COLEMAN FUNG

RISK MANAGEMENT

RESEARCH CENTER

Working Paper # 2007 -04

Optimal Spot Market Inventory Strategies in the Presence of Cost and Price Risk

Xin Guo, University of California Berkeley
Philip Kaminsky, University of California Berkeley
Pascal Tomecek, J.P. Morgan
M. Yuen, University of California Berkeley

October 7, 2007

University of California Berkeley

Optimal Spot Market Inventory Strategies in the Presence of Cost and Price Risk

Xin Guo, Philip Kaminsky* Pascal Tomecek† M. Yuen‡

UC Berkeley *J.P.Morgan* *UC Berkeley*

October 7, 2007

Abstract

We consider a firm facing random demand at the end of a single period of random length. At any time during the period, the firm can either increase or decrease inventory by buying or selling on a spot market where price fluctuates randomly over time, and the revenue the firm gets by meeting demand at the end of the period is a function of the spot market price at that time. We first demonstrate that this control problem is equivalent to a singular control problem of higher dimensions. We then use this insight combined with a novel control-theoretic approach to show that the optimal policy is completely characterized by a simple price-dependent two threshold policy. In a series of computational experiments, we explore the value of actively managing inventory during the period rather than making a purchase decision at the start of the period, and then waiting for demand.

*Department of Industrial Engineering and Operations Research, UC Berkeley, CA 94720-1777, USA. email: {xinguo, kaminsky}@ieor.berkeley.edu. Xin Guo's research is generously supported by OpenLink Fund at the Coleman Fung Risk Management Research Center at UC Berkeley.

†J.P. Morgan, New York, email: pascal.i.tomecek@jpmorgan.com.

‡Department of Industrial Engineering and Operations Research, UC Berkeley, CA 94720-1777, USA. email: ming @ieor.berkeley.edu.

1 Introduction

Spot market supply purchases are increasingly considered an important operational tool for firms facing the risk of higher than anticipated demand for goods (see, e.g., [29] and the references therein). For example, Hewlett-Packard manages the risks associated with electronic component procurement by utilizing a portfolio of long term and option contracts to cover likely demand, and procurement on the spot market if demand is higher than expected [5]. Indeed, there has been a recent stream of research focusing on determining an optimal mix of long term fixed commitment and options procurement contracts, given that after demand is realized, firms can if necessary procure on the spot market to meet demand. In these models, the spot market is typically employed if supply requirements exceed the contracted amount of the fixed commitment contract (or if the spot price happens to be lower than the exercise price of the procurement options).

In this paper, we argue that if effectively utilized, the spot market can be used to hedge against much more than just excess demand. In many cases, the spot market can be a powerful tool for hedging against both supply cost uncertainty and demand price uncertainty in the supply chain, even without an accompanying portfolio of supply contracts. To explore this concept, we develop a stylized model of a firm that has a random period of time to increase or decrease inventory by purchasing or selling on the spot market before facing a single demand of random magnitude, the revenue of which is a function of the spot market price at the time that the demand is realized. We demonstrate that in many cases, the firm can use purchases and sales on the supply spot market both to guard against low prices for its products, and high prices for the components that go into its products.

1.1 Literature Review and Our Work

There is a long history of research focusing on inventory strategies when the cost of the inventory is random, typically with the objective of minimizing inventory cost. Various researchers ([20],[27],[8],[7]), for example, considered versions of periodic review models where component costs (and sometimes other problem parameters) are Markov-modulated, usually demonstrating the optimality of state dependent basestock or (s,S) policies. Another stream of literature modeled deterministic demand and characterized optimal policies, either in a periodic model with random raw material prices([15],[4]), increasing ([12]), or decreasing ([33]) prices, or with one or two different levels of constant continuous time demand and occasional supply price discounts ([25],[13],[6]).

Finally, a growing stream of research considers the impact of a spot market on supply chain operations. [17] divided this work into two sets, one having to do with contract valuation (see references in [17]), and another with optimal procurement from the spot market. This category of work can be further subdivided into single and multi-period models. [28] considered a single period model where a supply contract is signed at the start of the period, demand is realized, and then the buyer can either make purchases on the spot market to meet demand, or salvage excess inventory. Optimal purchasing quantities were determined for this setting. [1] considered a similar setting, and focused on reserving an appropriate capacity level to meet demand. [34] studied a single period optimal procurement contract model in the presence of a spot market, and [14] considered a single period model designed to determine how the spot market impacts supply chain coordination. [11] investigated a single period model in which the firm must first select from a set of supply contracts, and then random demand is realized and the firm must meet that demand by utilizing either supply contracts or the spot market. Optimal contract selection and utilization policies were characterized. In [21], deterministic demand must be met after a deterministic time period, but the firm has a contract to procure supply on the spot market at some point before demand is realized, where spot market price is a continuous random process. The optimal purchase time was derived numerically. [35] considered a discrete time multi-period model in which the firm can either buy at a fixed price from a long-term supplier or buy from a spot market incurring variable purchasing cost and a fixed cost of using the spot market, and characterized the optimal purchasing policy. In [9], an optimal long term contract was compared to utilizing the spot market for a series of periods, where each period was essentially an independent news-vendor problem. [23] proposed a multi-period model in which a portfolio of supply contracts must be selected at the start of the horizon, and then in each period after demand is realized and the spot market price is observed, the decision to utilized contracts or buy on the spot market to meet demand must be made.

Other than the few exceptions noted above, these papers modeled the spot market with a single spot market price or a discretely realized series of spot market prices, and typically allowed one opportunity to buy or sell on the spot market following each demand realization. In contrast, we consider a continuous time model of spot market price evolution, and determine how the firm can buy and sell in the spot market repeatedly in order to guard against both supply cost uncertainty and demand price uncertainty.

In our work, we model the inventory level of the firm, Y_t, at time t with a pair of controls (ξ_t^+, ξ_t^-) so that $Y_t = Y_0 + \xi_t^+ - \xi_t^-$. Here ξ_t^+ and ξ_t^- are non-decreasing processes and represent respectively the total accumulated inventory ordered and sold by time t starting from time 0. We assume that the price of each unit of inventory is stochastic and is a Brownian motion process as in [17] and [21]. We also assume that the time until the (single) demand arrives, as well as the amount of that demand, is random. The revenue associated with the demand is assumed to be a function of the amount of that demand and the spot market price at the time when the demand arrives. In addition to the running holding cost, there are costs whenever an inventory level is increased or decreased by selling or buying at the spot market: adjustment cost is a function of the spot price and the amount of the adjustment, plus possibly a transaction cost. The cost could be negative when selling the inventory representing a savage value for the inventory. Subject to the cost structure, the goal is to maximize the expected discounted profit. To facilitate the analysis, we assume no fixed cost and focus on explicitly characterizing the optimal policy.

In particular, we show that the optimal inventory policy depends both on the spot price and inventory level, and that it is in principle a simple and continuous (F, G) policy. Given a spot price p and inventory level z, if (p, z) falls between $(F(z), G(z))$, no action is taken; if (p, z) falls above $F(z)$ (below $G(z)$), the inventory level is reduced to $F(z)$ (raised to $G(z)$).

Our technique is closely related to a stream of research ([3, 2, 10, 18, 19, 32, 31, 26]) focusing on continuous time inventory models via impulse controls (i.e. with a fixed cost) or singular controls (i.e. without a fixed cost) formulation. Most of these papers (with the exception of [2] where the demand process is Poisson) considered a one product inventory model where the inventory level is a controlled Brownian motion. That is, the inventory level without intervention is a Brownian motion, and the continuous adjustment of the inventory level is additive to the Brownian motion and incurs a linear cost (plus a possible fixed cost). Subject to an additional holding cost and shortage penalty, the objective in these papers was either to minimize the expected discounted cost or the average cost ([3, 26]) over an infinite time horizon. Except for [32, 26], most of the models assumed little constraints on the inventory level besides restricting it to the positive real line. Assuming a fixed cost, [10] proved the existence of an optimal (d, D, U, u) policy for this system: do nothing when inventory is in the region of (d, u), and adjust the inventory level to D (or U) whenever the inventory level falls to d (or rises to u). This optimal policy and the solution structure were more explicitly characterized under various scenarios in [18, 19, 32, 31, 26].

The main contribution of our paper is best discussed in light of several crucial elements underlying all previous control-theoretic inventory analysis. Firstly, the price of the inventory was assumed to be constant so that the cost of the inventory control would be linear. Secondly, the inventory control was additive to a Brownian motion, and as a result the inventory level was either unconstrained on the positive real line or an infinite penalty cost was needed to ensure an upper bound on the inventory level ([32, 26]). These two characteristics ensured that the control problem was one-dimensional and facilitated the analysis of the value function. The solution approach was to apply the Dynamic Programming principle and solve some form of Hamilton-Jacobi-Bellman equations or Quasi-Variational-Inequalities, with a priori assumptions on the regularity conditions. In contrast, in our model the adjustment cost is no longer linear and depends on the spot price, the transaction cost, and the amount of adjustment; the inventory control variable is modeled directly, and is no longer necessarily additive to the underlying Brownian motion process. Thus, constraints on the inventory level are modeled directly by the minimum and maximum capacity of the inventory, and can be easily extended to more detailed constrained inventory levels without further technical difficulty. In essence, the introduction of price dynamics leads to a higher dimensional singular control problem for which previous analysis cannot be directly generalized. The derivation in this paper is thus based on a new solution approach, which allows us to bypass the (possible) non-regularity of the value functions. The key idea is to break down the two-dimensional control problem by "slicing" it into pieces of one-dimensional problems, each of which is an explicitly solvable two-state switching problem, and to show that this re-parametrization is valid by the notion of "consistency" established in [16].

In the next section, we formally introduce our model. In Section 3, we develop explicit analytical expressions for the optimal policy for this model. In Section 4, we computationally explore the implications of our results.

2 Model and Preliminary Analysis

2.1 The Model

We consider a firm that purchases supply from a spot market in which the price of the supply component fluctuates over time. At a random time τ, the firm faces a random customer demand D. The firm meets demand if possible and charges with an exogenous price which is a function of the spot market price, and then salvages any excess inventory. At any time $t \in [0, \tau)$, the firm can instantaneously increase inventory (up to some upper bound on capacity $b < \infty$) or instantaneously decrease inventory down to some lower bound on inventory $a \geq 0$. However, the firm cannot buy inventory to satisfy demand at time τ, and the firm can only buy inventory a finite number of times in a finite interval. Net gain at time τ is from selling inventory to arriving customers and liquidating excess inventory, if any, as well as any additional penalty associated with not meeting demand, and thus is a function of the selling price and the inventory level at time τ, and the demand distribution. Moreover, at any time $t \in [0, \tau)$, inventory increase is associated with the purchase price of per unit at the supply spot market price (P_t), plus possibly additional proportional transaction cost (K^+). Similarly, inventory reduction is associated with the spot market price P_t, minus possibly additional proportional transaction cost (K^-). Finally, there is a running holding cost for each unit of inventory (C_h).

To capture these scenarios in mathematical terms, we start with probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and assume that the arrival time of the request, τ, is exponentially distribution with rate λ (so that the average arrival time is $1/\lambda$). D, the amount of commodity demanded at time is described by a distribution function F_D. Meanwhile, the spot market price $(P_t)_{t\geq 0}$ is stochastic and its dynamics are governed by a geometric Brownian motion such that [1]

$$dP_t = P_t(\mu dt + \sqrt{2}\sigma dW_t), \quad P_0 = p. \tag{1}$$

Here W_t is the standard Brownian motion on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and μ and σ represent respectively the expected spot market price appreciation and the potential price risk. We express the net gain at request time τ by $H(Y_\tau, D)P_\tau$, where $H(Y_\tau, D)$ represents the revenue multiplier associated with selling each unit of the inventory, as well as a possible penalty associated with each unit of unmet demand. Specifically,

$$H(y, D) = \alpha \min(D, y) + \alpha_o(y - D)^+ - \alpha_u(D - y)^+, \tag{2}$$

where $\alpha \geq 1$ is the earning price multiplier for each unit of met demand, $\alpha_u \geq 0$ is the penalty price multiplier for each unit the firm is short, and $0 \leq \alpha_o \leq 1$ is the fraction of price the firm is able to get by salvaging excess inventory.

To define admissible inventory policies, we specify the filtration $\mathbb{F}$ representing the information on which inventory decisions are based. Given λ and the distribution of D, it is clear that $\mathbb{F} = (\mathcal{F}_t)_{t\geq 0}$ is the filtration generated by P_t. Given $\mathbb{F}$, we define a pair of left-continuous with right limit, adapted, and non-decreasing processes ξ_t^+ and ξ_t^- to be the cumulative increases and decreases in supply inventory (purchases and sales, respectively) up to time t. Therefore, Y_t, the inventory level at time $t \in [0, \tau)$, is given by

$$Y_t = y + \xi_t^+ - \xi_t^-, \tag{3}$$

where y is the initial inventory amount.

To be consistent with the restriction that the firm can only purchase supply inventory on the spot market a finite number of times in a finite interval, Y_t is finite variation process. Meanwhile, for uniqueness of expression (3), (ξ^+, ξ^-) are supported on disjoint sets. Furthermore, ξ^+ and ξ^- are adapted to $\mathbb{F}$ implying that the firm is not clairvoyant. Y is left-continuous, capturing the restriction that the commodity cannot be purchased at time τ to satisfy demand. Also, note that given the upper and lower bounds on capacity discussed above, there exists $0 \leq a < b < \infty$ such that an admissible control policy must satisfy $Y_t \in [a, b]$ for all $t \leq \tau$. Finally, we assume $\mathbb{E}\left[\int_0^\infty e^{-\rho t} d\xi_t^+ + \int_0^\infty e^{-\rho t} d\xi_t^-\right] < \infty$.

To account for the time value between $[0, \tau]$, we define $r \geq 0$ to be a discount rate. Thus, at time $t \in [0, \tau)$, increases in the inventory incur a cost $-e^{-rt}(P_t + K^+)d\xi_t^+$, and decreases in the inventory generate revenue $e^{-rt}(P_t - K^-)d\xi_t^-$. In addition, assuming a running holding cost C_h for each unit of inventory, the holding cost between $(t, t + dt) \subset [0, \tau)$ is $e^{-rt}C_h Y_t dt$.

[1]The extra term $\sqrt{2}$ is for notational convenience in the main text.

Given the model outlined above and any admissible control policy (ξ^+, ξ^-), the expected return to the firm is:

$$\begin{aligned} J(p, y; \xi^+, \xi^-) = & \text{ payoff at transaction time } \tau - \text{running holding cost between } [0, \tau] \\ & - \text{cost of inventory control (via buying and selling) between } [0, \tau] \\ = & \ \mathbb{E}\bigg[e^{-r\tau} H(Y_\tau, D) P_\tau - \int_0^\tau e^{-rt} C_h Y_t dt \\ & - \int_0^\tau e^{-rt}(P_t + K^+) d\xi_t^+ + \int_0^\tau e^{-rt}(P_t - K^-) d\xi_t^-\bigg]. \end{aligned}$$

The firm's goal is to manage inventory in order to maximize the expected discounted value of over all possible admissible control policy (ξ^+, ξ^-). Therefore, the optimization problem for the firm is

$$W(p, y) = \sup_{(\xi^+, \xi^-) \in \mathcal{A}_y} J(p, y; \xi^+, \xi^-). \tag{4}$$

subject to

$$\begin{aligned} & Y_t := y + \xi_t^+ - \xi_t^- \quad \in [a, b], \quad y \in [a, b], \\ & dP_t := \mu P_t dt + \sqrt{2}\sigma P_t dW_t, \quad P_0 := p > 0, \\ & C_h \in \mathbb{R}, \quad K^+ + K^- > 0; \end{aligned} \tag{5}$$

and the supremum is over the set of admissible strategies

$$\begin{aligned} \mathcal{A}_y := \big\{ & (\xi^+, \xi^-) : \xi^\pm \text{ are left continuous, non-decreasing processes,} \\ & y + \xi_t^+ - \xi_t^- \in [a, b], \quad \xi_0^\pm = 0; \\ & \mathbb{E}\left[\int_0^\infty e^{-\rho t} d\xi_t^+ + \int_0^\infty e^{-\rho t} d\xi_t^-\right] < \infty. \} \end{aligned} \tag{6}$$

2.2 Preliminary Analysis

Assuming that τ is independent of $\mathbb{F}$ and D is independent of both τ and $\mathbb{F}$, this one period optimization problem is in fact equivalent to the following singular control problem over an infinite time horizon. This equivalence is based on a simple conditioning argument.

That is,

Proposition 2.1.

$$W(p, y) = -(C_h + p)y + V(p, y).$$

Here,

$$V(p, y) = \sup_{(\xi^+, \xi^-) \in \mathcal{A}_y} \tilde{J}(p, y; \xi^+, \xi^-). \tag{7}$$

with

$$\tilde{J}(p,y;\xi^+,\xi^-) = \quad \mathbb{E}\left[\int_0^\infty e^{-(r+\lambda)t}\tilde{H}(Y_t)P_t dt - (K^+ + C_h)\int_0^\infty e^{-(r+\lambda)t} d\xi_t^+ \right.$$
$$\left. -(K^- - C_h)\int_0^\infty e^{-(r+\lambda)t} d\xi_t^-, \right], \tag{8}$$

subject to Eqn. (5) with

$$\begin{aligned}\tilde{H}(y) = \quad & \lambda\mathbb{E}[H(y,D)] - (r+\lambda-\mu)y \\ = \quad & (\alpha+\alpha_u-\alpha_o)\left[y(1-F_D(y)) - \int_y^\infty z f_D(z)dz\right] \\ & +(\alpha-\alpha_o)\mathbb{E}[D] + (\alpha+\mu-r-\lambda)y. \end{aligned}\tag{9}$$

In order to keep the discounted value of p_τ finite, we assume throughout the paper that $r+\lambda<\mu$, which ensures the finiteness of the value function.

Proposition 2.2. *(Finiteness of Value Function) If $r+\lambda>\mu$, $V(p,y)\le \eta Mx+b-a$, where $M=\sup_{y\in[a,b]}|H(y)|<\infty$.*

Proof. Let $x>0$ and $y\in[a,b]$ be given. Since $\lambda+r>\mu$ we have

$$\mathbb{E}\left[\int_0^\infty e^{-\rho t}[H(Y_t)P_t]dt\right] \le \mathbb{E}\left[\int_0^\infty e^{-\rho t}[MP_t]dt\right] \le \eta Mx.$$

Note that for any given $(\xi^+,\xi^-)\in\mathcal{A}_y$, $a-y\le \xi_t^+-\xi_t^-\le b-y$. From integration by parts, for any $t>0$,

$$-\int_{[0,T)} e^{-\rho t} d\xi_t^+ \le -\int_{[0,T)} e^{-\rho t} d\xi_t^- + (y-a).$$

Which, together with $K^+ + K^- > 0$ and $K^+>0$, implies

$$\begin{aligned}&\mathbb{E}\left[-(K^+ + C_h)\int_0^\infty e^{-\rho t}d\xi_t^+ - (K^- - C_h)\int_0^\infty e^{-\rho t} d\xi_t^-\right] \\ &\le (y-a) - (K^+ + K^-)\mathbb{E}\left[\int_0^\infty e^{-\rho t} d\xi_t^-\right] \le b-a.\end{aligned}$$

Since these bounds are independent of the control, we have

$$V(p,y)\le \eta Mp + b - a < \infty.$$

□

Note that this optimization problem is a generalization of singular control problems studied extensively in [18]. Here, we incorporate the dynamics of the price process, and thus the singular control is two-dimensional in that the state space is (p, y). Intuitively, the function $\tilde{H}(\cdot)$ captures ultimate potential benefit of carrying inventory over time.

Moreover, to avoid an arbitrage opportunity in the market, we assume that $K^+ + K^- > 0$, and in addition, we can assume without loss of generality that $K^+ > 0$, and consider only bounded inventory level. Thus, throughout the paper, we make the following assumptions:

Standing Assumption

- $\rho := r + \lambda > \mu$.
- $K^+ + K^- > 0$, $K^+ > 0$.
- $[a, b]$ is bounded.

To simplify subsequent notation, we define the following:

- $\eta = \frac{1}{\rho - \mu}$
- $m < 0 < n$, and $n, m = \frac{-(b-\sigma^2) \pm \sqrt{(b-\sigma^2)^2 + 4\sigma^2 r}}{2\sigma^2}$.

3 Derivation and Main Results

Next, we explicitly derive $V(p, y)$ and the corresponding optimal control policy.

Let us start with some intuition, and follow the traditional dynamic programming approach (see, e.g., [10] or [18]). The optimization problem has the state space $\{p, z\}$. Given any price p and the inventory level z at time 0, there are three options: do nothing; increase the inventory by purchasing on the spot market; or reduce the inventory by selling on the spot market.

If a quantity is purchased on the spot market, the inventory level jumps from z to $z + \eta$, thus the value function is at least as good as choosing over all possible jump size η. That is,

$$V(p, z) \geq \sup_{\eta}(-(K^+ + C_h)\eta + V(p, z + \eta)),$$

leading to $V_y(p, y) \leq K^+ + C_h$ (where $V_y(\cdot, \cdot)$ indicates the derivative of the value function V with respect to y). Similarly we see $V_y(p, y) \geq -K^+ + C_h$. Meanwhile, if no action is taken between time 0 and an infinitesimal amount of time dt, then expressing the value function at time 0 in terms of the value function at time dt though dynamic programming and Ito's lemma (as in [10]) yields $\sigma^2 p^2 V_{pp}(p,y) + \mu p V_p(p,y) - rV(p,y) + \tilde{H}(p,y) \leq 0$. Combining these observations, we get the following (quasi)-Variational Inequalities

$$\max\{\sigma^2 p^2 V_{pp}(p,y) + \mu p V_p(p,y) - rV(p,y) + \tilde{H}(p,y),$$
$$V_y(p,y) - K^+ - C_h, -V_y(p,y) - K^- + C_h\} = 0. \quad (10)$$

That is, the optimal policy (if it exists) can be characterized by explicitly finding the action and continuation regions where

$$\begin{cases} \mathcal{S}_0 \text{ (Inventory increase)} = & \{(p,y) : V_y(p,y) = -K^- + C_h\}, \\ \mathcal{S}_1 \text{ (Inventory decresse)} = & \{(p,y) : V_y(p,y) = K^+ + C_h\}, \\ \mathcal{C} \text{ (No action)} = & \{(p,y) : V_{y^-}(p,y) > -K^- + C_h, V_{y^+}(p,y) < K^+ + C_h, \\ & \quad \sigma^2 p^2 V_{pp}(p,y) + \mu p V_p(p,y) - rV(p,y) + \tilde{H}(p,y) = 0\}. \end{cases}$$

Taking this intuition one step further, one would expect a state-dependent two-threshold policy, where inventory is lowered if it is above the upper threshold, and increased if it is below the lower threshold (see, e.g, [10]). The thresholds are state dependent due to the dynamic model for P_t. That is, we would expect a downsizing region for inventory: $\{(p,z) : p \geq G(z)\}$, an ordering region: $\{(p,z) : p \leq F(z)\}$, and a (continuation) no-action region: $\{(p,z) : F(z) < p < G(z)\}$.

However, in order to formalize this intuition to a complete characterization of the optimal policy and the value function, one in general would assume *a priori* smoothness for the value function and the boundary to solve the QVI. Unfortunately, the regularity conditions for this two-dimensional control problem do not hold in general. (See counter-examples in [16]). Indeed, as we shall see in our analysis, the value function may not be C^1 in p (although it is C^1 in y) and F, G may not be continuous.

Thus, instead of solving the QVI directly, we adopt a different approach by translating the singular control problem into a switching control problem, following [16]. The key idea is that by fixing each level z_0, we effectively will be solving for the $F(z_0), G(z_0)$ policy by solving a one-dimensional switching control problem. The switching control problem is a two-state switching problem, where switching from 0 to 1 corresponds to inventory increases and switching from 1 to 0 represents inventory reduction. In order for this approach to work for all z, meaning we can break down the two-dimensional control problem by slicing it into pieces of one-dimensional problem, we need to make sure the resulting control policies at different levels of z are "consistent". Intuitively, this consistency requires that for a given price p at level z_0, if it is optimal to reduce the inventory level, then it is also optimal to reduce the inventory level given the same p and a higher level $z(> z_0)$.

3.1 Derivation of $V(p, y)$

First, we establish the following results by direct computation.

Lemma 3.1. *$\tilde{H}(y)$ is concave y for ANY distribution of F_D with finite expectation. In particular,*

$$\tilde{H}(t, y_2) - \tilde{H}(t, y_1) = \int_{y_1}^{y_2} \tilde{h}(z)dz$$

with $\tilde{h}(y)$ decreasing in y, and

$$\tilde{h}(y) := \lambda[(\alpha + \alpha_u - \alpha_o)[1 - F_D(y)] + \alpha_o + \mu - \rho. \tag{11}$$

Furthermore, for any $z \in [a, b]$,

$$\mathbb{E}\left[\int_0^\infty |e^{\rho t}\tilde{H}(Y_t)P_t|dt < \infty\right], \quad \mathbb{E}\left[\int_0^\infty |e^{\rho t}\tilde{h}(Y_t)|dt < \infty\right].$$

It is worth mentioning that $\tilde{h}(y)$, the derivative of the modified revenue multiplier function, intuitively represents the impact on business of increasing or decreasing inventory levels. We shall see that this $\tilde{h}(\cdot)$ is a key quantity for characterizing optimal policies.

This lemma enables us to proceed as established in [16].

Step 1. Consider a corresponding switching control problem between two regimes 0 and 1: for a given inventory level z, switching from state 0 to 1 corresponds to inventory increase and switching from state 1 to 0 corresponds to inventory decrease. The cost for inventory change is given by $K^+ + C_h$ and $-K^- + C_h$, and the benefit of being at state 1 is accumulated at rate $\tilde{h}(y)$. Then, when there exists an consistent collection of switching controls so that the resulting singular controls are integrable, then we have

$$V(p, y) = \eta\tilde{H}(a)p + \int_y^b v_0(p, z)dz + \int_a^y v_1(p, z)dz.$$

where v_0 and v_1 are the corresponding value functions for switching controls which can be described analytically as follows:

Proposition 3.2. *v_0 and v_1 are the unique C^1 viscosity solutions with linear growth condition to the following system of variational inequalities:*

$$\min\left\{-\mathcal{L}v_0(p, z),\ v_0(p, z) - v_1(p, z) + K^+ + C_h\right\} = 0, \tag{12}$$

$$\min\left\{-\mathcal{L}v_1(p, z) - \tilde{h}(p, z),\ v_1(p, z) - v_0(p, z) + K^- - C_h\right\} = 0, \tag{13}$$

with boundary conditions $v_0(0^+, z) = 0$ and $v_1(0^+, z) = \max\{-K^- + C_h, 0\}$. Here $\mathcal{L}$ is the generator of the diffusion X^x, killed at rate ρ, given by $\mathcal{L}u(x, z) = \sigma^2 u_{xx}(x, z) + \mu u_x(x, z) - \rho u(x, z)$.

Step 2: Derivation of v_0, v_1.

To solve for v_0, v_1, we see by modifying the argument in [22, Theorem 3.1] that for any given $z \in [a, b]$ and $k \in \{0, 1\}$, an optimal switching control exists and can be described in terms of switching regions: there exist $0 < F(z) < G(z) < \infty$ such that it is optimal to switch from regime 0 to regime 1 (to increase the inventory at level z) when $P_t \in [G(z), \infty)$, and to switch from regime 1 to regime 0 (decrease the inventory at level z) when $P_t \in [0, F(z)]$. Furthermore, based on [22, Theorem 4.2], we see that for each $z \in [a, b]$, the switching regions are described in terms of $F(z)$ and $G(z)$, which take values in $(0, \infty]$. By the regularity of the value functions, $F(z)$ and $G(z)$ can be explicitly derived as follows.

Case I: $K^- - C_h \geq 0$**.** First, for each $z \in (a, b)$ such that $\tilde{h}(z) = 0$, it is never optimal to do anything, so we take $F(z) = \infty = G(z)$, and $v_0(x, z) = 0 = v_1(x, z)$.

Secondly, for z such that $\tilde{h}(z) > 0$, $G(z) < \infty$ and it is optimal to switch from regime 0 to regime 1 (to increase the inventory at level z) when $P_t \in [G(z), \infty)$. Since $K^- - C_h \geq 0$, it is never optimal to switch from regime 1 to regime 0 (i.e. $F(z) = \infty$). Furthermore, we have

$$v_0(x, z) = \begin{cases} A(z)p^n, & x < G(z), \\ \eta h(z)x - (K^+ + C_h), & x \geq G(z), \end{cases}$$

$$v_1(x, z) = \eta h(z)x.$$

Since v_0 is C^1 at $G(z)$, we get

$$\begin{cases} A(z)G(z)^n & = \eta h(z)G(z) - (K^+ + C_h), \\ nA(z)G(z)^{n-1} & = \eta h(z). \end{cases}$$

That is,

$$\begin{cases} G(z) & = \nu h(z)^{-1}, \\ A(z) & = \frac{K^+ + C_h}{(n-1)} G(z)^{-n} = \frac{K^+ + C_h}{(n-1)} \nu^{-n} h(z)^n, \end{cases}$$

where $\nu = (K^+ + C_h)\sigma^2 n(1 - m)$.

Finally, when $h(z) < 0$, it is optimal to switch from regime 1 to regime 0 (reduce inventory at level z) when $P_t \in [F(z), \infty)$. Since $K^+ + C_h > 0$, it is never optimal to switch from regime 0 to regime 1 (i.e. $G(z) = \infty$). The derivation of the value function proceeds analogously to the derivation for the case of $h(z) > 0$.

Case II: $K^- - C_h < 0$**.**

First of all, for each $z \in (a, b)$ such that $\tilde{h}(z) \leq 0$, it is always optimal to reduce the inventory because $(K^- - C_h) < 0$. That is, $F(z) = \infty = G(z)$. In this case, clearly $v_0(x, z) = 0$ and $v_1(x, z) = -K^- + C_h$.

Next, for each $z \in (a,b)$ such that $\tilde{h}(z) > 0$, it is optimal to switch from regime 0 to regime 1 (to increase in the inventory at level z) when $P_t \in [G(z), \infty)$, and to switch from regime 1 to regime 0 when $P_t \in (0, F(z)]$, where $0 < F(z) < G(z) < \infty$.

Moreover, v_0 and v_1 are given by

$$v_0(x,z) = \begin{cases} A(z)p^n, & x < G(z), \\ B(z)p^m + \eta x\tilde{h}(z) - (K^+ + C_h), & x \geq G(z), \end{cases}$$

$$v_1(x,z) = \begin{cases} A(z)p^n - (K^- - C_h), & x \leq F(z), \\ B(z)p^m + \eta x\tilde{h}(z), & x > F(z). \end{cases}$$

Smoothness of $v(x,z)$ at $x = G(z)$ and $x = F(z)$ leads to

$$\begin{cases} A(z)G(z)^n & = B(z)G(z)^m + \eta G(z)\tilde{h}(z) - (K^+ + C_h), \\ nA(z)G(z)^{n-1} & = mB(z)G(z)^{m-1} + \eta\tilde{h}(z), \\ A(z)F(z)^n & = B(z)F(z)^m + \eta F(z)\tilde{h}(z) + (K^- - C_h), \\ nA(z)F(z)^{n-1} & = mB(z)F(z)^{m-1} + \eta\tilde{h}(z). \end{cases} \tag{14}$$

Eliminating $A(z)$ and $B(z)$ from (14) yields

$$\begin{cases} (K^+ + C_h)G(z)^{-m} + (K^- - C_h)F(z)^{-m} & = \frac{-m}{(1-m)\rho}\tilde{h}(z)(G(z)^{1-m} - F(z)^{1-m}), \\ (K^+ + C_h)G(z)^{-n} + (K^- - C_h)F(z)^{-n} & = \frac{n}{(n-1)\rho}\tilde{h}(z)(G(z)^{1-n} - F(z)^{1-n}). \end{cases} \tag{15}$$

Since the viscosity solutions to the variational inequalities are unique and C^1, for every z there is a unique solution $F(z) < G(z)$ to (15). Let $\kappa(z) = F(z)\tilde{h}(z)$, $\nu(z) = G(z)\tilde{h}(z)$, then the following system of equations for $\kappa(z)$ and $\nu(z)$ is guaranteed to have a unique solution for each z:

$$\begin{cases} (K^+ + C_h)\nu(z)^{-m} + (K^- - C_h)\kappa(z)^{-m} & = \frac{-m}{(1-m)\rho}(\nu(z)^{1-m} - \kappa(z)^{1-m}), \\ (K^+ + C_h)\nu(z)^{-n} + (K^- - C_h)\kappa(z)^{-n} & = \frac{n}{(n-1)\rho}(\nu(z)^{1-n} - \kappa(z)^{1-n}). \end{cases}$$

Moreover, these equations depend on z only through $\nu(z)$ and $\kappa(z)$, implying that there exist unique constants κ, ν such that $\kappa(z) \equiv \kappa$ and $\nu(z) \equiv \nu$ for all z. Hence $F(z) = \kappa\tilde{h}(z)^{-1}$, $G(z) = \nu\tilde{h}(z)^{-1}$, with $\kappa < \nu$ being the unique solutions to

$$\begin{cases} \frac{1}{1-m}\left[\nu^{1-m} - \kappa^{1-m}\right] & = -\frac{\rho}{m}\left[(K^+ + C_h)\nu^{-m} + (K^- - C_h)\kappa^{-m}\right], \\ \frac{1}{n-1}\left[\nu^{1-n} - \kappa^{1-n}\right] & = \frac{\rho}{n}\left[(K^+ + C_h)\nu^{-n} + (K^- - C_h)\kappa^{-n}\right]. \end{cases}$$

Given $F(z)$ and $G(z)$, $A(z)$ and $B(z)$ are solved from Eq. (14),

$$\begin{cases} B(z) & = -\frac{G(z)^{-m}}{n-m}\left(\frac{G(z)\tilde{h}(z)}{\sigma^2(1-m)} - n(K^+ + C_h)\right) = -\frac{F(z)^{-m}}{n-m}\left(\frac{F(z)\tilde{h}(z)}{\sigma^2(1-m)} + n(K^- - C_h)\right), \\ A(z) & = \frac{G(z)^{-n}}{n-m}\left(\frac{G(z)\tilde{h}(z)}{\sigma^2(n-1)} + m(K^+ + C_h)\right) = \frac{F(z)^{-n}}{n-m}\left(\frac{F(z)\tilde{h}(z)}{\sigma^2(n-1)} - m(K^- - C_h)\right). \end{cases}$$

Step 3. F and G are increasing in z, and therefore by definition this collection of optimal switching controls is consistent. This consistent collection of optimal switching control corresponds to an admissible singular control $(\hat{\xi}^+, \hat{\xi}^-) \in \mathcal{A}_y$. Moreover, since $\mathcal{I}$ is bounded, it is integrable following [16, Theorem 3.10] as

$$\lim_{t\to\infty} \mathbb{E}\left[e^{-\rho t} G^{-1}(M_t)\right] = 0.$$

(See also Lemma 1 and Eqn. (23) in [24]).

As a result,

$$v_0(p,z) = \begin{cases} A(z)p^n, & p < G(z), \\ B(z)p^m + \eta\tilde{h}(z)p - K^+ - C_h, & p \geq G(z), \end{cases}$$

$$v_1(p,z) = \begin{cases} A(z)p^n - K^- + C_h, & p \leq F(z), \\ B(z)p^m + \eta\tilde{h}(z)p, & p > F(z). \end{cases}$$

Step 4. Combining these results, we see that the ordering region is given by $\{(p,z) : p \geq G(z)\}$ and the downsizing region by $\{(p,z) : p \leq F(z)\}$. It is optimal to take no action when when (X_t, Y_t) is in the continuation region, given by $\{(p,z) : F(z) < x < G(z)\}$. If (p,y) is in the ordering (or downsizing) region, then a jump is exerted at time zero to make $Y_{0+} = G^{-1}(p)$ (or $Y_{0+} = F^{-1}(p)$).

Finally, by [16, Theorem 3.10], we have

$$V(p,y) = \eta\tilde{H}(a)p + \int_y^b v_0(p,z)dz + \int_a^y v_1(p,z)dz.$$

3.2 Main Result

In summary, we see that the optimal value function is characterized below for two distinct cases. In the first, $K^- - C_h \geq 0$, implying that the proportional loss incurred upon selling inventory is greater than the gain from reduced future holding cost. In the second, $K^- - C_h < 0$, implying that reducing holding cost dominates the transaction cost.

Theorem 3.3. [Optimal value function for $K^- - C_h \geq 0$]

$$V(p,y) = \eta\tilde{H}(a)p + \int_a^y v_1(p,z)dz + \int_y^b v_0(p,z)dz, \tag{16}$$

where v_0 and v_1 are given by

1. *For each $z \in (a,b)$ such that $\tilde{h}(z) = 0 : v_0(p,z) = v_1(p,z) = 0.$*

2. For each $z \in (a,b)$ such that $\tilde{h}(z) > 0$:

$$\begin{cases} v_0(p,z) &= \begin{cases} A(z)p^n, & p < G(z), \\ \eta\tilde{h}(z)p - K^+ - C_h, & p \geq G(z), \end{cases} \\ v_1(p,z) &= \eta\tilde{h}(z)p, \end{cases}$$

where $G(z) = \nu\tilde{h}(z)^{-1}$, and $A(z) = \frac{K^+ + C_h}{(n-1)} G^{-n}(z)$, with $\nu = (K^+ + C_h)\sigma^2 n(1-m)$.

3. For each $z \in (a,b)$ such that $h(z) < 0$:

$$\begin{cases} v_0(p,z) &= 0, \\ v_1(p,z) &= \begin{cases} B(z)p^n + \eta\tilde{h}(z)p, & p < F(z), \\ -K^- + C_h, & p \geq F(z), \end{cases} \end{cases}$$

where $F(z) = -\frac{\kappa}{\tilde{h}(z)}$, and $B(z) = \frac{K^- - C_h}{(n-1)} \kappa^{-n} F^{-n}(z)$, with $\kappa = (K^- - C_h)\sigma^2 n(1-m)$.

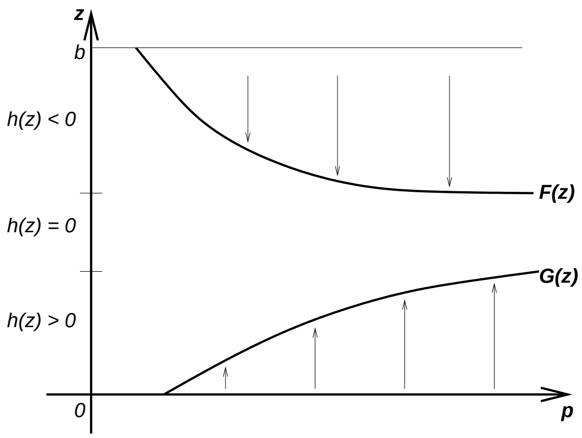

Figure 1: Policy when $K^- - C_h \geq 0$, with $F(z) = -\frac{\kappa}{\tilde{h}(z)}$ and $G(z) = \nu\tilde{h}(z)^{-1}$.

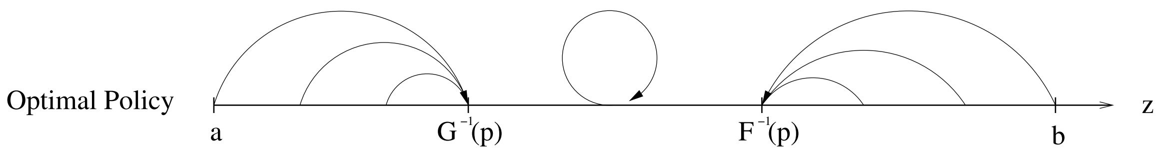

Figure 2: Illustration of the two-threshold order policy for fixed price p when $K^- - C_h \geq 0$.

Theorem 3.4. [Optimal value function for $K^- - C_h < 0$]

$$V(p,y) = \eta\tilde{H}(a)p + \int_a^y v_1(p,z)dz + \int_y^b v_0(p,z)dz, \tag{17}$$

where v_0 and v_1 are given by

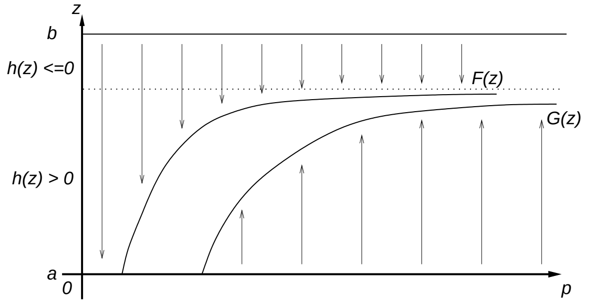

Figure 3: Policy when $K^- - C_h < 0$.

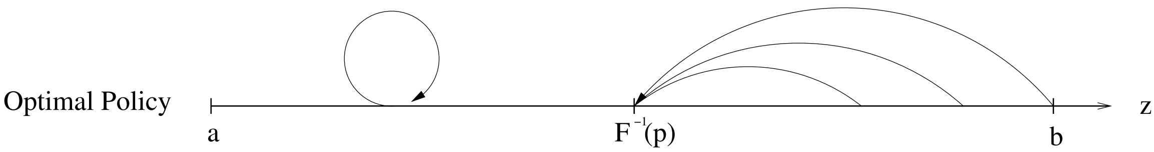

Figure 4: Illustration of the one-threshold order policy for fixed (low) price p and when $K^- - C_h < 0$.

1. *For each $z \in (a,b)$ such that $\tilde{h}(z) \le 0$: $v_0(p,z) = 0, v_1(p,z) = -K^-$.*

2. *For each $z \in (a,b)$ such that $\tilde{h}(z) > 0$:*

$$v_0(p,z) = \begin{cases} A(z)p^n, & p < G(z), \\ B(z)p^m + \eta\tilde{h}(z)p - K^+ - C_h, & p \ge G(z), \end{cases} \tag{18}$$

$$v_1(p,z) = \begin{cases} A(z)p^n - K^- + C_h, & p \le F(z), \\ B(z)p^m + \eta\tilde{h}(z)p, & p > F(z). \end{cases} \tag{19}$$

Here

$$A(z) = \frac{\tilde{h}(z)^n}{(n-m)\nu^n}\left(\frac{\nu}{\sigma^2(n-1)} + m(K^+ + C_h)\right) \tag{20}$$

$$B(z) = \frac{-\tilde{h}(z)^m}{(n-m)\nu^m}\left(\frac{\nu}{\sigma^2(1-m)} - n(K^+ + C_h)\right). \tag{21}$$

The functions F and G are non-decreasing with

$$F(z) = \frac{\kappa}{\tilde{h}(z)} \quad \text{and} \quad G(z) = \frac{\nu}{\tilde{h}(z)}, \tag{22}$$

where $\kappa < \nu$ are the unique solutions to

$$\frac{1}{1-m}\left[\nu^{1-m} - \kappa^{1-m}\right] = -\frac{\rho}{m}\left[(K^+ + C_h)\nu^{-m} + (K^- - C_h)\kappa^{-m}\right], \tag{23}$$

$$\frac{1}{n-1}\left[\nu^{1-n} - \kappa^{1-n}\right] = \frac{\rho}{n}\left[(K^+ + C_h)\nu^{-n} + (K^- - C_h)\kappa^{-n}\right]. \tag{24}$$

Theorem 3.5. **[Optimal control for** $K^- - C_h \geq 0$**]** *For each* $z \in (a, b)$*, the optimal control is described in terms of* $F(z)$ *and* $G(z)$ *from Theorem 3.3 such that*

- *For* z *such that* $\tilde{h}(z) > 0$*, it is optimal to increase inventory past level* z *when* $P_t^p \in [G(z), \infty)$*, and never decreases.*
- *When* $\tilde{h}(z) < 0$*, it is optimal to decrease below inventory level* z *when* $P_t^p \in [F(z), \infty)$*, and it is never optimal to increase. When* $\tilde{h}(z) = 0$*, it is optimal to do nothing (i.e.* $F(z) = \infty = G(z)$*).*

Theorem 3.6. **[Optimal control for** $K^- - C_h < 0$**]** *For each* $z \in (a, b)$*, the optimal control is described in terms of* $F(z)$ *and* $G(z)$ *from Theorem 3.4 such that*

- *For* z *such that* $\tilde{h}(z) > 0$*, it is optimal to increase inventory past level* z *when* $P_t^p \in [G(z), \infty)$*, and to decrease invenotry below level* z *when* $P_t^p \in (0, F(z)]$*.*
- *For* z *such that* $\tilde{h}(z) \leq 0$*, it is always optimal to decrease inventory level.*

The optimal policy is illustrated in Figure 1 and Figure 2 for the $K^- \geq C_h$ case, and Figure 3 and Figure 4 for the $K^- < C_h$ case, with corresponding numerical examples with specific parameter values in Figure 5 and Figure 6. In Figure 3 (and Figure 4), $K^- < C_h$, implying a relatively high holding cost. For low prices, inventory is decreased regardless of the value of $\tilde{h}(z)$, and for high prices, inventory is increased or decreased as necessary. For intermediate prices, in general, no action is taken if inventory is low enough, but otherwise it is decreased (where Figure 4 illustrates this last case). In contrast, when $K^- \geq C_h$ as in Figure 1 (and Figure 2), the relatively low holding cost introduces a different policy: above a certain threshold price, except around the $\tilde{h}(z) = 0$ region, inventory is typically decreased for negative $h(z)$ values, and increased for positive $h(z)$ values, as illustrated in Figure 2. In general, depending on the holding cost C_h and the cost of selling K^-, reducing holding cost is a key driver, so conditions must be more favorable before inventory is increased, and it is more likely that inventory will be decreased.

Remark 3.7. *We emphasize that these results are quite general, and indeed hold for any* $\tilde{H}(\cdot)$ *function that is concave. In particular, when* $\tilde{H}$ *is continuously differentiable, strictly increasing, strictly concave, we will have the regularity condition for* F *and* G *and for the value function, as postulated in the current literature.*

4 Computational Experiments and Observations

The purpose of our computational experiments are to assess when it might be valuable for a firm to actively utilize the spot market to guard against cost and price risks. To that end, we compared a variety of simulation runs in which the market met the conditions of our model, and where two different inventory management strategies were employed. The first of these was the optimal policy as described in this paper, and the second of these was a modified version of the traditional newsvendor solution. We elected to use this modified newsvendor approach as a reasonable proxy for how a manager who is not interested in repeatedly buying and selling on the spot market might manage the system.

For our newsvendor solution, since D is independent from τ and the price process, p_t, we modified the standard news-vendor approach and raised the inventory level to the critical fractile y^* as follows:

$$y^* = \min\left\{b, \max\left\{a, F^{-1}\left(\frac{E[p_\tau] - (p_0 + K^+ + C_h * E[\tau])}{E[p_\tau]}\right)\right\}\right\},$$

where $E[p_\tau]$ is the expected selling price, and $p_0 + K^+ + C_h * E[\tau]$ is the expect inventory stock-up cost. Moreover, in our case, $E[\tau] = 1/\lambda$, $E[p_\tau] = p_0 * \lambda/(\lambda - \mu)$, therefore:

$$y^* = \min\left\{b, \max\left\{a, F^{-1}\left(1 - \frac{(\lambda - \mu) * (p_0 + K^+ + C_h/\lambda)}{(p_0 * \lambda)}\right)\right\}\right\}.$$

4.1 $K^- \geq C_h$

For the $K^- \geq C_h$ simulation runs, we used the following parameters: price multiplier $\alpha = 1.3$, no penalty cost or salvage value ($\alpha_0 = \alpha_u = 0$, a price process with $\mu = 0.1$ and $\sigma = 1$, Lognormal demand with mean 5 and standard deviation 0.5, time until demand arrives exponentially distributed with rate $\lambda = 1$, additional proportional transaction cost for buying $K^+ = 1$, additional proportional transaction cost for selling $K^- = 2$, minimum inventory $a = 50$ and maximum inventory $b = 200$. For these parameter values, we considered all four combinations of two starting inventory levels ($y = 50, 150$) and two holding costs ($C_h = 0, 1$), and varied initial prices between 5 and 50.

For each parameter, we completed 25000 simulation runs, and in Tables 1 - 4, we report average profit for the optimal strategy W, average profit for the newsvendor strategy W^{nv}, improvement of the optimal strategy over the newsvendor strategy, coefficient of variation of the optimal strategy (CoV W, the ratio of standard deviation to average), and coefficient of variation of the newsvendor strategy (CoV W^{nv}). In Figure 5, we illustrate the optimal policy for one set of parameters.

We see that in each case, the optimal strategy is significantly better than the newsvendor strategy, particularly when the initial price is low. As the price increases, the advantage of the optimal strategy decreases, but always remains significant ranging from a high of several thousand percent better, to at low of 27 percent better. This advantage is particularly dramatic for runs with high starting inventory and high holding cost, or low starting inventory

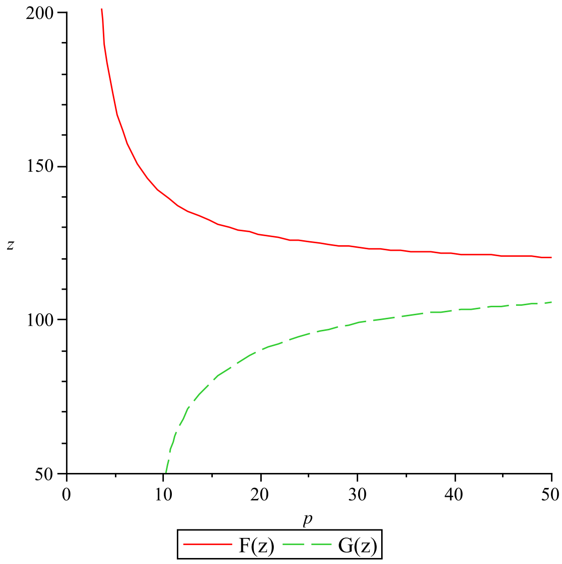

Figure 5: Example policy when $K^- - C_h \geq 0$, with $F(z) = -\frac{\kappa}{\tilde{h}(z)}$ and $G(z) = \frac{\nu}{\tilde{h}(z)}$. Here $\alpha = 1.3, \alpha_0 = 0, y = 50, C_h = 1, \mu = 0.1, \sigma = 1/\sqrt{2}, [a, b] = [50, 250], K^+ = 1, K^- = 2, \lambda = 1$, $D \backsim Lognormal(5, 0.5)$.

p	5	10	15	20	25	30	35	40	45	50
W	35	205	398	597	797	999	1201	1403	1605	1808
W^{nv}	16	133	249	365	545	728	906	1081	1255	1429
$\frac{W-W^{nv}}{W^{nv}}$	114%	54%	60%	63%	46%	37%	33%	29.76%	28%	27%
CoV W	59.3	21.0	16.4	14.7	13.8	13.2	12.9	12.6	12.4	12.2
CoV W^{nv}	61.6	15.2	12.2	11.1	10.6	10.3	10.1	10.0	9.9	9.8

Table 1: $K^- \geq C_h$,$C_h = 1$, $K^- = 2$, $y = 50$

and low holding cost, as these seem to be the cases where aggressive and active management of inventory leads to the greatest benefits. In addition to being significantly more profitable, when initial prices are low for some combinations of starting inventory and holding cost (either low starting inventory and no holding cost, or high starting inventory and holding cost), the coefficient of variation of the optimal strategy is less than that of the newsvendor strategy, suggesting higher returns at lower risk.

4.2 $K^- < C_h$

We also considered the case of $K^- < C_h$, to explore whether or not the qualitatively different shapes of the curves would lead to different results. We used the same parameters for our simulation, except that we only considered a holding cost $C_h = 1$ and $K^- = .5$. As before,

p	5	10	15	20	25	30	35	40	45	50
W	-111	84	285	488	691	894	1098	1302	1506	1710
W^{nv}	-284	-168	-51	65	266	462	648	829	1007	1184
$\frac{W-W^{nv}}{W^{nv}}$	NA	NA	NA	653%	160%	94%	69%	57%	50%	44%
CoV W	21.1	54.7	24.1	18.7	16.5	15.3	14.5	14.0	13.6	13.3
CoV W^{nv}	NA	NA	NA	62.4	21.7	16.2	14.1	13.0	12.3	11.8

Table 2: $K^- \geq C_h$,$C_h = 1$, $K^- = 2$, $y = 150$

p	5	10	15	20	25	30	35	40	45	50
W	133	302	478	655	833	1011	1189	1367	1545	1724
W^{nv}	104	209	369	526	681	835	989	1143	1297	1450
$\frac{W-W^{nv}}{W^{nv}}$	27%	45%	29%	25%	22%	21%	20%	20%	19%	19%
CoV W	13.5	14.0	12.0	11.3	10.9	10.7	10.6	10.5	10.4	10.3
CoV W^{nv}	6.5	6.5	6.8	6.9	6.9	6.9	6.9	6.9	6.9	6.9

Table 3: $K^- \geq C_h$,$C_h = 0$, $K^- = 2$, $y = 50$

we considered starting inventory levels ($y = 50, 150$) and varied initial prices between 5 and 50. We report these results in Tables 5 and 6. Observe that the same general observations made for the $K^- \geq C_h$ case apply here. In Figure 6, we illustrate the optimal policy for one set of parameters.

4.3 Changes in λ

It can be observed from the formulation that both threshold $F^{-1}(p)$ and threshhold $G^{-1}(p)$ are monotonic with respect to λ. However, when $K^- > C_h$, $F^{-1}(p)$ and $G^{-1}(p)$ are decreasing with respect to λ, as illustrated in Figure 7. This is not surprising, because as the rate of demand arrival increases, there is less flexibility to stay "idle", and it becomes important to

p	5	10	15	20	25	30	35	40	45	50
W	157	322	491	666	843	1021	1200	1380	1559	1739
W^{nv}	-96	8	204	376	539	698	855	1012	1168	1323
$\frac{W-W^{nv}}{W^{nv}}$	NA	3730%	140%	77%	57%	46%	40%	36%	34%	31%
CoV W	10.5	10.0	9.6	9.4	9.2	9.1	9.0	8.9	8.8	8.8
CoV W^{nv}	NA	162.6	12.4	9.7	8.8	8.3	8.0	7.8	7.7	7.6

Table 4: $K^- \geq C_h$,$C_h = 0$, $K^- = 2$, $y = 150$

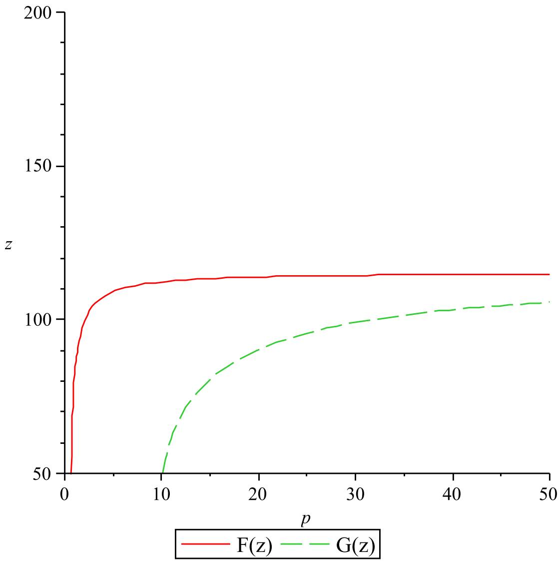

Figure 6: Example policy when $K^- - C_h < 0$, with $F(z) = \frac{\kappa}{\tilde{h}(z)}$ and $G(z) = \frac{\nu}{\tilde{h}(z)}$. Here $\alpha = 1.3, \alpha_0 = 0, y = 50, C_h = 1, \mu = 0.1, \sigma = 1/\sqrt{2}, [a, b] = [50, 250], K^+ = 1, K^- = 0.5, \lambda = 1$, $D \backsim Lognormal(5, 0.5)$.

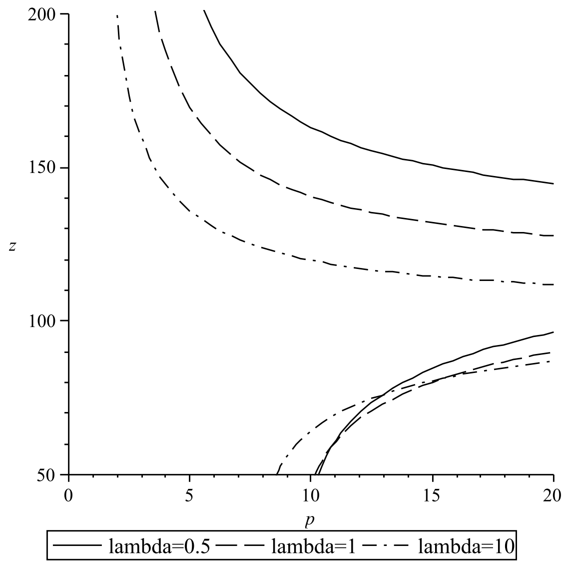

Figure 7: Threshold changes with respect to change of λ (with other parameters identical to those for Figure 5)

p	5	10	15	20	25	30	35	40	45	50
W	13	160	327	500	675	851	1027	1204	1380	1557
W^{nv}	4	108	212	317	475	637	794	949	1103	1258
$\frac{W-W^{nv}}{W^{nv}}$	233%	48%	54%	58%	42%	34%	29%	27%	25%	24%
CoV W	102.0	17.8	13.4	11.8	11.0	10.5	10.2	10.0	9.8	9.7
CoV W^{nv}	171.9	12.6	9.6	8.6	8.2	7.9	7.8	7.7	7.6	7.5

Table 5: $K^- < C_h$,$C_h = 1$, $K^- = 0.5$, $y = 50$

p	5	10	15	20	25	30	35	40	45	50
W	-88	88	266	445	625	804	984	1164	1345	1525
W^{nv}	-146	-41	63	167	338	507	669	828	985	1142
$\frac{W-W^{nv}}{W^{nv}}$	NA	NA	323%	166%	85%	59%	47%	41%	36%	34%
CoV W	NA	33.9	17.0	13.6	12.1	11.3	10.8	10.4	10.2	10.0
CoV W^{nv}	NA	NA	32.5	16.3	11.5	10.0	9.2	8.8	8.5	8.3

Table 6: $K^- < C_h$,$C_h = 1$, $K^- = 0.5$, $y = 150$

respond more quickly to price changes by either selling excess inventory or buying to adjust inadequate inventory levels.

5 Future Research

We have completely characterized the optimal policy for a firm facing random demand after a random period of time of being able to buy and sell on the spot market. In computational tests, we observed that this policy performs significantly better than a version of the traditional newsboy policy (utilized as a proxy for a reasonable inventory management policy for firms not interested in trading on the spot market), most notably when market prices are relatively low.

Although the results presented in this paper provide insight into the value to a firm of effectively utilizing the spot market, and contribute to the state of the art in continuous time inventory control, there are significant extensions possible to this work from both technical and modelling perspectives.

For example, the addition of a fixed inventory ordering cost changes the singular control to a more difficult impulse control problem. It will be interesting to see if the analogous state dependent version of (d, D, U, u) policy still holds for this two-dimensional problem, and whether the regularity property holds as well. Additionally, the price process could be modelled by stochastic processes other than a Brownian motion. For instance, it would be interesting to explore whether or not the two-threshold policy holds for the case of a mean-reverting process. More sophisticated constraints on the inventory, such as a requirement that the inventory either be 0 or above some minimum level, are in theory not more difficult by our solution approach, but it would be interesting to complete this analysis. Finally, multi-period models, with multiple demand opportunities and inventory carried between periods will be significantly more difficult to analyze, but may yield interesting insights on effective inventory management in the presence of a spot market.

References

[1] Akella, R. V.F. Araman, and J. Kleinknecht. 2001. B2B Markets: Procurement and Supplier Risk Management. Geunes H.E.R.J, P.M. Pardalos, and E. Romeijn (eds.). Supply Chain Management: Models, Applications, and Research Directions. Kluwer: Dordrecht.

[2] Archibald, A. 1981. Continuous Review (s, S) Policies with Lost Sales. Management Science. 27, 1171-1177.

[3] Bather, J. A. 1966. A Continous Time Inventory Model. Jounal of Applied Probability. 3, 538-549.

[4] Berling, P. Holding Inventory with Stochastically Mean-Reverting Purchase Price. Working Paper, Department of Industrial and Management Logistics, Lund University.

[5] Billington, C. 2002. HP Cuts Risk With Portfolio Approach. Purchasing Magazine Online, February 21st, 2002.

[6] Chaouc, B. 2006. Inventory Control and Periodic Price Discounting Campaigns. Naval Research Logistics, Vol. 54. Published Online 14 July 2006 in Wiley InterScience.

[7] Chen, F. and J-S. Song. 2001. Optimal Polices for Multiechelon Inventory Problems with Markov-Modulated Demand. Operations Resarch, Vol. 49, No. 2, 226-234.

[8] Cheng, F. and S.P. Sethi. 1999. Optimality of State-Dependent (s,S) Policies in Inventory Models with Markov-Modulated Demand and Lost Sales. Production and Operations Management, Vol. 8, No. 2, 183-192.

[9] Cohen, M. and N. Agrawal. 1999. An Analytical Comparison of Long and Short Term Contracts. IIE Transactions, Vol. 31, 783-796.

[10] Constantinides, G.M. and S. Richard. 1978. Existence of Optimal Simple Policies for Discounted-Cost Inventory and Cash Management in Continuous Time. Operations Research. Vol. 26, No. 4, 620-636.

[11] Fu, Q., C-Y Lee, and C-P Teo. 2006. Procurement Risk Management Using Options: Random Spot Price and the Portfolio Effect. Working paper.

[12] Gavirneni, S. and T.E. Morton. 1999. Inventory Control Under Speculation: Myopic Heuristics and Exact Procedures. European Journal of Operational Research, Vol. 117, 211-221.

[13] Goh, M. and M. Sharafali. 2002. Price-Dependent Inventory Models with Discount Offers at Random Times. Production and Operations Management, Vol. 11, No. 2, 139-156.

[14] Golovachkina, N. and J. Bradley. 2002. Supplier-Manufacturer Relationships under Forced Compliance Contracts. Working paper, Johnson School of Management, Cornell University.

[15] Golabi, K. 1985. Optimal Inventory Policies when Ordering Prices are Random. Operations Research, Vol. 33, No. 3, 575-588.

[16] Guo, X. and P. Tomecek. 2007. Connections Between Singular Control and Optimal Switching, *SIAM Journal on Control and Optimization*, To appear.

[17] Haksoz, C. and S. Seshadri. 2007. Supply Chain Operations in the Presence of Spot Market: A Review with Discussion. Journal of the Operational Research Society advance online publication, 11 April 2007.

[18] Harrison, J. M. and M. Taksar. 1983. Instantaneous control of Brownian motion, *Math. Oper. Res.*, 8(3): 439-453.

[19] Harrison, J. M. and T. M. Selke, and A. J. Taylor. 1983. Impulse Control of Brownian Motion, *Mathematics of Operations Research*, 8(3): 454-466.

[20] Kalymon, B. A. 1971. Stochastic Prices in a Single-Item Inventory Purchasing Model. Operations Research, Vol. 19, No. 6, 1434-1458.

[21] Li, C-L and P. Kouvelis. 1999. Flexible and Risk-Sharing Contracts under Price Uncertainty. Management Science Vol. 45, No. 10, 1378-1398.

[22] Ly Vath, V and H. Pham. 2007. Explicit Solution to An optimal Awitching Problem In The Two-Regime Aase, SIAM Journal on Control and Optimization, Vol 46, No.2, 395-426.

[23] Martinez-de-Albeniz, V. and D. Simchi-Levi. 2005. A Portfolio Approach to Procurement Contracts. Production and Operations Management, Spring 2005, Vol 14, No. 1, 90-114.

[24] Merhi, A. and M. Zervos. A Model for Reversible Investment Capacity Expansion,, *SIAM Journal on Control and Optimization*, To appear, (2007).

[25] Moinzadeh, K. 1997. Replenishment and Stocking Policies for Inventory Systems with Random Deal Offerings. Management Science, Vol. 43, No. 3, 334-342.

[26] Ormeci, A., J. G. Dai, and J. V. Vate. 2006. Impulse Control of Brownian Motion: the Constrained Average Cost Case. Preprint, Georgia Institute of Technology.

[27] Sethi, S. P. and F. Cheng. 1997. Optimality of (s,S) Policies in Inventory Models with Markovian Demand. Operations Research, Vol. 45, No. 6, 931-939.

[28] Seifert, R. W., U. W. Thonemann, and W. H. Hausman. 2004. Optimal Procurement Strategies for Online Spot Markets. European Journal of Operational Research, Vol. 152, 781-799.

[29] Simchi-Levi, D., P. Kaminsky, and E. Simchi-Levi. 2008. Designing and Managing the Supply Chain. Boston: McGraw-Hill.

[30] Spinler S., A. Huchzermeier, and P. R. Kleindorfer 2002. The Valuation of Options on Capacity. Working paper, the Wharton School.

[31] Sulem, A. 1986. A Solvable One-Diemensional Model of a Diffusion Inventory System. Mathematics of Operations Research. Vol. 11, No. 1, 125-133.

[32] Taksar, M. I. 1985. Average Optimal Singular Sontrol and a related Stopping Problem. Mathematics of Operations Research. Vol. 10, 63-81.

[33] Wang, Y. 2001. The Optimality of Myopic Stocking Policies for Systems with Decreasing Prices. European Journal of Operational Research, Vol 133, 153-159.

[34] Wu D. J., P. R. Kleindorfer, and J. E. Zhang. 2002. Optimal Bidding and Contracting Strategies for Capital-Intensive Goods. European Journal of Operational Research, 137(3), 657-676.

[35] Yi, J. and A. Scheller-Wolf. 2003. Dual Sourcing from a Regular Supplier and a Spot Market. Working Paper, Carnegie Mellon Tepper School of Business.