Optimal Timing of Inventory Decisions with Price Uncertainty

Vishal Gaur*, Sridhar Seshadri†, Marti G. Subrahmanyam†

July 8, 2007

Abstract

What is the optimal timing of inventory investment for a firm when its forecasts of demand and price improve with time but are correlated with the prices of traded assets in the financial markets? We consider this problem using a single period inventory model where demand and price are realized at time T and the stocking decision may be made at any time in the interval $[0, T]$. The processes for the firm's value as well as that of the market evolve as geometric Brownian motions. We show that the right to make the optimal inventory decision is a modified American-style option. However, since the stochastic variables defining the forecasts of demand and price are not tradeable, we cannot use standard dynamic hedging arguments in the Black-Scholes-Merton sense. Therefore, we use a risk-adjusted valuation approach for incomplete markets to determine the optimal timing strategy.

Our model provides results regarding the value of the option to postpone inventory procurement as a function of the key parameters: the market price of risk, the volatilities of price and demand, the correlation between these two variables and the return on the market portfolio, and the procurement cost. We illustrate the empirical validity of our model by testing it on firms in the gold mining industry. Thus, we shows that our model provides new evidence on the correlation between days of inventory and price volatility.

*The Johnson School, Cornell University, Sage Hall, Ithaca NY 14850. E-mail: *vg77@cornell.edu*.

†Leonard N. Stern School of Business, New York University, 8-160, 44 West Fourth Street, New York, NY 10012. E-mail: *sseshadr@stern.nyu.edu, msubrahm@stern.nyu.edu*.

1 Introduction

The value of postponement of inventory stocking decision is well recognized in the operations management literature. Swaminathan and Lee (2003) discuss various mechanisms for enabling postponement that have been studied in the literature. Whang and Lee (1999) classify the sources of the value of postponement into two types: uncertainty resolution and forecast improvement. The first is the flexibility to delay product differentiation so that more recent information about demand can be used at the time of differentiation (Feitzinger and Lee (1997), Lee and Tang (1997), Swaminathan and Tayur (1998), Aviv and Federgruen (2001)). The second is the flexibility to respond to early sales because the forecast of demand improves dramatically after observing early sales, which can then be used to make more accurate inventory decisions (Murray and Silver (1966), Fisher and Raman (1996), Eppen and Iyer (1997), Kaminsky and Swaminathan (2001, 2004)). Postponement mechanisms are also used in the literature on the value of information sharing in supply chains to determine the sources of the value of information (Cachon and Fisher 2000).

While both the above sources of the value of the postponement decision rely on early observation of demand, the forecast of demand could also improve with time due to resolution of other uncertainties. For example, Caldentey and Haugh (2005), and Gaur and Seshadri (2005) study models in which the forecast of demand is correlated with economic indices. Thus, a firm may profitably use a postponement strategy even when early demand cannot be observed. In addition, the value of postponement may result not only from the resolution of demand uncertainty, but also from the resolution of price uncertainty. The literature on postponement treats prices as exogenous and constant, whereas in many practical applications, the prices of products are stochastic and the forecast of prices improve over time, affecting the value of postponement. Further, shocks to prices could be positively or negatively correlated with shocks to demand, and one might ask how the value of postponement is affected by this correlation structure. Finally, the cost of procurement may change over time as new information is revealed.

This paper presents a model to study the impact of the above factors on the comparative value of postponement versus the early commitment of the inventory stocking decision. We consider a single-period, single-product inventory model. The selling price and the demand for the product are stochastic and are interrelated through the price elasticity of demand and correlated noise processes. The firm is a price taker in the goods market as well as the financial market. Demand occurs at a fixed future date, T, while the forecasts of demand and price evolve continuously over time from time 0 to time T, enabling the firm to gain by way of more accurate forecasts if it postpones the inventory decision. The forecasts of price and demand are correlated with prices in the financial market, so that market information can be used to explicitly compute the cost of capital of the firm and use it to determine the optimal inventory and optimal timing decisions. The inventory purchase decision can be made at any time between 0 and T, possibly with an increase in unit costs as the ordering date gets closer to time T. In this model, we treat postponement as a compound option on the forecasts of price and demand of the product. The value of postponement at time

t is measured as the difference between the expected profit from postponing the inventory decision and the expected profit from committing to the inventory decision at the current time.

This model allows us to study the value of postponement in relation to economic factors other than the ones studied previously. For example, the value of postponement critically depends on the cost of postponement. In our model, we view this as the unit cost of procurement at a time that is closer to the selling season. More generally, this cost can be viewed as the cost of shortening the lead-time of procurement and/or processing, the cost of redesigning the product to permit delayed differentiation, developing new sources that are closer to the market, etc. We obtain conditions on the cost frontier under which postponement is optimal. This helps identify situations under which supply chain re-design can be most beneficial. Thus, we attempt to link the degree of postponement to economic viability determined under the prevailing market prices.

There are several other factors that can affect the value of postponement. The parameters of the demand distribution, i.e., its mean and variance, are the most obvious ones and have been explored in previous work, e.g., Fisher and Raman (1996). With respect to prices, we expect to find that the revelation of price information is almost as critical as the demand information, although for reasons other than matching demand and supply. First, when firms expect prices to be high, they are willing to take greater risks (overstock) than when prices are low. Thus, the evolution of price information enables the firm to exercise better control over its safety stock. Second, the update of price information gives the firm the option to delay stocking the product in case the price forecast is too low. We determine the value of this option as a component of the value of postponement, and also study how it is affected by the correlation between price and demand forecasts. Our modeling of price builds on prior research on inventory management for traded commodities, see for example Akella et al. (2002), Seifert et al. (2004), and Hakzos and Seshadri (2006). It differs from the literature on joint inventory and pricing decisions in the respect that we do not treat price as a decision variable, but as an exogenous stochastic variable correlated with demand. Thus, our model is more applicable in competitive industries where the firm has limited pricing power and acts as a price-taker.

Finally, we allow the price and the demand of the product to be correlated with asset prices in the financial market. Thus, we study how the aggregate risk aversion of investors as reflected in the market price of risk affects the willingness of the decision-maker to postpone the inventory decision. Several papers have studied the impact of risk-aversion on inventory decisions, see for example Eeckhoudt et al. (1995), Agrawal and Seshadri (2000a, 2000b), and Chen and Federgruen (2000). However, there have not been any studies on the effect of risk-aversion on the *timing* of the inventory decision. One might argue that when the market price of risk is high, the firm's cost of capital is high and it is willing to pay more for demand information compared to when the market price of risk is low. Therefore, the propensity to postpone the inventory decision increases in the degree of risk aversion of the market. Our analysis shows that this argument may have to be modified by the correlation of demand and price processes with the financial market. We find that the propensity of the risk-averse decision maker to postpone the inventory decision increases or decreases in the degree of risk-aversion depending on whether the demand and price processes are positively or negatively correlated with the return on the market portfolio. It should be emphasized, however, that we do not consider the viewpoint of an individual risk-averse decision maker in this paper. Rather, we consider the effect of aggregate risk aversion of investors as reflected in the market price of risk assuming that the firm has free access to the capital market.

The modeling of risk-aversion in our paper merits comparison with the previous literature in joint operations-finance models. There are two broad paradigms in this area. One paradigm is based on the work of Sandmo (1971), and treats the decision-maker as risk-averse with an increasing concave or a mean-variance utility function. The operations literature that follows this preference-based approach in decision-making and hedging contexts includes Caldentey and Haugh (2005), Dong and Liu (2005), Gaur and Seshadri (2005), Levy (1985), Mathur and Ritchken (1999), Perrakis and Ryan (1984), Ritchken (1985), Ritchken and Kuo (1989), and Van Mieghem (2003). The other paradigm is based on Constantinides (1978) and McDonald and Siegel (1985), and assumes that the firm is owned by risk-averse investors, so that the decision maker must use the appropriate risk-adjusted cost of capital to make optimal investment decisions. This cost of capital takes into account the aggregate risk aversion of the market through the market price of risk. We follow the latter approach in this paper since it is "descriptive of value-maximizing publicly-owned firms, and is widely used in the finance literature" (McDonald and Siegel (1985)). In previous research, Birge (2000) uses this approach to study the impact of risk on capacity planning models, Birge and Xu (2005) incorporate bankruptcy deadweight costs in this approach in the context of multi-period inventory models, and Kogut and Kulatilaka (1994), Huchzermeier and Cohen (1996), and Kouvelis (1999) use this approach to study the value of operational flexibility. While this approach treats the financial market as being complete, it can in principle be modified to accommodate approximate arbitrage (Bernardo and Ledoit (1999, 2000), Cochrane and Saa-Requejo (2000), Hansen and Jagannathan (1991), Ross (1976), Shanken (1992), Snow (1992), Stutzer (1993)) or incomplete markets (Gaur, Seshadri and Subrahmanyam (2005)).

We empirically validate our model using public financial data for 22 gold mining firms for the period 1995-2005. Our theoretical analysis yields a closed form formula for optimal inventory. We directly use this formula for estimation purposes. We find that the days of inventory of gold mining firms is correlated with linear and quadratic terms in gold price volatility due to the effect of price volatility on safety stock. We also find that days of inventory is negatively correlated with gross margin, as previously documented in the literature using data for other industries (Gaur et al. 2005, Roumianetsev and Netessine 2006). Thus, our empirical results support findings in the existing literature and provide new insights.

The rest of this paper is organized as follows: §2 describes our model, §3 presents the main results from the computation of the value function, and §4 presents comparative statics of the value function to analyze the determinants of the optimal choice between postponement and early commitment. Finally, §5 illustrates empirical insights from our model applied to the gold mining industry, and §6 concludes the paper. All proofs are provided in the appendix unless otherwise noted.

2 Model

We consider a single-period model of a firm selling a product with stochastic demand and a stochastic selling price. The firm is a price-taker in the product and capital markets. Time is indexed from 0 to T with the price P_T and the demand D_T being realized at time T, and the inventory decision being taken at a time instant $t \in [0, T]$ given forecasts of price and demand available at that time. Cash flows are discounted at a constant risk-free rate of interest r.

We allow the price forecast to evolve continuously over time according to a mean-reverting continuous time stochastic process:

$$dP_t = h(m - \log P_t)P_t dt + s_P P_t dz_P, \qquad t \in [0, T]. \tag{1}$$

Here, h is the speed of reversion, m determines the long-run mean price, s_P is the instantaneous volatility of price, P_t is a state variable representing the information about price available to the firm at time t, and z_P is a standard Brownian motion. At time t, the firm uses the process specification (1) and the information P_t to construct a forecast of price at time T. The actual price, P_T, is stochastic due to any number of environmental factors such as competition, imperfections in quality of the product, changes in customer preferences, etc., that are unknown at time t. However, we assume that the firm's forecast of price becomes increasingly accurate with time, with the forecast at time T equal to the actual realization P_T. Such an improvement in accuracy takes place due to the gradual revelation of information as modeled by (1). Our specification for price forecast is similar to standard models of convergence of futures prices of commodities to cash as constructed by Ross (1995) and Schwartz (1997).

We assume that demand is isoelastic in price, i.e., $D_T = aP_T^{-\eta}\epsilon_T$, where a is a scaling constant, η is the constant price elasticity of demand, and ϵ_T is a random noise in the scale of demand. As with price, we assume that the firm has a forecast of ϵ_T at time $t \in [0, T]$ evolving as a continuous time stochastic process specified as:

$$d\epsilon_t = \alpha_\epsilon \epsilon_t dt + s_\epsilon \epsilon_t dz_\epsilon, \qquad t \in [0, T]. \tag{2}$$

Here, ϵ_t and ϵ_T have interpretations analogous to P_t and P_T, and z_ϵ is a standard Brownian motion. We call this process as the demand forecast process. For convenience, we define $D_t = aP_t^{-\eta}\epsilon_t$. The forecasts of demand and price at time t can be written as $E[D_T|P_t, \epsilon_t]$ and $E[P_T|P_t, \epsilon_t]$; they may not be equal to D_t and P_t, respectively, due to the presence of drift rates in the forecast processes. Also note that we model price as mean-reverting but the noise process in demand as a geometric Brownian motion (GBM) as per conventional knowledge. Price forecast processes are generally known to be mean-reverting but demand forecast processes are not. A GBM representation for ϵ_t coupled with a multiplicative model for D_T is useful because it constrains demand to be always positive.

We assume that z_ϵ is correlated with z_P with correlation coefficient ρ, i.e., $\rho dt = dz_P dz_\epsilon$. Thus, in our model, demand changes with price in two ways, one along the demand curve due to the price elasticity of demand η, and second due to a scaling of the demand curve caused by the conditional distribution of ϵ_T changing with price. A positive or a negative correlation between z_ϵ and z_P is plausible in practice. For example, if new competition or a new technology emerges in the market, then the firm's realized demand and price may both be lower than the respective forecasts, resulting in a positive correlation between z_ϵ and z_P. On the other hand, a maturing of the product lifecycle may expand the market and lower prices simultaneously, introducing a negative correlation between z_ϵ and z_P.

We do not distinguish between finished goods and raw material inventory, and assume that the firm can process the items obtained from the supplier at negligible cost and in negligible time. The firm can make its stocking decision at any time $t \in [0, T]$. For an order to be placed at time t, we assume that the supplier can provide an order lead-time of $L = T - t$. The cost of purchase may increase with time as we get closer to the market date T due to the decrease in lead-time. Hence, we assume that the unit purchase cost is time-dependent, and is denoted as c_t. The unit cost c_t includes all costs incurred for holding the finished product until the time of sale at time T. For simplicity, we assume that the firm can make its inventory ordering decision only once in the entire period $[0, T]$, i.e., the firm is not allowed to accumulate inventory gradually over time either continuously or on a finite number of dates. Instead, there is only one purchase decision. Thus, the firm's decision variables are the timing of the inventory purchase and the amount of inventory to purchase. We also assume zero salvage value of inventory left over at time T.

Our model is appropriate in settings in which the supplier requires the firm to place a single order in the entire selling season. The model is also sufficient to study the implications of price and demand volatility and changes thereon for the value of postponement. Note that the model is based on the traditional finite horizon inventory model except for stochastic prices and the continuous evolution of forecasts. Thus, like the traditional inventory model, it can be generalized in a number of ways. It can be extended in a straightforward manner to the case when the firm can make a finite number of purchases over time. It can also be extended to accommodate a lower bound on the lead-time by modifying the set of dates feasible for making the stocking decision to the interval $[0, T - \bar{L}]$, where $\bar{L} < T$ is a constant. A finite non-zero production time can be accommodated in the model. Finally, non-zero salvage value can also be accommodated. Since selling prices are stochastic, the salvage value may be represented as a fraction of the selling price, which would yield formulas similar to the ones we obtain in the paper.

3 Optimal Solution

Let $V(t, q)$ denote the value at time t of the claim on future revenues of the firm given information (P_t, ϵ_t) and inventory level q. Clearly, $V(T, q) = P_T \min\{q, D_T\}$ and $V(t, q) = e^{-r(T-t)} E[P_T \min\{q, D_T\} | P_t, \epsilon_t]$. Let $\pi(t, q) \equiv V(t, q) - c_t q$ denote the expected profit at time t if inventory q is purchased at time t. Let $q^*(t)$ denote the optimal inventory decision at time t and $\pi^*(t)$ denote the optimal expected profit. Also, for time instants u and t such that $0 \le u \le t \le T$, let $Y(u, t) \equiv e^{-r(t-u)} E[\pi^*(t) | P_u, \epsilon_u]$ denote the value of the firm at time u if the optimal inventory decision is taken at a fixed future date t, and the information at time u is (P_u, ϵ_u). In the absence of market-based valuation, these functions can be obtained using dynamic programming computations, albeit with stochastic prices. We compute these functions in §3.1. Then, in §3.2, we introduce the correlation of demand and price with a marketed asset and show how the value functions change under market-based valuation. We present three limiting cases of our formulas in §3.3. It turns out that these formulas have a simple structure. They also encompass the range of possibilities of postponement and stocking decisions that can result in our model. Thus, they are useful for sensitivity analysis and illustration.

3.1 Risk neutral investors

To focus on the basic intuition behind the model, we assume initially that investors are risk neutral. Thus, the value of the firm is computed under the decision maker's subjective probability measure with no adjustment for the price of risk. Let x_t denote $\log P_t$ and y_t denote $\log \epsilon_t$. Applying Ito's Lemma to (1) and (2), we have

$$\begin{aligned} dx_t &= h(m - \frac{s_P^2}{2h} - x_t)dt + s_P dz_P, \\ dy_t &= (\alpha_\epsilon - s_\epsilon^2/2)dt + s_\epsilon dz_\epsilon, \end{aligned}$$

for $t \in [0, T]$. Given the information (x_t, y_t) at time t, let μ_x and μ_y denote the means of x_T and y_T, σ_x^2 and σ_y^2 denote the variances of x_T and y_T, and σ_{xy} denote the covariance of x_T and y_T. We have

$$\begin{aligned} &\mu_x = x_t e^{-h(T-t)} + (m - \tfrac{s_P^2}{2h})(1 - e^{-h(T-t)}), \quad && \sigma_x^2 = (1 - e^{-2h(T-t)})\tfrac{s_P^2}{2h}, \\ &\mu_y = y_t + (\alpha_\epsilon - s_\epsilon^2/2)(T-t), && \sigma_y^2 = s_\epsilon^2(T-t), \\ &\sigma_{xy} = (1 - e^{-h(T-t)})\tfrac{\rho s_P s_\epsilon}{h}. \end{aligned}$$

Here, the expression for σ_{xy} is notable. It reflects the fact that we are correlating a mean-reverting process for price with a geometric Brownian motion for ϵ. Since the variance of the price forecast process tends to a constant as $T - t$ increases, the covariance of x_t and y_t is also found to tend to a constant, but at a slower rate. Thus, we find that the correlation coefficient between x_t and y_t, ρ_{xy}, is time-variant even though the correlation coefficient between z_P and z_ϵ is a constant. For simplicity, we write ρ_{xy} as $\sigma_{xy}/(\sigma_x\sigma_y)$.

Given these parameters, (x_T, y_T) have a bivariate normal distribution at time t with the pdf

$$f(x_T, y_T) = \frac{1}{2\pi\sigma_x\sigma_y\sqrt{1-\rho_{xy}^2}} \exp\left[-\frac{1}{2(1-\rho_{xy}^2)}\left\{\left(\frac{x_T - \mu_x}{\sigma_x}\right)^2 - 2\rho_{xy}\frac{x_T - \mu_x}{\sigma_x}\frac{y_T - \mu_y}{\sigma_y} + \left(\frac{y_T - \mu_y}{\sigma_y}\right)^2\right\}\right].$$

Thus, the expected profit at time t, $\pi(t, q)$, can be written as a function of the order quantity q as:

$$\pi(t,q) \quad = \quad e^{-r(T-t)}\int_{-\infty}^{\infty}\int_{-\infty}^{y_q} P_T D_T f(x_T, y_T) dy_T dx_T + e^{-r(T-t)}\int_{-\infty}^{\infty}\int_{y_q}^{\infty} P_T q f(x_T, y_T) dy_T dx_T - c_t q.$$

Here, y_q is the smallest value of y_T at which a stockout occurs. This value is defined by the condition $D_T \le q$, i.e., $aP_T^{-\eta}e^{y_T} \le q$, which implies that $y_T \le \eta x_T + \log(q/a)$, and thus, $y_q = \eta x_T + \log(q/a)$. Notice that y_q is a function of not only q, but also the price P_T.

For the interpretation of our results, it is useful to note the following formulas. $E[P_T|P_t,\epsilon_t]$ and $E[D_T|P_t,\epsilon_t]$, respectively, are the conditional expectations of price and demand at time t. $E[P_TD_T|P_t,\epsilon_t]$ is the conditional expectation of the revenue at time t if the firm carried infinite inventory and no stockout occurred.

$$\begin{aligned}
E[P_T|P_t,\epsilon_t] &= e^{\mu_x+\sigma_x^2/2}, \\
E[D_T|P_t,\epsilon_t] &= ae^{\left\{\mu_y+\sigma_y^2/2-\eta\mu_x+\eta^2\sigma_x^2/2-\eta\sigma_{xy}\right\}}, \\
E[P_TD_T|P_t,\epsilon_t] &= ae^{\left\{\mu_y+\sigma_y^2/2+(1-\eta)\mu_x+(1-\eta)^2\sigma_x^2/2+(1-\eta)\sigma_{xy}\right\}}.
\end{aligned}$$

We solve for $\pi(t,q)$ using the joint distribution of x_T and y_T, and obtain a closed-form solution shown in the following lemma. In this lemma and hereafter, $\Phi(\cdot),\overline{\Phi}(\cdot)$ and $\phi(\cdot)$ denote the cdf, the complementary cdf and the pdf, respectively, of the standard normal distribution.

Lemma 1. *The expected profit at time t is given by*

$$\pi(t,q) = e^{-r(T-t)}E[P_TD_T|P_t,\epsilon_t]\Phi(d_2)+e^{-r(T-t)}E[P_T|P_t,\epsilon_t]q\overline{\Phi}(d_1)-c_tq,$$

where

$$\begin{aligned}
d_1 &= \frac{\log(q/a)-\mu_y-\sigma_{xy}+\eta(\mu_x+\sigma_x^2)}{\sigma_z} = \frac{\log q - E[P_TD_T|P_t,\epsilon_t]/E[P_T|P_t,\epsilon_t]+\sigma_z^2/2}{\sigma_z}, \\
d_2 &= d_1-\sigma_z, \\
\sigma_z &= \sqrt{\eta^2\sigma_x^2+\sigma_y^2-2\eta\sigma_{xy}}.
\end{aligned}$$

We interpret the above expression for $\pi(t,q)$ as follows. The first term represents the expected revenue when demand is less than the inventory level, and the second term represents the expected revenue when demand is greater than the inventory level. In the first term, we see that the expected revenue from states in which demand is less than the inventory level separates into a product of the maximum expected revenue and a proportional scaling factor determined by the inventory level. In the second term, the expected revenue similarly separates into the expectation of P_T at time t, the quantity sold in the event of a stockout, and the probability of stockout.

The parameter σ_z is the standard deviation of $y_T-\eta x_T$ given information (x_t,y_t) at time t. It is related to the volatility of demand because demand is given by $ae^{y_T-\eta x_T}$. The numerators of d_1 and d_2 are functions of the parameters μ_x and σ_x of the price forecast process because the volatility of price affects the probability of stockout. Even if the price elasticity, η, is zero, the expression for σ_z simplifies, but the term σ_{xy} persists in d_1 and d_2 due to the correlation between x_T and y_T.

It can easily be shown that $\pi(t,q)$ is concave in the inventory level q. Thus, the optimal inventory level and optimal expected profit are given by the following proposition.

Proposition 1. *The optimal inventory level is given by*

$$q^*(t) \quad = \quad a \exp\left\{\mu_y + \sigma_{xy} - \eta(\mu_x + \sigma_x^2) + d_1^* \sigma_z\right\} \quad = \quad \frac{E[P_T D_T | P_t, \epsilon_t] \exp(d_1^* \sigma_z - \sigma_z^2/2)}{E[P_T | P_t, \epsilon_t]},$$

where

$$d_1^* = \Phi^{-1}\left(\max\left[0, 1 - \frac{c_t}{e^{-r(T-t)} E[P_T | P_t, \epsilon_t]}\right]\right).$$

The optimal expected profit is given by

$$\pi^*(t) \quad = \quad e^{-r(T-t)} E[P_T D_T | P_t, \epsilon_t] \Phi\left(d_1^* - \sigma_z\right).$$

Here, d_1^* is the inverse of the critical fractile similar to a newsboy formula. The critical fractile (i.e., the newsvendor fractile) depends on the drift and volatility of price, but not on the parameters of demand. In the expression for $q^*(t)$, the term $a \exp\left\{\mu_y + \sigma_{xy} - \eta(\mu_x + \sigma_x^2)\right\}$ is a scale variable used to obtain the forecast of the mean demand at time t. It is equal to the ratio $E[P_T D_T | P_t, \epsilon_t]/E[P_T | P_t, \epsilon_t]$. The remaining term $e^{d_1^* \sigma_z}$ is the safety stock factor, which depends on the critical fractile and the uncertainty of demand. The expression for the optimal expected profit can be interpreted as the product of the expected revenue under infinite inventory and a profit margin factor. The profit margin factor is given by d_1^* and a penalty for demand uncertainty σ_z.

We now compute $Y(u, t)$ for $u \leq t$ by taking the expectation of $\pi^*(t)$ with respect to x_t and y_t given the information at time u, (x_u, y_u). The function Y will help us determine the value of postponing the stocking decision from time u to a later time t.

Proposition 2. *The time u value of the expected profit if optimal inventory decision is taken at time t is given by:*

$$Y(u,t) \quad = \quad e^{-r(T-u)} E[P_T D_T | P_u, \epsilon_u] \int_{\underline{x}_t}^{\infty} \Phi(d_1^{**} - \sigma_z) \frac{1}{\sqrt{2\pi}\sigma_{xut}} \exp\left[-\frac{1}{2\sigma_{xut}^2}\left(x_t - \mu_{xut} - k\sigma_{xut}^2\right)^2\right] dx_t.$$

Here, k is a constant as defined below, $\underline{x}_t$ is the smallest value of the logarithm of price at which it becomes profitable to sell, and $\mu_{xut}, \mu_{yut}, \sigma_{xut}$ and σ_{yut} are distribution parameters of x_t and y_t computed at time u. We have

$$\begin{aligned} k &= (1-\eta)e^{-h(T-t)} + \rho \frac{\sigma_y}{\sigma_x}, \\ \underline{x}_t &= \left[\log c_t + r(T-t) - \frac{\sigma_x^2}{2} - (m - \frac{s_P^2}{2h})(1 - e^{-h(T-t)})\right] e^{h(T-t)}. \end{aligned}$$

In the expression for $Y(u.t)$, the integration is over the values of x_t that yield prices higher than the cost of procurement. The lower limit of integration, $\underline{x}_t$, represents the option to shut down, which is available to the firm if it postpones the inventory decision from time u to time t. The integration over the values of y_t gets factored out as shown in the proof of this proposition.

For use in comparative statics analysis, we rewrite the expression for $Y(u,t)$ as follows by simplifying the integration with respect to x_t:

$$Y(u,t) \quad = \quad e^{-r(T-u)}E[P_T D_T | P_u, \epsilon_u] \int_{\underline{\xi}}^{\infty} \Phi(d_1^{**} - \sigma_z)\phi(\xi)d\xi,$$

where

$$\begin{aligned} \underline{\xi} &= \left[\left\{\log c_t + r(T-t) - \frac{\sigma_x^2}{2} - (m - \frac{s_P^2}{2h})(1 - e^{-h(T-t)})\right\} e^{h(T-t)} - \mu_{xut} - k\sigma_{xut}^2\right] \frac{1}{\sigma_{xut}} \\ d_1^{**} &= \Phi^{-1}\left(\max\left[0, 1 - \frac{c_t}{\exp\{-r(T-t) + (\mu_{xut} + \sigma_{xut}\xi)e^{-h(T-t)} + (m - s_P^2/2)(1 - e^{-h(T-t)}) + \sigma_x^2/2\}}\right]\right) \end{aligned}$$

3.2 Risk averse investors

We now consider the more realistic case of risk averse investors. We show that the structure of the results in the previous section carries over with a drift adjustment to take into account the effect of risk.

Suppose that the forecasts of demand and price are correlated with the price of the market portfolio consisting of all assets in the economy. Let r_m be the expected rate of return on the market portfolio, s_m^2 be the instantaneous variance of the rate of return on the market portfolio, and $\lambda = (r_m - r)/s_m$ be the price of risk. The price of the market portfolio, m_t, evolves according to a geometric Brownian motion,

$$dm_t = r_m m_t dt + s_m m_t dz_m.$$

Let ρ_{Pm} and $\rho_{\epsilon m}$ denote the correlation coefficients of z_P and z_ϵ, respectively, with z_m. We assume that D_t, P_t and ϵ_t are not traded in the financial market.[1] Therefore, we use an equilibrium model of asset pricing to determine the risk premium that risk-averse investors in the market place on the value of the firm. In order to price risk, we use the Intertemporal Capital Asset Pricing Model (ICAPM) of Merton (1973) in our analysis, following the approach of McDonald and Siegel (1985). Any alternative equilibrium model of pricing may be used instead.

Hereafter, we shall use the expressions $\pi(t,q), q^*(t), \pi^*(t)$ and $Y(u,t)$ to refer to respective values under risk-adjusted valuation. The corresponding expressions under risk-neutral valuation, as computed in §3.1, are obtained by setting $\lambda = 0$. Suppose that $\pi(t,q)$ is traded, i.e., the firm decides to purchase q units at time t and the stock of the firm is traded in the financial market over the subsequent time period $u \in [t,T]$. Let $g_{tq}(u)$ denote the value of the firm at time u. For an investor to willingly hold this claim in his portfolio, its expected rate of return must be given by ICAPM. Using Ito's Lemma, we find that the expected rate of return satisfies the PDE

$$\begin{aligned} &\frac{\partial g_{tq}}{\partial t} + \epsilon_t \frac{\partial g_{tq}}{\partial \epsilon_t}(\alpha_\epsilon - \lambda\rho_{\epsilon m}s_\epsilon) + P_t \frac{\partial g_{tq}}{\partial P_t}(hm - h\log P_t - \lambda\rho_{Pm}s_P) \\ &\qquad + \frac{1}{2}s_\epsilon^2\epsilon_t^2\frac{\partial^2 g_{tq}}{\partial \epsilon_t^2} + \frac{1}{2}s_P^2P_t^2\frac{\partial^2 g_{tq}}{\partial P_t^2} + \rho_{P\epsilon}s_P s_\epsilon P_t\epsilon_t \frac{\partial^2 g_{tq}}{\partial P_t \partial \epsilon_t} \quad = \quad r g_{tq}. \qquad (3) \end{aligned}$$

[1] If we allow D_t and P_t to be tradeable, then the cash flows of the firm are perfectly hedgeable and can be valued at the risk-free rate by constructing a dynamic hedge using the approach of Black and Scholes.

The value of $g_{tq}(u)$ is given by the solution to this PDE with the boundary condition $g_{tq}(T) = P_T \min\{q, D_T\} - e^{r(T-t)}c_t q$. Thus, the risk-adjusted value of the firm at time t is given by $\pi(t, q) = g_{tq}(t)$. The optimal inventory level $q^*(t)$ and optimal expected profit $\pi^*(t)$ under risk-adjusted valuation can now be obtained by maximizing $g_{tq}(t)$. The value of postponement at time 0 is then obtained by solving a similar pde for $Y(u, t)$ with the boundary condition $Y(t, t) = \pi^*(t)$ at time t. The following proposition gives the solution to the risk-adjusted valuation problem.

Proposition 3. *The values of $\pi(t, q), q^*(t), \pi^*(t)$ and $Y(u, t)$ are obtained from Lemma 1 and Propositions 1 and 2 by adjusting the drift rates of the price process and the ϵ process as follows:*

(i) in (1), replace m by $m^* = m - \lambda\rho_{Pm}s_P/h$;

(ii) in (2), replace α_ϵ by $\alpha^*_\epsilon = \alpha_\epsilon - \lambda\rho_{\epsilon m}s_\epsilon$.

Proof. We obtain the expressions for $\pi(t, q), q^*(t), \pi^*(t)$ and $Y(u, t)$ by adjusting drift rates as above. With some algebraic manipulation, it can be seen that these expressions solve the pde with the respective boundary conditions for each function.

Proposition 3 shows that the expected profit under risk-adjusted valuation is similar to that obtained in §3.1 under risk-neutral valuation of the newsboy model with stochastic price and demand. The point of difference between these models is that $\pi(t, q)$ is computed using risk-adjusted drift rates taking into account the risk premium on the demand forecast and the risk premium on the price forecast. Moreover, when we substitute risk-adjusted drift rates in the expression for q^* in Proposition 1, we find that the optimal inventory level depends on the risk premium. Interestingly, unlike the known results on valuation under the preference based approach for stochastic demand and constant prices (please see references on risk-aversion in §1), the optimal inventory level and expected profit need not decrease with increase in λ.

3.3 Limiting Cases

The expressions for the optimal inventory level and the expected profit stated in Propositions 1-3 yield three limiting cases depending on the values of the mean-reversion parameter h for price and the price elasticity of demand η. We examine these cases because they fit different application contexts and are useful for studying the effects of demand volatility, price volatility, procurement cost, price elasticity of demand and price of risk on the optimal inventory, expected profit and the value of postponement. In each case, we provide two sets of formulas, one for the optimal expected profit, $\pi^*(u)$, at time u if the inventory decision is taken at time u, and another for the optimal expected profit, $Y(u, t)$, at time u if the inventory decision is postponed to a later time t. These two sets of formulas shall be useful in §4 in analyzing the value of postponement at time u and in §5 for empirical analysis.

Case 1: Constant prices. Let h tend to infinity. In the limit, μ_x tends to a constant value m and σ_x and σ_{xy} tend to zero. Therefore, this limiting case gives a model with constant price, e^m, and stochastic demand. The resulting formulas for $\pi^*(u)$ and $Y(u,t)$ are as follows.

$$\pi^*(u) = e^{-r(T-u)} D_u e^{\alpha_\epsilon^*(T-u)+m} \Phi\left(d_1^* - s_\epsilon\sqrt{T-u}\right) \tag{4}$$

$$Y(u,t) = e^{-r(T-u)} D_u e^{\alpha_\epsilon^*(T-u)+m} \Phi\left(d_1^{**} - s_\epsilon\sqrt{T-t}\right) \tag{5}$$

Here and in the rest of the paper, we denote the newsvendor fractiles for inventory decisions at times u and t as d_1^* and d_1^{**}, respectively. In Case 1, we have

$$\begin{aligned} d_1^* &= \Phi^{-1}\left(\max\left[0, 1 - \frac{c_u}{e^{-r(T-u)}e^m}\right]\right) \\ d_1^{**} &= \Phi^{-1}\left(\max\left[0, 1 - \frac{c_t}{e^{-r(T-t)}e^m}\right]\right) \end{aligned}$$

Case 2: GBM prices with $\eta > 0$. Let h tend to 0 such that hm^* stays constant at a value denoted α_P^*.[2] In this limiting case, the price process is a geometric Brownian motion. Thus, our model reduces to one in which price and ϵ forecasts follow correlated GBM processes with a constant correlation coefficient, and demand is price elastic. The resulting formulas are given as:

$$\pi^*(u) = e^{-r(T-u)} P_u D_u e^{[(1-\eta)\alpha_P^* - \eta(1-\eta)s_p^2/2 + (1-\eta)\rho s_P s_\epsilon + \alpha_\epsilon^*](T-u)} \Phi\left(d_1^* - \sigma_z\sqrt{\frac{T-u}{T-t}}\right) \tag{6}$$

$$Y(u,t) = e^{-r(T-u)} P_u D_u e^{[(1-\eta)\alpha_P^* - \eta(1-\eta)s_p^2/2 + (1-\eta)\rho s_P s_\epsilon + \alpha_\epsilon^*](T-u)} \int_{\underline{\xi}}^{\infty} \Phi(d_1^{**} - \sigma_z)\phi(\xi)d\xi \tag{7}$$

where

$$\begin{aligned} \sigma_z^2 &= (\eta^2 s_P^2 + s_\epsilon^2 - 2\eta\rho s_P s_\epsilon)(T-t) \\ d_1^* &= \Phi^{-1}\left(\max\left[0, 1 - \frac{c_u}{e^{(-r+\alpha_P^*)(T-u)}P_u}\right]\right) \\ d_1^{**} &= \Phi^{-1}\left(\max\left[0, 1 - \frac{c_t}{\exp[(-r+\alpha_P^*)(T-u) + \{r + s_P^2/2 + \rho s_P s_\epsilon - \eta s_P^2\}(t-u) + \xi s_P\sqrt{t-u}]P_u}\right]\right) \\ \underline{\xi} &= \frac{1}{s_P\sqrt{t-u}}\left[\log(c_t/P_u) + (r - \alpha_P^*)(T-u) - \{r + s_P^2/2 + \rho s_P s_\epsilon - \eta s_P^2\}(t-u)\right] \end{aligned}$$

Case 3: GBM prices with $\eta = 0$. This case is a simplification of Case 2 obtained by setting the price elasticity of demand to 0. Thus, our model reduces to one in which demand and price forecasts follow correlated GBM processes with a constant correlation coefficient but no price elasticity of demand. We have

$$\pi^*(u) = e^{-r(T-u)} P_u D_u e^{(\alpha_P^* + \alpha_\epsilon^* + \rho s_P s_\epsilon)(T-u)} \Phi\left(d_1^* - s_\epsilon\sqrt{T-u}\right) \tag{8}$$

$$Y(u,t) = e^{-r(T-u)} P_u D_u e^{(\alpha_P^* + \alpha_\epsilon^* + \rho s_P s_\epsilon)(T-u)} \int_{\underline{\xi}}^{\infty} \Phi(d_1^{**} - s_\epsilon\sqrt{T-t})\phi(\xi)d\xi \tag{9}$$

[2]The condition that hm^* stay constant ensures that the drift rate of the price process is non-zero. It is not a necessary condition for the convergence of the price process to a GBM. If, instead, m^* stays constant as h tends to 0, we get a GBM with zero drift.

where

$$
\begin{aligned}
d_1^* &= \Phi^{-1}\left(\max\left[0, 1 - \frac{c_u}{e^{(-r+\alpha_P^*)(T-u)}P_u}\right]\right) \\
d_1^{**} &= \Phi^{-1}\left(\max\left[0, 1 - \frac{c_t}{\exp[(-r+\alpha_P^*)(T-u) + \{r + s_P^2/2 + \rho s_P s_\epsilon\}(t-u) + \xi s_P\sqrt{t-u}]P_u}\right]\right) \\
\underline{\xi} &= \frac{1}{s_P\sqrt{t-u}}\left[\log(c_t/P_u) + (r - \alpha_P^*)(T-u) - \{r + s_P^2/2 + \rho s_P s_\epsilon\}(t-u)\right]
\end{aligned}
$$

3.4 Comparison with American Options in a financial market

The analogy between our model and the financial option valuation theory is methodologically useful. Suppose that the firm buys inventory q at time t. Then the claim on the future revenues of the firm, $V(t, q)$, is equivalent to a European option on the demand process at exercise price q. Further, prior to time t, when the firm has to decide when to buy inventory, the firm's problem is equivalent to the optimal timing of the decision to acquire a European option. This problem can be considered to be an American option on the newsvendor payoffs, and $Y(u, t)$ is equal to the value of this American option. In particular, in our setting, there is an exercise decision in the following sense: postponing the inventory decision is similar to not exercising the option early, while making the inventory investment early is like exercising the option. Moreover, the intuition in our case can also be translated into the usual tradeoff between the insurance value of the option and the adjustment to the drift term.

However, there are significant differences between the option modeled by us and American-style financial options modeled in the academic and practitioner finance literature. Since the payoffs to the firm occur at time T and not at the time of exercising the option, the usual intuition about option being in the money or out of money does not apply. Moreover, in our case, the distribution of the underlying asset's payoffs *can* be affected by the decision maker by choosing a suitable inventory level. Thus, the decision maker is making two decisions, one the "real" decision that affects the payoffs and two, the optimal "exercise" decision based on the evolution of the price and demand processes. Yet another difference has to do with the way information is used not just to exercise options but also to take the optimal stocking decision.

4 Sensitivity Analysis and Value of Postponement

Our model answers a number of questions that are useful to suppliers and retailers. Using this model, a retailer can determine its optimal timing decision, i.e., whether to purchase inventory now or to postpone, and in what quantity to purchase inventory. A supplier can determine the lead-times and corresponding procurement costs to offer to retailers such that they will be economically viable options. We address three types of questions related to these issues: (i) how does the stochasticity of price affect the retailer's postponement decision; (ii) how does the cost of capital (in other words, the risk premium) affect the postponement decision; (iii) at what rate can the procurement cost increase over a horizon $[0, t]$ such that postponement remains optimal over it.

We compare the expressions for $\pi^*(u)$ and $Y(u,t)$ to make the choice between purchasing inventory at time u and postponing the purchase of inventory to a fixed future time t. In Case 1, i.e., with constant price, it follows from equations (4) and (5) that the firm should postpone the inventory decision from time u to a future time $t \in (u, T]$ if

$$\Phi\left(d_1^{**} - s_\epsilon\sqrt{T-t}\right) \geq \Phi\left(d_1^* - s_\epsilon\sqrt{T-u}\right), \tag{10}$$

otherwise it should make the inventory decision at time u. Due to the monotonicity of the function Φ, the comparison is between the arguments. Thus, postponement is optimal if

$$d_1^{**} - s_\epsilon\sqrt{T-t} \geq d_1^* - s_\epsilon\sqrt{T-u}. \tag{11}$$

Equation (10) gives all the factors that determine the postponement decision under constant price, namely, the costs c_t and c_u, and the demand volatility s_ϵ.

In Case 2, i.e., when the price forecast follows a GBM, the expressions for $\pi^*(u)$ and $Y(u,t)$ in (6) and (7) show that there is value in postponing the inventory decision from time u to time t if the following inequality holds:

$$\int_{\underline{\xi}}^{\infty} \Phi(d_1^{**} - \sigma\sqrt{T-t})\phi(\xi)d\xi \;\; \geq \;\; \Phi\left(d_1^* - \sigma\sqrt{T-u}\right) \tag{12}$$

In this case, the postponement decision now depends on not only the costs and s_ϵ, but also the price volatility parameters s_P and ρ and the price elasticity η.

The optimal choice between postponement and early purchase of inventory depends on the tradeoff between increase in unit purchase cost and the value of information. By committing to the inventory decision at time u, the firm incurs a possibly lower unit purchase cost of c_u compared to c_t at time t. However, the firm uses poorer forecasts of price and demand compared to those at time t. The value of information has three different sources. First, the expression for $Y(u,t)$ includes the option not to stock any product if the price P_t falls below a risk-adjusted cost $c_t e^{(r-\alpha_P^*)(T-t)}$. We call this the optionality effect (O). Second, the optimization of inventory at time u uses an updated newsvendor fractile d_1^{**} compared to that at time u. We call this the newsvendor effect (NV). Finally, the safety stock required at time t is smaller than that at time u. We call this the volatility effect (Vol).

Apart from providing answers to the three questions posed above, equations (11) and (12) show that the drift rate α_ϵ and the information about ϵ_u have no effect on the postponement decision at time u. Under constant prices, (11) implies that the optimal timing decision is path independent. Thus, the optimal time to procure inventory is independent of time u and can be fixed in advance. Under stochastic prices, (12) implies that the optimal postponement decision at time u is independent of the sample path of ϵ_u, but does depend on the value of P_u. Thus, the optimal time to procure inventory cannot be fixed in advance.

In §4.1, we analyze the effect of price volatility on the value of postponement. For this, we assume that $\lambda = 0$ and $c_t = c_u$ to avoid mixing the effects of risk premium and change in cost. In §4.2 and §4.3, we analyze the effects of risk premium and change in procurement cost, respectively, on the value of postponement.

4.1 Effect of price volatility on postponement

From (12), price volatility, s_P, affects the postponement decision through d_1^{**} and σ_z. Both the NV and Vol effects of s_P are non-linear due to the presence of quadratic terms in s_P; they are further modulated by price elasticity and the correlation between price and ϵ.

When the price elasticity η is zero, only the NV effect matters. Figure 1 shows for this case how the value of postponement varies as a function of s_P and ρ. Observe that the value of postponement is increasing in ρ but is non-monotone in s_P. When ρ is non-negative, the value of postponement increases with s_P. When ρ is negative, the value of postponement decreases with s_P for small values of s_P but increases with s_P for large values of s_P. Thus, we find that it is not always optimal to postpone the inventory decision when prices are stochastic.

The intuition for the above effects of price volatility on the value of postponement is as follows. First, as s_P increases, the expected price increases because of the term $e^{s_P^2/2}$. This causes the newsvendor fractile, and correspondingly, the optimal profit and inventory to increase. In simple terms, this term captures the ability of the retailer to respond to price changes by adjusting the newsvendor fractile to the updated price forecast. Second, as s_P increases, the term $\rho s_P s_\epsilon$ increases or decreases depending on whether ρ is positive or negative. This term captures the natural hedge available to the retailer if prices move in a direction opposite to that of demand, and it captures the enhanced value of price updates if prices move in the same direction as demand. If ϵ and price are negatively correlated, then the product of random effects due to ϵ and price is less variable, so that d_1^{**} decreases. Therefore, there is less value to be had from postponement by exploiting the new information. If ϵ and price are positively correlated, then the price and demand information affect payoffs in different states in the same direction. Therefore, postponement becomes more valuable.

When the price elasticity η is non-zero, both the NV and Vol effects matter. Figure 2 shows how the value of postponement varies as a function of η and s_P when $\rho = 0.5$. Observe that for each value of s_P, the value of postponement is decreasing in η when η is small and is increasing in η when η is large. To see the intuition for this result, note that d_1^{**} decreases with η through the linear term $-\eta s_P^2$ whereas σ_z increases with η through the quadratic term $\eta^2 s_P^2$. When η is small, the linear effect dominates the quadratic effect leading to a decrease in the value of postponement, but when η is large, the quadratic effect dominates leading to an increase in the value of postponement. We also note that the non-monotonicity in η becomes more prominent when ρ increases, but disappears when ρ is negative with a sufficiently large magnitude.

4.2 Effect of risk premium on postponement

In a somewhat counter-intuitive result, we find that the risk premium has no effect on the postponement decision if the selling price is constant. Notice from (11) that the price of risk (λ) and the correlation between the demand forecast and the market portfolio ($\rho_{\epsilon m}$) are both absent in the inequality. This implies that if the supplier quotes the same cost function c_t to two different buyers with differing risk premia, then the buyers will take identical postponement decisions! This result follows from the fact that the optimal timing decision is independent of the sample path. Thus, when the risk premium affects the discount rate from time T to time u, it has the same proportional effect on $Y(u,t)$ and $\pi^*(u)$, and therefore, has no bearing on the postponement decision.

In Cases 2 and 3, when price follows a GBM, the value of postponement depends on the risk premium of price but not on the risk premium of demand, since d_1^{**} is a function of α_P^* but not of α_ϵ^*. Our numerical analysis shows that the value of postponement increases as the risk premium of price increases. This implies that postponement is more valuable when price is more positively correlated with the market. It is useful to relate the correlation of price with the market with the financial hedging role of the firm's business. When price is negatively correlated with the market, then the firm's business provides a financial hedging role and thus, the investors use a negative risk premium to value its stock. In this case, the optimal profit is higher. However, the value of postponement is diminished since the investors are willing to take a greater price risk. When price is positively correlated with the market, then the investors use a positive risk premium to value its stock. In this case, the optimal profit is lower, but the value of postponement is enhanced since the investors are willing to take a lower price risk.

4.3 Effect of procurement cost

Equations (11) and (12) enable us to compute the maximum possible rate of increase in cost under which postponement is the optimal strategy. We call this rate as the *cost trajectory.*

Under constant prices, i.e., in Case 1, we obtain the following proposition giving sufficient conditions for the effect of procurement cost on the postponement decision.

Proposition 4. *Under the assumption that price* $= e^m > c_t e^{r(T-t)} \quad \forall t$*,*

(i) If $c_t < c_u e^{r(t-u)}$ *for any* $t > u$*, then postponement from time* u *to time* t *increases the expected profit.*

(ii) If c_t *increases exponentially at a rate higher than* r*, i.e.,* $c_t = c_u e^{(r+\delta)(t-u)}$ *for some* $\delta > 0$*, then postponement from time* u *to* T *is an optimal strategy if* s_ϵ *is sufficiently large.*

Proposition 4 shows that postponement is the optimal strategy at time u if the cost of procurement increases at the risk-free rate. It also indicates that when the volatility of demand s_ϵ increases, it shifts the cost trajectory such that postponement from u to t remains optimal even if the cost increases slightly faster than the risk-free rate. Thus, the proposition provides another way to look at the value of postponement as a function of various parameters. It implies that we can assess the value of postponement by identifying how the cost trajectory shifts with a change in the parameters. While Proposition 4 gives upper bounds on the rate of increase in cost, we can numerically compute the exact cost trajectory using (10). Figure 3 illustrates an example of this trajectory for different values of s_ϵ.

Under GBM prices, i.e., in Case 2, the following proposition characterizes the cost trajectory. The results for Case 3 follow by setting $\eta = 0$ in this proposition.

Proposition 5. *If* $c_t \leq c_u e^{(r+\rho s_P s_\epsilon - \eta s_P^2)(t-u)}$ *for* $t > u$*, then postponement from time* u *to* t *is an optimal strategy.*

The new insight from this proposition is that when the price is stochastic, the cost trajectory depends on all the parameters of the demand and price forecast processes, and it can increase at a rate faster or slower than the rate in Case 1 depending on these parameters. For example, we observe from Proposition 5 that if the price elasticity of demand is zero and the correlation between ϵ and price is positive, then postponement is an optimal strategy even if cost grows slightly faster than the risk-free rate. The result is the opposite if the correlation between ϵ and price is negative. This result is consistent with the effect of ρ discussed in §4.1.

In summary, we have shown the effects of price volatility, cost of capital and growth rate of procurement cost on the value of postponement. The equations for profit and optimal inventory also lend themselves to empirical testing as discussed in §5 below.

5 Empirical Illustration

In this section, we illustrate the insights from our model using data for gold mining firms. We show that the inventory holding measured by days of inventory of gold firms is positively correlated both with gross margin (as a proportion of sales revenue) and price volatility of gold. These results provide interesting new evidence on the effect of price volatility on inventory levels, which substantiates our model. The regression model used in this section is constructed from the formulas derived in §3. Thus, this section demonstrates the applicability of our model for empirical analysis. It should be noted that the empirical analysis presented in this paper is meant to be only an illustration. As such, we do not offer the detailed econometric tests that would be carried out in a fuller investigation.[3]

We select firms in the gold mining industry listed on the U.S. stock exchanges for the analysis supporting our illustration. The primary reason for this choice is that our model closely fits the metal extraction decision of a gold mining firm, which is a classical inventory timing decision on the amount of gold to extract, given extraction lead times, costs, and forecasts of demand and price. Moreover, gold mining firms are convenient for empirical purposes because gold is a homogeneous good and is traded actively in global markets. Thus, we can readily obtain historical data on gold futures prices, which can be used to estimate parameters of price volatility. In addition, we can also readily obtain quarterly historical data on the inventory of gold stockpiles, profit margin and realized sales from the financial statements of firms.

Standard & Poor's Compustat database, accessed through Wharton Research Data Services (WRDS), yielded 42 firms in the NAICS category 212221 (gold ore mining) for the period 1995-2005. Of these, 22 firms had quarterly data available for 4 or more years and a common fiscal year ending date of Dec 31.

[3] A larger study using broad-based and detailed data may be conducted to examine the effects of price and demand volatility on inventory levels and profits. However, this would be the subject of a different paper.

We avoided using data from firms with fewer that 4 years of observations in order to ensure that we have a reasonable time-series for each firm. We also excluded firms with different fiscal year end dates to avoid biases due to differences in year-end window-dressing of inventory. After using these filters, we obtain 633 observations at the firm-quarter level for 22 firms aligned by calendar quarters. For each firm i in quarterly time period t, let S_{it} be the sales revenue, CS_{it} the cost of goods sold, GM_{it} the gross margin, I_{it} denote the average inventory, and DI_{it} the number of days of inventory. We compute gross margin as the ratio $1 - CS_{it}/S_{it}$, average inventory as the average of opening and closing inventories (in dollars) for the quarter, and days of inventory as the number of days in the quarter $\times$ I_{it}/CS_{it}. Table 1 shows summary statistics for each firm for sales, gross profit, days of inventory and gross margin. We note that days of inventory and gross margin exhibited substantial variation during the sample period, for each firm. For example, Campbell Resources Inc. (stock ticker: 3CBLRF), a typical firm, has days of inventory varying between 43.4 and 109.6, and gross margin varying between 2% and 45% across our sample period. Figure 4 shows a time-series plot of the trend in days of inventory across all firms during 1995-2005. We find that days of inventory has an increasing time trend during this period, statistically significant at $p < 0.01$. But the trend is not uniform: it declines during the sub-periods 1998-2001 and 2004-2005.

We obtain data for daily closing prices of the current gold futures contract traded on the COMEX section of the New York Mercantile Exchange (NYMEX) from Datastream. We estimate the price volatility of gold s_P separately for each calendar quarter in order to allow it to vary across time. To compute an estimate of price volatility, $\hat{s}_{Pt}$, for each quarter t, we first compute the daily return on gold prices and then find its standard deviation across all the days in that quarter to get the daily volatility of gold prices for the quarter. Since we use data for a short time period of 11 years, we ignore mean reversion and treat gold prices as following a GBM.[4] Figure 5 shows plots of the quarterly average price and price volatility of gold during 1995-2005. The price varied in the range of $252.8 (Aug 25, 1999) to $529.3 (Dec 12, 2005) with an overall mean of $366.58. Price volatility varies between 4.0% and 23.9% with an upward trend during 1995-2005, statistically significant at $p < 0.01$. Gold price also increased from 1995 to 2005, but does not have a uniform trend - it decreased sharply from 1996 to 2000, and then increased subsequently.

Figures 4-5 suggest that the days of inventory of gold might be correlated with gold price and gold price volatility. **Vishal, I am not sure if one can see this, since the figures are plotted separately. Either we should plot all the variables on the same graph, or we should modify this statement suitably. I would prefer the former, with different colors.** We use our model to conduct a more formal regression analysis of these variables. For example, Figures 4-6 further suggest that the correlations among these variables might not be distinguishable from time trends. Hence, we need to adjust the estimation results for time trends to draw such a conclusion.

From Proposition 1, the optimal inventory level for inventory procured at time t to be sold at time T is

[4] We did not employ more sophisticated models, such as GARCH, to account for time-varying volatility, since the analysis here is meant only for illustrative purposes.

given by

$$q^*(t) = \frac{E[P_T D_T | P_t, \epsilon_t]}{E[P_T | P_t, \epsilon_t]} \exp\left(d_1^* \sigma_z - \frac{\sigma_z^2}{2}\right). \tag{13}$$

Or,

$$\frac{q^*(t) E[P_T | P_t, \epsilon_t]}{E[P_T D_T | P_t, \epsilon_t]} = \exp\left(d_1^* \sigma_z - \frac{\sigma_z^2}{2}\right).$$

We use days of inventory, DI_{it} as a proxy for the left hand side of the above equation for each firm i in quarter t. The reason is that the numerator of this term is the value of the inventory, scaled by the number of days in the year, and the denominator is the expected revenue for the year. For the right hand side, we construct a linear approximation as follows:

$$\begin{aligned} d_1^* \sigma_z - \frac{\sigma_z^2}{2} &= d_1^* \sqrt{\left(s_\epsilon^2 + \eta^2 s_P^2 - 2\eta\rho s_P s_\epsilon\right)(T-t)} - \frac{1}{2}\left(s_\epsilon^2 + \eta^2 s_P^2 - 2\eta\rho s_P s_\epsilon\right)(T-t) \\ &\approx d_1^* \sqrt{T-t}\, |s_\epsilon - \eta s_P| - \frac{1}{2}\left(s_\epsilon^2 + \eta^2 s_P^2 - 2\eta\rho s_P s_\epsilon\right)(T-t). \end{aligned}$$

We use $\hat{s}_{Pt}$ as the estimator of s_P in time period t. Based on the definition of d_1^*, we use $\Phi^{-1}(GM_{it})$ as a proxy for d_1^* for firm i in quarter t. For purposes of this illustration, we assume that s_ϵ, η, ρ and $T-t$ do not vary across firms or over time. These assumptions for s_ϵ, η and ρ are reasonable because all gold mining firms face the same industry demand, since gold is a homogenous good, without any distinctive characteristics that vary across firms. Also, the characteristics of this demand might not vary over the relatively short period of our data set. The assumption regarding $T-t$ implies that firms have a constant lead-time over the period of study and that they vary cross-sectionally only in terms of the level of safety stock. Under these assumptions, the values of s_ϵ, η, ρ and $T-t$ can be absorbed in the coefficients of the regression model. Thus, we obtain the following regression equation:

$$\log DI_{it} = a_i + b_1 \Phi^{-1}(GM_{it}) + b_2 \hat{s}_{Pt} + b_3 \Phi^{-1}(GM_{it}) \hat{s}_{Pt} + b_4 \hat{s}_{Pt}^2 + \nu_{it}. \tag{14}$$

Here, a_i denotes the fixed effect for firm i, ν_{it} denotes the error term, and the coefficients $b_1, \ldots, b_4$ are assumed to be identical across firms and time. We employ firm-wise fixed effects to control for unobserved variables that may vary across firms but not over time.

Our regression model (14) merits another comparison with the formula (13). Suppose that we had data on σ_z. Then, theoretically, we could directly estimate (13) upon taking logarithms. However, σ_z includes the effects of s_ϵ, price volatility and $T-t$. Thus, a regression based on σ_z does not help assess the effect of price volatility on days of inventory. The choice of explanatory variables in the regression model (14) achieve this objective.

The first column in Table 2 shows the estimation results for (14) obtained from generalized least squares estimation controlled for heteroscedasticity across firms. We observe that the coefficients of $\Phi^{-1}(GM_{it}), \hat{s}_{Pt}$ and $\hat{s}_{Pt}^2$ are estimated as 0.256, 65.438 and -3201.56, respectively. All three coefficients are statistically significant at the 95% confidence level. The coefficient of the interaction term, $\Phi^{-1}(GM_{it})\hat{s}_{Pt}$, is -9.882, but is not statistically significant.

Thus, days of inventory are increasing in gross margin. This result is consistent with the news vendor model, with our theoretical model and with evidence obtained in prior empirical research on other industries (see, for example, Gaur et al. (2005) and Roumianetsev and Netessine (2006)). Days of inventory also has strong correlations with linear and quadratic terms in price volatility, $\hat{s}_{Pt}$. It increases with $\hat{s}_{Pt}$ for small values of $\hat{s}_{Pt}$ due to the effect of the correlation, ρ, between the price and demand forecast processes. But it decreases for large values of $\hat{s}_{Pt}$ because the price elasticity of demand and the lognormal distribution cause expected demand to decrease in $\hat{s}^2_{Pt}$. The signs of these coefficients are consistent with our model if the correlation coefficient between price and ϵ, ρ, is greater than 0 and s_ϵ is larger than $\eta\hat{s}_{Pt}$.

As noted earlier, days of inventory, gold price and gold price volatility are all trending upwards with time. Furthermore, $\Phi^{-1}(GM_{it})$ is also trending upwards with time, which is expected because price is trending upwards and $\Phi^{-1}(GM_{it})$ is a function of price. To control for time trends, we estimate the following augmented model:

$$\log DI_{it} = a_i + b_1\Phi^{-1}(GM_{it}) + b_2\hat{s}_{Pt} + b_3\Phi^{-1}(GM_{it})\hat{s}_{Pt} + b_4\hat{s}^2_{Pt} + b_5 t + \nu_{it}. \tag{15}$$

This model includes all the terms in (14) plus a linear time trend, $b_5 t$.

The second column in Table 2 shows estimation results for (15). We find that days of inventory has statistically significant positive correlation with d_1^* even after controlling for time trend. Thus, we are able to disentangle the effect of d_1^* on days of inventory from time trends. Note that d_1^* and time are poorly correlated. In particular, gold price declined during 1996-2000 and increased thereafter. Correspondingly, days of inventory declined during an overlapping period 1998-2001. Thus, our result provides evidence that firms adjusted their inventories in response to changes in gross margin independently of time trends.

We also find that $\hat{s}_{Pt}$ and $\hat{s}^2_{Pt}$ have coefficients with the same sign as in (14), but lose statistical significance when t is added to the model. Furthermore, t is not statistically significant even though the R^2 of the model increases. This suggests collinearity between time and price volatility. We are unable to conclude whether the changes in days of inventory are associated with time trend or price volatility. Conservatively, we might conclude that price volatility has no effect on days of inventory, but some other variable correlated with time does. This explanation could have been a viable alternative if days of inventory were decreasing with time because improvements in technology and processes can cause days of inventory to decline over time. However, since days of inventory are increasing with time, we cannot reasonably imagine an alternative explanation for our results. Therefore, we believe that the effect might be attributed to price volatility.

In summary, this section shows that our model can be estimated empirically and is a reasonable fit for observed data in the gold mining industry. It shows that days of inventory of gold mining firms is correlated with price volatility as well as gross margin. The effect of price volatility has not been investigated in the literature thus far for any industry. In addition, our results for gross margin are obtained using a Φ^{-1} transformation of gross margin derived in our analytical model, whereas similar results in the existing literature are obtained using a log transformation. We tested a log transformation as well and obtained results consistent with those from a Φ^{-1} transformation. However, a Φ^{-1} transformation is theoretically more appealing because it conforms with our analytical model and it maps values of gross margin from $(0,1)$ to the real line whereas a log transformation's range is restricted to the negative real line.

The analysis of this section might be extended to other industries. It might be improved by investigating differences in coefficients across firms or by relaxing the assumptions on s_ϵ, η, ρ and $T-t$. In particular, s_ϵ and $T-t$ might vary across firms and time. They might influence the effect of price volatility on days of inventory.

6 Conclusions

The model presented in this paper combines insights from two distinct literatures on inventory timing decisions and real options. The former literature has recognized that demand uncertainty affects the optimal inventory timing decision, while it has largely ignored the consequences of price uncertainty, as well as the correlation between the demand and price variables, due to a shock to the economy. We explicitly consider and model the price uncertainty and correlation variables. We use the real options methodology to characterize the optimal inventory timing decision as the exercise decision for a complex, bivariate, American-style option. As is common in much of the recent real options literature, we explicitly recognize the non-hedgeability of the state variables in our framework, and employ an explicit adjustment for the market price of risk.

Although the option we identify bears some resemblance to an American-style "quanto" option, there are important differences between them in terms of the timing of the option payoffs due to exercise, the endogeneity of the process generating quantity, and hence profits, due to the optimizing decision being undertaken, and the role of the information state variable. Consequently, the comparative statics analysis of our "option" price is quite complex. The main result here is that the value of postponement can increase or decrease with changes in price volatility depending on the magnitude of the price elasticity of demand, the correlation between the price and demand forecast processes, and the correlations of the price forecast process with the financial market. We find that the correlation of the demand forecast process with the financial market has no effect on the optimal inventory timing decision. The formulas obtained by us yield a cost trajectory that can be used to characterize the inventory timing decision for a firm.

Our modeling of the optimal inventory decision is admittedly stylized in several respects. For example, the inventory stocking decision is made only once during the time period in question. We do not consider inventory carrying costs. We also use a single period inventory model. However, we believe that we have identified the principal intuition of the uncertainty in the price and demand variables, creating an optionality for the decision maker, which has value. The modeling of this option and the optimization of its value is the novelty of our approach in this paper.

The empirical illustration presented in this paper shows that the model has promise. It is applied here to an industry with a high degree of homogeneity in the product characteristics and transparency in terms of price. Our application shows that the model fits the real data with statistical significance and provides new insights on the impact of price volatility on days of inventory. The model can be extended to other industries with non-homogeneous products such as retailing or consumer goods manufacturing. It can also be extended to provide insights on differences in lead times across firms.

References

Agrawal, V., S. Seshadri. 2000a. Effect of Risk Aversion on Pricing and Order Quantity Decisions. *M&SOM.* 2(4) 410-423.

Agrawal, V., S. Seshadri. 2000b. Risk intermediation in supply chains. *IIE TRANSACTIONS* 32(9) 819-831.

Akella, R., V. Araman, J. Kleinknecht. 2002. B2B Markets: Procurement and Supplier Risk Management. Book Chapter in *Supply Chain Management- Applications and Algorithms* edited by J. Geunes, H. E. Romeijn, P. Kouvelis, Kluwer Academic Publishers.

Aviv, Y., A. Federgruen. 2001. Design for Postponement: A Comprehensive Characterization of its Benefits Under Unknown Demand Distributions. *Operations Research* 49(4) 578-598.

Bernardo, A. E., O. Ledoit. 2000. Gain, Loss and Asset Pricing. *Journal of Political Economy.* 108 144-172.

Birge, J. 2000. Option Methods for Incorporating Risk into Linear Capacity Planning Models. *M&SOM.* 2(1) 19-31.

Birge, J., X. Xu. 2005. Operational Decisions, Capital Structure, and Managerial Compensation: A Newsvendor Perspective. Working Paper, Northwestern Univ.

Caldentey, R., M. Haugh. 2005. Optimal Control and Hedging of Operations in the Presence of Financial Markets. *Mathematics of Operations Research,* (forthcoming).

Chen, F., A. Federgruen. 2000. Mean-Variance Analysis of Basic Inventory Models. Working Paper, Graduate School of Business, Columbia University.

Cochrane, J. H., J. Saa-Requejo. 2000. Beyond Arbitrage: Good-Deal Asset Pricing Bounds in Incomplete Markets. *Journal of Political Economy.* 108 79-119.

Constantinides, G. M. 1978. Market Risk Adjustment in Project Valuation. *Journal of Finance.* 33 603-616.

Dong, L., H. Liu. 2005. Equilibrium Forward Contracts on Nonstorable Commodities in the Presence of Market Power. Working paper, Washington Univ., St. Louis.

Eeckhoudt, L., C. Gollier, H. Schlesinger. 1995. The Risk-averse (and Prudent) Newsboy. *Management Sci.* 41(5) 786-794.

Eppen, G., A. Iyer. 1997. Improved fashion buying with Bayesian updates. *Operations Research.* 45(6) 805-819.

Feitzinger, E., H. L. Lee. 1997. Mass Customization at Hewlett Packard - The Power of Postponement. *Harvard Business Review,* Jan-Feb.

Fisher, M. L., A. Raman. 1996. Reducing the cost of demand uncertainty through accurate response to early sales. *Operations Research.* 44(1) 87-99.

Gaur, V., M. L. Fisher, A. Raman. 2005. An Econometric Analysis of Inventory Turnover Performance in Retail Services. *Management Sci.* 51(2) 181-194.

Gaur, V., S. Seshadri. 2005. Hedging Inventory Risk through Market Instruments. *M&SOM.* 7(2) 103-120.

Gaur, V., S. Seshadri, M. Subrahmanyam. 2005. Intermediation and Value Creation in an Incomplete Market: Implications for Securitization. Working Paper, NYU Stern School of Business.

Hakzos, C., S. Seshadri. 2006. Supply Chain Operations in the Presence of Spot Market: A Review with Discussion. Journal of Operational Research Society forthcoming.

Hansen, L. P., R. Jagannathan. 1991. Implications of Security Market Data for Models of Dynamic Economies. *Journal of Political Economy.* 99, 225-262.

Huchzermeier, A., M. Cohen. 1996. Valuing Operational Flexibility Under Exchange Rate Risk. *Operations Research* 44(1) 100-113.

Lee, H., C. S. Tang. 1997. Modeling the Costs and Benefits of Delayed Product Differentiation. *Management Science.* 43(1) 40-53.

Levy, H. 1985. Upper and Lower Bounds of Put and Call Option Value: Stochastic Dominance Approach. *Journal of Finance.* 40, 1197-1217.

Kaminsky, P., J. M. Swaminathan. 2001. Utilizing Forecast Band Refinement for Capacitated Production Planning. *M&SOM.* 3(1) 68-81.

Kaminsky, P., J. M. Swaminathan. 2004. Effective Heuristics for Capacitated Production Planning with Multi-Period Production and Demand with Forecast Band Refinement. *M&SOM.* 6(2) 184-194.

Kogut, B., N. Kulatilaka. 1994. Operating Flexibility, Global Manufacturing, and the Option Value of Multinational Network. *Management Sci.* 40(1) 123-139.

Kouvelis, P. 1999. Global Sourcing Strategies Under Exchange Rate Uncertainty, Chapter 20 in *Quantitative Models For Supply Chain Management.* Eds. S. Tayur, R. Ganeshan, M. Magazine. Kluwer Academic Publishers, Boston.

Mathur, K., P. Ritchken. 1999. Minimum Option Prices under Decreasing Absolute Risk Aversion. *Review of Derivative Research.* 3, 135-156.

McDonald, R. L., D. R. Siegel. 1985. Investment and the valuation of firms when there is an option to shut down. *International Economic Review.* 26(2) 331-349.

Murray, G., E. Silver. 1966. A Bayesian analysis of the style goods inventory problem. *Management Sci.* 12(11) 785–797.

Perrakis, S., P. Ryan. 1984. Option Pricing in Discrete Time. *Journal of Finance.* 39, 519-525.

Petruzzi, N. C., M. Dada. 1999. Pricing and the newsvendor problem: A review with extensions. *Oper. Res.* 47(2) 183194.

Ritchken, P. 1985. On Option Pricing Bounds. *Journal of Finance.* 40, 1219-1233.

Ritchken, P., S. Kuo. 1988. Option Pricing with Finite Revision Opportunities. *Journal of Finance.* 43, 301-308.

Ross, S. A. 1976. The arbitrage theory of capital asset pricing. *Journal of Economic Theory.* 13, 341-360.

Roumiantsev, S., S. Netessine. 2006. What can be learned from classical inventory models: a cross-industry empirical investigation. *M&SOM* (forthcoming).

Sandmo, A. 1971. On the Theory of the Competitive Firm Under Price Uncertainty. *American Economic Review.* 61 65-73.

Schwartz, E. 1997. The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging. *Journal of Finance* 52(3) 923-973.

Seifert, R. W., U. W. Thonemann, W. H. Hausman. 2004. Optimal procurement strategies for online spot markets. *European Journal of Operational Research.* 152 781-799.

Shanken, J. 1992. The Current State of the Arbitrage Pricing Theory. *Journal of Finance.* 47, 1569-1574.

Snow, K. N. 1992. Diagnosing Asset Pricing Models Using the Distribution of Returns. *Journal of Finance.* 46, 955-983.

Stutzer, M. 1993. A Bayesian Approach to Diagnosis of Asset Pricing Models. *Journal of Econometrics.* 68, 367-397.

Swaminathan, J., H. L. Lee. 2003. Design for Postponement. Chapter 5 in *Handbooks in Operations Res. and Management Sci., Vol. 11, Supply Chain Management: Design, Coordination and Operation.* A. G. de Kok, S. C. Graves (eds.), Elsevier B.V., 199-228.

Swaminathan, J., S. Tayur. 1998. Managing Broader Product Lines through Delayed Differentiation Using Vanilla Boxes. *Management Science.* 44 S161-S172.

Van Mieghem, J. 2003. Capacity Management, Investment and Hedging: Review and Recent Developments. *M&SOM.* 5(4) 269-302.

Whang, S., H. L. Lee. 1999. Value of Postponement. Chapter 4 in *Product Variety Management: Research Advances.* T. H. Ho, C. S. Tang (eds.), Kluwer Academic Publishers, 65-84.

Appendix

Proof of Lemma 1. In the following derivation, we suppress the subscript T in x_T and y_T for convenience.

Integration wrt y given x: Note that the conditional distribution of y given x is normal with mean $\hat{\mu}_y = \mu_y + \rho_{xy}\frac{\sigma_y}{\sigma_x}(x - \mu_x)$ and variance $\hat{\sigma}_y = \sigma_y\sqrt{1-\rho_{xy}^2}$. Hence, the inner integral in the first term can be written as:

$$\int_{-\infty}^{y_q} \frac{aP_T^{1-\eta}e^y}{2\pi\hat{\sigma}_y} \exp\left\{-\frac{1}{2}\left(\frac{y-\hat{\mu}_y}{\hat{\sigma}_y}\right)^2\right\} dy \quad = \quad aP_T^{1-\eta}e^{\hat{\mu}_y+\hat{\sigma}_y^2/2}\Phi\left(\frac{y_q-\hat{\mu}_y-\hat{\sigma}_y^2}{\hat{\sigma}_y}\right).$$

And the inner integral in the second term can be written as:

$$\int_{y_q}^{\infty} P_T q \exp\left\{-\frac{1}{2}\left(\frac{y-\hat{\mu}_y}{\hat{\sigma}_y}\right)^2\right\} dy \quad = \quad P_T q\overline{\Phi}\left(\frac{y_q-\hat{\mu}_y}{\hat{\sigma}_y}\right).$$

Thus, the expected profit is equal to

$$E_x\left[aP_T^{1-\eta}e^{\hat{\mu}_y+\hat{\sigma}_y^2/2}\Phi\left(\frac{y_q-\hat{\mu}_y-\hat{\sigma}_y^2}{\hat{\sigma}_y}\right) + P_T q\overline{\Phi}\left(\frac{y_q-\hat{\mu}_y}{\hat{\sigma}_y}\right)\right] - cq.$$

We simplify the first term as follows:

$$
\begin{aligned}
E_x\left[aP_T^{1-\eta}e^{\hat{\mu}_y+\hat{\sigma}_y^2/2}\Phi\left(\frac{y_q-\hat{\mu}_y-\hat{\sigma}_y^2}{\hat{\sigma}_y}\right)\right] &= E_x\left[ae^{(1-\eta)x}e^{\hat{\mu}_y+\hat{\sigma}_y^2/2}\Phi\left(\frac{y_q-\hat{\mu}_y-\hat{\sigma}_y^2}{\hat{\sigma}_y}\right)\right] \\
&= E_x\left[ae^{(1-\eta)x}e^{\mu_y+\rho_{xy}\frac{\sigma_y}{\sigma_x}(x-\mu_x)+\hat{\sigma}_y^2/2}\Phi\left(\frac{y_q-\hat{\mu}_y-\hat{\sigma}_y^2}{\hat{\sigma}_y}\right)\right] \\
&= E_x\left[ae^{(1-\eta+\rho_{xy}\frac{\sigma_y}{\sigma_x})x}e^{\mu_y-\rho_{xy}\frac{\sigma_y}{\sigma_x}\mu_x+\hat{\sigma}_y^2/2}\Phi\left(\frac{y_q-\hat{\mu}_y-\hat{\sigma}_y^2}{\hat{\sigma}_y}\right)\right] \\
&= a\exp\left\{\mu_y+\hat{\sigma}_y^2/2+\mu_x(1-\eta)+\frac{\sigma_x^2(1-\eta+\rho_{xy}\frac{\sigma_y}{\sigma_x})^2}{2}\right\} \\
&\quad \int_{-\infty}^{\infty}\frac{1}{\sqrt{2\pi}\sigma_x}\Phi\left(\frac{y_q-\hat{\mu}_y-\hat{\sigma}_y^2}{\hat{\sigma}_y}\right)\exp\left[-\frac{1}{2\sigma_x^2}\left(x-\mu_x-(1-\eta+\rho_{xy}\frac{\sigma_y}{\sigma_x})\sigma_x^2\right)^2\right]dx
\end{aligned}
$$

This expression can be further simplified by treating it as a convolution of two normally distributed random variables after a suitable transformation of variables. To see this, we write the expression inside $\Phi()$ as

$$
\hat{\sigma}_y\xi \;\leq\; y_q-\hat{\mu}_y-\hat{\sigma}_y^2 \;=\; \log(q/a)+\eta x-\mu_y-\rho_{xy}\frac{\sigma_y}{\sigma_x}(x-\mu_x)-\hat{\sigma}_y^2.
$$

where $\xi \sim N[0,1]$ and $x \sim N[\mu_x+(1-\eta+\rho_{xy}\frac{\sigma_y}{\sigma_x})\sigma_x^2, \sigma_x]$. Rearranging the terms, we get:

$$
\hat{\sigma}_y\xi+(\rho_{xy}\frac{\sigma_y}{\sigma_x}-\eta)x \;\leq\; \log(q/a)-\mu_y+\rho_{xy}\frac{\sigma_y}{\sigma_x}\mu_x-\hat{\sigma}_y^2.
$$

Let $w=\hat{\sigma}_y\xi+\left(\rho_{xy}\frac{\sigma_y}{\sigma_x}-\eta\right)x$. Then, w is normally distributed with mean and variance,

$$
\mu_w=\left(\rho_{xy}\frac{\sigma_y}{\sigma_x}-\eta\right)\left(\mu_x+(1-\eta+\rho_{xy}\frac{\sigma_y}{\sigma_x})\sigma_x^2\right), \qquad \sigma_w^2=\hat{\sigma}_y^2+\left(\rho_{xy}\frac{\sigma_y}{\sigma_x}-\eta\right)^2\sigma_x^2.
$$

Substituting these values into the integral, we get:

$$
\begin{aligned}
&E_x\left[aP_T^{1-\eta}e^{\hat{\mu}_y+\hat{\sigma}_y^2/2}\Phi\left(\frac{y_q-\hat{\mu}_y-\hat{\sigma}_y^2}{\hat{\sigma}_y}\right)\right] \\
&= a\exp\left\{\mu_y+\hat{\sigma}_y^2/2+\mu_x(1-\eta)+\frac{\sigma_x^2(1-\eta+\rho_{xy}\frac{\sigma_y}{\sigma_x})^2}{2}\right\}\Pr\left[w\leq\log(q/a)-\mu_y+\rho_{xy}\frac{\sigma_y}{\sigma_x}\mu_x-\hat{\sigma}_y^2\right] \\
&= a\exp\left\{\mu_y+\hat{\sigma}_y^2/2+\mu_x(1-\eta)+\frac{\sigma_x^2(1-\eta+\rho_{xy}\frac{\sigma_y}{\sigma_x})^2}{2}\right\}\Phi\left(\frac{\log(q/a)-\mu_y+\rho_{xy}\frac{\sigma_y}{\sigma_x}\mu_x-\hat{\sigma}_y^2-\mu_w}{\sigma_w}\right).
\end{aligned}
$$

Substituting for μ_w and σ_w and simplifying, we get the first term in Lemma 1.

Likewise, we simplify the second term as follows:

$$
\begin{aligned}
E_x\left[P_T\overline{\Phi}\left(\frac{y_q-\hat{\mu}_y}{\hat{\sigma}_y}\right)\right] &= E_x\left[e^x\overline{\Phi}\left(\frac{\eta x+\log(q/a)-\hat{\mu}_y}{\hat{\sigma}_y}\right)\right] \\
&= e^{\mu_x+\sigma_x^2/2}\int_{-\infty}^{\infty}\frac{1}{\sqrt{2\pi}\sigma_x}\overline{\Phi}\left(\frac{\eta x+\log(q/a)-\hat{\mu}_y}{\hat{\sigma}_y}\right)\exp\left[-\frac{1}{2}\left(\frac{x-\mu_x-\sigma_x^2}{\sigma_x}\right)^2\right]dx \\
&= e^{\mu_x+\sigma_x^2/2}\int_{-\infty}^{\infty}\frac{1}{\sqrt{2\pi}\sigma_x}\overline{\Phi}\left(\frac{\eta x+\log(q/a)-\rho_{xy}\frac{\sigma_y}{\sigma_x}(x-\mu_x)-\mu_y}{\hat{\sigma}_y}\right)\exp\left[-\frac{1}{2}\left(\frac{x-\mu_x-\sigma_x^2}{\sigma_x}\right)^2\right]dx \\
&= e^{\mu_x+\sigma_x^2/2}\Pr\left[\hat{\sigma}_y\xi+\left(\rho_{xy}\frac{\sigma_y}{\sigma_x}-\eta\right)x\geq\log(q/a)+\rho_{xy}\frac{\sigma_y}{\sigma_x}\mu_x-\mu_y\right],
\end{aligned}
$$

where $\xi \sim N[0,1]$ and $x \sim N[\mu_x + \sigma_x^2, \sigma_x]$. Let $z = \hat{\sigma}_y \xi + \left(\rho_{xy}\frac{\sigma_y}{\sigma_x} - \eta\right) x$. Then, z is normally distributed with mean and variance,

$$\mu_z = \left(\rho_{xy}\frac{\sigma_y}{\sigma_x} - \eta\right)(\mu_x + \sigma_x^2), \qquad \sigma_z^2 = \hat{\sigma}_y^2 + \left(\rho_{xy}\frac{\sigma_y}{\sigma_x} - \eta\right)^2 \sigma_x^2 = \eta^2\sigma_x^2 + \sigma_y^2 - 2\eta\rho_{xy}\sigma_x\sigma_y = \sigma_w^2.$$

Thus, the second term gives

$$e^{\mu_x + \sigma_x^2/2}\overline{\Phi}\left(\frac{\log(q/a) + \rho_{xy}\frac{\sigma_y}{\sigma_x}\mu_x - \mu_y - \mu_z}{\sigma_z}\right).$$

Substituting for μ_z and σ_z and simplifying, we get the second term in Lemma 1. QED

Proof of Proposition 1. Using the expression in Lemma 1, we take derivatives of the expected profit wrt q and simplify. Thus, we get the first and second order derivatives as:

$$\begin{aligned}
\frac{d\pi(t,q)}{dq} &= e^{\mu_x + \sigma_x^2/2}\overline{\Phi}(d_1) - c, \\
\frac{d^2\pi(t,q)}{dq^2} &= -e^{\mu_x + \sigma_x^2/2}\phi(d_1)\frac{1}{\sigma_z q}.
\end{aligned}$$

Clearly, $\pi(t,q)$ is strictly concave in q. The necessary and sufficient condition for a positive inventory level is that $e^{\mu_x + \sigma_x^2/2} > c$. Under this condition, the FOC reduces to:

$$\begin{aligned}
\log(q/a) &= \mu_y + \rho_{xy}\sigma_x\sigma_y - \eta(\mu_x + \sigma_x^2) + \sigma_z\Phi^{-1}\left(1 - \frac{c}{e^{\mu_x + \sigma_x^2/2}}\right). \\
\text{Or,} \qquad q^*(t) &= a\exp\left[\mu_y + \rho_{xy}\sigma_x\sigma_y - \eta(\mu_x + \sigma_x^2) + \sigma_z\Phi^{-1}\left(1 - \frac{c}{e^{\mu_x + \sigma_x^2/2}}\right)\right].
\end{aligned}$$

Otherwise, $q^*(t) = 0$. Substituting these values in $\pi(t,q)$, we obtain the required result. QED

Proof of Proposition 2. Given information (x_u, y_u) at time u, the distribution of (x_t, y_t) is bivariate normal. Let μ_{xut} and μ_{yut} denote the conditional means of x_t and y_t, σ_{xut}^2 and σ_{yut}^2 denote the conditional variances, and σ_{xyut}^2 denote the conditional covariance. We have:

$$\begin{aligned}
&\mu_{xut} = x_u e^{-h(t-u)} + (m - \tfrac{s_P^2}{2h})(1 - e^{-h(t-u)}), &&\sigma_{xut}^2 = (1 - e^{-2h(t-u)})\tfrac{s_P^2}{2h}, \\
&\mu_{yut} = y_u + (\alpha_\epsilon - s_\epsilon^2/2)(t-u), &&\sigma_{yut}^2 = s_\epsilon^2(t-u), \\
&\sigma_{xyut} = (1 - e^{-h(t-u)})\tfrac{\rho s_P s_\epsilon}{h}.
\end{aligned}$$

These expressions are analogous to those derived in §3.1 for the distribution of (x_T, y_T) given (x_t, y_t). We integrate $\pi^*(t)$ with respect to the conditional distribution of (x_t, y_t). This gives the expression for $Y(u,t)$ given in Proposition 2. Unlike Lemma 1 and Proposition 1, a closed form expression for $Y(u,t)$ cannot be obtained due to the finite lower limit $\underline{x}_t$ on the integral over values of x_t. QED

Proof of Proposition 4. From (10), postponement of the inventory decision from time u to time $t > u$ leads to an increase in value if

$$d_1^{**} - s_\epsilon\sqrt{T-t} \geq d_1^* - s_\epsilon\sqrt{T-u}.$$

Clearly, if $c_t \le c_u e^{r(t-u)}$, then $d_1^{**} > d_1^*$ so that the inequality is satisfied. This proves (i).

For (ii), note that the inequality is satisfied as an equality at $t = u$. Hence, we take derivatives with respect to t to evaluate the sign of the inequality for $t > u$. Note that the RHS is independent of t. The derivative of the LHS with respect to t gives

$$-\frac{\tilde{c}_t'}{e^m \phi(d_1^{**})} + \frac{s_\epsilon}{2\sqrt{T-t}},$$

where $\tilde{c}_t \equiv c_t e^{r(T-t)} = c_u e^{rT+\delta(t-u)}$. We have

$$\begin{aligned} -\frac{\tilde{c}_t'}{e^m \phi(d_1^{**})} + \frac{s_\epsilon}{2\sqrt{T-t}} &= -\frac{\delta c_u e^{rT+\delta(t-u)}}{e^m \phi(d_1^{**})} + \frac{s_\epsilon}{2\sqrt{T-t}} \\ &\ge -\frac{\delta c_u e^{(r+\delta)T}}{e^m \phi(d_1^{**})} + \frac{s_\epsilon}{2\sqrt{T}}. \end{aligned}$$

This gives that postponement is an optimal strategy at time u if

$$s_\epsilon \ge \frac{2\sqrt{T}\delta c_u e^{(r+\delta)T}}{e^m \phi(d_1^{**})}.$$

This gives a finite lower bound on s_ϵ since $\phi(d_1^{**})$ is bounded below by the minimum of $\phi(\Phi^{-1}(1 - c_u e^{-m}))$ and $\phi(\Phi^{-1}(1 - c_T e^{r(T-u)-m}))$. QED

Proof of Proposition 5. If we show that it is optimal to postpone for $c_t = c_u e^{(r+\rho s_P s_\epsilon - \eta s_P^2)(t-u)}$, then it will also be optimal to postpone for $c_t \le c_u e^{(r+\rho s_P s_\epsilon - \eta s_P^2)(t-u)}$. Hence, we assume that $c_t = c_u e^{(r+\rho s_P s_\epsilon - \eta s_P^2)(t-u)}$.

From the definition of $Y(u,t)$, it is clear that $\lim_{t\to u} Y(u,t) = \pi^*(u)$. Thus, we compute the sign of the derivative of $Y(u,t)$ with respect to t to compare the value of postponement with the value of early commitment. We have (note that the value of the integrand at the lower limit equals zero):

$$\begin{aligned} \frac{dY}{dt} &= e^{-r(T-u)} P_u D_u e^{[(1-\eta)\alpha_P^* - \eta(1-\eta)s_p^2/2 + (1-\eta)\rho s_P s_\epsilon + \alpha_\epsilon^*](T-u)} \\ &\quad \int_{\xi_L}^{\infty} \phi\left(d_1^{**} - \sigma\sqrt{T-t}\right) \left[\frac{dd_1^{**}}{dt} + \frac{\sigma}{2\sqrt{T-t}}\right] \phi(\xi) d\xi. \end{aligned}$$

Here, we define $\sigma = \sqrt{\eta^2 s_P^2 + s_\epsilon^2 - 2\eta\rho s_P s_\epsilon}$ so that $\sigma_z = \sigma\sqrt{T-t}$. We also compute dd_1^{**}/dt as

$$\frac{dd_1^{**}}{dt} = \frac{c_u e^{(r-\alpha_P^*)(T-u)}}{P_u \phi(d_1^{**})} \left[\frac{s_P^2}{2} + \frac{\xi s_P}{2\sqrt{t-u}}\right] e^{-\frac{s_P^2}{2}(t-u) - \xi s_P \sqrt{t-u}}.$$

Collecting all terms together, note that in dY/dt, the expression $e^{-r(T-u)} P_u D_u e^{[(1-\eta)\alpha_P^* - \eta(1-\eta)s_p^2/2 + (1-\eta)\rho s_P s_\epsilon + \alpha_\epsilon^*](T-u)}$ is non-negative. In addition,

$$\int_{\xi_L}^{\infty} \phi\left(d_1^{**} - \sigma\sqrt{T-t}\right) \frac{\sigma}{2\sqrt{T-t}} \phi(\xi) d\xi$$

is non-negative. Thus, it only remains to ascertain the sign of

$$\int_{\xi_L}^{\infty} \phi\left(d_1^{**} - \sigma\sqrt{T-t}\right) \frac{dd_1^{**}}{dt} \phi(\xi) d\xi. \tag{16}$$

Note that

$$\frac{\phi\left(d_1^{**} - s_\epsilon\sqrt{T-t}\right)}{\phi(d_1^{**})} = e^{d_1^{**} s_\epsilon \sqrt{T-t} - \frac{s_\epsilon^2}{2}(T-t)},$$

which is positive and increasing in ξ. Moreover, $\left[\frac{s_P^2}{2} + \frac{\xi s_P}{2\sqrt{t-u}}\right] e^{-\frac{s_P^2}{2}(t-u) - \xi s_P \sqrt{t-u}} \phi(\xi)$ is negative for $\xi < s_P\sqrt{t-u}$ and non-negative otherwise. Thus, the integrand in (16) is a product of positive and increasing weights with a function that changes sign from negative to positive. To determine whether the sign of (16) is positive, it is sufficient to ignore all multiplicative terms that are positive and increasing in ξ. Thus, we are left with

$$\begin{aligned} \int_{\xi_L}^{\infty} \left[\frac{s_P^2}{2} + \frac{\xi s_P}{2\sqrt{t-u}}\right] e^{-\frac{s_P^2}{2}(t-u) - \xi s_P\sqrt{t-u}} \phi(\xi) d\xi &= \frac{s_P}{2\sqrt{t-u}} \int_{\xi_L + s_P\sqrt{t-u}}^{\infty} \hat{\xi}\phi(\hat{\xi}) d\hat{\xi} \\ &\geq 0, \end{aligned}$$

where the first equation uses the change of variables $\hat{\xi} = \xi + s_P\sqrt{t-u}$. QED

Figure 1: Plot of value of postponement versus price volatility s_P for different values of correlation coefficient ρ between price forecast and ε forecast processes when price elasticity $\eta = 0$

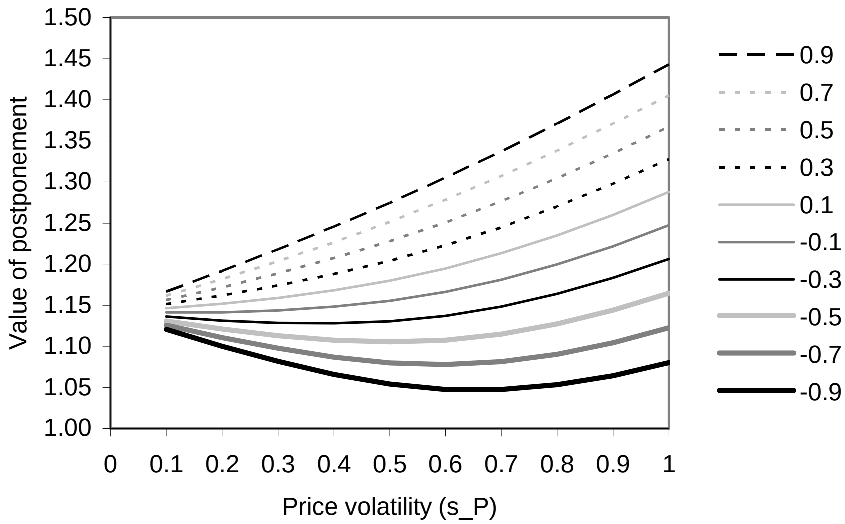

The parameters used to compute values for this graph are: time instant u=0 at which the decision whether or not to postpone is taken, t=0.5 at which inventory is purchased if we postpone at time u=0, and T=1 at which demand is realized; cost of purchase, $c_u = c_t = 0.4$; price information at time u, $P_u = 1$; $r = \alpha_P = \lambda = \rho_{Pm} = \eta = 0$; volatility of ε process, $s_\varepsilon = 0.5$. The correlation coefficient ρ varies from -0.9 to 0.9 in increments of 0.1 (only alternate lines are shown in the graph) and price volatility s_P varies from 0.1 to 1.0 in intervals of 0.1. The value of postponement is computed as the ratio of expected profit if the optimal inventory decision is postponed to time to that if the optimal inventory decision is taken at time u. It is optimal to postpone at time u if this ratio is greater than 1.

Figure 2: Plot of value of postponement versus price volatility s_P for different values of price elasticity η when the correlation coefficient ρ = 0.5.

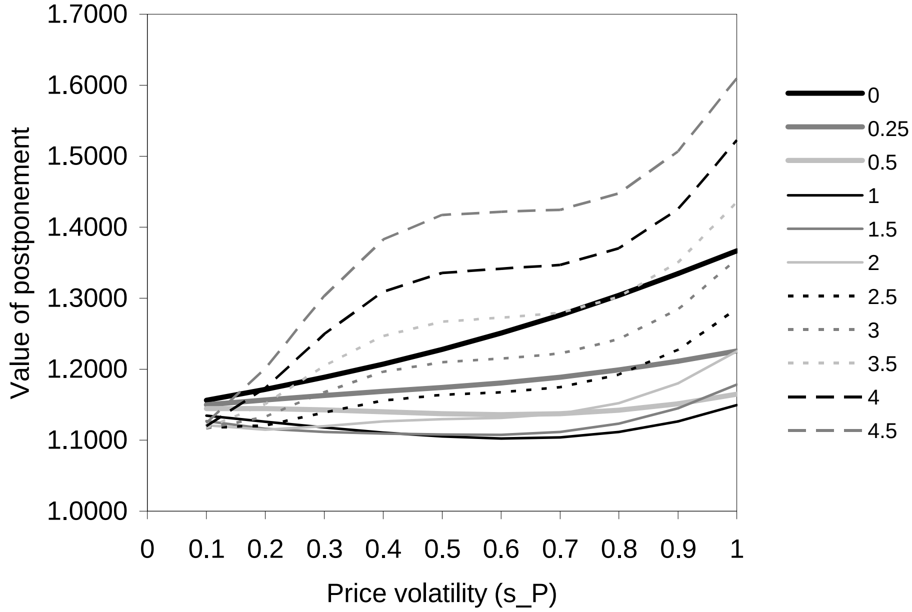

The parameters used to compute values for this graph are: time instant u=0 at which the decision whether or not to postpone is taken, t=0.5 at which inventory is purchased if we postpone at time u=0, and T=1 at which demand is realized; cost of purchase, $c_u = c_t = 0.4$; price information at time u, $P_u = 1$; $r = \alpha_P = \lambda = \rho_{Pm} = 0$; volatility of ε process, $s_\varepsilon = 0.5$; the correlation coefficient $\rho = 0.5$. The value of price elasticity of demand, η, varies from 0 to 4.5 in increments of 0.25 (only alternate lines are shown in the graph) and price volatility s_P varies from 0.1 to 1.0 in intervals of 0.1. The value of postponement is computed as the ratio of expected profit if the optimal inventory decision is postponed to time to that if the optimal inventory decision is taken at time u. It is optimal to postpone at time u if this ratio is greater than 1.

Figure 3: Plot of relative cost of procurement, c_t/c_0, as a function of time t for the case of constant prices

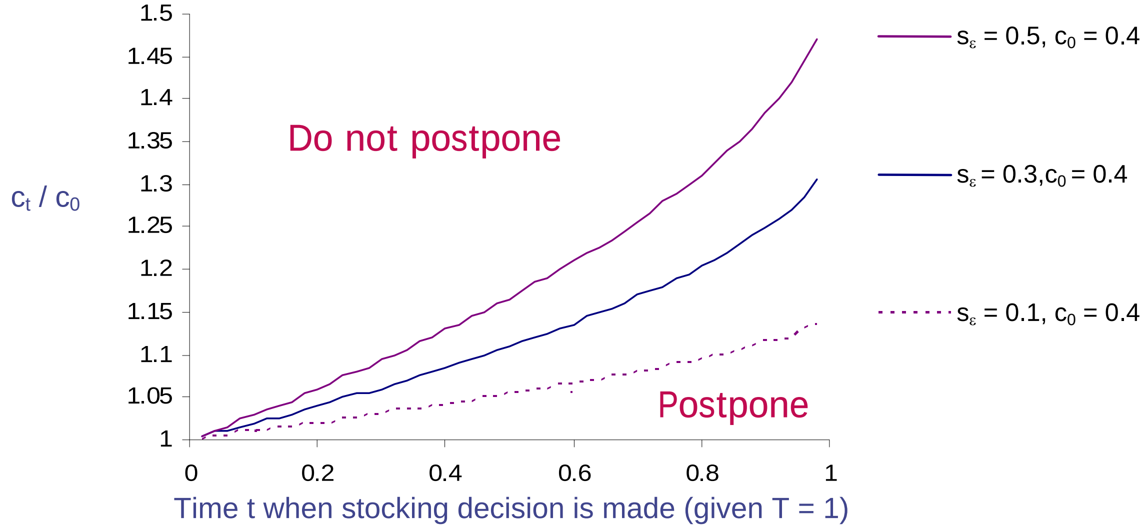

This graph depicts the decision of a firm at time 0 whether to stock inventory at time 0 or to postpone the stocking decision to a future time t. It shows the ratio of the cost of procurement at time t to the cost of procurement at time 0 at which the firm is indifferent between stocking inventory at time 0 and postponing the stocking decision to time t. The graph plots this ratio for t varying between 0 and 1 for different values of the volatility s_ε of the demand forecast process ε. The price is set to a constant value 1, the risk-free discount rate is r = 0.05, and demand is realized at time 1. The plot is independent of the drift rate and risk premium for the demand forecast process.

Figure 4: Time trend in days of inventory of gold mining firms from 1995 Q1 to 2005 Q4, estimated using a fixed effects model with dummy variables for each firm and each quarter

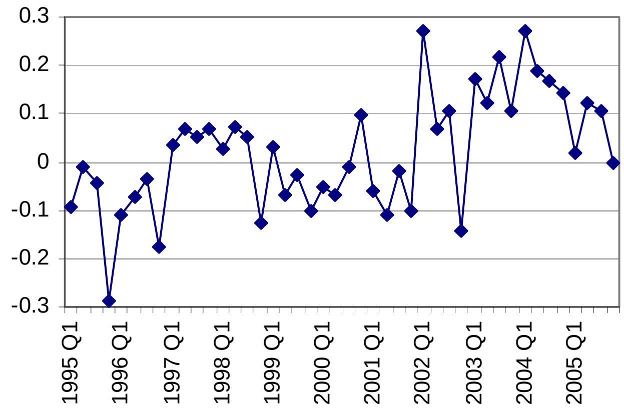

Figure 5: Time series plots of average price of gold and volatility of gold prices from 1995 Q1 to 2005 Q4

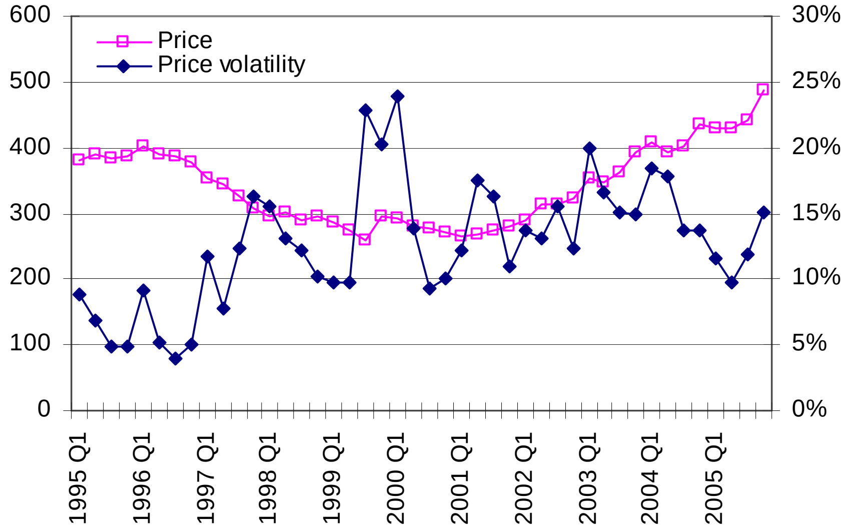

Note: Average price of gold for each quarter is measured as the average of daily closing prices in that quarter. Volatility of gold prices is measured by computing the standard deviation of daily returns for each quarter and multiplying by √252 to annualize the number.

Table 1: Summary statistics of quarterly days of inventory and gross margin for gold mining firms (NAICS code 212221) listed on NYSE/AMEX/NASDAQ for the period 1995-2005

Company name	Company Ticker Symbol	# of obs.	Median quarterly sales ($ million)	Median quarterly gross profit ($ million)	Days of Inventory			Gross Margin		
					Min	Max	Median	Min	Max	Median
Campbell Resources Inc	3CBLRF	19	9.4	1.1	43.37	109.58	68.46	0.02	0.45	0.22
Lionore Mining Intl Ltd	3LMGGF	22	47.6	22.7	16.36	69.24	34.56	0.04	0.58	0.37
Barrick Gold Corp	ABX	35	376.5	200.0	46.95	102.58	60.93	0.30	0.66	0.48
Alta Gold Co	ATGDQ	18	4.6	1.7	97.35	597.88	211.84	0.22	0.62	0.41
Anglogold Ashanti Ltd - ADR	AU	24	525.1	184.3	36.10	94.39	71.34	0.17	0.48	0.35
Compnia Minas Buenvntr - ADR	BVN	17	67.9	34.8	47.92	333.28	70.26	0.01	0.84	0.54
Cambior Inc	CBJ	34	65.4	17.4	12.15	70.67	44.63	0.08	0.36	0.26
Claude Resources Inc	CGR	37	4.8	1.4	9.32	349.52	204.71	0.02	0.47	0.28
Eldorado Gold Corp	EGO	17	9.2	2.9	63.80	194.42	87.54	0.20	0.47	0.33
Goldcorp Inc	GG	40	29.9	14.3	50.62	375.83	71.87	0.01	0.70	0.40
Glencairn Gold Corp	GLE	28	5.7	1.6	30.66	189.23	100.85	0.02	0.47	0.24
Glamis Gold Ltd	GLG	36	15.8	5.7	86.82	223.43	137.57	0.02	0.61	0.44
Golden Star Resources Ltd	GSS	19	13.1	4.4	78.74	167.22	124.24	0.07	0.62	0.36
Iamgold Corp	IAG	23	24.2	10.8	38.05	409.10	251.67	0.22	0.56	0.44
Kinross Gold Corp	KGC	37	68.8	20.4	63.47	119.39	84.70	0.11	0.41	0.31
Meridian Gold Inc	MDG	40	29.6	19.7	35.01	191.14	73.40	0.22	0.89	0.65
Newmont Mining Corp	NEM	43	424.2	170.3	74.29	203.13	113.39	0.22	0.54	0.41
Northgate Minerals Corp	NXG	20	34.5	9.3	22.49	69.50	40.78	0.04	0.53	0.30
Placer Dome Inc	PDG	40	307.5	134.0	53.34	96.84	75.01	0.11	0.60	0.41
Pacific Rim Mining Corp	PMU	19	8.3	2.6	69.28	259.45	124.25	0.03	0.58	0.29
Richmont Mines Inc	RIC	40	6.6	2.2	14.46	832.06	31.89	0.04	0.57	0.35
Sterlite Gold Ltd	SGDTF	25	4.3	1.0	59.38	171.93	102.68	0.03	0.75	0.29
Pooled Data		633	25.5	9.8	9.32	832.06	79.76	0.01	0.89	0.36

This table provides summary data on the sales, gross profit, days of inventory and gross margin, on a quarterly basis, for 22 gold mining firms listed on US exchanges, over the period 1995-2005. The first two columns provide the company name and ticker symbol, the third column provides the total number of observations in the dataset for each firm, and the remaining columns provide statistics for the variables over the sample period of quarterly observations. We compute gross profit as the difference between sales and cost of goods sold for the quarter. 18 of the 22 firms are listed on Compustat as active in Q4, 2005, while the others are listed as inactive. Firms with ticker symbols ATGDQ, EGO, PMU and PDG are not active.

Table 2: Estimates for the FGLS regression of log(DI_{it}) on $\Phi^{-1}(GM_{it})$, price standard deviation, price variance, and time trend

Coefficients' estimates for explanatory variables	Equation (14)		Equation (15)	
	Estimate	Std Err	Estimate	Std Err
$\Phi^{-1}(GM_{it})$	0.26**	0.11	0.25**	0.11
Price standard deviation, $\hat{s}_{Pt}$	4.12***	1.66	2.64	1.96
Interaction term, $\Phi^{-1}(GM_{it})\hat{s}^2_{Pt}$	-0.62	0.77	-0.66	0.77
Price variance, $\hat{s}^2_{Pt}$	-12.70**	6.12	-8.33	6.83
Linear time trend			0.003	0.002

We estimate the regression model using feasible generalized least squares (FGLS) procedure in order to control for group-wise heteroscedasticity across firms. We also tested a random effects model by treating the differences between firms as random effects rather than fixed effects. This model gave parameters' coefficients statistically indistinguishable from the fixed effects model. ***, ** denote statistical significance at 0.01, 0.05, respectively.