Montes-de-Oca. 1998. Application of average dynamic programming to inventory systems. *Mathematical Methods of Operations Research* **47** 451–471.

Veinott, A. F. 1964. Production planning with convex costs: A parametric study. *Management Science* **10**(3) 441–460.

Veinott, A. F., H. M. Wagner. 1965. Computing optimal (s, S) inventory policies. *Management Science* **11**(5) 525–552.

Veinott, A. F. 1966. The status of mathematical inventory theory. *Management Science* **12**(11) 745–777.

Webster, R. 1994. *Convexity.* Oxford University Press, Oxford, UK.

Wheeden, R. L., A. Zygmund. 1977. *Measure and Integral: An Introduction to Real Analysis.* Marcel Dekker, New York.

Yang, J., X. Qi, Y. Xia. 2005. A production-inventory system with Markovian capacity and outsourcing option. *Operations Research* **53**(2) 328–349.

Zheng, Y.-S. 1991. A simple proof for optimality of (s, S) policies in infinite-horizon inventory systems. *Journal of Applied Probability* **28** 802–810.

Zipkin, P. H. 2000. *Foundations of Inventory Management.* McGraw-Hill, Boston, MA.

CHAPTER 2: OPTIMAL COMPOSITION OF A RETAIL DISTRIBUTION FLEET WHEN SPOT CAPACITY IS AVAILABLE

2.1 Introduction

We address the fleet composition problem faced by a retail distribution firm, with particular attention to a major distributor of beverage products. Every working day, the distributor transports merchandise from a distribution center to a great number of customer sites using trucks of several sizes. Demand becomes known to the distributor just before the requested day of delivery, and the firm strives for a consistently high service level in which deliveries are not delayed. Vehicle routes are not fixed. The distributor experiences daily demand that is variable in terms of the total cubic volume as well as the average drop size. Customers are heterogeneous, ranging from large hypermarkets to small shops or restaurants. A customer's demand on a given day generally does not necessitate a particular size of vehicle to satisfy the demand. Common carriers can be hired on a spot basis to supplement the fleet as needed on a given day. We take the planning horizon to be one year in length for illustrative purposes. The goal is to balance the fixed cost of vehicles to be owned through the planning horizon against exposure to a high variable cost due to reliance on spot capacity. (We streamline our discussion in this chapter by referring to "owned" vehicles, though our reasoning is intended to apply similarly to typical leasing arrangements.)

To assist the distributor's decision of the proper fleet size and mix for next year, we propose an optimization model. The model revolves around the description of a day's demand in terms of the total cubic volume to be delivered and the total number of customers to visit. We submit that this compact characterization is useful for capturing the capabilities of different sizes of vehicles. Consider a specific, substantial cubic volume of merchandise for delivery on a particular day. If this demand is concentrated among only a few customer sites, then in the interest of minimizing variable cost we would tend to prefer utilizing a few large vehicles from our fleet. (Large vehicles tend to be more efficient than small vehicles in terms of variable cost, e.g. cost due to fuel and driver-hours, when moving a great volume of merchandise from one fixed location to another.) If the same cubic volume of demand is instead scattered evenly across a great many customer sites, then we are inclined to "parallel process" the deliveries by utilizing many small vehicles. In this chapter, we introduce a two-stage stochastic linear program with fixed recourse that is based on this two-dimensional characterization of demand, and we present an efficient algorithm facilitating solution of this model by a standard gradient-based method for continuous optimization.

There is a substantial current of published research concerning fleet composition decisions and the closely related problem of fleet sizing, incorporating a great variety of application contexts and modeling approaches. Some researchers take a more detail-oriented approach to modeling the day-to-day operations underlying the strategic decision at hand (e.g. Dell'Amico et al. 2007, Simão et al. 2008). Such models typically involve complex vehicle routing subproblems, and they are likely to be daunting in terms of computational effort and data requirements—particularly when considering a large customer base for deliveries, for which routes are not fixed from day to day, under a long planning horizon. Other models are less concerned with the intricacies of vehicle movements, likely making implementation easier (e.g. Wyatt 1961, Papier and Thonemann 2008).

In the wider context of capacity planning, there is a recent stream of literature introducing rather abstract capacity decision models whose applicability to fleet composition problems has been little studied. Eberly and Van Mieghem (1997) study optimal adjustments of multiple factor levels in a multiperiod, stochastic setting. Van Mieghem and Rudi (2002) introduce a multidimensional capacity investment model involving a more specific, but still quite general, model of resource processing. This more specific model is a recourse linear program within a two-stage stochastic program involving uncertain demand variables in the right-hand side; following earlier work, the authors illustrate a solution approach for the investment problem based on a decomposition of the space of possible demands.

In the domain of two-stage stochastic linear programming with fixed recourse, the decomposition just mentioned is fairly familiar (e.g. Birge and Louveaux 1997, pp. 170–171). The space of possible demands is dissected into regions such that the same basis is optimal throughout any given region. Such a construction facilitates various means of solution of the stochastic program; if the demand space has low dimension, then numerical integration might then be used to compute the first-stage cost and its gradient (or a subgradient). Essentially the same kind of decomposition is sought in parametric linear programming (e.g. Jones 2005).

Our approach brings together two key elements that have appeared only separately in the literature on fleet composition: the use of a two-stage stochastic programming formulation, and the description of a day's demand by the total quantity to be delivered and the total number of customers to visit. This synthesis allows us to capture our distributor's daily demand over an entire planning horizon by a bivariate probability distribution. In practice, this opens up the possibility of identifying a suitable distribution family, for present purposes reducing the task of forecasting a year's worth of daily demand to the estimation of a handful of distribution parameters. In the context of capacity decision models, we offer a specialized recourse structure for fleet composition within the framework of Eberly and Van Mieghem (1997) and Van Mieghem and Rudi (2002), and we offer an efficient algorithm producing a decomposition of the demand space given any number of vehicle types. This algorithm actually generates what we call a "definitive" collection of bases, which correspond to a dissection of the demand space facilitating solution of the stochastic program in the manner suggested above. In the context of stochastic linear programming and parametric linear programming, we name and explicitly define (for our model) the concept of a definitive collection of bases—which seems to have been only implicit in prior work—and we provide a tailored algorithm generating such a collection when the (recourse) linear program consists of minimization subject to two positive fractional covering constraints with variable upper bounds.

Our overarching contribution in this chapter, then, is a relatively simple and fast solution approach for the fleet composition problem faced by our prototypical distributor. Our research contributions in support of this overarching goal pertain to the type of model we propose, which is new in the context of fleet composition, and to our proposed method of solution. Our solution method, a standard gradient-based approach at surface level, relies on our specialized algorithm to facilitate rapid computation of the expected cost of a prospective fleet composition and its gradient. We articulate the concept of a "definitive" collection of bases, over whose feasible regions we perform double integration to obtain the expected second-stage cost or its gradient for any fleet composition. Our recourse linear program has exponentially many bases, but our algorithm generates a definitive set of $O(n^2)$ bases in $O(n^3)$ time, where n is the number of vehicle types.

The rest of the chapter is organized as follows. In Section 2.2, we review relevant literature in the contexts of fleet management, capacity planning, and optimization. In Section 2.3, we present our model, we make explicit certain key propositions underlying our model, and we elaborate on our model's relationship to newsvendor-type models in prior literature. In Section 2.4, we precisely define our concept of a definitive collection of bases, indicate the usefulness of such collections in solving our model, and present and validate our algorithm for generating definitive collections for our model. We offer concluding remarks in Section 2.5 and list references in Section 2.6.

2.2 Literature review

2.2.1 Fleet composition and related problems

Fleet composition has been an object of study in the operations research and management science literature since at least the 1950s. Early published work on this subject includes Kirby (1959) and, arguably, Dantzig and Fulkerson (1954). Studies have addressed fleets of cargo ships (Sigurd et al. 2005), warships (Crary et al. 2002), various classes of trucks (Gould 1969, Woods and Harris 1979, Ball et al. 1983, Wu et al. 2005), buses (Ceder 2005), locomotives (Nahapetyan et al. 2007, Godwin et al. 2008), cargo aircraft (Barnhart and Schneur 1996), airliners (Listes and Dekker 2005, Clark 2007), and automated guided vehicles (Hall et al. 2001, Vis et al. 2005), as well as rail freight cars (Papier and Thonemann 2008), different kinds of containers (Turnquist and Jordan 1986, Imai and Rivera 2001), barges and tugboats (Richetta and Larson 1997), and more. Some, such as List et al. (2003), have proposed general approaches to the problem that do not specify a particular mode of transportation. Some researchers approach the problem from the standpoint of drivers (i.e., the number and type of vehicle operators to employ), for instance Simão et al. (2008). Papers featuring substantial review and efforts at classification of literature in this area include Etezadi and Beasley (1983), Turnquist (1985), and Żak et al. (2008).

Key phrases in this current of research include "fleet size" (or "fleet sizing"), "fleet mix," "fleet composition," "fleet size and mix," "fleet planning," and "fleet design." "Fleet composition," the label that we adopt, has been inconsistently distinguished from "fleet sizing" in the literature, but the following definitions are in accord with much published work. A *fleet composition problem* is a decision problem where, for some set of vehicle types, we are to specify the number of vehicles of each type to include in the fleet. A *fleet sizing problem* or *fleet size problem* is one where we are merely to specify the total number of vehicles in the fleet. *Fleet size and mix* we take to be in the spirit of fleet composition. The clearest example of a fleet sizing problem would be one where there is a single vehicle type under consideration, and in a mathematical sense such a problem is obviously a special case of fleet composition. Note that, for a given fleet composition problem, it may well turn out that it is optimal to include *no* vehicles of a given type in the fleet. (Our definitions may be contrasted with those offered in Etezadi and Beasley 1983; cases where our classifications disagree with other authors' self-classifications include Fagerholt 1999 and Żak et al. 2008.)

Many authors' approaches to fleet composition involve detailed modeling of the operational level, including the routing of vehicles. Recent projects of this kind include Nahapetyan et al. (2007) for a major railroad company (involving rapid simulation of "trains, locomotives, terminals, and shops in an integrated framework" over many months of operation) and Simão et al. (2008) for a major truckload carrier (simulating "at a high level of detail the movements of over 6,000 drivers"). There is also a fairly standard variant of the classic vehicle routing problem that is designed for fleet composition, the *fleet size and mix vehicle routing problem.* This model was introduced in Golden et al. (1984a), and work since then is surveyed in Renaud and Boctor (2002). More recently, Dell'Amico et al. (2007) attack a generalization of this problem incorporating delivery time windows specific to each customer.

In our prototypical beverage distributor's context, a detail-oriented, routing-based approach is a daunting proposition: the customer base is large; there are several vehicle types; the planning horizon is long; demand is uncertain, and exhibits seasonal as well as day-to-day variation; routes are not fixed. Even though there are techniques aimed at lowering the computational burden in a routing-based approach (Etezadi and Beasley 1983, Golden et al. 1984b) or simplifying the spatial representation of demand (Klincewicz et al. 1990), the broad nature of the firm's business concerns (encompassing inventories and marketing as well as transportation of merchandise to customers) render simpler approaches especially attractive.

Our proposed model brings together two key elements that have appeared only separately in the literature on fleet composition. The first element is the use of a two-stage stochastic programming formulation. List et al. (2003) and Couillard and Martel (1990) are relatively recent examples of this type of model; the letter of Kirby (1959) can be considered their ancestor. Kirby (1959) suggests, in essence, a classic newsvendor approach to the problem of fleet sizing. He models demand by a random variable indicating the number of wagons (rail freight cars) required; wagons are hired when demand exceeds the fleet size. Wyatt (1961) and Alsbury (1972) extend this approach and suggest ways of approaching the fleet *composition* problem while keeping a one-dimensional characterization of demand. In our context, however, it not clear how to infer the individual utilization levels of the variously sized trucks given a scalar representation of demand. (Recall that our distributor's demand is variable in terms of the average daily drop size as well as the total cubic volume, and that the customer base is heterogeneous.)

The second key element of our model is the description of a day's demand by the total cubic volume to be delivered and the total number of customers to visit. This perspective appears to originate in the fleet composition approach of Eilon et al. (1971a, pp. 234–236). Subsequent models suggested by Etezadi and Beasley (1983) feature similarly compact characterizations of demand. These two papers offer deterministic mixed integer linear programming formulations that ask for a serial, daily-level forecast of demand across the planning horizon. However, they do not explicitly recognize that the serial order of demand is immaterial under their assumptions. Adopting a two-stage stochastic programming perspective, we open the door to an additional level of concision—capturing demand by a distribution function—while also accommodating uncertainty. We also simplify matters by dropping the requirement of integrality, which we believe to be justified considering factors such as the scale of our distributor's demand and the strategic nature of the decision at hand.

In closing, we note that we are leaving aside discussion of approaches to fleet composition and related problems based on certain continuous approximation methods (Diana et al. 2006, Smilowitz and Daganzo 2007) and queueing theory (Papier and Thonemann 2008, Żak et al. 2008). Useful as these techniques may be to enable simple implementations in this problem domain, we are not aware of their use in the literature to address our prototypical distributor's context. We leave elaboration of the usefulness of these techniques for our context as a direction for future research.

2.2.2 Capacity investment

This area of research is reviewed rather broadly in Van Mieghem (2003); we focus on a chain of literature originating with Eberly and Van Mieghem (1997). They introduce a model of investment in multiple resources in which the decisions concern the adjustment of resource levels across multiple periods in a stochastic environment. They argue that a particular type of control limit policy maximizes expected discounted profit when the cost of adjusting each individual resource is a kinked piecewise linear convex function and the operating profit function is concave in the resource levels (given any realized state of the world).

Our model qualifies as an instance of the Eberly and Van Mieghem (1997) model in which there is only one decision period, where the multiple resources at issue are the different vehicle types under consideration for inclusion in the fleet. In our beverage distributor's context, business concerns beyond transportation distract attention from the problem of planning the composition of the fleet multiple periods ahead, though future extension of our model to a multiperiod setting may nonetheless be useful. According to the authors, the optimal control limit policy (in our case, the optimal number of vehicles of each type to purchase) may be defined in terms of the gradient of the expected profit function. However, the highly abstract nature of their resource processing model precludes detailed advice for our context. Under our assumptions, their characterization of the optimal policy essentially amounts to the usual first-order condition for gradient-based convex optimization, and they offer no advice on how to calculate our gradient. Moreover, their focus on an abstract resource adjustment decision precludes detailed advice on modeling our resource processing mechanism. (The only explicit multiperiod use of their framework in a fleet-composition-related context appears to be a car rental fleet sizing model in Angelus and Porteus 2002.)

Within the framework of this 1997 paper, Harrison and Van Mieghem (1999) study a model in which the operating profit function is specified to be the solution of a linear program of the product mix type. This recourse program models production decisions taken on the basis of prior capacity investment decisions (i.e., available resources) and realized demand for multiple products. Given "kinked" linear resource adjustment cost functions, stationary parameters, and demand that is independent and identically distributed across periods, the problem collapses to a single-period problem in the same form. They classify this single-period problem, which is a two-stage stochastic programming model, as a *multidimensional newsvendor model.* We believe that our model is not a special case of their model; neither is our model more general. (After developing additional notation, we will contrast our model with theirs in detail in Section 2.3.2.) For an example problem, Harrison and Van Mieghem (1999) offer a parametric analysis involving identification of regions in the demand space throughout which there applies a constant optimal vector of shadow prices for capacity (degeneracy issues aside). They argue that the gradient of the expected operating profit function is equal to the sum of these optimal shadow prices weighted by the probabilities of the associated regions (assuming a probability distribution of demand having no point mass). To facilitate the solution of our model by a gradient-based approach, we will introduce an algorithm that in effect carries out this kind of decomposition of the demand space in our setting.

Van Mieghem and Rudi (2002) propose class of models called *newsvendor networks.* In their basic form, these models are single period capacity investment models formulated as two-stage stochastic programs featuring random demand for multiple products. Their resource processing model (i.e., the second stage) is more general than that of Harrison and Van Mieghem (1999), but more specific than that of Eberly and Van Mieghem (1997). Along with a capacity consumption matrix, the recourse linear program has a general input-output matrix translating input inventories (or stocks) into products. The authors show that this model captures several forms of production activities, including assembly, component commonality, input substitution or transshipment, resource flexibility, and simultaneous resource requirements. For an example problem, they identify regions in the demand space involved in a characterization of the gradient of the expected operating profit function as in Harrison and Van Mieghem (1999). Befitting the abstract nature of their model, they remark that particular instances may be substantially amenable to analytical solution, and they recommend optimization through simulation for other instances (i.e., estimating the gradient by solving the recourse linear program for a sample of demand vectors, as in Kim 2006). At least in terms of its mathematical structure, our model qualifies as a newsvendor network model. We show in Section 2.3.2 that our model (as we define it) can be transformed into the form of a newsvendor network; however, we argue that it is most natural to pose our recourse problem with two positive fractional covering constraints and some variables bounded above.

In the many published articles in this stream of capacity management research, one generally finds discussion of the structure of the decomposition of the demand space only for fairly specific cases. In particular, we are not aware of any substantial discussion of our type of recourse structure. This appears to be due largely to our perspective on demand for the service of transportation, highlighting two of its aggregate attributes that are *simultaneously* satisfied by a *single* activity. This perspective may be contrasted with a typical view of demand in terms of quantities of discrete products, in which a single activity produces a single product. (In our review of this stream of capacity investment research, the only model we have found for fleet composition is the model of Netessine et al. 2002, which also qualifies as a newsvendor network. Their model is in a car rental context, and it does not share our perspective on demand.) Furthermore, our desire to abstract away from a specific number of vehicle types with fixed characteristics invites a more generalized kind of analysis of a wide *class* of decompositions. By in effect generating this kind of decomposition, our algorithm reduces the temptation to re-solve the recourse linear program from scratch many times as in the simulation optimization method recommended for general problems by Van Mieghem and Rudi (2002). Since our demand space is only two-dimensional, we are able to recommend straightforward numerical integration instead of simulation for calculating the desired gradient; the appropriate optimal dual vector (and an optimal basis matrix inverse) can be pre-calculated for each region of integration before the gradient search algorithm begins.

Finally, given the newsvendor-type models under discussion, it is warranted to relate our model to other descendents of the classic newsvendor model, which serves as a foundation for models in capacity planning as well as inventory control. In particular, newsvendor network models such as ours may be distinguished from the "multi-product newsvendor" (or "multi-item newsvendor," or "newsstand") problems that appeared earlier in the literature. In these models, the focus is on dealing with complex ex-ante constraints (e.g. a budget constraint when selecting how many of each product to stock) while the ex-post constraints are trivial. For us, the main issue is a trivially constrained ex-ante capacity investment decision, which along with realized demand determines complex ex-post constraints on operations. Discussion of the basic newsvendor (or "newsboy") problem and its many variants can be found in Porteus (1990) and Khouja (1999).

2.2.3 Optimization

Our proposed model falls under the general category of two-stage stochastic linear programs with fixed and relatively complete recourse; a general reference on stochastic programming is Birge and Louveaux (1997). Our recourse program is a minimization problem with two positive fractional covering constraints, in which all coefficients are assumed positive, with some (but not all) variables bounded above.

Our main contribution in the context of optimization is our algorithm in Section 2.4.2 generating a "definitive" collection of bases of our recourse problem. We define a definitive collection of bases in the context of the analysis of our recourse linear program, which has varying parameters in the right-hand side. Such a collection $\mathcal{B}$ must satisfy three conditions: Optimality (each basis in $\mathcal{B}$ must be optimal for those parameter values where it is feasible), Covering (for each possible parameter value, there must be a feasible basis in $\mathcal{B}$), and Disjoint Interiors (for no pair of bases in $\mathcal{B}$ do the associated feasible regions have a common interior point). Having a definitive collection of bases of our model facilitates the rapid computation of the expected (first-stage or second-stage) cost and its gradient, so that one may then proceed with a standard gradient projection algorithm.

The usefulness of such collections as we generate, which goes beyond our present purpose of cost and gradient computation, is known to researchers in optimization. The issue of degeneracy complicates matters of conceptualization, but it is clear that the same general idea comes up in many contexts. In stochastic programming, collections such as ours are intimately related to the technique of "full decomposition" of the space of right-hand sides discussed in Birge and Louveaux (1997, pp. 170–171). "Bunching" techniques seek to employ in effect a partial decomposition (Birge and Louveaux 1997, pp. 169–174), and can be used to speed up the L-shaped method for solving stochastic linear programs with recourse (ibid., Chapter 5). These kinds of decompositions have been used to study the "distribution problem" of stochastic programming, in which one seeks to understand the distribution of the optimal objective value of a (linear) program having random parameters with known joint distribution (Foote 1980, Wets 1980). The same general kind of decomposition is sought in parametric linear programming (Jones 2005, from the perspective of optimal control). An early articulation of the theoretical underpinnings of such decompositions is the basis decomposition theorem of Walkup and Wets (1969).

As one would expect, algorithms producing these useful decompositions have been proposed. General algorithms are discussed in the studies of Foote (1980), Wets (1980), and Jones (2005). (Dual simplex pivoting in particular is a useful algorithmic element in situations where the right-hand side varies.) Studies aimed at producing such decompositions for particular classes of structured linear programs include Wallace (1986) and Filippi and Romanin-Jacur (2002). We are not aware of any general or specialized algorithms in prior literature guaranteeing polynomial upper bounds on time complexity or the number of bases to be generated for our recourse model. Thus, it seems appropriate to offer an algorithm tailored to our recourse structure and to identify its complexity. In Section 2.4.2, we offer arguments for the correctness and efficiency of our algorithm, which identifies $O(n^2)$ bases—out of exponentially many bases that can be defined—in $O(n^3)$ time, given n vehicle types.

We expect that our conception of a "definitive" collection of bases can be widened past the context of the particular model we study here. While the facts giving rise to our conception are known, it seems that our concept is something of an implicit idea in the literature, not having a concise term.

In closing, we note that our recourse linear program is structurally related to the problem of set covering in combinatorial optimization, and also to continuous models studied in the context of approximation algorithms (e.g. Fleischer 2004).

2.3 Our model

In Section 2.3.1, we define our model and articulate points to consider in making an assessment of our model's validity. In Section 2.3.2, we note that our model may be seen as a generalization of the classic newsvendor problem, and we show via a transformation that our model qualifies as a newsvendor network.

2.3.1 Definition

A type i of vehicle under consideration for ownership for the duration of the planning horizon is distinguished by four parameters: a fixed cost f_i, a variable cost v_i, a sites capacity s_i, and a cubic volume capacity c_i. The *fixed cost* is that cost incurred per working day due to including one vehicle of this type in the permanent fleet through the planning horizon, whether or not the vehicle is used. This quantity takes into account elements such as the vehicle's depreciation cost, taxes, fixed components of drivers' wages and insurance, the cost of parking, and the cost of maintenance performed at fixed time intervals, all amortized equally across the working days in the planning horizon. The *variable cost* is that additional cost per working day incurred due to fully utilizing one owned vehicle of this type, as against leaving it unused for the day. This includes, for example, the cost of fuel, variable components of drivers' wages, and the cost of maintenance performed at fixed mileage intervals as well as non-routine maintenance. The *cubic volume capacity* is the greatest volume of merchandise that a vehicle of this type can reliably deliver in one working day. The *sites capacity* is the greatest number of customer sites that a vehicle of this type can reliably complete in one working day. Note that full utilization of a vehicle may entail multiple loads in a given day. More broadly, the parameters v_i, s_i, c_i are not intrinsic to the vehicle; they depend on factors such as the spatial distribution of customers relative to the distribution center, customers' order sizes, and other factors.

A type i of vehicle available for hire on a daily (spot) basis during the planning horizon is distinguished by three parameters: a variable cost v_i, a cubic volume capacity c_i, and a sites capacity s_i. Under our assumptions, what makes a hired or spot vehicle different from an owned vehicle in our fleet is that a third party (most likely a common carrier) assumes responsibility for the fixed cost associated with long-term control of the vehicle. From the retail distributor's perspective, then, the cost incurred when relying on a specific type of hired capacity is essentially a variable cost that depends on the degree to which the capacity is utilized. The definitions of v_i, c_i, and s_i given above still apply. The method of measuring v_i in particular is likely to be different in the context of a spot vehicle, however, since the separate components of this cost are probably hidden.

Let $\mathcal{O}$ be the set of vehicle types under consideration for ownership, and let $\mathcal{H}$ be the set of vehicle types available for hire on a spot basis. We assume that there are n_1 vehicle types in $\mathcal{O}$ and n_2 vehicle types in $\mathcal{H}$, and we define $n = n_1 + n_2$. (To clear up a possible confusion: if e.g. a particular make and model of vehicle is available both for ownership and for hire, this should be counted as two distinct vehicle "types" in the sense relevant here—due to the different cost implications of the two means of access.) We allow a more compressed notation by defining $f \in \mathbb{R}_+^{n_1}$, $v \in \mathbb{R}_+^n$, $c \in \mathbb{R}_+^n$, and $s \in \mathbb{R}_+^n$ to be vectors consisting of the parameters introduced above.

Demand for transportation is characterized by two jointly distributed nonnegative random variables C and S, representing respectively the total cubic volume of merchandise to be delivered and the total number of distinct customer sites to be visited. This distribution may be interpreted as describing demand on a working day picked uniformly at random from the planning horizon. (We elaborate on this interpretation below.) We define the two-dimensional demand $D = (C, S)$.

We define the fleet composition vector $K \in \mathbb{R}_+^{n_1}$ to indicate the number of vehicles of each type we decide to include in our fleet; selecting K is the first-stage (or ex-ante) decision to be made in our model. Once the day's demand is known, we have the recourse (or second-stage, or ex-post) decision concerning how best to utilize owned and hired capacity to fulfill demand. We define the utilization vector $x \in \mathbb{R}^n_+$ indicating the number of vehicles of each type to utilize. Given a fleet composition K and realized demand $D = (C, S)$, we define the (optimal) second-stage cost:

$$\begin{aligned} z(K,D) := \quad \min \quad & v'x \\ \text{s.t.} \quad & c'x \geq C \\ & s'x \geq S \\ & x_i \leq K_i \quad i \in \mathcal{O} \\ & x_i \geq 0 \quad\;\; i \in \mathcal{O} \cup \mathcal{H} \end{aligned} \tag{1}$$

Observe that we have defined a continuous optimization problem. The constraint $x_i \leq K$ for $i \in \mathcal{O}$ reflects simply that we cannot utilize more of these vehicles than we own. Hired vehicle types, by contrast, are taken to have unlimited availability. The *expected* second-stage cost is

$$Z(K) := \mathbb{E}[z(K,D)]$$

where the expectation is taken with respect to the bivariate probability distribution. Finally, our formulation of the first-stage decision is simply

$$\begin{aligned} \min \quad & f'K + Z(K) \\ \text{s.t.} \quad & K \geq 0 \end{aligned}$$

which we also treat as a continuous optimization problem. If demand is known but exhibits variation (said variation being captured by the bivariate probability distribution), this formulation is equivalent to a deterministic minimization of costs over the entire planning horizon; if demand is uncertain, this formulation minimizes *expected* costs over the planning horizon.

We elaborate on the interpretation of the model when demand is uncertain. Let t index the working days in the planning horizon, and suppose the number of working days in the horizon is T. Let the random vector D be distributed as the demand on a working day picked uniformly at random from the planning horizon. In this context, we abbreviate by $P(t)$ the probability that day t is picked, which is by definition T^{-1}. Let the random vector D_t be distributed as the demand on day t. Observe that the conditional distribution of D, given that day t was selected, is equal to the distribution of D_t. Now the expected total cost incurred over the planning horizon, as a function of the capacity vector K, is equal to:

$$\textstyle\sum_{t=1}^{T}(f'K + \mathbb{E}[z(K, D_t)])$$

by linearity of expectation. It is equivalent to minimize the function obtained by multiplying the above by the positive constant T^{-1}. The function so obtained is equal to:

$$f'K + \textstyle\sum_{t=1}^{T} \mathbb{E}[z(K, D) \,|\, t]P(t)$$

which reduces to our objective function $f'K + \mathbb{E}[z(K, D)]$ by appeal to the law of total expectation.

Throughout the rest of the chapter, we assume for the sake of simplicity that each of the parameters f_i, v_i, c_i, and s_i is positive, and that n_1 and n_2 are positive. We also assume that the random variable $D = (C, S)$ has a (known) joint probability density function with finite second moments.

We now articulate various propositions underlying our model that must be considered in making a full assessment of our model's validity. A full defense of our model, considering further details of the application context, possible qualifications of certain propositions, mathematical properties of our model, and ways in which our model may be extended, is beyond the scope of this chapter. We do, however, offer a few comments on the propositions.

1. *Our continuous, linear model of fixed costs is valid.*

2. *Our continuous, linear model of resource processing costs, given realized demand, is valid.*

   Since in theory an optimal solution to our recourse problem might have only one demand constraint binding, we highlight the subtle proposition that *we incur utilization cost for (optimal) "surplus production" as though demand matched the total capacities of the vehicles we utilized.*

   We comment on the first two points. Regarding the general issue of the appropriateness of a continuous and linear modeling approach, we appeal to the large customer base of our prototypical beverage distributor, the strategic nature of the decision at hand, and the broad nature of the firm's business concerns (including inventories and marketing, for example, in addition to the transportation of merchandise to customers that is our focus in this study). In regard to our model more specifically, including the representation of demand by the total cubic volume and number of sites, we appeal to the intuitive considerations given in the introduction to this chapter, along with the distribution context described. Further research will be required to fully validate our proposed model.

3. *The fleet composition vector directly translates to the set of vehicles actually available for use day-to-day, independent of utilization decisions during the planning horizon.*

4. *Per-unit utilization costs and capacity parameters of owned vehicle types are constant throughout the planning horizon, and are independent of our fleet composition and day-to-day utilization decisions.*

5. *Hired vehicle types may be utilized in (arbitrarily) great quantities; their per-unit utilization costs and capacity parameters are constant throughout the planning horizon, and are independent of our fleet composition and day-to-day utilization decisions.*

6. *Deliveries are not delayed.* To put it another way, *there are no backlogs.*

7. *Demand is independent of the fleet composition and day-to-day utilization decisions.*

8. *Variable costs incurred at opposite ends of the planning horizon need not be discounted differently.*

9. *Our parameters, and the distribution of demand, can be measured with a level of accuracy appropriate to the decision at hand, considering the effort associated with this measurement.*

   In some situations (such as when the fleet is to be leased rather than owned), contractual agreements might greatly simplify the process of parameter estimation.

   For a published account of efforts at estimating fixed and variable cost parameters for three vehicle fleets, see Eilon et al. (1971b), which informed the exposition of our cost parameters above. In our view, further studies of this nature would be valuable additions to the literature.

2.3.2 Relation to other newsvendor-type models

Having dealt with the relationship of our model to other newsvendor-type models to some extent in our literature review (specifically Section 2.2.2), and having introduced the notation of our model above, we are now in a position to compare and contrast the models under discussion with greater precision. For simplicity, we maintain our assumption that each of our parameters f_i, v_i, c_i, and s_i is positive, and that n_1 and n_2 are positive.

First, we illustrate a reduction of our model to the traditional newsvendor problem for a special case. Suppose that in our model there is a single vehicle type available for ownership, $1 \in \mathcal{O}$, and a single type that will be available for hire on a spot basis, $2 \in \mathcal{H}$. (Here, $n_1 = n_2 = 1$.) Furthermore, suppose that $c_1 = c_2$ and $s_1 = s_2$. In this case, we can meaningfully view any given demand $D = (C, S)$ in terms of the number of vehicles required, regardless of type. Let us define this transformed demand as $D^* = \max\{C/c_1, S/s_1\} = \max\{C/c_2, S/s_2\}$. The cost incurred by purchasing K vehicles, after demand is realized, may now be expressed as

$$(f_1 + v_1)K + (-v_1)(K - D^*)^+ + v_2(D^* - K)^+$$

We recognize this as the form of a traditional newsvendor cost model with an ordering cost of $(f_1 + v_1)$, a salvage value of v_1, and a lost-sales penalty of v_2. We assume that $f_1 + v_1 < v_2$, so there is some incentive to own vehicles. Accordingly, assuming a differentiable probability density function, the optimal K is the critical fractile K^* such that $P(D^* \leq K^*) = (v_2 - f_1 - v_1)/(v_2 - v_1)$. A similar reduction is possible if we have the more general condition $c_1/s_1 = c_2/s_2$.

In the remainder of this section, we examine the relationship between our model and the multidimensional newsvendor model defined by Harrison and Van Mieghem (1999) as well as the newsvendor network model defined by Van Mieghem and Rudi (2002). We argue that our model qualifies as a newsvendor network, but that it is most natural to use the formulation we introduced in Section 2.3.1. We also argue that our model is not a special case of the model defined by Harrison and Van Mieghem (1999), though their model is also a newsvendor network. We adjust the notation of Van Mieghem and Rudi (2002) to avoid excessive conflict with our notation.

In the multidimensional newsvendor model of Harrison and Van Mieghem (1999), the second-stage cost is $\max\{p'x : Ax \le K, x \le D, x \ge 0\}$, where K is the vector of chosen capacity levels and D is the vector of realized demand. In the more general newsvendor network model of Van Mieghem and Rudi (2002), the form of the recourse problem is:

$$
\begin{aligned}
\max \quad & (r - c_A)'x - c_P'(D - R_D x) - c_H'(\bar{S} - R_S x) \\
\text{s.t.} \quad & R_S x \le \bar{S} \\
& R_D x \le D \\
& Ax \le K \\
& x \ge 0
\end{aligned}
$$

where R_S, R_D, and A are nonnegative matrices. (Though Harrison and Van Mieghem do not explicitly assume that A is nonnegative, it is reasonable to suppose that they did not intend for their model to be interpreted otherwise.) Here, $\bar{S}$ is a vector of "input stocks" chosen along with K in the first stage; each decision variable in the vector x corresponds to an "activity" that depletes input stocks and uses resources to produce demanded "outputs." The first-stage objective function to be maximized is the expected value of the optimal profit defined above, less the investment cost $c_S'\bar{S} + c_K'K$. In terms of mathematical structure, the presence of "resources" adds nothing; observe that any constraint from the set $Ax \le K$ could be subsumed into the set $R_S \le \bar{S}$ with a corresponding "holding cost" (c_H) component equal to 0, moving the associated investment cost from c_K to c_S. (Van Mieghem and Rudi 2002 go on to extend their model, creating a "dynamic newsvendor network" in which the distinction between input stocks and resources is much more substantial.) We submit that the main question becomes whether our recourse problem can be phrased in the form:

$$\begin{aligned} \max \quad & (r - c_A)'x - c_P'(D - R_D x) - c_H'(K - Ax) \\ \text{s.t.} \quad & R_D x \leq D \\ & Ax \leq K \\ & x \geq 0 \end{aligned}$$

where R_D and A are nonnegative matrices. If we take the natural step of interpreting D above as our two-dimensional demand $D = (C, S)$, it is likewise natural to introduce an "activity" i for each of our vehicle types with corresponding decision variable x_i and corresponding column (c_i, s_i) in R_D. (Note that this structure, reflecting a vehicle's ability to produce two "outputs" simultaneously, is not captured by the recourse problem of Harrison and Van Mieghem 1999.) Since we require that all demand be satisfied, but equality of utilized vehicle capacity and demand might not be optimal (or even feasible), we introduce for each vehicle type i additional "activities" with associated decision variables x_i^c and x_i^s; the corresponding columns in R_D are respectively $(c_i, 0)$ and $(0, s_i)$. For each owned vehicle type i, the corresponding row of A has 1 in the columns corresponding to x_i, x_i^c, and x_i^s, and 0 elsewhere; K is again interpreted as the fleet composition vector. We set $c_H = 0$ and $r = 0$; we define the components of c_A to be v_i for those entries corresponding to x_i, x_i^c, and x_i^s; and we define the entries of c_P to be strictly greater than the least marginal cost of fulfilling the corresponding dimension of demand by a spot vehicle type. Because the "penalty cost" is so severe, we are assured that both demand constraints will be binding at optimality—a spot vehicle type can always be used to pick up the slack while reducing the penalty. It can now be argued that any optimal solution to the newsvendor network recourse problem we have formulated can be translated into an optimal solution of the recourse problem in our model with the same cost (though the objective values have opposite signs). (To translate a solution from the newsvendor network, we set the utilization of vehicle type i in our model equal to $x_i + x_i^c + x_i^s$

from the newsvendor network solution.) Of course, the vector f representing our linear capacity investment cost can be translated directly to c_S or c_K. We conclude that our model can be put into the form of a newsvendor network, though doing so seems to require a rather contrived mechanism.

2.4 An algorithm facilitating solution of our model

In Section 2.4.1, we first reiterate our technical assumptions and note that our deterministic equivalent program is thereby assured of being a convex program, with the gradient defined everywhere in the interior of the feasible set. We then define bases of our recourse program (as well as "optimal" bases) by reference to an alternative formulation of the recourse program with two surplus variables. We define the feasible regions of our bases in the space of possible demands, and we introduce notation for the optimal dual vectors associated with optimal bases. We define our notion of a "definitive" collection of bases of our recourse program and elucidate the role of such a collection in a standard gradient search algorithm. In Section 2.4.2, we present our algorithm generating a definitive set of bases of our recourse program, and we argue for its correctness and polynomial complexity.

2.4.1 Preliminaries

We first reiterate our technical assumptions, which we have imposed for simplicity.

Assumption 1. Each of the parameters f_i, v_i, c_i, and s_i is positive, and n_1 and n_2 are positive.

Assumption 2. The nonnegative random variable $D = (C, S)$ has a known joint probability density function with finite second moments.

By our first assumption, our stochastic programming model has relatively complete recourse. Our optimization model, $\min f'K + Z(K)$ subject to $K \geq 0$, may therefore be viewed as a deterministic equivalent program, pushing the stochasticity within $Z(K)$ into the background of our attention; we do not need to add constraints on K to ensure that the second stage has an optimal solution for any possible demand vector. By our two assumptions together, we may invoke a result stated in Birge and Louveaux (1997, Theorem 6, pp. 90–91) to conclude that our deterministic equivalent program is a convex program, with the gradient of $f'K + Z(K)$ defined for all K in the interior of the feasible set.

The major work to be done involves construction of a special set of bases of our recourse problem, reformulated with two surplus variables x_c and x_s to produce equality constraints (rather than covering constraints). We identify a basis by partitioning the decision variables (formally, their indices) into three sets: (1) a set B comprised of two basic indices, (2) a set L of indices of variables held at their lower bounds (i.e., zero), and (3) a set U of indices of variables held at their upper bounds, chosen from $\mathcal{O}$. We describe a basis as optimal if (1) its associated solution is feasible (i.e., basic variables x_{i_1} and x_{i_2} are nonnegative and, where applicable, are bounded above by K_{i_1} and K_{i_2} respectively), (2) the reduced cost $\bar{v}_i$ is nonnegative for each $i \in L$, and (3) the reduced cost $\bar{v}_i$ is nonpositive for each $i \in U$. (We define $\bar{v}_i = v_i - v'_B B^{-1} A_i$, where B^{-1} is interpreted as the inverse of the square matrix composed of the two basic columns from the system of equalities in the formulation above, and A_i is the selected nonbasic column of that system. For the theory behind our definition of bases using variables with upper and lower bounds, see e.g. Bertsimas and Tsitsiklis 1997, especially the exercise on p. 135.) For each basis b we define the associated feasible regions in the space of possible demands:

$$\Omega_b = \{(C, S, K) \in \mathbb{R}_+^{2+n_1} : b \text{ is feasible for } (C, S, K)\}$$

$$\Omega_b(K) = \{(C, S) \in \mathbb{R}_+^2 : (C, S, K) \in \Omega_b\}$$

A basis b that is optimal for some point in Ω_b is optimal for every point in this region. The dual of our recourse problem (as originally formulated) may be expressed as follows:

$$\begin{array}{rlll} \max & C\mu_1 + S\mu_2 + K'\lambda & & \\ \text{s.t.} & c_i\mu_1 + s_i\mu_2 + \lambda_i \le v_i & i \in \mathcal{O} & \\ & c_i\mu_1 + s_i\mu_2 \le v_i & i \in \mathcal{H} & \\ & \mu_j \ge 0 & j \in \{1,2\} & \\ & \lambda_i \le 0 & i \in \mathcal{O} & \end{array} \tag{2}$$

If b is optimal for Ω_b, then a uniquely optimal dual vector λ_b obtains throughout the interior of Ω_b. Whenever $K_i = 0$, the ith component of λ_b may be interpreted as the *right* partial derivative of the optimal cost with respect to K_i at demand points in the interior of $\Omega_b(K)$.

We desire a set $\mathcal{B}$ of bases such that three key conditions are satisfied:

1. (Optimality) Each basis $b \in \mathcal{B}$ is optimal for Ω_b.

2. (Covering) $\bigcup_{b\in\mathcal{B}} \Omega_b(K) = \mathbb{R}^2_+$ for each $K \ge 0$.

3. (Disjoint Interiors) For each $K \ge 0$ and basis pair $b_1, b_2 \in \mathcal{B}$, the interiors of $\Omega_{b_1}(K)$ and $\Omega_{b_2}(K)$ are disjoint.

We call a set satisfying these conditions a *definitive* collection of bases. The value, for our purposes, of a definitive set is due in large part to the following result. The result follows from Proposition 1 of Van Mieghem and Rudi (2002), but we provide our own proof.

Proposition. Let $\mathcal{B}$ be a definitive set of bases, suppose Assumptions 1 and 2 hold. Then for each $K \ge 0$, we have:

$$\nabla Z(K) = \sum_{b\in\mathcal{B}} \lambda_b P(\Omega_b(K))$$

Whenever $K_i = 0$, the ith component of this sum should be interpreted as a right partial derivative.

Proof. Let $K \geq 0$ be given. Consider a point $D = (C, S)$ that lies in the interior of the region $\Omega_b(K)$ for some $b \in \mathcal{B}$. Supposing $K_i > 0$, let $\{\delta_m : m \geq 1\}$ be a sequence of real numbers converging to zero; if $K_i = 0$, let the sequence be positive as well. Since b is an optimal basis, the sequence of elements $\delta_m^{-1}[z(K + \delta_m e_i, D) - z(K, D)]$ will eventually become constant and equal to its limit $\lambda_{b,i}$. (Here $e_i \in \mathbb{R}^{n_1}$ is the unit vector with ith component equal to one, and $\lambda_{b,i}$ is the ith component of λ_b.) Since $z(\cdot, \cdot)$ is Lipschitz continuous (in fact, piecewise linear), the family of functions $g_D(\delta) = \delta^{-1}[z(K + \delta e_i, D) - z(K, D)]$ is uniformly bounded. Because D has a probability density function, the set of boundary points of the regions $\Omega_b(K)$ for all $b \in \mathcal{B}$ must have probability mass zero. Applying the bounded convergence theorem, $\frac{\partial}{\partial K_i}\mathbb{E}[z(K, D)] = \mathbb{E}[\frac{\partial}{\partial K_i} z(K, D)]$, where we understand $\frac{\partial}{\partial K_i}$ as the right partial derivative if $K_i = 0$. Since $\mathcal{B}$ is definitive, we also have $\mathbb{E}[\frac{\partial}{\partial K_i} z(K, D)] = \sum_{b \in \mathcal{B}} \lambda_{b,i} P(\Omega_b(K))$. $\square$

Having an algorithm for generating a definitive set of bases of our recourse problem, we may pursue a standard gradient projection method for optimization. (Provision for projection of the improving direction—and limitation of the maximum step size—are required merely due to our nonnegativity constraints on K.) Computation of the objective function and its gradient in this context may be facilitated as follows. Before initiating the search algorithm, generate a definitive collection $\mathcal{B}$ using our algorithm. For each basis $b \in \mathcal{B}$, store the corresponding basis matrix inverse B^{-1} and the vector λ_b along with the defining sets B, L, and U. (The elements of λ_b are defined for reference within the presentation of our algorithm below.) Now suppose in the course of our gradient search algorithm we are given $K \geq 0$ and we require the gradient of the objective function. For each basis $b \in \mathcal{B}$, we obtain $P(\Omega_b(K))$ through numerical integration of the density function over the appropriate region; the desired gradient is then $f + \sum_{b\in\mathcal{B}} \lambda_b P(\Omega_b(K))$, where the ith component is actually the right partial derivative if $K_i = 0$. Finally, suppose we are given $K \geq 0$ and we require the value of the objective function for our search algorithm. For this end, we perform numerical integration over the demand space to calculate the expectation of $f'K + z(K, D)$ with respect to the density function. In doing so, we evaluate the second-stage cost $z(K, D)$ at a given point $D = (C, S)$ using the sets B, L, and U and the stored matrix B^{-1} corresponding to a basis $b \in \mathcal{B}$ for which $D \in \Omega_b(K)$. We suggest performing this integration over the regions $\Omega_b(K)$ in sequence, so that we always know immediately which basis parameters apply for a given demand point falling in the current region. For the purpose of numerical integration, we can reduce the (theoretically unbounded) demand space to a large bounded region, such that nearly all of the probability mass is captured. On determining the best way to perform the numerical integrations we suggest, one may find some discussions and further references in the works by Foote (1980), Wets (1980), and Birge and Louveaux (1997, especially pp. 286–288). Due to the strategic nature of the decision of interest and the low dimension of the parameter space, we submit that sophisticated numerical integration methods may not be crucial.

2.4.2 Efficient generation of "definitive" collections for our model

Below, we offer our algorithm for generating a definitive set $\mathcal{B}$. The algorithm takes the liberty of discarding from the formulation certain vehicle types that will not (or need not) be present in an optimal fleet composition. The algorithm is based on a representation of the variable-cost efficiency of a vehicle type by the vector $(c_i/v_i, s_i/v_i)$, which we will call ϕ_i. By Assumption 1, each of these vectors is well-defined and has both components strictly positive.

Algorithm.

1. (Eliminate clearly uneconomical or unnecessary vehicle types.) For each vehicle

type i_1, determine whether there is a second vehicle type i_2 such that $\phi_{i_1} \leq \phi_{i_2}$, $f_{i_1}/c_{i_1} \geq f_{i_2}/c_{i_2}$, and $f_{i_1}/s_{i_1} \geq f_{i_2}/s_{i_2}$. If so, eliminate type i_1. In this context, spot vehicle types are understood to have $f_i = 0$; ties may be broken arbitrarily. (To fulfill a given quantity of demand with i_2 instead of i_1, the variable cost and the prerequisite fixed cost will both be lower. Once this step is complete, we are assured that no two efficiency vectors ϕ_{i_1} and ϕ_{i_2} are identical. We assume that $n_1 > 0$ even after all possible eliminations have been performed; otherwise, the basic tradeoff giving rise to our model does not apply.)

2. (Generate bases allowing for a surplus of sites capacity.) Order the vehicle types so that $c_{i_k}/v_{i_k} \geq c_{i_{k+1}}/v_{i_{k+1}}$ for all $k \in \{1, \ldots, n-1\}$, where $s_{i_k}/v_{i_k} > s_{i_{k+1}}/v_{i_{k+1}}$ if $c_{i_k}/v_{i_k} = c_{i_{k+1}}/v_{i_{k+1}}$. The types are now in decreasing order of variable-cost efficiency at handling cubic volume, with ties broken in favor of the type that is more efficient at handling customer sites. For each ℓ from 1 to $\min\{k : i_k \in \mathcal{H}\}$, generate the basis defined by $B = \{i_\ell, s\}$, $L = \{i_k : k > \ell\} \cup \{c\}$, and $U = \{i_k : k < \ell\}$.

   The optimal dual vector we associate with any given basis generated here has $\mu_1 = v_{i_\ell}/c_{i_\ell}$, $\mu_2 = 0$, $\lambda_i = 0$ for all $i \in \mathcal{O} \setminus U$, and $\lambda_i = v_i - c_i\mu_1$ for all $i \in U$.

3. (Generate bases allowing for a surplus of cubic volume capacity—this step is symmetrical with the preceding step.) Order the vehicle types so that $s_{i_k}/v_{i_k} \geq s_{i_{k+1}}/v_{i_{k+1}}$ for all $k \in \{1, \ldots, n-1\}$, where $c_{i_k}/v_{i_k} > c_{i_{k+1}}/v_{i_{k+1}}$ if $s_{i_k}/v_{i_k} = s_{i_{k+1}}/v_{i_{k+1}}$. The types are now in decreasing order of variable-cost efficiency at handling customer sites, with ties broken in favor of the type that is more efficient at handling cubic volume. For each ℓ from 1 to $\min\{k : i_k \in \mathcal{H}\}$, generate the basis defined by $B = \{c, i_\ell\}$, $L = \{i_k : k > \ell\} \cup \{s\}$, and $U = \{i_k : k < \ell\}$.

   The optimal dual vector we associate with any given basis generated here has

$\mu_1 = 0$, $\mu_2 = v_{i_\ell}/s_{i_\ell}$, $\lambda_i = 0$ for all $i \in \mathcal{O} \setminus U$, and $\lambda_i = v_i - s_i\mu_2$ for all $i \in U$.

4. (Generate bases featuring two basic vehicle types.) Perform the following for each pair of vehicle types $\{i_1, i_2\}$. If $\phi_{i_1} \leq \phi_{i_2}$ or $\phi_{i_1} \geq \phi_{i_2}$, continue with the next pair to be considered. Otherwise, calculate the values μ_1 and μ_2 such that the line $\mu_1 y_1 + \mu_2 y_2 = 1$ in $\mathbb{R}^2_+$ intersects both ϕ_{i_1} and ϕ_{i_2}. Identify the set V of vectors ϕ_i with $i \notin \{i_1, i_2\}$ such that $\mu_1(c_i/v_i) + \mu_2(s_i/v_i) > 1$ or such that ϕ_i is a convex combination of ϕ_{i_1} and ϕ_{i_2}. If $V \cap \mathcal{H}$ is nonempty, continue with the next pair to be considered. Otherwise, generate the basis with $B = \{i_1, i_2\}$, $U = V$, and all remaining indices (including c and s) assigned to L.

   The optimal dual vector we associate with any given basis generated here has $\lambda_i = 0$ for all $i \in \mathcal{O} \setminus U$ and $\lambda_i = v_i - c_i\mu_1 - s_i\mu_2$ for all $i \in U$.

Our main technical result is the following:

Theorem. Suppose Assumption 1 holds. Our algorithm generates a definitive set of no more than $\binom{n}{2} + 2n_1 + 2$ bases in $O(n^3)$ time.

Note that we are in effect discarding an exponential number of bases, due to the structure of our recourse linear program. This theorem is proved by means of four lemmas. The first lemma establishes the algorithm's computational complexity. The other three lemmas establish the three key conditions defining a definitive collection of bases of our model (given in Section 2.4.1 above).

Lemma 1. Suppose Assumption 1 holds. Our algorithm generates no more than $\binom{n}{2} + 2n_1 + 2$ bases in $O(n^3)$ time.

Proof. We may quickly verify that Steps 1–3 of our algorithm require no more than $O(n^2)$ time. A straightforward implementation of our algorithm requires $O(n^3)$ time due to Step 4, which may perform quadratically many linear-time operations. (For pairs of vectors ϕ_i, we may determine for each remaining vector which simply defined region it belongs to.) No bases are generated in Step 1. In Step 2, the greatest number of bases that may be generated is $n_1 + 1$. For this bound to be reached, all owned vehicle types must have greater variable-cost efficiency at handling cubic volume than the best hired type by this measure. In Step 3, the focus is instead on customer sites, and similarly the greatest number of bases that may be generated is $n_1 + 1$. In Step 4, we generate at most one basis for each pair of vehicle types, so an upper bound on the number of bases produced here is $\binom{n}{2}$. $\square$

Lemma 2. Suppose Assumption 1 holds. The bases generated by our algorithm satisfy the Optimality Condition.

Proof. Using our assumption that each c_i and s_i is positive, we easily see that $\Omega_b(1)$ has nonempty interior. Selecting (C, S) from this interior, the corresponding basic solution is primal nondegenerate, and it will suffice to show that this solution is also optimal. If b was generated in Step 2, then $B = \{i_\ell, s\}$ for some vehicle type i_ℓ, and we construct the dual solution with $\mu_1 = v_{i_\ell}/c_{i_\ell}$, $\mu_2 = 0$, $\lambda_i = 0$ for all $i \in \mathcal{O} \setminus U$, and $\lambda_i = v_i - c_i\mu_1$ for all $i \in U$. We may easily verify that this dual solution is feasible and satisfies complementary slackness conditions with the primal solution corresponding to basis b, establishing optimality. If instead b was generated in Step 3, the situation is symmetrical with the previous case; here, $B = \{c, i_\ell\}$ for some vehicle type i_ℓ, and our dual solution has $\mu_1 = 0$, $\mu_2 = v_{i_\ell}/s_{i_\ell}$, $\lambda_i = 0$ for all $i \in \mathcal{O} \setminus U$, and $\lambda_i = v_i - s_i\mu_2$ for all $i \in U$. Finally, if b was generated in Step 4, then $B = \{i_1, i_2\}$ for some vehicle types i_1 and i_2. We let μ_1 and μ_2 be those (unique) values calculated by our algorithm for this pair of vehicle types; we set $\lambda_i = 0$ for all $i \in \mathcal{O} \setminus U$ and $\lambda_i = v_i - c_i\mu_1 - s_i\mu_2$ for all $i \in U$. Without much trouble, we may confirm that this dual solution is feasible and the complementary slackness conditions hold. $\square$

Lemma 3. Suppose Assumption 1 holds. The bases generated by our algorithm satisfy the Covering Condition.

Proof. By positivity of n_1, n_2, and each c_i and s_i, the second-stage program is feasible and admits an optimal basic solution for each $(C, S, K) \in \mathbb{R}_+^{2+n_1}$. By the theory of linear programming, a basic solution to our recourse problem (2) has at most two variables that are not at their bounds. By our assumption that v is positive, an optimal basic solution will never have x_c and x_s away from their bounds. If (C, S, K) is such that an optimal basic solution exists with x_s and x_{i_ℓ} away from their bounds for some vehicle type i_ℓ, then complementary slackness dictates that $\mu_1 = v_{i_\ell}/c_{i_\ell}$, that vehicle types i with $v_i/c_i < \mu_1$ must be owned types with $x_i = K_i$, and that vehicle types i with $v_i/c_i > \mu_1$ must have $x_i = 0$. Here when we have multiple vehicle types with $v_i/c_i = \mu_1$, optimality is maintained by utilizing them in decreasing order of variable-cost efficiency with respect to sites with the last type used designated basic; we also maintain optimality by ensuring for each type i with $K_i = 0$ that $i \in U$ if $v_i/c_i < \mu_1$ and $i \in L$ if $v_i/c_i > \mu_1$, ending up with a basis we generated in Step 2. We may argue similarly that some basis generated in Step 3 is optimal when (C, S, K) is such that an optimal basic solution exists with x_c and x_{i_ℓ} away from their bounds for some vehicle type i_ℓ. Finally, if (C, S, K) is such that an optimal basic solution exists with x_{i_1} and x_{i_2} away from their bounds for the pair of vehicle types i_1 and i_2, then complementary slackness dictates that μ_1 and μ_2 are as calculated in Step 4 for this basic pair, that vehicle types i with ϕ_i strictly above this line must be owned types with $x_i = K_i$, and that vehicle types i with ϕ_i strictly below this line must have $x_i = 0$. Here when we have more than two vehicle types with ϕ_i on this line, then we may arrive at a basic pair and corresponding optimal utilizations consistent with Step 4 by systematically decreasing the utilization of the "outermost" types on this line while increasing utilization of "inner" types on the line; we also maintain optimality by ensuring for each type i with $K_i = 0$ that $i \in U$ if ϕ_i lies above the line and $i \in L$ if ϕ_i lies below the line, ending up with a basis we generated in Step 4. For each $K \geq 0$, the foregoing cases establish for all but a set of Lebesgue measure zero the covering of R_+^2 by $\bigcup_{b\in\mathcal{B}} \Omega_b(K)$. Since this is a finite covering by closed sets, it follows that this union in fact covers all of the demand space $\mathbb{R}_+^2$. $\square$

Lemma 4. Suppose Assumption 1 holds. The bases generated by our algorithm satisfy the Disjoint Interiors Condition.

Proof. Let $K \geq 0$ be given, and let (C, S) be an interior point of some region $\Omega_b(K)$ for a basis b our algorithm has generated. If b was generated in Step 2, with $B = \{i_\ell, s\}$ for some vehicle type i_ℓ, by complementary slackness we must have $\mu_1 = v_{i_\ell}/c_{i_\ell} > 0$ and $\mu_2 = 0$ at dual optimality. If (C, S) is to be an interior point of $\Omega_{b^*}(K)$ for some other basis b^* generated by our algorithm, these unique dual values imply that b^* must have $B^* = \{i_{\ell^*}, s\}$ with $v_{i_{\ell^*}}/c_{i_{\ell^*}} = v_{i_\ell}/c_{i_\ell}$. Without loss of generality, we may assume that type i_ℓ comes before type i_{ℓ^*} in the ordering of Step 2. But now any point in the interior of $\Omega_b(K)$ involves strictly less than $\sum_{i\in U} c_i K_i + c_{i_\ell} K_{i_\ell}$ units of cubic volume, while any point in the interior of $\Omega_{b^*}(K)$ involves strictly more than this quantity of cubic volume since $(U \cup \{i_\ell\}) \subseteq U^*$ by construction. We may argue similarly for the case where b was generated in Step 3, where $B = \{c, i_\ell\}$ for some vehicle type i_ℓ and the uniquely optimal dual values are $\mu_1 = 0$ and $\mu_2 = v_{i_\ell}/s_{i_\ell} > 0$. To conclude, we consider the case where b was generated in Step 4, with $B = \{i_1, i_2\}$ for some vehicle types i_1 and i_2. Here, complementary slackness implies that the uniquely optimal dual values of μ_1 and μ_2 are as calculated for this basic pair in Step 4. If another basis b^* generated in Step 4 has $B = \{i_3, i_4\}$ and the interior points of $\Omega_{b^*}(K)$ share the optimal values of μ_1 and μ_2, then we must have ϕ_{i_1}, ϕ_{i_2}, ϕ_{i_3}, and ϕ_{i_4} colinear. To support the desired conclusion that (C, S) cannot be an interior point of $\Omega_{b^*}(K)$, we describe the geometric structure of the regions corresponding to all bases generated in Step 4 using these values of μ_1 and μ_2. Given μ_1 and μ_2, we may without loss of generality assume that the set of vectors falling on the critical line of Step 4 is $\{\phi_1, \phi_2, \ldots, \phi_m\}$ with these vectors appearing on the line in the given order from left to right. For more than one basis to be generated, we must have $m > 2$.

The "canonical" case involves $\{1, \ldots, m\} \subseteq \mathcal{O}$ with $K_i > 0$ for $i \in \{1, \ldots, m\}$, in which $\binom{m}{2}$ bases are generated; the corresponding regions are parallelograms arranged adjacent to each other in a kind of modified binomial tree structure. Here, the "top level" region corresponding to $B = \{1, m\}$ shares sides with the two regions in the next lower level corresponding to $B = \{1, m-1\}$ and $B = \{2, m\}$. These two regions share sides with the three regions in the next lower level (if one exists), and so on for $(m-1)$ levels total. At the lowest level, pairs of regions with $B = \{i, i+1\}$ and $B = \{i+1, i+2\}$ for $i \in \{1, \ldots, m-2\}$ share a side; these adjacency relationships are in effect added on to a binomial tree structure. Non-canonical cases may be seen as limiting cases, in which opposite sides of certain parallelograms take zero length (when some $i \in \{1, \ldots, m\}$ have $K_i = 0$) or infinite length (when some $i \in \{1, \ldots, m\}$ are in $\mathcal{H}$). $\square$

2.5 Conclusion

In this chapter, we have introduced a relatively simple and fast solution approach for the fleet composition problem faced by a retail distribution firm, focusing on the context of a major beverage distributor. In support of this overarching goal, we introduced a fleet composition model with a novel combination of characteristics, and we performed technical analysis facilitating solution of the model.

We close by indicating a few logical avenues for further research. Our model's validity could be examined more deeply in a study providing details of its practical implementation in a real-life trial. Such an investigation should clarify the degree to which certain extensions of our approach are necessary and feasible. It should also yield insights regarding the best methods of computing solutions to our model. Another possible investigation would further illuminate the connection of the idea of a "definitive" collection of bases to the wider contexts of stochastic and parametric programming.

2.6 References

Alsbury, P. 1972. The vehicle fleet mix. *International Journal of Physical Distribution* **2**(3) 123–125.

Angelus, A., E. L. Porteus. 2002. Simultaneous capacity and production management of short-life-cycle, produce-to-stock goods under stochastic demand. *Management Science* **48**(3) 399–413.

Ball, M. O., B. L. Golden, A. A. Assad, L. D. Bodin. 1983. Planning for truck fleet size in the presence of a common-carrier option. *Decision Sciences* **14**(1) 103–120.

Barnhart, C., R. R. Schneur. 1996. Air network design for express shipment service. *Operations Research* **44**(6) 852–863.

Bertsimas, D., J. N. Tsitsiklis. 1997. *Introduction to Linear Optimization.* Athena Scientific, Belmont, MA.

Birge, J. R., F. Louveaux. 1997. *Introduction to Stochastic Programming.* Springer, New York.

Ceder, A. 2005. Estimation of fleet size for variable bus schedules. *Transportation Research Record* **1903** 3–10.

Clark, P. 2007. *Buying the Big Jets: Fleet Planning for Airlines,* 2nd ed. Ashgate, Aldershot, UK.

Couillard, J., A. Martel. 1990. Vehicle fleet planning in the road transportation industry. *IEEE Transactions on Engineering Management* **37**(1) 31–36.

Crary, M., L. K. Nozick, L. R. Whitaker. 2002. Sizing the U.S. destroyer fleet. *European Journal of Operational Research* **136** 680–695.

Dantzig, G. B., D. R. Fulkerson. 1954. Minimizing the number of tankers to meet a fixed schedule. *Naval Research Logistics Quarterly* **1** 217–222.

Dell'Amico, M., M. Monaci, C. Pagani, D. Vigo. 2007. Heuristic approaches for the fleet size and mix vehicle routing problem with time windows. *Transportation Science* **41**(4) 516–526.

Diana, M., M. M. Dessouky, N. Xia. 2006. A model for the fleet sizing of demand responsive transportation services with time windows. *Transportation Research Part B* **40** 651–666

Eberly, J. C., J. A. Van Mieghem. 1997. Multi-factor dynamic investment under uncertainty. *Journal of Economic Theory* **75** 345–387.

Eilon, S., C. D. T. Watson-Gandy, N. Christofides. 1971a. *Distribution Management: Mathematical Modelling and Practical Analysis.* Charles Griffin & Company Limited, London.

Eilon, S., C. D. T. Watson-Gandy, A. Heilbron. 1971b. A vehicle fleet costs more. *International Journal of Physical Distribution* **1** 126–132.

Etezadi, T., J. E. Beasley. 1983. Vehicle fleet composition. *Journal of the Operational Research Society* **34** 87–91.

Fagerholt, K. 1999. Optimal fleet design in a ship routing problem. *International Transactions in Operational Research* **6** 453–464.

Filippi, C., G. Romanin-Jacur. 2002. Multiparametric demand transportation problem. *European Journal of Operational Research* **139** 206–219.

Fleischer, L. 2004. A fast approximation scheme for fractional covering problems with variable upper bounds. *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms,* New Orleans, LA. SIAM, Philadelphia, PA, 1001–1010.

Foote, B. L. 1980. A discussion of the properties of a basic simplex algorithm for generating the decision regions of Bereanu. M. A. H. Dempster, ed. *Stochastic Programming,* Academic Press, London, 207–222.

Godwin, T., R. Gopalan, T. T. Narendran. 2008. Tactical locomotive fleet sizing for freight train operations. *Transportation Research Part E* **44** 440–454.

Golden, B., A. Assad, L. Levy, E. Gheysens. 1984a. The fleet size and mix vehicle routing problem. *Computers & Operations Research* **11**(1) 49–66.

Golden, B. L., F. J. Gheysens, A. A. Assad. 1984b. On solving the vehicle fleet size and mix problem. J. P. Brans, ed. *Operational Research '84.* North-Holland, Amsterdam, 721–734.

Gould, J. 1969. The size and composition of a road transport fleet. *Operational Research Quarterly* **20** 81–92.

Hall, N. G., C. Sriskandarajah, T. Ganesharajah. 2001. Operational decisions in AGV-served flowshop loops: Fleet sizing and decomposition. *Annals of Operations Research* **107** 189–209.

Harrison, J. M., J. A. Van Mieghem. 1999. Multi-resource investment strategies: Operational hedging under demand uncertainty. *European Journal of Operational Research* **113** 17–29.

Imai, A., F. Rivera. 2001. Strategic fleet size planning for maritime refrigerated containers. *Maritime Policy and Management* **28**(4) 361–374.

Jones, C. N. 2005. *Polyhedral Tools for Control.* Ph.D. dissertation, Control Group, Department of Engineering, University of Cambridge, Cambridge, UK.

Khouja, M. 1999. The single-period (news-vendor) problem: Literature review and suggestions for future research. *Omega* **27** 537–553.

Kim, S. 2006. Gradient-based simulation optimization. L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, R. M. Fujimoto, eds. *Proceedings of the 2006 Winter Simulation Conference,* 159–167.

Kirby, D. 1959. Is your fleet the right size? *Operational Research Quarterly* **10** 252.

Klincewicz, J. G., H. Luss, M. G. Pilcher. 1990. Fleet size planning when outside carrier services are available. *Transportation Science* **24**(3) 169–182.

List, G. F., B. W. Wood, L. K. Nozick, M. A. Turnquist, D. A. Jones, E. A. Kjeldgaard, C. R. Lawton. 2003. Robust optimization for fleet planning under uncertainty. *Transportation Research Part E* **39** 209–227.

Listes, O., R. Dekker. 2005. A scenario aggregation-based approach for determining a robust airline fleet composition for dynamic capacity allocation. *Transportation Science* **39**(3) 367–382.

Nahapetyan, A., R. K. Ahuja, F. Z. Sargut, A. John, K. Somani. 2007. A simulation/optimization framework for locomotive planning. C. Liebchen, R. K. Ahuja, J. A. Mesa, eds. *ATMOS 2007 – 7th Workshop on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems,* IBFI, Dagstuhl, Germany, 259–276.

Netessine, S., G. Dobson, R. A. Shumsky. 2002. Flexible service capacity: Optimal investment and the impact of demand correlation. *Operations Research* **50**(2) 375–388.

Papier, F., U. W. Thonemann. 2008. Queuing models for sizing and structuring rental fleets. *Transportation Science* **42**(3) 302–317.

Porteus, E. L. 1990. Stochastic inventory theory. D. P. Heyman, M. J. Sobel, eds. *Stochastic Models,* Elsevier Science, Amsterdam, 605–652.

Renaud, J., F. F. Boctor. 2002. A sweep-based algorithm for the fleet size and mix vehicle routing problem. *European Journal of Operational Research* **140** 618–628.

Richetta, O., R. C. Larson. 1997. Modeling the increased complexity of New York City's refuse marine transport system. *Transportation Science* **31**(3) 272–293.

Sigurd, M. M., N. L. Ulstein, B. Nygreen, D. M. Ryan. 2005. Ship scheduling with recurring visits and visit separation requirements. G. Desaulniers, J. Desrosiers, M. M. Solomon, eds. *Column Generation,* Springer, New York, 225–245.

Simão, H. P., J. Day, A. P. George, T. Gifford, J. Nienow, W. B. Powell. 2008. An approximate dynamic programming algorithm for large-scale fleet management: A case application. *Transportation Science,* published online before print August 15, 2008.

Smilowitz, K. R., C. F. Daganzo. 2007. Continuum approximation techniques for the design of integrated package distribution systems. *Networks* **50**(3) 183–196.

Turnquist, M. A. 1985. Research opportunities in transportation system characteristics and operations. *Transportation Research Part A* **19**(5–6) 357–366.

Van Mieghem, J. A., N. Rudi. 2002. Newsvendor networks: Inventory management and capacity investment with discretionary activities. *Manufacturing & Service Operations Management* **4**(4) 313–335.

Van Mieghem, J. A. 2003. Capacity management, investment, and hedging: Review and recent developments. *Manufacturing & Service Operations Management* **5**(4) 269–302.

Vis, I. F. A., R. B. M. de Koster, M. W. P. Savelsbergh. 2005. Minimum vehicle fleet size under time-window constraints at a container terminal. *Transportation Science* **39**(2) 249–260.

Walkup, D. W., R. J.-B. Wets. 1969. Lifting projections of convex polyhedra. *Pacific Journal of Mathematics* **28**(2) 465–475.

Wallace, S. W. 1986. Decomposing the requirement space of a transportation problem into polyhedral cones. *Mathematical Programming Study* **28** 29–47.

Wets, R. J.-B. 1980. The distribution problem and its relation to other problems in stochastic programming. M. A. H. Dempster, ed. *Stochastic Programming,* Academic Press, London, 245–262.

Woods, D. G., F. C. Harris. 1979. Truck fleet composition for concrete distribution. *International Journal of Physical Distribution and Materials Management* **10** 3–14.

Wu, P., J. C. Hartman, G. R. Wilson. 2005. An integrated model and solution approach for fleet sizing with heterogeneous assets. *Transportation Science* **39**(1) 87–103.

Wyatt, J. K. 1961. Optimal fleet size. *Operational Research Quarterly* **12** 186–187.

Żak, J., A. Redmer, P. Sawicki. 2008. Multiple objective optimization of the fleet sizing problem for road freight transportation. *Journal of Advanced Transportation* **42**(4) 379–427.