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

Arnoud den Boer* Ohad Perry† Bert Zwart‡

December 18, 2012

Abstract

We study a single-product pricing and inventory model in which the price of the cost 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 of inventory occurs at a deterministic rate, which is determined by the sell price of the product. Hence, at any time, the controller has to simultaneously decide the selling price of the product and whether to order or not, taking into account the current cost price of the product and the inventory level. We consider two policies, derive the associated steady state distributions and cost functionals, and apply these to study these policies.

1 Introduction

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

*Department Stochastic Operations Research, University of Twente, P.O. Box 217, 7500 AE Enschede; The Netherlands. Email: a.v.denboer@utwente.nl

†Industrial Engineering and Management Sciences McCormick School of Engineering Northwestern University, Evanston, IL 60208. Email: ohad.perry@northwestern.edu

‡Centrum Wiskunde & Informatica, Science Park 123, Amsterdam, The Netherlands. Email: bert.zwart@cwi.nl

This is in accordance with current practice of dynamic pricing, where sell prices are not fixed quantities, but may change continuously.

Determining optimal order policies and sell prices are often treated as separate problems, but it is intuitively not difficult to see that it may be beneficial to consider these problems simultaneously. For example, if the cost price of new items currently is high, it may be profitable to increase the sell price, such that the moment at which all inventory is sold-out, is delayed. This increases the probability that in the mean time, the cost price of new items decreases, such that new items can be ordered at considerable lower costs.

We study a model with continuous inventory, in continuous time. 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. The cost price at which new items can be acquired is modeled as a finite-state Markov chain, where each state represents a different cost price. Every time an order is placed the seller pays some fixed order costs K, and any moment that the inventory-level is $x > 0$, the seller faces holding costs at a rate $h(x)$. Ordered items are assumed to arrive instantaneously. The inventory-level decreases at a deterministic demand rate, which depends on the sell price.

To maintain tractability, we make a number of assumptions on the cost price process, the order policy, and the sell price policy. In particular, we initially assume that the cost price varies between two prices only, a low and a high cost price. For the order policy, we study two variants of the well-known (s, S)-policy. In the first order policy, OP1, $S - x$ items are ordered if the inventory-level x is at or below s and at the same time the cost price is low. If the inventory hits zero and the cost 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, OP2, is similar to OP1 except that never orders are placed when the cost price is high. If the inventory-level hits zero, the seller waits until the cost price becomes low, at which moment he orders S items. We assume that the seller uses a sell price policy of the following type: if the inventory-level exceeds q, a high sell price p_h is charged. Otherwise, a low sell-price p_l is charged. Here $q \geq 0$, $0 < p_l < p_h$ are decision variables.

We consider the pricing-and-inventory problem in stationarity. Under mild assumptions on the relation between demand rate and sell price, we show that the joint process of inventory level and cost price admits a unique stationary distribution. For a fixed order policy OP1 or OP2, we derive balance equations for the stationary distribution of the inventory-level process, and calculate an explicit solution. This enables us to calculate the long-run profit for both policies, as function of (s, S, Q, p_l, p_h, q) in case of OP1, and (s, S, p_l, p_h, q) in case of OP2. To determine the optimal values of the decision variables, one needs to solve a (rather complicated) non-convex non-linear optimization problem.

We conduct a numerical study to compare the performance of OP1 and OP2. We also compare them to a standard (s, S)-policy OP0, which does not take into account the random nature of the cost price process. By studying several instances, it turns out that OP1 in general performs better or equal than both OP2 and OP0. The difference in performance, especially between OP1 and OP0, can be quite large. This shows that it is beneficial to take into account random changes in the cost prices. The policies OP0 and OP2 have no clear 'best': for some instances, the first is outperformed by the latter, while for other instances it is the other way around. We also study the sensitivity of the profit functions with respect to changes in the parameters.

The remainder of this paper is organized as follows: In §2 we describe the model and motivate the structure of the (s, S) (or (s, Q, S)) control policies. In §3 we develop the steady-state equations for the content level process. Those equations are then applied in §**??** in a numerical study, as described above. In §**??** we extend the model and consider cases in which the cost price of the item changes after non-exponential random time in states, and we also consider lead times.

2 Model and Assumptions

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. Following the terminology in [**?**] and [**?**], we refer to cost price as the "state of the world". In particular, the cost 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) cost price. It is thus clear that the "state-of-the-world" process W may effect the decision of the controller whether or not to buy at each decision epoch in order to replenish his inventory.

Let $C := \{C(t) : t \geq 0\}$ denote the content-level process. 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 $q_1 < q_2 < \cdots < q_{k-1} = 0$, 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.

For simplicity of the exposition, in this study we restrict attention to a model consisting of two sell prices, so that only one threshold q should be determined, although we do not rule out the cases in which $q = S$ or $q = 0$, so that only one sell price is employed. Generalizing the problem to more sell prices is straightforward. We are hence looking for a threshold q (that should be optimized) such that, whenever $C > q$, the sale price is p_l (low), and is p_h (high) whenever $C \le q$. 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 $W \le q$ implies a demand rate d_h.

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 cost 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 cost price and zero lead-time assumptions, it is desirable to place most of the orders, if not all of them, when the cost 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 $W(t) = w_\lambda$, i.e., 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 $Q \le S$ or (ii) wait for the cost price to change from w_λ to w_μ.

We thus consider two natural ordering policies:

Order Policy 1 (OP1). Determine two levels s and S. If the content level C hits s and at the same time the cost price is low, i.e., $C(t-) = s$ and $W(t-) = w_\mu$, then place an order of size $S - s$ (so that $W(t) = S$. If, on the other hand, upon hitting level s the cost price is high, i.e., $C(t-) = s$ and $W(t-) = w_\lambda$, then wait until either (i) the cost 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 (OP2). Similarly to OP1, except that never place an order while the cost price is high, i.e., whenever $W = w_\lambda$. When level 0 is hit (and it can only be reached during expensive periods) wait until the cost price changes to cheap (w_μ), at which point order up to level S. Note that, under OP2, 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, etc. 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 OP1 the control is determined by the decision variables (s, S, q, Q, p_l, p_h), while under OP2 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 OP1, and $C_2 := \{C_2(t) : t \geq 0\}$, denote the content-level process under OP2. 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).

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 OP1 and OP2, respectively. For $x \in [0, S]$ let $d_1(p_1(x))$ and $d_2(p_2(x))$ denote the respective demand functions. With an abuse of notation (based on our assumption about the relation between the price and the demand), we treat $d_i(\cdot)$ as a function of $x \in [0, S]$, denoted as $d_i(x)$, $i = 1, 2$.

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

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

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 $0 < x \le S$, if the input is shut off, i.e., 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 $X_i := \{X_i(t) : t \ge 0\} := \{(C_i(t), W(t)) : t \ge 0\}$ is a two-dimensional Markov process with state space $\mathcal{S} := [0, S] \times \{w_\lambda, w_\mu\}$. Since X_i is a Markov process on a general state space, the existence of a unique stationary distribution is not immediate. However, it is simple to show that X is regenerative and posseses 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 regenerative and admits a unique stationary distribution.*

Proof. First, it is easy to see that X will return to state $x^* := (S, w_\mu)$ in finite time, given our assumptions on the model. In particular, the expected return time to state x^* is finite. Moreover, X has a nonlattice distribution. That is easy to see in OP2, since X spends an exponential amount of time with mean $1/\lambda$ in state $(0, w_\lambda)$ (and by Assumption 1, X will reach that state with probability 1). That is also easy to see if OP1 is employed, since then there are random jumps each time C_1 hits level s during an expensive period, and W changes to "cheap" before C_1 hits level 0. ∎

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 $y > S - s$, i.e., that the content level can go below level s. However, OP2 requires that the content level can reach level zero in finite time. ∎

3.1 Steady-State Balance 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), i.e., $C(\infty)$ 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, e.g., [**?**] 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 OP1, 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, i.e., 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 OP2.

The main difficulty in our model is in determining the long-run rate of jumps from level s, i.e., 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 3.2 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, i.e.,

$$\pi_2 := P(C_2(\infty) = 0) > 0. \tag{2}$$

Theorem 3.1. (**balance equations**) *The steady state density $f_1(x)$ of C_1 satisfies one of the following* *integral equations, depending on whether $s \le Q$ or $s > Q$:*

$$\underline{\textit{If } s \le Q\textit{:}} \qquad d_1(x)f_1(x) \;=\; \begin{cases} \lambda \int_0^x f_1(w)\,dw + d_1(0)f_1(0), & 0 \le x < s, \\[2ex] \lambda \int_0^s f_1(w)\,dw + d_1(0)f_1(0) + k_1, & s \le x < Q, \\[2ex] \lambda \int_0^s f_1(w)\,dw + k_1, & Q \le x \le S. \end{cases}$$

$$\underline{\textit{If } s > Q\textit{:}} \qquad d_1(x)f_1(x) \;=\; \begin{cases} \lambda \int_0^x f_1(w)\,dw + d_1(0)f_1(0), & 0 \le x < Q, \\[2ex] \lambda \int_0^s f_1(w)\,dw, & Q \le x < s, \\[2ex] \lambda \int_0^s f_1(w)\,dw + \tilde{k}_1, & s \le x \le S. \end{cases} \tag{3}$$

The steady-state density $f_2(x)$ of C_2 satisfies the integral equation

$$d_2(x)f_2(x) = \begin{cases} \lambda \int_0^x f_2(w)\,dw + \lambda\pi_2, & 0 \le x < s, \\[2ex] \lambda \int_0^s f_2(w)\,dw + \lambda\pi_2 + k_2, & s \le x \le S. \end{cases} \tag{4}$$

Proof. We explain only the the integral equation for f_1 in (3) for the case $s \le Q$. The other equations are derived similarly. The steady state distribution of C_1 is absolutely continuous in $[0, S]$ with density $f_1(x)$, and $d_1(x)f_1(x)$ in the left-hand side is the long-run rate of downcrossings of level x. Thus, in steady state, the right-hand side of (3) represents the long-run rate of upcrossings of level x. To see this, assume that $C_1(0) \stackrel{\mathrm{d}}{=} C_1(\infty)$, namely, $C_1(0)$ has the steady-state distribution of the content level. That makes C_1 a stationary process, so that $C_1(t) \stackrel{\mathrm{d}}{=} C_1(\infty)$ 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:

(i) $0 \le C_1(\tau-) < x < s$. 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 PASTA 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 $\lambda \int_0^x f_1(w)dw + d(0)f_1(0)$.

(ii) $0 \le C_1(\tau-) \le s$ 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 3.2 below).

(iii) $0 \le C_1(\tau-) \le s$ 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 (3) and (4) in terms of the constants k_1, $\tilde{k}_1$ and k_2. These constants are computed in Lemma 3.2 below.

Solution of f_1: Let $F_1(x) := \int_0^x f_1(s)ds$ denote the cumulative distribution function (cdf), related to 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

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

$$F_1(x) = -\frac{c_0}{\lambda} + C_1 e^{\lambda D_1(x)}, \quad x \in [0, s),$$

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

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

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

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

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

Now, for $x \in [Q, S]$, $d_1(x)f_1(x)$ in this region is constant.

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 the unknown constant k_1.

Solution of f_2: Using simple arguments, as those for f_1, we get:

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

where $F_2(s) = \pi_2(e^{\lambda D_2(s)} - 1)$ and π_2 is obtained via the normalizing condition $\int_0^S f_2(w)\, dw = 1 - \pi_2$. ∎

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 OP1 and OP2, 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 OP1, given that the last jump prior to hitting s was to level Q (which under OP1 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 3.1 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 $\mathbf{1}\{s < Q\}$ be the indicator function which equals 1 if $s < Q$ and 0 otherwise.

Lemma 3.1.

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

Proof. For simplicity, we say that W is at state 0 if $W = w_\lambda$, and at state 1 if $W = w_\mu$, $t \geq 0$. Since the CTMC C has only two states, we can use the uniformization method; see, e.g., §II in [**?**], so that all transitions are generated by a single Poisson process. In particular, we consider a uniformized version of C, which spends an exponential amount of time with rate $\lambda + \mu$ in either state. Let $P_t(i, j)$ denote the transition operator of C, and $P(i, j)$ the transition probabilities of the discrete-time Markov chain (DTMC) associated with the uniformized version of C, $i, j = 0, 1$.

Note that $P(0, 0) = P(1, 0) = \lambda/(\lambda + \mu)$, so that the transition probabilities to 0 are the same for all n;

$P^n(i,0) = \lambda/(\lambda+\mu)$, $n \geq 1$, $i = 0,1$. Hence,

$$P_t(0,0) = \sum_{n=0}^{\infty} P^n(0,0) e^{-(\lambda+\mu)t} \frac{[(\lambda+\mu)t]^n}{n!} = \frac{\lambda}{\lambda+\mu} + \frac{\mu}{\lambda+\mu} e^{-(\lambda+\mu)t}.$$

The result for $\theta_1(s,S)$ follows by replacing t with $D_1(S) - D_1(s)$, namely with the time it takes the content level to reach s, starting in level S.

The proof for $\gamma_1(s,Q)$ is similar. However, level s can be reached, after starting at level Q, only if $s < Q$. Hence, the indicator function in the expression. ∎

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

Lemma 3.2. *Consider $x \in (s,S]$. Then the long-run rate of upcrossings of level x under OP1 is given by k_1 if $s \leq Q$ and $\tilde{k}_1$ if $s \geq Q$. It is given by k_2 under OP2, where*

$$\begin{aligned} &k_1 := \gamma_1(s,Q) d_1(0) f_1(0) + \theta_1(s,S) d_1(S) f_1(S) \quad \textit{and} \quad \tilde{k}_1 := \theta_1(s,S) d_1(S) f_1(S), \\ &k_2 := \theta_2(s,S) d_2(s) f_2(s). \end{aligned} \tag{5}$$

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 Discrete-Time Markov Chain (DTMC) with two states – S and Q. The transition matrix of that DTMC at jump epochs is

$$P := \begin{bmatrix} P_{S,S} & P_{S,Q} \\ \\ P_{Q,S} & P_{Q,Q} \end{bmatrix} = \begin{bmatrix} \theta_1 + (1-\theta_1)(1-e^{-\lambda D_1(s)}) & (1-\theta_1)e^{-\lambda D_1(s)} \\ \\ 1-(1-\gamma_1)e^{-\lambda D_1(s)} & (1-\gamma_1)e^{-\lambda D_1(s)} \end{bmatrix}. \tag{6}$$

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 (6).

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

$$\nu_S = \frac{1-(1-\gamma_1)e^{-\lambda D_1(s)}}{1-(\theta_1-\gamma_1)e^{-\lambda D_1(s)}} \quad \text{and} \quad \nu_Q = 1-\nu_S, \tag{7}$$

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

$$k_1 = (\nu_S\theta_1 + \nu_Q\gamma_1)d_1(s)f_1(s) \tag{8}$$

is the long run rate of jumps from level s.

We next show that the expression for k_1 in (5) gives the same expression as in (8): From (3) (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 (5), we rewrite k_1 to get

$$k_1 = \frac{\gamma_1c_0 + \theta_1c_0e^{\lambda D_1(s)} - \theta_1c_0}{1-\theta_1}. \tag{9}$$

It is then a matter of simple algebra to show that the expression for k_1 in (9) 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 (7). 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, i.e., k_1 in (9) is equal to the expression (8). This proves the claim. ∎

Remark 3.2. The terms for the constants in Lemma 3.2 can be guesses. To see that, consider k_1 and note that we can compute its value by conditioning on the last jump prior to hitting s (during a cheap period), 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 (5). Similar reasonings give us the expressions for $\tilde{k}_1$ and k_2. ∎

3.4 Profit Functions under OP1 and OP2

We can use the solutions for f_1 and f_2 and compute the long-run profit functions for both policies. We denote by $R_1 := R_1(s, S, Q, p_l, p_h, q)$ the long-run average profit function generated by OP1, and by $R_2 := R_2(s, S, p_l, p_h, q)$ the long-run profit function generated by OP2. The expressions for the steady-state profit functions R_1 and R_2 are as follows:

$$\begin{aligned} R_1 = & \int_0^S [p(w)d_1(w) - h(w)]f_1(w)dw - [K + w_\mu(S-s)]k_1 \\ & - \lambda \int_0^s [K + w_\mu(S-w)]f_1(w)dw - (K + w_\lambda Q)d_1(0)f_1(0) \end{aligned} \tag{10}$$

and

$$\begin{aligned} R_2 = & \int_0^S [p(w)d_2(w) - h(w)]f_2(w)dw - [K + w_\mu(S-s)]k_2 \\ & - \lambda \int_0^s [K + w_\mu(S-w)]f_2(w)dw - (K + w_\mu S)\lambda\pi_2 - a\frac{d(0)f_2(0)}{\lambda}. \end{aligned} \tag{11}$$

We now explain the expressions in (10) and (11):

- The first terms on the right hand sides, $\int_0^S [p(w)d_i(w) - h(w)]f_i(w)dw$, $i = 1, 2$, 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 $C(t) = w_\mu$, i.e., 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 Theorem 3.1.
- The average ordering costs $\lambda \int_0^s [K + w_\mu(S-w)]f_i(w)dw$, $i = 1, 2$, 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 OP2 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 OP1, 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 OP2, 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 conduct numerical experiments to assess the behavior of different order policies. We use a linear demand model $d(p) = 50 - p$, with $p_l = 0$ and $p_h = 50 - 10^{-3}$, and linear holding costs $h(x) = h \cdot x$, for $h > 0$.

In the following plots we visualize 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 OP0, OP1, and OP2.

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. OP2 performs slightly better than OP1, and both outperform OP0. Table **??** lists the optimal profit and decision variables for the order policies OP0, OP1, and OP2. Figure **??** 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 OP1 does not depend on a; for OP2, the optimal profit is decreasing in a.

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. OP1 performs slightly better than OP0, and both outperform OP2. Table **??** lists the optimal profit and decision variables for the order policies

Figure 1: Sensitivity analysis for scenario 1

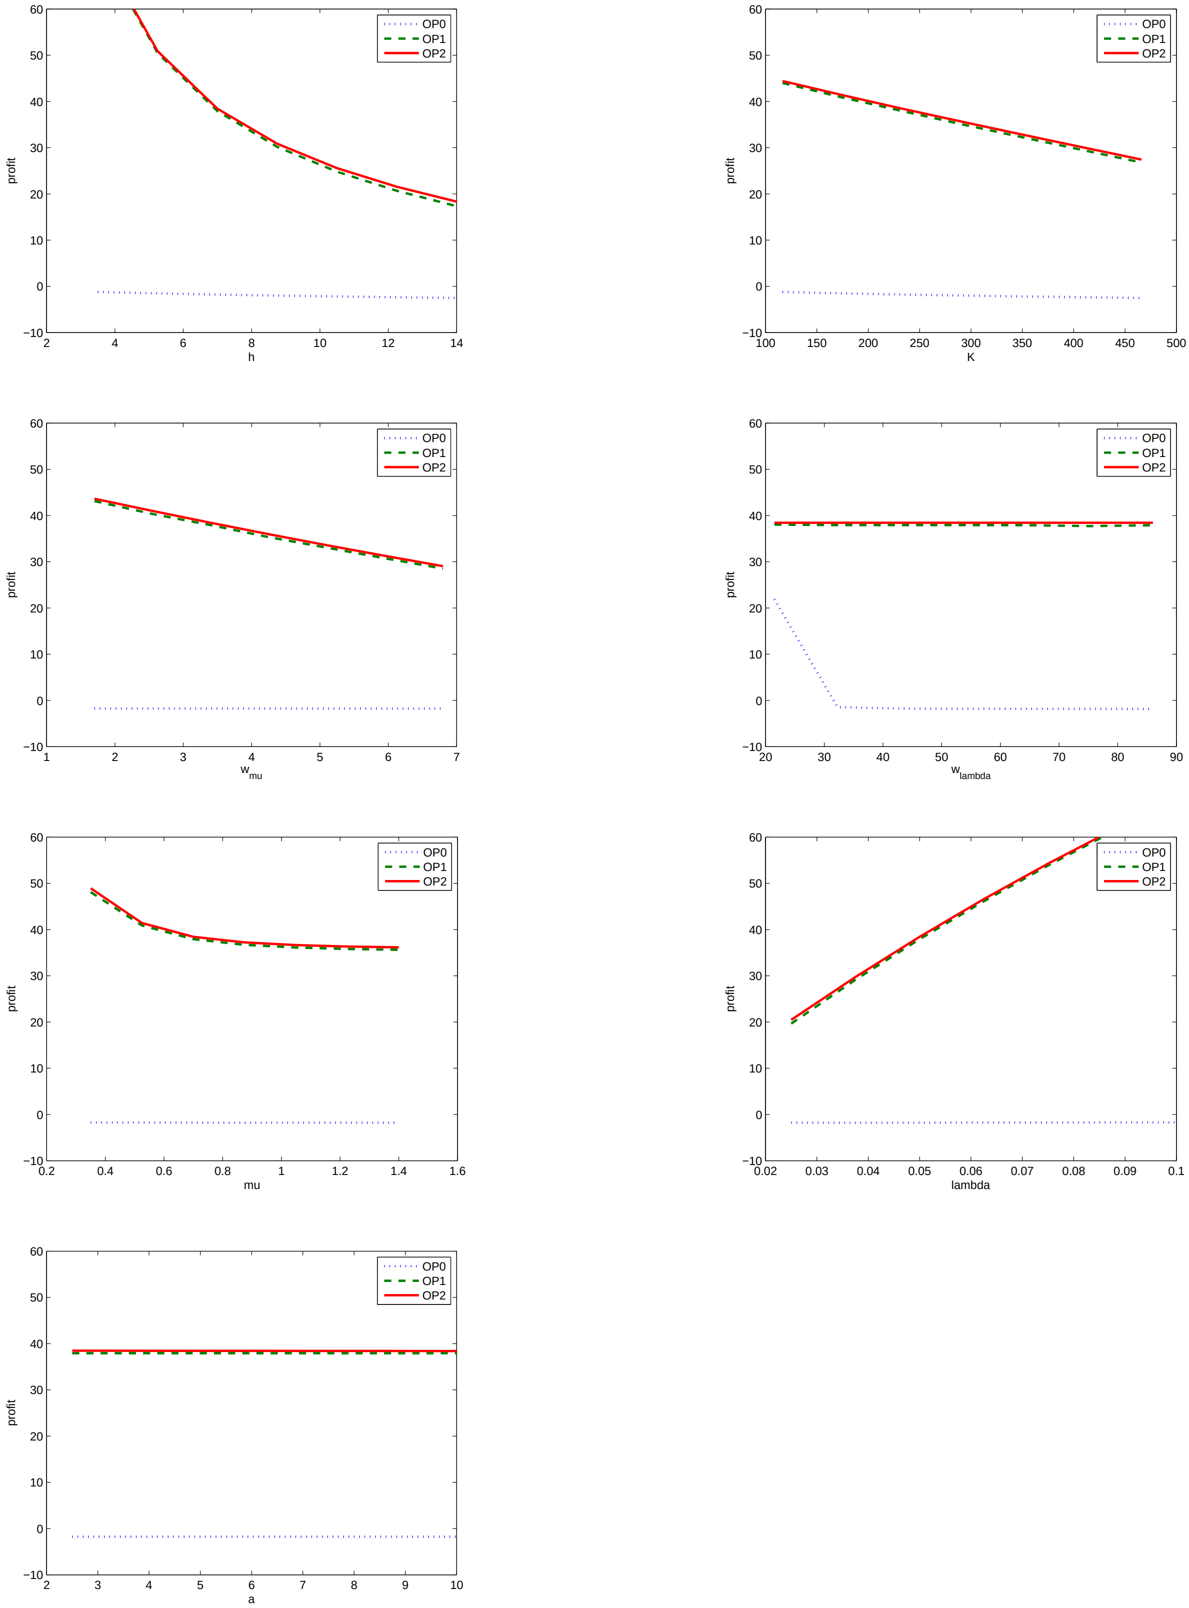

Table 1: Profit and optimal solution under different order policies, for scenario 1

Order Policy	Profit	p_l	p_h	q	s	Q	S
OP0	-1.75936	46.7862	49.999	0.251794	5.4192 E-10		3.20171
OP1	37.9172	33.0997	49.999	0.0865384	6.00827	0.0865491	61.0472
OP2	38.4475	33.0997	49.999	0.0104782	5.9419		60.9695

OP0, OP1, and OP2. Figure **??** 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 OP1 does not depend on a; for OP2, the optimal profit is decreasing in a.

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	68.9299	37.9017	40.3715	9.51125	0		21.4643
OP1	69.1156	37.7775	40.3724	9.99395	3.7814 E-9	20.5741	23.5341
OP2	38.8532	37.3205	49.999	0.00756244	0.00756244		25.0628

5 Generalizations

In this section we present two generalizations for the basic model analyzed above. 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, i.e., 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.

We now consider the case in which one of the periods, either the cheap or the expensive period, follows a phase-type distribution. For simplicity of exposition, we take the exact distribution to have two exponential phases, but our arguments extend directly to more general phase-type distributions. The model can be extended by considering expensive non-exponential periods, or cheap non-exponential periods. Here, we will consider the latter case. Specifically, assume that the cheap period is exponentially distributed with rate μ, but the law of the expensive period is $Erlang(2, \lambda)$. Our analysis for this case is different than before:

Figure 2: Sensitivity analysis for scenario 2

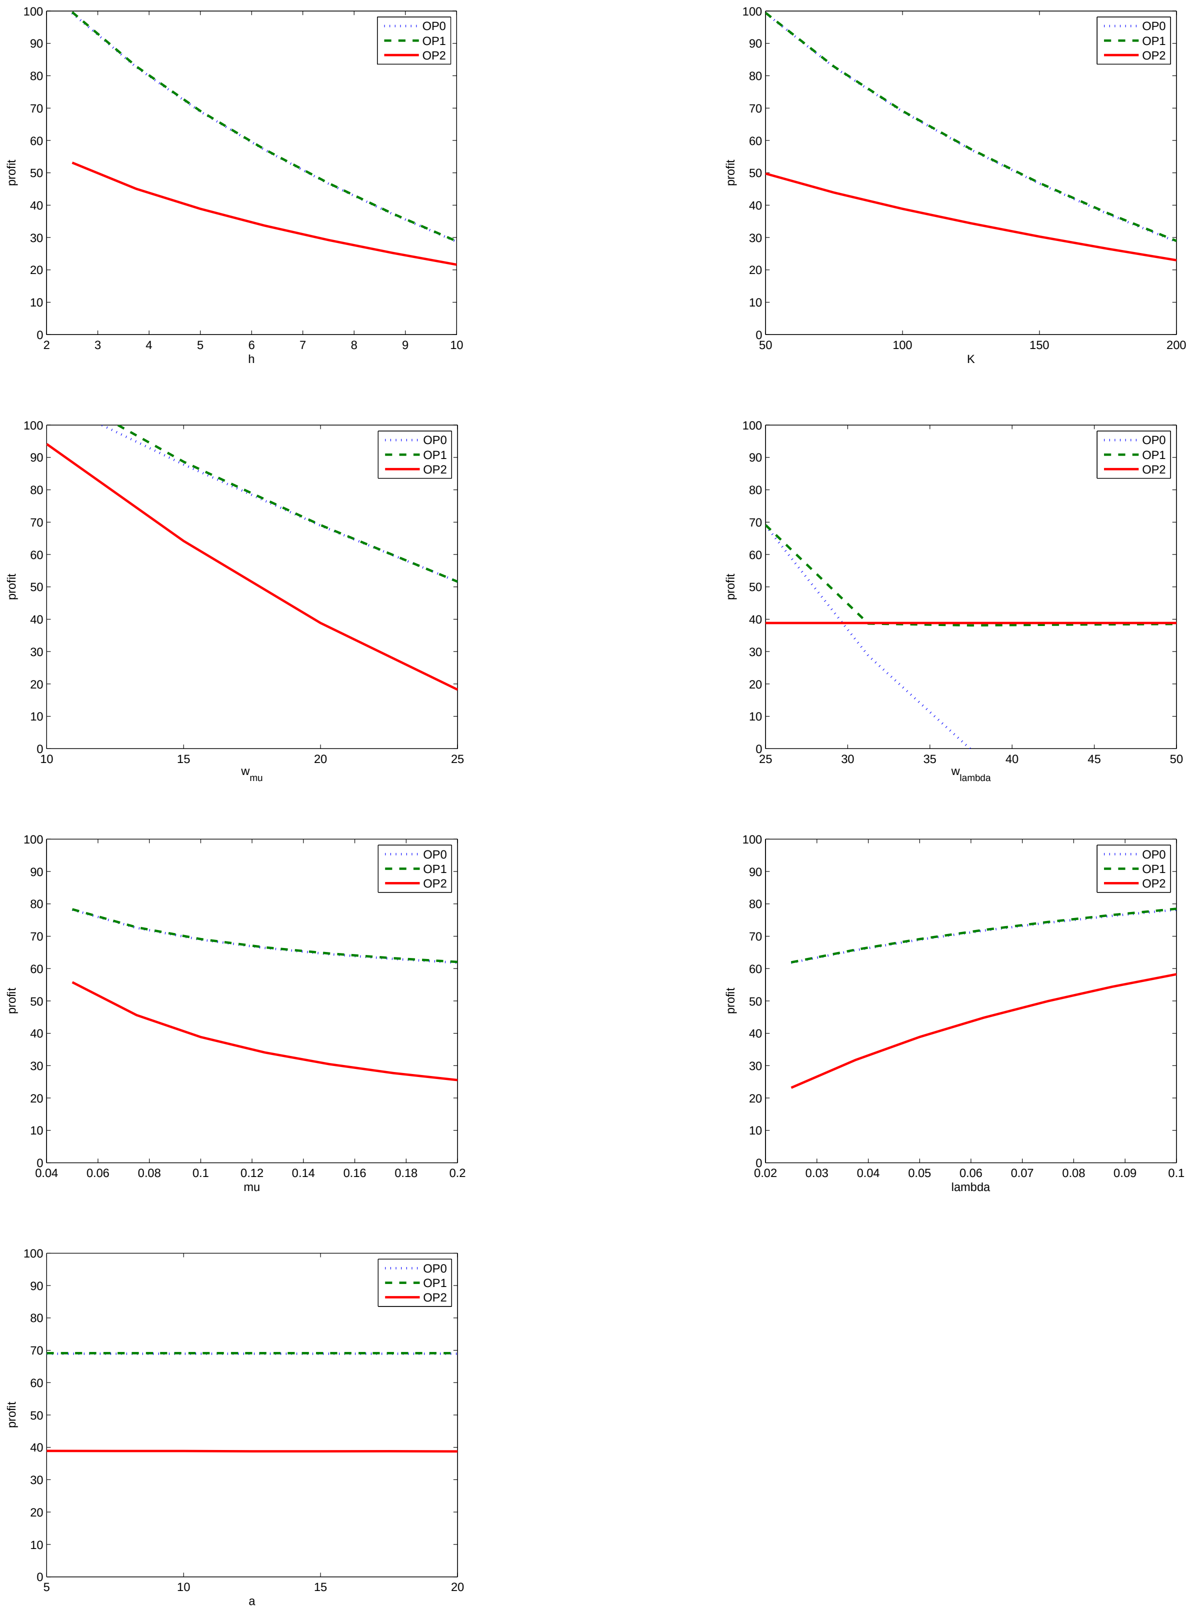

Instead of equalizing the number of up and down crossings for each level x, we compare the number of downcrossings of a certain level with the number of downcrossings of another level.

We designate the probabilities that level s is downcrossed by the first phase and the second phase of the expensive period, respectively, by p_1 and p_2. In the next theorem we introduce the balance equation of the content level where p_1 and p_2 will be computed in the sequel.

Theorem 5.1.

$$a(x)f(x) = \begin{cases} a(s)f(s)\left[(1+\lambda[D(s)-D(x)])\, e^{-\lambda[D(s)-D(x)]}p_1 + p_2 e^{-\lambda[D(s)-D(x)]}\right] & 0 < x < s, \\ a(S)f(S) & s \le x < S, \end{cases}$$

Proof. (i) $s < x < S$. 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 x, so that $a(x)f(x) = a(S)f(S)$.

(ii) $0 < x < s$. For every x we mark a downcrossing of level s as a downcrossing of type 1 if no jump occurs after the latter downcrossing, and a downcrossing of level x. Otherwise, the latter downcrossing is of type 2. It is clear that the long-run average number of downcrossings of level x is equal to the long-run average number of downcrossings of type 1. The probability of a type-1 downcrossing is

$$\left[p_1\,(1+\lambda[D(s)-D(x)])\, e^{-\lambda[D(s)-D(x)]} + p_2 e^{-\lambda[A(s)-A(x)]}\right],$$

since with probability p_1 level s is downcrossed during the first phase of the expensive period, and with probability

$$(1+\lambda[D(s)-D(x)])$$

no jump occurs between the latter two downcrossings. Multiplying together, we get that the latter probability is

$$p_1\,(1+\lambda[D(s)-D(x)])\, e^{-\lambda[D(s)-D(x)]}.$$

Similarly, with probability p_2 level s is downcrossed during the second phase of the expensive period, and with probability $e^{-\lambda[D(s)-D(x)]}$ no jump occurs between the latter two downcrossings. □

It remains to compute p_1 and p_2. To that end, we construct an auxiliary process $\chi := \{\chi(t) : t \ge 0\}$, where

$$\chi(t) := t + S_1 + \cdots + S_{N(t)}, \qquad \chi(0) = 0,$$

where the S_i's are iid random variables having Laplace transforms

$$\tilde{G}(\alpha) = \frac{\mu}{\mu + \alpha} \cdot \frac{\lambda}{\lambda + \alpha}$$

and $\{N(t) : t \geq 0\}$ is a Poisson process with rate μ. In particular, $\sum_{j=1}^{N(t)} S_j$ is a compound Poisson process and χ is a non-decreasing process that increases either linearly, at rate 1, between jumps, or by positive jumps of (random) size S, where S is distributed as a sum of two independent exponential random variables: one with rate μ and the other with rate λ.

We can think of each jump of χ as having two phases: The first phase is distributed exponentially with rate μ, and the second exponentially with rate λ. 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 the well-known Wald's martingale

$$M_\alpha(t) := \frac{e^{-\alpha\chi(t)}}{E[e^{-\alpha\chi(t)}]} = e^{-\alpha\chi(t) - \varphi(\alpha)t}, \quad \alpha > \max(-\lambda, -\mu) \tag{12}$$

where

$$\varphi(\alpha) := -\left[\alpha + \mu[1 - \frac{\mu}{\mu + \alpha} \cdot \frac{\lambda}{\lambda + \alpha}]\right]. \tag{13}$$

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

$$1 = E[e^{-\alpha\chi(\tau) - \varphi(\alpha)\tau}]. \tag{14}$$

It follows from the memoryless property of the exponential random variable that the stopping time τ and the martingale $M_\alpha(\tau)$ are conditionally independent given the phase in which level $D(S) - D(s)$ is upcrossed. Specifically, if level $D(S) - D(s)$ is upcrossed by the first phase of the jump, then $M_\alpha(\tau) = D(S) - D(s) + X_\mu + X_\lambda$, where X_λ and X_μ denote two independent exponential random variables having respective rate λ and μ. If level $D(S) - D(s)$ is upcrossed by the first phase of the jump, then $M_\alpha(\tau) = D(S) - D(s) + X_\lambda$. Finally, if level $D(S) - D(s)$ is upcrossed by the continuous drift of χ, then $M_\alpha(\tau) = D(S) - D(s)$.

Let B_0, B_1 and B_2 be the events that level $D(S) - D(s)$ is reached by the drift, upcrossed by the first phase of the jump and upcrossed by the second phase of the jump, respectively. Then by (**??**),

$$\begin{aligned} 1 &= E[e^{-\alpha Y(\tau)-\varphi(\alpha)\tau}\mathbf{1}_{B_0}] + E[e^{-\alpha Y(\tau)-\varphi(\alpha)\tau}\mathbf{1}_{B_1}] + E[e^{-\alpha Y(\tau)-\varphi(\alpha)\tau}\mathbf{1}_{B_2}] \\ &= e^{-\alpha(D(S)-D(s))}E[e^{-\varphi(\alpha)\tau}\mathbf{1}_{B_0}] \\ &\quad +\tfrac{\mu}{\mu+\alpha}\cdot\tfrac{\lambda}{\lambda+\alpha}\cdot e^{-\alpha(D(S)-D(s))}E[e^{-\varphi(\alpha)\tau}\mathbf{1}_{B_1}] \\ &\quad +\tfrac{\lambda}{\lambda+\alpha}\cdot e^{-\alpha(D(S)-D(s))}E[e^{-\varphi(\alpha)\tau}\mathbf{1}_{B_2}] \end{aligned} \tag{15}$$

where the second step is implied by the above conditional independence and the memoryless property.

Lemma 5.1. *We have*

$$E[\mathbf{1}_{B_0}] = 1 - p_1 - p_2, \quad E[\mathbf{1}_{B_1}] = p_2 \quad \textit{and} \quad E[\mathbf{1}_{B_2}] = p_1.$$

Proof. Take the projection of χ on the process axis. Then, $1-p_1-p_2$ is the conditional probability that level s will be reached during a cheap period, given the period is cheap at level S; p_1 is the conditional probability that level s will be reached during the second phase of the expensive period given the same event; and if at level S the period is cheap, level s will be reached during the first phase of the expensive period. p_2 is the conditional probability that level s will be reached during the first phase of the expensive period given the same event. ☐

As τ is bounded ($0 < \tau < D(S) - D(s)$) the restricted transforms $E[e^{-\varphi(\alpha)\tau}\mathbf{1}_{B_0}]$, $E[e^{-\varphi(\alpha)\tau}\mathbf{1}_{B_1}]$ and $E[e^{-\varphi(\alpha)\tau}\mathbf{1}_{B_2}]$ are analytic functions on the entire complex plane. Obviously, we want to pick those value of α in (**??**) for which $\varphi(\alpha) = 0$. By (**??**), $\varphi(\alpha) = 0$ holds for $\alpha = 0$ and the roots of the quadratic equation

$$\alpha^2 + \alpha(2\lambda + \mu) + \lambda\mu + \lambda + \mu = 0. \tag{16}$$

Inserting the roots of (**??**) into (**??**) yields the two equations for $i = 1, 2$.

$$\begin{aligned} 1 &= e^{-\alpha_i(D(S)-D(s))}(1 - p_1 - p_2) \\ &\quad +\tfrac{\mu}{\mu+\alpha_i}\cdot\tfrac{\lambda}{\lambda+\alpha_i}e^{-\alpha_i(D(S)-D(s))}p_1 \\ &\quad +\tfrac{\lambda}{\lambda+\alpha_i}e^{-\alpha_i(D(S)-D(s))}p_2. \end{aligned}$$

To solve for f according to the balance equation in Theorem 1, we use the normalizing condition and the fact that $d(x)f(x)$ is a continuous function at $x = s$.

Remark 5.1. Clearly, computing the probability that, at the downcrossing of level s at a specific future time the state of the world W is at a particular state, becomes hard as the number of states of the process W increases. However, even if explicit computations are impossible, one can solve the Kolmogorov backward or forward equations for the generator matrix of W numerically to compute the desired probabilities.

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 $\exp(\eta)$ period 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 $\exp(\eta)$ period.

Theorem 5.2. *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*

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

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

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

(ii) $s_0 \leq x < s_1$. 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 $\exp(\eta)$ 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 cost price is changed from expensive to cheap an order will be placed and it will take an $\exp(\eta)$ 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) $s_1 \leq x < S$. 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 $\eta[\gamma + (1-\gamma)(1 - e^{-\lambda[D(s_1)-D(s_0)]})][F(s_1) - F(s_0)]$. $\square$

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 cost 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 cost price is expensive and the state is S. Then the solution (α_1, α_2) to the equations

$$(\alpha_1, \alpha_2) \begin{pmatrix} p_{cc} & 1 - p_{cc} \\ 1 - p_{ee} & p_{ee} \end{pmatrix} = (\alpha_1, \alpha_2) \quad \text{and} \quad \alpha_1 + \alpha_2 = 1,$$

is the solution of the conditional steady state probability - α_1 (α_2) that level s_1 is downcrossed during a cheap period (expensive period), given that at the starting point, i.e., at level S, the cost price is cheap (expensive). Finally

$$\gamma = \alpha_1 p_{cc} + \alpha_2 (1 - p_{ee}).$$

Computing p_{cc} and p_{ee} is similar to the computations in Lemma 3.1.

Acknowledgments

This research is conducted while the first and second author were affiliated at CWI, and is made possible by an NWO VIDI grant. The third author is also affiliated with VU University Amsterdam, Eurandom and Georgia Tech.

References

[1] H.S. Ahn, M. Gumus and P. Kaminsky (2007). Pricing and manufacturing decisions when demand is a function of prices in multiple periods, *Oper. Res.* 55(6), 1039-1057.

[2] S. Asmussen (2003). *Applied Probability and Queues*, 2nd ed., Springer, New York.

[3] P.H. Brill (2008). *Level crossing methods in stochastic models* (2008). Springer Verlag.

[4] S. Browne and P. Zipkin (1991). Inventory models with continuous, stochastic demands. *Ann. Appl. Prob.* 1(3), 419–435.

[5] B. Chaouch (2007). Inventory Control and Periodic Price Discounting Campaigns. *Nav. Res. Logist.* 54(1), 94–108.

[6] J.W. Cohen (1977). On up- and downcrossings. *Journal of Appl. Prob.* 14(2), 405–410.

[7] M. Goh and M. Sharafali (2002). Price-dependent inventory model with discount offers at random times. *Prod. Oper. Manage.* 11(2), 139–156.

[8] K. Moinzadeh (1997). Replenishment and stock policies for inventory systems with random deal offerings. *Man. Sci.* 43(3). 334–342.

[9] J. S. Song and P. Zipkin (1993). Inventory Control in a Fluctuating Demand Environment. *Oper. Res.* 4(2), 351–370.