

The Impact of Financial Turbulence on Inventory Control

Lucy G. Chen • Robin O. Roundy • Rachel Q. Zhang • Lawrence W. Robinson

** Business School, National University of Singapore, Singapore, 117592*

*** School of Operations Research and Industrial Engineering,
Cornell University, Ithaca, NY 14853*

**** Department of Industrial Engineering and Logistics Management
Hong Kong University of Science and Technology
Clear Water Bay, Kowloon, Hong Kong*

***** Johnson Graduate School of Management, Cornell University, Ithaca, NY 14853
bizcg@nus.edu.sg • robin@orie.cornell.edu • rzhang@ust.hk • lwr2@cornell.edu*

Abstract

In global supply chains, both the operational costs and discount factors of a firm may be strongly affected by unpredictable national and international economic forces. In this paper, we incorporate such financial uncertainties into a single-product, periodic-review, finite-horizon stochastic inventory system by modeling both the discount factors and operational costs as stochastic processes that evolve as financial uncertainties are realized. We establish conditions under which (s, S) inventory policies are optimal when discount factors and costs are stochastic and correlated both to one another and over time. Most of these conditions are related to speculative motives for inventory and backorders. In addition, we provide an illustrative example of the benefits of incorporating stochastic discount factors and operational costs by looking at four types of Mexican firms that conduct sizeable business with the U.S. The likelihood of financial instability increases dramatically during the months surrounding Mexican presidential elections. A numerical case study based on Mexican financial data shows that both the business type and the stockout protocol have large impacts on inventory decisions and on the cost penalty for ignoring the dynamic nature of the financial environment (cost penalties of up to 78%). To the best of our knowledge, ours is the first inventory model that incorporates stochastic discount factors and operational costs that evolve in response to changing financial conditions.

1. Introduction

In this paper, we consider inventory decisions for firms that conduct business internationally. For instance, a firm may buy raw material and produce a product abroad (domestically), and then sell it domestically (abroad). In anticipation of near-term instabilities such as wars, political events, and

economic cycles, firms need to understand the impact of such risks and adjust their major operational decisions, *e.g.*, production planning and inventory decisions, accordingly. In almost all research and practice to date, it has been generally assumed that operational costs as well as the discount factors, are fixed and known in advance. Under these conditions, it is well established that variants of (Q, r) and (s, S) policies are optimal under most conditions.

However, it has long been recognized in the finance literature that risks in both international and domestic markets can have a significant impact on a firm's operational costs and the discount factors that a firm uses for evaluating its future cash flows and making operational decisions. According to studies on international asset pricing (such as Stulz (1994)), both national market indices (*e.g.*, interest rates and domestic market return) and international market indices (*e.g.*, exchange rates, returns of certain international markets and interest rates in markets in which the firm has assets invested) can be effectively modeled as random variables. As a result, both the operational cost parameters (*e.g.*, inventory holding costs, backorder costs, and purchasing prices) and the discount factor should be treated as stochastic processes that evolve as these financial indices are realized over time.

Researchers in operations have of course recognized the risks associated with inventory. Some have tried to incorporate risks of future cash flows into inventory decision-making in deterministic inventory models. For instance, Beranek (1967) investigated the impact of the timing of the loan repayments on inventory carrying costs and hence the optimal ordering policy. Grubbstrom (1980) illustrated ways to correctly determine the capital costs of work-in-progress and finished goods inventory, clarifying the relationships between the payment structure and the stocks of finished goods and work-in-progress.

Some excellent work has been done on incorporating risks of holding inventory by specifying the opportunity cost of capital within a stochastic inventory model. Anvari (1987) and Kim and Chung (1989) both showed that the level of inventory determines the risk and hence the opportunity cost of capital. They concluded that the opportunity cost of capital for investments in inventory is an increasing function of the inventory level, and is higher for firms facing more risky demand. Kim and Chung (1989) also showed that the optimal inventory level decreases as the demand volatility increases. Singhal *et al.* (1994) analyzed the effect of risks on lot sizes and reorder points. They showed that an increase in demand riskiness leads to a lower reorder point and a smaller lot size with relatively short replenishment lead times. However, when lead times are long, increase in demand riskiness still decreases the reorder point, but may result in a larger lot size. They also concluded that

cost minimization models with a fixed opportunity cost of capital have a large penalty in overall costs, compared with a model where the opportunity cost of capital is adjusted for the risk of cash flows from inventory decisions.

Recently, several inventory researchers have used both operational and financial hedging to ameliorate the risks associated with demand and exchange rates. Van Mieghem (2003) studied ways to mitigate risks in capacity investment by using operational and financial hedging. Van Mieghem (2004) studied three canonical networks: a product-dedicated network, a sequential network featuring resource-sharing, and a parallel network that allows for resource substitution. He showed that resource levels may actually increase under a general utility function with risk-aversion. He also showed that the optimal levels of flexible technology, substitutable inventory, and resource imbalance do increase in risk aversion with operational hedging. Gaur and Seshadri (2005) used options to construct an optimal hedging transaction that minimizes profit variance and increases the expected utility for a risk-averse decision maker, when demand is correlated with the price of a financial asset. They showed that a risk-averse decision maker will order more inventory when he is able to hedge. At a more macro level, Axaroglou and Kouvelis (2001) investigated the ownership structure of a firm's foreign subsidiaries who supply a foreign market, and showed that exchange rates are an important factor for this decision. Despite the excellence of these works, they did not address the question how firms should adjust their inventory replenishment policies in expecting the potential large financial turbulence in the near future.

The research that is most related to ours is conducted by Ding *et al.* (2004) and Kazaz *et al.* (2005). Both papers examine the impact of exchange-rate changes on the optimal capacity investment decision in a global manufacturing network. While Kazaz *et al.* (2005) mainly looks at the operational hedging, Ding *et al.* (2004) explores integrated operational and financial hedging, and they found that firms using financial hedging are less sensitive to demand and exchange rate volatilities, and that allocation and financial hedging enable a firm to exercise more control over its profit variance. While these literatures carefully study the impact of exchange-rate changes on the inventory decision making, they ignore the impact of the volatile discount factor, which is one focus of our paper. Our paper also differs from these papers in the problem setting – while these papers study a one-period problem, we look at a finite horizon problem, which is a sensible model to examine especially if the financial turbulence lasts several review periods.

In this paper, we incorporate both domestic and international risk into a single-product, periodic-review, finite-horizon stochastic inventory system by allowing both the discount factors and operational costs to be stochastic and to evolve over time as financial information is realized. We propose a risk-factor model for the stochastic discount factors, and analyze the significance of the risks in inventory replenishment decisions for four different types of firms under three different stockout protocols.

On a theoretical level, we establish conditions under which (s, S) inventory policies are optimal when discount factors, costs, and demands are correlated and stochastic (we also provide, through a numerical case study, different types of replenishment policies that are optimal if these conditions fail to hold). These conditions result from speculative motives for inventory and backorders in the presence of financial turbulence. In the case of cyclic periods of significant financial risks (such as the Mexican presidential elections), we develop methodologies for using historical data to obtain a model of the manner in which discount factors evolve.

In addition, we illustrate the benefits of incorporating stochastic discount factors and operational costs through an example of four types of Mexican firms that conduct sizeable business with the U.S. The likelihood of financial instability increases dramatically during the months before and after Mexican presidential elections. A numerical case study based on Mexican financial data shows that both the business type and stockout protocol have large and very different impacts on inventory decisions, and that the cost penalty for ignoring the stochastic nature of the financial environment can be severe (cost penalties of up to 78%).

Our work provides models and algorithms that will enable firms who operate internationally to better understand and control risks and to improve their profitability in the face of periods of financial instability. To the best of our knowledge, this is the first paper that incorporates into inventory models stochastic discount factors and operational costs that evolve in response to changing financial conditions.

The paper is organized as follows. In Section 2 we introduce an inventory model with stochastic discount factors and cost parameters. In Section 3 we establish conditions under which the (s, S) policy is optimal for this model, and discuss the implications of these conditions. In Section 4, based on empirical study on asset pricing, we identify financial risk factors that affect discount factors and operational costs in a global setting and propose models for estimating the stochastic discount factors that are updated from period by period as financial information is revealed. In Section 5, we conduct

a numerical case study on the optimal ordering policy and calculate the cost penalties for ignoring the dynamic nature of the discount factors and costs. In Section 6 we summarize the paper.

2. Problem Formulation

2.1 Model Description

We consider a periodic review inventory system with a single product. At the beginning of each period, purchasing decisions and payments to the suppliers are made. Demand then occurs and available inventory is delivered to customers. Holding and stockout costs are incurred, and revenues are realized.

We consider both the backorder case and the lost sales case (referred to as LS). Under backorder, we consider both the case in which the customers pay when they place orders, referred to as the “pay-to-order” case (PTO), and the case in which they pay when their orders are delivered, called the “pay-to-delivery” case (PTD). In both cases, we assume that corporate policy calls for demand to be met as soon as inventory becomes available. Customers pay for the product in their domestic currency, and all costs and revenues are converted to the firm’s domestic currency at the time at which they are incurred. For example, for a Chinese firm importing goods from the United States, the costs the firm incurs in the U.S. are immediately converted from U.S. dollars to Chinese yuan, and Chinese customers pay for the product in yuan. On the other hand, for a Chinese firm exporting goods to the U.S., many of the firm’s costs are incurred in China, and are denominated in terms of yuan. The revenues from customers and all costs incurred in the U.S. are immediately converted from U.S. dollars to Chinese yuan.

We first introduce the following notation that will be used in our model. We use boldfaced letters to represent vectors whose dimensions will be clear from the context.

- T = the length of the planning horizon,
- t = the period index, starting with period 1,
- L = replenishment lead time (assumed to be zero in the LS case),
- x_t = inventory position at the beginning of period t , before ordering,

$$\begin{aligned}
y_t &= \text{inventory position at the beginning of period } t, \text{ after ordering,} \\
D_t &= \text{demand in period } t, \\
D_{t,w} &= \sum_{j=t}^w D_j, \text{ total demand from period } t \text{ to } w, \\
\mathbf{D}_{t:w} &= (D_t, D_{t+1}, \dots, D_w), \text{ vector of demands from period } t \text{ to } w, \\
\mathbf{Z}_t &= \text{vector of financial information at the beginning of period } t, \\
\mathbf{z}_t &= \text{a realization of } \mathbf{Z}_t, \\
\mathbf{Z}_{t:w} &= (\mathbf{Z}_t, \dots, \mathbf{Z}_w), \text{ vector of financial information from period } t \text{ to } w.
\end{aligned}$$

The Financial Information Vector As we will discuss in more detail later, the financial information vector \mathbf{Z}_t may include exchange rates, stock market indices, interest rates, and other firm-specific information (*e.g.*, the change in a firm's stock price), all of which evolve over time. We will assume that $\{\mathbf{Z}_t, 1 \leq t \leq T + 1\}$ is a Markov chain, possibly with an infinite number of states, and is not necessarily time-homogeneous. In fact, we are particularly interested in cases that have identifiable time periods during which the financial instability is much higher than normal.

In this paper, t will usually refer to the current time period, so that \mathbf{z}_τ has already been observed for $\tau \leq t$. The distributions of $\mathbf{Z}_{t+1:T+1} = (\mathbf{Z}_{t+1}, \mathbf{Z}_{t+2}, \dots, \mathbf{Z}_{T+1})$ and $\mathbf{D}_{t:T+1} = (D_t, D_{t+1}, \dots, D_T)$ may be conditioned on the current, observed vector \mathbf{z}_t . We allow demand to be correlated with the financial information, as is clearly the case for many goods. However, we recognize that there are many consumer goods, such as light bulbs, whose demand can be assumed to be independent of the financial information.

Stochastic Discount Factors Since a discount factor is used to calculate the present value, at the beginning of period t , of a cash flow incurred at the beginning of period $t + 1$, it is natural to model it as a function of both \mathbf{z}_t and \mathbf{Z}_{t+1} and denote this factor as $\beta_t(\mathbf{z}_t, \mathbf{Z}_{t+1})$. Let $t \leq w$. Then the discount factor that converts cash flows at the end of period w to equivalent flows at the beginning of period t is a function of both \mathbf{z}_t and $\mathbf{Z}_{t+1:w+1}$, and is defined as $\beta_{t,w}(\mathbf{z}_t, \mathbf{Z}_{t+1:w+1})$ where

$$\beta_{t,w}(\mathbf{z}_t, \mathbf{Z}_{t+1:w+1}) = \prod_{j=t}^w \beta_j(\mathbf{Z}_j, \mathbf{Z}_{j+1}).$$

Stochastic Cost Parameters Since ordering costs, both fixed and variable, are incurred at the

beginning of a period, we model them as functions of the realized financial information. We assume that salvage is only possible at the end of the planning horizon, so that the salvage value is a random variable determined by the end-of-horizon financial information. Note that the fixed and variable purchasing cost may vary from period to period due to fluctuation in the exchange rate. Thus we define

$$\begin{aligned}
K_t(\mathbf{z}_t) &= \text{fixed ordering (and/or setup) cost at the beginning of period } t, \\
c_t(\mathbf{z}_t) &= \text{variable purchasing (and/or production) cost at the beginning of period } t, \\
v(\mathbf{Z}_{T+1}) &= \text{unit salvage value for on-hand inventory at the end of the planning horizon.}
\end{aligned}$$

In practice, holding costs, stockout costs and revenue for period t are probably incurred at various times throughout the time period. So for added flexibility, we model them as functions of both \mathbf{z}_t and \mathbf{Z}_{t+1} . Note also that even though the underlying selling price may be quite stable in the market in which a product is sold, the amount of revenue a firm is receiving in domestic currency can vary from period to period as the financial information (*e.g.*, exchange rate) fluctuates. Thus, we denote

$$\begin{aligned}
h_t(\mathbf{z}_t, \mathbf{Z}_{t+1}) &= \text{holding cost for carrying one unit of inventory from period } t \text{ to } t+1, \\
&\quad \text{valued at the end of period } t, \\
\pi_t(\mathbf{z}_t, \mathbf{Z}_{t+1}) &= \text{penalty cost (including the goodwill cost) for one unit of backorders} \\
&\quad \text{or lost sales at the end of period } t, \text{ valued at the end of period } t, \\
p_t(\mathbf{z}_t, \mathbf{Z}_{t+1}) &= \text{selling price in period } t, \text{ valued at the end of period } t.
\end{aligned}$$

To simplify the presentation, we will often make the dependencies on the financial information vectors implicit. That is, we will often write K_t , c_t , h_t , p_t , π_t , D_t , β_t , $\beta_{t,w}$ and v . All of these quantities are always non-negative, except possibly v . Note that $\beta_{t,w} = \beta_t \beta_{t+1,w}$. We also define $A^+ = \max\{0, A\}$, $A^- = (-A)^+$, $A \wedge B = \min\{A, B\}$, $A \vee B = \max\{A, B\}$, and

$$\delta(a) = \begin{cases} 1 & \text{if } a > 0, \\ 0 & \text{otherwise.} \end{cases}$$

2.2 Model Formulation

Let $G_t(x_t, y_t, \mathbf{z}_t, \mathbf{Z}_{t+1:t+L+1}, \mathbf{D}_{t:t+L})$ denote the cost incurred in period $t + L$, given the state (x_t, \mathbf{z}_t) , inventory decision in period t , y_t ($\geq x_t$), and the random evolution of financial and demand information between periods. Then, for $1 \leq t \leq T - L$,

$$\begin{aligned}
& G_t(x_t, y_t, \mathbf{z}_t, \mathbf{Z}_{t+1:t+L+1}, \mathbf{D}_{t:t+L}) \\
&= h_{t+L} \times (y_t - D_{t,t+L})^+ + \pi_{t+L} \times (y_t - D_{t,t+L})^- \\
&\quad - p_{t+L} \times \begin{cases} y_t \wedge D_{t,t+L}, & \text{if LS,} \\ (D_{t+L} + (y_{t-1} - D_{t-1,t+L-1})^- - (y_t - D_{t,t+L})^-), & \text{if PTD,} \\ D_{t+L}, & \text{if PTO,} \end{cases} \\
&= h_{t+L} \times (y_t - D_{t,t+L})^+ - p_{t+L} D_{t+L} \\
&\quad + \begin{cases} (\pi_{t+L} + p_{t+L})(y_t - D_{t,t+L})^-, & \text{if LS,} \\ ((\pi_{t+L} + p_{t+L})(y_t - D_{t,t+L})^- - p_{t+L}(y_{t-1} - D_{t-1,t+L-1})^-), & \text{if PTD,} \\ \pi_{t+L}(y_t - D_{t,t+L})^-, & \text{if PTO.} \end{cases}
\end{aligned}$$

In deriving these expressions, we applied the equality $y_t \wedge D_{t,t+L} = D_{t,t+L} - (y_t - D_{t,t+L})^-$, and made use of our assumption that $L = 0$ under lost sales. The number of units delivered to a client in period $t+L$ (*i.e.*, the term multiplied by p_{t+L} under PTD in the first expression above) merits some discussion. For case PTD, the potential sales in period $t + L$ is the demand in that period plus the backlog (if any) at the end of the previous period $t + L - 1$, $D_{t+L} + (y_{t-1} - D_{t-1,t+L-1})^-$. The actual sales will be the potential sales minus the ending backlog, $(y_t - D_{t,t+L})^-$. In other words, sales in period $t + L$ will equal demand minus the change in backlog, $D_{t+L} + (y_{t-1} - D_{t-1,t+L-1})^- - (y_t - D_{t,t+L})^-$.

Letting $\beta_{T+1} = 1$ and defining

$$\hat{\pi}_t = \pi_t + \begin{cases} p_t, & \text{if LS,} \\ p_t - \beta_{t+1} p_{t+1} & \text{if PTD,} \\ 0, & \text{if PTO,} \end{cases} \quad (1)$$

we can view the inventory problem as one with penalty costs $\hat{\pi}_t$ and a single period cost in period $t + L$, $1 \leq t \leq T - L$, of

$$\hat{G}_t(y_t, \mathbf{z}_t, \mathbf{Z}_{t+1:t+L+2}, \mathbf{D}_{t:t+L}) = h_{t+L} (y_t - D_{t,t+L})^+ + \hat{\pi}_{t+L} (y_t - D_{t,t+L})^- - p_{t+L} D_{t+L}. \quad (2)$$

Note that in the PTD case, the penalty cost $\hat{\pi}_t$ is a function of both $p_t(\mathbf{z}_t, \mathbf{Z}_{t+1})$ and $p_{t+1}(\mathbf{Z}_{t+1:t+2})$, and may be negative for some realizations of the financial information vectors. (We will define p_{T+1} and the end of horizon costs below.) The dynamic programming formulation for the expected discounted cost at the beginning of period t , $1 \leq t \leq T - L$, is given by

$$f_t(x_t, \mathbf{z}_t) = -c_t x_t + \min_{y_t \geq x_t} \{K_t \delta(y_t - x_t) + V_t(y_t, \mathbf{z}_t)\}, \quad (3)$$

$$\begin{aligned} V_t(y_t, \mathbf{z}_t) &= c_t y_t + E_{\mathbf{D}_{t:t+L}, \mathbf{Z}_{t+1:t+L+2}} [\beta_{t,t+L} \widehat{G}_t(y_t, \mathbf{z}_t, \mathbf{Z}_{t+1:t+L+2}, \mathbf{D}_{t:t+L}) \\ &\quad + \beta_t f_{t+1}(\gamma(y_t - D_t), \mathbf{Z}_{t+1})] \end{aligned} \quad (4)$$

where $\gamma(x) = x$ under backlogging, and $\gamma(x) = x^+$ under lost sales.

We now define the end-of-horizon cost that initiates the recursion in (3) and (4). We assume that all outstanding backorders are filled via a zero-lead-time purchase made at the end of period T at the unit cost of $c_{T-L+1} = c_{T-L+1}(\mathbf{Z}_{T+1})$, from the same suppliers or from some secondary market. (The time index of c_{T-L+1} is chosen for notational convenience.) In the PTD case, the end-of-horizon revenue is taken at the beginning of period $T + 1$ at $p_{T+1} = p_{T+1}(\mathbf{Z}_{T+1})$ per unit and is already included in \widehat{G}_{T-L} . Therefore,

$$\begin{aligned} f_{T-L+1}(x_{T-L+1}, \mathbf{z}_{T-L+1}) &= E_{D_{T-L+1,T}, \mathbf{Z}_{T-L+2:T+1}} \left\{ \beta_{T-L+1,T} \left[-v \times (x_{T-L+1} - D_{T-L+1,T})^+ \right. \right. \\ &\quad \left. \left. + \begin{cases} 0, & \text{if lost sales,} \\ c_{T-L+1} \times (x_{T-L+1} - D_{T-L+1,T})^-, & \text{if backorder} \end{cases} \right] \right\}. \end{aligned} \quad (5)$$

3. Structural Analysis

If all cost parameters and the discount factors are known in advance, it is well known that an (s, S) policy is optimal for the above problem, provided that a few side conditions are satisfied. For example, one requirement for the optimality of (s, S) policies is the monotonicity of the fixed ordering cost, $K_t \geq \beta_t K_{t+1}$. This condition is necessary for the K -convexity property to hold, and it may fail when financial markets are dynamic. Financial turbulence can cause a number of other requirements for the optimality of (s, S) policies to fail as well.

In this section, we will provide in Theorem 1 sufficient conditions for an (s, S) policy to be optimal and discuss the implications of these conditions. These conditions will help companies understand

how their decision-making should be adjusted when facing financial risks. We begin this section by discussing some important properties of the functions defined in the previous section, which impact the optimality of (s, S) policies.

3.1 Some Properties

We start with the finiteness and continuity of the functions \widehat{G}_t , f_t and V_t , and then discuss the asymptotic properties of functions with linear growth rates.

3.1.1 Finiteness and Continuity

So far we have allowed the functions defined in Section 2.2 to take on values in $\Re \cup \{\pm\infty\}$. We now consider the finiteness and continuity of the functions defined above. Let $u_t = E\{\beta_{t,t+L}[h_{t+L} \vee \beta_{t+L+1} p_{t+L+1} \vee (\pi_{t+L} + p_{t+L})]\}$ for $1 \leq t \leq T - L$, and let $U_{T-L+1} = E\{\beta_{T-L+1,T}[|v| \vee c_{T-L+1}]\}$. The following finiteness conditions are assumed to hold throughout this paper.

Assumption 1 *We will assume the following.*

- $E[\beta_{t,r} u_r]$, $E\{\beta_{t,w} D_w E[\beta_{w+1,r} u_r]\}$, $E[\beta_{t,T-L} U_{T-L+1}]$ and $E\{\beta_{t,w} D_w E[\beta_{w+1,T-L} U_{T-L+1}]\}$ are finite for all \mathbf{z}_t and all $1 \leq t \leq w + 1$, $1 \leq w \leq r \leq T - L$.
- *There exists a policy ω that has a finite expected cost. In other words, the following properties hold. Define f_t^ω and V_t^ω as the functions analogous to f_t and V_t in (3) and (4), corresponding to ω . (This means that f_t^ω and V_t^ω satisfy (4), and they satisfy (3) if we replace the minimum over y_t with the value of y_t that ω defines.) Then $f_t^\omega(x_t, \mathbf{z}_t)$ and $V_t^\omega(y_t, \mathbf{z}_t)$ are finite for all finite values x_t and y_t , for all \mathbf{z}_t .*

At first glance, it might appear that the first of these assumptions implies the second one. In fact, if the discount factors are less than or equal to 1 (*i.e.*, if $P(\beta_w \leq 1 | \mathbf{z}_t) = \mathbf{1}$ for all $w \geq t$ and for all (t, \mathbf{z}_t)), then the statement of the first assumption can be simplified, and it is easy to prove that the first assumption implies the second. In Section 5, we will consider one interesting application in which β_t is frequently greater than 1.

3.1.2 Asymptotic Properties of Functions with Linear Growth Rates

The functions we will deal with are all asymptotically linear. We now define some basic properties of functions whose asymptotic growth rates are linear.

Definition 1 *We say that*

1. $f(x)$ is $[a, \cdot]$ -divergent if $\lim_{x \rightarrow -\infty} [f(x) - ax] = \infty$;
2. $f(x)$ is $[\cdot, b]$ -divergent if $\lim_{x \rightarrow \infty} [f(x) - bx] = \infty$; and
3. $f(x)$ is $[a, b]$ -divergent if it is both $[a, \cdot]$ -divergent and $[\cdot, b]$ -divergent.

Definition 2 *We say that*

1. $f(x)$ is $[a, \cdot]$ -asymptotic if, for all $\epsilon > 0$, $f(x)$ is $[a + \epsilon, \cdot]$ -divergent and $-f(x)$ is $[-a + \epsilon, \cdot]$ -divergent;
2. $f(x)$ is $[\cdot, b]$ -asymptotic if, for all $\epsilon > 0$, $f(x)$ is $[\cdot, b - \epsilon]$ -divergent and $-f(x)$ is $[\cdot, -b - \epsilon]$ -divergent;
3. $f(x)$ is $[a, b]$ -asymptotic if it is $[a, \cdot]$ -asymptotic and $[\cdot, b]$ -asymptotic.

3.2 Sufficient Properties for the Optimality of (s, S) Policies

Let

$$S_t(\mathbf{z}_t) = \arg \min \{V_t(y_t, \mathbf{z}_t)\}, \quad (6)$$

$$s_t(\mathbf{z}_t) = \inf \{y_t : K_t(\mathbf{z}_t) + V_t(S_t(\mathbf{z}_t), \mathbf{z}_t) \geq V_t(y_t, \mathbf{z}_t)\}. \quad (7)$$

Since $S_t(\mathbf{z}_t)$ satisfies the inequality in (7), $s_t(\mathbf{z}_t)$ is well-defined and $S_t(\mathbf{z}_t) \geq s_t(\mathbf{z}_t)$. Also note that $s_t(\mathbf{z}_t) \geq 0$ in case LS. Technically, the domain of $V_t(\cdot, \mathbf{z}_t)$ is limited to real numbers, but we relax this limitation to allow both $S_t(\mathbf{z}_t)$ and $s_t(\mathbf{z}_t)$ to be infinite. In the event that $s_t(\mathbf{z}_t) = -\infty$, no order is placed in period t . In times of financial turmoil this might happen, and we view $(s_t(\mathbf{z}_t), S_t(\mathbf{z}_t))$ policies with $s_t(\mathbf{z}_t) = -\infty$ favorably. On the other hand, if $S_t(\mathbf{z}_t) = \infty$, then all orders placed are for infinite amounts of inventory and an infinite profit may be possible. We adopt the point of view that such a

problem is not well-posed. In this paper, the phrase *an (s, S) policy is optimal* implies that $S_t(\mathbf{z}_t) < \infty$ in (6). To simplify the presentation, we will often make the time index and the dependence on the financial information implicit, that is, we will often write (s, S) in stead of $(s_t(z_t), S_t(Z_t))$.

The following lemma provides the sufficient conditions for the optimality of the (s, S) policy.

Lemma 1 *Suppose that $V_t(\cdot, \mathbf{z}_t)$ is K_t -convex and $[\cdot, 0]$ -divergent. Then in period t with information set \mathbf{z}_t , an (s, S) policy is optimal. In cases PTD and PTO, if in addition $V_t(\cdot, \mathbf{z}_t)$ is $[0, \cdot]$ -divergent, the optimal (s, S) policies satisfy $s > -\infty$.*

Proof: See the Appendix.

The optimality of an (s, S) policy is thus reduced to the K_t -convexity and $[a, b]$ -divergence of the functions $V_t(\cdot, \mathbf{z}_t)$. We consider these two properties in turn, starting with $[a, b]$ -divergence.

3.2.1 $[a, b]$ -Divergence of the Functions $V_t(\cdot, \mathbf{z}_t)$

Let

$$\mathcal{A}_t = \mathcal{A}_t(\mathbf{z}_t) = \begin{cases} c_t + E_{\mathbf{Z}_{t+1:t+L+2}} \{-\beta_{t,t+L} \hat{\pi}_{t+L} + \beta_t [-c_{t+1} + \mathcal{A}_{t+1}^+(\mathbf{Z}_{t+1})]\}, & \text{if backorder,} \\ c_t - E_{\mathbf{Z}_{t+1:t+L+2}}(\beta_{t,t+L} \hat{\pi}_{t+L}), & \text{if lost sales,} \end{cases} \quad (8)$$

$$\mathcal{B}_t = \mathcal{B}_t(\mathbf{z}_t) = c_t + E_{\mathbf{Z}_{t+1:T+1}} \left[\sum_{j=t+L}^T \beta_{t,j} h_j - \beta_{t,T} v \right], \quad (9)$$

for $1 \leq t \leq T - L$, with $\mathcal{A}_{T-L+1} = 0$. We will see in Lemma 2 below that $V_t(\cdot, \mathbf{z}_t)$ is $[\mathcal{A}_t, \mathcal{B}_t]$ -asymptotic, so we hope that $\mathcal{A}_t < 0 < \mathcal{B}_t$. In the backorder cases (PTD and PTO), \mathcal{A}_t is the expected cost differential, as of the beginning of period t , between the option of ordering one unit in period t at c_t , and the option of carrying one unit of backorder in period $t + L$ at the modified backorder cost $\hat{\pi}_{t+L}$ and ordering one unit in period $t + 1$ at c_{t+1} , or later if that would be cheaper (*i.e.*, if $\mathcal{A}_{t+1} > 0$). In case LS for $t < T$, \mathcal{A}_t is irrelevant because the domain of $V_t(\cdot, \mathbf{z}_t)$ is \mathfrak{R}^+ . It is not difficult to see that \mathcal{B}_t is the expected cost of ordering one unit in period t , carrying it from period $t + L$ to the end of the horizon, and salvaging it. If $\mathcal{B}_t < 0$ for some t and \mathbf{z}_t then an infinite amount of inventory should be ordered, resulting in an expected cost of $-\infty$, and the problem is ill posed. Hence, we require that $\mathcal{B}_t > 0$ for all realizations \mathbf{z}_t of \mathbf{Z}_t , for $1 \leq t \leq T - L$.

Lemma 2 *$V_t(\cdot, \mathbf{z}_t)$ is $[\mathcal{A}_t, \mathcal{B}_t]$ -asymptotic in cases PTD and PTO, and $[\cdot, \mathcal{B}_t]$ -asymptotic in case LS.*

Proof: See the Appendix.

3.2.2 The K_t -Convexity of $V_t(\cdot, \mathbf{z}_t)$

We now consider the property required for the functions $V_t(\cdot, \mathbf{z}_t)$ to be K_t -convex.

Lemma 3 $V_t(y_t, \mathbf{z}_t)$ is K_t -convex if the following conditions hold. For all possible financial information vectors \mathbf{z}_{T-L} , $\mathcal{B}_{T-L} - \mathcal{A}_{T-L} \geq 0$. In addition, for all \mathbf{z}_j and all $j = t, \dots, T - L - 1$,

1. $h_{j+L} + \hat{\pi}_{j+L} \geq 0$ with backorders and $h_j + \hat{\pi}_j - c_{j+1} \geq 0$ with lost sales, and
2. $K_j(\mathbf{z}_j) \geq E_{\mathbf{Z}_{j+1}}[\beta_j(\mathbf{z}_j, \mathbf{Z}_{j+1})K_{j+1}(\mathbf{Z}_{j+1})]$.

Proof: See the Appendix.

3.3 Sufficient Conditions for (s, S) Policies to be Optimal

Combining Lemmas 1, 2 and 3, we can summarize the conditions for (s, S) policies to be optimal in the following theorem.

Theorem 1 An (s, S) policy is optimal in period t , if

1. $\mathcal{B}_j > 0$ for all realizations \mathbf{z}_j and all $j = t, \dots, T - L$,
2. $\mathcal{B}_{T-L} - \mathcal{A}_{T-L} \geq 0$ for all possible realizations \mathbf{z}_{T-L} , and
3. for all realizations \mathbf{z}_j and all $j = t, \dots, T - L - 1$,
 - (a) $h_{j+L} + \hat{\pi}_{j+L} \geq 0$ under backlogging and $h_j + \hat{\pi}_j - c_{j+1} \geq 0$ under lost sales; and
 - (b) $K_j(\mathbf{z}_j) \geq E_{\mathbf{Z}_{j+1}}[\beta_j(\mathbf{z}_j, \mathbf{Z}_{j+1})K_{j+1}(\mathbf{Z}_{j+1})]$.

If in addition $\mathcal{A}_t < 0$, then $s > -\infty$.

The conditions under which (s, S) policies will be optimal, as listed in Theorem 1, are related to speculative motives for inventory and backorders, with the exception of condition 3(b). We now discuss the implications of each condition in detail. When financial factors are stationary, all of the conditions are very common assumptions that hold for most real-world inventory systems. With potentially turbulent financial markets, we will see that Conditions 1 and 2 can safely be assumed to hold, while

Conditions 3a and 3b may fail. The un-labeled condition, $\mathcal{A}_t < 0$, is less important – rather than preventing an (s, S) policy from being optimal, it ensures that if the inventory level is low enough, a positive amount of inventory will be ordered. In the examples in Section 5 derived from Mexican presidential elections, Conditions 3a and 3b, and the un-labeled condition, all fail frequently.

Condition 1, $\mathcal{B}_t > 0$: This condition should hold if the end-of-horizon salvage values are set appropriately. If $\mathcal{B}_t \leq 0$, it would be optimal to order an infinite amount of inventory in period t , hold it, and salvage it at the end of the horizon. In that case, the problem is ill-posed, probably because the salvage value v was mis-specified.

We have a suggestion for an appropriate, stochastic value for v . In practice, the inventory system will almost certainly continue to operate after period T , and the salvage value v that minimizes the end-of-horizon effect should satisfy

$$v(\mathbf{z}_{1:T+1}) < \min \left\{ \left[c_t(\mathbf{z}_t) + \sum_{j=t+L}^T \beta_{t,j}(\mathbf{z}_{t:j+1}) h_j(\mathbf{z}_{j:j+1}) \right] / \beta_{t,T}(\mathbf{z}_{t:T+1}) : 1 \leq t \leq T - L \right\} \quad (10)$$

for all $\mathbf{z}_{1:T+1}$, viewed from the beginning of period $T + 1$. (10) ensures that $\mathcal{B}_t > 0$. Note that v will probably be a function of the entire historical financial information $(\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_{T+1})$ from the entire planning horizon, not merely a function of \mathbf{z}_{T+1} as we have formally assumed. However, we can assume without loss of generality that \mathbf{z}_{T+1} contains all of the information $(\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_{T+1})$, so all of our results continue to hold.

Condition 2, $\mathcal{B}_{T-L} - \mathcal{A}_{T-L} \geq 0$: This condition generally holds of its own accord. However, it can fail if the end-of-horizon parameters are not appropriate. In that case, the modeler should restrict the applicability of our model to appropriate end-of-horizon parameters that ensure this condition to hold. More specifically, since $L = 0$ in the lost sales case, $\mathcal{A}_T = c_T - E_{\mathbf{Z}_{T+1}}[\beta_T(p_T + \pi_T)] \leq 0$ if it is cheaper to purchase a unit in period T and deliver it to a client than it is to lose the sale. If p_T and π_T are observed at the beginning of period T , then \mathcal{A}_T is naturally negative and we have $\mathcal{B}_T - \mathcal{A}_T \geq 0$. However, in the real world, p_T and π_T relate to cash flows that are realized at various times within period T . Consequently, p_T and π_T may be dependent on \mathbf{Z}_{T+1} as well as \mathbf{z}_T , in which case $\mathcal{A}_T > 0$ can conceivably occur. In that case, $\mathcal{B}_{T-L} - \mathcal{A}_{T-L} \geq 0$ if it is appropriate to specify the salvage value

of v such that

$$v(\mathbf{z}_{T:T+1}) \leq h_T(\mathbf{z}_{T:T+1}) + \hat{\pi}_T(\mathbf{z}_{T:T+1})$$

and, in addition, (10) holds.

Under backlogging, $\mathcal{A}_{T-L} = c_{T-L} - E_{\mathbf{z}_{T-L+1:T+1}}[\beta_{T-L,T}(c_{T-L+1} + \hat{\pi}_T)] \leq 0$ is equivalent to saying that it is cheaper to purchase a unit in period $T - L$ and deliver it to a client in period T than it is to buy and deliver the unit at the end of period T at cost of $c_{T-L+1} + \pi_T$. (Recall that c_{T-L+1} is a unique, zero-lead-time purchase opportunity that is available only at the end of period T .) Like v , only certain values of c_{T-L+1} will ensure $E(\beta_{T-L,T} c_{T-L+1}) \geq c_{T-L} - E(\beta_{T-L,T} \hat{\pi}_T)$ and to reduce the end-of-horizon effect. We recommend defining c_{T-L+1} so that

$$c_{T-L+1}(\mathbf{z}_{T-L:T+1}) \geq c_{T-L}(\mathbf{z}_{T-L})/\beta_{T-L+1,T}(\mathbf{z}_{T-L:T+1}) - \hat{\pi}_T(\mathbf{z}_T, \mathbf{z}_{T+1}).$$

This potentially makes c_{T-L+1} a function of $\mathbf{z}_{T-L:T+1}$ rather than a function of \mathbf{z}_{T+1} as we originally assumed, but as we mentioned before, we can assume without loss of generality that the information in $\mathbf{z}_{T-L:T+1}$ is included in \mathbf{z}_{T+1} .

Condition 3a, $h_t + \hat{\pi}_t \geq 0$ (backlogging) and $h_t + \hat{\pi}_t - c_{t+1} \geq 0$ (lost sales): This condition deals with the situation in which we have a unit of demand in period t , and the inventory to meet that demand. However, under backlogging we choose to carry the inventory to period $t + 1$ and deliver it to the client then. In the lost sales case, we choose to lose the sale and use the unit of inventory to displace a unit that we would otherwise have purchased in period $t + 1$. The condition states that such decisions are not preferable. The condition always holds for PTO systems ($\hat{\pi}_t = \pi_t$), and is a very sensible assumption in most circumstances. However as we will see in Subsection 5.5.2, when there is strong potential for financial turbulence this condition can fail, especially for PTD systems.

Condition 3b, $K_t(\mathbf{z}_t) \geq E_{\mathbf{z}_{t+1}}[\beta_t(\mathbf{z}_t, \mathbf{z}_{t+1})K_{t+1}(\mathbf{z}_{t+1})]$: This condition clearly holds when there are no fixed costs; *i.e.*, when $K_t = 0$ for all t . When fixed costs are present, financial instability can cause it to fail, so that (s, S) policies may not be optimal.

Condition $\mathcal{A}_t < 0$: This condition is not needed for an (s, S) policy to be optimal. It guarantees that $s > -\infty$; *i.e.*, that in period t there is an inventory level so low that an order would be placed. Under lost sales, the comments in the first paragraph of our discussion of Condition 2 regarding $\mathcal{A}_T < 0$ apply. The condition generally holds, but depending on the assumptions made by the modeler, dynamic financial markets might cause it to fail.

Under backlogging, if $\mathcal{A}_t \geq 0$ then for some time period $\tau \geq t$, $\mathcal{A}_\tau \geq 0 \geq \mathcal{A}_{\tau+1}$. This implies that $c_\tau - E[\beta_{\tau, \tau+L} \pi_{\tau+L} + \beta_\tau c_{\tau+1}] \geq 0$. In other words, the expected cost of filling a demand in period $\tau + L$ by making a purchase in period $\tau + 1$ and delivering it to the client one period late, is smaller than the expected cost of making the purchase in period τ and delivering it to the client on time. Although $\mathcal{A}_t < 0$ for all t is a sensible and very common assumption, dynamic financial markets may cause it to fail.

4. Modeling the Stochastic Discount Factors β_t

In Sections 2 and 3, we discussed an inventory model that incorporates dynamic financial information, and provided the conditions under which the traditional inventory (s, S) policy is optimal. In this section, we introduce methods to explicitly model the stochastic discount factors β_t introduced in Section 2. For completeness, we start by presenting two basic models and then discuss conventional approaches for pricing financial securities.

Estimating the discount factor β_t is very challenging as it necessarily includes subjective beliefs about future cash flows. There are two different types of discount factors: a project-specific discount factor and a firm-wide discount factor. As a function of the risks that a particular project faces, the project-specific discount factor is generally calculated as a modification to the firm-wide discount factor. Since the risks associated with inventory are viewed as conventional risks, a firm-wide discount factor is used in practice. Hence, we focus on firm-wide discount factors and drop the word “firm-wide” for ease of presentation.

There are several simple ways to model discount factors. The simplest one, which has been used in many inventory models, is to assign it a constant value that applies to the entire time horizon under consideration (*e.g.*, $\beta_t(\mathbf{z}_t) = 0.95$ for all t and all \mathbf{z}_t). A generalization of this approach allows the rates to vary over time, although they are still assumed to be known in advance (*i.e.*, β_t is a

function of t , but not of \mathbf{z}_t). A third approach, rarely seen in the inventory literature, is to make the discount factor a deterministic function of the current financial information. In this approach, $\beta_t(\mathbf{z}_t)$ is known at the beginning of period t , and is a function of \mathbf{z}_t rather than of $(\mathbf{z}_t, \mathbf{Z}_{t+1})$ as we assume it to be. For example, one might simply let $\beta_t(\mathbf{z}_t) = 1/R_t(\mathbf{z}_t)$, where $R_t(\mathbf{z}_t)$ is the expected return of the firm in period t . As a second example, a firm conducting international business might let $\beta_t(\mathbf{z}_t) = \frac{1}{1+i_t^d} = \frac{e_t}{r_t^F(1+i_t^f)}$ where i_t^d is the domestic interest rate, e_t is the exchange rate, i_t^f the interest rate in the foreign country, and r_t^F the forward exchange rate for $t + 1$ given in period t . In that case, $\mathbf{z}_t = (i_t^d)$ or $\mathbf{z}_t = (e_t, i_t^f, r_t^F)$; see Ross *et al.* (2001). Although this state-dependent approach takes financial information into account, it does not fully capture the stochastic nature of the discount factors.

4.1 Pricing a Security

In the finance literature, the discount factor $\beta_t(\mathbf{z}_t, \mathbf{Z}_{t+1})$ for period t is usually interpreted as the reciprocal of one plus the required rate of return of the firm in period t , which is clearly random at the start of period t . Although the question of how to practically measure the intrinsic return of a firm is still an active research question, the standard academic approach is to estimate it as the relative price change of the firm's common stock over the period. Therefore, in this section, we review the most prominent developments in asset pricing.

The Capital Asset Pricing Model (CAPM) and Arbitrage Pricing Theory (APT) are two most popular models for security pricing. The CAPM specifies the asset's expected return to be an affine function of the overall market risk. APT, in contrast, assumes that a stock's return is a linear function of macroeconomic factors, some of which may be unique to that firm. So APT differs from CAPM in its assumptions and in its explanation of the risk factors associated with the risk of an asset. In our modeling of discount factors, we will identify risk factors based on empirical studies of asset prices, and adopt an APT type model.

In international markets, research results reveal that while exchange rate exposure does not, on average, appear to earn a risk premium in the U.S. stock market (Jorion 1990, 1991), it does earn a significant risk premium in other markets, especially in developing countries (Bailey and Chuang 1995). In the context of international asset pricing, market segmentation is another issue. For assets in a market that is fully integrated with the world market, national market risks might diminish when

world market risk is incorporated. This leads to the necessity of testing (statistically) whether the return of a national market portfolio is a significant risk factor that merits a risk premium. Political risk is yet another dimension of risk that needs to be considered (Dumas 2003). Generally, when pricing a security using the APT approach, one first identifies possible risk factors, and then performs statistical tests to determine which ones should remain in the model.

4.2 Modelling the Discount Factors

We propose three simple models for obtaining the discount factors, all of which are APT-type linear models. The first model relates a firm's return R_t to relevant risk factors $\mathbf{g}_t(\mathbf{z}_t, \mathbf{Z}_{t+1})$, functions of relevant financial information \mathbf{z}_t and \mathbf{Z}_{t+1} . The other two models relate the discount factor β_t to risk factors directly.

1. In the first model, we set

$$R_t = \mathbf{a}^T \mathbf{g}_t, \quad (11)$$

and $\beta_t = 1/(1 + R_t)$. Here $\mathbf{a}^T = (a_0, a_1, \dots, a_N)$ is a vector of constants and are estimated using historical data, and $\mathbf{g}_t = (1, g_{1t}, \dots, g_{Nt})^T$ are the risk factors.

2. Due to the intrinsic relationship between a firm's return and its discount factor, the discount factor should be affected by the same risk factors. Therefore, we can relate the discount factor and the financial risk factors directly as

$$\beta_t = \mathbf{b}^T \mathbf{g}_t, \quad (12)$$

where $\mathbf{b}^T = (b_0, b_1, \dots, b_N)$ is a vector of constants.

3. Instead of applying the APT model directly to the discount factors in an additive fashion, we might need to use a multiplicative model for the discount factors, by assuming a log linear relationship between the discount factors and the risk measures. That is,

$$\log(\beta_t) = \mathbf{c}^T \log(\mathbf{g}_t), \quad (13)$$

for some constant vector $\mathbf{c} = (c_0, c_1, \dots, c_N)^T$. Here $\log(\mathbf{g}_t) = (1, \log(g_{1t}), \dots, \log(g_{Nt}))$.

Suggested by Bailey and Chuang (1995), we propose the following five risk factors for Mexican firms that do business with U.S. firms:

$$\begin{aligned}
g_{1t}(\mathbf{z}_t, \mathbf{Z}_{t+1}) &= \frac{Z_{1,t+1}}{z_{1,t}} = \text{relative change of the exchange rate at period } t+1 \text{ over period } t, \\
g_{2t}(\mathbf{z}_t, \mathbf{Z}_{t+1}) &= \frac{Z_{2,t+1}}{z_{2,t}} = \text{one plus world market return for period } t, \\
g_{3t}(\mathbf{z}_t, \mathbf{Z}_{t+1}) &= \frac{Z_{3,t+1}}{z_{3,t}} = \text{one plus Mexican market return for period } t, \\
g_{4t}(\mathbf{z}_t) &= z_{4t} = \text{Mexican domestic interest rate at the beginning of period } t, \\
g_{5t}(\mathbf{z}_t) &= z_{5t} = \text{U.S. interest rate at the beginning of period } t,
\end{aligned}$$

and \mathbf{Z}_t is given by

$$\begin{aligned}
Z_{1t} &= \text{exchange rate at the beginning of period } t, \\
Z_{2t} &= \text{world market index at the beginning of period } t, \\
Z_{3t} &= \text{Mexican market index at the beginning of period } t, \\
Z_{4t} &= \text{Mexican domestic interest rate at the beginning of period } t, \\
Z_{5t} &= \text{U.S. interest rate at the beginning of period } t \text{ (foreign to Mexico)}.
\end{aligned}$$

Although Bailey and Chung (1995) also considered political risk factors, we consider only factors that reflect economic risks, because the political risks will be embedded in the economic factors. For example, in recent decades the political turbulence of the presidential election in Mexico actually changes the national economy (FOCAL 1999), which is reflected in the economic risk.

4.3 Modelling The Evolution of β_t

In order to model the evolution of the discount factor β_t over time, we need to model the evolution of the risk factors g_{it} over time. Standard approaches to time series modelling are applicable. The first step in modelling g_{it} is to test for autocorrelation. If there is no autocorrelation, the sequence $\{g_{it} : t = 1, 2, \dots\}$ can be modelled as observations from an iid random variable. In our experience, the sample autocorrelation functions generally have an exponentially decreasing appearance, so we apply

the simple AR(1) model. In the additive AR(1) model we update g_{it} using

$$g_{i,t+1} = \phi_0 + \phi_1 g_{it} + \xi_{it}, \quad (14)$$

and in the log-linear version of the AR(1) model we have

$$\log(g_{i,t+1}) = \alpha_0 + \alpha_1 \log(g_{it}) + \log(\eta_{it}). \quad (15)$$

Here, both the $\{\xi_{it} : t = 1, 2, \dots\}$ and $\{\log(\eta_{it}) : t = 1, 2, \dots\}$ are iid. random variables with mean 0. Note that when the g_{it} 's are iid, ϕ_1 (or α_1) is zero.

5. Numerical Case Study

There are a great many different types of instability affecting financial markets, as well as structures of a company's global supply chain, in addition to the range of cost structures, lead times and demand characteristics that differentiate one company from another. Thus a comprehensive numerical study is well beyond the scope of this paper. Instead in this section, we perform an illustrative case study. Our case study considers four different types of Mexican firms, differentiated by the (simple) architectures of their supply chains, which operate in the United States and/or Mexico. They are a *Local Firm*, an *Importer*, an *Exporter*, and a Mexican-owned *Subsidiary* that operates in the U.S.. Here, Mexico is referred to as the domestic country and the U.S. as the foreign country. These firms anticipate and respond to an observed fact: that Mexican financial markets are turbulent in the months surrounding their presidential elections. This fact motivated our research. We examine how the inventory decisions made by these firms change because of financial uncertainty during the months near an election.

The purpose of this numerical study is three-fold. First, we study how firms should adjust their inventory replenishment decisions in expecting the financial uncertainty during the months surrounding an election. Second, we examine what ordering policies are optimal when conditions discussed in Section 2 fail to hold. And third, we study the cost penalty of ignoring the dynamic nature of the financial environment.

We test two different inventory control policies. The first one is the *optimal policy*, computed via dynamic programming (DP). In the DP model, the state contains all of the currently-available

financial information, which has an impact on current and future discount factors and exchange rates (see Subsection 5.2 below). The second policy was motivated by conversations with Latin American businessmen. These people think very carefully about exchange rates, and do not hesitate to temporarily alter their inventory management policies if they anticipate a swing in exchange rates in the near future. However they seem to be unconscious of the fact that discount factors should impact their operational decisions in the short-to-medium term. This observation led us to the *constant-discount-factor policy*. This policy fully understands and reacts to the probabilistic model that governs future exchange rates, but it optimizes its actions under the inaccurate assumption that the discount factor will remain constant over the planning horizon. Algorithmically, this is a DP algorithm which uses the same state space as the optimal policy. It has two value functions, one of which determines the policy it will follow under the inaccurate, constant-discount-factor assumption, and one which determines the costs that will actually be incurred due to the use of the constant discount factor. Neither the constant-discount-factor policy nor the optimal policy is guaranteed to be an (s, S) policy in all cases. Therefore, the minimizations in the DP recursions are done by complete enumeration.

We begin this section by making some comments on the financial context for our case study. In Subsections 5.2 - 5.3 we discuss technical aspects of our numerical test environment, including the significance tests we performed on our risk factors, distributional assumptions regarding the financial variables, and operational costs. Subsection 5.4 summarizes our test environment. In Subsection 5.5 we discuss the degree to which the conditions for the optimality of (s, S) policies hold within our test environment (Mexican firms and presidential election periods). In Subsection 5.6 we optimize the operations of the supply chains, examine optimal inventory policies and costs when the impact of financial turbulence is properly accounted for (and when it is not), and explore general implications for global supply chain management.

5.1 The Financial Environment

The periods of economic turbulence that Mexico has gone through in the last few decades are reflected in large measure in the fluctuation of exchange rates. The monthly change in the exchange rate between the Mexican peso and the U.S. dollar from February 1976 to May 2004 is plotted in Figure 1. As one can see, there are significant spikes, representing sharp devaluations of the Mexican peso and economic crises. A majority, but not all, of these spikes occur during a time period starting four months before

the presidential election and ending two months after the presidential inauguration (March 1976 - Feb 1977, and every 6 years thereafter), called the *electoral season*. The months in an electoral season are called *election months*; all other months are called *regular months*. We define quarters to start in the months March, June, September and December, and talk about *election quarters* and *regular quarters*. As exchange rates change, the costs and prices of goods will move in desirable or undesirable

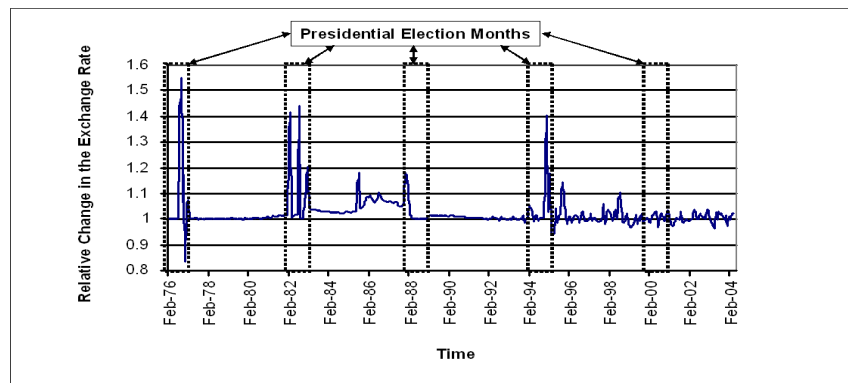


Figure 1: Relative Change in the Exchange Rate

directions, and it is important that firms anticipate these changes and plan ahead. For instance, a Mexican importer may want to order more than she otherwise would, if she expects a devaluation of the Mexican peso.

During times of financial turbulence in Mexico, banks tend to dramatically decrease the number of loans that they approve. Consequently, the effective cost of capital increases dramatically, much more strongly than the official domestic interest rates suggest. This probably explains why the domestic interest rate g_{At} is not statistically significant (see Subsection 5.2 below), even though the presence of discount factors greater than one during election quarters indicates that capital is very expensive (see Table 2).

5.2 The Risk Factors: Selection and Evolution

In this subsection we summarize our approach to the selection, and to modeling the evolution over time, of the risk factors. A more detailed treatment of this material is found in the Appendix. We start by determining the significance of the five risk factors proposed in Section 4.2, using the log linear model for the discount factor, $\log(\beta_t) = \mathbf{c}^T \log(\mathbf{g}_t)$ (see Model 3 in Subsection 4.2). We then create a statistical model of the evolution of the discount factors. To estimate the parameters \mathbf{c}^T of

the log linear model, we use historical data on the return of one Mexican firm, the DESC group, which has four divisions that do business in both Mexico and the U.S.: chemical, automotive, food, and real estate. Stepwise regression showed that the exchange rate risk (g_{1t}) and Mexican domestic market risk (g_{3t}) were significant at the 95% level, and that the constant factor $g_{0t} \equiv 1$, g_{2t} and g_{4t} were not significant. Consequently, we use g_{1t} and g_{3t} as the risk factors. Table 1 gives the coefficients \mathbf{c}^T for the quarterly discount factor that we obtained in this manner.

Table 1: Predicting the Quarterly Discount Factor β_t : Coefficients

		Change in the Exchange Rate (g_1)	Mexican Market Return Plus One (g_3)
Election Quarters	Coefficients	1.18	-0.477
Regular Quarters	Coefficients	0.531	-0.943

We now model the evolution over time of the risk factors g_{1t} and g_{3t} , following the approach of Subsection 4.3. Based on statistical tests of autocorrelation, we decided to model g_{3t} , and g_{1t} during election quarters, as iid random variables. During regular quarters, we model the change in the exchange rate g_{1t} using the standard logarithmic AR(1) autoregressive model $\log(g_{1,t+1}) = \alpha_0 + \alpha_1 \log(g_{1t}) + \log(\eta_{1t})$; see (15). We assume that the $\log(\eta_{1t})$'s are iid with mean 0. Using regression, we find that $\alpha_0 = 0.00609$ and $\alpha_1 = 0.675$, showing $\log(g_{1t})$ is a stable process.

To complete the model we need distributions for the following four iid random variables: η_{1t} and g_{3t} during the regular quarters, and g_{1t} and g_{3t} during election quarters. We achieve this by discretizing the logarithms of these random variables and creating histograms. This approach gives rise to Table 2 in which, for each quarter t , we give the mean and standard deviation of the stochastic discount factor $\beta_t = \beta_t(\mathbf{Z}_t, \mathbf{Z}_{t+1})$ (the ordered pair $(\mathbf{Z}_t, \mathbf{Z}_{t+1})$ is random). Here quarters 1, 2, and 7 are regular quarters, and quarter 3 to 6 are election quarters. Even in the regular quarters, there is a positive probability that the discount factor is greater than one.

Table 2: Mean and Standard Deviation of the Discount Factor Over Time

Quarter	1	2	3	4	5	6	7
Mean	0.971	0.973	1.203	1.203	1.203	1.203	0.905
Standard Deviation	0.285	0.287	0.286	0.286	0.286	0.286	0.292

5.3 The Operational Costs

In reality, operational costs and revenues can be affected by many factors and may be difficult to model. For our case study we will assume that the costs and selling prices have known values which, over the planning horizon, are constant in the markets in which they occur. The costs incurred in Mexico are denoted as (K, c, h, π, p) and those in the U.S. are referred to as (K', c', h', π', p') . We assume that in each quarter the firm converts all costs and revenues into Mexican pesos immediately. Consequently the values of the dollar-denominated costs and prices are proportional to the exchange rate z_{1t} in a given quarter t . Table 3 lists the cost parameters used under the different scenarios. Note that by assumption, the fixed and variable ordering costs are incurred in the market where a product is made, but the holding costs, stockout costs and revenues are realized in the market where the product is sold.

Table 3: Cost Parameters for the Four Firms

	K_t	c_t	h_t	π_t	p_t
Local Firm	K	c	h	π	p
Importer	$K'z_{1t}$	$c'z_{1t}$	h	π	p
Exporter	K	c	$h'Z_{1,t+1}$	$\pi'Z_{1,t+1}$	$p'Z_{1,t+1}$
Subsidiary	$K'z_{1t}$	$c'z_{1t}$	$h'Z_{1,t+1}$	$\pi'Z_{1,t+1}$	$p'Z_{1,t+1}$

5.4 Summary of the Test Environment for the Case Study

In this case study we consider four hypothetical Mexican firms, differentiated by their supply chain architecture. The Local Firm operates entirely in Mexico, the Importer buys or manufactures in the U.S. and sells in Mexico, the Exporter buys or manufactures in Mexico and sells in the U.S., and the Subsidiary is Mexican-owned and operates in the U.S.

In our numerical test environment we use the risk factors g_{1t} (relative change of the exchange rate at quarter $t+1$ over quarter t) and g_{3t} (one plus Mexican market return for quarter t). The risk factor g_{1t} during the regular quarters evolves according to the logarithmic autoregressive AR(1) model of (15), using the coefficients and the distribution for the noise factor $\log(\eta_{1t})$ described in Subsection 5.2. The risk factor g_{3t} during regular quarters, and the risk factors g_{1t} and g_{3t} during election quarters, follow the iid distributions described in Subsection 5.2.

In each quarter the risk factors g_{1t} and g_{3t} determine the exchange rate directly, and they govern the discount factors *via* (13), with the coefficients that were obtained in Subsection 5.2. The exchange rate impacts the cost parameters (see Table 3). The fixed costs K and K' are set to zero. The other cost parameters are $c = 2$, $h = 0.1$, $\pi = 0.8$, $p = 6$, $c' = 0.2$, $h' = 0.01$, $\pi' = 0.08$ and $p' = 0.6$.

Our case study considers a 7-quarter (*i.e.*, 7 time period) problem with zero lead time ($L = 0$), zero starting inventory ($x_1 = 0$), and a discrete uniform demand distribution over $[0, 5]$ in each quarter. The demands are stochastically independent of the financial variables. The electoral season consists of quarters 3 through 6, and election quarters have high financial instability. Quarters 1, 2 and 7 are regular quarters that have lower financial volatility. At the end of quarter 7, any positive inventory will be sold back to the supplier at the salvage value $v(\mathbf{Z}_8)$. If there is negative inventory, the firm is required to fill all backordered demand by a purchase at the unit price $c_{T-L+1}(\mathbf{Z}_8)$, and receives revenue in the PTD case. The salvage value $v(\mathbf{Z}_8)$ and the purchase price $c_{T-L+1}(\mathbf{Z}_8)$ at the end of quarter 8 are determined as specified in Section 3. We test two inventory control policies, the optimal policy and the constant-discount-factor policy, which we described in the third paragraph of Section 5.

5.5 Test of the Conditions in Theorem 1

In this subsection we test the optimality conditions listed in Theorem 1. These conditions determine whether an (s, S) policy is optimal, and whether or not $s > -\infty$. We characterize the cost functions V_t when they are not convex. Several of the observations that we will make have interesting implications. We test the conditions in this subsection, and we discuss the wider implications of this data in Subsection 5.6.

To briefly summarize the conclusions of this subsection, in our case study, we find that if the fixed cost is positive then the value function is seldom K -convex, so an (s, S) policy is seldom to be optimal. If the fixed cost is $K = 0$, a convex value function means that an order-up-to- S policy is optimal. This always happens in the PTO case, and it generally happens in the LS case (except for the Importer after the electoral season). In the PTD case, convexity will almost certainly fail in at least one time period, except for the Importer. When convexity does fail, the value functions have a specific structure, which has interesting operational implications.

Theorem 1 has four conditions. Since we will limit the applicability of our model to the cost

parameters $v(\mathbf{Z}_{T+1})$ and $c_{T-L+1}(\mathbf{Z}_{T+1})$ that ensure the value function V_t is convex in the last quarter, we only test V_t for K -convexity in the first 6 quarters. Also, we limit attention to problems with $\mathcal{B}_t > 0$, which holds if the salvage value $v(\mathbf{Z}_{T+1})$ is an appropriate one. These leave us with three conditions to test: $K_t \geq E(\beta_t K_{t+1})$, $h_t + \hat{\pi}_t \geq 0$ under backlogging and $h_t + \hat{\pi}_t - c_{t+1} \geq 0$ under LS and $\mathcal{A}_t < 0$.

5.5.1 The condition $K_t \geq E(\beta_t K_{t+1})$.

This condition is essential for K -convexity. In the financial environment we study, as one might expect, this condition can easily be violated. In Table 4, because of the randomness that exists, $E(\beta_t K_{t+1})$ is clearly effected by the state at the start of quarter t . Two other factors contribute to large percentages of failure observed. First, during the electoral season the mean discount factor is greater than one (see Table 2). Second, during regular quarters the peso-to-dollar exchange rate is stochastic, and it almost never decreases. Hence, the constant dollar-denominated fixed cost that the Importer and the Subsidiary face have an expectation that increases as a function of time, when converted to Mexican pesos.

Because this condition is violated, it is likely that the structure of an optimal policy is complex in form. Further research is needed on the structure and computation of optimal policies when financial instability is present, and on the impact of financial instability on the performance of popular heuristics. For our case study we will assume that the fixed ordering cost is zero, so this condition holds.

Table 4: Percentage of states in which $K_t \geq E(\beta_t K_{t+1})$ fails

Quarters t	Local Firm	Importer	Exporter	Subsidiary
1	0	100%	0	100%
2	37.5%	62.5%	37.5%	62.5%
3-6	100%	100%	100%	100%

5.5.2 The condition $h_t + \hat{\pi}_t \geq 0$ under backorders, and $h_t + \hat{\pi}_t - c_{t+1} \geq 0$ under lost sales.

This condition and the previous one are sufficient for the value functions V_t functions to be K -convex and hence, for an (s, S) policy to be optimal. In the absence of fixed costs, this condition is sufficient for an order-up-to policy to be optimal. The condition always holds for the PTO (pay to order) case.

In the PTD (pay to delivery) case, the condition $h_t + \hat{\pi}_t \geq 0$ always holds in quarters 1 and 7, but it never holds in quarters 2-6. In other words, the condition fails when any of the parameters in $\hat{\pi}_t = \pi_t + p_t - \beta_{t+1}p_{t+1}$ has an index in an election quarter. To get a better understanding of what is happening, for the PTD case all of the functions V_t that we observed fall into one of three different categories: (1) Monotonically increasing functions (called *Increasing*, for which no order is placed; *i.e.*, an order-up-to- S policy where $S = -\infty$ is optimal), (2) Functions that are convex but not monotone (called *Convex*, for which an order-up-to- S policy with finite S is optimal), and (3) Functions that first increase, then decrease, and then increase again (called *Up-down-up*, for which no order-up-to- S policy is optimal).

Table 5 shows the percentage of states in each quarter when the cost function is Increasing, Convex and Up-down-up for the four different firms, under PTD. For the Up-down-up functions, we also give the local optimal order-up-to level. The implications of these observations are discussed in Subsection 5.6.2.

Table 5: Percentage of the States Having Different Types of Cost Functions under PTD
“Local Minimum” is the inventory level that locally minimizes the Up-down-up functions

Type of Firm	Cost Functions	Quarters t						
		1	2	3	4	5	6	7
Local Firm	Increasing	0	87.5%	0	0	0	0	0
	Convex	100%	0	0	0	0	100%	100%
	Up-down-up	0	12.5%	100%	100%	100%	0	0
	Local minimum	-	13	12	9	7	-	-
Importer	Increasing	0	50%	0	0	0	0	0
	Convex	100%	37.5%	100%	99.75%	86.77%	100%	100%
	Up-down-up	0	12.5%	0	0.25%	13.23%	0	0
	Local minimum	-	16	-	11	8	-	-
Exporter	Increasing	100%	100%	100%	95.38%	25.77%	0	0
	Convex	0	0	0	0	0	100%	100%
	Up-down-up	0	0	0	4.62%	74.23%	0	0
	Local minimum	-	-	-	9	6	-	-
Subsidiary	Increasing	100%	100%	0	0	0	0	0
	Convex	0	0	0	0	0	100%	100%
	Up-down-up	0	0	100%	100%	100%	0	0
	Local minimum	-	-	14	11	8	-	-

Table 6 shows the percentage of states in each quarter when the condition $h_t + \hat{\pi}_t - c_{t+1} \geq 0$ fails to hold for the LS (lost sales) case. In our computational experiments, under LS, whenever the cost

functions fail to be convex they are still quasi-convex; hence an order-up-to policy remains optimal. These experiments are admittedly limited, but as far as they go, they indicate that if $K_t \equiv 0$ then order-up-to- S policies are likely to be optimal in many, if not most, lost sales settings.

Table 6: Percentage of States in Which $h_t + \hat{\pi}_t - c_{t+1} \geq 0$ Fails for the LS Case

Quarters t	Local Firm	Importer	Exporter	Subsidiary
1-3	0	0	0	0
4	0	0.2%	0	0
5	0	9.69%	0	0
6	0	22.43%	4.28%	0

5.5.3 The condition $\mathcal{A}_t < 0$.

Recall that this condition is not required for an (s, S) policy to be optimal. In the backorder cases it ensures that when an (s, S) policy is optimal, an (s, S) policy with $s > -\infty$ is optimal. For the backlogging cases, we chose $c_{T-L+1}(\mathbf{Z}_{T+1})$ so that the condition holds for the final quarter; we need only test the condition for quarters 1 through 6. In the PTO case the condition holds for all firms and all states. For the PTD case the condition holds when Table 5 above indicates that the value function is Convex, and it fails when the value function is either Increasing or Up-down-up.

In Case LS the condition $\mathcal{A}_t < 0$ lacks importance. The corresponding phenomenon is an order-up-to level $S = 0$. Table 7 gives the percentage of states in which this occurs.

Table 7: Percentage of States Having Order-up-to Levels of Zero for the LS Case

Quarter	1	2	3	4	5	6	7
Local Firm and Subsidiary	0	0	0	0	0	0	0
Importer	0	0	0	0	1.66%	15.07%	32.99%
Exporter	0	0	0	0	0	2.14%	25.90%

5.6 Implications of Stochastic Discount Factors

In this subsection we explore the implications of using sub-optimal inventory control policies in environments where financial turbulence is likely to occur. We do this by first comparing the cost of the optimal policy and the cost of the constant-discount-rate policy (see paragraph 3 of Section 5). Then

we look at the inventory stocking decisions made by the optimal policy. We consider four illustrative firms, under our three different stockout protocols.

Figure 2 shows the cost penalty of using a constant discount factor, which is the ratio of the cost of the constant-discount-ratio policy to the optimal cost, minus 1. That penalty can be as high as 78%. The cost penalty for all of the firms under the three different stockout protocols is very high. It is especially high for the Exporter and the Subsidiary in the PTD case. This is apparently because, under the PTD stockout protocol, we can control the timing of cash inflows, which has a strong impact on profitability. That opportunity does not exist under the PTO and LS protocols. Under PTD, this opportunity brings less benefit to the Local Firm and the Importer than it does to the Exporter and the Subsidiary, for which the revenue is in U.S. dollars, whose value is likely to increase during the election quarters. It is interesting that the Local Firm pays a heavy cost penalty for using deterministic discount factors under both the LS and PTO protocols – a penalty that is entirely caused by the discount factor, because the Local Firm is not directly impacted by the exchange rate.

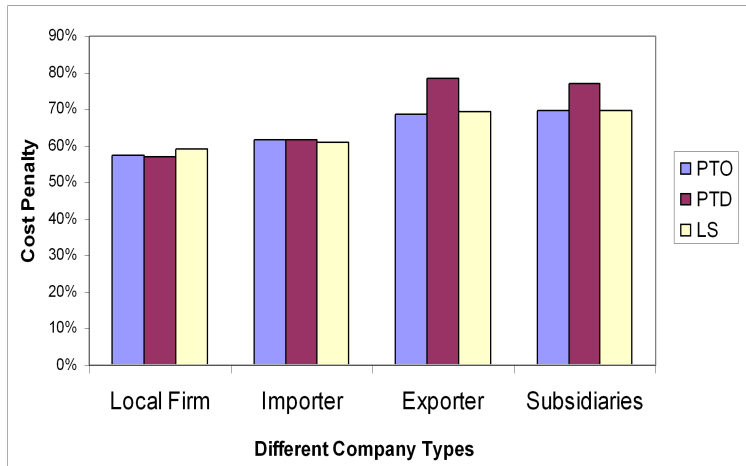


Figure 2: Cost Penalty for Using a Constant Discount Rate

5.6.1 The Pay-to-Order (PTO) Stockout Protocol

We now consider the ordering behavior of the four firms under the three stockout protocols. We begin with the PTO stockout protocol. Recall from Section 5.5.3 that an order-up-to policy with finite S is optimal in all states. Recall that the maximum demand in a quarter is five and that the lead time is zero. Thus the speculative motives would be the only reason why the optimal echelon inventory would

ever exceed five. Also note that under PTO the revenues are not controllable. Therefore the impact of stochastic discount factors on supply chain management decisions arises from the costs (measured in pesos), not from the revenues.

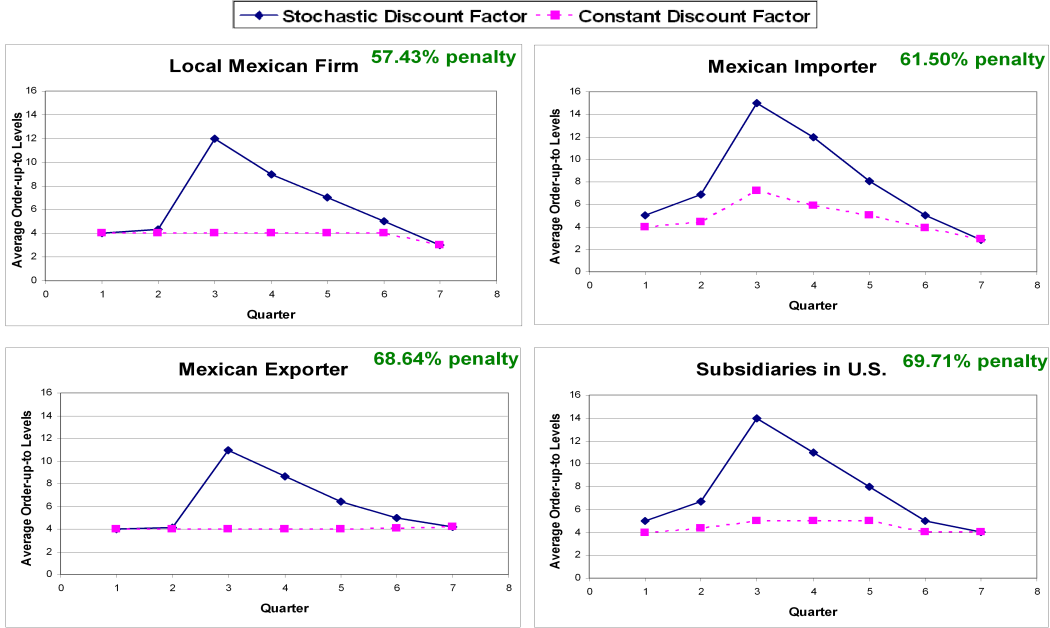


Figure 3: Average Order-up-to Levels under PTO

Figure 3 shows the average of the optimal order-up-to levels (over all the financial states in each period) for the four firms under PTO. All firms have optimal order-up-to levels that are much higher than those obtained using a constant discount factor. As the electoral season draws to a close, the factors that cause these discrepancies become less likely to occur, and the optimal order-up-to levels descend towards the levels that one would expect to see under a constant discount factor. These swings in the order-up-to levels are the result of three different factors, listed in order of their apparent importance.

- Large Discount Factors Impact Acquisition Costs** for all four firms. During the electoral season there is a strong possibility that the cost of capital will rise strongly. That makes a peso more valuable later on than it is at the beginning of the electoral season. Therefore, all of the firms order a lot of inventory at the beginning of quarter 3. This is the only factor that impacts the Local Firm. Figure 3 indicates that it causes an increase of approximately 8 (from 4 to 12) in the order-up-to level in quarter 3.

- **Devaluation Impacts Acquisition Costs.** There is a strong possibility of a devaluation of the peso during the electoral season, making production in the U.S. more expensive. For the Importer and the Subsidiary, this creates another incentive to stock up on inventory before the election season begins. Figure 3 seems to indicate that this factor increases the order-up-to levels of the Importer and the Subsidiary in quarter 3 by approximately 3 units (comparing the right column of Figure 3 to the left column).
- **Devaluation Impacts Holding Costs.** A devaluation of the peso would increase the cost of holding inventory in the U.S. This gives the Exporter and the Subsidiary an incentive to decrease their inventory levels. Figure 3 indicates that this factor decreases the order-up-to levels of the Exporter and the Subsidiary in quarter 3 by approximately 1 unit (comparing the top row of Figure 3 to the bottom row).

5.6.2 The Pay-to-Delivery (PTD) Stockout Protocol

Figure 4 shows the average order-up-to levels for the four different firms under PTD. The PTD case is more complex because ordering policies impact the timing of the cash inflows in addition to the operating costs. Table 5 indicates that under PTD the value function is often Up-down-up, meaning that it has two local minima; we compute the optimal policy by enumeration. The averages of the resulting optimal order-up-to levels are found in Figure 4.

The ordering decisions under PTD are quite different from those under PTO. For all firms except the Importer, the dominant goal is to delay income until the last quarter of the electoral season, in order to maximize the value of the revenue obtained. Discount factors make this goal attractive for all firms; exchange rates increase its attractiveness for the Exporter and the Subsidiary. The importance of deferring revenue dominates other considerations, such as the attractiveness of buying inventory when it is cheap (before the electoral season begins), and the cost of carrying the inventory.

In this context, Table 5 is interesting. It shows that for the Local Firm and the Subsidiary, in most quarters that are part of the electoral season, and in most states, the value function is Up-down-up. That means that there is a threshold – an inventory level at which the value function is increasing and is equal to the local minimum. If the inventory level is above the threshold then an order-up-to- S policy is optimal. Furthermore, Table 5 indicates that the order-up-to level (*i.e.*, the Local Minimum)

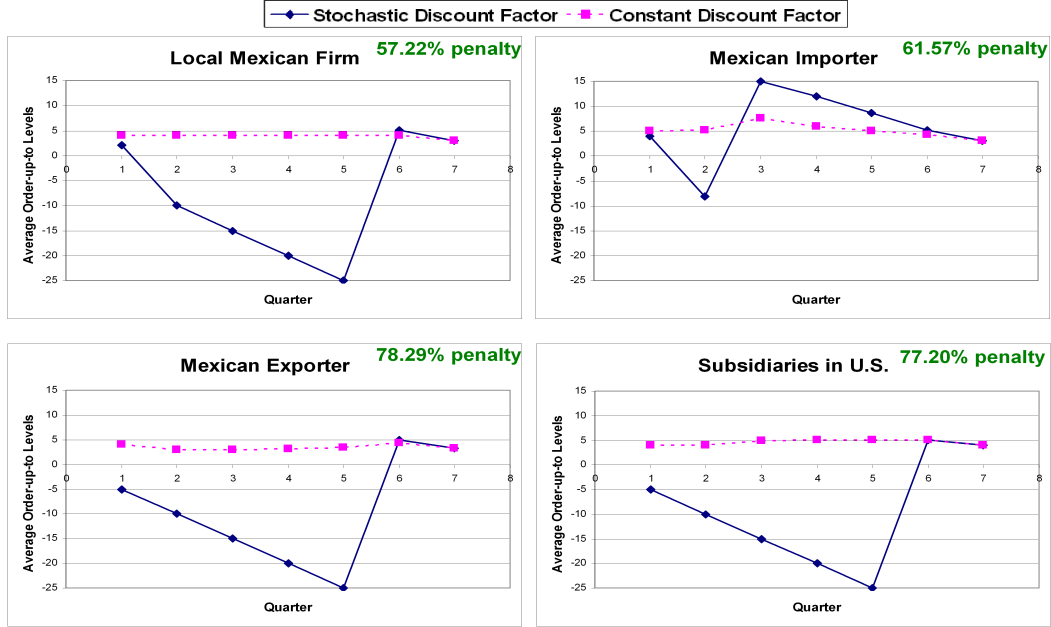


Figure 4: Average Order-up-to Levels under PTD

is chosen based on the same logic that applies in the PTO case – buy the inventory soon, before the discount factor causes the effective acquisition cost to rise. On the other hand, if the net inventory level is below the threshold then there are existing backorders or there is the potential to accumulate backorders, and in the optimal policy we do not order anything until quarter 6, in order to maximize the value of the revenue obtained from these backorders. For the Local Firm and the Subsidiary, comparing the local minima from Table 5 and the average order-up-to levels of Figure 4, it is clear that in periods 1 and 2 the optimal policy manages the inventory system so that during the electoral season the net inventory level will almost always be below the threshold, and both purchases and revenues will be deferred. (For the Exporter, Up-down-up value functions are less prevalent because the exchange rates create an added incentive to delay the revenue.)

An Importer operating with the PTD stockout protocol uses a hybrid strategy. She is influenced by the factor that dominated optimal policies in the PTO case (buy inventory soon, before the discount factor and exchange rate cause acquisition costs to rise), and the factor that determined behavior for the other three firms in the PTD case (delay income, to maximize its value). However the economics of buying in the U.S. and selling in Mexico result in a very different balance between these factors. In quarter 2 the importance of delaying income dominates, and no purchases are made. But when the electoral season begins (the beginning of quarter 3) the balance shifts, the combined impact of the

discount factor and the exchange rate on acquisition costs starts to dominate, and the Importer buys large amounts of inventory while it is still relatively inexpensive. In Table 5 we see that the value functions are almost all convex, an indication of the fact that this strategy is consistently followed in nearly all states and time periods.

5.6.3 The Lost Sales (LS) Stockout Protocol

Figure 5 shows the order-up-to levels for the four firms under lost sales. Like the PTO protocol, under lost sales the firm cannot use supply chain operations to shift revenue in time. The results are similar to the PTO case, both in behavior, and in the reasons for that behavior. There are two moderately interesting differences. First, since unsatisfied demand is lost, firms have a stronger incentive to stock inventory. This is indicated by the fact that the constant-discount model usually stocks to a level of 5, rather than 4 in the PTO case. Secondly, during the electoral season devaluations may have rendered the Importer’s business unprofitable, by increasing the cost of acquiring inventory in the U.S. Consequently, in quarter 7, in 33% of the states the Importer has an order-up-to level of zero (see Table 7). In the backorder cases the Importer must meet all of the demand, so this does not happen.

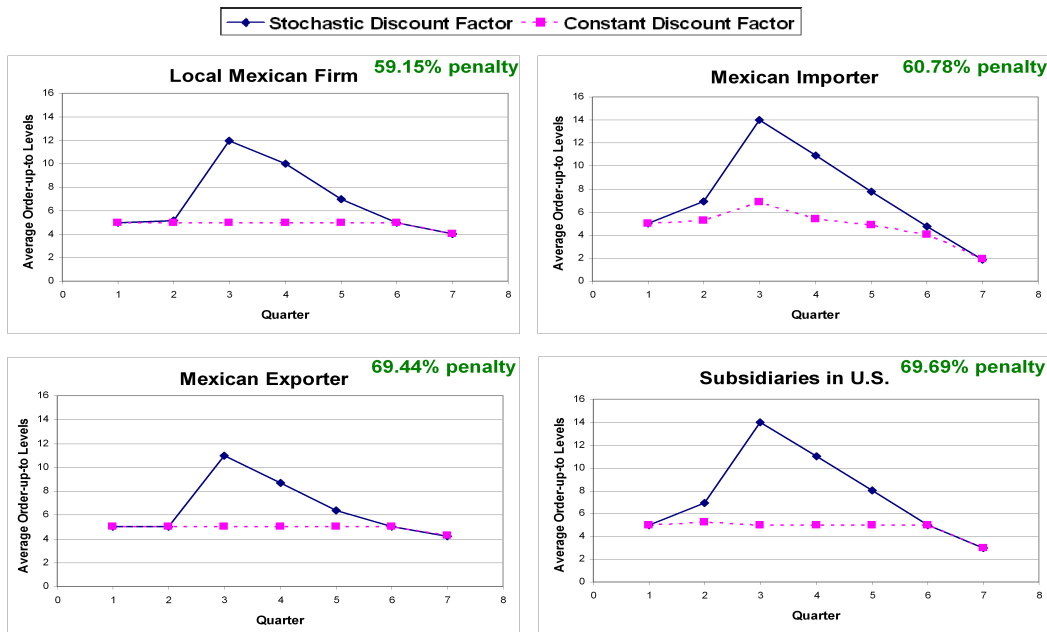


Figure 5: Average Order-up-to Levels under LS

6. Conclusion

In this paper we challenge the use of constant discount factors in traditional inventory models. On the theoretical level, we prove that under certain circumstances, a state dependent (s, S) ordering policy will be optimal. In stochastic financial markets, when an (s, S) policy is optimal, neither s nor S is necessarily finite. We also provide the conditions under which both s and S will be finite.

In environments that are affected by cyclical periods of economic uncertainty and economic stability, we propose a data-driven approach for modeling the distribution, and the evolution over time, of stochastic discount factors. Our approach is based on the arbitrage pricing theory, and the following concept, which is common in the finance literature: the (firm-wide) discount factor should be the reciprocal of one plus the rate of return of the firm.

We present a case study based on four Mexican firms, which are affected by the financial instability that often surrounds presidential elections in Mexico. During the electoral season the discount factor has a mean that is greater than one, meaning that the internal rate of return has a negative mean. This is because in Mexico, capital often becomes very expensive or temporarily unavailable during periods of financial turbulence; hence, a peso in the future is worth more than a peso is worth today. The four types of firms all reacted to the potential for financial turbulence very strongly, and in very different ways. We observed the stockpiling of large amounts of inventory, and the use of supply chain operations to defer revenue for extended periods of time. The primary drivers of this behavior are the potential for currency devaluations, and the discount factors.

In traditional inventory models, when unfilled demand is backordered, it does not matter really whether the client pays when the order is placed or when the inventory is delivered. Any difference can be compensated for by changing the backorder cost. However when discount factors and other financial information are stochastic and non-stationary, it can make a substantial difference. Suppliers that are paid when the inventory is delivered can use supply chain operations to change the timing of their cash inflows – at times, with great effect. From the buyer’s perspective, a PTO (pay-to-order) protocol creates a much more desirable set of supplier incentives than a pay-to-delivery (PTD) discipline does.

A comprehensive study (or set of studies) of supply chain operations in times of predictable financial turbulence would be very valuable. In addition to a much more comprehensive set of scenarios, there are other very important considerations that remain to be studied. For example, we have assumed that

under the PTD protocol, demand is met whenever inventory is available. Companies that are more concerned with profit than service might make a different decision. We also assumed that all costs and revenues have constant values in the local currency, and that they are converted to the firm's domestic currency at the time they are incurred (see Subsection 5.3). Once again, different assumptions could be modeled. Finally, when there are cyclic time intervals that oscillate between periods of high and low financial risk (like the Mexican presidential elections) we have shown how to use historical data to obtain an evolutionary model of discount factors. However for one-time events (such as the Iraq War), the question of how to model the manner in which discount factors and other financial variables might change is of course much harder.

This research is especially useful for global firms facing predictable financial instabilities in the near future. By using a stochastic discount factor rather than a constant discount factor, such a firm could potentially save a lot of money.

References

- Adler, M. and B. Dumas. 1983. International Portfolio Choice and Corporation Finance: A Synthesis. *The Journal of Finance*, 38(3), 925-984.
- Anvari, M. 1987. Optimal Criteria and Risk in Inventory Models: The Case of the Newsboy Problem. *The Journal of Operational Research Society*, 38(7), 625-632.
- Axaroglou, K. and P. Kouvelis. 2001. Exchange Rates and Ownership Structure of Foreign Subsidiaries: An Empirical Investigation. Working paper.
- Bailey, W. and Y. Chuang. 1995. Exchange Rate Fluctuations, Political Risk, and Stock Returns: Some Evidence from an Emerging Market. *Journal of Financial and Quantitative Analysis*, 30(4), 541-561.
- Beranek, W. 1967. Financial Implications of Lot-size Inventory Models. *Management Science*, 13(8), 401-408.
- Canadian Foundation for the Americas (FOCAL). 1999. "Mexico's Presidential Election". *FOCAL Policy Paper (FPP-99-04)*, ISBN: 1-896301-30-4
- Cochrane, J. H. 2001. *Asset Pricing*. Princeton University Press.

- Ding, Q., L. Dong and P. Kouvelis. 2004. On the Integration of Production and Financial Hedging Decisions in Global Market. Working paper.
- Dumas, B. 2003. Which “going price” is right? *INSEAD Quarterly*, 6, 19-21.
- Errunza, V. and E. Losq. 1985. International Asset Pricing under Mild Segmentation: Theory and Test. *The Journal of Finance*, 40(1), 105-124.
- Gaur, V., S. Seshadri. 2005. Hedging Inventory Risk through Market Instruments. *Manufacturing & Service Operations Management*, 7(2), 103-120.
- Grubbstrom, R W. 1980. A Principle for Determining the Correct Capital Costs of Work-in-progress and Inventory. *International Journal of Production Research*, 18(2), 259-271. *The Journal of Finance*, 52(2), 557-590.
- Jorion, P. 1990. The Exchange Rate Exposure of U.S. Multinationals. *Journal of Business*, 63(3), 331-345.
- Jorion, P. 1991. The Pricing of Exchange Rate Risk in the Stock Market. *Journal of Financial and Quantitative Analysis*, 26(3), 363-376.
- Kazaz, B., M. Dada and H. Moskowitz. 2005. Global Production Planning Under Exchange-Rate Uncertainty. *Management Science*, 51(7), 1101-1119
- Kim, Y. H. and K. H. Chung. 1989. Inventory Management under Uncertainty: A Financial Theory for the Transactions Motive. *Managerial and Decision Economics*, 10(4), 291-298.
- Ross, S. A., R. W. Westerfield and J. F. Jaffe. 2001. *Corporate Finance*, McGraw-Hill Irwin.
- Scarf, H. E. 1959. The Optimality of (S,s) Policies in the Dynamic Inventory Problem. Chapter 13, in *Mathematical Methods in the Social Sciences*, ed. by K.J. Arrow, S. Karlin and P. Suppes. Stanford University Press.
- Singhal, V. R., A. S. Raturi and J. Bryant. 1994. On Incorporating Business Risk into Continuous Review Inventory Models. *European Journal Of Operational Research*, 75(1), 136-150.
- Stulz, R. M. 1994. International Portfolio Choice and Asset Pricing: An Integrative Survey. NBER Working Paper No. 4645.
- Whitin, T. M. 1954. Inventory Control Research: A survey. *Management Science*, 1(1), 32-40.

- Van Mieghem, J. A. 2003. Capacity Management, Investment and Hedging: Review and Recent Developments. *Manufacturing & Service Operations Management*, 5(4), 269-302.
- Van Mieghem, J. A. 2004. Risk-Averse Newsvendor Networks: Resource Sharing, Substitution and Operational Hedging. Working paper.
- Zhang, X. 2001. Specification Tests of International Asset Pricing Models. *Forthcoming in Journal of International Money and Finance*.

Appendix

In this Appendix, we state and prove all the lemmas (Lemma 1 to Lemma 3). We also state and prove two lemmas that we used in the proofs of Lemmas 1 and 2: Lemma 4 and Lemma 5. For the ease of exposition, we put Lemmas 4 and 5 ahead of Lemmas 1, 2, and 3. We also state some properties of $[a, b]$ -divergent and $[a, b]$ -asymptotic functions, the definition of which is also found in that paper. In addition, we provide more details on Section 5.2 of the paper.

Property 1 *The following properties of $[a, b]$ -divergent and $[a, b]$ -asymptotic functions are easily verified.*

1. *If $f_i(x)$ is $[a_i, b_i]$ -divergent and $c_i \geq 0$, then $\sum_i c_i f_i(x)$ is $\left[\sum_i c_i a_i, \sum_i c_i b_i \right]$ -divergent.*
2. *If $f_i(x)$ is $[a_i, b_i]$ -asymptotic, then $\sum_i c_i f_i(x)$ is $\left[\sum_i c_i a_i, \sum_i c_i b_i \right]$ -asymptotic.*

Lemma 4 *For every time period t and every information set \mathbf{z}_t the following claims hold.*

- *The function $E_{\mathbf{D}_{t:t+L}, \mathbf{z}_{t+1:t+L+2}}[\beta_{t,t+L} \widehat{G}_t(y_t, \mathbf{z}_t, \mathbf{Z}_{t+1:t+L+2}, \mathbf{D}_{t:t+L})]$ is a well-defined Lipschitz continuous function of y_t . It is finite if y_t is finite.*
- *$f_t(x_t, \mathbf{z}_t)$ is a Lipschitz continuous function in x_t . If x_t is finite, then f_t is either finite or equal to $-\infty$, i.e., $f_t \in \mathfrak{R} \cup \{-\infty\}$.*
- *$V_t(y_t, \mathbf{z}_t)$ is a Lipschitz continuous function in y_t . If y_t is finite, then $V_t \in \mathfrak{R} \cup \{-\infty\}$.*
- *If the minimum in (4) is attained at a finite y_s for all $s \geq t$ and all \mathbf{z}_s , then V_t and f_t are finite when x_t, y_t are finite.*

Proof: We prove a series of claims, the first of which has to do with the quarter- t cost function \widehat{G}_t . Let $1 \leq \tau \leq t$. The Finiteness Assumption implies that for every \mathbf{z}_τ , $E_{\mathbf{D}_{t:t+L}, \mathbf{Z}_{\tau+1:t+L+2}}[\beta_{\tau,t+L} \widehat{G}_t(y_t, \mathbf{z}_t, \mathbf{Z}_{t+1:t+L+2}, \mathbf{D}_{t:t+L})]$ exists and is finite when $y_t = 0$. Equations (1) and (2) imply that the derivative of \widehat{G}_t , before taking any expectations, is between $-(\pi_t + p_t)$ and $h_t \vee \beta_{t+1} p_{t+1}$. By the Finiteness Assumption, if either $\tau = t$ and y_t is finite, or if $y_t = (x_1 - D_{1,t-1})^+$, then $E_{\mathbf{D}_{t:t+L}, \mathbf{Z}_{\tau+1:t+L+2}}[\beta_{\tau,t+L} \widehat{G}_t(y_t, \mathbf{z}_t, \mathbf{Z}_{t+1:t+L+2}, \mathbf{D}_{t:t+L})]$ is $(E_{\mathbf{Z}_{\tau+1:t}}[\beta_{\tau,t} u_t])$ -Lipschitz continuous for every \mathbf{z}_τ , and the first assertion of the Lemma holds.

Our second claim has to do with the end-of-horizon cost function f_{T-L+1} . Let $1 \leq \tau \leq T - L + 1$. The Finiteness Assumption implies that $E_{\mathbf{Z}_{\tau+1:T-L+1}}[\beta_{\tau,T-L} f_{T-L+1}]$ is finite when $x_{T-L+1} = 0$. It is easily verified that the quantity in brackets in (5) has a derivative whose absolute value is at most $|v| + c_{T-L+1}$. Consequently, by the Finiteness Assumption, $E_{\mathbf{Z}_{\tau+1:T-L+1}}[\beta_{\tau,T-L} f_{T-L+1}]$ is $(E_{\mathbf{Z}_{\tau+1:T-L+1}}[\beta_{\tau,T-L} U_{T-L+1}])$ -Lipschitz-continuous in x_{T-L+1} , and hence $E_{\mathbf{Z}_{\tau+1:T-L+1}}[\beta_{\tau,T-L} f_{T-L+1}]$ is finite if x_{T-L+1} is finite. Setting $\tau = T - L + 1$ we see that the second and fourth assertions of the Lemma hold for f_{T-L+1} . Also note that f_{T-L+1} is policy-independent.

Our third claim is that the order-up-to-zero policy, which we will call ζ , has finite expected cost. Let f_t^ζ and V_t^ζ be analogous to f_t and V_t in (3) and (4), corresponding to ζ . Specifically, f_t^ζ and V_t^ζ satisfy (4), and they satisfy (3) if we replace the minimum over y_t with x_t^+ . We claim that $f_t^\zeta(x_t, \mathbf{z}_t)$ and $V_t^\zeta(y_t, \mathbf{z}_t)$ are finite for all finite values x_t and y_t , for all \mathbf{z}_t . To prove the claim note that under ζ , $y_t = (x_1 - D_{1,t-1})^+$. Hence $V_t^\zeta(y_t, \mathbf{z}_t) = c_t y_t + \sum_{t:\tau \leq t \leq T-L} E_{\mathbf{D}_{\tau:t+L}, \mathbf{Z}_{\tau+1:t+L+2}}[\beta_{\tau,t+L} \widehat{G}_t((x_1 - D_{1,t-1})^+, \mathbf{z}_t, \mathbf{Z}_{t+1:t+L+2}, \mathbf{D}_{t:t+L})]$ + $E_{\mathbf{D}_{\tau:T-L}, \mathbf{Z}_{\tau+1:T-L+1}}[\beta_{\tau,T-L} f_{T-L+1}((x_1 - D_{1,T-L})^+, \mathbf{Z}_{T-L+1})]$. This is finite by our first two claims. Substituting $y_t \leftarrow x_t^+$ in (3) we see that $f_t^\zeta(x_t, \mathbf{z}_t)$ is also finite for finite y_t .

We now prove the lemma by induction. For $1 \leq t \leq T$ we define $U_t = 2c_t + u_t + E[\beta_t U_{t+1}]$. The Finiteness Assumption implies that U_t is finite for all t . Our inductive hypothesis consists of the following affirmations.

1. V_t is $(U_t - c_t)$ -Lipschitz-continuous in y_t , the second and fourth assertions of the Lemma hold for V_t , and $V_t \leq V_t^\zeta$, for $1 \leq t \leq T - L$ and all \mathbf{z}_t .
2. f_t is U_t -Lipschitz-continuous in x_t , the third and fourth assertions of the Lemma hold for f_t , and $f_t \leq f_t^\zeta$, for $1 \leq t \leq T - L + 1$ and all \mathbf{z}_t .

The second claim initializes the induction by establishing Affirmation 2 for $t = T - L + 1$. Assume that we have proven Affirmation 2 for $t + 1$. By (4), $V_t \leq V_t^\zeta$, so $V_t \in \mathfrak{R} \cup \{-\infty\}$ when y_t is finite. Because $U_t = 2c_t + u_t + E[\beta_t U_{t+1}]$ is finite, V_t is $(U_t - c_t)$ -Lipschitz-continuous. We see that Affirmation 1 holds for t .

We now assume that Affirmation 1 holds for t and prove Affirmation 2. By (3), $f_t \leq f_t^\zeta$, so $V_t \in \mathfrak{R} \cup \{-\infty\}$ when x_t is finite. In (3), $\min_{y_t > x_t} \{K_t \delta(y_t - x_t) + V_t(y_t, \mathbf{z}_t)\} = K_t + \min_{y_t > x_t} V_t(y_t, \mathbf{z}_t)$ and $\{K_t \delta(0) + V_t(x_t, \mathbf{z}_t)\} = V_t(x_t, \mathbf{z}_t)$ are both $(U_t - c_t)$ -Lipschitz-continuous in x_t , by Affirmation 2. Therefore their minimum is $(U_t - c_t)$ -Lipschitz-continuous in x_t , and f_t is U_t -Lipschitz-continuous. At this point Affirmation 2 follows readily. \diamond

We consider a function to be Lipschitz-continuous if it is equal to $-\infty$ everywhere.

Lemma 5 *Let $f(x)$ be $[a, b]$ -asymptotic, and let W and D be correlated random variables with $E(|W|) > 0$ and either $W \geq 0$ or $W \leq 0$. Then $E[Wf(x - D)]$ is $[E(W)a, E(W)b]$ -asymptotic if the expectations exist, and if $E[WD]$ and $E[Wf(x - D)]$ exist and are finite for all finite x .*

Proof: Since $-f(x)$ is $[-a, -b]$ -asymptotic, we can change the signs of both W and f without altering the claim. Therefore, we assume that $W \geq 0$ and $E(W) > 0$. Let $\zeta \in \{-1, 1\}$ be a constant. For every $\epsilon > 0$ and $A \geq 0$, there is a v such that for all $x, x \geq v$, we have $A + (b - \epsilon)x \leq f(x) \leq -A + (b + \epsilon)x$. By case analysis ($\zeta = 1, -1$), we see that $A \leq \zeta [f(x) - (b - \zeta \epsilon)x]$. Then

$$\begin{aligned}
& E\{W \zeta [f(x - D) - (b - \zeta \epsilon)x]\} - E\{W \zeta [f(x - D) - (b - \zeta \epsilon)x] 1(x - D \leq v)\} \\
&= E\{W \zeta [f(x - D) - (b - \zeta \epsilon)x] 1(x - D > v)\} \\
&= E\{W \zeta [f(x - D) - (b - \zeta \epsilon)(x - D)] 1(x - D > v)\} - E\{W \zeta D (b - \zeta \epsilon) 1(x - D > v)\} \\
&\geq E[WA 1(x - D > v)] - E[W|D| (|b| + \epsilon)] \\
&\geq A E[W] - A E[W 1(x - D \leq v)] - (|b| + \epsilon) E[W|D|].
\end{aligned}$$

Since both W and $W|D|$ have finite means, the third term is a finite constant and the second term converges to 0 as $x \rightarrow \infty$. Consider the left-hand side of the inequality. Since $f(x)$ is linearly bounded, there exist constants A' and A'' such that

$$E\{|W \zeta [f(x - D) - (b - \zeta \epsilon)x]| 1(x - D \leq v)\}$$

$$\begin{aligned}
&\leq E\{W [(A' + A''|x - D|) + (|b| + \epsilon)x] 1(x - D \leq v)\} \\
&\leq E\{W [(A' + A''|D|) + (A'' + |b| + \epsilon)x] 1(x - D \leq v)\} \\
&\leq E\{W [(A' + A''|D|) + (A'' + |b| + \epsilon)(D + v)] 1(x - D \leq v)\}.
\end{aligned}$$

Because the expected values of W and $W|D|$ are both finite, the dominated convergence theorem applies, and this expression converges to 0 as $x \rightarrow \infty$.

We have proven that $\liminf_{x \rightarrow \infty} E\{W \zeta [f(x - D) - (b - \zeta \epsilon)x]\} \geq A E[W] - (|b| + \epsilon)E[W|D|]$. Since $E(W) > 0$ and this is true for all $A \geq 0$, $E[W \zeta f(x - D)] - E[W \zeta (b - \zeta \epsilon)x]$ diverges as $x \rightarrow \infty$. Similarly, $E[W \zeta f(x - D)] - E[W \zeta (a - \zeta \epsilon)x]$ diverges as $x \rightarrow -\infty$. Considering the cases $\zeta = 1$ and $\zeta = -1$, and recalling that this holds for all $\epsilon > 0$, we see that $E[W f(x - D)]$ is $[E(W)a, E(W)b]$ -asymptotic. \diamond

One consequence of Lemma 5 is that $E_{\mathbf{D}_{t:t+L}, \mathbf{Z}_{t+1:t+L+2}}[\beta_{t,t+L} \widehat{G}_t(y_t, \mathbf{z}_t, \mathbf{Z}_{t+1:t+L+2}, \mathbf{D}_{t:t+L})]$ is $[E_{\mathbf{Z}_{t+1:t+L+2}}(-\widehat{\pi}_{t+L}), E_{\mathbf{Z}_{t+1:t+L+2}}(h_{t+L})]$ -asymptotic.

Lemma 1 *Suppose that $V_t(\cdot, \mathbf{z}_t)$ is K_t -convex and $[\cdot, 0]$ -divergent. Then in period t with information set \mathbf{z}_t , an (s, S) policy is optimal. In cases PTD and PTO, if in addition $V_t(\cdot, \mathbf{z}_t)$ is $[0, \cdot]$ -divergent, the optimal (s, S) policies satisfy $s > -\infty$.*

Proof: For a given information set \mathbf{z}_t , suppose that whenever $x < y$ we have $V_t(x, \mathbf{z}_t) \leq V_t(y, \mathbf{z}_t) + K_t$. Then it is optimal to never order, and $s_t(\mathbf{z}_t) = -\infty$, i.e., an $(s, S) = (-\infty, S_t(\mathbf{z}_t))$ policy is optimal. Furthermore, $V_t(\cdot, \mathbf{z}_t)$ is not $[0, \cdot]$ -divergent.

On the other hand, if $V_t(\cdot, \mathbf{z}_t)$ is K_t -convex and $[\cdot, 0]$ -divergent, and if $V_t(x, \mathbf{z}_t) > V_t(y, \mathbf{z}_t) + K_t$ for some $x < y$, then $S_t(\mathbf{z}_t) < \infty$ and the proof becomes classical. There is no local maximum \underline{S}_t of V_t such that $\underline{S}_t < y$ and $V_t(\underline{S}_t, \mathbf{z}_t) > V_t(y, \mathbf{z}_t) + K_t$, because the existence of \underline{S}_t would violate the K_t -convexity of V_t . Therefore $\{x : x < y \text{ and } V_t(x, \mathbf{z}_t) > V_t(y, \mathbf{z}_t) + K_t\}$ is a non-empty connected set containing $-\infty$, and V_t is non-increasing on this set. Consequently $S_t(\mathbf{z}_t)$ and $s_t(\mathbf{z}_t)$ exist, $s_t(\mathbf{z}_t) > -\infty$, and if the starting inventory level is less than $S_t(\mathbf{z}_t)$, then the $(s_t(\mathbf{z}_t), S_t(\mathbf{z}_t))$ policy is optimal. Also, we have proven that $S_t(\mathbf{z}_t) > x$ whenever $x < y$ and $V_t(x, \mathbf{z}_t) > V_t(y, \mathbf{z}_t) + K_t$. Thus, if the starting inventory level is greater than $S_t(\mathbf{z}_t)$, then it is optimal not to order; i.e., the $(s_t(\mathbf{z}_t), S_t(\mathbf{z}_t))$ policy is optimal.

\diamond

Lemma 2 $V_t(\cdot, \mathbf{z}_t)$ is $[\mathcal{A}_t, \mathcal{B}_t]$ -asymptotic in cases PTD and PTO, and $[\cdot, \mathcal{B}_t]$ -asymptotic in case LS.

Proof: We prove the lemma by induction on t starting from period $T-L$. It is easy to show that $c_{T-L}y$ is $[c_{T-L}, c_{T-L}]$ -asymptotic and \widehat{G}_{T-L} is $[-\widehat{\pi}_T, h_T]$ -asymptotic. By Lemma 5, $f_{T-L+1}(x_{T-L+1}, \mathbf{z}_{T-L+1})$ is $[E_{\mathbf{z}_{T-L+2:T+1}}(-\beta_{T-L+1,T} c_{T-L+1}), E_{\mathbf{z}_{T-L+2:T+1}}(-\beta_{T-L+1,T} v)]$ -asymptotic in the backorder cases, and $[\cdot, E_{\mathbf{z}_{T-L+2:T+1}}(-\beta_{T-L+1,T} v)]$ -asymptotic in case LS. By (4), Property 1 and Lemma 5, V_{T-L} is $[\mathcal{A}_{T-L}, \mathcal{B}_{T-L}]$ -asymptotic ($[\cdot, \mathcal{B}_{T-L}]$ -asymptotic in case LS), and the lemma holds for $t = T - L$.

Assume that the lemma holds for $t + 1$, and recall that $\mathcal{B}_{t+1} > 0$ by assumption. Clearly, the minimum in (3) is $[\cdot, \mathcal{B}_{t+1}]$ -asymptotic, and f_{t+1} is $[\cdot, \mathcal{B}_{t+1} - c_{t+1}]$ -asymptotic. In the backorder cases (PTD and PTO), if $\mathcal{A}_{t+1} < 0$ then there is a finite y_{t+1} that minimizes the term in braces in (3). If $\mathcal{A}_{t+1} \geq 0$, then there may not exist a finite minimizer y_{t+1} . In either case the minimum in (3) is $[(\mathcal{A}_{t+1})^+, \mathcal{B}_{t+1}]$ -asymptotic, and hence f_{t+1} is $[(\mathcal{A}_{t+1})^+ - c_{t+1}, \mathcal{B}_{t+1} - c_{t+1}]$ -asymptotic.

By Property 1 and Lemma 5, and because $c_t y_t$ is $[c_t, c_t]$ -asymptotic and \widehat{G}_t is $[-\widehat{\pi}_{t+L}, h_{t+L}]$ -asymptotic, V_t is $[\mathcal{A}_t, \mathcal{B}_t]$ -asymptotic ($[\cdot, \mathcal{B}_t]$ -asymptotic in case LS), and the lemma holds for period t .

◇

Lemma 3 $V_t(y_t, \mathbf{z}_t)$ is K_t -convex if the following conditions hold. For all possible financial information vectors \mathbf{z}_{T-L} , $\mathcal{B}_{T-L} - \mathcal{A}_{T-L} \geq 0$. In addition, for all \mathbf{z}_j and all $j = t, \dots, T - L - 1$,

1. $h_{j+L} + \widehat{\pi}_{j+L} \geq 0$ with backorders and $h_j + \widehat{\pi}_j - c_{j+1} \geq 0$ with lost sales, and
2. $K_j(\mathbf{z}_j) \geq E_{\mathbf{z}_{j+1}}[\beta_j(\mathbf{z}_j, \mathbf{z}_{j+1})K_{j+1}(\mathbf{z}_{j+1})]$.

Proof: We first consider the backorder cases. In these cases, the condition $\mathcal{B}_{T-L} - \mathcal{A}_{T-L} \geq 0$ implies that $E_{\mathbf{z}_{T-L+1:T+1}}[\beta_{T-L,T}(h_T + \widehat{\pi}_T + c_{T-L+1} - v)] \geq 0$ and hence

$$\begin{aligned}
& V_{T-L}(y_{T-L}, \mathbf{z}_{T-L}) \\
&= c_{T-L}y_{T-L} + E_{\mathbf{D}_{T-L:T}, \mathbf{z}_{T-L+1:T+1}}[\beta_{T-L,T}\widehat{G}_{T-L}(y_{T-L}, \mathbf{z}_{T-L}, \mathbf{z}_{T-L+1:T+1}, \mathbf{D}_{T-L:T}) \\
&\quad + \beta_{T-L}f_{T-L+1}(y_{T-L} - D_{T-L}, \mathbf{z}_{T-L+1})] \\
&= c_{T-L}y_{T-L} + E_{\mathbf{z}_{T-L+1:T+1}}\{E_{\mathbf{D}_{T-L:T}}[\beta_{T-L,T}\widehat{G}_{T-L}(y_{T-L}, \mathbf{z}_{T-L}, \mathbf{z}_{T-L+1:T+1}, \mathbf{D}_{T-L:T}) \\
&\quad + \beta_{T-L}f_{T-L+1}(y_{T-L} - D_{T-L}, \mathbf{z}_{T-L+1})|\mathbf{z}_{T-L+1:T+1}]\} \\
&= c_{T-L}y_{T-L} + E_{\mathbf{z}_{T-L+1:T+1}}\{\beta_{T-L,T}E_{\mathbf{D}_{T-L:T}}[(h_T - v)(y_{T-L} - D_{T-L,T})^+
\end{aligned}$$

$$+(\hat{\pi}_T + c_{T-L+1})(y_{T-L} - D_{T-L,T})^- - p_T D_T | \mathbf{Z}_{T-L+1:T+1} \},$$

which is convex. Furthermore, the condition $h_{t+L} + \hat{\pi}_{t+L} \geq 0$ affirms that \hat{G}_t is convex. The proof is completed by induction on t , in the classical manner, based on (3), (4), and the condition $K_t(\mathbf{z}_t) \geq E_{\mathbf{Z}_{t+1}}[\beta_t(\mathbf{z}_t, \mathbf{Z}_{t+1}) K_{t+1}(\mathbf{Z}_{t+1})]$.

Under lost sales, $L = 0$ and

$$\begin{aligned} V_T(y_T, \mathbf{z}_T) &= c_T y_T + E_{D_T, \mathbf{Z}_{T+1}}[\beta_T \hat{G}_T(y_T, \mathbf{z}_T, \mathbf{Z}_{T+1}, D_T) + \beta_T f_{T+1}((y_T - D_T)^+, \mathbf{Z}_{T+1})] \\ &= c_T y_T + E_{\mathbf{Z}_{T+1}}\{\beta_T E_{D_T}[\hat{G}_T(y_T, \mathbf{z}_T, \mathbf{Z}_{T+1}, D_T) + f_{T+1}((y_T - D_T)^+, \mathbf{Z}_{T+1})] | \mathbf{Z}_{T+1}\} \\ &= c_T y_T + E_{\mathbf{Z}_{T+1}}\{\beta_T E_{D_T}[(h_T - v)(y_T - D_T)^+ + \hat{\pi}_T (y_T - D_T)^- - p_T D_T] | \mathbf{Z}_{T+1}\}. \end{aligned}$$

Because $\mathcal{B}_T - \mathcal{A}_T = E_{\mathbf{Z}_{T+1}}[\beta_T(h_T + \hat{\pi}_T - v)] \geq 0$, $V_T(\cdot, \mathbf{z}_T)$ is convex.

The inductive step of the proof for the lost sales case proceeds as follows. Assume that $V_{t+1}(y_{t+1}, \mathbf{z}_{t+1})$ is K_{t+1} -convex. In (3), the classical logic implies that

$$F_{t+1}(x_{t+1}, \mathbf{z}_{t+1}) = f_{t+1}(x_{t+1}, \mathbf{z}_{t+1}) - c_{t+1} x_{t+1} = \min_{y_{t+1} \geq x_{t+1}} \{K_{t+1} \delta(y_{t+1} - x_{t+1}) + V_{t+1}(y_{t+1}, \mathbf{z}_{t+1})\}$$

is a K_{t+1} -convex function of x_{t+1} . Using the properties of $V_{t+1}(x_{t+1}, \mathbf{z}_{t+1})$ and $F_{t+1}(x_{t+1}, \mathbf{z}_{t+1})$ discussed in the second paragraph of the proof of Lemma 1, and considering the cases $s_{t+1} > 0$, $s_{t+1} = 0 < S_{t+1}$ and $S_{t+1} = 0$, we can show that $F_{t+1}(x_{t+1}^+, \mathbf{z}_{t+1})$ is a K_{t+1} -convex function of x_{t+1} , where x_{t+1} can be either positive or negative. In case LS, $\gamma(x_{t+1}) = (x_{t+1})^+$, so we can write $f_{t+1}(\gamma(x_{t+1}), \mathbf{z}_{t+1}) = -c_{t+1} x_{t+1}^+ + F_{t+1}(x_{t+1}^+, \mathbf{z}_{t+1})$.

Now consider equation (4) for period t . The term in brackets is

$$\begin{aligned} &\beta_t \hat{G}_t(y_t, \mathbf{z}_t, \mathbf{Z}_{t+1}, D_t) + \beta_t f_{t+1}((y_t - D_t)^+, \mathbf{Z}_{t+1}) \\ &= \beta_t [(h_t - c_{t+1})(y_t - D_t)^+ + \hat{\pi}_t (y_t - D_t)^- + F_{t+1}((y_t - D_t)^+, \mathbf{Z}_{t+1}) - p_t D_t]. \end{aligned}$$

The condition that $h_t + \hat{\pi}_t - c_{t+1} \geq 0$ guarantees that $(h_t - c_{t+1})(y_t - D_t)^+ + \hat{\pi}_t (y_t - D_t)^-$ is convex. Consequently, the term in brackets is K_{t+1} -convex. The condition $K_t(\mathbf{z}_t) \geq E_{\mathbf{Z}_{t+1}}[\beta_t(\mathbf{z}_t, \mathbf{Z}_{t+1}) K_{t+1}(\mathbf{Z}_{t+1})]$ implies that the expectation of this term over \mathbf{Z}_{t+1} , conditioned on \mathbf{z}_t , is K_t -convex and the inductive step of the proof for the lost sales case is completed. \diamond

Details for Section 5.2 – The Risk Factors: Selection and Evolution

In this section we give more detail on our approach to the selection and evolution of the risk factors, summarized in Subsection 5.2. We start with the regression that we used in to determine the significance of the five risk factors suggested in the finance literature for firms located in Mexico. Recall that we used the log linear model for the discount factor, $\log(\beta_t) = \mathbf{c}^T \log(\mathbf{g}_t)$ (see Model 3 in Subsection 4.2. To estimate the parameters \mathbf{c}^T of the log linear model, we use historical data on the return of the DESC group, which does business in both Mexico and the U.S. in four divisions: chemical, automotive, food, and real estate. We collected monthly exchange rates for the time period from January 1976 to May 2004, the monthly Mexican market index from January 1981 to May 2004, and all other data from January 1988 to May 2004. (All data are obtained from MSCI and the DATASTREAM database from the library at the Johnson Graduate School of Management, Cornell University. Since the returns of the DESC group before 1988 are not available, the regression uses data from January 1988 on.) Using the monthly data, stepwise regression showed that the exchange rate risk (g_{1t}) and Mexican domestic market risk (g_{3t}) were significant at the 95% level, and that the constant factor $g_{0t} \equiv 1$, g_{2t} and g_{4t} were not significant. Consequently, we use g_{1t} and g_{3t} as the risk factors.

Note that the regressions described above need to be done separately for election months and regular months, and that they have more statistical power if they are based on monthly data rather than quarterly data. However, to keep the total number of states in the DP model under control, optimization models use time periods that are one quarter long. For regular months we did monthly regressions to determine which financial factors are significant, and quarterly regressions to obtain regression coefficients that are compatible with our optimization models. For election months there was not enough quarterly data to make the second step (the estimation of coefficients) meaningful. Therefore we did all of our regressions using monthly data, and we converted the resulting model of a monthly discount factor into a model for a quarterly discount factor analytically by doing convolutions. Table 8 gives the coefficients \mathbf{c}^T for the quarterly discount factor that we obtained in this manner, and the corresponding significant levels.

We looked at data from other Mexican companies as well, but among the Mexican companies we were able to identify that do substantial amounts of business in the U.S., the DESC group has the most complete series of published corporate returns. Therefore in our case study, we make the assumption

		Change in the Exchange Rate (g_1)	Mexican Market Return Plus One (g_3)
Election Quarters	Coefficients	1.18	-0.477
	Significance	99.6% ^a	95.5% ^a
Regular Quarters	Coefficients	0.531	-0.943
	Significance	78% ^b	99.99% ^c

Table 8: Predicting the Quarterly Discount Factor β_t : Coefficients and Significance Levels for Each Risk Factor

^aFrom a regression using monthly data.

^bFrom a regression using quarterly data. Using monthly data the significance is 99.99%.

^cFrom a regression using quarterly data.

that all four of the firms have returns, and (consequently) discount factors, whose relationship to the five financial risk factors is similar to that of the DESC group.

We now model the evolution over time of the risk factors g_{1t} and g_{3t} , following the approach of Subsection 4.3. For the reasons given in the second paragraph of this subsection, we work with monthly data during election months, and quarterly data during regular quarters. We start by testing whether there is autocorrelation in g_{1t} and g_{3t} . There is statistically significant autocorrelation in g_{1t} (change in the exchange rate) during regular quarters, but not in g_{3t} (Mexican market return). During election months neither g_{1t} nor g_{3t} displayed any autocorrelation. Therefore we model g_{3t} , and g_{1t} during election months as iid random variables.

For the regular quarters, we model the change in the exchange rate g_{1t} using the standard logarithmic AR(1) autoregressive model $\log(g_{1,t+1}) = \alpha_0 + \alpha_1 \log(g_{1t}) + \log(\eta_{1t})$ (see (15)). We assume that the $\log(\eta_{1t})$'s are iid with mean 0. Using regression, we find that $\alpha_0 = 0.00609$ and $\alpha_1 = 0.675$. Hence $\log(g_{1t})$ is a stable process.

To complete the model we need distributions for the following four iid random variables: η_{1t} and g_{3t} during the regular quarters, and g_{1t} and g_{3t} during election months. We accomplish this by discretizing the logarithms of these random variables and creating histograms. We use quarterly historical data from regular months for $\log(\eta_{1t})$ and $\log(g_{3t})$, and monthly historical data from election months for $\log(g_{1t})$ and $\log(g_{3t})$. The resulting histograms appear in Figure 6. Note that during election months the values are more spread out than they are in regular quarters, indicating a higher probability of financial turbulence during the election months.

As we mentioned earlier, our optimization algorithm uses quarterly time periods. For regular quarters, the histograms for the quarterly random variables $\log(\eta_{1t})$ and $\log(g_{3t})$ in Figure 6 are used directly as distributions. For election months, we need quarterly distributions for $\log(g_{1t})$ and $\log(g_{3t})$, but the histograms of $\log(g_{1t})$ and $\log(g_{3t})$ in Figure 6 are monthly. These histograms are converted to quarterly distributions through numerical convolution.

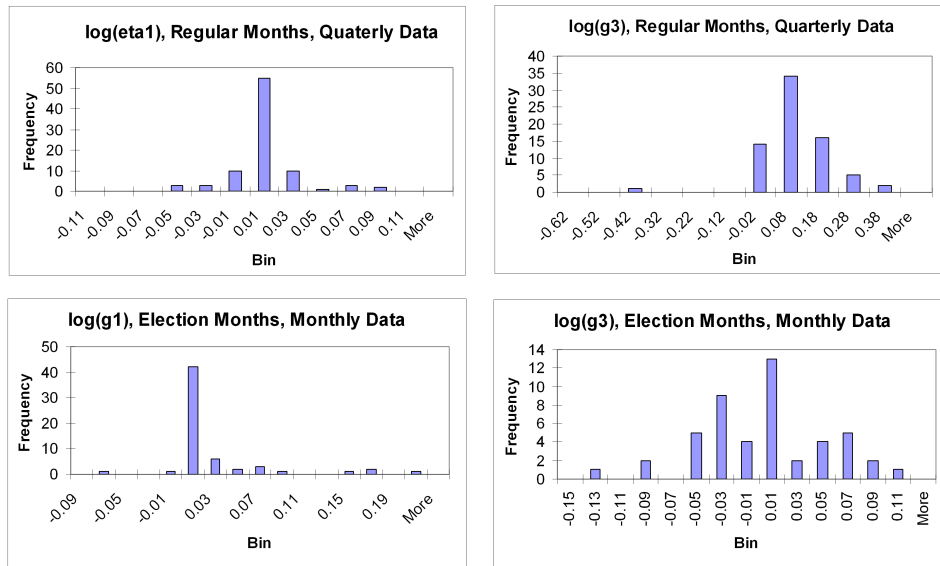


Figure 6: Histograms.

$$g_{1t} = \text{relative change in peso-per-dollar exchange rate,}$$

$$g_{3t} = (1 + \text{Mexican market return}).$$

As we mentioned in the main body of the paper, this subsection and Subsection 4.2 give rise to Table 2 in which, for each quarter t , we give the mean and standard deviation of the stochastic discount factor $\beta_t = \beta_t(\mathbf{Z}_t, \mathbf{Z}_{t+1})$ (the ordered pair $(\mathbf{Z}_t, \mathbf{Z}_{t+1})$ is random).