Dynamic Pricing Policies for an Inventory Model with Random Windows of Opportunities

Arnoud den Boer,[1,2] Ohad Perry,[3] Bert Zwart[2]

[1] *Korteweg-de Vries Institute for Mathematics, University of Amsterdam*

[2] *Amsterdam Business School, University of Amsterdam*

[3] *Industrial Engineering and Management Sciences, Northwestern University, Evanston, Illinois 60208*

Received 11 February 2016; revised 23 March 2017; accepted 24 March 2017
DOI 10.1002/nav.21737
Published online 00 00 2017 in Wiley Online Library (wileyonlinelibrary.com).

Abstract: We study a single-product fluid-inventory model in which the procurement price of the product fluctuates according to a continuous time Markov chain. We assume that a fixed order price, in addition to state-dependent holding costs are incurred, and that the depletion rate of inventory is determined by the sell price of the product. Hence, at any time the controller has to simultaneously decide on the selling price of the product and whether to order or not, taking into account the current procurement price and the inventory level. In particular, the controller is faced with the question of how to best exploit the random time windows in which the procurement price is low. We consider two policies, derive the associated steady-state distributions and cost functionals, and apply those cost functionals to study the two policies. Naval Research Logistics 00: 000–000, 2017

Keywords: single-product fluid-inventory model; optimal order policies; stochastic systems

1. INTRODUCTION

We consider a continuous review, single product, pricing-and-inventory problem in a random environment, where the purpose of the seller is to maximize his expected profit by determining an order policy and sell prices. At the procurement side, the seller faces randomly fluctuating prices at which he can acquire new items, but also holding costs and fixed order costs. Based on these quantities, the seller needs to decide when to order new items, and how many. At the sales side, in accordance with current practice if dynamic pricing, the seller can change the sell price at any moment.

More specifically, the procurement price at which new items can be acquired is modeled as a finite-state Markov chain, where each state represents a different procurement price. Every time an order is placed, the seller pays some fixed order cost K, and any moment that the inventory-level is $x > 0$, the seller faces holding costs at a rate $h(x)$. We initially assume that ordered items arrive instantaneously and later generalize the model to include exponential lead times. The seller needs to determine an order policy (when to order new items, and how many), and a sell price policy (which sell price to charge at which moment), in order to optimize the expected profit.

Determining optimal order policies and sell prices is typically treated as separate problems, but it is intuitively clear that it may be beneficial to consider these problems simultaneously. For example, if the procurement price of new items is time-dependent and is currently high, it may be profitable to increase the sell price, so that the moment at which all inventory is sold-out is delayed. This increases the probability that, in the mean time, the procurement price of new items decreases, so that new items can be ordered at considerably lower costs.

A Fluid Analysis

The complexity of the stochastic model we consider renders exact analysis prohibitively hard, and we therefore resort to fluid approximations. In our setting here, the fluid approximation is a piecewise-continuous process having a deterministic evolution between jump epochs.

Fluid models are prevalent approximations for complex stochastic systems in general, and queueing and inventory systems in particular. In a fluid inventory model, each item becomes an "atom" in a *continuous* content process, and

Correspondence to: Ohad Perry (ohad.perry@northwestern.edu)

the random depletion rate of inventory is replaced by the mean demand rate, ignoring variability. Thus, fluid analysis is appropriate whenever the number of items sold between decision epochs, as well as the number of items ordered, is large, such that each single item is relatively negligible.

By aggregating the effect of a large number of events, a fluid inventory model provides the time-dependent average behavior of the system. Although this view can often be made rigorous (under appropriate regularity conditions) by proving that the fluid model arises as a functional weak law of large numbers for the stochastic inventory process under study (see, e.g., [10] and [21]), fluid dynamics are typically assumed directly in inventory systems, without any reference to an underlying discrete system. For example, one of the most fundamental inventory models - the (purely deterministic) *economic order quantity* (EOQ) model—can be thought of as a fluid approximation for a stochastic inventory system. For a recent application, see [19] and references therein.

Other than their tractability and relative simplicity, fluid approximations have the advantage of being less sensitive to distributional properties of the stochastic system they approximate, making fluid analysis robust to variations in the modeling assumptions. However, we emphasize at the outset that our fluid model is itself a stochastic process, whose randomness emanates from that of the environment. We elaborate below.

A Class of Stationary Markovian Policies

Our goal is to find an effective control for the inventory process. Since the random environment must clearly be taken into account when making ordering decisions, the fluid model is a stochastic process. To facilitate long-run analysis, *we limit attention to a class of stationary and Markovian policies*, so that, in particular, the fluid model is a stationary Markov process under the control (namely, it achieves a unique steady state). It is well known that, if a Markov process has a continuous segment, then it is either a transformation of a Brownian motion or it is deterministic [13]. Since our model has nonincreasing sample paths between jumps, it must exhibit a deterministic motion between jump epochs, so that it is a *piecewise-deterministic Markov* process, as in [15].

Now, in order for the controlled fluid process to be Markov, decisions must be made based on the current state of the process and the environment. Furthermore, for the process to be ergodic (so that a unique stationary and limiting distribution exists), the fluid process must be regenerative. Therefore, the class of controls we consider is continuous-review policies of (s, S) type. It is significant that, in addition to being mathematically attractive, these policies are also easy to employ and are prevalent in practice; see, for example, [25]. Note also that, with no lead times, (s, S) policies are equivalent to (r, Q) policies in which a fixed quantity $Q > 0$ is ordered whenever the inventory level falls below some level $r \geq 0$.

In particular, for the order policy, we study two variants of an (s, S)-policy. In the first order policy, which we denote by OP_1, $S-s$ items are ordered if the inventory-level is at or below some $s > 0$ and at the same time the procurement price is low. If the inventory hits zero and the procurement price is high, Q items are ordered. Here s, S, Q are decision variables, with $0 \leq s < S$, $0 < Q \leq S$. The second order policy, denote by OP_2, orders are never placed when the procurement price is high. If the inventory-level hits zero, the seller waits until the procurement price becomes low, at which moment he orders S items. We allow the sell price to change with the inventory level. Note that either policy takes advantage of the low procurement price, and we thus refer to the (random) time periods of low procurement price as *random windows of opportunity* (for the seller).

We consider the pricing-and-inventory problem in stationarity. Under mild assumptions on the relation between the demand rate and sell price, we show that the joint process of inventory level and procurement price admits a unique stationary distribution. For a fixed order policy OP_1 or OP_2, we derive balance equations for the stationary distribution of the inventory-level process, as in [5] and [14], from which the stationary distribution can be numerically computed. This enables us to express the long-run profit for both policies, as function of $(s, S, Q, p(\cdot))$ in case of OP_1, and $(s, S, p(\cdot))$ in case of OP_2, where $p(\cdot) : [0, S] \rightarrow \mathbb{R}_+$ is the sell price function.

For practical purposes, one uses a piecewise constant sell-price function $p(\cdot)$. For example, if $p(\cdot)$ has only two values, p_l (for "low" price) and p_h (for a "high" price), then whenever the inventory level exceeds some (switching) threshold q, a sell price p_h is charged, whereas a sell price p_l is charged whenever the inventory level is below q. In that case, (p_l, p_h, q) is a vector of finite decision variables which satisfy $0 < p_l < p_h$ and $0 \leq q \leq S$; see Section 2 below for an elaboration.

To determine the optimal values of the decision variables, one needs to solve a (rather complicated) nonconvex nonlinear optimization problem. We conduct a numerical study to compare the performance of OP_1 and OP_2. We also compare them to a standard (s, S)-policy OP0, which does not take into account the random nature of the procurement price process. By studying several instances, it turns out that OP_1 in general performs better or equal than both OP_2 and OP0. The difference in performance, especially between OP_1 and OP0, can be quite large. This shows that it is beneficial to take into account random changes in the procurement prices. As should be expected, the policies OP0 and OP_2 have no clear "best": for some instances, the first is outperformed by the latter, while for other instances it is the other way around.

1.1. Related Literature

Threshold policies have been shown to be optimal in numerous settings, including under Markovian-demand environments. A detailed development and literature review is found in [4]. An early result by Iglehart and and Karlin [20] considers a discrete-time inventory model with demand that is governed by a discrete-time Markov chain (DTMC). In particular, at each period the demand distribution is set by the state of the DTMC. Song and Zipkin [30] analyze a continuous-time inventory model having Markov-modulated Poisson demand and backlogging, and prove that a state-dependent (s, S) policy is optimal under the assumption of fixed ordering costs. See also [11], which considers a discrete-time version of the problem, and [12] which extends the model to include lost sales. Related to [30] are the two papers [2] and [3], which consider EOQ-type models with Markov-modulated demand process. A multistage serial inventory model with Markov modulated demand in stage 1 is analyzed in [9] and it is shown that, under linear holding and ordering costs, an echelon base-stock with state-dependent order-up-to levels policy is optimal.

The paper [6] considers a fluid-inventory model in which the demand rate changes according to a CTMC; whenever the inventory content hits 0 an order of size Q_i is placed if the governing CTMC is at state i, $i = 1, 2$. See also [8] which considers a fluid inventory model in which the procurement prices change according to an exogenous CTMC and, unlike our model, this also affects the demand. (We assume that demand is affected by the sell price, which is controlled by the decision maker.) In [28], the authors consider an inventory model which replenishes at a constant deterministic rate, but decreases randomly (via jumps) when demand arrives (according to a compound renewal process); implying that the demand at each arrival epoch is relatively very large. In contrast, in our model demand arrives continuously and is infinitesimal in the fluid model, and the orders are relatively large, causing the jumps in the fluid model; see also [29].

Threshold policies for EOQ-type fluid inventory models are considered in [18] and [24]. In both references it is assumed that discounts are offered by a supplier to a reseller randomly in accordance with a Poisson process, but there are no "windows" openings. In these cases, the reseller has to decide at any discount epoch whether to replenish his inventory or not. In [27] a firm purchases a product in an auction in order to satisfy its own demand in each period. In particular, each period consists of two phases: in the first phase, the firm participates in N auctions, and in the second it sells the purchased products in its own market. The probability of the firm winning an auction is assumed to be a function of its own bid as well as of the number of its opponents in the auction. In particular, there is no a-priori fixed procurement price, although the firm does have some control over that price via its bidding strategy.

Price-Regulated Demand

Starting with the seminal work of Naor [26], a standard assumption in the economic analysis of queues is that customers' arrival rate to a service system is completely determined by the price and expected reward of joining the system to get served. For example, one often assumes that the potential arrival rate, known as the *market size*, is a constant Λ, and the arrival process to the system is a nonhomogeneous Poisson process having an instantaneous rate $\Lambda f(p(t))$ at time t when the price is $p(t)$, where for $0 \le p_m < p_M \le \infty$, $f : [p_m, p_M) \to [0, 1]$ is a known *market response function*. In particular, $f(p(t))$ is the fraction of customers that are expected to join the system when the price is $p(t) \in [p_m, p_M)$; see, for example, [1, 16, 22] and references therein. The market response function is determined by the probability that a generic arrival will choose to buy a product at the given sell price. In practice, that probability, and thus the response function, are not known in complete certainty, although they can typically be evaluated via past demand data; see [17] and the reference therein. In the aggregate time-dependent (functional) average approximation that the fluid model provides, the market size and response function completely determine the depletion rate of the fluid content for any given sell price.

Organization

The remainder of this article is organized as follows: In Section 2, we describe the model and motivate the structure of the control policies. In Section 3, we develop the steady-state equations for the content level process. Those equations are then applied in Section 4 in a numerical study, as described above. In Section 5, we extend the model and consider cases in which the procurement price of the item changes after non-exponential random time in states, and we also consider lead times.

2. THE MODEL

We consider a fluid inventory model of one product with zero lead time of the (s, S) type, operating in a stochastically changing cost environment. We use $C := \{C(t) : t \ge 0\}$ to denote the content-level process, assumed to be right continuous with jumps at ordering epochs. As there are no lead times, the right continuity of C implies that, if t is a jump epoch, then $C(t-) < C(t)$, where $C(t-)$ denotes the left limit at t.

Following the terminology in [7] and [30], we refer to procurement price as the "state of the world." In particular, the procurement price of the product changes according to a two-state continuous-time Markov chain (CTMC) $W := \{W(t) : t \geq 0\}$, with W attaining two values: w_λ (high) and w_μ (low). Naturally, w_λ is strictly larger than w_μ. (Otherwise, the state of the world is irrelevant.) More specifically, W moves between the two states w_λ and w_μ, and remains at w_λ for an exponential amount of time with rate λ, and in w_μ for an exponential amount of time with rate μ. When $W = w_\lambda$ the controller faces a regular (expensive) price, and when $W = w_\mu$, the controller faces a discounted (cheap) procurement price. It is thus clear that the "state-of-the-world" process W may affect the decision of the controller whether or not to buy at each decision epoch in order to replenish his inventory.

We assume that a holding cost is incurred at rate $h(x)dx$ whenever $C(t)=x,\ t \geq 0$, and that a fixed set-up cost K is incurred when an order is placed, independent of the order size. In addition, we assume that the demand rate is a known one-to-one and onto function of the sell price. Under this assumption, the controller can dynamically regulate the release rate of inventory by changing the sell price. There can be several policies for determining the sell price. In this study, we focus on the state of the content level C. More precisely, since the more inventory present, the higher instantaneous holding cost is paid, the controller has an incentive to drain inventory at a higher rate when C is high, by lowering the sell price. In the continuous settings, the optimal release rate may change continuously as a deterministic function of C, so that infinitely many pricing policies can be applied. For practical purposes, the optimal pricing policy can be approximated by searching for a finite set of sell prices $p_1 < p_2 < \cdots < p_k$ (with k fixed) and thresholds $q_1 > q_2 > \cdots > q_{k-1}$, such that the sell price is p_i at time t if $q_{i-1} < C(t) < q_i,\ i = 1, 2 \ldots, k-1$. Clearly, as the number of decision variables increases, the optimization problem becomes more complicated.

In the simple (*s, S*) model, the optimal control is comprised of two factors: when to place an order (in the sense of fixing s) and how much to order (fixing level S). Thus, if the procurement price was always w_μ we would have been looking for a level s such that, whenever the content-level process C hits s, an order of size S–s is placed. In light of the randomness of the procurement price and zero lead-time assumptions, it is desirable to place most of the orders, if not all of them, when the procurement price is w_μ. In particular, the distinction between "most" and "all" depends on whether it is optimal to place an order whenever both $C(t)=0$ and $W(t) = w_\lambda$, that is, whenever the content level drops to zero at the time of an expensive cost-price period. In that case, one should consider two options: (i) order up to level $Q \leq S$ or (ii) wait for the procurement price to change from w_λ to w_μ.

We thus consider two natural ordering policies:

Order Policy 1 (OP$_1$)

Determine two levels s and S. If the content level C hits s and at the same time the procurement price is low, that is, $C(t-) = s$ **and** $W(t-) = w_\mu$, then place an order of size S–s (so that $C(t)=S$. If, conversely, upon hitting level s the procurement price is high, that is, $C(t-) = s$ and $W(t-) = w_\lambda$, then wait until either (*i*) the procurement price changes to w_μ, at which point order up to S, or (*ii*) the content level hits 0, at which point order up to level Q, where $Q \leq S$

Order Policy 2 (OP$_2$)

Similarly to OP$_1$, except that never place an order while the procurement price is high, that is, whenever $W = w_\lambda$. When level 0 is hit (and it can only be reached during expensive periods) wait until the procurement price changes to cheap (w_μ), at which point order up to level S. Note that, under OP$_2$, there is no extra level Q (alternatively, $Q \equiv S$).

We further assume that there is a cost incurred for letting C stay at state 0 for an interval. This cost can be due to unsatisfied demand and loss of good will of customers and so forth. In particular, if $C(t)=0$ on some interval $[t_1, t_2]$, then a cost $a(t_2 - t_1)$ is incurred.

To fully describe the control, we need also to characterize the threshold q and the sell prices p_l and p_h. That is, under OP$_1$ the control is determined by the decision variables (s, S, q, Q, p_l, p_h), while under OP$_2$ the control is determined by the decision variables (s, S, q, p_l, p_h). Alternatively, because of the equivalence between the sell prices and the demand rate, we can replace p_l and p_h by d_l and d_h, respectively.

To distinguish between the two policies, we let $C_1 := \{C_1(t) : t \geq 0\}$ denote the content-level process under OP$_1$, and $C_2 := \{C_2(t) : t \geq 0\}$, denote the content-level process under OP$_2$. We still use the notation C in discussions in which no specific process is considered (if the same is true for both C_1 and C_2).

2.1. The Fluid Process Achieved via Asymptotic Considerations

In our model, the content-level process C is assumed to decrease deterministically and continuously in between orders, with the instantaneous decrease rate determined by the sell price. To see that this assumption follows from standard assumptions in the literature (as was reviewed in Section 1.1 above), recall that a fluid inventory model is achieved as a relaxation to a stochastic system, building on asymptotic considerations, as described in Section 1. In particular, our model is appropriate as an approximation for a large inventory system in which large orders are placed rarely relative to the interarrival times of customers that purchase those items.

For example, orders of several hundreds of items may be made once every few weeks, and tens of items are sold daily. We now provide a quick overview of the arguments that can be used to prove the limiting result.

Formally, the fluid model can be achieved as a stochastic-process limit of a sequence of stochastic inventory systems indexed by n. With each $n \geq 1$, there are associated parameters s^n, S^n, q^n, and Q^n that increase linearly with n, so that, for example, $S^n/n \to S > 0$ as $n \to \infty$, and an arrival process of customers purchasing the items with an arrival rate that increases proportionally to n as well.

Let $C^n := \{C^n : n \geq 1\}$ denote the content level process in system n, and let $\bar{C}^n(t) := C(nt)/n$, $t \geq 0$. The fluid process C is an approximation for C^n in the sense that $\bar{C}^n(t) \approx nC(t)$ for large-enough n. Observe that the fluid approximates a large system C^n by accelerating time by a factor of n in $\bar{C}^n$. This implies that the state-of-the-world process should be "slowed-down" (relative to the demand process) in order for it to have the same time scale as the interorder times. Specifically, in system n, the procurement price evolves according to a CTMC $W^n := \{W^n(t) : t \geq 0\}$, which spends an exponentially-distributed amount of time in state w_λ and in w_μ with mean $n\lambda$ and $n\mu$, respectively.

It remains to describe the demand process that leads to the deterministic demand rate in the fluid model. To this end, assume that the arrival process of customers constitutes a Poisson process with some rate $\Lambda > 0$. Assume further that each customer has a private valuation for the item under consideration, and that the valuations of customers are random variables that are independent across the customers and are identically distributed. Let V denote a generic random variable that has the customers' valuation distribution, and let F_V denote its cumulative distribution function (cdf). Then an arrival will purchase an item with probability $F_V^c(p(x)) := P(V > p(x))$ when the price is $p(x)$, so that $F_V^c(p(x))$ is the proportion of all arrivals that purchase items when the price is $p(x)$. Given these assumptions, the instantaneous *demand rate* when the content is x and the price is $p(x)$, is

$$d(x) := \Lambda F_V^c(p(x)), \tag{1}$$

so that the *demand process* is

$$\mathcal{N}\left(\int_0^t \Lambda F_V^c(p(x))dx\right), \quad t \geq 0,$$

where $\mathcal{N}$ is a unit-rate Poisson process. It follows from the functional strong-law of large number for the Poisson process (e.g., section 3.2 in [32]) and the continuity of the composition mapping at continuous limits (e.g., [31], Theorem 13.2.1]) that, uniformly over compact intervals,

$$\begin{aligned}\frac{\mathcal{N}\circ \int_0^t n\Lambda F_V^c(p(x))dx}{n} &:= \frac{\mathcal{N}\left(\int_0^t n\Lambda F_V^c(p(x))dx\right)}{n} \\ &\to \int_0^t \Lambda F_V^c(p(x))dx \\ &\quad \text{as } n \to \infty \quad w.p.1.\end{aligned}$$

In particular, we obtain that the fluid model $C(t)$ is differentiable at all its continuity points, namely, at all points in which no orders arrive, and evolves according to the differential equation $C'(x) = -d(x)$ between orders, for $d(x)$ in (1). We thus see that the random demand process in the prelimit is replaced with a deterministic demand process in the fluid limit.

3. STEADY-STATE ANALYSIS

We will analyze the inventory system in stationarity. Hence, we need to argue that a unique stationary distribution indeed exists for our system. We will analyze a system having a general demand-rate function, which allows for a general pricing policy analysis in our setting. Let $p_1 : [0, S] \to \mathbb{R}_+$ and $p_2 : [0, S] \to \mathbb{R}_+$ be the pricing policies under OP_1 and OP_2, respectively. For $x \in [0, S]$, let $d_1(x) := d(p_1(x))$ and $d_2(x) := d(p_2(x))$ denote the respective demand functions.

We make the following assumption, which will be shown to ensure that the system possesses a unique stationary distribution. Let

$$D_i(x) := \int_0^x \frac{1}{d_i(y)}\, dy, \quad 0 \leq x \leq S. \tag{2}$$

ASSUMPTION 1: The pricing policy employed is such that $D_i(S) < \infty$ for $i = 1, 2$.

Note that $D_i(x)$ is the time to reach level 0 from level x, for all $0 < x \leq S$, if the input is shut off, that is, if there are no new inventory orders during $D_i(x)$ time units. Then Assumption 1 simply states that the content level can reach state 0 in finite time, provided no new orders are placed during the time interval $[0, D_i(S)]$ and $C_i(0) = S$. This assumption holds trivially whenever d_i is a simple function, $i = 1, 2$, which is the case amenable to numerical studies and optimizations.

Note that, for $i = 1, 2$, the content level C_i is not Markov, but

$$X_i := \{X_i(t) : t \geq 0\} := \{(C_i(t), W(t)) : t \geq 0\}$$

is a two-dimensional Markov process with state space $\mathcal{S} := [0, S] \times \{w_\lambda, w_\mu\}$. It is simple to show that X is regenerative and possesses a unique stationary distribution. Let $W(\infty)$ denote a random variable having the stationary distribution of the process W, and let $C_i(\infty)$ be a random variable having the stationary distribution of C_i, $i = 1, 2$. Then

$X_i(\infty) := (C_i(\infty), W(\infty))$ is a random variable with the stationary distribution of the process X_i, $i = 1, 2$. All these random variables exist by the following theorem.

PROPOSITION 3.1: If Assumption 1 holds, then for $i = 1$, 2, the joint process $X_i = (C_i, W)$ is a (possibly delayed) regenerative process admitting a unique stationary distribution.

PROOF: It is easy to see that X_i, $i = 1, 2$, is nonlattice, and will return to state $x^* := (S, w_\mu)$ in finite expected time, given our assumptions on the model. In particular, let E_μ and E_λ denote two generic exponential random variables representing the times that W spends in each of its states w_μ and w_λ, respectively, and let T denote the return time of C_1 to S, and take $X_1(0) = (S, w_\mu)$. Then

$$\begin{aligned} E[T] &= E[T|W(D(S-s)) = w_\mu]P(W(D(S-s) = w_\mu) \\ &\quad + E[T|W(D(S-s)) = w_\lambda]P(W(D(S-s)) = w_\lambda) \\ &\leq D(S-s) + E[T|W(D(S-s)) = w_\lambda]. \end{aligned}$$

Now, if at time $D(S-s)$ the state of the world is w_λ, then the content process will either jump back to S if $E_\mu \leq s$, namely, with probability $1 - e^{-\mu D(s)}$, or it will jump to Q if $E_\mu > s$, that is, with probability $e^{-\mu D(s)}$. Therefore, letting T_Q denote the time to return to x^* when starting in (Q, w_λ), the second term in the right-hand side of the equality above satisfies

$$\begin{aligned} E[T|W(D(S-s)) = w_\lambda] &\leq D(S)(1 - e^{-\mu D(s)}) \\ &\quad + E[T_Q]e^{-\mu D(s)}. \end{aligned}$$

Observe that T_Q is a geometric sum with success probability $P(E_\lambda > D(Q)) = e^{-\lambda D(Q)}$ of random variables, where each of the random variables in the sum is bounded from above by $D(Q)$ w.p.1. The statement of the proposition follows for X_1 from Assumption 1. Similar arguments can be employed to prove the result for X_2. □

REMARK 3.1: It is clear from the arguments in the proof of Proposition 3.1 that it is sufficient to assume that $D_1(y) < \infty$ for some $y > S - s$, that is, the content level can go below level s. However, OP_2 requires that the content level can reach level zero in finite time.

3.1. Steady-State Equations

We now compute the unique stationary distribution of the processes C_1 and C_2. In some models, simplifications occur due to a form of asymptotic independence between the content level C and the "world" process W (using our notation), that is, $C(\infty)$ is independent of $W(\infty)$, so that the stationary distribution of X is the product of the stationary distributions of C and W. Such is the case, for example, when W is a "well-behaved" Markov process which determines the demand process; see, for example, [7] and references therein. However, such simplification cannot be expected to hold in our model, since the position of $C(t)$ contains significant information on the value of $W(t)$ at each t, even when the joint process X is stationary (that is, if $X(t)$ is distributed as $X(\infty)$ for all $t \geq 0$). For example, if $C(t) < s$, then necessarily $W(t) = w_\lambda$. However, there is still simplification in our case, which stems from the fact that the world process W does not depend on the content level C, and can be analyzed separately. We can thus find the stationary distribution of C by computing relevant stationary quantities of W.

We next introduce integral representations for the steady-state density functions of the content level process. Let $f_1 : [0, S] \to \mathbb{R}_+$ and $f_2 : [0, S] \to \mathbb{R}_+$ denote the steady-state density functions of C_1 and C_2, respectively. The next theorem provides an integral representation for the steady-state densities f_1 and f_2. We present two equations for the density under OP_1, for the two cases $s < Q$ and $s \geq Q$.

Consider the case $s < Q$, and take $x > s$. Let k_1 denote the long-run rate of upcrossings of level x, that is, the long-run average number of jumps from s to S. For the case $s \geq Q$, let $\tilde{k}_1$ denote the long-run rate of upcrossing of level x, $s \leq x \leq S$. We denote by k_2 the long-run rate of upcrossings of level x, $x \geq s$, caused by jumps from level s under OP_2.

The main difficulty in our model is in determining the long-run rate of jumps from level s, that is, the values of k_1, $\tilde{k}_1$, and k_2. We first present the integral equations for the steady-state densities without specifying these constants: their values are computed in Lemma 3.4 below, after the solutions to the steady-state densities, and their respective cdf's are computed in terms of these constants.

Let π_2 denote the atom at 0 of the stationary content level C_2, that is,

$$\pi_2 := P(C_2(\infty) = 0) > 0. \tag{3}$$

LEMMA 3.1 (integral equations for steady-state densities) The steady-state densities $f_1(x)$ of C_1 and $f_2(x)$ of C_2 exist.

Furthermore, $f_1(x)$ satisfies one of the following integral equations, depending on whether $s \leq Q$ or $s > Q$:

$$\underline{If\ s \leq Q} : d_1(x) f_1(x) = \begin{cases} \lambda \int_0^x f_1(w)\,dw + d_1(0) f_1(0), & 0 \leq x < s, \\ \lambda \int_0^s f_1(w)\,dw + d_1(0) f_1(0) + k_1, & s \leq x < Q, \\ \lambda \int_0^s f_1(w)\,dw + k_1, & Q \leq x \leq S. \end{cases}$$

$$\underline{If\ s > Q} : \ d_1(x) f_1(x) = \begin{cases} \lambda \int_0^x f_1(w)\,dw + d_1(0) f_1(0), & 0 \leq x < Q, \\ \lambda \int_0^s f_1(w)\,dw, & Q \leq x < s, \\ \lambda \int_0^s f_1(w)\,dw + \tilde{k}_1, & s \leq x \leq S. \end{cases} \tag{4}$$

The steady-state density $f_2(x)$ satisfies the integral equation

$$d_2(x)f_2(x) = \begin{cases} \lambda\int_0^x f_2(w)\,dw + \lambda\pi_2, & 0 \le x < s, \\ \lambda\int_0^s f_2(w)\,dw + \lambda\pi_2 + k_2, & s \le x \le S. \end{cases} \quad (5)$$

PROOF: Existence of the stationary densities follows from Corollary 4.1 in [23]. We explain the derivation of the integral equation for f_1 for the case $s \le Q$. The other equations are derived similarly.

First, $d_1(x)f_1(x)$ in the left-hand side of (4) is the long-run rate of downcrossings level x, while the right-hand side represents the long-run rate of upcrossings of level x, $0 \le x \le S$. In steady state, the rate of downcrossing must equal to rate of upcrossing that level, which is what (4) states. For a rigorous definition of "rate" we again refer to [23]. Here, the meaning will become clear from the proof. To see this, assume that $C_1(0) \stackrel{d}{=} C_1(\infty)$, namely, $C_1(0)$ has the steady-state distribution of the content level. That makes C_1 a stationary process, so that $C_1(t) \stackrel{d}{=} C_1(\infty)$ for all $t \ge 0$. Let τ be an arbitrary point of a jump. Since jumps can only occur when $0 \le C_1 \le s$, we separate the analysis into three cases as follows:

1. $0 \le C_1(\tau-) < x < s$. The last jump in the cycle brings the content level up to level Q, and the other jumps, if any, bring the content to level S (where $S \ge Q$). Thus, if $C_1(\tau-) > 0$, τ is a beginning of a cheap period and $C_1(\tau) = S$. If $C_1(\tau-) = 0$, then τ is a time of depletion and $C_1(\tau) = Q$. Both types of jumps imply that the jump is an upcrossing of level x. Since the expensive period is exponentially distributed with rate λ, it follows by the well-known PASTA (Poisson Arrivals See Time Average [33]) property that if $C_1(\tau-) > 0$, then $C_1(\tau-)$ and C_1 are equal in distribution, and the rate at which level x is upcrossed is λ. The rate at which $C_1(\tau-) = 0$ is $d(0)f_1(0)$. Thus, the rate at which level x is upcrossed is $\lambda\int_0^x f_1(w)dw + d(0)f_1(0)$.
2. $0 \le C_1(\tau-) \le s$ and $s \le x < Q$. Again, every jump is an upcrossing of level x. However, in addition to the previous case (i), there is also a possibility to jump above level x from level s (when level s is reached during a cheap period). That long-run rate is denoted by k_1 (and will be computed in Lemma 3.4 below).
3. $0 \le C_1(\tau-) \le s$ and $Q \le x \le S$. In this case, level x cannot be upcrossed by a jump from level 0. Thus, the rate $d_1(0)f_1(0)$ is removed. □

The arguments for f_1 in the case $s > Q$ and for f_2 are similar. (Note however that f_2 has an atom π_2 at level 0.)

3.2. Solutions to f_1 and f_2

We solve for f_1 and f_2 in (4) and (5) in terms of unknowns k_1, $\tilde{k}_1$ and k_2 whose values are determined by the transient distribution of the state of the world process W. We compute these unknowns explicitly in Lemma 3.4.

LEMMA 3.2 (Steady-state distribution). The steady-state density functions f_1 and f_2 satisfy

$$f_1(x) = \begin{cases} \frac{c_0}{d_1(x)}e^{\lambda D_1(x)}, & 0 < x < s, \\ (c_0e^{\lambda D_1(s)} + k_1)D_1(x), & s \le x < Q, \\ (c_0e^{\lambda D_1(s)} + k_1)D_1(Q)/d_1(x), & Q \le x \le S, \end{cases}$$

and

$$f_2(x) = \begin{cases} \frac{\lambda\pi_2}{d_2(x)}e^{\lambda D_2(x)}, & 0 < x < s, \\ (\lambda F_2(s) + \lambda\pi_2 + k_2)D_2(x), & s \le x < S, \end{cases}$$

where the constant k_1, $\tilde{k}_1$, and k_2 are given in Lemma 3.4 below, and c_0 *and* π_2 are the unique constants for which

$$\int_0^S f_1(s)dx = 1 \quad \text{and} \quad \pi_2 := 1 - \int_0^S f_2(x)dx.$$

PROOF: Let $F_1(x) := \int_0^x f_1(s)ds$ denote the cdf, associated with the density f_1. Let $c_0 := d_1(0)f_1(0)$. For $0 \le x < s$, we write $f_1(x) - \lambda/d(x)F_1(x) = c_0/d_1(x)$. Then, multiplying that equation by $exp\{-\lambda D_1(x)\}$ and integrating (recall that $\frac{d}{dx}D_1(x) = 1/d_1(x)$), we get

$$e^{-\lambda D_1(x)}F_1(x) = \int_0^x \frac{c_0}{d_1(s)}e^{-\lambda D_1(s)}ds = -\frac{c_0}{\lambda}e^{-\lambda D_1(x)} + C_1,$$

so that

$$F_1(x) = -\frac{c_0}{\lambda} + C_1e^{\lambda D_1(x)}, \quad x \in [0, s),$$

for some constant C_1. Using the initial condition $F_1(0) = 0$ (and $D_1(0) = 0$), we see that $C_1 = c_0/\lambda$. It follows that

$$F_1(x) = \frac{c_0}{\lambda}(e^{\lambda D_1(x)} - 1), \quad 0 \le x < s,$$

so that

$$f_1(x) = \frac{c_0}{d_1(x)}e^{\lambda D_1(x)}, \quad 0 \le x < s.$$

Hence,

$$f_1(s-) = \frac{c_0}{d_1(s)}e^{\lambda D_1(s)} \quad \text{and} \quad F_1(s) = \frac{c_0}{\lambda}[e^{\lambda D_1(s)} - 1].$$

Next, consider $x \in [s, Q)$. Then

$$\begin{aligned} d_1(x)f_1(x) &= \lambda F_1(s) + c_0 + k_1 \\ &= c_0 e^{\lambda D_1(s)} + k_1. \end{aligned}$$

Since the right-hand side of (4) over $[Q, S]$ is a constant, the expression for f_1 follows. Finally, the constant c_0 is obtained by applying the normalization condition $\int_0^S f_1(x)\,dx = 1$, and is given in terms of k_1.

The solution for f_2 is computed similarly. □

3.3. Jumps from Level *s*

It remains to find the constants k_1, $\tilde{k}_1$, and k_2. To that end, we define the following conditional probabilities: Let $\theta_1(s, S)$ and $\theta_2(s, S)$ denote the conditional probabilities that level s is downcrossed during a cheap period, under OP_1 and OP_2, respectively, given that the last jump prior to hitting s was to level S. Let $\gamma_1(s, Q)$ denote the conditional probability that level s is downcrossed during a cheap period under OP_1, given that the last jump prior to hitting s was to level Q (which under OP_1 corresponds to the beginning of a regenerative cycle). The closed-form expressions for $\theta_1(s, S)$, $\theta_2(s, S)$ and $\gamma_1(s, Q)$ are computed in Lemma 3.3 below. These expressions depend only on the (known) parameters of the cost process C, and on the function D.

Observe that $\gamma_1(s, Q) = 0$ if $Q < s$. Let $1\{s < Q\}$ be the indicator function which equals 1 if $s < Q$ and 0 otherwise. The proof of the following lemma is straightforward, and is thus omitted.

LEMMA 3.3:

$$\theta_1(s, S) = \theta_2(s, S) = \frac{\lambda}{\lambda+\mu} + \frac{\mu}{\lambda+\mu} e^{-(\lambda+\mu)[D_1(S)-D_1(s)]},$$
$$\gamma_1(s, Q) = \left(\frac{\lambda}{\lambda+\mu} - \frac{\lambda}{\lambda+\mu} e^{-(\lambda+\mu)[D_1(Q)-D_1(s)]}\right) 1\{s < Q\}.$$

In the next lemma, we express the constants k_1, $\tilde{k}_1$ and k_2.

LEMMA 3.4: Consider $x \in (s, S]$. Then the long-run rate of upcrossings of level x under OP_1 is given by k_1 if $s \le Q$ and $\tilde{k}_1$ if $s \ge Q$. It is given by k_2 under OP_2, where

$$\begin{aligned} k_1 &:= \gamma_1(s, Q) d_1(0) f_1(0) + \theta_1(s, S) d_1(S) f_1(S), \\ \tilde{k}_1 &:= \theta_1(s, S) d_1(S) f_1(S), \\ k_2 &:= \theta_2(s, S) d_2(s) f_2(s), \end{aligned} \tag{6}$$

for $\gamma_1(s, Q)$, $\theta_1(s, S)$ and $\theta_2(s, S)$ in Lemma 3.3.

PROOF: We find k_1. The computations of $\tilde{k}_1$ and k_2 are similar. (See also Remark 3.2 below.) Consider the state of the content level immediately after a jump. Clearly, the process between jumps is a DTMC with two states – S and Q. The transition matrix of that DTMC at jump epochs is

$$\begin{aligned} P := &\begin{bmatrix} P_{S,S} & P_{S,Q} \\ P_{Q,S} & P_{Q,Q} \end{bmatrix} \\ = &\begin{bmatrix} \theta_1 + (1-\theta_1)(1-e^{-\lambda D_1(s)}) & (1-\theta_1)e^{-\lambda D_1(s)} \\ 1-(1-\gamma_1)e^{-\lambda D_1(s)} & (1-\gamma_1)e^{-\lambda D_1(s)} \end{bmatrix}. \end{aligned} \tag{7}$$

We now explain the entries of the transition matrix, starting with the first row. The content level jumps to state S only when the environment is cheap. There are two possibilities to make a transition from S to S: Either the content level started at S and arrived at level s during a cheap period, in which case there is a jump immediately back to level S—this event occurs with probability θ_1. Else, the content level arrives at level s during an expensive period and there is no jump at s, but the expensive period is terminated before the content level reaches level 0. The probability of that latter event is $(1-\theta_1)(1-e^{-\lambda D_1(s)})$. This explains the first row of the transition matrix (7).

Turning to the second row, recall that the content level reaches level 0 only when the environment is expensive, in which case the content level jumps to level Q. Thus, the DTMC at jumps epochs moves from Q to Q only if level s was reached during an expensive period, and the environment remained expensive till the content level reached 0. The event occurs with probability $P_{Q,Q} = (1-\gamma_1)e^{-\lambda D_1(s)}$. To see why, note that $1-\gamma_1$ is the probability of reaching s at "expensive", given that the last jump was to Q, and $e^{-\lambda D_1(s)}$ is the probability that the environment did not change to "cheap" after level s was downcrossed, and before level 0 was reached.

We denote the stationary probabilities of the above Markov chain by ν_S and ν_Q, with $\nu := (\nu_S, \nu_Q)$. Calculating $\nu P = \nu$ and $\nu_S + \nu_Q = 1$ gives

$$\nu_S = \frac{1-(1-\gamma_1)e^{-\lambda D_1(s)}}{1-(\theta_1-\gamma_1)e^{-\lambda D_1(s)}} \quad \text{and} \quad \nu_Q = 1 - \nu_S, \tag{8}$$

where ν_S and ν_Q are interpreted as the limiting proportion of jumps to levels S and Q, respectively. Hence,

$$k_1 = (\nu_S\theta_1 + \nu_Q\gamma_1) d_1(s) f_1(s) \tag{9}$$

is the long run rate of jumps from level s.

We next show that the expression for k_1 in (6) gives the same expression as in (9): From (4) (the case $s < Q$) we see that

$$d_1(0)f_1(0) = d_1(S)f_1(S) - d_1(s)f_1(s) =: c_0,$$

and from the solution to f_1 we see that $d_1(s)f_1(s) = c_0e^{\lambda D(s)} + k_1$. Substituting for $d_1(0)f_1(0)$ and $d_1(S)f_1(S)$ in the expression for k_1 in (6), we rewrite k_1 to get

$$k_1 = \frac{\gamma_1 c_0 + \theta_1 c_0 e^{\lambda D_1(s)} - \theta_1 c_0}{1-\theta_1}. \tag{10}$$

It is then a matter of simple algebra to show that the expression for k_1 in (10) is equal to

$$(1-\nu_S\theta_1 - \nu_Q\gamma_1)^{-1}(\nu_S\theta_1 + \nu_Q\gamma_1)c_0e^{\lambda D_1(s)},$$

for ν_S and ν_Q in (8). We now use the solution for f_1 once more to replace $c_0e^{\lambda D_1(s)}$. In particular, from $c_0e^{\lambda D_1(s)} = d_1(s)f_1(s) - k_1$ we get the desired equality, that is, k_1 in (10) is equal to the expression (9). This proves the claim. □

REMARK 3.2: The values of the terms in (6) have an intuitive interpretation. For example, the value of k_1 can be computed by conditioning on the last jump prior to hitting s, namely we condition on whether we started at level Q or S, where these conditional probabilities are $\gamma_1(s, Q)$ and $\theta_1(s, S)$, respectively. Then the long-run rate of hitting s, when starting in Q, is also the long-run rate of hitting level 0 from above, which is equal to $d_1(0)f_1(0)$. The long-run rate of hitting s when starting in S, is the long-run rate of downcrossing S, which is equal to $d(S)f_1(S)$. This logic gives the expression for k_1 in (6). Similar reasonings give us the expressions for $\tilde{k}_1$ and k_2.

3.4. Profit Functions Under OP$_1$ and OP$_2$

We can use the solutions for f_1 and f_2 and compute the long-run profit functions for both policies. We denote by $R_1 := R_1(s, S, Q, p(\cdot))$ the long-run average profit function generated by OP$_1$, and by $R_2 := R_2(s, S, p(\cdot))$ the long-run profit function generated by OP$_2$.

PROPOSITION 3.2: We have

$$\begin{aligned} R_1 = &\int_0^S [p(w)d_1(w) - h(w)]f_1(w)dw \\ &- [K + w_\mu(S-s)]k_1 - \lambda\int_0^s [K + w_\mu(S-w)] \\ &\times f_1(w)dw - (K + w_\lambda Q)d_1(0)f_1(0) \end{aligned} \tag{11}$$

and

$$\begin{aligned} R_2 = &\int_0^S [p(w)d_2(w) - h(w)]f_2(w)dw \\ &- [K + w_\mu(S-s)]k_2 - \lambda\int_0^s [K + w_\mu(S-w)] \\ &\times f_2(w)dw - (K + w_\mu S)\lambda\pi_2 - a\frac{d(0)f_2(0)}{\lambda}, \end{aligned} \tag{12}$$

for $\pi_2 = 1 - \int_0^S f_2(x)dx$ in (3).

PROOF: The first terms on the right-hand sides of (11) and (12), $\int_0^S [p(w)d_i(w) - h(w)]f_i(w)dw$, $i = 1, 2$, are the average income flowing into the system, since $[p(w)d_i(w) - h(w)]dw$ is the infinitesimal flow into the system whenever the content level is w.

The cost $[K + w_\mu(S-s)]$ is incurred every time level s is downcrossed and $W(t) = w_\mu$, that is, the state of the world is "cheap." Conditioning on the state of the content level just after the last jump, gives the long-run rate of downcrossing level s during a cheap period, as explained in the proof of Lemma 3.1.

The average ordering costs (Textranslationfailed), $i = 1, 2$, are paid after level s is downcrossed during an expensive period and the next cheap period starts before the content level drops to 0. The fact that the expensive period is exponentially distributed with rate λ implies that cheap periods arrive in accordance with a Poisson process with rate λ. Hence, the conditional ordering cost, given that the state is w, is $K + w_\mu(S-w)$ and the deconditioning is taken with respect to the steady state density by PASTA.

The last term on the right-hand side of R_1 is the ordering cost when the content level drops to 0 during an expensive period and an immediate order of size Q is placed. Again, $d(0)f_1(0)$ is the long-run average number of hitting level 0 from above.

The last two terms on the right-hand side of R_2 are associated with the atom of C at state 0. First, under OP$_2$ the controller will wait for the next cheap period to arrive, and then will place an order of size S. The rate of those ordering costs is $\lambda\pi_2$ by PASTA. Second, there is a cost $a(t_2 - t_1)$ for staying at state 0 over the interval $[t_1, t_2]$. Since the long-run average time between two hits of level 0 is $d(0)f_2(0)$, we have by renewal reward that

$$\frac{1/\lambda}{1/(d(0)f_2(0))} = \frac{d(0)f_2(0)}{\lambda}$$

is the long-run proportion of time spent in state 0. □

Under OP$_1$, the average ordering cost is $K + w_\mu E(S - C_1)$ when $W = w_\mu$, but the last order of each cycle is placed in an expensive period with the ordering cost being $K + w_\mu E(S - C_1)$. Under OP$_2$, all orders are placed in cheap periods with the expected ordering cost being $K + w_\mu E(S - C_1)$. In particular, the set-up cost of the last order in the cycle is $K + w_\mu S$.

4. NUMERICAL STUDY

We consider models with two sell prices, denoted by p_h and p_l, so that only one threshold q for switching from the high price p_h to the lower one p_l should be determined. Note that $q = S$ or $q = 0$ is possible, in which case only one sell price is employed. Letting d_l and d_h denote the demand rate whenever the sale price is p_l and p_h, respectively, we have that $C > q$ implies a demand rate d_l, and $C \leq q$ implies a demand rate d_h. We use a linear demand model $d(p) = 50 - p$, with domain $[0, 50 - 10^{-3}]$, and linear holding costs $h(x) = hx$, for $h > 0$. In the plots in Fig. 1, we plot the sensitivity of the optimal profit with respect to changes in one of the parameters $(h, K, w_\mu, w_\lambda, \mu, \lambda, a)$. For different parameter values, we calculate the optimal (p_h, p_l, q, s, Q, S) under the policies OP_1 and OP_2. We also compare their performance with an *(s, S)* policy, denoted by OP0, under which an order of size $S-s$ is placed whenever the content process hits level s. In particular, under OP0 the parameters (s, S, p_h, p_l) are optimized without taking the stochastic fluctuations of the procurement price into account. Instead, the procurement price under OP0 is taken to be weighted average of w_μ and w_λ.

4.1. Scenario 1: $(h, K, w_\mu, w_\lambda, \mu, \lambda, a) = (7, 233, 3.4, 43, 0.7, 0.05, 5)$

In this scenario, the cheap periods are relatively rare, with a very cheap price. OP_2 performs slightly better than OP_1, and both outperform OP0. Table 1 lists the optimal profit and decision variables for the order policies OP0, OP_1, and OP_2. Figure 1 shows sensitivity of the optimal profits w.r.t. changes in the parameters $(h, K, w_\mu, w_\lambda, \mu, \lambda, a)$. For all policies, the profit is decreasing in *h, K*, w_μ, w_λ, and μ, and increasing in λ. The profit of OP0 and OP_1 does not depend on a; for OP_2, the optimal profit is decreasing in a. Clearly, taking the fluctuating procurement prices into consideration make a big difference, as no profit can be make under OP0.

4.2. Scenario 2: $(h, K, w_\mu, w_\lambda, \mu, \lambda, a) = (5, 100, 20, 25, 0.1, 0.05, 1)$

Here the difference between cheap and expensive price is less extreme, and cheap periods last longer. OP_1 performs slightly better than OP0, and both outperform OP_2. Table 2 lists the optimal profit and decision variables for the order policies OP0, OP_1, and OP_2. Figure 2 shows sensitivity of the optimal profits w.r.t. changes in the parameters $(h, K, w_\mu, w_\lambda, \mu, \lambda, a)$. For all policies, the profit is decreasing in *h, K*, w_μ, w_λ, and μ, and increasing in λ. The profit of OP0 and OP_1 does not depend on a; for OP_2, the optimal profit is decreasing in a.

5. GENERALIZATIONS

In this section, we present two generalizations for the basic model analyzed above for the OP_2 policy. We first consider a model having the same structure as the basic model, but with a random environment process that is more general. We then consider a model with exponential lead times, that is, when there is a positive random time from the moment an order is made by the controller until the commodity arrives.

5.1. Phase-Type Expensive Periods

Our analysis can be extended to the case in which one of the periods, either the cheap or the expensive period, follows a phase-type distribution; in particular, the state-of-the-world process W evolves as a CTMC with more than two states. We employ simple martingale arguments (the optional stopping theorem for an appropriate Wald's martingale; see below)

For simplicity of exposition, we consider a model in which the expensive period is distributed as the sum of two independent exponential random variables. Specifically, assume that the cheap period is exponentially distributed with rate μ, and that the expensive period is a sum of two independent exponential random variables X_1 and X_2, with X_i having mean $1/\lambda_i$, $i = 1, 2$. We refer to X_1 and X_2 as the first and second phase of the expensive period, respectively.

We designate the probabilities that level s is downcrossed at stationarity by the first phase and the second phase of the expensive period, respectively, by ξ_1 and ξ_2. In the next theorem, we introduce the balance equation of the content level in terms of $f(S)$, ξ_1, and ξ_2. The computations of these quantities is carried out below.

LEMMA 5.1: *For $T := D(s) - D(x)$ it holds that*

$$d(x)f(x) = \begin{cases} d(S)f(S)\Big[\xi_1\left(\frac{\lambda_1}{\lambda_1-\lambda_2}e^{-\lambda_2 T} - \frac{\lambda_2}{\lambda_2-\lambda_1}e^{-\lambda_1 T}\right) \\ \quad + \xi_2 e^{-\lambda_2 T}\Big] \quad 0 < x < s, \\ d(S)f(S) \quad s \leq x \leq S, \end{cases}$$

PROOF:

(i) $s \leq x \leq S$. In this region, every downcrossing of level x is followed by a downcrossing of level S with no jump in between. Thus, the long-run average number of downcrossings of level x is equal to that of the long run average number of downcrossings of level S, so that $d(x)f(x) = d(S)f(S)$.

(ii) $0 < x < s$. For every x, we mark a downcrossing of level s as a downcrossing of type i if level s is downcrossed during phase i of the expensive period, $i = 1, 2$. It follows from the definition that ξ_1 is the probability that the time until the next jump is a convolution of two exponential random variables with rates λ_1 and λ_2. Similarly, ξ_2 is the probability that the time to the next jump is exponential with rate λ_2.

The probability of a type-1 downcrossing is therefore the probability that the next jump will occur of two independent exponential random variables with rates λ_1 and λ_2, is

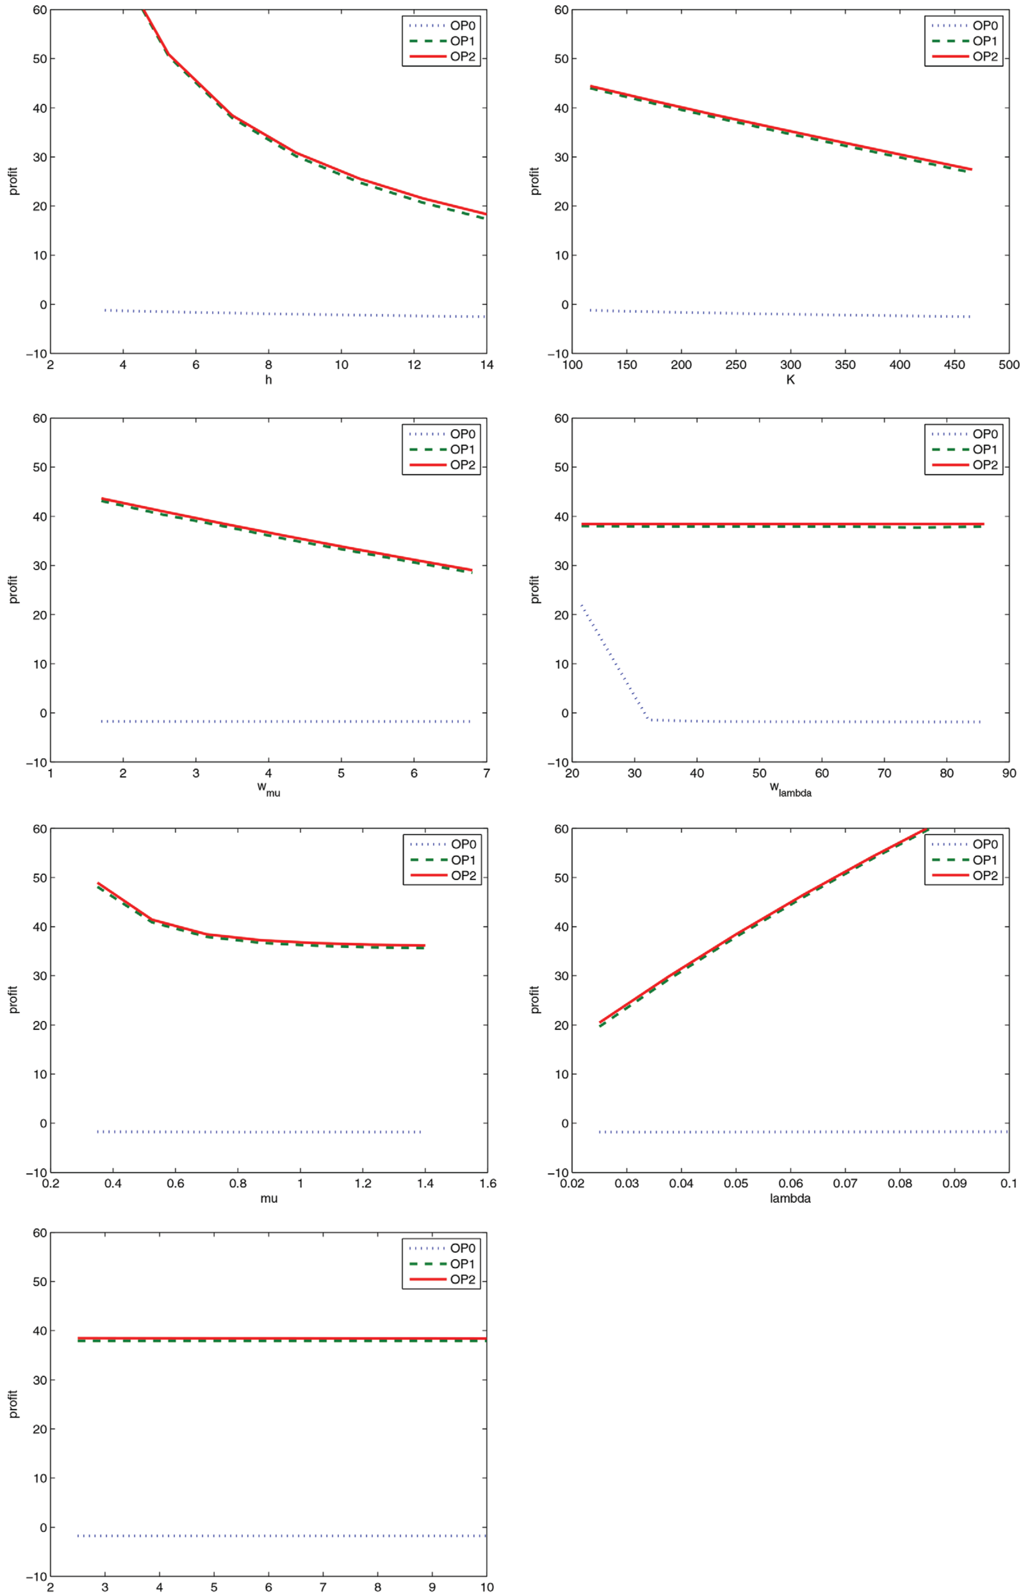

Figure 1. Sensitivity analysis for scenario 1. [Color figure can be viewed at wileyonlinelibrary.com]

Table 1. Profit and optimal solution under different order policies for scenario 1

Order policy	Profit	p_l	p_h	q	s	Q	S
OP0	−1.76	46.79	50.00	0.25	0		3.20
OP_1	37.92	33.10	50.00	0.09	6.01	0.09	61.05
OP_2	38.45	33.10	50.00	0.01	5.94		60.97

Table 2. Profit and optimal solution under different order policies for scenario 2

Order policy	Profit	p_l	p_h	q	s	Q	S
OP0	68.93	37.90	40.37	9.51	0		21.46
OP_1	69.12	37.78	40.37	9.99	0	20.5741	23.53
OP_2	38.85	37.32	50.00	0.01	0.01		25.06

larger than $T := D(s) - D(x)$, conditional on level s being downcrossed during the first phase X_1. This gives the first expression in the square brackets. Similarly, with probability ξ_2 level s is downcrossed during the second phase of the expensive period, and with probability $e^{-\lambda_2 T}$ no jump occurs between the latter two downcrossings. □

Computing ξ_1 and ξ_2

To compute the probabilities ξ_1 and ξ_2, we construct an auxiliary proces

$$\chi(t) := t + S_1 + \cdots + S_{N(t)}, \quad t \geq 0, \quad \text{with } \chi(0) = 0,$$

where S_i, $i \geq 1$, are independent and identically distributed random variables, each having the distribution of the expensive period, in particular, the Laplace transform of each S_i is

$$\tilde{G}(\alpha) = \frac{\lambda_1}{\lambda_1 + \alpha} \cdot \frac{\lambda_2}{\lambda_2 + \alpha},$$

and $\{N(t) : t \geq 0\}$ is a Poisson process with rate μ. Then $\sum_{j=1}^{N(t)} S_j$ is a compound Poisson process and χ is a nondecreasing process that increases either linearly at rate 1 between jumps, or by positive jumps of (random) size S, where S is a generic random variable with the distribution of S_1.

We can think of each jump of χ as having two phases: The first phase is distributed exponentially with rate λ_1, and the second exponentially with rate λ_2. The process χ can thus leave the interval $[0, D(S) - D(s))$ in three ways: (i) attaining the boundary point $D(S) - D(s)$ on a linear segment of the path, (ii) upcrossing level $D(S) - D(s)$ by the first phase of the jump, and (iii) upcrossing level $D(S) - D(s)$ by the second phase of the jump.

Define the stopping time

$$\tau := \inf\{t > 0 : \chi(t) \geq D(S) - D(s)\}$$

and consider Wald's martingale associated with the process χ

$$M_\alpha(t) := \frac{e^{-\alpha\chi(t)}}{E[e^{-\alpha\chi(t)}]} = e^{-\alpha\chi(t) - \varphi(\alpha)t}, \quad \alpha \notin \{-\lambda_1, -\lambda_2\}, \tag{13}$$

where

$$\varphi(\alpha) := -\left[\alpha + \mu\left(1 - \frac{\lambda_1}{\lambda_1 + \alpha} \cdot \frac{\lambda_2}{\lambda_2 + \alpha}\right)\right]. \tag{14}$$

Clearly $M_\alpha(t)$ is bounded, so the optional stopping theorem can be applied, yielding $E[M_\alpha(0)] = E[M_\alpha(\tau)]$, so that

$$1 = E[e^{-\alpha\chi(\tau) - \varphi(\alpha)\tau}]. \tag{15}$$

Let B_0, B_1, and B_2 be the events that χ reaches level $D(S) - D(s)$ by the drift, upcrossed by the first phase of the jump and upcrossed by the second phase of the jump, respectively. It follows from the memoryless property of the exponential random variable that if level $D(S) - D(s)$ is upcrossed by the first phase of the jump, then $\chi(\tau) \stackrel{d}{=} D(S) - D(s) + X_1 + X_2$, where $\stackrel{d}{=}$ denotes equality in distribution. If level $D(S) - D(s)$ is upcrossed by the second phase of the jump, then $\chi(\tau) \stackrel{d}{=} D(S) - D(s) + X_2$. Finally, if level $D(S) - D(s)$ is upcrossed by the continuous drift of χ, then $\chi(\tau) \stackrel{d}{=} D(S) - D(s)$. Hence, by (15),

$$\begin{aligned} 1 &= \sum_{i=0}^{2} E[e^{-\alpha\chi(\tau) - \varphi(\alpha)\tau} 1_{B_i}] \\ &= e^{-\alpha(D(S)-D(s))} E[e^{-\varphi(\alpha)\tau} 1_{B_0}] \\ &\quad + \frac{\lambda_1}{\lambda_1 + \alpha} \frac{\lambda_2}{\lambda_2 + \alpha} e^{-\alpha(D(S)-D(s))} E[e^{-\varphi(\alpha)\tau} 1_{B_1}] \\ &\quad + \frac{\lambda_2}{\lambda_2 + \alpha} e^{-\alpha(D(S)-D(s))} E[e^{-\varphi(\alpha)\tau} 1_{B_2}], \end{aligned} \tag{16}$$

for $\varphi(\alpha)$ in (14).

To obtain the probabilities ξ_0, ξ_1, and ξ_2 we substitute $\varphi(\alpha) = 0$ in (16). It is easy to see from (14) that $\varphi(\alpha)$ has three roots, with one of them, denoted by α_0, being 0. The other two roots, denoted by α_1 and α_2, are the solutions to the quadratic equation

$$\alpha^2 + (\lambda_1 + \lambda_2 + \mu)\alpha + \lambda_1\mu + \lambda_2\mu + \lambda_1\lambda_2 = 0. \tag{17}$$

Inserting the roots α_i into (16) yields three equations with the three unknowns ξ_i, $i = 1, 2, 3$. Formally,

$$\begin{aligned} 1 &= e^{-\alpha_i(D(S)-D(s))} E[1_{B_0}] \\ &\quad + \frac{\lambda_1}{\lambda_1 + \alpha_i} \frac{\lambda_2}{\lambda_2 + \alpha_i} e^{-\alpha_i(D(S)-D(s))} E[1_{B_1}] \end{aligned}$$

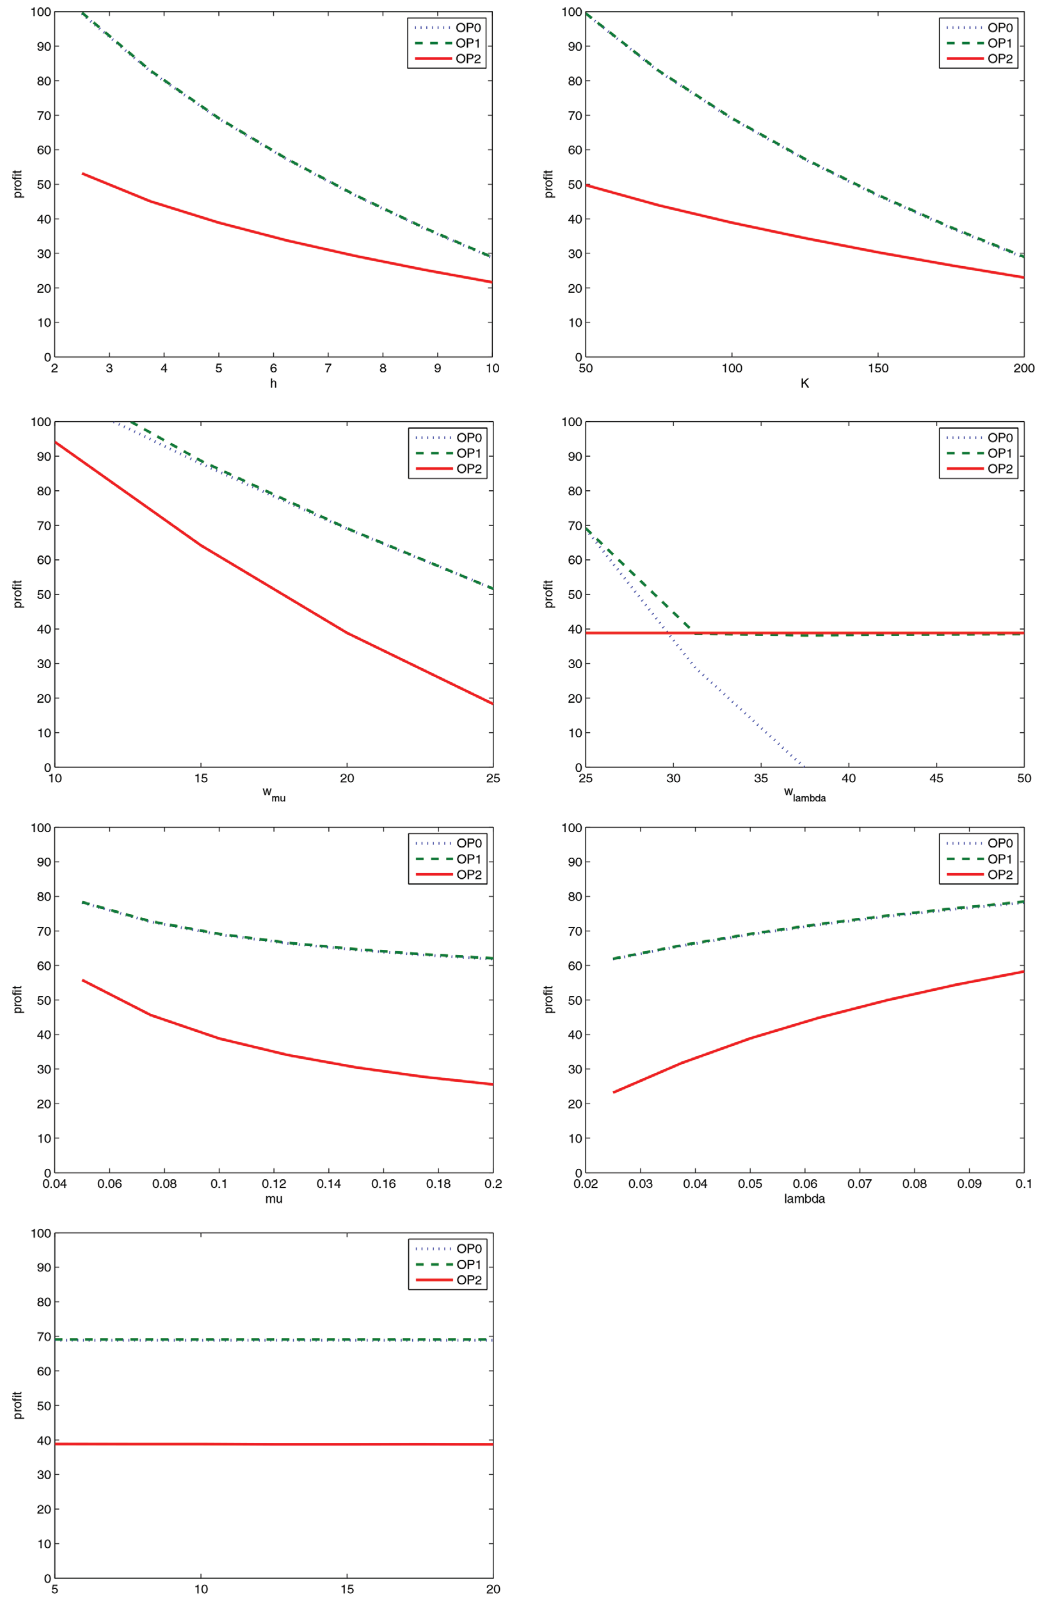

Figure 2. Sensitivity analysis for scenario 2. [Color figure can be viewed at wileyonlinelibrary.com]

$$
\begin{aligned}
&+ \frac{\lambda_2}{\lambda_2+\alpha_i} e^{-\alpha_i(D(S)-D(s))} E[1_{B_1}] \\
&= e^{-\alpha_i(D(S)-D(s))}\xi_0 + \frac{\lambda_1}{\lambda_1+\alpha_i}\frac{\lambda_2}{\lambda_2+\alpha_i} e^{-\alpha_i(D(S)-D(s))}\xi_1 \\
&+ \frac{\lambda_2}{\lambda_2+\alpha_i} e^{-\alpha_i(D(S)-D(s))}\xi_2, \qquad (18)
\end{aligned}
$$

where the equation associated with $\alpha_0 = 0$ gives the normalizing condition $\xi_0 + \xi_1 + \xi_2 = 1$.

Computing f(S)

With ξ_0, ξ_1, and ξ_2 in hand, we can express $f(x)$ in terms of $f(S)$ via the equations in Lemma 5.1. Finally, we use the normalizing condition

$$\int_0^S f(x)dx = 1 - \pi,$$

where π is the atom at 0. Specifically,

$$f(x) = \begin{cases} k\frac{d(s)}{d(x)} f(S) & 0 < x < s, \\ \frac{d(s)}{d(x)} f(S) & s \le x \le S, \end{cases} \qquad (19)$$

where

$$
\begin{aligned}
k = \xi_1 &\left(\frac{\lambda_1 e^{-\lambda_2(D(S)-D(s))}}{\lambda_1-\lambda_2} - \frac{\lambda_2 e^{-\lambda_2(D(S)-D(s))}}{\lambda_1-\lambda_2}\right) \\
&- \xi_2 e^{-\lambda_2(D(S)-D(s))}. \qquad (20)
\end{aligned}
$$

To compute π, let I be the interval of time in which the inventory system is empty. Then

$$\pi = d(0) f(0) E[I],$$

since $1/d(0)f(0)$ is the expected cycle length so that, by renewal theory, $E[I]d(0)f(0)$ is the long run proportion of time that the inventory system is empty.

To compute $E[I]$, we consider the possibilities of reaching level 0.

1. Level s is reached during the first phase of the expensive period and from here during the next $D(s)$ time units the first phase is not changed. The probability of the latter event is $\xi_1 e^{-\lambda_1 D(s)}$. Once level 0 is reached, the expected time until an order is placed is $1/\lambda_1 + 1/\lambda_2$.
2. Level s is reached during the first phase of the expensive period, but during the next $D(s)$ time units the first phase end, so that level 0 is reached during the second phase. The probability of the latter event is

$$
\begin{aligned}
\xi_1 &\int_0^{D(s)} \lambda_1 e^{-\lambda_1 u} e^{-\lambda_2(D(s)-u)} du \\
&= \frac{\xi_1\lambda_1}{\lambda_1-\lambda_2}\left(e^{-\lambda_2 D(s)} - e^{-\lambda_1 D(s)}\right).
\end{aligned}
$$

   Once level 0 is reached, the expected time at level 0 is $1/\lambda_2$.
3. Level s is reached during the second phase of the expensive period and during the next $D(s)$ time units the second phase does not end. The probability of the latter event is $\xi_2 e^{-\lambda_2 D(s)}$. Then level 0 is reached during the second phase, and the expected time at 0 is $1/\lambda_2$.

Therefore,

$$
\begin{aligned}
E[I] = &\left(\frac{1}{\lambda_1}+\frac{1}{\lambda_2}\right)\xi_1 e^{-\lambda_1 D(s)} \\
&+ \frac{1}{\lambda_2}\frac{\xi_1\lambda_1}{\lambda_1-\lambda_2}\left(e^{-\lambda_2 D(s)} - e^{-\lambda_1 D(s)}\right) \\
&+ \frac{1}{\lambda_2}\xi_2 e^{-\lambda_2 D(s)},
\end{aligned}
$$

so that

$$
\begin{aligned}
\pi = d(0)f(0)\Bigg[&\left(\frac{1}{\lambda_1}+\frac{1}{\lambda_2}\right)\xi_1 e^{-\lambda_1 D(s)} \\
&+ \frac{1}{\lambda_2}\frac{\xi_1\lambda_1}{\lambda_1-\lambda_2}\left(e^{-\lambda_2 D(s)} - e^{-\lambda_1 D(s)}\right) + \frac{1}{\lambda_2}\xi_2 e^{-\lambda_2 D(s)}\Bigg],
\end{aligned}
$$

where by (19),

$$f(0) = k\frac{d(s)}{d(0)} f(S),$$

for k in (20), and $f(S)$ is obtained via

$$1 - \pi = kd(s)f(S)\int_0^s \frac{1}{d(x)}dx + d(s)f(S)\int_s^S \frac{1}{d(x)}dx.$$

In particular,

$$f(S) = \frac{1-\pi}{kd(s)\int_0^s \frac{1}{d(x)}dx + d(s)\int_s^S \frac{1}{d(x)}dx}.$$

5.2. Exponential Lead Times

We assume exponential leadtime with parameter η. When there are positive leadtimes, it makes sense to modify the control by considering two levels in which, when downcrossed, the controller should place an order. We thus have three critical levels $0 < s_0 < s_1 < S$. The cycle starts with $C(0) = S$. Then the content level decreases at rate $d(x)$ without any jumps until it reaches level s_1. If level s_1 is reached during a cheap period, an order is placed and it takes an exponential amount of time with rate η until it arrives. Otherwise, if level s_1 is reached during an expensive period, no order is placed and the content level decreases until the expensive period is terminated and replaced by a cheap period or until level s_0 is reached. In any case, when level s_0 is reached (either during a cheap period or an expensive period) an order is placed and arrives after an exponential period of time with rate η.

THEOREM 5.1: Let *f(x)* denote the steady state density of the content level C, and let *F(x)* denote the corresponding cumulative distribution function. Then *f(x)* satisfies the integral equation

$$d(x)f(x)=\begin{cases}\eta F(x), \quad 0 \le x < s_0,\\ \eta F(s_0)+\eta[\gamma+(1-\gamma)(1-e^{-\lambda[D(s_1)-D(x)]})]\\ \quad [F(x)-F(s_0)], \quad s_0 \le x < s_1,\\ \eta F(s_0)+\eta[\gamma+(1-\gamma)(1-e^{-\lambda[D(s_1)-D(s_0)]})]\\ \quad [F(s_1)-F(s_0)], \quad s_1 \le x \le S,\end{cases}$$

where γ is the probability that level s_1 is downcrossed during the cheap period.

PROOF:

(i) $0 \le x < s_0$. In this region, the order is on its way. Since the lead time is exponentially distributed the arrival process can be interpreted as a Poisson process with rate η.

(ii) $s_0 \le x < s_1$. The jump may occur below s_0 or above s_0. If the content level is below s_0, jumps arrive with rate $\eta F(s_0)$. If the content level is above s_0, then there are two possibilities: With probability γ level s_1 is downcrossed during a cheap period and an order is placed immediately; it will arrive after an exp(η) period of time. With probability $1-\gamma$ level s_1 is downcrossed during an expensive period and no order is placed. However, if during the time period from downcrossing of level s_1 until level x is reached the procurement price is changed from expensive to cheap an order will be placed and it will take an exp(η) period until the order arrives (the probability of the latter event is $1-e^{-\lambda[D(s_1)-D(x)]}$). For either possibility, the probability that the jump occurs at some level between s_0 and x is $F(x)-F(s_0)$.

(iii) $s_1 \le x < S$. In this region, we note that no jumps starts when the content level is above level s_1. We thus have to distinguish between two possibilities. If the content level is below level s_0 the rate of the jumps is $\eta F(s_0)$. If the content level is above level s_0 the rate of the jumps is

$$\eta[\gamma+(1-\gamma)(1-e^{-\lambda[D(s_1)-D(s_0)]})][F(s_1)-F(s_0)].$$

□

To compute γ, we extend the argument of the previous section. Level S can be reached either during a cheap period or an expensive period. Since after every jump the content level is equal to S we define the embedded chain

$$P=\begin{pmatrix} p_{cc} & 1-p_{cc}\\ 1-p_{ee} & p_{ee}\end{pmatrix},$$

where p_{cc} is the conditional probability that the next jump occurs during a cheap period given that the present procurement price is cheap and the state is S. Similarly, p_{ee} is the conditional probability that the next jump occurs during an expensive period given that the present procurement price is expensive and the state is S. Then the solution (ζ_1, ζ_2) to the equations

$$(\zeta_1,\zeta_2)\begin{pmatrix} p_{cc} & 1-p_{cc}\\ 1-p_{ee} & p_{ee}\end{pmatrix}=(\zeta_1,\zeta_2) \quad \text{and} \quad \zeta_1+\zeta_2=1,$$

is the solution of the conditional steady state probability—ξ_1 (ξ_2) that level s_1 is downcrossed during a cheap period (expensive period), given that at the starting point, that is, at level S, the procurement price is cheap (expensive). Finally

$$\gamma=\zeta_1 p_{cc}+\zeta_2(1-p_{ee}).$$

Computing p_{cc} and p_{ee} is similar to the computations in Lemma 3.3.

6. SUMMARY

We considered Markovian fluid-inventory models operating in a random environment, governed by the procurement price. For the first model, in which the procurement price changes in accordance with a CTMC, two natural *(s, S)*-type policies were considered, and the respective stationary distributions for the random fluid content process were computed. In turn, those stationary distributions can be used to optimize the control parameters for each policy. Two generalizations, the first involving a more complex random price environment, and the second incorporating lead times, were also analyzed.

There are two immediate related problems to address. First, the complexity of the optimization problem requires developing an efficient algorithm to solve for the optimal solutions in more complex settings than those considered in our numerical examples here. Second, since it is unlikely that an overall optimal policy can be found, Brownian approximations for the underlying regenerative process may be developed to obtain asymptotically-optimal control mechanisms. We leave those problems for future research.

ACKNOWLEDGMENTS

This research is conducted while A.d.B. and O.P. were affiliated at CWI, and is made possible by an NWO VIDI grant. O.P. was supported by NSF grant CMMI 1436518.

REFERENCES

[1] B. Ata and S. Shneorson, Dynamic control of an M/M/1 service system with adjustable arrival and service rates, Manage Sci 52 (2006), 1778–1791.
[2] O. Berman, D. Perry, and W. Stadje, A fluid EOQ model with a two-state random environment, Probab Eng Inf Sci 20 (2006), 329–349.
[3] O. Boxma, D. Perry, and S. Zacks, A fluid EOQ model of perishable items with intermittent high and low demand rates, Math Oper Res 40 (2014), 390–402.
[4] D. Beyer, C. Feng, S.P. Sethi, and M. Taksar, Markovian demand inventory models, Springer, New York, 2010.
[5] P.H. Brill, Level crossing methods in stochastic models (2008), Springer Verlag, 2008.
[6] P.H. Brill and B.A. Chaouch, An EOQ model with random variations in demand, Manage Sci 41 (1995), 927–936.
[7] S. Browne and P. Zipkin, Inventory models with continuous, stochastic demands, Ann Appl Prob 1 (1991), 419–435.
[8] B. Chaouch, Inventory control and periodic price discounting campaigns, Nav Res Logist 54 (2007), 94–108.
[9] F. Chen and J.S. Song, Optimal policies for multiechelon inventory problems with Markov-modulated demand, Oper Res 49 (2001), 226–234.
[10] H. Chen and D.D. Yao, A fluid model for systems with random disruptions, Oper Res 40 (1992), 239–247.
[11] F. Cheng and S.P. Sethi, Optimality of (s, S) policies in inventory models with Markovian demand, Oper Res 45 (1997), 931–939.
[12] F. Cheng and S.P. Sethi, Optimality of state-dependent (s, S) policies in inventory models with Markov-modulated demand and lost sales, Prod Oper Manage 8 (1999), 183–192.
[13] E. Cinlar and J. Jacod, Representation of semimartingale Markov processes in terms of Wiener processes and Poisson random measures, Semin Stochastic Process (1981), 159–242. Birkhäuser.
[14] J.W. Cohen, On up- and downcrossings, J Appl Prob 14 (1977), 405–410.
[15] M.H.A. Davis, Piecewise-deterministic Markov processes: A general class of non-diffusion stochastic models, J R Stat Soc 46 (1984), 353–388.
[16] S. Dewan and H. Mendelson, User delay costs and internal pricing for a service facility, Manage Sci 36 (1990), 1502–1517.
[17] G. Gallego and G. Van Ryzin, Optimal dynamic pricing of inventories with stochastic demand over finite horizons, Manage Sci 40 (1994), 999–1020.
[18] M. Goh and M. Sharafali, Price-dependent inventory model with discount offers at random times, Prod Oper Manage 11 (2002), 139–156.
[19] M. Hu and Y. Yang, Modified echelon (r, Q) policies with guaranteed performance bounds for stochastic serial inventory systems, Oper Res 62 (2014), 812–828.
[20] D. Iglehart and S. Karlin, "Optimal policy for dynamic inventory process with nonstationary stochastic demands," in: K.J. Arrow, S. Karlin, and H. Scarf (Editors), Studies in applied probability and management science, Stanford University Press, Stanford, California, 1962, Chapter 8, pp. 127–148.
[21] O. Kella and W. Whitt, A storage model with a two-state random environment, Oper Res 40 (1992), S257–S262.
[22] H. Mendelson, Pricing computer services: Queueing effects, Commun ACM 28 (1985) 312–321.
[23] M. Miyazawa, Rate conservation laws: A survey, Queue Syst 15 (1994), 1–58.
[24] K. Moinzadeh, Replenishment and stock policies for inventory systems with random deal offerings, Manage Sci 43 (1997), 334–342.
[25] S. Nahmias, Production and operations analysis, McGraw-Hill/Irwin, New York, 2005.
[26] P. Naor, The regulation of queue size by levying tolls, Econometrica 37 (1969), 15–24.
[27] K.S. Puranam and M.N Katehakis, On optimal bidding and inventory control in sequential procurement auctions: The multi period case, Ann Oper Res 217 (2014), 447–462.
[28] J. Shi, M.N. Katehakis, and B. Melamed, Martingale methods for pricing inventory penalties under continuous replenishment and compound renewal demands, Ann Oper Res 208 (2013), 593–612.
[29] J. Shi, M.N. Katehakis, B. Melamed, and Y. Xia, Optimal continuous replenishment for a production-inventory system with compound Poisson demands and lost sales, Oper Res 62 (2014), 1048–1063.
[30] J.S. Song and P. Zipkin, Inventory control in a fluctuating demand environment, Oper Res 4 (1993), 351–370.
[31] W. Whitt, Stochastic process limits, Springer, New York, 2002.
[32] W. Whitt, Internet supplement to stochastic-process limits, 2002, Available at: http://www.columbia.edu/∼ww2040/supplement.html. Last accessed on February, 2017.
[33] R.W. Wolff, Stochastic modeling and the theory of queues, Prentice Hall, New Jersey, 1989.