Joint Inventory and Revenue Management with Removal Decisions

Alvaro Maggiar*
maggiara@amazon.com
Amazon.com

Ali Sadighian
alisadi@amazon.com
Amazon.com

August 14, 2017

Abstract

We study the problem of a retailer that maximizes profit through joint replenishment, pricing and removal decisions. This problem is motivated by the observation that retailers usually retain rights to remove inventory from their network either by returning it to the suppliers or through liquidation in the face of random demand and capacity constraints. We develop a tractable dynamic program by leveraging the concept of $L^\natural$-concavity that allows us to partially characterize the structure of optimal policies. Different modes of return rights are considered, and we explore extensions that consider lead-time, perishability and temporal demand dependency under capacity constraints. We additionally present illustrative numerical results.

Keywords— Inventory Management, Revenue Management, Dynamic Pricing, Dynamic Programming, Supermodularity, $L^\natural$-concavity, Capacity Management

1 Introduction

We consider in this paper the inventory management problem faced by a retailer who purchases and sells a variety of products. The retailer operates a periodic review of their inventory levels and adjusts them through different levers such as procurement, returns to suppliers, liquidation, and pricing, in order to maximize its net present value in the face of varying demand and costs.

The advent of online retailing has motivated a continuing and active stream of research that tackles the problem of simultaneous pricing and inventory management. Dynamic pricing has demonstrated significant benefits in revenue optimization by adjusting prices in response to inventory levels; with price and inventory shown to be strategic substitutes (Federgruen and Heching, 1999), so that raising one decreases the benefits of raising the other. Inventory management thus benefits from this more flexible interaction between a retailer and its customers. Retailers, however, also face their vendors on the other side of the supply chain, and this interaction should also be considered in the inventory management picture. The flexibility offered by dynamic over fixed pricing on the demand side can then somewhat be compared to that offered by vendor contracts on the supply side: both allow for hedging inventory management risk by lifting some rigidity in the exchanges, either with the end customer or the original source of the product. Most of the literature dealing with supply chain contracts focuses on single period problems and the coordination of the supply chain. The ability to return units to vendors in exchange for a (partial) refund, or simply to dispose of overstock inventory, is an important way for a retailer to hedge their inventory risk and subsequent bottom line, but also to manage their inventory in the face of capacity constraints. The latter is particularly important for retailers who offer products with varying seasonalities and whose shelf space is limited. In that case, the cost of carrying

*Corresponding author.

unsold inventory is not confined to its holding cost, but also to the lost opportunities prevented by the occupied space, which could have been used to carry more popular and profitable products instead.

We extend joint inventory and pricing models to incorporate return and liquidation channels. Return rights are usually negotiated as part as of vendor agreements, which specifies in particular the quantity of units purchased that can be returned, as well as the corresponding return (or buy-back) value. As a result, decisions must be taken not only based on the inventory level, but also on the level of returnable inventory. We present a joint replenishment, pricing and removal inventory management framework that relies on concepts recently ported from discrete convex analysis to inventory management with success, namely $L^\natural$-concavity. We derive its structural properties, discuss practical implementation details and present some illustrative numerical results. The numerical implementation section provides a detailed exposition of the tools and algorithms that practitioners would be interested in for real life implementations of inventory management problems.

The rest of this paper is structured as follows: we review in Section 2 the existing relevant literature and give more context to the problem, the model is developed in Section 3, following which we discuss extensions of the models in Section 4 and capacity management in Section 5; implementation details are then considered in Section 6 and numerical results presented in Section 7.

2 Literature Review

This paper belongs to a rich stream of literature devoted to the problem of joint inventory and pricing management. Early models of coordinated inventory control and pricing originating with Whitin (1955) and Mills (1959) consider a single period and deterministic setting, extended by such work as Karlin and Carr (1962) and Zabel (1970) to the case a stochastic demand. Petruzzi and Dada (1999) offer a review of single period pricing and inventory problems in the context of the seminal newsvendor problem.

A first foray into multiple period and stochastic problems is found in Zabel (1972). The area was later revived through the work of Federgruen and Heching (1999) who derived the structural properties of such periodic-review joint replenishment and pricing inventory problems, proving the optimality of a so-called base-stock list-price policy. Several extensions were then developed, notably to take into account a fixed ordering cost in Chen and Simchi-Levi (2004a), Chen and Simchi-Levi (2004b) or Huh and Janakiraman (2008); in the case of supply uncertainty in Li and Zheng (2006) and Feng (2010); and with batch ordering in Yang et al. (2014). Reviews of this extensive field can be found in Elmaghraby and Keskinocak (2003), Chan et al. (2004) and Chen and Simchi-Levi (2012).

In parallel to the research on joint pricing and replenishment, results pertaining to $L^\natural$-concave functions (Murota, 2003) were ported successfully from their field of discrete convex analysis to that of inventory management. The concept of $L^\natural$-concavity is related to supermodularity, which is widely used in inventory management to derive structural properties of optimal policies. $L^\natural$-concavity represents a stronger property than supermodularity and was first used in the inventory management literature in Lu and Song (2005) and in Zipkin (2008a) who revisited the lost-sales inventory model, inspiring a number of subsequent papers using the same tools to analyze inventory models (e.g. Huh and Janakiraman (2010); Pang et al. (2012); Gong and Chao (2013); Chen et al. (2014b)). A number of these latter papers demonstrate the applicability of $L^\natural$-concavity to joint replenishment and pricing problems. In a recent paper, Chen (2017) underlines the usefulness of $L^\natural$-convexity in inventory management.

A contribution of our paper is the addition of returnability considerations in a multiple periods inventory problem. Return rights have been considered in one form or another for single period problems in a number of papers, usually as a means of coordinating the supply chain between a vendor and a retailer. The seminal work of Pasternack (1985) introduced buy-back contracts and showed that coordination can be reached. A wider class of contracts is studied in Cachon (2003), although most of these contracts involved a fixed price newsvendor problem. Emmons and Gilbert (1998) extend the work of Pasternack (1985) for price dependent demand models. Other similar single-period models are studied in Padmanabhan and Png (1997), Taylor (2001) and Tsay (2002).

3 Joint Inventory, Pricing and Return Management

This section is devoted to our basic models. Note that we present some mathematical tools related to $L^\natural$-concavity in Appendix A, as well as some precisions regarding results concerning the preservation of supermodularity. We first describe the setting of the problem in Section 3.1. We then lay out the mathematical formulation of the basic models in Sections 3.2, 3.3 and 3.4. These models differ in the type of vendor contracts used. We first consider in Section 3.2 our main model in the case of absolute bounds on the return quantity allowed for each purchase (as described for example in (Ozlem, 2003, ch.2) for a newsvendor formulation), meaning that whenever a purchase order is made, all units are returnable for a given refund up to a specified limit. Section 3.3 on the other hand deals with the case of a fractional return right, so that a fractional amount of any purchase is returnable. Such a model was presented for instance in Pasternack (1985). The two models differ in the strength of the results obtained, justifying the focus on the former. The model with absolute bounds allows for the application of results related to $L^\natural$-concavity (defined in Appendix A), which are stronger than those obtained for fractional return rights where only supermodularity results can be derived. Finally, owing to the fact that an important fraction of products are either fully returnable or non-returnable, we consider the special case of full returnability, or equivalently only liquidability, in Section 3.4.

3.1 Problem Description

The problem considered is a finite-time, periodic-review, joint inventory and pricing problem wherein a retailer makes the following decisions at the beginning of each period $t = 0, \dots, T-1$: (i) purchasing quantity, (ii) return and liquidation quantities, and (iii) price at which to sell the product. Returns and liquidations differ in the value that can be extracted from discarding units on-hand. Return rights are usually negotiated with vendors upon placement of a purchasing order and tend to offer a certain insurance against low sales by allowing units to be returned to the vendors for a (usually high) fraction of the purchasing cost. The contract, however, stipulates a maximum number of units that can be returned. Different contract types exist and we consider two versions: one with an absolute bound for each purchase (Section 3.2), and a second that allows for a fraction of all purchased units to be returned (Section 3.3). Liquidation on the other hand is a disposition channel offered by a third party willing to take any amount of the inventory for pennies on the dollar. The definition of liquidation can be here enlarged to also encapsulate the case where no such avenue is available and simply means discarding inventory, potentially at a cost. We additionally consider the important particular case of full returnability in Section 3.4.

We assume in this paper that the stochastic demand possesses an additive form. More specifically, we suppose that the random demand at time t depends on the prevailing price p through the following relation:

$$D_t(p) = d_t(p) + \varepsilon_t, \tag{1}$$

where d_t is a deterministic function, strictly decreasing, and ε_t a random perturbation with mean 0 and support $[A, B]$, $A < 0 < B \leq +\infty$. Letting $\bar{p}$ denote the price upper bound, we require that $D_t(\bar{p}) + A \geq 0$. Note in particular that d_t represents the mean demand. This additive demand model is common in the literature and has been used extensively in related joint inventory-pricing problems, e.g. Chen and Simchi-Levi (2004a); Feng (2010); Allon and Zeevi (2011); Pang et al. (2012); Chen et al. (2014b); Yang et al. (2014).
Further let $p_t := d_t^{-1}$ be the inverse demand function and $r_t(d, x)$ be the expected revenue function, where x is the inventory level. It follows that $r_t(d, x) = r_t(d) = dp_t(d)$ in the backlogging case and $r_t(d, x) = d\mathbb{E}[\min(D_t(d), x)]$ in the lost-sales case. We then make the following assumption, focusing on the backlogging case:

Assumption 1
For all $t = 0, \dots, T$, the inverse function of d_t, denoted p_t, is continuous and strictly decreasing. Furthermore, the expected revenue $r_t(d) = dp_t(d)$ is concave in d.

This assumption, also fairly common in the literature, can be found in Chen and Simchi-Levi (2004a), Talluri and Van Ryzin (2006) and Pang et al. (2012), for example.

Remark 1
We focus in this paper on the backlogging case. The results extend to the case of lost sales so long as the revenue function remains concave. However, when lost sales are assumed, the revenue function $r_t(d,x) = p_t(d)\mathbb{E}[\min(D_t(d),x)]$ is no longer guaranteed to be concave in d, let alone jointly concave in d and x. The concavity of $r_t(d,x)$ will depend on the specific demand model used and discussions can be found in Kocabiyikoglu and Popescu (2011) and Chen et al. (2014b). In particular, Proposition 3 in Chen et al. (2014b) provides sufficient conditions for $r_t(d,x)$ to be $L^\natural$-concave.

Rewards and costs in future periods are discounted by a discount factor $\gamma \in [0,1]$. Let:

x_t = inventory level at the beginning of period t, before decisions,
y_t = returnable inventory level at the beginning of period t, before decisions,
$\mathbf{z}_t$ = (x_t, y_t), state of the system at the beginning of period t,
c_t = unit purchasing cost in period t,
$h_t^+(x)$ = inventory cost incurred in period t when its ending inventory position is x,
$h_t^-(x)$ = stockout cost incurred in period t,
s = return value,
l = liquidation value,
q_r = quantity of returnable units returned (>0) or purchased (<0),
q_{nr} = quantity of units to be liquidated (>0) or non-returnable units to be purchased (<0).

In each period $t = 1,\ldots,T-1$, the sequence of events is the following:

1. Based on the inventory level x_t and returnable level y_t, the retailer decides on: (i) a quantity q_r of returnable units to be purchased or returned, (ii) a quantity q_{nr} of non-returnable units to be purchased or liquidated, (iii) a target mean demand d (equivalently a price $p = p_t(d)$).
2. A random demand $D_t(d) = d + \varepsilon_t$ realizes.
3. Stockouts are satisfied in the period in which they occur at a cost greater than c_t and any remaining inventory is carried over to the next period, incurring a holding cost.

The choice made as to how to handle stockouts is to have them satisfied in the period in which they occur at a cost greater than c_t. This choice is motivated by the flexibility offered by such modeling in approximating behaviors that lie in between full backlogs and lost sales. In particular, most products have a manufacturer's suggested retail price (MSRP) that corresponds to their base price, and the pricing decision made by the retailer is the amount of markdown to apply. In such a scenario, we may consider a per-unit stockout cost of the form $c_t + \alpha(p - c_t) + g$, where p is the base price of the item, $\alpha \in [0,1]$ the fraction of customers who will decide not to buy the item due to being out of stock, and g the loss of goodwill suffered because of the stockout, which can be interpreted as the resulting loss in customer lifetime value. In particular, we observe that a value of $\alpha = 1$ would approximate a lost-sale scenario, while $\alpha = 0$ approximates a full backlogging scenario.

To help establish the structural properties of the results, we make the following standard assumptions.

Assumption 2

(i) The holding and stockout cost functions $h_t^+(x)$ and $h_t^-(x)$ are convex,

(ii) For any time $t = 0,\ldots,T-1$, the purchasing cost is greater than the return value: $c_t > s$, $\forall t$,

(iii) The liquidation value l is less than the return value s: $l < s$.

For ease of exposition, we will consider piecewise linear forms for the holding and stockout cost functions:

$$h_t^+(x) = h_t^+ x^+, \qquad h_t^-(x) = h_t^- x^-, \qquad h_t(x) := h_t^+(x) + h_t^-(x),$$

where the per-unit stockout cost h_t^- is expressed as the sum of the purchasing cost c_t in period t, and supplemental cost k_t: $h_t^- = c_t + k_t$. The supplemental stockout cost can be interpreted in different ways, either as the extra cost of fulfilling unmet demand through an express order, or as representing the loss of goodwill and subsequent life-time value of the customer suffered from being unable to satisfy demand from on-hand inventory.

We also consider limiting the total number of units that can be purchased in any given period. Letting $-\underline{q}_t$ be this limit in period t, we will require the following assumptions to preserve the structure of our results.

Assumption 3

For any time t ($0 \le t < T$), at least one of the following conditions hold true:

1. $\underline{q}_t = -\infty$,
2. $c_t + h_t^+ \ge \gamma c_{t+1}$ and $(c_t + k_t) + h_t^+ \ge \gamma(c_{t+1} + k_{t+1})$.

Where $\underline{q}_t$ is the lower bound on $q_t^r + q_t^{nr}$ at time t, and γ is the discount factor.

The first condition simply states that there is no limit on the amount of purchasable units. The second condition states that it is neither advantageous to plan a backlog for the next period in the current period, nor buy for the next period in the current one. Satisfying one of these two conditions in any time period ensures that that the problem preserves the appropriate structure and that all our results hold. Note that these conditions are not the tightest and could be relaxed.

We further define three functions b_t^r, b_t^{nr} and b_t as follows:

$$b_t^r(q) := sq^+ - c_t q^-, \qquad b_t^{nr}(q) := lq^+ - c_t q^-, \qquad b_t(q, y) := sq^+ + (l - s)(q - y)^+ - c_t q^-.$$

b_t^r, b_t^{nr} represent the costs and revenues associated with purchasing and returning/liquidating returnable and non-returnable units, respectively. We will later show that it is never beneficial to liquidate a unit that can be returned instead, leading to the unified function $b_t(q, y)$.

The state variables x_t and y_t, which represent the inventory level and returnable inventory level, respectively, evolve according to the purchases, sales, returns and liquidations that take place. Any returned, liquidated or sold unit lowers the inventory level by one, while only returned units lower the returnable level. Similarly, non-returnable purchased units increase only the inventory level, while returnable purchased units increase both inventory and returnable levels. The manner in which the latter occurs is dependent on the returnability rights negotiated with the vendor, which are developed in the corresponding sections here-under. We also point out that inventory is fungible between purchases, even across vendors, so that it is not necessary to keep track of which order a unit came from in order to return it to the appropriate vendor, allowing us to only keep track of the overall returnable level.

Finally, we let $V_t(\mathbf{z}) = V_t(x, y)$ denote the maximum expected discounted profit for periods $t, \ldots, T$ when starting period t in state $\mathbf{z}$. We then make the following mild assumption on the terminal value function V_T.

Assumption 4

The terminal value function $V_T(\mathbf{z})$ is $L^\natural$-concave and the restrictions $V_T(x, .)$ are non-decreasing in y for any x with a slope less than or equal to s.

Most sensible terminal value functions will satisfy Assumption 4. In particular the functions $V_T(x, y) = 0$ or $V_T(x, y) = s\min(x, y) + l\max(x - y, 0)$ are viable choices.

3.2 Fixed Returnability

The first model we consider is one with absolute bounds on the amount of returnable units that can be purchased in any given period. Such a model has been considered in the literature, albeit in the context of the newsvendor model, in Ozlem (2003) and mentioned in Cachon (2003).

More precisely, we assume that in any period t, there exists a limit $-\underline{q}_t^r$ of returnable units that can be purchased. This limit is part of the negotiated contract with the vendor and we assume here that it is already known, although the model and proofs would carry over in the case of random limits. Ozlem (2003) for example suggest that the retailer share their demand forecast with the vendor and the threshold be set to some quantile of the demand distribution.

We prove a number of results about the value functions V_t and the structural properties of the optimal decisions. Using the concept of $L^\natural$-concavity we prove in this section that the value function V_t is $L^\natural$-concave (thus concave and supermodular) for all t and that the intuitive monotonicity properties of the optimal decisions hold, namely that:

- returns, liquidations and markdowns increase with inventory levels for a fixed returnable quota,
- returns increase and liquidations and markdowns decrease with increasing returnable quota for a fixed inventory level,
- total purchases increase with decreasing inventory levels,
- purchases of non-returnable units increase with increasing returnable levels for a fixed inventory level.

We additionally derive sensitivity bounds on those changes. The price to pay for the powerful results provided by the $L^\natural$-concavity is somewhat cumbersome proofs that require a number of reformulations and changes of variable to fit the required framework.

While the ultimate decision variables we are interested in are the quantity of purchased, returned, liquidated and expected sold units (note that the last one is equivalent to price), we need to transform the problem to make it amenable to being analyzed through the lens of $L^\natural$-concavity. This often implies expressing variables of interest as the difference between two auxiliary variables. To that purpose, we let $x_+ := x - q_{nr}$, $x_{++} := x_+ - d$ and $\hat{x}$ so that $x_+ - \hat{x}$ represents the amount of demand satisfied from inventory. Further define the following sets (dropping time indices for convenience):

$$A := \left\{(\mathbf{z}, x_+, x_{++}, q_r) : \ \underline{d} \le x_+ - x_{++} \le \bar{d}, \ q_r - y \le 0, \ \underline{q}^r \le q_r, \ 0 \le x_+, \ \max(0, \underline{q} - \underline{q}^r) \le x - x_+ \right\},$$
$$A_\varepsilon := \{(\mathbf{z}, x_+, x_{++}, q_r, \hat{x}) : \ (\mathbf{z}, x_+, x_{++}, q_r) \in A, \ 0 \le \hat{x} - q_r, x_{++} - \hat{x} \le \varepsilon, \ x_+ - \hat{x} \ge 0\}$$

The constraints enforced in set A are rather self-explanatory. The additional constraints enforced in set A_ε correspond to: i) the constraint that after purchases and returns, and fulfilling demand from on-hand inventory, the inventory level cannot be negative; ii) the constraint that we cannot fulfill more demand from on-hand inventory than the actual demand realization; and iii) the constraint that we cannot fulfill a negative amount of demand.

It follows from Proposition A.1 that the sets A and A_ε are $L^\natural$-convex for any ε. These sets represent the different constraints the model abides by. The optimality equation that characterizes the problem is then given by the following:

$$V_t(\mathbf{z}) := \max_{(x_+, x_{++}, q_r):(\mathbf{z}, x_+, x_{++}, q_r) \in A} f_t(\mathbf{z}, x_+, x_{++}, q_r), \tag{2}$$

$$\text{where } f_t(\mathbf{z}, x_+, x_{++}, q_r) := r(x_{++} - x_+) + b_t^r(q_r) + b_t^{nr}(x - x_+) + \mathbb{E}\left[u_t(\mathbf{z}, x_+, x_{++}, q_r | \varepsilon_t)\right], \tag{3}$$
$$u_t(\mathbf{z}, x_+, x_{++}, q_r | \varepsilon_t) := -h_t^-(x_{++} - q_r - \varepsilon_t) + \max_{\hat{x}:(\mathbf{z}, x_+, x_{++}, q_r, \hat{x}) \in A_{\varepsilon_t}} g_t(\mathbf{z}, q_r, \hat{x}),$$
$$g_t(\mathbf{z}, q_r, \hat{x}) := \gamma V_{t+1}(\hat{x} - q_r, y - q_r) - h_t^+(\hat{x} - q_r).$$

It also follows from Assumption 3 that the optimal $\hat{x}$ above is given by $\hat{x} = \max(0, x_+ - (d + \varepsilon_t)) = \max(0, x_{++} - \varepsilon_t)$; in other words, we always fulfill demand to the maximum possible extent.

Theorem 3.1
Suppose V_T is $L^\natural$-concave, then V_t, f_t, g_t and u_t are all $L^\natural$-concave for all t (and for any value of ε in the case of u.)

Proof 1
The proof is by induction. Suppose then that $V_{t+1}(\mathbf{z})$ is $L^\natural$-concave. It follows directly from the convexity of h^+ (and thus concavity of $-h^+$), Proposition A.3 (a), (e) and (c) that $g_t(\mathbf{z}, q_r, \hat{x})$ is $L^\natural$-concave.

Next, the $L^\natural$-concavity of $u_t(\mathbf{z}, x_+, x_{++}, q_r|\varepsilon_t)$ for any ε is established by Proposition A.3 (c) and (d), whence we also deduce the $L^\natural$-concavity of $\mathbb{E}\left[u_t(\mathbf{z}, x_+, x_{++}, q_r|\varepsilon_t)\right]$ through Proposition A.3 (b).

Proposition A.3 (e) and (a) then yield the $L^\natural$-concavity of $f_t(\mathbf{z}, x_+, x_{++}, q_r)$.

A final application of Proposition A.3 (d) gives the $L^\natural$-concavity of $V_t(\mathbf{z})$.

The following corollary allows us to simplify the formulation (3) by lumping together returnable and non-returnable units.

Corollary 1
Suppose V_T satisfies Assumption 4, then for all $t = 0, \ldots, T-1$, V_t satisfies the same properties and it is never optimal to liquidate a unit that could be returned instead. Additionally, it is never optimal to purchase a non-returnable unit when a returnable unit can be purchased.

Letting $x^+ := (x - D_t(d) - q)^+$ and $y^+ := y - \min(\max(q, -\underline{q}^r), y)$, we can then formulate the optimality equation as follows:

$$V_t(x,y) = \max_{\underline{d} \le d \le \bar{d}, \underline{q} \le q} \; r_t(d) + b_t(q,y) - \mathbb{E}\left[h_t^-(x - D_t(d) - q)\right] - \mathbb{E}\left[h_t^+(x^+)\right] + \gamma \mathbb{E}\left[V_{t+1}(x^+, y^+)\right]. \tag{4}$$

We now derive the monotonicity properties and bounds of the optimal decisions. The main tool in achieving this is Lemma A.1 that relies on the $L^\natural$-concavity of the objective function in the optimality equation.

Lemma 3.1
For any $t = 0, \ldots, T-1$ and any $\omega > 0$, the following results hold:

(a) $d(x,y)$ is nondecreasing in x and nonincreasing in y and the following inequalities hold:

$$d(x,y) \le d(x+\omega, y) \le d(x,y) + \omega$$
$$d(x,y) \ge d(x, y+\omega) \ge d(x,y) - \omega,$$

(b) $q_r(\mathbf{z})$ is nondecreasing in $\mathbf{z}$ and the following inequalities hold:

$$q_r(\mathbf{z}) \le q_r(\mathbf{z} + \omega\mathbf{e}) \le q_r(\mathbf{z}) + \omega.$$

(c) $q_{nr}(x,y)$ is nondecreasing in x and nonincreasing in y and the following inequalities hold:

$$q_{nr}(x,y) \le q_{nr}(x+\omega, y) \le q_{nr}(x,y) + \omega$$
$$q_{nr}(x,y) \ge q_{nr}(x, y+\omega) \ge q_{nr}(x,y) - \omega,$$

Proof 2
See Appendix B. •

Lemma 3.1 describes how optimal decisions vary with $\mathbf{z}$. These monotonicity properties are consistent with intuition and we observe in particular the following:

- returns, liquidations and markdowns increase with inventory levels for a fixed returnable quota,
- returns increase and liquidations and markdowns decrease with increasing returnable quota for a fixed inventory level,
- total purchases increase with decreasing inventory levels,
- purchases of non-returnable units increase with increasing returnable levels for a fixed inventory level.

3.3 Fractional Returnability

We consider here a model in which the retailer benefits from fractional return rights. By fractional return rights, we mean that for any amount of units purchased from the vendor, a fraction α of them will be returnable. This behavior differs from the assumption used in Section 3.2 where all purchased units were returnable up to a threshold. Obviously, a null fraction would imply that units are not returnable and the problem would simplify to the combined pricing and inventory problem of Federgruen and Heching (1999) in the absence of a liquidation channel. On the other end of the spectrum, a fraction of 1 corresponds to the case where all units are returnable. Both extreme cases are important and admit the same modeling, only swapping return and liquidation values s and l. These particular cases are studied in Section 3.4.

This returnability assumption introduces a relation between returnable and non-returnable units, for purchased units. In particular, we established in Lemma 3.1 that $q_r(x,y)$ was nondecreasing in y, while $q_{nr}(x,y)$ was nonincreasing in y in the case of absolute bounds on the returnable units; in other words they evolve in opposite directions for increasing values of y. However, because we are now enforcing a relation between returnable and non-returnable units purchased, it is clear that the results obtained for fractional return rules will differ.

A major difference between fractional and absolute return rules is that the structure on which the optimization problems are framed for the fractional returns is no longer a lattice, which causes the loss of $L^\natural$-concavity properties and additional difficulties in the establishment of the results. We show however that two key properties, namely supermodularity and concavity, of the value function are preserved. We recall that $L^\natural$-concave functions are in particular supermodular and concave, so that the results for fractional returns are in some sense weaker than the ones obtained for absolute returns.

Because we can no longer separate returnable from non-returnable units in the same way as in the absolute return rights case, we now let q_r and q_{nr} represent only the returned and liquidated units, but they no longer stand for the purchased units as well. Consequently, we now have $q_r \geq 0$ and $q_{nr} \geq 0$. The opposite of the purchased units will be denoted by $q \leq 0$, so that the purchased units are given by $-q$. We also slightly amend the notation from Section 3.2 so that $\tilde{x} := x - q_r - q$, $x_+ := \tilde{x} - q_{nr}$, $x_{++} := x_+ - d$ and $\tilde{y} = y - q_r - \alpha q$. Additionally define the following sets:

$$\begin{aligned} D &:= \left\{\mathbf{v} :\ \underline{d} \leq x_+ - x_{++} \leq \bar{d},\ \tilde{x} - x_+ \geq 0,\ \tilde{x} \geq 0,\ x_+ \geq 0,\ \tilde{y} \geq 0,\ q \leq 0,\ q_r \geq 0\right\}, \\ S_\varepsilon(x_+, x_{++}, \tilde{y}) &:= \{\hat{x} :\ \hat{x} \geq 0,\ x_{++} - \hat{x} \leq \varepsilon,\ x_+ - \hat{x} \geq 0\}, \\ A &:= \{\mathbf{v} \in D : \mathbf{A}\mathbf{v} = \mathbf{z}\} \end{aligned}$$

where:

$$\mathbf{v} = (\tilde{x}, x_+, x_{++}, \tilde{y}, q_r, q)^T, \qquad \mathbf{z} = (x, y)^T, \qquad \mathbf{A} = \begin{pmatrix} 1 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & \alpha \end{pmatrix}.$$

Note that the matrix equality corresponds to the following set of constraints:

$$\begin{cases} \tilde{x} + q_r + q & = x \\ \tilde{y} + q_r + \alpha q & = y \end{cases}.$$

The first set of inequalities in the definition of the set D represent the constraints on the mean demand (equivalently the price), the second states that the number of liquidated units is non-negative, the third inequality enforces the fact that we cannot return more units than we have on-hand and the fourth that we cannot return more units than we have returnable units.

With this notation, the optimality equation that characterizes the problem is given by the following:

$$\begin{aligned} V_t(x,y) &:= \max_{\mathbf{v}\in A} f_t(\mathbf{v}), \\ \text{where } f_t(\mathbf{v}) &:= r_t(x_{++}-x_+) + cq + sq_r + l(\tilde{x}-x_+) + \mathbb{E}\left[u_t(x_+,x_{++},\tilde{y}|\varepsilon_t)\right], \\ u_t(x_+,x_{++},\tilde{y}|\varepsilon_t) &:= -h_t^-(x_{++}-\varepsilon_t) + \max_{\hat{x}\in S_{\varepsilon_t}(x_+,x_{++},\tilde{y})} g_t(\hat{x},\tilde{y}), \\ g_t(\hat{x},\tilde{y}) &:= \gamma V_{t+1}(\hat{x},\tilde{y}) - h_t^+(\hat{x}) \end{aligned} \tag{5}$$

Similarly to the case of absolute return rights, it follows from Assumption 3 that the optimal $\hat{x}$ above is given by $\hat{x} = \max(0, x_+ - (d+\varepsilon_t)) = \max(0, x_{++} - \varepsilon_t)$; in other words, we always fulfill demand to the maximum possible extent.

The following theorem is the equivalent to Theorem 3.1 in the absolute return rights case, except that the weaker properties being preserved in the fractional case are concavity and supermodularity.

Theorem 3.2
Suppose V_T is concave and supermodular, then V_t, f_t, g_t and u_t are all concave and supermodular for all t (and for any value of ε in the case of u.)

Proof 3
The proof is by induction. The result holds by assumption for $t = T$, so suppose now that V_{t+1} is concave and supermodular. The concavity and supermodularity of g_t then follows immediately for the convexity of h_t^+. The concavity of the part being maximized in the expression of u_t is a result of (Simchi-Levi et al., 2014, Proposition 2.1.15 (b)) and the concavity of u_t then follows from the convexity of h_t^-. Similarly, the supermodularity of the part being maximized is a result of (Simchi-Levi et al., 2014, Proposition 2.2.5 (e)), h_t^- is convex and thus submodular, yielding that u_t is supermodular. We then deduce the concavity and supermodularity of $\mathbb{E}[u_t(x_+,x_{++},\tilde{y}|\varepsilon_t)]$ from (Simchi-Levi et al., 2014, Proposition 2.1.3 (e)) and (Simchi-Levi et al., 2014, Proposition 2.2.5 (d)), respectively. The concavity of r_t and linearity of the other components of f_t then ensure that f_t is also concave and supermodular.

Finally, to establish the concavity and supermodularity of V_t, we apply Proposition A.4, observing that B is here the identity matrix and $A \geq 0$, which is sufficient to satisfy the conditions of the theorem. •

The monotonicity of the optimal decisions is however more arduous. As mentioned in Section A.2, we cannot apply (Chen et al., 2013, Remark 2) for which we produce a counter-example in Appendix A.3. The study of the monotonicity of the decisions with respect to the state variables is complicated by the presence of constraints in the optimization problem (5) that prevent the use of standard monotone comparative statics results (Milgrom and Shannon, 1994; Topkis, 2011), which require that the feasible set be a lattice. The extension of monotone comparative statics to non-lattice sets has been studied, e.g. Quah (2007) and Barthel et al. (2015), by defining a $\mathcal{C}$-flexible order on sets instead of the typical strong set order. Unfortunately, the present feasible set does not satisfy those conditions either and we must resort to a more tedious sensitivity analysis to derive the desired results (Fiacco and Ishizuka, 1990; Strulovici and Weber, 2010). Consequently, we omit such proofs in this paper.

3.4 Full and Null Returnability Rights

An important particular case of the previous models is that of full and null returnability rights. This corresponds to the cases where either units are all returnable, or none is, respectively. These two cases are modeled identically, differing only by the value of removing a unit, which is s in the former and l in the latter. Consequently, we formulate the problem with full returnability, keeping in mind that the formulation would be identical in the case of null returnability. In this setting, it is unnecessary to keep track of the returnable quota and the state space reduces to the inventory level x_t at time t. The problem is similar to the combined pricing and inventory control problem of Federgruen and Heching (1999) with the addition of a removal channel.

We again let $q \leq 0$ be the opposite of the purchased units and $q_r \geq 0$ the number of returned units. We let here $\tilde{x} := x - q$ be the state after purchases have been made, $x_+ := x - q_r$ be the state after returns have occurred, and $x_{++} := x_+ - d$. Further defined the following set:

$$A := \left\{(x, \tilde{x}, x_+, x_{++}) : \underline{d} \leq x_+ - x_{++} \leq \bar{d},\ \tilde{x} \geq x,\ x_+ \leq \tilde{x}\right\}.$$

The problem can then be formulated as follows:

$$V_t(x) := cx + W_t(x), \tag{6}$$

$$\text{where } W_t(x) := \max_{(\tilde{x}, x_+, x_{++}) : (x, \tilde{x}, x_+, x_{++}) \in A} f_t(\tilde{x}, x_+, x_{++}), \tag{7}$$

$$f_t(\tilde{x}, x_+, x_{++}) := r_t(x_+ - x_{++}) - c_t\tilde{x} + s(\tilde{x} - x_+) + w_t(x_{++}),$$

$$w_t(x_{++}) := -\mathbb{E}\left[h_t(x_{++} - \varepsilon_t)\right] + \gamma\mathbb{E}\left[V_{t+1}\left((x_{++} - \varepsilon_t)^+\right)\right]$$

The arguments developed in Section 3.2 can be applied identically to show that provided $V_T(x)$ is concave, we have that for any t, $f_t(\tilde{x}, x_+, x_{++})$ is $L^\natural$-concave and V_t is concave.

Theorem 3.3
Suppose V_T is concave, then V_t, W_t and w_t are concave, and f_t is $L^\natural$-concave for all t.

We may also derive the following monotonicity results for the decision variables:

Lemma 3.2
For any $t = 0, \ldots, T-1$, $d(x)$, $q(x)$ and $q_r(x)$ are nondecreasing in x and the following inequalities hold for any $\omega > 0$:

$$d(x) \leq d(x + \omega) \leq d(x) + \omega,$$
$$q(x) \leq q(x + \omega) \leq q(x) + \omega,$$
$$q_r(x) \leq q_r(x + \omega) \leq q_r(x) + \omega.$$

Federgruen and Heching (1999) prove the optimality of a base-stock list-price policy for their combined pricing and inventory control problem. When a removal channel is considered, the optimal policy admits a symmetric counterpart that leads to what can be described as an *interval-stock list-prices* policy. An interval-stock list-prices policy is characterized by two stock level and price pairs, or equivalently stock level and expected demand pairs $(\underline{x}^*, \underline{d}^*)$ and $(\bar{x}^*, \bar{d}^*)$, with $\underline{x}^* \leq \bar{x}^*$ and $\underline{d}^* \leq \bar{d}^*$. Then, if the inventory level is below $\underline{x}^*$, it is increased to the stock level $\underline{x}^*$ and the expected demand is set to $\underline{d}^*$; if the inventory level is above the stock level $\bar{x}^*$, it is decreased to the stock level $\bar{x}^*$ and the expected demand is set to $\bar{d}^*$; and if the inventory level is between $\underline{x}^*$ and $\bar{x}^*$, no order or removal is triggered and the expected demand is set to $d(x)$ such that $\underline{d}^* \leq d(x) \leq \bar{d}^*$ and $d(x)$ is non-decreasing in x.

Theorem 3.4
The optimal policy for the full or null returnability joint pricing and inventory management problem is an interval-stock list-prices policy.

Proof 4
See Appendix C. •

4 Extensions

We present in this section a number of extensions of our joint inventory, pricing and return management. We focus on the model with absolute return rights of Section 3.2 given its structural superiority over that with fractional return rights. Its $L^\natural$-concavity is key in rather readily extending the results to allow for perishability, lead-times and demand correlation, as we will see below.

4.1 Lead Times

The model presented in Section 3 assumes that ordered units arrive immediately. When lead times are sufficiently small, they may indeed be disregarded; however in many practical cases the assumption is unrealistic, for example for retailers that procure from overseas manufacturers.

In the absence of pricing decisions, inventory models with fixed lead times (also known as time lags) have been studied in Karlin and Scarf (1958) and Scarf (1959); and stochastic lead times have been considered in Kaplan (1970); Nahmias (1979); Ehrhardt (1984). Some of these results were revisited in Zipkin (2008a) in the context of $L^\natural$-concavity.

When pricing is considered the structure of joint inventory and pricing problems with lead times was partially derived in Pang et al. (2012) using $L^\natural$-concavity, following the work of Zipkin (2008a). Their results can directly be ported to our model by using the same state space description as those papers, to which we add the returnability quota that need not be differentiated by lead time considering the assumption of inventory fungibility. Consequently, for a fixed lead time L the state of the system is described by a vector $\mathbf{z} = (x_0, x_1, \ldots, x_{L-1}, y)$, where for any $0 \le l \le L-1$, x_l describes the sum of the inventory on-hand and scheduled to arrive within the next l periods. In particular, x_0 represents the on-hand inventory and x_{l-1} the inventory position. y describes the returnable quota, as before. All the results derived in Section 3.2 and in Pang et al. (2012) about the preservation of $L^\natural$-concavity and the monotonicity of the optimal decision variables can then be shown to hold in a similar manner. Zipkin (2008a) also discusses how the results can be used in the case of stochastic lead times.

4.2 Perishability

A number of consumable products have a maximum shelf-life and expire after some amount of time. The inventory problem for perishable products, first considered in Fries (1975) and Nahmias (1975), has been studied in the case of joint inventory and pricing decisions in Chen et al. (2014b). Their work also uses $L^\natural$-concavity, and their results can thus easily be ported to our model. Letting l be lifetime of the product, the formulation calls for a state vector $\mathbf{z} = (x_1, x_2, \ldots, x_{l-1}, y)$, where x_i is the total amount of inventory with residual lifetimes no more than i periods, $i = 1, \ldots, l-1$. With this notation, and similarly to the extension to lead times, most of the results derived in Section 3.2 and in Chen et al. (2014b) about the preservation of $L^\natural$-concavity and the monotonicity of the optimal decision variables hold. There are nonetheless some different expected and observed behaviors because not all units possess the same value depending on their remaining life. As a result, it might at times be beneficial to simultaneously markdown, dispose of and purchase units. Furthermore, returnability for perishable products likely calls for a more sophisticated approach, since it is likely that only units with a minimum remaining shelf-life can be returned or liquidated. Such additional features may be considered in future work.

4.3 Dependent Demand

Most multiple period inventory problems assume that the demand realization period over period are uncorrelated. This assumption is of little realism for many products, especially products subject to seasonality and popularity effects such as toys, books or media. Inglehart and Karlin (1962) introduced a model with a Markov-modulated demand process characterized by a *state of the world* w on which the demand distribution depends. Such a model has further been studied in Song and Zipkin (1993) and Sethi and Cheng (1997); and Beyer et al. (2010) offer a more detailed exposition of the problem.

Our framework readily accommodates such parametrization of the demand evolution. We consider instead a slightly modified version of this demand evolution model. To take into account a sense of the demand correlation from period to period, we use the variable w to capture the state of the demand. This state variable is not cardinal, but rather a label or descriptor of the state of the world at time t, and the demand realization at time t is then conditioned on this state w_t. Additionally, there exists a transition function Ξ_t mapping demand realizations at time t to demand states at time $t+1$. Ξ_t captures the correlation from period to period. For example, if we assume a positive correlation from period to period and the demand realization d_t at time t is high, $\Xi_t(d_t)$ will map to a state w_{t+1} corresponding to a state of high demand. The optimality equation formulated in Section 3.2 is then modified as follows:

$$\begin{aligned} V_t(\mathbf{z}, w) &:= \max_{(x_+, x_{++}, q_r):(\mathbf{z}, x_+, x_{++}, q_r)\in A} f_t^w(\mathbf{z}, x_+, x_{++}, q_r), \\ \text{where } f_t^w(\mathbf{z}, x_+, x_{++}, q_r) &:= r(x_{++} - x_+) + b_t^r(q_r) + b_t^{nr}(x - x_+) + \mathbb{E}\left[u_t^w(\mathbf{z}, x_+, x_{++}, q_r|\varepsilon_t^w)\right], \\ u_t^w(\mathbf{z}, x_+, x_{++}, q_r|\varepsilon_t^w) &:= -h_t^-(x_{++} - q_r - \varepsilon_t^w) + \max_{\hat{x}:(\mathbf{z}, x_+, x_{++}, q_r, \hat{x})\in A_{\varepsilon_t^w}} g_t^w(\mathbf{z}, q_r, \hat{x}), \\ g_t^w(\mathbf{z}, q_r, \hat{x}) &:= \gamma V_{t+1}(\hat{x} - q_r, y - q_r, \Xi_t(\varepsilon_t^w)) - h_t^+(\hat{x} - q_r). \end{aligned}$$

It is easily seen that all the structural results derived in Section 3.2 still hold when we augment the state space with the state variable w.

When we consider a finite number of possible demand realizations, as is the case in many practical implementations where demand distributions are represented with only finite number of quantiles (e.g. Raz and Porteus (2006)), the information of the conditional demand distributions at time t can be captured by a probability matrix P_t whose rows correspond to the states at time t and the columns to the possible demand realizations. The elements P_t^{ij} of the matrix are then the probability of occurrence of the j-th demand realization given state i.

Example 4.1

We examine a case with 3 possible demand states: low (l), medium (m) and high (h); and in which demand is modeled through six possible realizations in each periods ($d^i,\ i = 1, \ldots, 6$) with the first two leading to a state of low demand, the next two to a state of medium demand and the last two to a state of high demand.

We may for example assume that given a state, it is twice as likely that we will end up in the same state as in one of the other two states. Such an assumption would lead to the following probability matrix P_t and mapping function Ξ_t:

$$P_t = \begin{pmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{8} & \frac{1}{8} & \frac{1}{8} & \frac{1}{8} \\ \frac{1}{8} & \frac{1}{8} & \frac{1}{4} & \frac{1}{4} & \frac{1}{8} & \frac{1}{8} \\ \frac{1}{8} & \frac{1}{8} & \frac{1}{8} & \frac{1}{8} & \frac{1}{4} & \frac{1}{4} \end{pmatrix} \qquad \Xi_t = \begin{pmatrix} l & l & m & m & h & h \end{pmatrix}.$$

Here we abuse notations to let $\Xi_t(d^i) = (\Xi_t)_i$.

Practitioners are often equipped with forecasts for each period that need to be woven together using some dependence information. If a high quantile of the demand distribution realizes in one period, higher quantiles are more likely to realize in the following period if we assume a positive correlation period over period. The formulation above is well suited for this purpose. The problem faced is then that of quantifying the dependence of the demand from one period to the next, as well as using it in a way to generate the probability transition matrices described above. A simple way to quantify the dependence between demands D_t and D_{t+1} in periods t and $t+1$, respectively, is through their correlation, from which we can then derive the probability transition matrices. A description of the process that uses copulas can be found in Appendix D.

5 Inventory Management under Capacity Constraints

Large retailers offer a wide selection of products on their platforms. Availability and short delivery times are paramount to their business, requiring that products be appropriately procured and stored in their warehouses to ensure satisfactory service levels. The total storage capacity often represents a constraint for the retailers and products compete for the available space. Additionally, inbound flows, in other words the total number or volume of incoming units, might also be constrained in any given period by the available staffing resources. Capacity management thus plays a central part in the inventory management problems that they face and necessitates at times aggressive removal of slow and small margin products in favor of more profitable ones, and a curbing of buying decisions to prevent violating the capacity constraints. Part of the capacity management is also long term planning of the supporting infrastructure to assess the pertinence of purchasing additional resources.

5.1 Capacity Management

The problem of multi-product inventory management in the presence of linking constraints is relatively well-studied in the literature, mostly in the context of deterministic demand, e.g. (Hadley and Whitin, 1963, sec. 6-4), Evans (1967), or in the context of stochastic demand but as a deterministic constraint, e.g. Smith and Agrawal (2000). In our framework, the on-hand inventory levels in each period are random variables that depend upon previous purchasing decisions and demand realizations. In the deterministic setting, the common approach is to dualize the constraint and bring it to the objective function using Lagrange multipliers. Constraints involving stochastic variables are usually dealt with as chance constraints (see Shapiro et al. (2014)). Imposing the constraints to hold for any path realization would result in particularly conservative decisions, especially in the presence of lead times, but would also fail to be appropriate given that the capacity constraints are often soft in that if the capacity of a storage facility is violated, delivering trucks or pallets can be made to temporarily wait outside before being brought in. A simple way to model the constraints is to express them in expectation, which allows for them to be dualized, just as in the deterministic case.

Consider the notationally simplified multi-period formulation of our inventory problem, which results from the following definitions:

State $\mathbf{z}_t$: represents the inventory state at time t. Depending on the level of complexity of the considered model, this can include the inventory level, the number of units purchased in the past and yet to arrive, etc...

Policy $\mathbf{u}_t(\mathbf{z}_t)$: represents the policy applied at time t as a function of the state $\mathbf{z}_t$. A policy consists of a set of decisions to be applied, such as units to be purchased, returned or liquidated, or markdown to be applied.

Randomness $\boldsymbol{\xi}_t$: represents the various sources of randomness in time period t. Such sources can be demand or lead times, for example.

Reward function $\pi_t(\mathbf{z}_t, \mathbf{u}_t(\mathbf{z}_t), \boldsymbol{\xi}_t)$: represents the profit generated in time period t as a function of the inventory state, policy and realization of the random variables. It can include purchasing costs, selling revenue, holding cost, shipping costs, etc...

Transition function $\mathbf{f}(\mathbf{z}_t, \mathbf{u}_t(\mathbf{z}_t), \boldsymbol{\xi}_t)$: represents the transition function that characterizes the dynamics of the problem.

The problem to be optimized at time t can then be expressed as follows:

$$V_t(\mathbf{z}_t) := \max_{\mathbf{u}\in U} \ \mathbb{E}\left[\sum_{s=t}^{\infty} \gamma^{s-t} \pi_s(\mathbf{z}_s, \mathbf{u}_s(\mathbf{z}_s), \boldsymbol{\xi}_s)\right]$$
$$s.t.\ \mathbf{z}_{s+1} = \mathbf{f}(\mathbf{z}_s, \mathbf{u}_s(\mathbf{z}_s), \boldsymbol{\xi}_s)$$

Its dynamic programming formulation reads:

$$\begin{aligned} V_t(\mathbf{z}_t) := \max_{\mathbf{u} \in U} \ \mathbb{E}\left[\pi_t(\mathbf{z}_t, \mathbf{u}_t(\mathbf{z}_t), \boldsymbol{\xi}_t) + \gamma V_{t+1}(\mathbf{z}_{t+1})\right] \\ s.t. \ \mathbf{z}_{t+1} = \mathbf{f}(\mathbf{z}_t, \mathbf{u}_t(\mathbf{z}_t), \boldsymbol{\xi}_t) \end{aligned}$$

Consider now the case of products $a \in \mathcal{A}$ that share linking constraints of the form:

$$\sum_{a \in \mathcal{A}} \mathbf{g}_t^a(\mathbf{z}_t^a, \mathbf{u}_t^a(\mathbf{z}_t^a)) \leq \mathbf{K}_t.$$

These constraints depend explicitly on the state variables in future periods, which are random variables. We express them in expectation and incorporate them to the original problem, which now reads:

$$\begin{aligned} \max_{\mathbf{u} \in U} \ \mathbb{E}\left[\sum_{a \in \mathcal{A}} \sum_{s=t}^{\infty} \pi_s^a(\mathbf{z}_s^a, \mathbf{u}_s^a(\mathbf{z}_s^a), \boldsymbol{\xi}_s^a)\right] \\ s.t. \ \mathbf{z}_{s+1}^a = \mathbf{f}^a(\mathbf{z}_s^a, \mathbf{u}_s^a(\mathbf{z}_s^a), \boldsymbol{\xi}_s^a) \\ \mathbb{E}\left[\sum_{a \in \mathcal{A}} \mathbf{g}_s^a(\mathbf{z}_s^a, \mathbf{u}_s^a(\mathbf{z}_s^a))\right] \leq \mathbf{K}_s \end{aligned}$$

By dualizing the constraints and bringing them into the objective function with the help of a set of dual variables $\boldsymbol{\lambda}$, we can restate the problem as a minimizing problem over the dual variables $\boldsymbol{\lambda}$, which is separable for a fixed value of $\boldsymbol{\lambda}$:

$$\min_{\boldsymbol{\lambda} \geq \mathbf{0}} \sum_{a \in \mathcal{A}} V_t^a(\mathbf{z}_t^a; \boldsymbol{\lambda}) \tag{8}$$

where (dropping the product superscript a for clarity)

$$\begin{aligned} V_t(\mathbf{z}_t; \boldsymbol{\lambda}) := \max_{\mathbf{u} \in U_t} \ \mathbb{E}\left[\pi_t^{\boldsymbol{\lambda}}(\mathbf{x}_t, \mathbf{u}_t(\mathbf{z}_t), \boldsymbol{\xi}_t) + V_t(\mathbf{z}_{t+1}; \boldsymbol{\lambda})\right] \\ s.t. \ \mathbf{z}_{t+1} = \mathbf{f}(\mathbf{z}_t, \mathbf{u}_t(\mathbf{z}_t), \boldsymbol{\xi}_t) \end{aligned}$$

and

$$\pi_t^{\boldsymbol{\lambda}}(\mathbf{z}_t, \mathbf{u}_t(\mathbf{z}_t), \boldsymbol{\xi}_t) := \pi_t(\mathbf{z}_t, \mathbf{u}_t(\mathbf{z}_t), \boldsymbol{\xi}_t) - \boldsymbol{\lambda}_t^T(\mathbf{g}_t(\mathbf{z}_t, \mathbf{u}_t(\mathbf{z}_t)) - \mathbf{K}_s).$$

The expression of constraints as expectations can be motivated by different versions of the law of large numbers. Consider for example the following capacity constraint:

$$\sum_{a \in \mathcal{A}} v^a x_t^a \leq K_t^v,$$

where v^a is the volume of product a and K_t^v the available storage capacity in period t. If we assume for example that the demand distributions are uncorrelated across products, then the random variables x_t^a are uncorrelated, and assuming further that they have finite mean and uniformly bounded variance, which should hold in any realistic scenario, then Tchebychev's theorem (Gnedenko, 2005) yields that the total volume $\sum_{a \in \mathcal{A}} v^a x_t^a$ at time t converges in probability to the expected volume $\sum_{a \in \mathcal{A}} v^a \mathbb{E}[x_t^a]$ as the number of products goes to infinity.

Remark 2
The dualization of a storage capacity constraint is akin to increasing the per-period holding costs so that the linear holding cost function in period t reads $h_t^+(x_t^a) = (h_t^+ + \lambda_{t+1} v^a) x_t^{a+}$.

5.2 Capacity Planning

The counterpart of the preceding section is the problem of capacity expansion. The capacity management described above is essential in satisfying the infrastructural constraints faced by the retailer, but can also help guide their long-term planning and support the decision to expand the capacity. Capacity expansion has been widely studied and is the subject of an extensive literature, see Manne (1967) or Luss (1982) for some of the earlier work. The question here is somewhat more elementary and is simply to evaluate the benefits of building additional storage facilities. Provided the problem formulated by (8) is constrained to begin with, it can be solved for increasingly large values of the capacity constraint, thus building a curve that measures the trade-off between additional capacity and Net Present Value (NPV) gains, which can then be evaluated against the curve measuring the cost of increasing capacity. Figure 1 illustrates the idea by plotting these curves, showing that in the illustrative curve, it would be beneficial to open an additional storage facility.

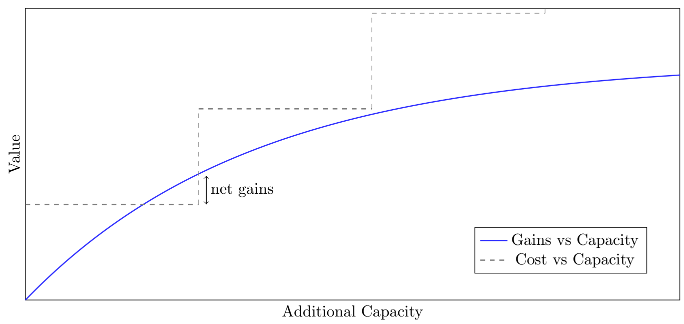

Figure 1: Illustrative example of the capacity expansion problem.

6 Implementation

We discuss in this section some important aspects of the practical implementation of the Stochastic Dynamic Program corresponding to inventory management problems possessing the $L^\natural$-concavity property such as our fixed returnability model derived in Section 3.2 or the one with full (or null) returnability rights of Section 3.4.

6.1 State Discretization and Interpolation

The inventory management problem is modeled as a continuous state stochastic dynamic program with state parameters $\mathbf{z}$ contained in an interval of $\mathbb{R}^n$ for any time period t, where n depends on whether we consider returnability, lead times or perishability for example. Implementing the dynamic program requires that the state space be discretized, which can only be done to some extent to avoid the curse of dimensionality. This discretization creates a grid on which an initially continuous function is approximated by extending its discrete approximation by interpolation for states falling in between grid points. We discuss in Section 6.1.1 the pitfalls usually associated with discretization and interpolation, and show in Section 6.1.2 how the $L^\natural$-concavity of the value function is key in overcoming them and implementing an efficient dynamic program of the problem.

6.1.1 Discretization and Pitfalls

Discretizing the state space often has the negative drawback of being unstable in that it does not preserve some properties of the value function such as convexity (Cai and Judd, 2013). This represents a major problem because it prevents most structural properties of the problem from holding and renders the optimization procedures much more difficult to implement. Preserving the structural properties of the function requires that the discrete approximation of the original function be extendable, in some way, to a function with the same properties.

Consider for example the quadratic and concave function $f(\mathbf{z}) = \mathbf{z}^T Q \mathbf{z}$, where $Q = \begin{bmatrix} -1 & 0.7 \\ 0.7 & -1 \end{bmatrix}$. Figure 2 shows the contour plots of the original function and its bilinear interpolation from its values at integer points. We observe that the bilinear interpolation - which in spite of its name does not produce a linear function - causes undulations in the resulting interpolated surface. These artifacts are not anodyne and lead to spurious local minima. As a consequence, the monotone properties of the optimal policies are lost and the artificial local optima can lead to suboptimal decisions.

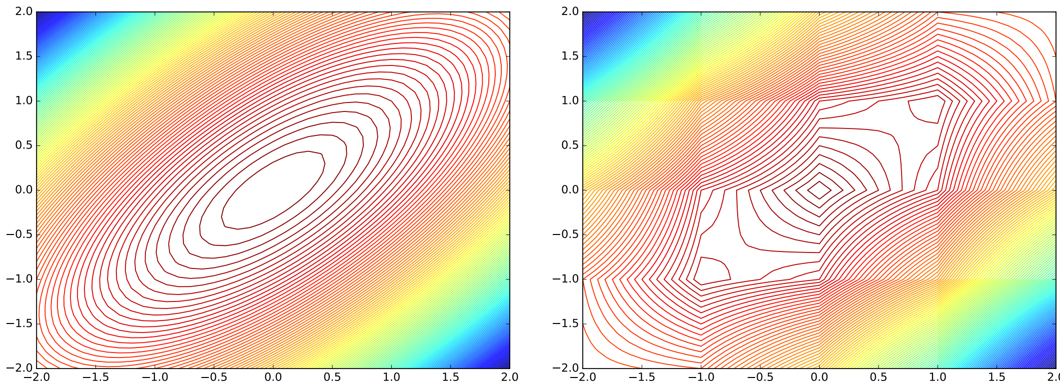

Figure 2: Contour maps of the example function: (left) original, (right) bilinear interpolation.

6.1.2 Extension of $L^\natural$-concave functions

The previous section discussed the issues inherent to the necessary discretization of a dynamic program. It turns out that the $L^\natural$-concave functions yield naturally well-behaved extensions, which are in turn $L^\natural$-concave. The extensions of a discrete $L^\natural$-concave function to a continuous $L^\natural$-concave follows from the following lemma proved in Murota (2003) and reproduced in Chen et al. (2014a).

Lemma 6.1 – Murota (2003); Chen et al. (2014a)
Given a discrete $L^\natural$-concave function f, its global extension can be obtained as follows: first obtain its Lovász Extension for every unit hypercube in its domain, and then paste all these extensions together.

This lemma relies on the Lovász extension (Lovász, 1983), which is defined as follows:

Definition 6.1 – Lovász extension
For a function $f : \{0,1\}^N \to \mathbb{R}$ defined on the unit hypercube, its Lovász extension $f^L : [0,1]^N \to \mathbb{R}$ is defined by

$$f^L(\mathbf{x}) = \sum_{i=0}^{n} (x_{\sigma(i)} - x_{\sigma(i+1)}) f(S_i),$$

where σ is a permutation such that $x_{\sigma(1)} \geq \ldots \geq x_{\sigma(n)}$, $S_i = \{\sigma(1), \ldots, \sigma(i)\}$, $x_{\sigma(0)} = 1$ and $x_{\sigma(n+1)} = 0$.

(Note that if we think of S_i as a point on the hypercube, its coordinates are obtained by setting 1 to positions $\sigma(k)$, $k = 1, \ldots, i$ and 0 elsewhere.)

The Lovász extension is easy to visualize on the two-dimensional square as it consists in splitting it along its antidiagonal and fitting a linear plane on each obtained triangle, which is illustrated in Figure 3 (left). Lemma 6.1 then states that in order to obtain the continuous extension of a discrete $L^\natural$-concave function, we need only apply the above Lovász extension to each unit hypercube in its domain, as illustrated in Figure 3 (right). Chen et al. (2014a) discuss the approximation errors resulting from such discretizations in dynamic programs.

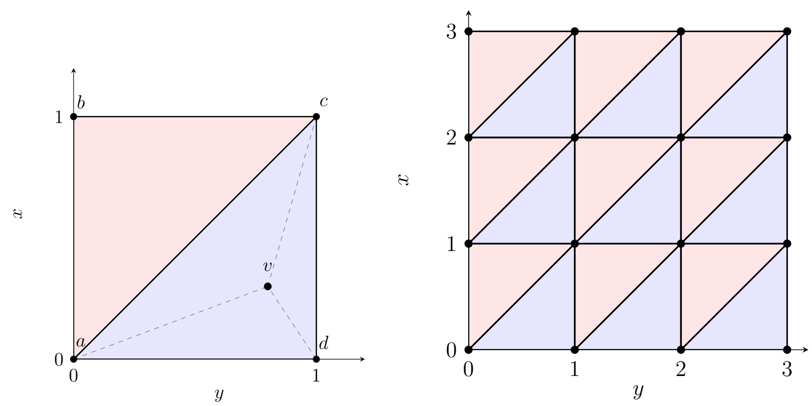

Figure 3: Illustration of the Lovász extension on the two-dimensional square (left) and from $\{0,1,2,3\}^2$ to $[0,3]^2$ (right).

As an illustration, since the example function used in Section 6.1.1 is $L^\natural$-concave, we plot in Figure 4 its contour plot using the $L^\natural$-concave extension instead of a bilinear interpolation. We observe that the function is better behaved and that its concavity is preserved. (In fact its $L^\natural$-concavity is preserved). We detail the algorithm in Appendix 2.

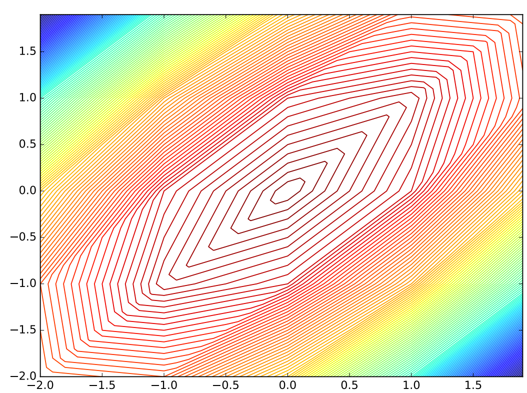

Figure 4: Lovász extension of the example function from Section 6.1.1.

Remark 3
Lovász (1983) also proved that a function is supermodular (resp. submodular) if and only if its Lovász extension is concave (resp. convex). Consequently, the Lovász extension can more generally be used as an interpolation method in the implementation of problems whose value functions are supermodular and concave (or submodular and convex), which are many in the inventory management literature. In particular, it is still useful in the case of the fractional returnability considered in Section 3.3.

6.2 Demand Discretization

6.2.1 Discretization

The formulation of the problem requires in most cases - that is unless we use a deterministic method - the evaluation of an expectation tied to the randomness in the demand in each period. A common practical approximation (e.g. Raz and Porteus (2006)) of the demand is obtained by selecting a number K of equally-probable quantiles of the demand and averaging the values corresponding to each of those demand realizations. Thus, an expression of the form $\mathbb{E}[g(D_t)]$ is approximated as:

$$\mathbb{E}[g(D_t)] \approx \frac{1}{K} \sum_{q=1}^{K} g(D_t^q), \tag{9}$$

where D_t^q is a quantile realization of the demand D_t. Note that these quantiles are themselves dependent on the decision variable d_t, which sets the expected demand over period t and since we assume an additive model, the quantiles need to be lifted accordingly.

6.2.2 Change of Probability

The Markov-Modulated formulation of the problem presented in Section 4.3 makes use of conditional demand distributions. The initial (unconditional) distribution for a given period t corresponds to the demand forecast for that period considering the initial expected demand and is given in the form of a number K of equally likely quantile realizations $\{D_t^q\}_{q=1,\dots,K}$. Any expectation over the demand realizations is then computed using the approximation given by (9).

Conditional distributions given a demand state w_t are generated by reweighing the likelihood of each quantile. The reweighing is specified by the list of matrices $\{P_t\}_t$ discussed in Section 4.3. Each row of the matrix P_t corresponds to one of the states at time t and contains the weights of each quantile for that state, effectively operating a change of probability. Any expectation over the demand given w_t is then given by:

$$\mathbb{E}_{\varepsilon_t^w}[g(D_t)] = \sum_{q=1}^{K} P_t^w(q) g(D_t^q).$$

6.3 Optimization

6.3.1 Algorithm

Each state of the DP requires solving an optimization problem given by the optimality equation (2), for which we proved that the objective function is $L^\natural$-concave. An optimality criterion for $L^\natural$-concave functions is given by the following theorem from (Murota, 2003, Theorem 7.14).

Theorem 6.1
Let f be an $L^\natural$-concave function from $\mathbb{Z}^n$ to $\mathbb{R}$ and $\mathbf{x}^* \in \mathrm{dom}(f)$. Then, $\mathbf{x}^*$ is a global maximizer of f on $\mathbb{Z}^n$ if and only if:

$$f(\mathbf{x}^*) \geq f(\mathbf{x}^* \pm \mathbf{p}), \quad \forall \mathbf{p} \in \{0,1\}^n.$$

Figure 5 illustrates Theorem 6.1 by highlighting the points at which the function needs to be evaluated to ensure optimality of a point x^* for a two-dimensional problem.

It follows from the above that a simple optimization algorithm to solve the optimization problem is given by the greedy algorithm described in Algorithm 1, based on the steepest descent scaling algorithm of Murota (2003), and which is in essence a pattern-search algorithm.

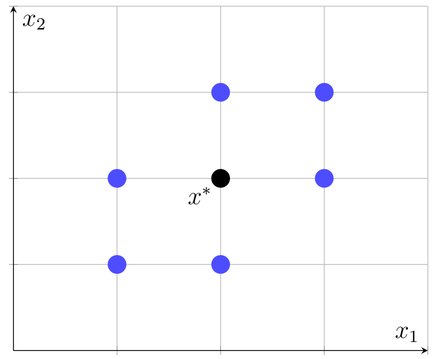

Figure 5: Illustration of the optimality condition for a two-dimensional $L^\natural$-concave function.

Algorithm 1 – Greedy algorithm

let $f(\mathbf{z})$ be an $L^\natural$-concave function on $\mathbb{R}^n$, and $\underline{\mathbf{z}}$ and $\bar{\mathbf{z}}$ represent the lower and upper bounds of the variables and z_i the i-th component of $\mathbf{z}$.

Step 0

let $\delta := \min\left(\max_i \frac{\bar{z}_i - \underline{z}_i}{2}, \bar{\delta}\right)$ and $\varepsilon > 0$, where $\bar{\delta} > 0$ is an upper bound on δ and ε a tolerance level.

let $\mathbf{z} = \frac{\bar{\mathbf{z}} + \underline{\mathbf{z}}}{2}$.

Step 1

while $\delta > \varepsilon$:
- **for each** $\mathbf{p} \in \{0,1\}^n$:
  - **for each** $s \in \{-1,1\}$:
    - **let** $\mathbf{z}' = \mathbf{z} + s\delta\mathbf{p}$,
    - **if** $f(\mathbf{z}') > f(\mathbf{z})$, **let** $\mathbf{z} \leftarrow \mathbf{z}'$ and go back to the first step in the loop.
- **let** $\delta \leftarrow \frac{\delta}{2}$ and go back to the first step in the loop.

This algorithm can be improved in many different ways such as favoring directions that have yielded progress first in the step direction selection.

Alternatively, recall that the dimensionality of the problem (2) could be reduced, resulting in (4). The trade-off for the removed dimension is the loss of $L^\natural$-concavity, resulting in mere concavity. Given that the function are not necessarily differentiable, the optimization of the resulting problem is not necessarily much easier.

6.3.2 Bounds

The above algorithm makes use of bounds on the variables. The $L^\natural$-concavity of the function being optimized plays an important role in reducing the search space of the optimization problem by allowing us to apply the bounds derived in Lemma 3.1.

7 Numerical Results

The models described in Section 3 and the extensions discussed in Section 4 have been implemented and we present in this section the results obtained on an illustrative example.

7.1 Example Data

We consider an example that runs over 40 periods at the beginning of which the purchasing, pricing and removal decisions are made. Note that scenarios in which the decisions do not have the same periodicity, say purchases can only be made every third period, are easily accommodated by setting the constraints appropriately.

7.1.1 Demand Model

We recall that the demand model used in this paper is an additive model of the form:

$$D_t(p) = d_t(p) + \varepsilon_t.$$

According to Assumption 1, we require from the demand model that the revenue function $dp_t(d)$ be concave. Common demand models such as linear, log-linear or isoelastic demand models satisfy this requirement. For our purposes, we choose a log-linear model, so that:

$$d_t(p) = \mu_t e^{-\beta(p-p_0)},$$

where μ_t is the expected forecasted demand at the base price p_0 (MRSP) in period t. β is a model parameter so that $\mathcal{E} = -\beta p_0$ represents the elasticity of the product at the base price.

Additionally, we need to specify the distribution of the noise ε_t. We assume that $\mu_t + \varepsilon_t$ is Gamma-distributed with mean μ_t and a coefficient of variation of 1.

7.1.2 Product Data and Assumptions

Parameters We consider an example product whose parameters are given in Table 1 here-under. We

p_0	c	s	l	h^+	k	$\mathcal{E}$
90	75	30	5	2	15.5	-2

Table 1: Parameters used for the examples.

recall that p_0 is the product's base price, c the cost, s the refund value, l the liquidation value, h^+ the holding cost per period, k the penalty for backordered sale and ε the elasticity of the product at the base price.

Mean Demand It remains to specify the forecasted expected demand μ_t in each period. We consider a constant expected demand of 50 per-period with a peak around the 20th period. A plot of the forecasted expected demand is presented in Figure 6, and we recall that we assume a Gamma distributed demand with coefficient of variation of 1.

Terminal Value We assume that at the end of the horizon, on-hand inventory is fully returned and liquidated.

Pricing Constraints We constrain the price of the product to always be less than or equal to its base price. Equivalently, the demand decision is always greater than or equal to the forecasted expected demand at the base price.

7.1.3 Dynamic Program Discretization

The periods can be thought of as weeks and we use a discount factor $\gamma = 0.9984$, which approximately corresponds to 8% per year.
The demand distributions are approximated through the use of 99 quantiles.

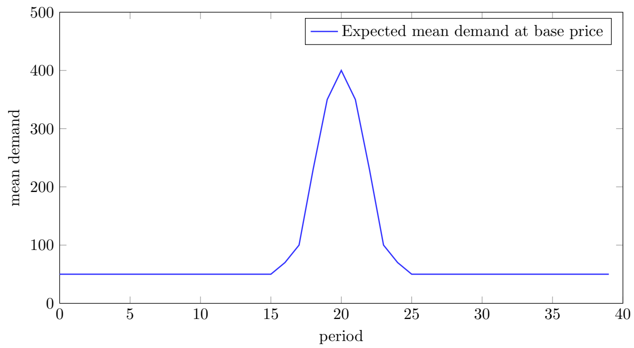

Figure 6: Mean demand at base price for the example product.

7.2 Vanilla Model

We present here results obtained with the basic models. Because the state spaces for the fixed and fractional returnability models are three-dimensional (inventory level, returnable level, and time), they are not well suited to being displayed on paper. We thus only present here partial results for the fixed returnability case, with more detailed and animated results being displayed in the electronic companion.

7.2.1 Fixed Returnability

We consider first the results obtained on the fixed returnability model, where we assume that the limit on the amount of returnable units that can be purchased in any period is given by the median of the forecasted demand at the base price. Figure 7 presents the results obtained at a fixed period $t = 24$ as four panels that describe the following:

- the top-left figure is a contour map of the NPV with inventory level (x) as the ordinate and the returnable level (y) as the abscissa,
- the top-right figure shows the evolution of the mean demand over time and the dot places the current period and mean demand on the figure,
- the bottom-left figure displays a contour map of the total purchasing and removal decision as a function of the state variable (recall that a negative value corresponds to a rebuy decision). The axes are as in the top-left figure the inventory level (x) as the ordinate and the returnable level (y) as the abscissa,
- the bottom-right figure displays a contour map of the markdown decision (expressed as a price cut on the base price p_0) as a function of the state variable.

We observe that the numerical results coincide with the theoretical properties derived in Section 3.2. In particular, we have the following:

- It can be checked that the NPV surface is $L^\natural$-concave,
- Purchased units decrease with inventory level, while returns and liquidations increase,
- Returns increase with the returnable level, while liquidations decrease,
- Markdowns increase with the inventory level and decrease with the returnable level.

The top left corner of the figure containing the purchasing and removal decisions shows some right angles in the iso-level curves. These correspond to the point at which liquidations start occurring. If we fix the returnable level, we observe that as the inventory level increases the iso-level curve runs

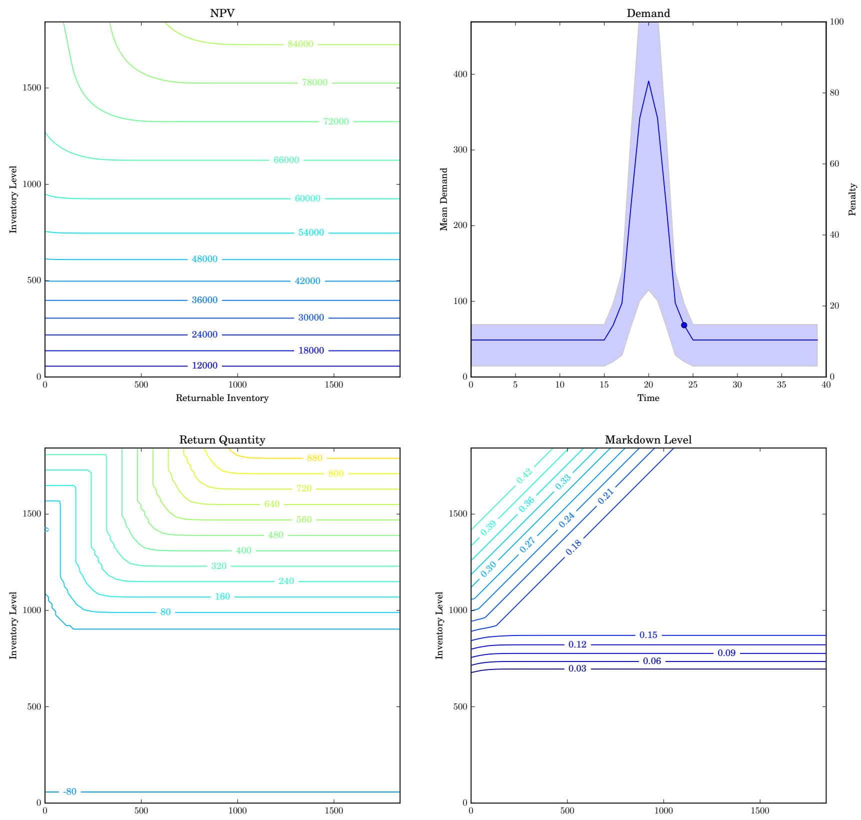

Figure 7: Summarizing graphs of the example presented in Section 7.2.1 at $t = 24$

along a vertical line, signifying that the entire returnable quota is being used, before hitting a point at which liquidations are triggered, resulting in this right angle.
The full animated results are presented in the electronic companion, where we also present results obtained with slightly different parameters (return value of $60 and holding cost of $1) that emphasize how buying decisions vary with inventory and returnable levels.

7.2.2 Full Returnability

When considering the full returnability case, the dimension of the results decreases to two, allowing us to present time-varying decisions. We showed in Section 3.4 that this model admitted order-up-to and remove-down-to levels that help us describe the optimal policies, which are presented for the considered example in Figure 8. We additionally plot the level at which markdowns start occurring. We recall that we applied the constraint that the price could only be marked down from the base price p_0.

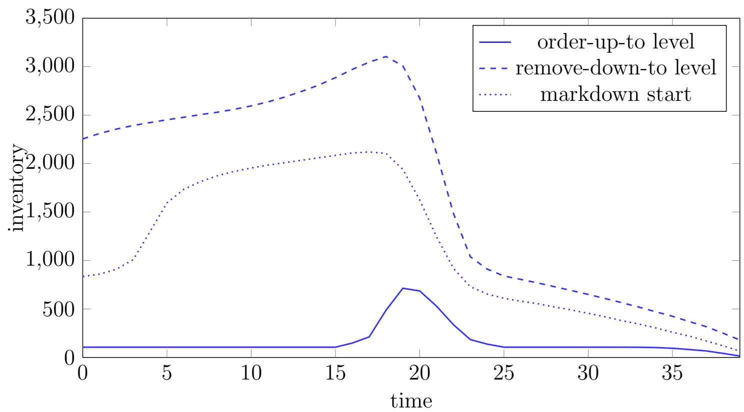

Figure 8: Order-up-to, remove-down-to, and markdown start levels for the considered example.

We observe that the remove-down-to and markdown levels increase as we approach the peak, before quickly decreasing as we near the end of the horizon and the value of higher inventory levels lowers. For a fixed time t, we may also illustrate the interval-stock list-prices policy by plotting the target inventory level and expected demand as a function of the starting inventory level. We plot these in Figure 9 for $t = 21$.

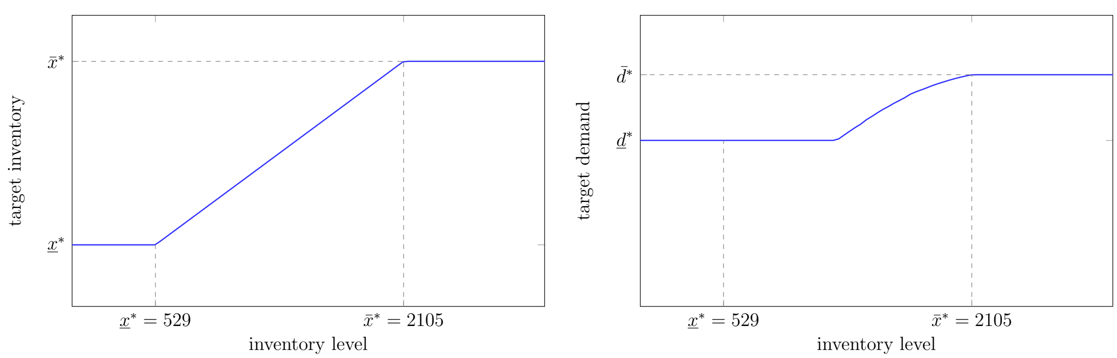

Figure 9: Target inventory level and expected demand in period $t = 21$ as a function of the starting inventory.

We easily identify on those figures the order-up-to and remove-down-to levels $\underline{x}^*$ and $\bar{x}^*$, as well as the list-price demands $\underline{d}^*$ and $\bar{d}^*$. We further observe that because of the imposed constraint that the price can only be marked down from the base price p_0, the expected demand does not start increasing at $\underline{x}^*$, but at some intermediate point between $\underline{x}^*$ and $\bar{x}^*$.

7.3 Capacity Constraints

We discussed in Section 5 the importance of a capacity constraint management system for large retailers. We proposed the handling of such constraints through their dualization, leading quite simply to increased holding costs during the constrained periods. We here illustrate the impact of such solution on the example in the fully returnable case. We may for example suppose that the demand peak experienced by the considered product is shared by many other ones, leading to capacity violations. Accordingly, this results in penalties being applied in the concerned weeks. Suppose that the product has a unit volume and that the capacity constraints lead to penalties $\lambda_{18} = 1$, $\lambda_{19} = 4.8$, $\lambda_{18} = 5.7$, $\lambda_{18} = 5$, $\lambda_{18} = 1$. Figure 10 compares the order-up-to, remove-down-to and markdown levels of the unconstrained scenario (already shown in Figure 8) with the newly constrained levels.

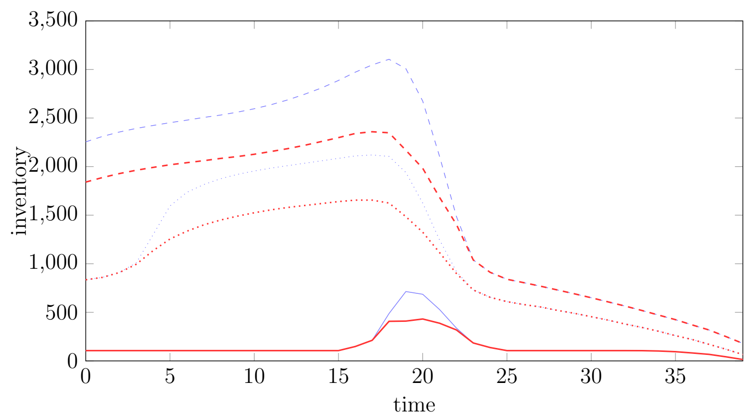

Figure 10: Comparison of constrained (thick) and unconstrained (thin) order-up-to (full), remove-down-to (dashed), and markdown start (dotted) levels.

We observe that the effect of the penalties is twofold:

1. They reduce the remove-down-to and markdown start levels, causing a more aggressive removal and markdown behavior even long before the constrained periods.

2. They curb the purchasing decisions in the affected periods.

These implications are in line with the anticipated and desired behavior. By returning and removing units that are unlikely to sell or would provide less profit than other products, we free up space that can be put to better use by more profitable products. Likewise, the reduced buying decisions help guarantee that we do not exceed the capacity allowance.

7.4 Correlated Demand

We now investigate the impact of demand correlation on the values and decisions of the problem. We detailed in Section 4.3 how such effects could be incorporated in the model, and in particular how we could use correlation information, as expanded upon in Appendix D. We thus consider the full returnability case and assume that demand correlation period over period is constant throughout the dynamic program horizon. We then run our algorithm for different values of the correlation, ranging from -1 to 1, and inspect some of the characteristic values of the problem, such as the order-up-to and remove-down-to levels, as well as the relative change in NPV for different levels of starting inventory. The NPV values are normalized by their value for a null correlation, corresponding to the default scenario and most common setting in inventory management problems. This helps us assess the degradation of the decisions and values caused by neglecting to take into account demand correlation. We plot in Figure 11 these normalized NPV curves as a function of correlation and observe that the relative errors in NPV tend to increase as the correlation increases in magnitude. These effects can grow quite large, especially as we approach the limiting values on both ends of the correlation spectrum. We note that while the shape of the normalized plots is here a “smile”, other examples have lead to downward smiles instead. It is thus difficult to fully characterize the effect of demand correlation, although experiments have shown that these can grow large in magnitude and a proper accounting of demand correlation is often crucial to a satisfactory management of popular and fast-selling products.

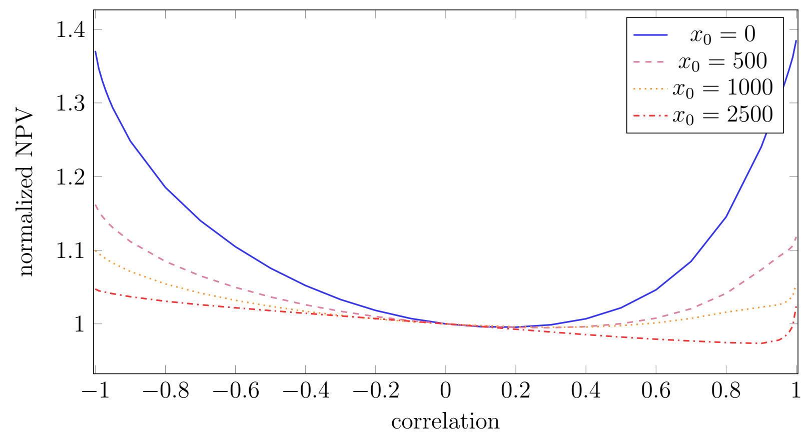

Figure 11: Impact of demand correlation on NPV at time 0 for different levels of starting inventory.

Figure 12 shows the evolution of the order-up-to and remove-down-to levels with demand correlation. We observe here again that they are heavily affected by correlation and can display wild swings when the correlation gets increasingly close to the extremities.

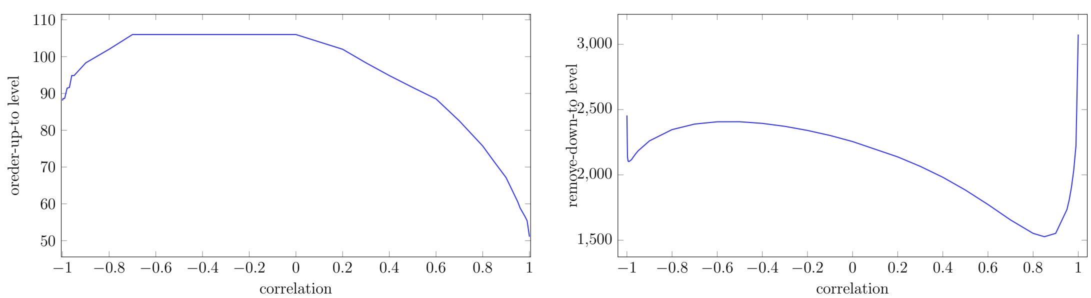

Figure 12: Impact of demand correlation on the order-up-to and remove-down-to levels.

7.5 Perishability

We mentioned in Section 4.2 that perishability could easily be incorporated in our framework, following the work of Chen et al. (2014b). We have thus implemented such a perishability model assuming for presentation purposes a shelf-life of 3 periods and full returnability rights. This allows to present the results as a function of the on-hand inventory with one remaining period of shelf-life and the on-hand inventory with two remaining periods of shelf-life. Given the perishability of the units, we decreased the salvage value to 10. Figure 13 presents the optimal values and decisions at a given period. We observe that for a fixed on-hand inventory level, the decisions can vary greatly depending on the makeup of the inventory, with order quantities higher for inventory composed of a majority of units in their last period of life. This can be explained by the fact that perishing units will bring down the inventory level in the next period, increasing the marginal value of the last bought unit compared to a similarly purchased unit in the absence of perishing units. This allows for more aggressive buys when large amounts of the on-hand inventory are about to perish. On the other hand, in order to extract as much value as possible from perishing units, markdowns are also more aggressive. We thus end up in a situation in which we are more aggressive both on our markdown and our purchasing behavior.

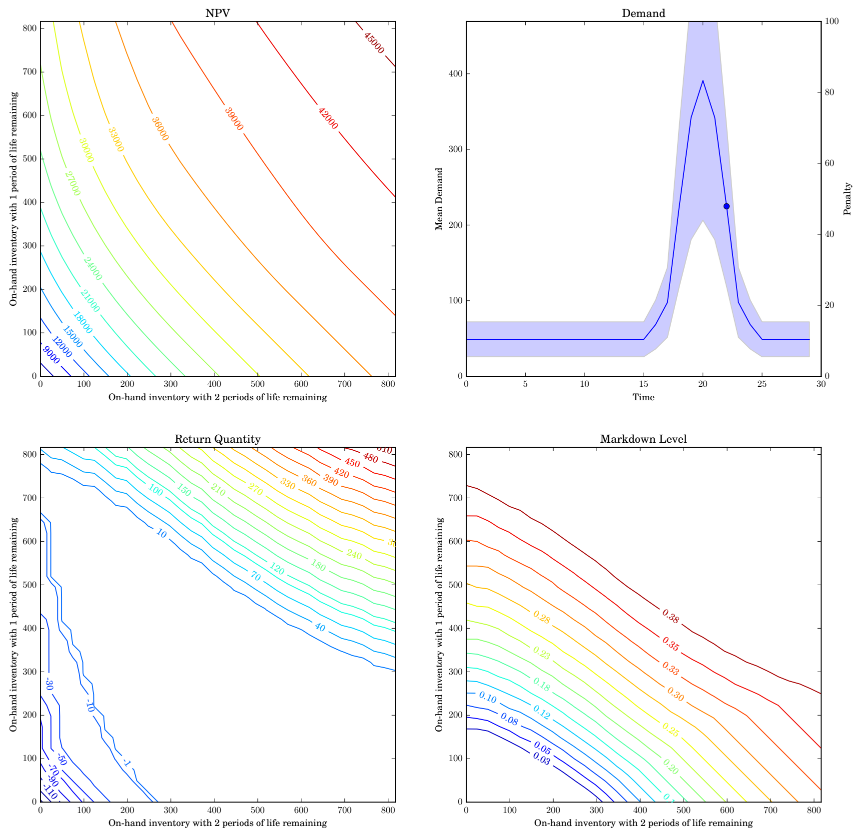

Figure 13: Illustration of the scenario with perishability.

7.6 Lead Times

The last illustration we present is that of the implementation of our framework in the context of an inventory management problem with fixed lead times. We mentioned in Section 4.1 that lead times can easily be implemented in our framework and how to do so. We first present some results for the traditional problem of lead times with lost sales to demonstrate the usefulness of $L^\natural$-concavity, not only in the derivation theoretical properties, but also in the implementation of the models; before presenting results in the context of our full model.

7.6.1 Lead Times with Lost Sales

When a full backlogging policy is assumed for excess demand, Karlin and Scarf (1958) show that the optimal policy in any period depends only on the total net inventory position so that the state space can be collapsed and reduced to a one-dimensional state. On the other hand, if a lost-sales policy is assumed, they show that no optimal policy is a function of the sum of the on-hand and pipeline inventory, preventing the same trick from applying. To circumvent the dimensionality issue, a number of bounds and heuristics have been developed over the years for the lost sales scenario. Zipkin (2008b) summarizes and compares many of these, which include a myopic policy Morton (1971), vector base-stock policies based on the bounds derived in Morton (1969) and a dual-balancing policy Levi et al. (2008). More recent heuristics include a quadratic approximation of cost functions Sun et al. (2014) and a modified myopic policy Brown and Smith (2014).

We illustrate the exact formulation of the problem in the case of a two-period lead time. Note that following Remark 1, we turned off the dynamic pricing in favor of a fixed pricing to guarantee the preservation of the structural results. We also turned off returns to present results for the more traditional setting. Figure 14 shows the resulting purchasing decisions as a function of on-hand and pipeline inventory.

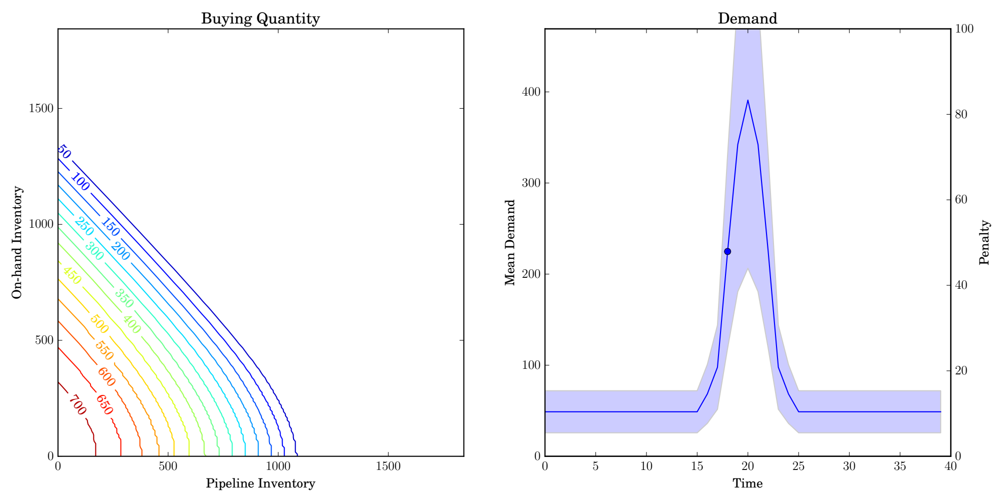

Figure 14: Buying decisions as a function of on-hand and pipeline inventory in the case of a two-period lead time with lost sales.

7.6.2 Full Model

We here return to the full model, albeit in the context of full returnability for plotting purposes. Figure 15 shows the optimal decisions and values at a given period of the dynamic program in the presence of a two-period lead time.

8 Conclusion

We presented in this paper theoretical and practical results pertaining to a joint replenishment, pricing and removal management problem. This problem was motivated by the challenges that a large, possibly online, retailer faces in their inventory planning, especially in the presence of capacity constraints and for whom capacity management during peak periods of the year is of paramount importance. The contribution of this paper is twofold: from a theoretical perspective, we extended previously studied models to consider removal channels, and derived structural properties thereof; while from a practical

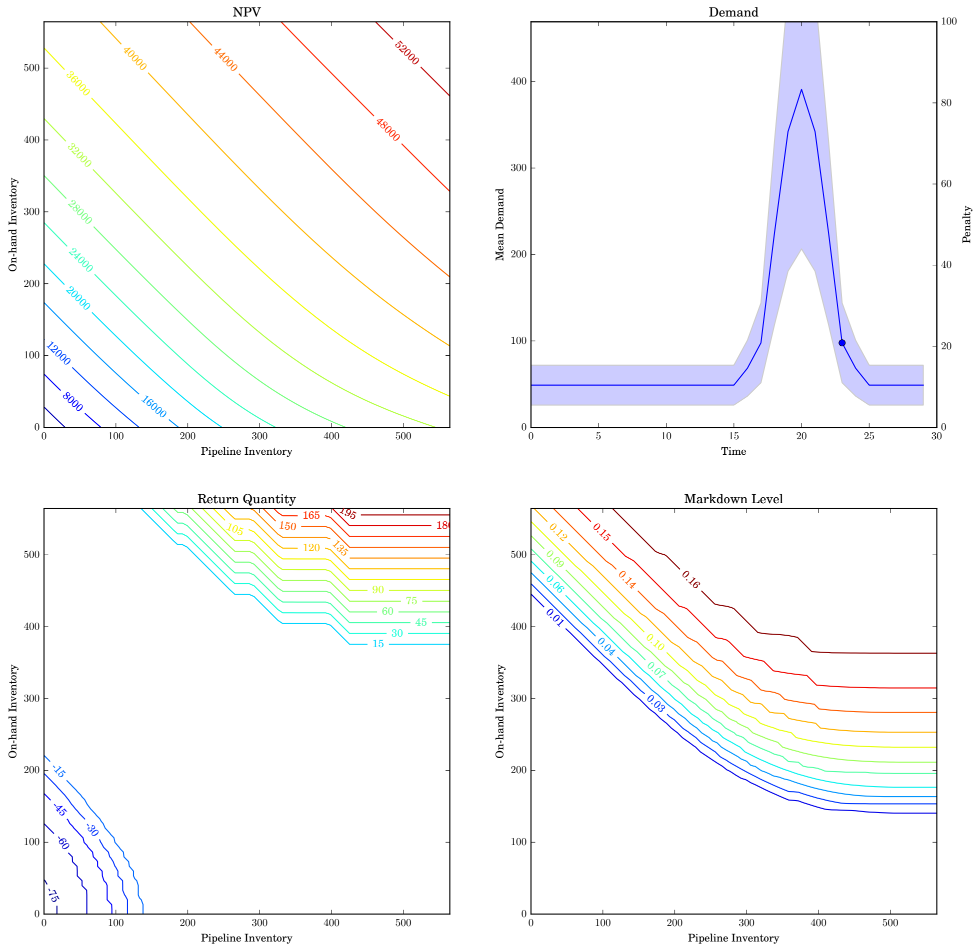

Figure 15: Results in the case of a two-period lead time.

point of view we detailed and illustrated the implementation of such models. A number of additional extensions to the baseline framework, such as lead times, demand correlation or perishability were also covered. A key aspect of the paper was underlying the usefulness, both theoretical and practical, of the concept of $L^\natural$-concavity that has garnered a lot of interest in the inventory management literature over the past few years.

The present paper and results signal a number of future research directions. We noted that until now, the literature dealing with vendor returns was focused on the single-period problem and the possibility of coordinating the supply chain. While we assumed in this paper that parameters were exogenous, it will be interesting to investigate the possibility of coordinating with vendors and jointly optimize costs, revenue refunds and limits; or simply evaluating the value of increasing these parameters to negotiate contracts that yield more value to the retailer. Similarly, the extension to multiple vendors that offer different lead time/return value/return limit trade-offs is of interest. Some details of the model are also to be improved, such as the expiration of return rights, which in spite of the fungibility of inventory across vendors, cannot be assumed to be valid for undetermined lengths of time.

Acknowledgments

We benefited from the help and input of a number of people whom we thank for their contributions. In particular, we would like to acknowledge the following people: Hernan Awad, Patrick Bajari, Deepak Bhatia, Mahmoud Bishr, Travis Brayak, Yazhe Feng, Ioannis Giannakakis, Ivan Hu, XiaoChang Hu, Salal Humair, Chao Liu, Jayanth Marasanapalle, Ahmed Saadany, Anna Sáez de Tejada Cuenca, Tolga Seyhan, Tao Shan, Jun Tong, Yuliang Qiao, Di Wu.

References

G. Allon and A. Zeevi. A Note on the Relationship Among Capacity, Pricing, and Inventory in a Make-to-Stock System. *Production and Operations Management*, 20(1):143–151, 2011. 3

A.-C. Barthel, T. Sabarwal, et al. Directional Monotone Comparative Statics. Technical report, University of Kansas, Department of Economics, 2015. 9, 36

D. Beyer, F. Cheng, S. P. Sethi, and M. Taksar. *Markovian Demand Inventory Models*. Springer, 2010. 11

D. B. Brown and J. E. Smith. Information Relaxations, Duality, and Convex Stochastic Dynamic Programs. *Operations Research*, 62(6):1394–1415, 2014. 27

G. P. Cachon. Supply Chain Coordination with Contracts. *Handbooks in operations research and management science*, 11:227–339, 2003. 2, 6

Y. Cai and K. L. Judd. Shape-Preserving Dynamic Programming. *Mathematical Methods of Operations Research*, 77(3):407–421, 2013. 15

M. C. Cario and B. L. Nelson. Modeling and Generating Random Vectors with Arbitrary Marginal Distributions and Correlation Matrix. Technical report, Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL, 1997. 42

L. M. Chan, Z. M. Shen, D. Simchi-Levi, and J. L. Swann. Coordination of Pricing and Inventory Decisions: A Survey and Classification. In *Handbook of Quantitative Supply Chain Analysis*, pages 335–392. Springer, 2004. 2

W. Chen, M. Dawande, and G. Janakiraman. Fixed-dimensional stochastic dynamic programs: An approximation scheme and an inventory application. *Operations Research*, 62(1):81–103, 2014a. 16, 17

X. Chen. $L^\natural$-Convexity and Its Applications in Operations. *Frontiers of Engineering Management*, forthcoming, 2017. 2

X. Chen and D. Simchi-Levi. Coordinating Inventory Control and Pricing Strategies with Random Demand and Fixed Ordering Cost: The Finite Horizon Case. *Operations Research*, 52(6):887–896, 2004a. 2, 3, 4

X. Chen and D. Simchi-Levi. 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, 2004b. 2

X. Chen and D. Simchi-Levi. Pricing and Inventory Management. In Ö. Özer and R. Phillips, editors, *The Oxford Handbook of Pricing Management*, chapter 30. Oxford University Press, United Kingdom, 11 2012. 2

X. Chen, P. Hu, and S. He. Preservation of Supermodularity in Two Dimensional Parametric Optimization Problems and its Applications. Technical report, UIUC, 2012. URL `http://publish.illinois.edu/xinchen/files/2012/12/ConvSubM.pdf`. 36

X. Chen, P. Hu, and S. He. Technical Note - Preservation of Supermodularity in Parametric Optimization Problems with Nonlattice Structures. *Operations Research*, 61(5):1166–1173, 2013. 9, 33, 36, 37, 38

X. Chen, Z. Pang, and L. Pan. Coordinating Inventory Control and Pricing Strategies for Perishable Products. *Operations Research*, 62(2):284–300, 2014b. 2, 3, 4, 11, 25

R. Ehrhardt. (s, S) Policies for a Dynamic Inventory Model with Stochastic Lead Times. *Operations Research*, 32(1):121–132, 1984. 11

W. Elmaghraby and P. Keskinocak. Dynamic Pricing in the Presence of Inventory Considerations: Research Overview, Current Practices, and Future Directions. *Management science*, 49(10):1287–1309, 2003. 2

H. Emmons and S. M. Gilbert. The Role of Returns Policies in Pricing and Inventory Decisions for Catalogue Goods. *Management science*, 44(2):276–283, 1998. 2

R. V. Evans. Inventory Control of a Multiproduct System with a Limited Production Resource. *Naval Research Logistics (NRL)*, 14(2):173–184, 1967. 13

A. Federgruen and A. Heching. Combined Pricing and Inventory Control Under Uncertainty. *Operations Research*, 47(3):454–475, 1999. 1, 2, 8, 9, 10, 40

Q. Feng. Integrating Dynamic Pricing and Replenishment Decisions Under Supply Capacity Uncertainty. *Management Science*, 56(12):2154–2172, 2010. 2, 3

A. V. Fiacco and Y. Ishizuka. Sensitivity and Stability Analysis for Nonlinear Programming. *Annals of Operations Research*, 27(1):215–235, 1990. 9

B. E. Fries. Optimal Ordering Policy for a Perishable Commodity with Fixed Lifetime. *Operations Research*, 23(1):46–61, 1975. 11

S. Fujishige. *Submodular Functions and Optimization*, volume 58. Elsevier, 2005. 33

B. V. Gnedenko. *The Theory of Probability and the Elements of Statistics*. American Mathematical Society, 5th edition, 2005. 14

X. Gong and X. Chao. Optimal Control Policy for Capacitated Inventory Systems with Remanufacturing. *Operations Research*, 61(3):603–611, 2013. 2

G. Hadley and T. M. Whitin. *Analysis of Inventory Systems*. Prentice Hall, 1963. 13

W. T. Huh and G. Janakiraman. (s, S) Optimality in Joint Inventory-pricing Control: An Alternate Approach. *Operations Research*, 56(3):783–790, 2008. 2

W. T. Huh and G. Janakiraman. On the Optimal Policy Structure in Serial Inventory Systems with Lost Sales. *Operations Research*, 58(2):486–491, 2010. 2

D. L. Inglehart and S. Karlin. Optimal Policy for Dynamic Inventory Process with Nonstationary Stochastic Demands. In K. J. Arrow, S. Karlin, and H. Scarf, editors, *Studies in Applied Probability and Management Science*, chapter 8, pages 127–147. Stanford University Press, Stanford, CA, 1962. 11

R. S. Kaplan. A Dynamic Inventory Model with Stochastic Lead Times. *Management Science*, 16(7): 491–507, 1970. 11

S. Karlin and C. R. Carr. Prices and Optimal Inventory Policy. In K. J. Arrow, S. Karlin, and H. Scarf, editors, *Studies in Applied Probability and Management Science*, chapter 10, pages 159–172. Stanford University Press, Stanford, CA, 1962. 2

S. Karlin and H. Scarf. Inventory Models of the Arrow-Harris-Marschak Type with Time Lag. In K. J. Arrow, S. Karlin, and H. Scarf, editors, *Studies in the Mathematical Theory of Inventory and Production*, chapter 10, pages 155–178. Stanford University Press, Stanford, CA, 1958. 11, 27

A. Kocabiyikoglu and I. Popescu. An Elasticity Approach to the Newsvendor with Price-Sensitive Demand. *Operations Research*, 59(2):301–312, 2011. 4

R. Levi, G. Janakiraman, and M. Nagarajan. A 2-Approximation Algorithm for Stochastic Inventory Control Models with Lost Sales. *Mathematics of Operations Research*, 33(2):351–374, 2008. 27

Q. Li and S. Zheng. Joint Inventory Replenishment and Pricing Control for Systems with Uncertain Yield and Demand. *Operations Research*, 54(4):696–705, 2006. 2

L. Lovász. Submodular functions and convexity. In *Mathematical Programming: The State of the Art*, pages 235–257. Springer, 1983. 16, 17

Y. Lu and J.-S. Song. Order-based Cost Optimization in Assemble-to-order Systems. *Operations Research*, 53(1):151–169, 2005. 2

H. Luss. Operations Research and Capacity Expansion Problems: A Survey. *Operations research*, 30 (5):907–947, 1982. 15

A. S. Manne. *Investments for Capacity Expansion: Size, Location, and Time-Phasing*, volume 5. Mit Press, 1967. 15

P. Milgrom and C. Shannon. Monotone Comparative Statics. *Econometrica: Journal of the Econometric Society*, pages 157–180, 1994. 9

E. S. Mills. Uncertainty and Price Theory. *The Quarterly Journal of Economics*, 73(1):116–130, 1959. 2

T. E. Morton. Bounds on the Solution of the Lagged Optimal Inventory Equation with no Demand Backlogging and Proportional Costs. *SIAM review*, 11(4):572–596, 1969. 27

T. E. Morton. The Near-Myopic Nature of the Lagged-Proportional-Cost Inventory Problem with Lost Sales. *Operations Research*, 19(7):1708–1716, 1971. 27

K. Murota. *Discrete Convex Analysis*. SIAM, 2003. 2, 16, 18, 33

S. Nahmias. Optimal Ordering Policies for Perishable Inventory – II. *Operations Research*, 23(4): 735–749, 1975. 11

S. Nahmias. Simple Approximations for a Variety of Dynamic Leadtime Lost-Sales Inventory Models. *Operations Research*, 27(5):904–924, 1979. 11

M. Ozlem. *Managing Supply Contracts and Inventory Risks in a Supply Chain*. PhD thesis, Stanford University, 2003. 3, 6

V. Padmanabhan and I. P. Png. Manufacturer's Return Policies and Retail Competition. *Marketing Science*, 16(1):81–94, 1997. 2

Z. Pang, F. Y. Chen, and Y. Feng. Technical Note - A Note on the Structure of Joint Inventory-Pricing Control with Leadtimes. *Operations Research*, 60(3):581–587, 2012. 2, 3, 4, 11

B. A. Pasternack. Optimal Pricing and Return Policies for Perishable Commodities. *Marketing Science*, 4(2):166–176, 1985. 2, 3

N. C. Petruzzi and M. Dada. Pricing and the Newsvendor Problem: A Review with Extensions. *Operations research*, 47(2):183–194, 1999. 2

J. K.-H. Quah. The Comparative Statics of Constrained Optimization Problems. *Econometrica*, 75 (2):401–431, 2007. 9, 36

G. Raz and E. L. Porteus. A Fractiles Perspective to the Joint Price/Quantity Newsvendor Model. *Management science*, 52(11):1764–1777, 2006. 12, 18

H. Scarf. The Optimality of (s, S) Policies in the Dynamic Inventory Problem. Technical Report 11, Stanford University, April 1959. 11

S. P. Sethi and F. Cheng. Optimality of (s, S) Policies in Inventory Models with Markovian Demand. *Operations Research*, 45(6):931–939, 1997. 11

A. Shapiro, D. Dentcheva, and A. Ruszczyanski. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM - Society for Industrial and Applied Mathematics, 2nd edition, 2014. 13

D. Simchi-Levi, X. Chen, and J. Bramel. *The Logic of Logistics: Theory, Algorithms, and Applications for Logistics Management*. Springer Science & Business Media, third edition, 2014. 9, 33

A. Sklar. Random Variables, Joint Distribution Functions, and Copulas. *Kybernetika*, 9(6):449–460, 1973. 41

S. A. Smith and N. Agrawal. Management of Multi-Item Retail Inventory Systems with Demand Substitution. *Operations Research*, 48(1):50–64, 2000. 13

J.-S. Song and P. Zipkin. Inventory Control in a Fluctuating Demand Environment. *Operations Research*, 41(2):351–370, 1993. 11

B. H. Strulovici and T. A. Weber. Generalized Monotonicity Analysis. *Economic theory*, 43(3):377–406, 2010. 9

P. Sun, K. Wang, and P. Zipkin. Quadratic Approximation of Cost Functions in Lost Sales and Perishable Inventory Control Problems. Technical report, Fuqua School of Business, Duke University, Durham, NC, 2014. 27

K. T. Talluri and G. J. Van Ryzin. *The Theory and Practice of Revenue Management*, volume 68. Springer Science & Business Media, 2006. 4

T. A. Taylor. Channel Coordination under Price Protection, Midlife Returns, and End-of-Life Returns in Dynamic Markets. *Management Science*, 47(9):1220–1234, 2001. 2

D. M. Topkis. *Supermodularity and Complementarity*. Princeton University Press, 2011. 9, 36

A. A. Tsay. Managing Retail Channel Overstock: Markdown Money and Return Policies. *Journal of Retailing*, 77(4):457–492, 2002. 2

T. M. Whitin. Inventory Control and Price Theory. *Management science*, 2(1):61–68, 1955. 2

Y. Yang, Y. Chen, and Y. Zhou. Coordinating Inventory Control and Pricing Strategies Under Batch Ordering. *Operations Research*, 62(1):25–37, 2014. 2, 3

E. Zabel. Monopoly and Uncertainty. *The Review of Economic Studies*, 37(2):205–219, 1970. 2

E. Zabel. Multiperiod Monopoly under Uncertainty. *Journal of Economic Theory*, 5(3):524–536, 1972. 2

P. Zipkin. On the Structure of Lost-Sales Inventory Models. *Operations Research*, 56(4):937–944, 2008a. 2, 11

P. Zipkin. Old and New Methods for Lost-Sales Inventory Systems. *Operations Research*, 56(5): 1256–1263, 2008b. 27

A Mathematical Preliminaries

We reproduce in this section some definitions and technical results related to supermodularity and $L^\natural$-concavity. These are mostly taken from (Murota, 2003, ch.5-7), (Simchi-Levi et al., 2014, ch.2) and (Fujishige, 2005, ch.16). We additionally amend some results presented in Chen et al. (2013) pertaining to the preservation of supermodularity with nonlattice structures.

A.1 Supermodularity and $L^\natural$-concavity

Definition A.1 – Supermodularity
Suppose X is a lattice in $\mathbb{R}^n$ and a function $f : X \to \mathbb{R}$. The function f is supermodular on the set X if, for any $\mathbf{x}, \mathbf{x}' \in X$,

$$f(\mathbf{x}) + f(\mathbf{x}') \leq f(\mathbf{x} \vee \mathbf{x}') + f(\mathbf{x} \wedge \mathbf{x}'). \tag{10}$$

($\vee$ and $\wedge$ denote the component-wise maximum and minimum, respectively.)

We next present a definition of a $L^\natural$-convex set.

Definition A.2 – $L^\natural$-convex set
A set V is an $L^\natural$-convex set if:

$$\mathbf{p}, \mathbf{q} \in V \implies (\mathbf{p} - \alpha\mathbf{e}) \vee \mathbf{q},\ \mathbf{p} \wedge (\mathbf{q} + \alpha\mathbf{e}) \in V \quad (\alpha \in \mathbb{R}_+),$$

where $\mathbf{e}$ is a vector of ones.

The following is a characterization of $L^\natural$-convex sets, an example of which is shown in Figure 16.

Proposition A.1
A set with a representation $\{\mathbf{x} \in \mathbb{R}^n : \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}, x_i - x_j \leq v_{ij}, \forall i \neq j\}$, is $L^\natural$-convex in the space $\mathbb{R}^n$, where $\mathbf{l}, \mathbf{u} \in \mathbb{R}^n$ and $v_{ij} \in \mathbb{R}(i \neq j)$. In fact, any closed $L^\natural$-convex set in the space $\mathbb{R}^n$ can have such a representation.

$L^\natural$-concave functions are defined as follows. (We recall that $\bar{\mathbb{R}} = \mathbb{R} \cup +\infty$.)

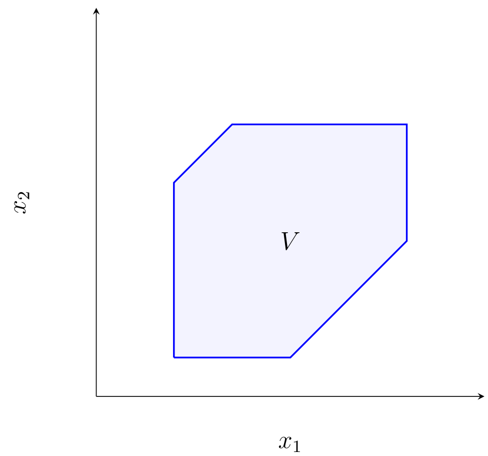

Figure 16: Example of a $L^\natural$-convex set V in $\mathbb{R}^2$.

Definition A.3 – $L^\natural$-concave function

A function $f : \mathbb{R}^n \to \bar{\mathbb{R}}$ is $L^\natural$-concave if for any $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^n$, $\alpha \in \mathbb{R}_+$:

$$f(\mathbf{x}) + f(\mathbf{x}') \leq f((\mathbf{x} + \alpha\mathbf{e}) \wedge \mathbf{x}') + f(\mathbf{x} \vee (\mathbf{x}' - \alpha\mathbf{e})).$$

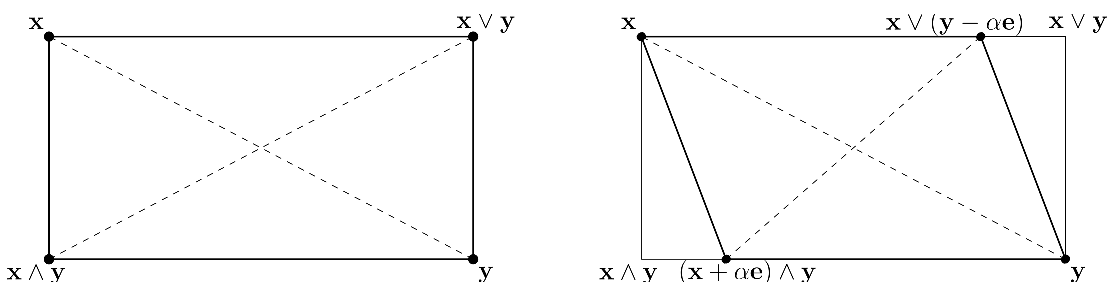

Figure 17: Comparison of supermodularity (left) and $L^\natural$-concavity (right).

Figure 17 compares the definitions of supermodularity and $L^\natural$-concavity. The following equivalent characterization of $L^\natural$-concavity is often easier to work with.

Proposition A.2

A function $f : \mathbb{R}^n \to \bar{\mathbb{R}}$ is $L^\natural$-concave if and only if $g(\mathbf{x}, \xi) := f(\mathbf{x} - \xi\mathbf{e})$ is supermodular on $(\mathbf{x}, \xi) \in \mathbb{R}^n \times \mathbb{R}$.

$L^\natural$-concave functions have many interesting properties. In particular they are concave, supermodular and the Hessian of a smooth $L^\natural$-concave function is diagonally dominant. Additionally, it can be shown that the restriction of a $L^\natural$-concave function to a $L^\natural$-convex set is also $L^\natural$-concave. An example of a quadratic $L^\natural$-concave function is presented in Figure 18.

$L^\natural$-concavity is preserved through a number of operations, the most relevant of which are reproduced below. We use $\mathbf{e}$ to denote a vector of ones of appropriate size.

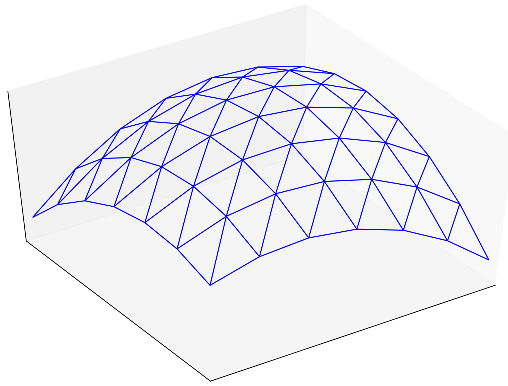

Figure 18: Example of a quadratic $L^\natural$-concave function

Proposition A.3

(a) Any nonnegative linear combination of $L^\natural$-concave functions is $L^\natural$-concave.

(b) Assume that a function $f(.,.)$ is defined on the product space $\mathbb{R}^n \times \mathbb{R}^m$. If $f(.,\mathbf{y})$ is $L^\natural$-concave for any given $\mathbf{y} \in \mathbb{R}^m$, then for a random vector $\boldsymbol{\xi}$ in $\mathbb{R}^m$, $\mathbb{E}_{\boldsymbol{\xi}}\left[f(\mathbf{x},\boldsymbol{\xi})\right]$ is $L^\natural$-concave, provided it is well-defined.

(c) If $f : \mathbb{R}^n \to \bar{\mathbb{R}}$ is an $L^\natural$-concave function, then $g : \mathbb{R}^n \times \mathbb{R} \to \bar{\mathbb{R}}$ defined by $g(\mathbf{x}, \xi) = f(\mathbf{x} - \xi\mathbf{e})$ is also $L^\natural$-concave.

(d) Assume that A is an $L^\natural$-convex set of $\mathbb{R}^n \times \mathbb{R}$ and $f(.,.) : \mathbb{R}^n \times \mathbb{R}^m \to \bar{\mathbb{R}}$ is an $L^\natural$-concave. Then the function

$$g(\mathbf{x}) = \sup_{\mathbf{y}:(\mathbf{x},\mathbf{y})\in A} f(\mathbf{x},\mathbf{y})$$

is also $L^\natural$-concave over $\mathbb{R}^n$ if $g(\mathbf{x}) \neq \infty$ for any $\mathbf{x} \in \mathbb{R}^n$.

(e) Given any univariate concave functions $g_i : \mathbb{R} \to \bar{\mathbb{R}}$ $(i = 1, \dots, n)$ and $h_{ij} : \mathbb{R} \to \bar{\mathbb{R}}$ $(i \neq j)$, the function f defined by:

$$f(\mathbf{x}) := \sum_{i=1}^{n} g_i(x_i) + \sum_{i \neq j} h_{ij}(x_i - x_j)$$

is $L^\natural$-concave. As a special case, any linear function is $L^\natural$-concave.

The following lemma is important to analyze the structural properties of problems that admit an $L^\natural$-concave objective function.

Lemma A.1
Let $g(\mathbf{x}, \xi) : \mathbb{R}^n \times \mathbb{R} \to \bar{\mathbb{R}}$ be $L^\natural$-concave, and let $\xi(\mathbf{x})$ be the smallest optimal optimal solution (assuming existence) of the optimization problem $f(\mathbf{x}) = \max_{\xi\in\mathbb{R}} g(\mathbf{x}, \xi)$ for any $\mathbf{x} \in dom(f)$. Then $\xi(\mathbf{x})$ is nondecreasing in $\mathbf{x} \in dom(f)$, and $\xi(\mathbf{x} + \omega\mathbf{e}) \leq \xi(\mathbf{x}) + \omega$ for any $\omega > 0$ with $\omega \in \mathbb{R}$ and $\xi(\mathbf{x}) + \omega \in dom(f)$.

A.2 Preservation of Supermodularity

The model with fractional returns considered in Section 3.3 fails to preserve $L^\natural$-concavity. Additionally, the sets on which the optimization problems are formulated are not even lattices, preventing the application of the standard results preserving supermodularity (Topkis, 2011, Sec. 2.7) and comparative statics (Topkis, 2011, Sec. 2.8). The preservation of supermodularity in two dimensional parametric optimization problems has been studied by Chen et al. (2013) where preservation of supermodularity is proved for a class of problems whose domain is not a lattice. The following is a restatement of the generalized version of their result with the corrected assumption that the matrix $B^T B$ be *inverse monotone* instead of $\mathcal{L}_0$, in addition to an assumption we believe is missing from the original paper. We recall that an $\mathcal{L}_0$-matrix is defined as a square matrix with nonnegative diagonal entries and nonpositive off-diagonal entries.

Proposition A.4
Let A be an $m \times n$ matrix and B be an $m \times 2$ matrix, such that $B^T B$ is an inverse monotone matrix and $B^T A \geq 0$ (where the inequality is understood element-wise). Let $\mathbf{D}$ be a closed convex sublattice of $\mathbb{R}^n$ and $\mathbf{S} = \{\mathbf{x} : A\mathbf{y} = B\mathbf{x} \text{ for some } \mathbf{y} \in \mathbf{D}\}$.
Further assume that $BB^+\mathbf{z} = \mathbf{z}$ for any $\mathbf{z} \in \mathbf{A}(\mathbf{D})$, where B^+ is the generalized inverse of B.
Consider the following optimization problem:

$$f(\mathbf{x}) = \max_y \{g(\mathbf{y}) : \ A\mathbf{y} = B\mathbf{x}, \ \mathbf{y} \in \mathbf{D}\}.$$

If g is concave and supermodular on $\mathbf{D}$, then so is f on $\mathbf{S}$.

Proof 5
The proof of the original result is missing from Chen et al. (2013), but present in the working paper Chen et al. (2012), to which we refer.
The proof contains the assumption that $B^T B$ is an $\mathcal{L}_0$-matrix, from which it is inferred that $(B^T B)^{-1} \geq 0$. That statement does not necessarily hold, though. To ensure that $(B^T B)^{-1} \geq 0$, we need to assume instead that $B^T B$ is inverse monotone.
Additionally, a step of the proof is as follows:

> For any $\mathbf{x}$, $\tilde{\mathbf{x}} \in \mathbf{S}$, let $\mathbf{y}$, $\tilde{\mathbf{y}}$ be the corresponding optimal solutions. Since $\mathbf{y} \wedge \tilde{\mathbf{y}}, \mathbf{y} \vee \tilde{\mathbf{y}} \in \mathbf{D}$, there exist $\mathbf{a}, \mathbf{b} \in \mathbf{S}$ such that $A(\mathbf{y} \wedge \tilde{\mathbf{y}}) = B\mathbf{a}$ and $A(\mathbf{y} \vee \tilde{\mathbf{y}}) = B\mathbf{b}$.

The existence of $\mathbf{a}$ and $\mathbf{b}$ as in the above quote does not, however, follow from $\mathbf{y} \wedge \tilde{\mathbf{y}}, \mathbf{y} \vee \tilde{\mathbf{y}} \in \mathbf{D}$. Letting $\mathbf{z}_a := A(\mathbf{y} \wedge \tilde{\mathbf{y}})$ and $\mathbf{z}_b := A(\mathbf{y} \vee \tilde{\mathbf{y}})$, in order to guarantee that there exist $\mathbf{a}$ and $\mathbf{b}$ such that $B\mathbf{a} = \mathbf{z}_a$ and $B\mathbf{b} = \mathbf{z}_b$, we need $BB^+\mathbf{z}_a = \mathbf{z}_a$ and $BB^+\mathbf{z}_b = \mathbf{z}_b$ instead. •

Additionally, we are interested in the monotonicity of the optimal set associated with the problem considered in Proposition A.4. The monotonicity of the optimal sets of problems not defined on lattices is investigated for concave and supermodular functions in Quah (2007); Barthel et al. (2015). Remark 2 in Chen et al. (2013) states that if the equality constraint is replaced by an inequality and B is the identity matrix, then the monotone statics results from Quah (2007) may be applied. Unfortunately, we present a counter-example in the next section (Appendix A.3) that disproves the claim, preventing us from applying the result.

A.3 Counterexample to (Chen et al., 2013, Remark 2)

The comparative statics results from Quah (2007) states that if C' dominates C in the $\mathcal{C}$-flexible order and g is a concave and supermodular function, then the set $\arg\max_{y \in C'} g(y)$ dominates the set $\arg\max_{y \in C} g(y)$ in the $\mathcal{C}$-flexible order. In particular this implies that for any $x \in \arg\max_{y \in C} g(y)$, there exists $x' \in \arg\max_{y \in C'} g(y)$ such that $x' \geq x$.

We recall that C' dominates C in the $\mathcal{C}_i$-flexible order if for any $x \in C$ and $x' \in C''$, with $x_i > x'_i$, there is some $\lambda \in [0,1]$ such that $x \wedge x' + \lambda v'_x$ is in C and $x \vee x' - \lambda v_x$ is in C', where $v_x := x \vee x' - x$. Figure 19 illustrates this relation.

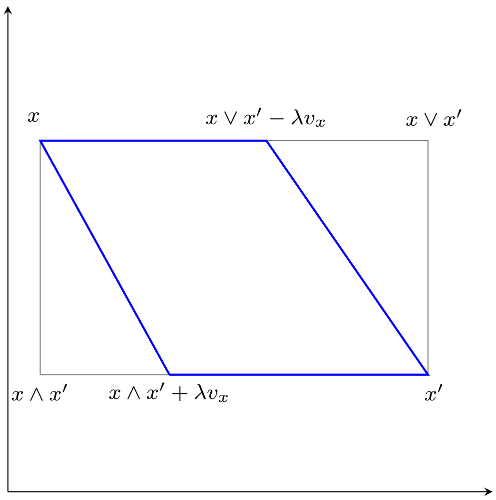

Figure 19: Illustration of $\mathcal{C}$-flexible order.

Consider now the optimization problem:

$$f(\mathbf{x}) = \max_{y} \{g(\mathbf{y}) : \; A\mathbf{y} \leq \mathbf{x}\}.$$

where $\mathbf{D}$ is a closed convex lattice and A is the following matrix:

$$A = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}.$$

Remark 2 in Chen et al. (2013) states that the above monotone statics result may be applied to this problem. However, letting $\mathbf{x}^1 = (a, b)$, $\mathbf{x}^2 = (a, b+1)$, $\mathbf{Y}_1 := \{\mathbf{y} : A\mathbf{y} \leq \mathbf{x}^1\}$ and $\mathbf{Y}_2 := \{\mathbf{y} : A\mathbf{y} \leq \mathbf{x}^2\}$, we observe that $\mathbf{x}^2 \geq \mathbf{x}^1$ but $\mathbf{Y}_2$ does not dominate $\mathbf{Y}_1$ in the $\mathcal{C}$-flexible order. Figure 20 illustrates how this relation fails to hold. By choosing points $\mathbf{y} \in \mathbf{Y}_1$ such that $y_1 + y_2 = a$ and $\mathbf{y}' \in \mathbf{Y}_2$ such that $y_1' + y_2' = a$ and $y_2' > b$, we observe that the only value of $\lambda \in [0, 1]$ that yields $\mathbf{y} \vee \mathbf{y}' - \lambda \mathbf{v}_y \in \mathbf{Y}_2$ is 1, but that leads to $\mathbf{y} \wedge \mathbf{y}' + \lambda \mathbf{v}_y = \mathbf{y}' \notin \mathbf{Y}_1$.

The failure of the monotonicity of the optimal sets can be further tested by letting $g(\mathbf{y}) := -(y_1 - 10)^2 - (y_2 - 5)^2$. For $\mathbf{x}^1 = (10, 5)$, it is easy to see that the optimal solution is given by $\mathbf{y}^1 = (5, 5)$, while for $\mathbf{x}^2 = (10, 7.5)$ the optimal solution is given by $\mathbf{y}^2 = (7.5, 2.5)$. In this example, we have $\mathbf{x}^2 \geq \mathbf{x}^1$ but not $\mathbf{y}^2 \geq \mathbf{y}^1$.

B Proof of Lemma 3.1

Let:

$$\begin{aligned}
&\mathbf{z} := (x, y), \\
&(x_+(\mathbf{z}), x_{++}(\mathbf{z}), q_r(\mathbf{z})) := \text{the least optimal solution of the optimality equation (2)}, \\
&{q_r}^*(\mathbf{z}, x_+, x_{++}) := \operatorname*{argmax}_{q_r : \underline{q}^r \leq q_r \leq y} f_t(\mathbf{z}, x_+, x_{++}, q_r) \text{ when } (\mathbf{z}, x_+, x_{++}) \text{ is given}, \\
&x_{++}^*(\mathbf{z}, x_+) := \operatorname*{argmax}_{x_{++} : \underline{d} \leq x_+ - x_{++} \leq \bar{d}} f_t(x, y, x_+, x_{++}, {q_r}^*(\mathbf{z}, x_+, x_{++})) \text{ when } (\mathbf{z}, x_+) \text{ is given}, \\
&x_+^*(\mathbf{z}) := \operatorname*{argmax}_{x_+ : \max(0, \underline{q} - \underline{q}^r) \leq x - x_+, 0 \leq x_+} f_t(x, y, x_+, x_{++}^*(\mathbf{z}, x_+), {q_r}^*(\mathbf{z}, x_+, x_{++}^*(\mathbf{z}, x_+))) \text{ when } \mathbf{z} \text{ is given}.
\end{aligned}$$

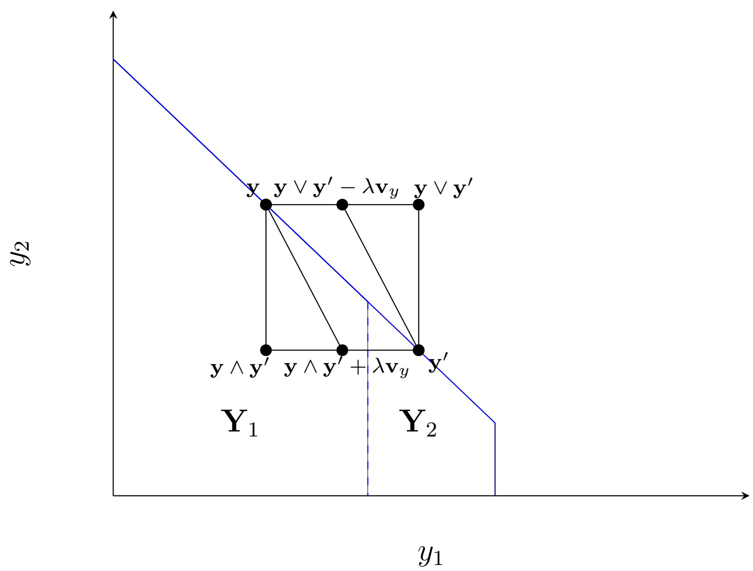

Figure 20: Illustration of the counterexample to (Chen et al., 2013, Remark 2).

Each of the above assignments should be understood as the least element of the set when it greater than a singleton.
In other words, $(x_+(\mathbf{z}), x_{++}(\mathbf{z}), q_r(\mathbf{z}))$ is the optimal policy in state $\mathbf{z} = (x, y)$ (and time t, which we omit to keep the notation light). Observe then that we also have $q_r(\mathbf{z}) = {q_r}^*(\mathbf{z}, x_+(\mathbf{z}), x_{++}(\mathbf{z}))$, $x_{++}(\mathbf{z}) = x^*_{++}(\mathbf{z}, x_+(\mathbf{z}))$, and $x_+(\mathbf{z}) = x^*_+(\mathbf{z})$, so that the least element of the set of optimal solutions can be obtained by solving an equivalent sequential decision optimization problem giving the following reformulation:

$$\begin{aligned}
V_t(\mathbf{z}) &= \max_{x_+:\max(0,\underline{q}-\underline{q}^r)\le x-x_+,0\le x_+} \psi_t(\mathbf{z}, x_+), \\
\text{where } \psi_t(\mathbf{z}, x_+) &= \max_{x_{++}:\underline{d}\le x_+-x_{++}\le\bar{d}} \phi_t(\mathbf{z}, x_+, x_{++}), \\
\phi_t(\mathbf{z}, x_+, x_{++}) &= \max_{q_r:\underline{q}^r\le q_r\le y} f_t(\mathbf{z}, x_+, x_{++}, q_r).
\end{aligned}$$

We then derive the following monotonicity properties of the optimal policies.

Lemma B.1
For any $t = 0, \ldots, T-1$ and any $\omega > 0$, the following results hold:

(a) $x_+(\mathbf{z})$ is nondecreasing in $\mathbf{z}$, and $x_+(\mathbf{z}) \le x_+(\mathbf{z}+\omega\mathbf{e}) \le x_+(\mathbf{z}) + \omega$,

(b) $x_{++}(\mathbf{z})$ is nondecreasing in $\mathbf{z}$, and $x_{++}(\mathbf{z}) \le x_{++}(\mathbf{z}+\omega\mathbf{e}) \le x_{++}(\mathbf{z}) + \omega$,

(c) $q_r(\mathbf{z})$ is nondecreasing in $\mathbf{z}$, and $q_r(\mathbf{z}) \le q_r(\mathbf{z}+\omega\mathbf{e}) \le q_r(\mathbf{z}) + \omega$.

Proof 6
(a) results from a direct application of Lemma A.1.
Consider now (b): applying Lemma A.1 implies that $x^*_{++}(\mathbf{z}, x_+)$ is non-decreasing in $(\mathbf{z}, x_+)$ and that

$$x^*_{++}(\mathbf{z}+\omega\mathbf{e}, x_+ + \omega) \le x^*_{++}(\mathbf{z}, x_+) + \omega. \tag{11}$$

Consequently, we have the following string of inequalities:

$$\begin{aligned} x_{++}(\mathbf{z}) &= x^*_{++}(\mathbf{z}, x_+(\mathbf{z})) \\ &\leq x^*_{++}(\mathbf{z}+\omega\mathbf{e}, x_+(\mathbf{z}+\omega\mathbf{e})) \\ &\leq x^*_{++}(\mathbf{z}+\omega\mathbf{e}, x_+(\mathbf{z})+\omega) \\ &\leq x^*_{++}(\mathbf{z}, x_+(\mathbf{z}))+\omega \\ &= x_{++}(\mathbf{z})+\omega, \end{aligned}$$

where the first equality hods by definition of $x_{++}(\mathbf{z})$, the first inequality from the non-decreasing properties of $x_+(\mathbf{z})$ and $x^*_{++}(\mathbf{z}, x_+)$, the second inequality from (a) and the non-decreasing property of x^*_{++}, the third inequality from (11), and the final equality by definition. Note that in the second line, $x^*_{++}(\mathbf{z}+\omega\mathbf{e}, x_+(\mathbf{z}+\omega\mathbf{e})) = x_{++}(\mathbf{z}+\omega\mathbf{e})$.
The proof for (c) is similar. •

We may then restate the results of Lemma B.1 in terms of the variables of interest, namely d, q_r and q_{nr} to obtain the results presented in Lemma 3.1.

Proof 7
Proof of Lemma 3.1. (b) corresponds to Lemma B.1 (c).
We now prove the first string of inequalities in (c).
Let $\omega > 0$, we have from Lemma B.1 (a) the following inequalities:

$$\begin{aligned} & x_+(x,y) \leq x_+(x+\omega, y) \leq x_+(x,y)+\omega \\ \Leftrightarrow\quad & x - q_{nr}(x,y) \leq x+\omega - q_{nr}(x+\omega,y) \leq x - q_{nr}(x,y)+\omega \\ \Leftrightarrow\quad & q_{nr}+\omega \geq q_{nr}(x+\omega,y) \geq q_{nr}(x,y) \end{aligned}$$

which proves the first string of inequalities in (c) and that q_{nr} is nondecreasing in x.
All other inequalities are proved similarly. •

C Proof of Theorem 3.4

Consider the optimization problem defined in (7) and its associated Lagrangian $\mathcal{L}(\tilde{x}, x_+, x_{++}, \boldsymbol{\lambda})$, with $\boldsymbol{\lambda} \leq \mathbf{0}$, where λ_1 is associated with the constraint $\tilde{x} \geq x$, λ_2 with $\tilde{x} \geq x_+$, λ_3 with $x_+ - x_{++} \geq \underline{d}$ and λ_4 with $x_+ - x_{++} \leq \bar{d}$. The first-order necessary conditions require that $\nabla\mathcal{L}(\tilde{x}, x_+, x_{++}, \boldsymbol{\lambda}) = \mathbf{0}$, and thus in particular that:

$$\begin{aligned} (s-c_t) - \lambda_1 - \lambda_2 &= 0, \\ r'_t(x_+ - x_{++}) - s + \lambda_2 - \lambda_3 + \lambda_4 &= 0, \\ -r'_t(x_+ - x_{++}) + w'_t(x_{++}) + \lambda_3 - \lambda_4 &= 0 \end{aligned}$$

We branch out on the possible combinations of λ_1 and λ_2:

$\lambda_1 = \lambda_2 = 0$**:** This would require $s = c_t$, which contradicts Assumption 2 (ii) that $s < c_t$.

$\lambda_1 = 0$ **and** $\lambda_2 > 0$**:** $\lambda_2 > 0$ implies by the complementary conditions that $x_+ = \tilde{x} =: \underline{x}^*$ and thus that $q_r = 0$. Substituting in the first-order conditions above, we find that the optimal x_+ and x_{++} should satisfy $w'_t(x_{++}) = c_t$, and $r'_t(x_+ - x_{++}) = c_t + \lambda_3 - \lambda_4$. In particular, since the objective function is concave in x_+, it follows that the optimal demand is given by $\underline{d}^* := \min\left(\max(r'^{-1}_t(c_t), \underline{d}), \bar{d}\right)$.

$\lambda_1 > 0$ **and** $\lambda_2 = 0$**:** $\lambda_1 > 0$ implies by the complementary conditions that $\tilde{x} = x$ and thus that $q = 0$. Substituting in the first-order conditions above, we find that the optimal $x_+ =: \bar{x}^*$ and x_{++} should satisfy $w'_t(x_{++}) = s$, and $r'_t(x_+ - x_{++}) = s + \lambda_3 - \lambda_4$. In particular, since the objective function is concave in x_+, it follows that the optimal demand is given by $\bar{d}^* := \min\left(\max(r'^{-1}_t(s), \underline{d}), \bar{d}\right)$.

$\lambda_1 > 0$ **and** $\lambda_2 > 0$**:** The complementary conditions imply here that $x = \tilde{x} = x_+$, so that $q = q_r = 0$, and x_{++} should satisfy $w'_t(x_{++}) = s - \lambda_2$.

Assume temporarily that all the inverse mappings $w_t'^{-1}(c_t)$, $r_t'^{-1}(c_t)$, $w_t'^{-1}(s)$ and $r_t'^{-1}(s)$ exist. Following the above analysis, for any value of x, there exist three points satisfying the KKT conditions, corresponding to the following values of q^*, q_r^* and d^*:

- $(\underline{x}^* - x, 0, \underline{d}^*)$,
- $(0, x - \bar{x}^*, 0, \bar{d}^*)$,
- $(0, 0, \tilde{d})$, for some $\tilde{d} \in [\underline{d}^*, \bar{d}^*]$.

Note that the first solution corresponds to the solution of the problem where the constraints $\tilde{x} \geq x$ and $x_+ \leq \tilde{x}$ are omitted. Similarly, the second solution is obtained by omitting only the constraint $x_+ \leq \tilde{x}$. Consequently, so long as feasibility allows, the first solution is preferred to the second, which in turn is preferred to the third. This allows us to synthesize the results as follows:

$$(q^*, q_r^*, d^*) = \begin{cases} (\underline{x}^* - x, 0, \underline{d}^*) & \text{if } x < \underline{x}^*, \\ (0, 0, \tilde{d}), \text{ for some } \tilde{d} \in [\underline{d}^*, \bar{d}^*] & \text{if } \underline{x}^* \leq x \leq \bar{x}^*, \\ (0, x - \bar{x}^*, \bar{d}^*) & \text{if } x > \underline{x}^*, \end{cases}$$

This proves the result.

Finally, some minor assumptions on the holding cost functions and terminal values can be imposed to guarantee that the inverse mappings do actually exist, in a manner similar to Federgruen and Heching (1999). Nevertheless, the structure of the results can still be made to hold by allowing the threshold values to take on infinite values.

D Generation of Probability Matrices for the Correlated Demand Model

To implement the model with correlated demand described in Section 4.3, we need to provide or compute the probability matrices that characterize the evolution of the demand through time. The implementation models the stochastic demand through a number of quantiles, which are then mapped to states. At one end of the demand state modeling is a model with as many states as quantiles, one state for each quantile, while at the other end all the states are collapsed into a single state, resulting in a independent demand model. The transition probability matrices inform us in each time period t on the probability of each quantile realization given that we are in a demand state w_t.

To make things more explicit, consider the case where we have as many states as quantiles and we are using 4 quantiles to represent the demand. The four states correspond to the following events:

$$S_t^1 = \left\{-\infty \leq D_t \leq d_t^{25}\right\}, \quad S_t^2 = \left\{d_t^{25} < D_t \leq d_t^{50}\right\}, \quad S_t^3 = \left\{d_t^{50} \leq D_t \leq d_t^{75}\right\}, \quad S_t^4 = \left\{d_t^{75} \leq D_t \leq +\infty\right\},$$

where d_t^q is the q-th quantile of D_t. The transition probability matrix P_t reads:

$$P_t = \begin{array}{c} \\ S_t^1 \\ S_t^2 \\ S_t^3 \\ S_t^4 \end{array} \begin{array}{c} \begin{array}{cccc} S_{t+1}^1 & S_{t+1}^2 & S_{t+1}^3 & S_{t+1}^4 \end{array} \\ \begin{bmatrix} p_{11} & p_{12} & p_{13} & p_{14} \\ p_{21} & p_{22} & p_{23} & p_{24} \\ p_{31} & p_{32} & p_{33} & p_{34} \\ p_{41} & p_{42} & p_{43} & p_{44} \end{bmatrix} \end{array},$$

where $p_{ij} = \mathbb{P}\left[S_{t+1}^j | S_t^i\right] = \frac{\mathbb{P}\left[S_{t+1}^j, S_t^i\right]}{\mathbb{P}[S_t^i]}$.

In order to derive the elements p_{ij}, we would need access to the joint distribution of D_t and D_{t+1}. While we know the marginal distributions corresponding to D_t and D_{t+1}, as well as the cumulative distribution $D_{t,t+1}$, characterizing the joint distribution is quite challenging and not necessarily well-defined. Consequently, some gaps need to be filled and additional structure inherited from some other object, which can be achieved through the use of copulas.

Theorem D.1 – Sklar's Theorem Sklar (1973)
For $n \geq 2$, let F be an n-dimensional distribution function with marginal distributions $F_1, \ldots, F_n$. Then there exists an n-copula C such that

$$F(x_1, \ldots, x_n) = C(F_1(x_1), ..., F_n(x_n)) \tag{12}$$

for all n-tuples $(x_1, \ldots, x_n)$ in $\mathbb{R}^n$. Conversely, let C be an n-copula and $(F_1, \ldots, F_n)$ an n-tuple of 1-dimensional distribution functions. Let an n-place real function F be defined via (12). Then F is an n-dimensional distribution function with 1-margins $F_1, \ldots, F_n$.

Sklar's Theorem relates the joint distribution function with a copula. It states that a joint distribution function together with the marginal distributions defines a copula (unique when the distributions are continuous). Conversely, given a copula and the marginals, we define a cumulative distribution function. The latter case is of interest to us since we do not possess the joint distribution function of D_t and D_{t+1}. Instead, we know only their marginal distributions, together with an idea of their correlation. We may then consider using a copula that uses the correlation information, which together with the marginal distribution functions will help us approximate the unknown joint distribution.

One of the most common, and certainly easier to implement, copulas is the Gaussian copula C_R^g for a given corelation matrix R. This also plays nicely with the amount of information at our disposal, namely the correlation between the demand distributions. The Gaussian copula is given by:

$$C_R^g(u_1, \ldots, u_n) = \Phi_R\left(\Phi^{-1}(u_1), \ldots, \Phi^{-1}(u_n)\right), \tag{13}$$

where Φ_R is the cumulative distribution function of the multivariate Gaussian distribution with correlation matrix R and Φ^{-1} is the inverse cumulative distribution of the univariate standard normal.

It follows from (13) that the cumulative distribution $F_{t,t+1}$ that results from using the Gaussian copula in conjunction with the marginal distributions F_t and F_{t+1} of D_t and D_{t+1}, respectively, yields the following equality:

$$F_{t,t+1}(d_t^{q_1}, d_{t+1}^{q_2}) = \Phi_\rho\left(n^{q_1}, n^{q_2}\right),$$

where d_t^q is the q-th quantile of D_t and n^q the q-th quantile of a standard normal distribution. Grossly and abusively stated, this implies that the probability of observing the q_2-th quantile of D_{t+1} given that we observed the q_1-th quantile of D_t is the same as observing the q_2-th quantile of a standard normal given that we observed the q_1-th quantile of a standard normal with which it has a correlation of ρ. The implication of this observation is that we will inherit the transition probability matrix from the standard normal case. Thus, we have:

$$\mathbb{P}[d_{t+1}^{q_1} \leq D_{t+1} \leq d_{t+1}^{q_2} | D_t] = \mathbb{P}[n^{q_1} \leq N_2 \leq n^{q_2} | N_1],$$

where N_1 and N_2 are two standard normal distributions with a correlation of ρ.

Example D.1
Consider the scenario depicted in Example 4.1 where the demand is modeled through 4 quantiles and there are as many demand states as quantiles. Further assume that the correlation between consecutive demand realizations is 0.5. The transition probability matrix is then given by:

$$P_t = \begin{array}{c} \\ S_t^1 \\ S_t^2 \\ S_t^3 \\ S_t^4 \end{array} \begin{array}{c} \begin{array}{cccc} S_{t+1}^1 & S_{t+1}^2 & S_{t+1}^3 & S_{t+1}^4 \end{array} \\ \begin{bmatrix} 0.48 & 0.28 & 0.17 & 0.07 \\ 0.28 & 0.30 & 0.26 & 0.17 \\ 0.17 & 0.26 & 0.30 & 0.28 \\ 0.07 & 0.17 & 0.28 & 0.48 \end{bmatrix} \end{array},$$

If instead, we only considered two demand states, say low and high, corresponding to quantiles below and above the 50th, respectively, but still four quantiles; we would only need collapsing the corresponding rows and reweighing appropriately:

$$P_t = \begin{array}{c} \\ S_t^1 \\ S_t^2 \end{array} \begin{array}{c} \begin{array}{cccc} S_{t+1}^1 & S_{t+1}^2 & S_{t+1}^3 & S_{t+1}^4 \end{array} \\ \left[\begin{array}{cccc} 0.38 & 0.29 & 0.21 & 0.12 \\ 0.12 & 0.21 & 0.29 & 0.38 \end{array}\right] \end{array}.$$

Remark 4
The implied correlation between D_t and D_{t+1} obtained after the use of the Gaussian copula C_ρ^g will, in most cases, not be equal to ρ. An exact process requires finding the right correlation to be used for the copula in order to attain the desired correlation. A reference on this process can be found in Cario and Nelson (1997).

E Lovász Extension

We describe in this section the algorithm to produce the Lovász extension that is used to interpolate the value functions in the implementation of our dynamic programs.

Algorithm 2 – Lovász Extension
Let f be a function defined on $\delta\mathbb{Z}^n$, where $\delta > 0$ is the step size of the grid on which f is defined. Let $\mathbf{x} \in \mathbb{R}^n$ at which we want to extend f.

Step 0

let

$$\underline{\mathbf{x}} := \delta \left\lfloor \frac{\mathbf{x}}{\delta} \right\rfloor, \qquad \bar{\mathbf{x}} := \delta \left\lceil \frac{\mathbf{x}}{\delta} \right\rceil, \qquad \mathbf{w} := \frac{\mathbf{x} - \underline{\mathbf{x}}}{\delta},$$

where operations are defined element-wise. (Note that $\underline{\mathbf{x}}$ and $\bar{\mathbf{x}}$ are grid points.)
let σ be a permutation on $\{1, \ldots, n\}$ that orders $\{w_1, \ldots, w_n\}$, the elements of $\mathbf{w}$, in decreasing order.

Step 1

let $\tilde{\mathbf{x}} := \underline{\mathbf{x}}$
let $v := (1 - w_{\sigma_1}) f(\tilde{\mathbf{x}})$
for $i = 1, \ldots, n$:
  $\tilde{x}_{\sigma(i)} = \bar{x}_{\sigma(i)}$
  if $i < n$:
    $\tilde{w} = w_{\sigma(i)} - w_{\sigma(i+1)}$
  else if $i = n$:
    $\tilde{w} = w_{\sigma(i)}$
  $v = v + \tilde{w} f(\tilde{\mathbf{x}})$,
return v

F Example Results

We present in this section some more detailed results corresponding to the example presented in Section 7. Note that in order for the animations to be displayed properly, the document needs to be opened with Adobe Acrobat Reader.

F.1 Fixed Returnability

F.1.1 Default Scenario

We present in Figure 21 the results obtained with the parameters described in Table 1.

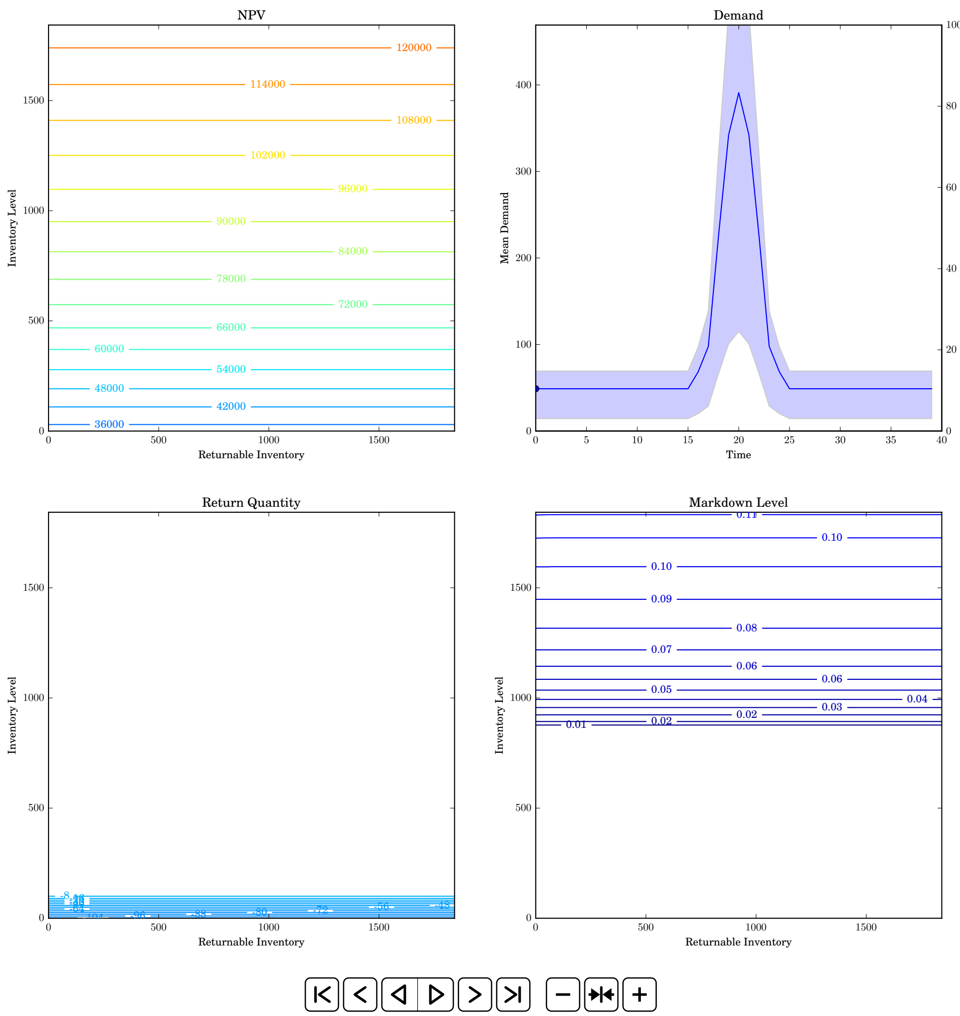

Figure 21: Summarizing graphs of the example in the fixed returnability case.

F.1.2 Modified Scenario

The animation in Figure 21 makes it seem as though there is an order up to level independent of the inventory and returnable level. To illustrate how this is not the case, we change the parameters so that the return value s becomes \$60 and the holding cost h \$1. Figure 22 shows the resulting behavior. We observe for example in period 19 how different make-ups of inventory lead to different buying behaviors. In particular, we observe that having more returnable units leads to higher order up to levels since it allows the retailer to take on more risk. Similarly, when inventory level is high, the order up to level is higher for lower returnable levels to stock up on returnable units since the quota of such purchasable units is fixed in any given period.

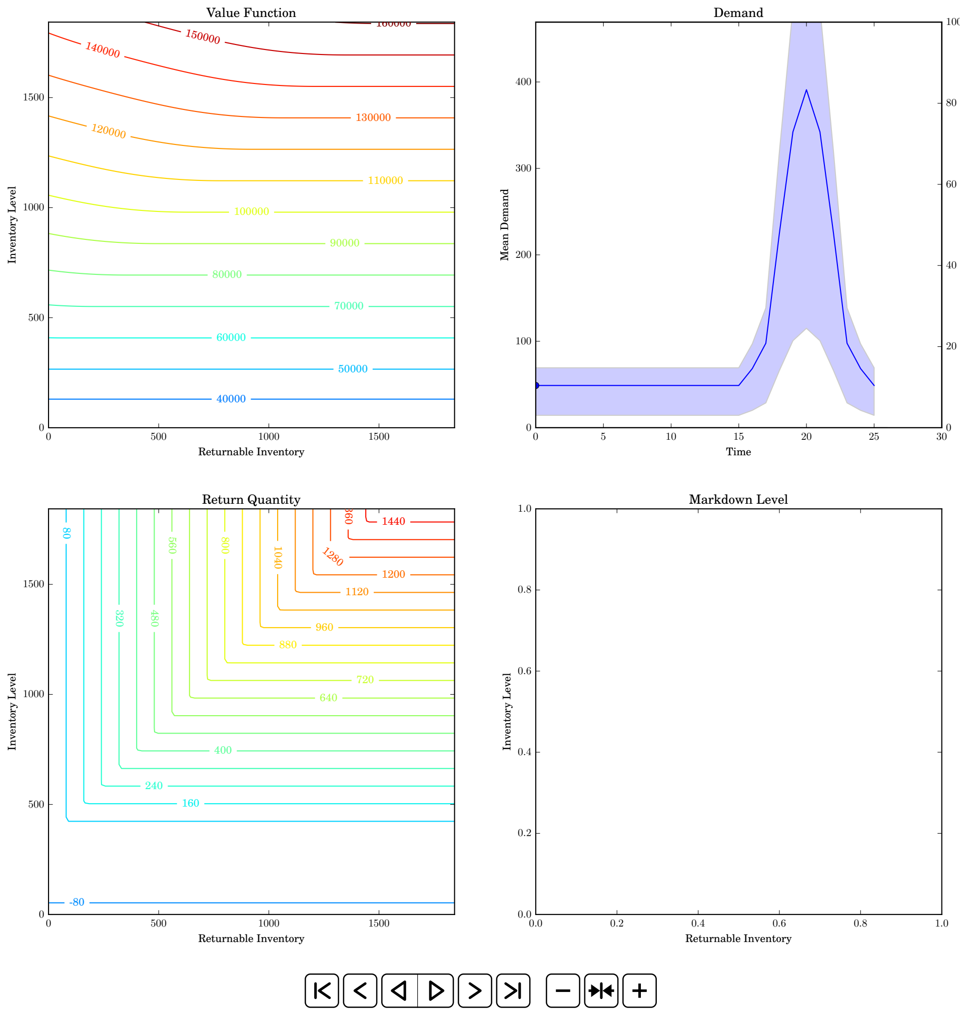

Figure 22: Summarizing graphs of the example in the fixed returnability case with modified parameters.

F.2 Fixed Returnability with Capacity Constraint

Figure 23 presents the results for the fixed returnability case in the presence of capacity constraints, as described in Section 7.3 for the case of full returnability. When compared to Figure 21, we observe that the inclusion of penalties causes a prolonged markdown season before the peak and more aggressive markdowns and returns coming down from it.

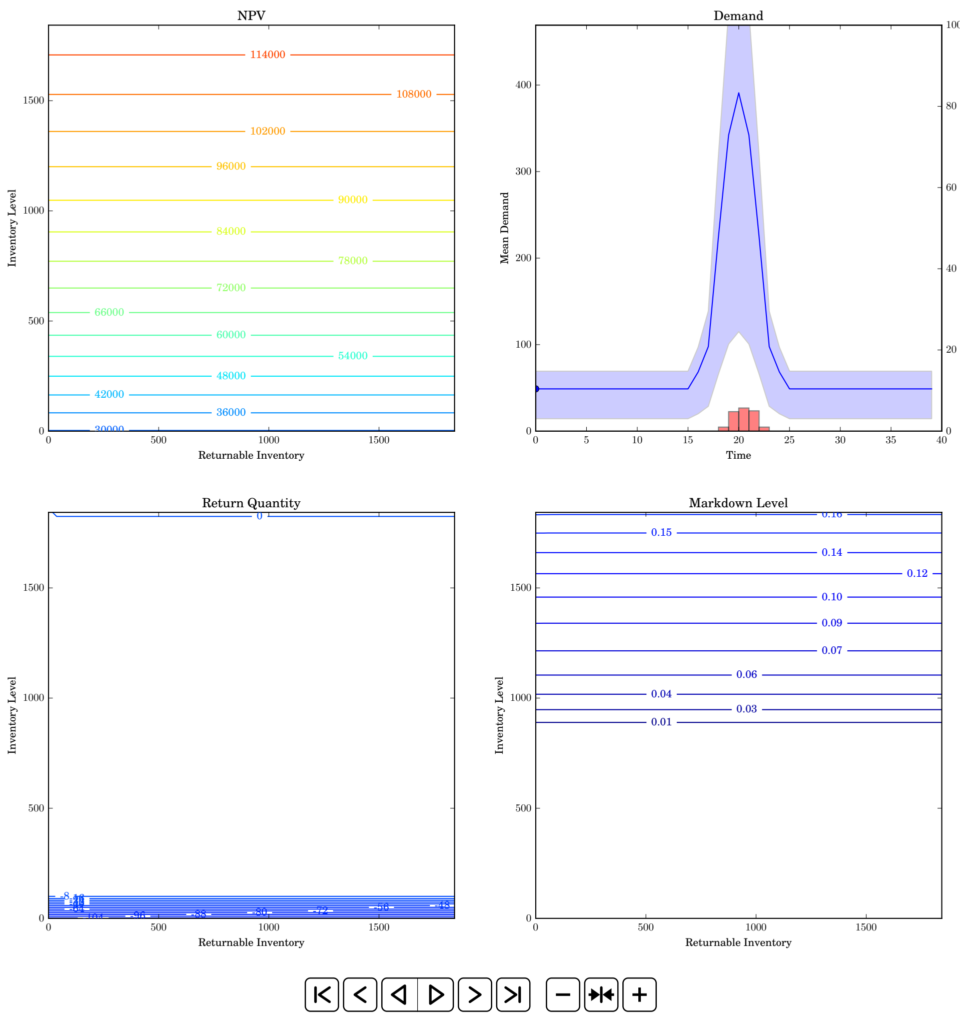

Figure 23: Summarizing graphs of the example in the fixed returnability case with capacity constraints.