European Journal of Operational Research 204 (2010) 421–438

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Production, Manufacturing and Logistics

Optimal policies for inventory systems with finite capacity and partially observed Markov-modulated demand and supply processes

Kenan Arifoğlu, Süleyman Özekici *

Department of Industrial Engineering, Koç University, 34450 Sarıyer-İstanbul, Turkey

ARTICLE INFO

Article history:
Received 4 February 2008
Accepted 28 October 2009
Available online 6 November 2009

Keywords:
Random yield
Fixed capacity
Random environment
Modified inflated base-stock policy
Dynamic programming
Sufficient statistics
POMDP

ABSTRACT

We analyze a single-item periodic-review inventory system with random yield and finite capacity operating in a random environment. The primary objective is to extend the model of Gallego and Hu (2004) to the more general case when the environment is only partially observable. Although our analysis is specific to inventory systems, it can also be applied to production systems by replacing the fixed capacity supplier with a fixed capacity producer. Using sufficient statistics, we consider single-period, multiple-period and infinite-period problems to show that a state-dependent modified inflated base-stock policy is optimal. Moreover, we show that the multiple-period cost converges to the infinite-period cost as the length of the planning horizon increases.

1. Introduction

We consider a single-item periodic-review inventory system where the yield is random in the sense that only a random proportion of the order quantity is received. Moreover, the supplier or producer has fixed or limited capacity, and the system operates in a partially observed random environment. An inventory system is not isolated from the environment in which it operates. Like demand, supply and costs are quite sensitive to external (social, economical, political and/or natural) factors. For this purpose, as in similar studies concerning random environments, we model the outside world as a discrete-state, time-dependent Markov chain. However, we also suppose that the inventory manager (IM) cannot directly observe the state of this Markov chain (real environment). Instead, the IM creates his own (observed) environment based on his observations. We further assume that these observations are incomplete because the environment is defined by a changing set of known and unknown factors that are subject to fluctuations; we assume as a result that it is not possible either to identify the complete set of factors, or to follow their fluctuations. Hence, one way or another, information about the true state of the environment is neither complete nor perfect, and the IM has to base his decisions on such incomplete and imperfect information. However, unlike the majority of papers in the literature, we model the outside environment by a hidden process that is not perfectly observable by the IM. In particular, we suppose that the IM observes a process that depends on the real environmental process in a probabilistic fashion. Although the hidden process has a Markovian structure, the observed process does not necessarily satisfy the Markov property. Therefore, the observed environment gives only partial information regarding the real (unobserved) environment. In the literature, this type of a decision process in which there is an internal Markov process that can only be partially observed through another process, is called a partially observed Markov decision process (POMDP).

In this paper, using POMDP, we analyze a periodic-review inventory model with random yield and fixed capacity in order to identify the optimal policy. The model enables us to consider the effect of the randomly changing environment on supply, demand and cost parameters, which in turn influence the ordering decisions of the IM. There are several instances when uncertainty as to the true state of the environment can cause the IM to react with inventory policies or actions that are different from situations when the knowledge of the environment is perfect. In fact one could argue that in most cases the IM's knowledge of the environmental factor that impact demand and supply processes, and even the cost parameters, is imperfect. This paper attempts to model precisely such situations, and compares the optimal inventory policies to those when the true state of the environment is revealed to the IM with absolute certainty.

In the inventory system that we consider, the IM of a retailer places inventory replenishment orders with a supplier who has fixed capacity; hence independent of the quantity requested by

---

* Corresponding author.
E-mail address: sozekici@ku.edu.tr (S. Özekici).

0377-2217/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2009.10.029

the IM, the supplier can only ship up to a specified amount. Moreover, the supply is subject to yield randomness. Hence, a random proportion of the delivery is received at the store due to possible problems in production, transportation and other reasons. The supplier may even be unavailable in which case nothing is delivered. In effect, the amount received towards a specific order placed by the IM is a random quantity.

Detailed discussions on both periodic-review and continuous-review inventory models with random supply can be found in [24]. Random supply models in stationary environments that are closely related to ours are analyzed in [9,12,6,23]. We can classify current random supply models in the literature in two groups, namely: random yield and random capacity. Random yield models do not necessarily lead to nice characterizations on the optimal policy. Henig and Gerchak [12] show that a "nonorder-up-to" policy is optimal if there is random proportional yield. In this policy structure, there is a predetermined inventory level under which an order is always given; however, unlike the base-stock policy, this order does not necessarily bring the inventory level up to a constant base-stock level. Instead, it increases the inventory level above the predetermined inventory level. Because of this particular structure, Zipkin [25] calls it an "inflated" base-stock policy. Random capacity models, on the other hand, often lead to optimality of well-known ordering structures. Federgruen and Zipkin [9] analyze an inventory problem with fixed production capacity and show that a "modified" base-stock policy is optimal. Using this policy, it is optimal to order up to the base-stock if the fixed capacity is sufficient; otherwise, one should order as much as the limited capacity. A model considering random capacity is given in [6] who show that a base-stock policy is optimal. Wang and Gerchak [23] combine random proportional yield and random capacity and show the optimality of "inflated" base-stock policies.

This paper models random supply, random environment and POMDP in the context of inventory management. Earlier papers considering the random environment focus only on the effect of the changing environment on demand. Among many others, Karlin and Fabens [15], Iglehart and Karlin [14], Song and Zipkin [20], and Sethi and Cheng [18] provide examples along this direction. For example, Song and Zipkin [20] modeled the outside world as a Markov chain and assumed that demand in successive periods are dependent on the state of this Markov chain. In a later study, Sethi and Cheng [18] also incorporated fluctuating demand environment into their model using Markov chains, and found the most general setting under which an environment-dependent (s, S) policy is optimal.

Later, starting with mid-1990s, researchers introduced models where not only the demand, but the supply is also affected by the fluctuating environment. For example, Song and Zipkin [21] analyzed the effect of the fluctuating environment on supply by using a Markov chain approach. They show that the optimal policy has the same structure as in standard models, but that its parameters change dynamically to reflect current supply conditions. Another paper which considers the possible effect of fluctuating environment on supply as well as on demand and cost parameters is Ö zekici and Parlar [17]. They assume that the supplier is either available or unavailable when the order is placed so that the order is either totally satisfied or in the other extreme, remains entirely unfulfilled. In addition, Erdem and Özekici [8] extend Özekici and Parlar [17] and assume that the supplier is always available, but that its capacity is random and in turn dependent on the state of the environment. Furthermore, Arifoğlu and Özekici [1] extends both Özekici and Parlar [17] and Erdem and Özekici [8] by considering a more general framework in which there is a supplier with random capacity and a transporter with random availability. As a result of their analysis, they show that an environment-dependent base-stock policy is optimal. Finally, Gallego and Hu [11] analyze inventory problems with random yield and finite capacity in a random environment. They assume that the capacity of the supplier is finite and the retailer receives a random proportion of the amount produced. Moreover, they distinguish between demand and supply environments by using two Markov chains: one for the demand environment and one for the supply environment. As a result of their analysis, they show that environment-dependent modified inflated base-stock policy is optimal.

Our primary objective is to extend the model and results in [11] by considering the case where there is imperfect information and the environment is only partially observed. Furthermore, we assume that all costs are also dependent on the observed environment. Among all sources of randomness in an inventory model, demand is by far the most important one. Precision in estimating future demand is extremely important for the IM since it has a huge impact on inventory levels and costs. The state of an economy, for example, is a factor that affects demand, among others. The demand levels go down probabilistically in a recession, while just the opposite is true in an expanding economy. Managerial decisions on inventory control depend very much on this classification. However, there is often imperfection in determining this state and experts look at certain observed economic indicators to make judgements. This approach imposes another source of randomness between the observed and real states. However, reduced uncertainty in estimating the demand for real states can erase the effect of this additional randomness with better information. Technically speaking, the formulation with imperfect information requires a more complicated model and challenging computations. However, the structure of the optimal policies remain intact under reasonable conditions. This structure now depends, of course, on the information collected in time.

POMDPs have a wide range of application areas such as machine maintenance and replacement, human learning and instruction, medical diagnosis and decision making, and search for moving objects. We refer our reader to Smallwood and Sondik [19], and Monahan [16] for a detailed discussion on these processes. Despite extensive research on POMDPs, they have not been as widely employed in inventory models. Treharne and Sox [22] is an exception where they assume that the demand environment is randomly depicted by a Markov chain; however, the state of this Markov chain is only partially observed. They do not consider the supply side and assume that the capacity of the supplier is infinite. They show that state-dependent base-stock policy is optimal where the state is defined as the inventory position and the conditional distribution of true environmental state. Another paper which applies POMDPs is Bensoussan et al. [3] who study three different models: information delay, filtered newsvendor and zero balance walk. The one that is related to our model is the information delay model. In this model, they analyze the case where the IM cannot observe the current inventory level due to information delay. Instead, he can observe the inventory level of a prior period. Finally, they show that state-dependent base-stock policy is optimal for the information delay model. However, neither Treharne and Sox [22] nor Bensoussan et al. [3] consider the supply side. Therefore, we also extend Treharne and Sox [22] by allowing randomness in yield and partial observation of the supply environment.

The organization of this paper is as follows. In the next section, we introduce our notation, describe the basic model, and state our assumptions. In Section 3, we analyze the single-period problem and present our results. Next, in Section 4, we focus on the multiple-period problem. Then, in Section 5, we study the infinite-period problem. We discuss the effects of the observation level and different demand/supply combinations on the optimal policies via numerical illustrations in Section 6. Moreover, we discuss convergence of policy parameters and optimal cost function in Section 6. Finally, in Section 7, we conclude and state possible extensions