



Optimal Control

On the optimal control of partially observed inventory systems

Alain Bensoussan, Metin Çakanyıldırım, Suresh P. Sethi

*International Center for Decision and Risk Analysis, School of Management, P.O. Box 830688, SM 30,
University of Texas at Dallas, Richardson, TX 75083-0688, USA*

Received 1 July 2005; accepted 19 July 2005

Presented by Alain Bensoussan

Abstract

This Note introduces recent developments in the analysis of inventory systems with partial observations. The states of these systems are typically conditional distributions, which evolve in infinite dimensional spaces over time. Our analysis involves introducing unnormalized probabilities to transform nonlinear state transition equations to linear ones. With the linear equations, the existence of the optimal feedback policies are proved for two models where demand and inventory are partially observed. In a third model where the current inventory is not observed but a past inventory level is fully observed, a sufficient statistic is provided to serve as a state. The last model serves as an example where a partially observed model has a finite dimensional state. In that model, we also establish the optimality of the basestock policies, hence generalizing the corresponding classical models with full information. *To cite this article: A. Bensoussan et al., C. R. Acad. Sci. Paris, Ser. I 341 (2005).*

© 2005 Académie des sciences. Published by Elsevier SAS. All rights reserved.

Résumé

Contrôle optimal des stocks avec information partielle. On présente un certain nombre de modèles de systèmes de stocks avec information partielle. Ils sont formalisés comme des problèmes de contrôle où l'état est une probabilité conditionnelle, dans un espace de dimension infinie. On introduit des probabilités non normalisées, permettant de transformer des équations non linéaires en équations linéaires. On peut alors montrer l'existence de feedbacks optimaux pour deux modèles où la demande et le stock sont partiellement observables. Dans un troisième modèle, le stock n'est pas observé, mais un stock antérieur est observé. Une statistique exhaustive est obtenue, et l'état est de dimension finie. On établit l'optimalité des politiques « stock de base », généralisant les modèles classiques avec information complète. *Pour citer cet article : A. Bensoussan et al., C. R. Acad. Sci. Paris, Ser. I 341 (2005).*

© 2005 Académie des sciences. Published by Elsevier SAS. All rights reserved.

Version française abrégée

Bien que la gestion des stocks ait été l'objet d'une littérature considérable, la situation où l'information est partielle n'a pratiquement pas été abordée jusqu'à présent. On montre, sur un certain nombre de modèles, comment l'idée très féconde des probabilités conditionnelles non normalisées, introduites dans le filtrage non linéaire par Zakaï peut être adaptée et permet d'établir l'existence de politiques optimales. Les problèmes à traiter sont du type suivant

E-mail addresses: alain.bensoussan@utdallas.edu (A. Bensoussan), metin@utdallas.edu (M. Çakanyıldırım), sethi@utdallas.edu (S.P. Sethi).

$$\pi_{t+1}(x) = \mathbb{1}_{z_t=q_t} \frac{\int_{q_t}^{\infty} \pi_t(\xi) p(x|\xi) d\xi}{\int_{q_t}^{\infty} \pi_t(\xi) d\xi} + \mathbb{1}_{z_t < q_t} p(x|z_t), \quad (1)$$

où $p(x|\xi)$ est la probabilité de transition d'une chaîne de Markov, z_t est le processus observé, q_t le contrôle adapté à \mathcal{Z}_{t-1} et $\pi_t(x)$ est une probabilité conditionnelle qui représente l'état du système. L'équation de Bellman associée est donnée par

$$V(\pi) = \min_q \left\{ \int L(x, q) \pi(x) dx + \alpha V \left(\frac{\int_q^{\infty} p(\cdot|\xi) \pi(\xi) d\xi}{\int_q^{\infty} \pi(\xi) d\xi} \right) \int_q^{\infty} \pi(\xi) d\xi + \alpha \int_0^q V(p(\cdot|\xi)) \pi(\xi) d\xi \right\}. \quad (2)$$

La linéarisation est obtenue par les Éqs. (8)–(10) ci dessous. Un principe de contraction peut s'appliquer à (8) et permet de résoudre aussi (2), et d'obtenir un feedback optimal. Un problème beaucoup plus complexe est décrit par les Éqs. (19), (22), (23). Là aussi, la linéarisation simplifie considérablement, même si le principe de contraction ne s'applique plus. On prouve l'existence d'une solution maximale qui s'interprète comme la fonction valeur d'un problème de contrôle. On construit une suite de contrôles feedback formant une suite minimisante (le feedback optimal peut ne pas être atteint).

1. Introduction

Inventory control is one of the most important topics in operations research. A systematic analysis of inventory problems began with the development of the classical EOQ formula of Harris in 1913 [9] and a significant plateau was reached in 1958 by Arrow, Karlin and Scarf [1]. A critical assumption in this vast literature has been that the level of inventory at any given time is fully observed. Some of the most celebrated results such as the optimality of base-stock or (s, S) policies have been obtained under the assumption of full observation. Most of the well-known inventory policies are not only non-optimal, but are also not applicable in the partial observation environment. While one works with a finite dimensional state space in the full observation case, one usually has to deal with an infinite dimensional state space in the partial observation setting. In particular, in the full observation cases the inventory level is often the state variable, whereas with the partial observations, the system state is typically the conditional distribution of the inventory level given the partial observations.

It is the purpose of this paper to introduce some techniques, reminiscent of the Zakai equation [11] in stochastic control, for the analysis of inventory control with partial observations. The Zakai equation uses a transformation that changes the highly nonlinear Kushner equation [10] for evolution of probability distributions over time into a linear equation. While the transformation does not remove the infinite dimensionality, the linearity permits a number of important control problems with partial observations to be solved. On the other hand, some inventory models can admit a finite dimensional sufficient statistic. For a specific inventory model, this statistic turns out to be the known inventory level in an earlier period plus all of the inventory orders made since that period. When there exists a sufficient statistic, the well-known optimal ordering policies are more likely to remain valid.

In Sections 2 and 3, we introduce two separate inventory models, both of which benefit from the derivation of a Zakai-type equation. In Section 4, we provide an inventory model which admits a sufficient statistic. Sections 2–4 are respectively based on [5,3,4] and [6].

2. Partially observed demand due to censoring

Although many inventory models assume a static and a known demand distribution, this distribution in reality can change over time and so it is likely to be unknown. The distribution needs to be updated based on the demand observations. Ding et al. [7] study a multiperiod newsvendor model with Bayesian updating of demands. Specifically, the demand is observed *only* when it is less than the inventory level, i.e., unmet demand is *censored*. Consequently, the state of the system is the distribution of the demand characterized by a parameter, which is updated in each period based on the partial observation of the demand at that time. Ding et al. assume that the demands are independently and identically distributed. On the contrary, in this section we do not make any assumption on the demand distribution or correlation between the demands in different periods, except that it is a Markov process. Furthermore, we develop a Zakai-type equation [11] for the evolution of the conditional probability distribution of the demand over time. This

facilitates the analysis of the dynamic programming equation for the problem. We establish that the value function is the unique solution of the DP equation and that there exists an optimal feedback policy for the problem.

We now formally define our problem. Let (Ω, \mathcal{F}, P) be the probability space and let $t \geq 1$ be the indices for the periods. Let $D_t \geq 0$ denote the demand occurring at the beginning of period t . The demand is modelled by a Markov process with the transition probabilities given by $p(x, \xi) := P(D_{t+1} = x | D_t = \xi)$. The inventory available to satisfy the demand D_t , or a part thereof, is called q_t . Like [7], we assume that the excess inventory at the end of each period is salvaged and the unmet demand in a period is lost. Thus, q_t can also be thought as the order placed and delivered at the beginning of period t before the demand D_t arrives. Then the amount z_t of sales is given by

$$z_t := \min\{D_t, q_t\}. \tag{3}$$

When $D_t < q_t$, the demand is met and therefore observed. On the other hand, when $D_t \geq q_t$, the inventory is not sufficient to meet the demand in period t . Then, the amount of sales is only q_t and $D_t - q_t$ is the unmet demand. When the demand is not met, the unmet demand is not observed by the inventory manager (IM). Let \mathcal{Z}_t denote the sigma algebra $\mathcal{Z}_t := \sigma(\{z_1, \dots, z_t\})$. Since $\mathcal{Z}_t \subseteq \sigma(\{D_1, \dots, D_t\})$, we have a partially observed inventory model.

Let the function $L(D, q)$, which depends on the demand D and the available inventory level q , denote the one-period cost function:

$$L(D, q) = \begin{cases} cq - h(q - D) & \text{if } D \leq q \\ cq + b(D - q) & \text{if } q \leq D \end{cases}, \tag{4}$$

where h, c and b are, respectively, the salvage value per unit, the ordering cost per unit, and the shortage cost per unit. It is reasonable to assume that $0 \leq h < c < b$.

With the discount factor $0 < \alpha < 1$ and with \mathbf{q} defining the sequence of order quantities $\mathbf{q} = \{q_1, q_2, \dots\}$, with q_t adapted to \mathcal{Z}_{t-1} , our objective is to minimize

$$J(\mathbf{q}) := \sum_{t=1}^{\infty} \alpha^{t-1} \mathbb{E}L(D_t, q_t). \tag{5}$$

We take the initial distribution of D_1 as given to start the process. Note that q_1 is determined with certainty at the beginning of period 1. This completes the definition of our problem.

We start by studying the evolution of the demand distribution. Let π_t denote the conditional probability density function of D_t . Since π_{t+1} is updated from π_t and \mathcal{Z}_t , it is obvious that π_{t+1} is \mathcal{Z}_t -measurable. When an integral is taken over $[0, \infty)$, we suppress the limits to save on the notation.

Using probabilistic arguments, we obtain the random density measure π_{t+1} in terms of π_t :

$$\pi_{t+1}(x) = \mathbb{1}_{z_t=q_t} \frac{\int_{q_t}^{\infty} \pi_t(\xi) p(x|\xi) d\xi}{\int_{q_t}^{\infty} \pi_t(\xi) d\xi} + \mathbb{1}_{z_t < q_t} p(x|z_t). \tag{6}$$

With this distribution as the state variable, we now proceed to derive the dynamic programming equation for our problem.

2.1. Bellman equation with normalized and unnormalized probabilities

We begin with the expected total costs occurring in the current period t , i.e., $\mathbb{E}L(D_t, q_t) = \int L(x, q_t) \pi_t(x) dx$. Eq. (6) implies that the future costs should have two terms depending on whether the demand is observed. When the demand is not observed, the future costs depend on the first term on the right-hand side of (6). Otherwise, they depend on the second term. Summing up the current and future costs that result in by ordering q yields the terms inside the curly brackets in the Bellman equation below:

$$V(\pi) = \min_q \left\{ \int L(x, q) \pi(x) dx + \alpha V \left(\frac{\int_q^{\infty} p(\cdot|\xi) \pi(\xi) d\xi}{\int_q^{\infty} \pi(\xi) d\xi} \right) \int_q^{\infty} \pi(\xi) d\xi + \alpha \int_0^q V(p(\cdot|\xi)) \pi(\xi) d\xi \right\}. \tag{7}$$

This equation yields the optimal q as a function of the current knowledge summarized in π about the demand. The state transition equations can be linearized as we illustrate next.

We begin with (6) and develop what is known as the unnormalized probability. Its dynamics is similar to the Zakai equation obtained in the filtering literature [11].

Taking a cue from (6), we define $\rho_t(x)$ by the recursive linear equation

$$\rho_{t+1}(x) = \mathbb{1}_{z_t=q_t} \int_{q_t}^{\infty} p(x|\xi)\rho_t(\xi) d\xi + \mathbb{1}_{z_t < q_t} p(x|z_t) \quad \text{for } t \geq 1, \tag{8}$$

with the initial value $\rho_1(x) = \rho(x) = \pi(x)$. Eq. (8) corresponds to the Zakai equation for systems with diffusions in [11] and [2]. Also set $\lambda_t := \int \rho_t(x) dx$. Then we have $\lambda_1 = 1$ and $\lambda_{t+1} = \mathbb{1}_{z_t=q_t} \int_{q_t}^{\infty} \rho_t(\xi) d\xi + \mathbb{1}_{z_t < q_t}$ for $t \geq 1$, which follows directly from integrating the left-hand side of (8) over $[0, \infty)$. Moreover, $\rho_t(x) = \pi_t(x)\lambda_t$, where λ_t is a weighting factor which allows us to recover normalized probabilities.

We define, for any positive L^1 function ρ ,

$$W(\rho) := V\left(\frac{\rho}{\int \rho(x) dx}\right) \int \rho(x) dx. \tag{9}$$

We use (7) and (9) to obtain a recursive equation for W :

$$W(\rho) = \min_q \left\{ \int L(x, y)\rho(x) dx + \alpha W\left(\int_q^{\infty} p(\cdot|\xi)\rho(\xi) d\xi\right) + \alpha \int_0^q W(p(\cdot|\xi))\rho(\xi) d\xi \right\}. \tag{10}$$

From the linearity in ρ of the first term inside the minimum above, we can see that $W(0) = 0$ and W is homogenous of degree 1, i.e., $W(a\rho) = aW(\rho)$ for $a > 0$. The right-hand side in the Bellman equation (10) has linear growth, which facilitates the derivation of the existence results below.

2.2. Existence of an optimal feedback solution and monotone value functions

First we show that there exists a unique solution of the DP equation. For this, we define some normed linear spaces. Let $\mathcal{H} := \{\rho \in L^1(\mathfrak{R}^+): \int x|\rho(x)| dx < \infty\}$, where $L^1(\mathfrak{R}^+)$ is the space of integrable functions whose domain is the set of nonnegative real numbers. \mathcal{H} is a subspace of $L^1(\mathfrak{R}^+)$. Since we are interested in nonnegative unnormalized probabilities, we define $\mathcal{H}^+ = \{\rho \in \mathcal{H} \mid \rho \geq 0\}$, a closed subset of \mathcal{H} with the norm $\|\rho\| = \int |\rho(x)| dx + \int x|\rho(x)| dx$. Let us also define the space $\mathcal{B} = \{\phi(\rho): \mathcal{H}^+ \rightarrow \mathfrak{R} \mid \sup_{x>0} |\phi(\rho)|/\|\rho\| < \infty\}$ with the norm $\|\phi\|_{\mathcal{B}} = \sup_{\rho \in \mathcal{H}^+} |\phi(\rho)|/\|\rho\|$.

We assume that for any $\mu \in \mathcal{B}$, $\int x \int p(x|\xi)\mu(\xi) d\xi dx \leq \bar{c} \int \xi \mu(\xi) d\xi$ with $\alpha \bar{c} < 1$. We also define the map $T(W)$ as

$$T(W)(\rho) := \min_q \left\{ \int L(x, y)\rho(x) dx + \alpha W\left(\int_q^{\infty} p(\cdot|\xi)\rho(\xi) d\xi\right) + \alpha \int_0^q W(p(\cdot|\xi))\rho(\xi) d\xi \right\}. \tag{11}$$

The next theorem summarizes the results.

Theorem 2.1. *If there exists a solution $W(\rho)$ of (10), it is in \mathcal{B} . Furthermore, $\|T(W) - T(\tilde{W})\|_{\mathcal{B}} \leq \alpha \max\{1, \bar{c}\} \times \|W - \tilde{W}\|_{\mathcal{B}}$. Thus, there exists one and only one solution $W(\rho)$ of the Bellman equation (10). Moreover, $W(\rho)$ is continuous at each $\rho \in \mathcal{H}^+$ and there exists an optimal feedback policy.*

Using a value iteration based on (10), we can deduce that $W(\rho)$ is monotone increasing and concave in ρ . More interestingly, the total discounted cost $V(\pi)$ is smaller when the demand is believed to be smaller in the hazard rate order. We also say that linear operator $\mathcal{P}(\pi)$ defined by $\mathcal{P}(\pi) := \int p(\cdot|\xi)\pi(\xi) d\xi$ is hazard rate order preserving if $\mathcal{P}(\pi) \leq \mathcal{P}(\pi')$ holds for all π and π' with $\pi \leq \pi'$.

Theorem 2.2. *$V(\pi) \leq V(\pi')$ if π is smaller than π' in the hazard rate order and $\mathcal{P}(\pi)$ is hazard rate order preserving.*

3. Partially observed inventory due to zero-balance walk

In this section, the demand has a known distribution but the inventory level I_t at the beginning of period t is given by

$$I_{t+1} = (I_t + q_t - D_t)^+ \quad \text{for } t \geq 1. \tag{12}$$

In every period t , demand D_t is met, to the extent possible, from the on-hand stock $I_t + q_t$. When the demand is met entirely, inventory holding costs apply to the remaining inventory. Otherwise, there are lost sales costs. We assume demand D_t to be i.i.d. A generic demand is denoted by D , which is i.i.d. with each D_t . Let f denote the density and F denote the cumulative distribution of D . Let $\bar{F} = 1 - F$.

In classical inventory settings, the inventory level I_t is observed, and is used to determine the order quantity q_t in period t . In this section, the IM does not automatically observe the inventory level due to transaction errors, misplaced inventories, spoilage or production yield. The inventory level is observed only when there is no physical inventory. When there is inventory, only the event that the inventory is positive is observed. We continue to use z_t to denote the observed signal, which in this context is

$$z_t := \mathbb{1}_{I_t=0}, \quad t \geq 0. \tag{13}$$

The signal z_t is a discrete-time Markov Chain with the state space $\{0, 1\}$: 1 means an empty inventory shelf and 0 means a nonempty shelf. This observation process mimics what is known as ‘zero-balance walk’ [8] at some companies where employees walk around the shelves to identify the stocked-out items and verify the inventory levels for those items.

The order q_t is adapted to $\mathcal{Z}_t := \sigma(\{z_j : 1 \leq j \leq t\})$. Clearly $\mathcal{Z}_t \subset \mathcal{F}_t := \sigma(\{I_j : 1 \leq j \leq t\})$. Given a stationary cost function $c(I_t, q_t)$ that depends on the inventory level I_t and the order size q_t , and with \mathbf{q} defining the admissible order quantities, the total discounted cost is defined by

$$J(\zeta, \pi, \mathbf{q}) := \sum_{t=1}^{\infty} \alpha^{t-1} \mathbb{E}c(I_t, q_t). \tag{14}$$

The initial conditions are a pair $(\zeta, \pi(x))$, where ζ is 1 or 0. If ζ is 1, then $I_1 = 0$. If ζ is 0, then $I_1 > 0$ and $\pi(\cdot)$ is the probability distribution of I_1 . We look for q_t , adapted to $\mathcal{Z}_t, t \geq 0$, to minimize $J(\zeta, \pi, \mathbf{q})$.

3.1. Evolution of state probabilities

We now develop the conditional probability density $\pi_t(\cdot)$ of I_t given \mathcal{Z}_{t-1} and $I_t > 0$. Derivation of π_t in this case is more involved than in Section 2, so we furnish more details. By definition, $\int_0^x \pi_t(y) dy = P(I_t \leq x | \mathcal{Z}_{t-1}, I_t > 0)$. Since the event $[I_t = 0]$ is observable, conditional probabilities are needed only when $I_t > 0$. We can obtain

$$\int_0^{\infty} \phi(x) \pi_t(x) dx = \mathbb{E}[\phi(I_t) | \mathcal{Z}_{t-1}, I_t > 0] = \frac{\mathbb{E}[\phi(I_t) \mathbb{1}_{I_t>0} | \mathcal{Z}_{t-1}]}{\mathbb{E}[\mathbb{1}_{I_t>0} | \mathcal{Z}_{t-1}]}, \tag{15}$$

where ϕ is an arbitrary test function. In order to obtain a recursive expression for π_t in terms of π_{t-1} , we begin with expressing $\mathbb{E}(\phi(I_t) | \mathcal{Z}_t)$ in terms of conditional expectations with respect to \mathcal{Z}_{t-1} in the next two lemmas.

Lemma 3.1.

$$\mathbb{E}(\phi(I_t) | \mathcal{Z}_t) = \mathbb{1}_{I_t=0} \phi(0) + \mathbb{1}_{I_t>0} \frac{\mathbb{E}(\phi(I_t) \mathbb{1}_{I_t>0} | \mathcal{Z}_{t-1})}{P(I_t > 0 | \mathcal{Z}_{t-1})} = \mathbb{1}_{I_t=0} \phi(0) + \mathbb{1}_{I_t>0} \mathbb{E}(\phi(I_t) | \mathcal{Z}_{t-1}, I_t > 0). \tag{16}$$

Lemma 3.2.

$$\begin{aligned} \mathbb{E}(\phi(I_t) | \mathcal{Z}_t) \mathbb{1}_{I_t>0} &= \mathbb{1}_{I_{t-1}=0} \frac{\int_0^{\infty} \phi(z) f(q_{t-1} - z) \mathbb{1}_{q_{t-1} \geq z} dz}{F(q_{t-1})} \\ &+ \mathbb{1}_{I_{t-1}>0} \frac{\int_0^{\infty} \phi(z) \int_{(z-q_{t-1})^+}^{\infty} f(y + q_{t-1} - z) \pi_{t-1}(y) dy dz}{\int_0^{\infty} F(y + q_{t-1}) \pi_{t-1}(y) dy}. \end{aligned} \tag{17}$$

Going back to the conditional probability π_t , we may write $\mathbb{E}(\phi(I_t)|\mathcal{Z}_t) = \mathbb{1}_{I_t=0}\phi(0) + \mathbb{1}_{I_t>0} \int_0^\infty \phi(z)\pi_t(z) dz$. Equating the second term on the right-hand side of this equality to (17), we get

$$\pi_t(x) = \mathbb{1}_{I_{t-1}=0} \left\{ \frac{f(q_{t-1} - x)\mathbb{1}_{x \leq q_{t-1}}}{F(q_{t-1})} \right\} + \mathbb{1}_{I_{t-1}>0} \left\{ \frac{\int_{(x-q_{t-1})^+}^\infty f(y + q_{t-1} - x)\pi_{t-1}(y) dy}{\int_0^\infty F(y + q_{t-1})\pi_{t-1}(y) dy} \right\}. \tag{18}$$

Equivalently, the conditional probability evolves according to the highly nonlinear equation

$$\pi_t(x) = z_{t-1} \frac{f(q_{t-1} - x)\mathbb{1}_{x < q_{t-1}}}{F(q_{t-1})} + (1 - z_{t-1}) \frac{\int_{(x-q_{t-1})^+}^\infty f(y + q_{t-1} - x)\pi_{t-1}(y) dy}{\int_0^\infty F(q_{t-1} + y)\pi_{t-1}(y) dy},$$

$$t \geq 2, \pi_1(x) = \pi(x), \tag{19}$$

which corresponds to the Kushner equation [10] in our inventory context.

For linearizing (19), we set $\rho_t(x) := \lambda_t \pi_t(x)$. On account of this weighting, $\rho_t(x)$ becomes the unnormalized probability and it evolves according to the linear equation

$$\rho_t(x) = z_{t-1} f(q_{t-1} - x)\mathbb{1}_{x < q_{t-1}} + (1 - z_{t-1}) \int_{(x-q_{t-1})^+}^\infty f(y + q_{t-1} - x)\rho_{t-1}(y) dy, \quad \rho_1(x) = \pi(x). \tag{20}$$

By integrating both sides of $\rho_t(x) = \lambda_t \pi_t(x)$, we get $\lambda_t = z_{t-1} F(q_{t-1}) + (1 - z_{t-1}) \lambda_{t-1} \int F(q_{t-1} + y)\pi_{t-1}(y) dy$ which is a linear equation defining λ_t recursively starting with $\lambda_1 = 1$. However, note that λ_t depends on π_{t-1} . The normalized probabilities can easily be computed from the unnormalized probabilities by $\pi_t(x) = \rho_t(x) / \int \rho_t(x) dx$.

The evolution equations can be written in the operator form in the space \mathcal{H} and its dual space $\mathcal{H}_* = \{\phi: \sup_{x>0} |\phi(x)|/(1+x) < \infty\}$. Furthermore, $\langle \rho, \phi \rangle = \int_0^\infty \rho(x)\phi(x) dx$ for $\rho \in \mathcal{H}, \phi \in \mathcal{H}_*$.

For any scalar $q > 0$, define the linear operator Φ from \mathcal{H} to \mathcal{H} as $\Phi(q, \rho)(x) = \int_{(x-q)^+}^\infty f(y + q - x)\rho(y) dy$. Note that $\Phi(q, \delta)(x) = f(q - x)\mathbb{1}_{x < q}$ so $\Phi(0, \delta)(x) = 0$ for the Dirac delta function δ . Define the nonlinear operator Ψ as

$$\Psi(q, \rho) = \frac{\Phi(q, \rho)}{\langle \Phi(q, \rho), 1 \rangle}. \tag{21}$$

With these notations and the initial conditions $\pi_1 = \rho_1 = \pi$, we can write (19) and (20) in the operator form: $\pi_t = z_{t-1}\Psi(q_t, \delta) + (1 - z_{t-1})\Psi(q_t, \pi_{t-1})$ and $\rho_t = z_{t-1}\Phi(q_t, \delta) + (1 - z_{t-1})\Phi(q_t, \rho_{t-1})$. We emphasize that the second equality is linear, while the first is nonlinear.

3.2. The Bellman equation

We write $\pi_t[\mathbf{q}]$ to emphasize the dependence of the state π_t on the control policy. We assume that $c(I_t, q_t)$ has linear growth in I_t for every fixed q_t , i.e., $c(\cdot, q_t) \in \mathcal{H}_*$. The cost function is:

$$J(\zeta, \pi, \mathbf{q}) = \sum_{t=1}^\infty \alpha^{t-1} \mathbb{E}[\mathbb{E}[c(\cdot, q_t)|\mathcal{Z}_t]] = \sum_{t=1}^\infty \alpha^{t-1} \mathbb{E}\{z_t c(0, q_t) + (1 - z_t)c(\cdot, q_t), \pi_t[\mathbf{q}]\}$$

where $\pi_t[\mathbf{q}]$ is given by (19). Recall that the initial conditions $\zeta_1 = \zeta \in \{0, 1\}$ and $\pi_1 = \pi$ are given. We now study only the discounted infinite horizon cost, so the time index t is suppressed. We define the value function $V(\zeta, \pi) := \inf_{\mathbf{q}} J(\zeta, \pi, \mathbf{q})$.

If we write $v := V(1, \pi)$ which, in fact, is not dependent on π , and $V(\pi) := V(0, \pi)$, then we obtain the following system:

$$V(\pi) = \inf_q \left\{ \langle c(\cdot, q), \pi \rangle + \alpha v \int \bar{F}(y + q)\pi(y) dy + \alpha V(\Psi(q, \pi))F(y + q)\pi(y) dy \right\}, \tag{22}$$

$$v = \inf_q \{c(0, q) + \alpha v \bar{F}(q) + \alpha V(\Psi(q, \delta))F(q)\}. \tag{23}$$

The study of the system in (22), (23) simplifies considerably when working with the unnormalized probability $\rho \in \mathcal{H}^+$. Towards that end, we define a new value function $Z(\cdot): Z(\rho) := V(\rho/\lambda)\lambda$ and $\lambda := \int \rho(x) dx$. We obtain the following new system of equations:

$$Z(\rho) = \inf_q \left\{ c(\cdot, q), \rho \right\} + \alpha v \int \bar{F}(y + q)\rho(y) dy + \alpha Z(\Phi(q, \rho)), \tag{24}$$

$$v = \inf_q \left\{ c(0, q) + \alpha v \bar{F}(q) + \alpha Z(\Phi(q, \delta)) \right\}. \tag{25}$$

The pair $(v, Z(\rho))$ is the solution of (24), (25). We have $Z(0) = 0$ and Z is homogenous of degree 1. Unlike the operator Ψ , Φ is a linear operator.

3.3. Existence of a solution to the Bellman equation

For the existence results, we bound the single period cost. Suppose that positive constants c, c_0, c_1 , and h are such that $cq < c(x, q) \leq c_0 + c_1q + hx$ for $x \geq 0$, where c_0 can be interpreted as the maximum expected lost sales cost that can be incurred in a period. Indeed, we set $c_0 = c(0, 0)$.

We need some short-hand notation: Define the function $K : \mathfrak{R} \times \mathcal{H} \times \mathfrak{R} \times \mathcal{B} \rightarrow \mathfrak{R}$ as $K(q, \rho; v, Z) := \langle c(\cdot, q), \rho(\cdot) \rangle + \alpha v \int \bar{F}(y + q)\rho(y) dy + \alpha Z(\Phi(q, \rho))$. For $\rho = \delta$, $K(q, \delta; v, Z) = c(0, q) + \alpha v \bar{F}(q) + \alpha Z(\Phi(q, \delta))$. Define the map $T : \mathfrak{R} \times \mathcal{B} \rightarrow \mathfrak{R} \times \mathcal{B}$ as $T(v; Z(\rho)) := (\inf_q K(q, \delta; v, Z); \inf_q K(q, \rho; v, Z))$. Define $(Z_0(\rho), v_0)$ as the value function when $q = 0$. Then, we have respectively $\langle c(\cdot, 0), \rho \rangle + \alpha v_0 \int \bar{F}(y)\rho(y) dy + \alpha Z_0(\Phi(0, \rho)) = Z_0(\rho)$ and $v_0 = c_0 + \alpha v_0$.

Lemma 3.3. Z_0 exists and is uniquely defined in \mathcal{B} .

Note that $T(v_0; Z_0(\rho)) \leq (v_0; Z_0(\rho))$ and $T(v; Z(\rho)) \leq (v; Z(\rho))$. These inspire a value iteration scheme which yields (\bar{v}, \bar{Z}) .

Theorem 3.4. (\bar{v}, \bar{Z}) is a solution of (24), (25) and $0 \leq \bar{v} \leq v_0, 0 \leq \bar{Z}(\rho) \leq Z_0(\rho)$. Moreover, it is the maximal solution satisfying $(v; Z) = T(v; Z)$. Also $\bar{Z}(\pi) = \inf_{\mathbf{q}} J(0, \pi, \mathbf{q})$ and $\bar{v} = \inf_{\mathbf{q}} J(1, \pi, \mathbf{q})$.

Owing to the last statement, $\bar{Z}(\pi)$ and \bar{v} are interpreted as the infima of the costs. This, however, does not imply the existence of a feedback policy unless the order quantity q is bounded. When the optimal q is bounded in terms of ρ , the bound can be arbitrarily large as $\rho \rightarrow 0$. Because of this, we choose to assume a bound on q in addition to the cost bounds in the previous section. Let the maximum production capacity be m and let the corresponding Z and v be denoted by Z^m and v^m . Then (24), (25) is written as

$$Z^m(\rho) = \inf_{q \leq m} \left\{ \langle c(\cdot, q), \rho \rangle + \alpha v^m \int \bar{F}(y + q)\rho(y) dy + \alpha Z^m(\Phi(q, \rho)) \right\}, \tag{26}$$

$$v^m = \inf_{q \leq m} \left\{ c(0, q) + \alpha v^m \bar{F}(q) + \alpha Z^m(\Phi(q, \delta)) \right\}.$$

We can check that constants A^m and B^m exist such that $|Z^m(\rho) - Z^m(\rho')| \leq A^m \int |\rho(y) - \rho'(y)| dy + B^m \int y|\rho(y) - \rho'(y)| dy$ for any two $\rho, \rho' \in \mathcal{H}$. Therefore, Z^m is Lipschitz continuous on \mathcal{H} . This additional smoothness property allows us to establish the uniqueness of a solution to the system in (26) in the absence of a contraction property on T as in Theorem 2.1. The next result validates the monotone iterative process. That is, (v^m, Z^m) minimizes the total discounted cost.

Proposition 3.5. The solution (v^m, Z^m) of (26) is the minimum total discounted cost, i.e., $Z^m(\pi) = \inf_{\mathbf{q}: q_t \leq m} J(0, \pi, \mathbf{q})$ and $v^m = \inf_{\mathbf{q}: q_t \leq m} J(1, \delta, \mathbf{q})$.

Since $Z^m(\pi)$ and v^m are defined as a solution of (26) and they are given by the infima in Proposition 3.5, both $Z^m(\pi)$ and v^m are unique. As m increases, we have $\inf_{\mathbf{q}: q_t \leq m} J(0, \pi, \mathbf{q}) \downarrow \inf_{\mathbf{q}} J(0, \pi, \mathbf{q})$ and $\inf_{\mathbf{q}: q_t \leq m} J(1, \pi, \mathbf{q}) \downarrow \inf_{\mathbf{q}} J(1, \pi, \mathbf{q})$. These imply the last statement in Theorem 3.4.

4. Partially observed inventory due to information delays

Here we consider partial observability arising from information delays. The current inventory level is not observed by the IM. Instead, he observes the exact inventory level of a prior period. This model enables us to find a sufficient

statistic. It is important to allow for backordering of the immediately unmet demand to obtain the sufficient statistic. Hence, $I_{t+1} = I_t + q_t - D_t$ for $t \geq 1$. We suppose that the demands D_t are i.i.d. and use D^i to denote the sum of i demands. To initialize the inventory to start its evolution, we assume that I_1 is known at $t = 1$.

Let $\theta \geq 0$ be the amount of the delay and $\beta_t := \max\{t - \theta, 1\}$. The observed signal $z_t = I_{\beta_t}$ at time t denotes the last inventory level observed in period t . The signal observed at the beginning of period t generates the sigma algebra $\mathcal{Z}_t = \sigma(\{z_1; z_2; \dots; z_t\})$. However, the underlying and unobserved state at time t is I_t and generates the sigma algebra \mathcal{F}_t . Clearly, $\mathcal{Z}_t \subseteq \mathcal{F}_t$. The expression for the total discounted cost is $J(\mathbf{q}) := \sum_{t=1}^T \alpha^{t-1} \mathbb{E}c(I_t, q_t)$. The objective is to minimize $J(\mathbf{q})$ over the class of admissible orders \mathbf{q} adapted to \mathcal{Z}_t .

We define the *reference inventory position* as $x_t := z_t + \sum_{j=\beta_t}^{t-1} q_j$, which can be obtained from the information in \mathcal{Z}_t . The evolution equation for x_t is $x_{t+1} = x_t + q_t - \mathbb{1}_{\theta \leq t-1} D_{t-\theta}$. Moreover, the unobserved I_t can be expressed in terms of the observed x_t , i.e., $I_t = x_t - \sum_{j=\beta_t}^{t-1} D_j$.

By a sequence of conditioning arguments, we obtain $\mathbb{E}c(I_t, q_t) = \mathbb{E}c^{t-\beta_t}(x_t, q_t)$, where $c^i(x, q) := \mathbb{E}c(x - D^i, q)$. Thus, the single-period costs can be computed in terms of the reference inventory position and the objective becomes $\sum_{t=1}^T \alpha^{t-1} \mathbb{E}c^{t-\beta_t}(x_t, q_t)$. This shows that x_t is a sufficient statistic for our problem. Moreover, we have a problem in the form of a standard inventory problem with the state x_t , with the single-period nonstationary cost function $c^{t-\beta_t}(x_t, q_t)$, and with the sigma algebra \mathcal{Z}_t . Therefore, we can directly write the associated DP equation starting with $V_{T+1}(x) = 0$:

$$V_t(x) = \begin{cases} \inf_{q \geq 0} c^{t-1}(x, q) + \alpha V_{t+1}(x + q) & \text{for } t \leq \theta, \\ \inf_{q \geq 0} c^\theta(x, q) + \alpha \mathbb{E}V_{t+1}(x + q - D) & \text{for } t \geq \theta + 1. \end{cases}$$

The one-period cost $c^i(x, q)$ can typically be written as $cq + \mathbb{E}H(x - D^i)$, where the first term $c > 0$ (second term $H(\cdot)$) represents the ordering (inventory) cost. Suppose that the inventory cost is convex, then a base-stock type policy is optimal:

$$q_t^*(x) = (u_t^* - x)^+, \quad \text{where } u_t^* := \begin{cases} \arg \min_u cu + V_{t+1}(u) & \text{for } t \leq \theta, \\ \arg \min_u cu + \mathbb{E}V_{t+1}(u - D) & \text{for } t \geq \theta + 1. \end{cases}$$

We can extend our results to the case of delays given to be a Markov process. In that case, the reference inventory position along with the value of the latest delay observation and the age of this observation are sufficient statistics for finding the optimal order quantities. These sufficient statistics are of dimension three whereas the state space of the associated DP is in general infinite dimensional. The optimal ordering policy remains to be of base stock type with respect to the reference inventory position. However, the base stock levels depend now on the value of the latest delay observation and the age of this observation.

References

- [1] K.J. Arrow, S. Karlin, H.E. Scarf, *Studies in the Mathematical Theory of Inventory and Production*, Stanford University Press, California, 1958.
- [2] A. Bensoussan, *Stochastic Control of Partially Observable Systems*, Cambridge University Press, Cambridge, UK, 1992.
- [3] A. Bensoussan, M. Çakanyıldırım, S.P. Sethi, Partially observed inventory systems: The case of zero balance walk, Working paper SOM 200548, School of Management, University of Texas at Dallas, TX, 2004.
- [4] A. Bensoussan, M. Çakanyıldırım, S.P. Sethi, Optimality of base stock and (s, S) policies for inventory problems with information delays, Working paper SOM 200547, School of Management, University of Texas at Dallas, TX, 2005.
- [5] A. Bensoussan, M. Çakanyıldırım, S.P. Sethi, A multiperiod newsvendor problem with partially observed demand, Working paper SOM 200550, School of Management, University of Texas at Dallas, TX, 2005.
- [6] A. Bensoussan, M. Çakanyıldırım, S.P. Sethi, Optimal ordering policies for inventory problems with dynamic information delays, Working paper 200560, School of Management, University of Texas at Dallas, TX, 2005.
- [7] X. Ding, M.L. Puterman, A. Bisi, The censored newsvendor and the optimal acquisition of information, *Oper. Res.* 50 (3) (2002) 517–527.
- [8] M.L. Fisher, A. Raman, A.S. McClelland, Rocket science retailing is almost here: Are you ready?, *Harvard Business Review* (July–August 2000) 115–124.
- [9] F.W. Harris, Operations and cost, in: *Factory Management Series*, A.-W. Shaw Company, Chicago, IL, 1913, pp. 48–52 (Chapter IV).
- [10] H.J. Kushner, Dynamic equations for nonlinear filtering, *J. Differential Equations* 3 (1967) 179–190.
- [11] M. Zakai, On the optimal filtering of diffusion processes, *Z. Wahrsch. Verw. Gebiete* 11 (1969) 230–243.