Joint Pricing and Inventory Decisions for Substitutable and Perishable Products under Demand Uncertainty

Fei Fang* Tri-Dung Nguyen† Christine S.M. Currie ‡

August 17, 2020

Abstract

In this paper, we develop a model for dynamic pricing and inventory decisions for multiple substitutable and perishable products under a multiple-period lifetime. Retailers place orders at the beginning of the first period, and products will be sold at full price during that period. All leftover products will then be carried over to the subsequent periods and sold at discounted prices. Demands for leftover products are assumed to follow a linear stochastic model depending on the discounted prices of the substitutable products. The optimal order quantities and prices are obtained by maximizing the total expected profit over the lifetime, taking into account the revenue, the backorder cost and the holding cost. We provide analytical properties of the optimal policy such as the concavity of the value functions and then utilise these in the numerical scheme for finding the optimal prices and ordering quantities. Experimental results are reported through a case study of a high-street fashion company, demonstrating the benefits of considering pricing and inventory decisions simultaneously for substitutable products.

Keywords: Pricing, inventory, demand uncertainty, dynamic programming, perishable products, substitution.

*CORMSIS – Centre for Operational Research, Management Sciences and Information Systems, University Of Southampton, UK, f.fang@soton.ac.uk.

†Mathematical Sciences, Business School and CORMSIS, University of Southampton, Southampton, SO17 1BJ, United Kingdom, T.D.Nguyen@soton.ac.uk.

‡Mathematical Sciences, University of Southampton, Southampton, SO17 1BJ, United Kingdom, Christine.Currie@soton.ac.uk.

1 Introduction

Retailers frequently offer a wide selection of products, and customers will often be satisfied with a substitute if their preferred product is unavailable or its price is not attractive enough. This substitution effect makes it necessary to consider multiple products simultaneously when making price and inventory decisions. The prices charged for products will impact demand, having a knock-on effect on the optimal inventory policy. This interdependence means that pricing and inventory decisions should be made simultaneously, which has become more prevalent in recent years (e.g. Zhu and Thonemann (2009); Rajan et al. (1992); Abad (1996); Federgruen and Heching (1999); Gilbert (2000)). For perishable products, retailers will typically follow a markdown policy over the product's lifetime. Incorporating this effect by developing a pricing and inventory policy that extends over multiple periods and accounts for products being substitutable helps to prevent avoidable wastage.

We consider the situation in which a retailer is selling an identifiable set of substitutable products and where initial prices are set to the supplier's recommended retail prices. The retailer's objective is to maximise the profit obtained from selling the products, accounting for backorder and holding costs. Decision variables are the prices to charge in subsequent periods, which we refer to as the *markdown prices*, and the optimal initial order quantity. Demands for products are uncertain and dependent on the individual prices of all of the products within the substitutable set, an effect that we term *cross-price dependency*.

There is a rich literature on joint pricing and ordering decision problems. Rajan et al. (1992) study a joint pricing and inventory policy problem for a monopolistic retailer. In this model, the price is allowed to vary over the inventory cycle and no backlogging is permitted. Abad (1996) presents a model of dynamic pricing and lot-sizing by a retailer who sells a perishable product, where the demand is allowed to be partially backordered. They find that retailers can benefit from the backordering strategy to control costs on highly perishable products. Federgruen and Heching (1999) consider the problem of finding pricing and inventory replenishment strategies under demand uncertainty. In their model, both finite and infinite horizon models are studied and two options are considered for price changes: (i) the price can be changed arbitrarily (upwards or downwards); and (ii) the price can only be decreased. They develop the optimal policy by analysing the properties of both models. Li et al. (2009) study a pricing and inventory control problem for a perishable product with fixed lifetime. They formulate the problem as a Markov decision process and develop an optimal base-stock-list-price policy for products. Rujing (2007) studies the joint pricing and inventory decisions for a single perishable product with a multiple-period lifetime. They assume that decision makers determine the optimal order quantity at the beginning of the first period and determine a pricing policy dependent on the remaining inventory and the remaining lifetime.

The problem has some cross-over with dynamic pricing research, which has seen an increased focus on pricing models for multiple products since 2006, as discussed in the review article by Chen and Chen (2015). The majority of these models assume that arrivals follow a Poisson process and, where there is no clear preference order amongst the products (horizontal differentiation rather than vertical differentiation), customer choice is typically described using a random utility model (e.g. Akcay et al. (2010); Dong et al. (2009); Li and Graves (2012); Suh and Aydin (2011)). These models do not incorporate inventory decisions, assuming the available inventory to be fixed in advance and, being principally designed for the sale of transport and leisure services, do not tend to include either holding or backorder costs.

Despite the extensive literature in this area, to the best of our knowledge, this is the first work considering the joint pricing and inventory control problem across substitutable and perishable products over multiple periods taking account of holding and backorder costs. Gilbert (2000) studies a joint pricing and production schedule problem for a set of items but where a single price is used for the entire planning horizon. Demand is assumed to be seasonal and price dependent but cross-price effects on demand are not considered. A retailer faces cross-price effects when the demand for each product depends on the prices of other products in the market. Zhu and Thonemann (2009) do incorporate cross-price effects in their study, considering a joint pricing and inventory control periodic review problem for a retailer across two similar products. They derive the optimum joint pricing and inventory policy and find that the retailer can significantly improve profits by managing the two products jointly as opposed to independently, particularly under high cross-price demand elasticity. They also find that the retailer can improve profits by using dynamic pricing when demand is non-stationary. In both of these examples, the products are assumed to be non-perishable and consequently policies do not need to consider the shelf-life of products. Smith and Agrawal (2000) ignore the effect of price but use a newsboy formulation to develop optimal inventory policies that account for substitution effects. This joint selection of stock levels for substitutable products is shown to have a significant impact on the optimal inventory policies and profit margins.

The key contributions of our work are:

- We develop a stochastic dynamic programming model to study the joint pricing and inventory decision problem for substitutable and perishable products with stochastic demand. The model is introduced in Section 2.

- We derive analytical properties of the optimal decisions in Section 3.1. We show that the value functions are concave under mild conditions, which allows us to develop an efficient numerical scheme for finding the optimal joint ordering and pricing policy in Section 3.3. It should be noted that, while the concavity property has been shown previously in the inventory control literature, we show that the proof techniques in some of the existing work need to be more rigorous. Existing work has also generated these results in models that do not involve all the complexities of substitutability, perishability, multi-periods, and multiple products that we do. We provide a discussion of this in Section 3.2.

- The numerical results are based on a case study with real data, and are presented in Section 4, where we demonstrate the benefits of the strategies proposed.

2 Model formulation

We consider a setting in which decision makers, e.g., retailers, make pricing and ordering decisions over substitutable products at the beginning of the first period and pricing decisions during the subsequent discount periods. Here, the selling season is split into distinct time periods, where the prices on offer are constant during each period. Retailers are unlikely to want the duration of the markdown period to be equal to the time that the product is on sale at full price. As a result, the lengths of the time periods may not be equal. All leftover products will be carried over to the next period while all leftover products at the end of the last period will be disposed of. The ordering quantities and prices in each period are the decision variables in this joint inventory and pricing model. Note that all of the results are still applicable if the initial price is fixed by the supplier.

Demands for the products in each period are assumed to be cross-price dependent among substitutable products, with the addition of error terms to account for the uncertainty. The objective is to maximize the total expected profit over the planning horizon. In this research, we analyze the optimal solution-structure of the problem and show that retailers can significantly improve profits by managing these substitutable products simultaneously.

Let T be the number of periods and let m be the number of products. Let p_{it}, x_{it}, d_{it} be the price, the stock available, and the demand of product $i = 1, \ldots, m$ at time $t = 1, \ldots, T$, respectively. The retailer needs to decide on the initial level of stock available $x_{i1}, i = 1, .., m$, at the beginning of the time horizon as well as the set the prices (p_{it}), $t = 1, \dots, T$. Here, we assume that T, m take positive integer values, x_{it} takes integer values while p_{it} and d_{it} take non-negative real values. As soon as a product is out of stock, we no longer make the price decision and we can set it to zero in the subsequent periods. In what follows we use p_t, x_t and d_t to refer to the price, stock and demand vectors respectively at time t.

All products are substitutable, which means the demand for product i not only depends on its own price, but also on the prices of all other substitutable products. We assume that the demand follows a linear model as follows:

$$d_{it} = b_{it} + \sum_{j=1}^{m} a_{ij} p_{jt} + \varepsilon_{it}, \quad \forall t = 1, \dots, T, \quad \forall i = 1, \dots, m,$$

where b_{it} and a_{ij} are known parameters, and where ε_{it} is the uncertainty term, typically modelled as standard Gaussian noise, in the stochastic demand function. We assume $a_{ii} < 0, \forall i$, which means any increase in the price of a product would lead to a decrease in the demand for that product. We also assume $a_{ij} \geq 0, \forall j \neq i$, which implies the substitutability among the products; that is, an increase in the price of a product would drive customers to choosing other substitutable products. In other words, we can also interpret a_{ii} as the demand i loss per unit price increase of product i and a_{ij} as the demand i increase per unit price increase of product j, i.e., the price substitution rate. We deliberately choose a simple parametric model that satisfies our assumptions about customer behaviour to describe the demand. This gives us some clarity when understanding the results of the optimisation but recent work in dynamic pricing, such as by Besbes and Zeevi (2015), also suggests that the "price of misspecification" of linear models may not be as high as previously thought, particularly if models are recalibrated regularly. Similar models have been used previously in the literature in Yu (2015) and Kim and Bell (2011).

Let $x_t = \left[x_{1t}, \dots, x_{mt}\right]^\top$ be the (column) vector of left-over products at period t. Similarly, let p_t, d_t, ε_t be the vectors of prices, demands and error terms at period t.

The demand function can be rewritten in vector notation as

$$d_t = b_t + \mathrm{A} p_t + \varepsilon_t, \tag{1}$$

where $\boldsymbol{b_t} = \left[b_{1t}, \dots, b_{mt}\right]^\top$ is the constant vector representing the expected demand when all prices are set to zero, and where the matrix $A = (a_{ij})_{i,j=1,\dots,m}$ models the substitution relationship among the products. We also assume that A is a negatively diagonal dominant matrix, i.e.,

$$-a_{ii} \geq \sum_{j,j\neq i} a_{ij} \quad \text{and} \quad a_{ii} \geq \sum_{j,j\neq i} a_{ji}.$$

The negatively diagonal dominant matrix A guarantees that when all the prices are increased, the total demand will be decreased. In this paper, we assume that the demand function is linear on the prices for convenience as well as for our numerical tests but all the results presented in the paper are still applicable for other demand functions $d(p)$ as long as the expected revenue $\mathbb{E}[pd(p)]$ is a concave function on the price p. This also implies that we can relax constraints such as $a_{ij} \geq 0, \forall j \neq i$ stated above as long as the matrix A is negative semidefinite.

Note that here we assume A to be time-independent for a clearer exposition but the model can be extended to the time-dependent case. Note also that, if $a_{ij} = 0, \forall i \neq j$, then the products are 'unrelated' and the model to be described will become separable; that is the retailer can solve the optimal joint pricing and inventory for each product separately. All the theoretical results to follow are still applicable.

At period t, the realised demand given the price p_t is as given in Equation (1) and the revenue is $d_t^\top p_t$. If the stock available x_{it} exceeds the demand d_{it}, then the left-over stock is $x_{i,t+1} = x_{it} - d_{it}$. In that case, the retailer has to pay the holding cost $h_i(x_{it} - d_{it})^+$ to carry these over to the next period, where $h_i \geq 0$ is a given per-unit, per-time holding cost for product i, and where, for each $y \in \mathbb{R}$, $(y)^+ = \max(0, y)$ denotes the positive quantity in y and $(y)^- = -\min(0, y)$ denotes the negative quantity in y.

Conversely, if the demand d_{it} exceeds the stock availability, the retailer would need to backorder from another source at unit price π_i. We assume that $\pi_i \geq p_{i,max}$, which is defined as the maximum price for product i, to enforce the fact that backordering is costly. Here, as the products are marked down in the subsequent periods, it is possible to set $p_{i,max} = p_{i1}$, which is the price set at the first period. We assume a situation in which all stockouts are fulfilled through backordering. Other authors have instead considered partial backordering for inventory problems. For example, Pentico and Drake (2009) assume that a proportion β of stockouts will be backordered and the remainder lost. Nagarajan and Rajagopalan (2008) consider the case of substitutable products and assume that a fixed proportion of stockouts will instead purchase the substitute product. While we acknowledge that incorporating partial backordering will be useful in some practical situations, we leave this for future work.

Let $V_t(x_t)$ be the optimal expected profit-to-go from period t to T given stock availability at period t being x_t. We assume that the products need to be disposed of after T periods unless backordering is required; that is the leftover value function after period T can be written as

$$V_{T+1}(x_{T+1}) := -\sum_{i=1}^{m} \pi_i (x_{i,T+1})^{-}.$$

Let $I_t = \{i \in 1, \ldots, m \mid x_{it} > 0\}$ and $J_t = \{j \in 1, \ldots, m \mid x_{jt} \leq 0\}$. Then for $j \in J$, the retailer is out-of-stock for product j and accepts to pay the backordering cost $-\pi_j(x_{jt})^{-}$. In addition, for all $j \in J_t$, since we no longer sell product j in the subsequent periods, we simply set $p_{jk} = 0$ and assume $d_{jk} = 0, \forall k = t, \ldots, T$.

For the remaining leftover products in set I_t, let us denote $p_{I_t,t}$ to be the vector of prices and let $d_{I_t,t}$ be the corresponding demand vectors. We also assume that the demands for these products still follow a linear model $d_{it} = b_{it} + \sum_{j \in I_t} a_{ij} p_{jt} + \varepsilon_{it}$. For convenience, we use the same set of parameters (A, b) here but it is possible to have these dependent on the set of remaining products I_t.

In general, we have the following recursive formulation:

$$V_t(x_t) = -\sum_{j \in J_t} \pi_j (x_{jt})^{-} + \max_{p_{I_t,t}} \underbrace{\mathbb{E}\left[p_{I_t,t}^{\top} d_{I_t,t} - h \sum_{i \in I_t} (x_{it} - d_{it})^{+} + \gamma V_{t+1}(x_t - d_t)\right]}_{G_t(x_t, p_t)}, \tag{2}$$

where $G_t(x_t, p_t)$ denotes as the expected profit-to-go given the stock available x_t and the discounting price p_t.

This is called a Bellman optimality equation that represents the dynamic of the multi-period joint inventory and pricing problem over substitutable products. At each period t, the retailer aims to set the price vector p_t to maximise the total expected profit for period t, i.e. the quantity, $p_t^{\top} d_t - h \sum_{i=1}^{m} (x_{it} - d_{it})^{+}$, plus the profit-to-go $V_t(x_t - d_t)$ from period $(t+1)$ to T given the stock available $(x_t - d_t)$, discounted by factor γ.

For a product $i = 1, \ldots, m$ with $x_{it} < 0$, the corresponding profit-to-go is $-\pi_i(x_{it})^{-}$, which is negative to denote backordering cost. Here, we note that the selling cycle ends as soon as the stock has been exhausted, except for periods where demand exceeds supply and backordering is required. Note that reordering is only used to recover the mismatch between supply and demand.

Remark 1: Readers may recognise some unconventional notation here; in particular (a) we allow the stock level x_t to be negative, in which case a backorder cost is incurred, and (b) we settle the backordering cost as a part of the profit-to-go right from the beginning to handle cases when the demand from the previous period exceeds the capacity (and hence backordering is required). The main reason for doing this is for convenience in deriving the concavity of the profit-to-go functions, which will become clearer when we show the analytical proofs in Section 3.1. We also provide a comparison between our chosen model and existing joint pricing and inventory models in the literature in Section 6.1.

Remark 2: If we wish for the backordering cost $\pi_i(x_{it})^-$ not to be discounted from period $(t+1)$ to period t by a factor γ, we can replace this with $\frac{\pi_i}{\gamma}(x_{it})^-$ such that γ will be cancelled out under the term $\gamma V_{t+1}(x_t - d_t)$ in (2).

Without the holding and backordering costs and without constraining on non-negative stock availability, we can show that system (4) follows a linear-quadratic dynamic programming model, which is well studied in the literature, see for example Bertsekas (2005). In this case, the optimal dynamic decision on p_t is linear in the state vectors x_t. However, due to the holding and backordering costs, we need to pay particular attention to the sign (and magnitude) of the state and the decision variables (x_t, p_t). The optimal decisions no longer follow a linear policy and solving the problem now becomes more challenging. We will resolve this challenge by showing the concavity of both the expected profit $G_t(x_t, p_t)$ and the profit-to-go functions $V_t(x_t)$ in the next section.

Once we have worked backward to find the profit-to-go function $V_1(x_1)$, we can solve for the optimal initial ordering x_1. Let the ordering cost to be charged at the beginning of the first period be $\boldsymbol{c} = (c_1, ...c_m)$. Then the decision maker would solve:

$$\min_{x_1 \geq 0} \; V_1(x_1) - c^\top x_1.$$

3 Finding the Optimal Pricing and Inventory Policy

In this section, we first derive some theoretical properties on the concavity of the profit-to-go functions. We also discuss some common mistakes in the literature when deriving this for similar multi-period joint inventory and pricing problems. Finally, we use the concavity property to develop a numerical scheme for finding the optimal pricing and inventory policy.

3.1 Analytical Properties of Optimal Pricing and Inventory Policy

In this section, we present the concavity properties of the optimal pricing and inventory policy.

Lemma 1. *For $t = 1, \ldots, T$, if $V_{t+1}(x_{t+1})$ is concave with respect to x_{t+1}, then the expected profit $G_t(x_t, p_t)$ is jointly concave with respect to (x_t, p_t).*

Proof. The total profit is formulated as: $G_t(x_t, p_t) = F_t(x_t, p_t) + \gamma V_t(x_{t+1})$, where

$$\begin{aligned} F_t(x_t, p_t) &= \sum_{i=1}^{m} \left(p_{it}\mathbb{E}d_{it} - h_i\mathbb{E}(x_{it} - d_{it})^+\right) \\ &= p_t^\top A p_t + b_t^\top p_t - \sum_{i=1}^{m} \left(h_i\mathbb{E}(x_{it} - d_{it})^+\right). \end{aligned}$$

We first prove that $F_t(x_t, p_t)$ is jointly concave with respect to (x_t, p_t). The first quadratic term $p_t^\top A p_t + b_t^\top p_t$ is concave with respect to p_t due to the negatively diagonal dominance property of matrix A.

Considering now the second term, let us introduce auxiliary variables

$$z_{it} = x_{it} - (b_{it} + \sum_{j=1}^{n} a_{ij}p_{jt}).$$

Then the holding cost for product i can be rewritten as:

$$L_{it} := h_i\mathbb{E}(x_{it} - d_{it})^+ = h_i\mathbb{E}(z_{it} - \xi_{it})^+.$$

We find that L_{it} is a piecewise linear and convex function with respect to z_{it} as both the Expectation operator and the $(\cdot)^+$ operator preserve convexity. Therefore, L_{it} is jointly convex with respect to (x_t, p_t) since z_{it} is a linear function of (x_t, p_t). In summary, $F_t(x_t, p_t)$ is now proven to be jointly concave with respect to (x_t, p_t).

Finally, $V_{t+1}(x_{t+1})$ is jointly concave with respect to (x_t, p_t) since $x_{i,t+1}$ is a linear function of (x_t, p_t) while $V_{t+1}(x_{t+1})$ is concave with respect to x_{t+1} by the lemma's hypothesis.

Altogether, we show that G_t is jointly concave with respect to (x_t, p_t) if $V_{t+1}(x_{t+1})$ is concave with respect to x_{t+1} . □

Remark 4: Note that, for the case of $t = 1$, although the concavity of $G_1(x_1, p_1)$ is derived (under the appropriate assumption), we assume that p_1 is given in our setting.

We will first prove that if $V_{t+1}(x_{t+1})$ is concave with respect to x_{t+1}, then $V_t(x_t)$ is jointly concave with respect to x_t. We first make the following assumption.

Assumption 1. *(i)The backordering cost for each product is not smaller than its maximum sale price; that is, $\pi_i \geq p_{i,\max}$ for all i. (ii) The uncertain demand $d(p)$ is non-negative for all feasible price p.*

The first part of Assumption 1 essentially states that backordering is costly. This assumption will also be used to show the concavity of the profit-to-go function in Section 3.

Lemma 2. *Under Assumption 1, if $V_{t+1}(x_{t+1})$ is concave with respect to x_{t+1}, then $V_t(x_t)$ is concave with respect to x_t.*

The proof makes use of the following proposition.

Proposition 1 (Proposition B-4, Concavity reservation Under Maximization, in Hayman and Sobel (1984))**.** *Let Y be a nonempty set with A_y a nonempty set for each $y \in Y$. Let $C = \{(y,z) : z \in A_y, y \in Y\}$, let J be a real-valued function on C, and define*

$$f(y) = \inf\{J(y,z) : z \in A_y\}, y \in Y$$

If C is a convex set and J is a convex function on C, then f is a convex function on any convex subset of $Y^* = \{Y : y \in Y, f(y) > -\infty\}$.

Remark 5: It is easy to derive a similar version of Proposition 1 where the convexity of J is replaced by concavity and the infimum operator is replaced by the supremum (or maximum) operator.

Proof of Lemma 2. Under the assumption that $V_{t+1}(x_{t+1})$ is concave with respect to x_{t+1}, the condition in Lemma 2 holds and it follows that the total profit $G_t(x_t, p_t)$ is jointly concave with respect to (x_t, p_t). According to Proposition 1, the concavity is preserved under maximization of $G_t(x_t, p_t)$ and hence the second part of Equation (2), i.e., $\max_{p_{I_t,t}} \mathbb{E}\left[p_{I_t,t}^\top d_{I_t,t} - h\sum_{i \in I_t}(x_{it} - d_{it})^+ + \gamma V_{t+1}(x_t - d_t)\right]$, is concave on the non-negative orthant of x_t. We are left to show that $V_t(x_t)$ is concave with respect to x_t over the entire domain of x_t.

For those $x_{it} < 0$, the profit-to-go is $-\pi_i(x_{it})^-$ which is a linear function with slope π_i. Under Assumption 1, the revenue from selling the products is at most $\sum_{i=1}^m \pi_i x_{it}^+$ and hence the hyperplane $\pi_i(x_{it})$ is a super-gradient of the profit-to-go function at $x_{it} = 0$. Thus, the profit-to-go function as shown in Equation (2) can be viewed as the minimum of two concave functions: $\pi \sum_{i=1}^m x_{it}$; and a truncated concave function. This implies that the profit-to-go function is a concave function over the entire domain of x_t.

□

We now can state the main theorem.

Theorem 3. *The following statements hold:*

1. *The expected profit $G_t(x_t, p_t)$ is jointly concave with respect to (x_t, p_t), $\forall t = 1...T$*

2. The maximum expected profit $V_t(x_t)$ is concave with respect to x_t, $\forall t = 1...T$.

Proof. We will use an induction method and results from Lemmas 1 and 2 to prove this theorem. In the last period, we have

$$V_{T+1}(x_{T+1}) := -\sum_{i=1}^{m} \pi_i (x_{i,T+1})^-,$$

which is a concave function on x_{T+1}. Thus, $G_T(x_T, p_T)$ is a jointly concave function on (x_T, p_T) by Lemma 1. As a result, $V_T(x_T)$ is a concave function on x_T by Lemma 2. Thus, the results stated in the theorem hold for $t = T$. Suppose the statements hold for $(t+1)$, i.e., $G_{t+1}(x_{t+1}, p_{t+1})$ is jointly concave with respect to (x_{t+1}, p_{t+1}) and $V_{t+1}(x_{t+1})$ is concave with respect to x_{t+1}. We will also prove that these hold for t. Indeed, since $V_{t+1}(x_{t+1})$ is concave with respect to x_{t+1}, by Lemma 1, we have $G_t(x_t, p_t)$ is jointly concave on (x_t, p_t). Similarly, by Lemma 2, we have $V_t(x_t)$ is concave on x_t. Thus, the statements hold for t. This completes the proof. □

3.2 Numerical Scheme for Finding the Optimal Policy

The concavity properties of $G_t(x_t, p_t)$ and $V_t(x_t)$ for each $t = 1, \ldots, T$ allow us to design an efficient algorithm to solve for the optimal price vector p_t and the initial ordering vector x_1. Specifically, we follow standard methods from dynamic programming to work backward from $t = (T+1)$ to $t = 1$ when we calculate $V_t(x_t)$. To make things easier, we assume that x_t only takes discrete values from a given set. We believe this to be a reasonable assumption since the stock availability should only take integer and bounded values. Here, we note that the Gaussian assumption on the error terms of the demand model is used only when we utilise the least square regression to estimate the parameters for the numerical study. The theoretical results apply for both discrete and continuous demand distributions and only requires the error terms to have zero means.

This means that the functional form of $V_t(x_t)$ can be viewed as a look up table on t and x_t. Our task is simply to fill in an $(m+1)$-dimensional table where the first dimension is for t and the remaining m dimensions for x_t. We will fill in the table from $t = (T+1)$ and work backward to $t = 1$.

We already know that $V_{T+1}(x_{T+1}) = -\sum_{i=1}^{m} \pi_i (x_{i,T+1})^-$. Suppose that we have already determined $V_{t+1}(x_{t+1})$, in the form of a look up table, for some t and we need to work out $V_t(x_t)$. For each x_t, we need to solve Equation (2). The first part of $-\sum_{j \in J_t} \pi_j (x_{jt})^-$ is easy to compute. The second part involves maximizing the function $G_t(x_t, p_t)$ over the decision vector p_t. In our numerical scheme, the expectation is replaced by the sample approximation over 1000 random ξ_t. Given the concavity property of $G_t(x_t, p_t)$ as derived in the previous section, its maximization over p_t can be computed efficiently using a number of possible numerical schemes such as the downhill simplex method in Nelder and Mead (1965), which is a derivative-free method that guarantees convergence if designed appropriately. For simplicity, we take advantage of the concavity and the fact that the decision variables p_t only take discrete values in reality and design a simple binary-search method for optimising $G_t(x_t, p_t)$ over p_t. Specifically, we start the search from a coarse grid of p_t and find an optimal solution $\tilde{p}_t$. By concavity, we know that the global optimal solution p_t^* must lie within the neighbour grids around $\tilde{p}_t$. We therefore narrow the search on the neighbourhood over a finer grid. The process is repeated several times until we find the solution to be of acceptable accuracy (in the numerical experiments, we set the prices to the nearest ten pence.)

4 Numerical Results

In this section, we demonstrate how a retailer can improve profitability by making the pricing and inventory decisions together. This is demonstrated through 110 randomly generated instances and through a case study of a joint pricing and inventory problem, considering substitution between two similar products. The case we worked on is the sale of a new design of jeans in the winter collection (2015-2016). We collect data from the sales of the old design with two different washes (stone washed and fabric dyed) in the previous winter sales (2014-2015). The stone washed jeans are deemed to be product number 1 and the fabric dyed are deemed to be product number 2.

We are provided with the ordering cost from the supplier, the number of sales and prices for each wash in the full price selling period (Nov-Dec) with c_1= £26, c_2= £28. Holding costs and backorder costs are provided by the revenue manager and are assumed to be the same in both periods with $h = \pounds 10$, $\pi = \pounds 152$. The maximum prices are also set at $p_{i,max} = \pi$ to be consistent with Assumption (1).

We are also provided with six sets of observations for the second period, where in each set we have the prices of the two products with different washes; and the number of sales of each. Other information we have in the data includes the ordering cost, the serial number, the description of the products, the store name and additional fields not relevant to the optimisation we describe here. With these six observations, while quite limited, we are able to employ a linear regression and to provide estimates of the parameters in the demand function.

We run linear regression models over these historical observations (including the prices and sales of the two products) to estimate the demands as linear functions of the prices. The following are the resulting estimated parameters: $a_{11} = -7.5$, $b_1 = 789$, $a_{21} = 1.2$; $a_{22} = -1.6$, $b_2 = 105$, $a_{12} = 0.6$. We use the same set of parameters for each of the time period $t = 1, 2, 3$.

Using results from linear regression, we also assume that the error terms ε in the second period follow normal distributions with mean 0 and standard deviations of 81 and 30 for the two products, respectively; that is $\varepsilon_{1,2} \sim N(0, 81), \varepsilon_{2,2} \sim N(0, 30)$, where $N(\mu, \sigma^2)$ denotes a normal distribution with mean μ and variance σ^2. The standard deviations are obtained as the Root Mean Squared Error (MSE) in the regression model, which represents a frequently used measure of the differences between values (sample and population values) predicted by a model or an estimator and the values actually observed.

With these, the demand functions for the products are:

$$d_{12} = 789 - 7.5p_{12} + 0.6p_{22} + \xi_{12},$$

$$d_{22} = 105 - 1.6p_{22} + 1.2p_{12} + \xi_{22}.$$

4.1 Benefits of Jointly Making Inventory and Pricing Decisions

In this part, we demonstrate the benefits of jointly making pricing and inventory decisions by comparing it with two other benchmark strategies that a retailer can follow, namely to consider inventory optimization and pricing optimisation in a decoupled manner. We assume that a manager has separate tools for inventory optimisation and pricing optimization and applies them sequentially. This results in three approaches, which for convenience we label as:

A1 : **Inventory-Price-Inventory** In this case, we first use the average price $\bar{p} = p_{max}/2$ to carry out inventory optimization and obtain x^1. We then obtain the optimal prices corresponding to inventory x^1 and finally using those optimal prices carry out a further inventory optimization.

A2 : **Price-Inventory-Price.** We assume the manager optimises the prices initially assuming average inventory levels, before using these optimal prices and finding the corresponding optimal inventory and finally re-optimising the prices.

A3 : **Joint pricing and inventory optimization** (where both x and p are decision variables). In this case, the manager makes joint pricing and inventory decisions.

Under A3, the optimal ordering quantities for the two products are 748 and 199 correspondingly. This corresponds to an optimal total expected profit-to-go over three periods of £22,997. The optimal prices for the first product over the three periods are £73, £65 and £56, respectively while those for the second product are £83, £73 and £63. Under A1, the total expected profit-to-go is £20,155 and the optimal quantities are 1000 and 241, which both exceed the optimal ordering quantities under A3. Under A2, the total expected profit-to-go is £20,749 and the corresponding optimal ordering quantities for the two products are 845 and 311, respectively, both of which also exceed the optimal ordering quantities under A3.

Comparing between approaches A1 and A3, we can observe that making joint inventory and pricing decisions will lead to an increase of around 12.36% in profit compared to inventory optimisation only. Similarly, we can observe an increase of around 9.78% in profit through changing from A2 to A3.

4.2 Benefits of Considering Substitution

In this study, we recommend that managers account for the effects of substitutions when making inventory and pricing decisions. We quantify the benefits of this recommendation in what follows by comparing our method with a benchmark where the retailer makes independent decisions on the (substitutable) products. Essentially, this involves comparing between two models:

B1 Without considering substitution, demand is assumed to follow: $d_i = \beta_i - \alpha_i p_i + \xi_i', \forall i$.

B2 Considering substitution, demand is assumed to follow: $d_i = b_i - a_i p_i + l_j p_j + \xi_i, \forall i$.

In model B2, the price substitution has been taken into account between products 1 and 2; this policy is the optimal policy we are proposing. In model B1, we assume that the demand of a product depends only on its own price; we term this the independent policy.

Due to the limited number of observations in our dataset, and in order to obtain more reliable results, as well as to test model B1, we will take the following steps in our experimental set up. We use the multi-variate demand function in (1) to generate 100 historical pairs of prices. This provides us with the corresponding 100 pairs of historical demands. To consider B1, we now run linear regression on each individual product, i.e., for each $i = 1, 2$, we run least square regression of historical demand d_i over prices p_i in order to obtain β_i and α_i. Next, we find the optimal decisions (order quantities and prices) for both models. Here, we note that this step is still the same for model B2, while that of model B1 involves running two separable single product cases, one for each product. While the total optimal profit-to-go is still £22,997 for model B2 as before, the total profit-to-go for case B1 is now reduced to £22,588. The loss is essentially because there are cross-effects on the demands between the two products that model B1 ignores. While the loss in the case is not too much, this is due to the relative scale between parameters l_j and a_i. In the next section, we will explore how the benefit of considering substitution is related to the magnitude of the price cross-effect.

4.3 Simulated Examples

In this section, we generate 110 random instances and explore further the results in the previous case study. Specifically, we assume the matrix A in the demand function (1) has the following structure

$$A = \begin{bmatrix} -1 & \delta \\ \delta & -1 \end{bmatrix},$$

where δ can take 11 increasing values from $0, 0.05, 1, \ldots, 0.5$. Here, δ represents the level of the price cross-effects between the products. For example, $\delta = 0$ means the two products are uncorrelated and we expect that there is no loss in not considering product substitution, while larger δ implies a higher differentiation between models B1 and B2, as presented in the previous section.

We also use 10 different random seeds and this together with the 11 possibilities for δ give us 110 randomly generated instances. For each fixed random seed, we generate random ϵ and random b. Here, b receives integer values between 10 and 20.

Similar to the case study, we set π as the nearest integer value of $p_i^*, i = 1, 2$, where $p^* = -1/2A^{-1}b$ is the optimal price if there were no holding and backorder costs. The maximum price is set as π for consistency with Assumption 1. The original ordering cost is set as $c = 0.3p^*$ while the holding cost is set as $h = 0.3 \min(p_1^*, p_2^*)$. We set the maximum ordering quantity to be $T \max(b_1, b_2)$, where $T = 3$ is the number of periods, and where $\max(b_1, b_2)$ represents a reasonable estimate of the maximum ordering quantity (given the demand function (1)). The rest of the simulation setup is similar to that described in Sections 4.1 and 4.2.

Figure 1 shows the percentage loss of not considering substitution, i.e., using model B1 instead of model B2, for different choices of δ. The boxplots are generated over the 10 random seeds. With $\delta = 0$, the prices set by one product do not affect the demand of the other and the two products are considered unrelated or non-substitutable. In this case, there is no difference in results for models B1 and B2. We find that the losses generally increase with the increase of δ to around 9% at $\delta = 0.5$. The results show that it is beneficial to consider substitution if there is evidence of cross price effects between products.

Figure 2 shows the percentage loss incurred by not considering inventory and pricing jointly, i.e., using the optimal pricing and optimal inventory in a decoupled manner. The boxplots are generated over the 110 instances, i.e., 10 random seeds and 11 different values of δ. We find that the losses are around 9.6% for case A1 and are around 7.9% for case A2. In some instances, the losses could be up to around 30%. We note here that models A1 and A2 have been considered 'smart' in the sense that, although decoupled, the manager still consider them sequentially. The losses could have been higher if

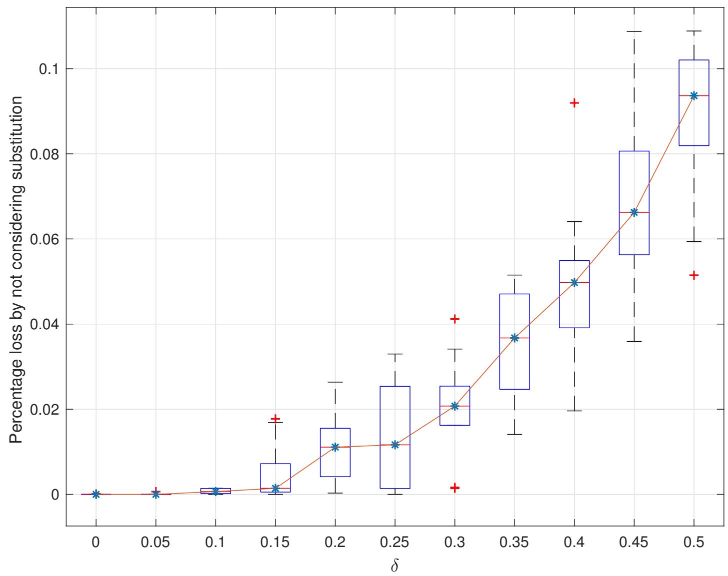

Figure 1: Benefits of considering substitution

we had only considered pricing (by taking, say the average inventory) or only inventory (by taking the average price). The results show that it is beneficial to jointly optimise pricing and inventory in the same model.

5 Conclusion

In this paper, we develop a stochastic dynamic programming model for joint pricing and ordering decisions over substitutable products. The framework is applicable to settings where there are multiple discounting periods and where the values of perishable products decrease quickly such as in fast fashion industries. Once the ordering decision is made at the beginning of the first period, no replenishment is considered but discounted retail prices might be offered in the subsequent periods. We show that the expected profit-to-go is jointly concave with respect to the retail prices and the initial ordering under mild assumptions. We also point out some common mistakes in the literature with regards to the concavity analysis. We utilise this concavity property and develop an efficient numerical scheme for computing the optimal ordering and pricing policy. In the numerical results, we consider a case study of fashion products and demonstrate that by jointly making pricing and inventory decisions, the retailers can benefit significantly. In addition, by considering substitutions between similar products, it is possible to obtain higher profits, which increase as the level of substitution increases.

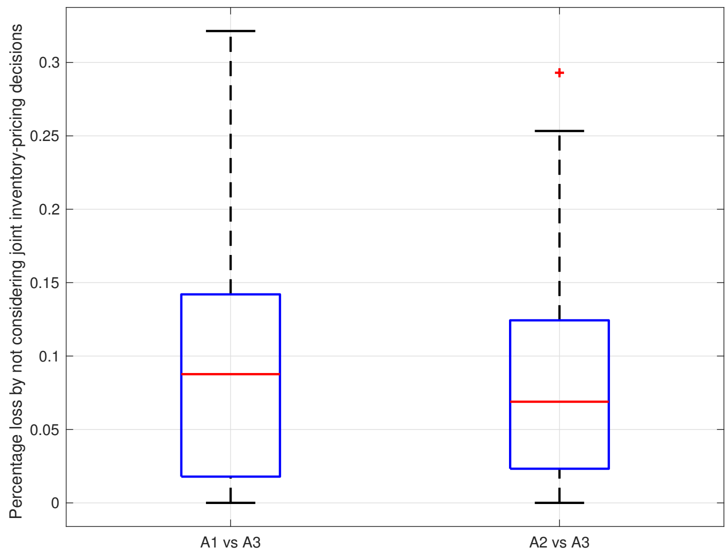

Figure 2: Benefits of considering joint inventory-pricing decisions

Acknowledgement

The authors would like to thank the Editor and the three anonymous reviewers for their valuable comments and detailed suggestions on how to improve the manuscript. The second author acknowledges the funding support from the Engineering and Physical Sciences Research Council (grant EP/P021042/1).

References

Abad, P., 1996. Optimal pricing and lot-sizing under conditions of perishability and partial backordering. Management Science 42 (8), 1093–1104.

Akcay, Y., Natarajan, H., Xu, S., 2010. Joint dynamic princing of multiple perishable products under consumer choice. Management Science 56 (8), 1345–1361.

Bertsekas, D., 2005. Dynamic programming and optimal control. Vol. 1. Athena scientific Belmont, MA.

Besbes, O., Zeevi, A., 2015. On the (surprising) sufficiency of linear models for dynamic pricing with demand learning. Management Science 61 (4), 723–739.

Ceryan, O., Sahin, O., Duenyas, I., 2013. Dynamic pricing of substitutable products in the presence of capacity flexibility. Manufacturing & Service Operations Management 15 (1), 86–101.

Chen, L.-M., Sapra, A., 2013. Joint inventory and pricing decisions for perishable products with two-period lifetime. Naval Research Logistics (NRL) 60 (5), 343–366.

Chen, M., Chen, Z.-L., 2015. Recent developments in dynamic pricing research: Multiple products, competition, and limited demand information. Production and Operations Management 24 (5), 704–731.

Dong, L., Kouvelis, P., Tian, Z., 2009. Dynamic pricing and inventory control of substitute products. Manufacturing and Service Operations Management 11 (2), 317–339.

Federgruen, A., Heching, A., 1999. Combined pricing and inventory control under uncertainty. Operations research 47 (3), 454–475.

Gilbert, S., 2000. Coordination of pricing and multiple-period production across multiple constant priced goods. Management Science 46 (12), 1602–1616.

Hayman, D., Sobel, M., 1984. Stochastic models in operations research, vol ii.

Kim, S.-W., Bell, P. C., 2011. Optimal pricing and production decisions in the presence of symmetrical and asymmetrical substitution. Omega 39 (5), 528–538.

Li, H., Graves, S., 2012. Pricing decisions during inter-generational product transition. Production and Operations Management 21 (1), 14–28.

Li, Y., Lim, A., Rodrigues, B., 2009. Note—pricing and inventory control for a perishable product. Manufacturing & Service Operations Management 11 (3), 538–542.

Nagarajan, M., Rajagopalan, R. S., 2008. Inventory models for substitutable products: Optimal policies and heuristics. Management Science 54 (8), 1453–1466.

Nelder, J., Mead, R., 1965. A simplex method for function minimization. The Computer Journal 7 (4), 308–313.

Pentico, D. W., Drake, M. J., 2009. The deterministic eoq with partial backordering: A new approach. European Journal of Operational Research 194, 102–113.

Rajan, A., Steinberg, R., Steinberg, R., 1992. Dynamic pricing and ordering decisions by a monopolist. Management Science 38 (2), 240–262.

Rujing, L., 2007. Joint pricing and ordering decisions for perishable products. Ph.D. thesis, National University of Singapore.

Smith, S., Agrawal, N., 2000. Management of multi-item retail inventory systems with demand substitution. Operations Research 48 (1), 50–64.

Suh, M., Aydin, G., 2011. Dynamic pricing of substitutable products with limited inventories under logit demand. IIE Transactions 43 (5), 323–331.

Yu, M., 2015. Joint pricing and inventory decisions under resource constraints and price-based substitution. Journal of Revenue & Pricing Management 14 (1), 19–27.

Zhu, K., Thonemann, U. W., 2009. Coordination of pricing and inventory control across products. Naval Research Logistics (NRL) 56 (2), 175–190.

6 Appendices

6.1 Alternative Dynamic Joint Pricing and Inventory Models

In this section, we present alternative dynamic joint pricing and inventory models in the literature. The first model, which commonly appears in the literature such as in Bertsekas (2005); Zhu and Thonemann (2009), is presented below

$$\phi_t(x_t) = \max_{p_t} \; \mathbb{E}\left[p_t^\top d_t - h\sum_{i=1}^{m}(x_{it} - d_{it})^+ - \sum_{i=1}^{m}\pi_i(x_{it} - d_{it})^- + \gamma\phi_{t+1}(x_t - d_t)\right], \tag{3}$$

where $\phi_t(x_t)$ denotes the profit-to-go function over real values of the stock level x_t.

In this model, x_t can take *negative values* and in that case the selling cycle can continue thanks to the capability to order more stocks, not only to recover the mismatch between the supply and demand but also for future sales if we are not in the very last selling period. It is also assumed that the products need to be disposed of after T periods unless backordering is needed; that is the leftover value function after period T can be written as $\phi_{T+1}(x_{T+1}) := 0$.

This model is appropriate for settings where the repeated reordering can take place during the course of sale. Given the fact that this paper considers short season products and that the ordering takes place only at the beginning and potentially in the last iteration when demand exceeds inventory, Chen and Sapra (2013) present the following alternative model.

$$\Phi_t(x_t) = \max_{p_t} \; \mathbb{E}\left[p_t^\top d_t - h\sum_i (x_{it}-d_{it})^+ - \sum_i \pi_i (x_{it}-d_{it})^- + \gamma\Phi_{t+1}((x_t-d_t)^+)\right], \tag{4}$$

where $\Phi_t(x_t)$ denotes the profit-to-go function over stock level x_t.

In this model, it is implied that $x_t \geq 0$, that is the stock availability should not take a negative value and the selling cycle ends as soon as the stock has been exhausted, except for periods where demand exceeds supply and backordering is required. Note that backordering is only used to recover the mismatch between supply and demand. In this model, it is also assumed that the products need to be disposed of after T periods, i.e., $\Phi_{T+1}(x_{T+1}) := 0$. Model (4) is also mentioned as an extension of Model (3) in Ceryan et al. (2013).

We note that the difference between Models (3) and (4) is in the last term where the future cost-to-go is written as $\Phi_{t+1}((x_t - d_t)^+)$ instead of $\phi_{t+1}(x_t - d_t)$. The discrepancy between the two models is due to the fact that their domains on the stock availability are defined differently.

While Model (4) is more appropriate for our application, it leads to some technical difficulties in the analysis of the profit-to-go function and the optimal policy. Specifically, it would be a mistake if we attempt to follow Zhu and Thonemann (2009) to show that 'the Hessian matrix' of $\Phi_t(x_t)$ is negative semidefinite. This is not the right approach since $V_t(x_t)$ can be non-differentiable. In fact, the involvement of the holding and backordering costs as well as the signed constraints over (x_t, p_t) results in $G_t(x_t, p_t)$ and $V_t(x_t)$ having a piece-wise shape that not be differentiable at the joints of these pieces. In (Chen and Sapra, 2013, page 357), the authors use an induction argument, i.e., the concavity of the profit-to-go function at period $t+1$, to show the negative semidefinite property of the Hessian matrix of the profit-to-go function at period t (In the proof of Theorem 1, Part 2B). The derivation essentially takes for granted that the Hessian matrix of the profit-to-go function at period $t+1$ exists without a rigorous proof. Here, we note that the concavity of the profit-to-go function does not imply the existence of its Hessian matrix. Similarly, in Ceryan et al. (2013), the authors consider mainly Model (3) and claim that all results can be extended to Model (4) but we believe that some further conditions are perhaps needed for showing the twice-differentiability property.

If we wish to use Model (4), we would need to be extra careful on handling the sign of $(x_t - d_t)$ on studying the concavity of $\Phi_{t+1}(x_{t+1})$, or equivalently the concavity of $\Phi_{t+1}(x_t - d_t)$. While it is possible for $\Phi_{t+1}(x_{t+1})$ to be concave over non-negative x_{t+1}, it is not concave over the entire $\mathbb{R}^m$ domain unless we incorporate the backordering cost into $\Phi_{t+1}(x_{t+1})$ (in addition to assuming that the backordering price is higher than the sale price). Therefore, $\Phi_{t+1}(x_t - d_t)$ might not be concave over x_t (or even over

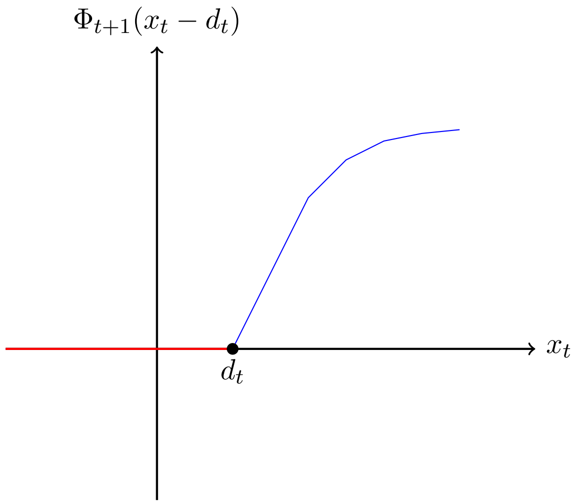

(a) Φ_{t+1} is non-concave over x_t if $\pi = 0$

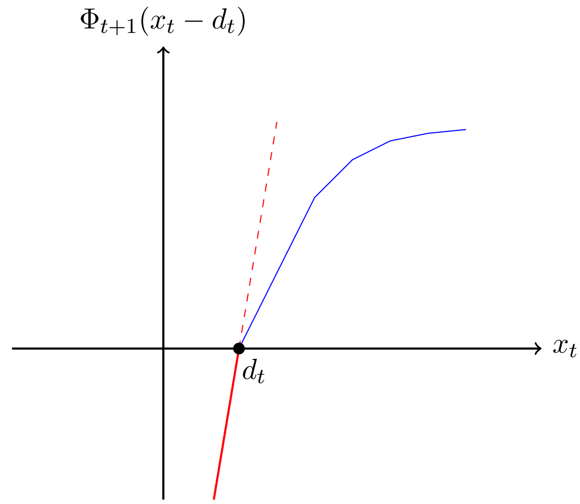

(b) Φ_{t+1} is concave over x_t if $\pi > p$

Figure 3: Concavity of Profit-to-go functions

non-negative x_t) as this depends also on the sign of $(x_t - d_t)$.

Figure 1a shows an example where $\Phi_{t+1}(x_t - d_t)$ is concave over $x_t \geq d_t$ but this does not translate to the concavity of $\Phi_{t+1}(x_t - d_t)$ over all x_t if we do not take the backordering cost into account. Figure 1b shows how the inclusion of the backordering cost resolves this issue. Note, however, that this only holds in the case where the backordering price is higher than the sale price. In both cases, we deliberately show the profit-to-go function to be piece-wise concave (where each piece could be linear or non-linear) and is generally non-differentiable at the joints.

6.2 Comparison Between Models (2) and (4)

Due to the technical difficulties with analysing Model (4) directly as presented in Section 6.1, we develop Model (2) in Section 2. Intuitively, the two models are quite similar except for the rearrangement of the backordering cost terms and the explicit handling of the set of remaining stocks in Model (2). We now show that the two models share common features as formally stated in the following Lemma.

Lemma 4. *Under Assumption* (1)*, the following results hold*

(a) $\Phi_t(0) = 0$ *for all* $t = 1, \ldots, T$.

(b) In the case $m = 1$*, we have* $V_t(x_t) = \Phi_t(x_t)$ *for all* $t = 1, \ldots, T$*, and* $x_t \geq 0$.

The fact that $\Phi_t(0) = 0$ is important as it is fair to expect that the profit-to-go is equal to zero when there is no stock available. The result of $\Phi_t(0) = 0, \forall t = 1, \ldots, T$ is important to show that both Models (4) and (2) are 'nearly' equivalent as stated in part (b).

Proof. Part (a). We will prove the lemma by induction. For $t = T+1$ the result holds trivially by our construction of $\Phi_{T+1}(x_t) = 0, \forall x_t \geq 0$. Suppose the result stated in the lemma holds for $t+1$, i.e., $\Phi_{t+1}(0) = 0$. We need to prove that $\Phi_t(0) = 0$.

Using the Bellman Equation (4), we have

$$\Phi_t(0) = \max_{p_t} \mathbb{E}\left[p_t^\top d_t - h\sum_i (-d_{it})^+ - \sum_{j \in J_t} \pi_j(-d_{it})^- + \gamma\Phi_{t+1}((-d_t)^+)\right] = \max_{p_t} \mathbb{E}\left[(p_t - \pi_i)^\top d_t\right],$$

where the simplification in the second equality is due to the assumption that $d_{it} \geq 0$ and that $\Phi_{t+1}(0) = 0$.

Since $d_{it} \geq 0$ and since $p_i t \leq \pi_i$, the maximum is attained at $p_{it} = \pi_i$ and at which point we have $\Phi_t(0) = 0$.

Part (b). We will prove by induction. For $t = T+1$, the result holds trivially by the construction of $V_t(x_t)$ and $\Phi_t(x_t)$. Suppose the result stated in the lemma holds for $t+1$, i.e., $V_{t+1}(x_{t+1}) = \Phi_{t+1}(x_{t+1})$ for all $x_{t+1} \geq 0$. We will prove that $V_t(x_t) = \Phi_t(x_t)$ for all $t = 1, \ldots, T$, and $x_t \geq 0$.

For $x_t \geq 0$, we have

$$\begin{aligned}
V_t(x_t) &= \max_{p_t} \mathbb{E}\left[p_t d_t - h(x_t - d_t)^+ + \gamma V_{t+1}(x_t - d_t)\right] \\
&= \max_{p_t} \mathbb{E}\left[p_t d_t - h(x_t - d_t)^+\right] + \mathbb{E}\left[\gamma V_{t+1}(x_t - d_t)\right] \\
&= \max_{p_t} \mathbb{E}\left[p_t d_t - h(x_t - d_t)^+\right] + \gamma\int_{d_t \leq x_t} V_{t+1}(x_t - d_t)dF(d_t) + \gamma\int_{d_t > x_t} V_{t+1}(x_t - d_t)dF(d_t) \\
&= \max_{p_t} \mathbb{E}\left[p_t d_t - h(x_t - d_t)^+\right] + \gamma\int_{d_t \leq x_t} \Phi_{t+1}(x_t - d_t)dF(d_t) - \gamma\int_{d_t > x_t} \frac{\pi}{\gamma}(d_t - x_t)dF(d_t) \\
&= \max_{p_t} \mathbb{E}\left[p_t d_t - h(x_t - d_t)^+\right] + \gamma\int_{d_t \leq x_t} \Phi_{t+1}(x_t - d_t)^+ dF(d_t) \\
&\quad +\gamma\int_{d_t < x_t} \Phi_{t+1}(x_t - d_t)^+ dF(d_t) - \gamma\int_{d_t < x_t} \Phi_{t+1}(x_t - d_t)^+ dF(d_t) \\
&\quad -\int_{d_t \geq x_t} \pi(d_t - x_t)^+ dF(d_t) - \int_{d_t \leq x_t} \pi(d_t - x_t)^+ dF(d_t) + \int_{d_t \leq x_t} \pi(d_t - x_t)^+ dF(d_t) \\
&= \max_{p_t} \mathbb{E}\left[p_t d_t - h(x_t - d_t)^+\right] + \gamma\mathbb{E}\left[\Phi_{t+1}(x_t - d_t)^+\right] - \gamma\int_{d_t < x_t} \Phi_{t+1}(0)dF(d_t) \\
&\quad -\mathbb{E}\left[\pi(d_t - x_t)^+\right] \\
&= \max_{p_t} \mathbb{E}\left[p_t d_t - h(x_t - d_t)^+ - \pi(-x_{it} + d_{it})^+ + \gamma\Phi_{t+1}(x_t - d_{it})^+\right] - \mathbb{E}\left[\gamma\int_{d_t < x_t} \Phi_{t+1}(0)dF(d_t)\right] \\
&= \Phi_t(x_t).
\end{aligned}$$

□

We conjecture that part (b) can be extended to the case of $m > 1$ with further conditions at the boundary when some components of x_t are equal to zero. The proof is more involved as we have to divide the domain of $(x_t - d_t)$, which is in $\mathbb{R}^m$, into different orthants and derive the corresponding integrals.