Optimal Timing of Inventory Decisions with Price Uncertainty

Vishal Gaur[*], Nikolay Osadchiy[†], Sridhar Seshadri[‡], Marti G. Subrahmanyam[§]

October 14, 2015

Abstract

What is the optimal time for a firm to buy inventory to sell in a product market in which the selling price and demand are random and their forecasts improve with time? What is the value of order timing flexibility to the firm? What lead times would a supplier see? To answer these questions, we develop a continuous time inventory model where demand and price are realized at the horizon date T, and the stocking decision can be made at any time in the interval $[0, T]$ given progressively more accurate forecasts of price and demand and a time dependent purchasing cost. We show that the optimal timing of inventory ordering decision follows a simple threshold policy in the price variable with a possible option of non-purchasing, and is independent of the demand. Given this policy structure, we evaluate the benefits of timing flexibility using the best pre-committed order timing policy as the benchmark. We find that the time flexible ordering policy is particularly valuable in the high price volatility or low margin environments. Higher price volatility leads to a greater postponement in the purchasing decisions. In addition, decreasing profit margins lead to a higher lead time variability.

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

[†]Goizueta Business School, Emory University, Atlanta, GA 30309. E-mail: *nikolay.osadchiy@emory.edu*.

[‡]Indian School of Business, Hyderabad, India. E-mail: *sridhar_seshadri@isb.edu*.

[§]Leonard N. Stern School of Business, New York University, New York, NY 10012. E-mail: *msubrahm@stern.nyu.edu*.

1 Introduction

Consider a firm that sells in a product market in which the selling price and demand are random and their forecasts improve with time. What is the optimal time to buy inventory? Postponing the decision can give more accurate information about the demand and price, but will also increase the purchasing cost. This basic trade off is well recognized in the literature. However, the majority of the literature on postponement treats prices as exogenous and constant, whereas in many practical applications, the prices of products are stochastic and the forecasts of prices improve over time. In fact, in many industries from milk farming to construction, price fluctuations can even force companies out of business. Thus, postponement may be valuable due to the resolution of not only demand uncertainty, but also price uncertainty. Depending on the price forecast, a decision maker may choose whether to purchase inventory at that moment or wait. This issue has become considerably more important in recent years. During the economic recession of 2008, firms faced weakened demand for their products as well as low prices that were realized in the marketplace, and were adjusting their ordering decisions accordingly. Moreover, due to increased competition, firms are price-takers in many markets. Hence, they are susceptible to not only demand volatility but also price volatility.

This motivates our research questions. In the presence of demand and price uncertainty, what is the optimal procurement policy? Is there an optimal order timing, and if so, is it deterministic or stochastic? What is the value of timing flexibility in the inventory procurement? How are the order timing and the value of timing flexibility affected by the volatility of demand and price, and the correlation between the two?

To answer these questions, we present a continuous time inventory model where demand and price are realized at time T and the firm has an option to make a stocking decision at any time in the interval $[0, T]$. The selling price and the demand for the product are stochastic, and evolve as a mean-reverting and a geometric Brownian motion process, respectively. Further, they are interrelated through the price elasticity of demand and correlated noise processes. The correlation between the noise processes can occur if both of them are correlated with the state of the economy or with a return on a broad market portfolio. Finally, the cost of procurement may change over time reflecting shorter lead times or other factors. The firm is a price taker in the product market. Its 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 firm makes the stocking decision that maximizes its value adjusted for the riskiness of the inventory investment.

We make a distinction between *postponement* and *timing flexibility*. Postponement means that the order can be made at any pre-committed future time, whereas timing flexibility means that order can be made at any time as long as a certain condition is met. The first key insight from our paper is that a firm should not pre-commit to a fixed lead time in price-volatile environments. We find that the optimal timing of inventory procurement is flexible and stochastic. It is independent of demand forecast, and follows a simple threshold policy in the price variable. In the case of constant prices, this yields the optimal postponement time that balances the benefit of improved demand information and the increased cost of procurement. This is a static postponement scenario, studied extensively in the literature. In the case of random prices, the optimal order timing depends on the price sample path and is, therefore, uncertain. In some situations, if the price turns out to be low, the firm may find it optimal not to buy inventory at all. However, price volatility can also be beneficial to a firm if the price goes high enough to make the inventory investment profitable.

The second contribution of our paper is to provide insights into the value of order timing flexibility and its drivers. We compare the firm's value and the realized lead times under the time flexible ordering policy with the best static postponement policy. We find that, for the time flexible policy, the order-triggering price thresholds are non-monotone in time: they decrease at first despite increasing costs. The firm's values under the time flexible ordering policy and the optimal static postponement increase with price volatility. The difference between the two defines the value of timing flexibility. Overall, timing flexibility can increase the value of a firm by up to 30-40% compared to the optimal static postponement policy under price volatility.

The value of timing flexibility depends on profit margin, price volatility, demand volatility, and is mediated by the price mean-reversion strength, price demand elasticity and correlation, and the market price of risk. Timing flexibility is particularly valuable for products with low or even negative *ex-ante* profit margin. In absolute terms, its value is the highest when the *ex-ante* margin is zero, in other words, when the option to invest is neither in- nor out-of-the-money. In this condition, timing flexibility gives the firm the ability to react to the favourable price forecast. For products with *ex-ante* positive margin the value of timing flexibility increases with price volatility.

Firms have more incentive to postpone ordering decisions if prices are volatile and margins are low. In addition, the variability of order timing increases as *ex-ante* margins decrease. For products with *ex-ante* positive margins, order timing variability increases with the price volatility. To sum up, the value of timing flexibility comes from three sources: a volatility effect because the volatility of price and demand forecasts decreases with postponement, an optionality effect because of the option to forgo or postpone the purchase, and a newsvendor effect because the optimal purchase quantity adjusts to available price and demand information and time elapsed.

Our results not only inform the purchasing firm, but also a supplier. Volatile prices imply that ordering lead timing is random, as opposed to the classic postponement problem where prices are fixed and the optimal order timing is known in advance. If price volatility is high, the supplier should be prepared to offer shorter lead times, but not completely eliminate the early shipment option. It should also be prepared for a more variant order timing for more expensive products. In our numerical experiments, we observe firms postponing their purchases despite exponentially increasing costs. Thus, suppliers can charge exponentially increasing prices for quick response shipping options and increase their revenue. Overall, a volatile price environment can be a win-win for both buyers and suppliers.

2 Related Literature

This paper builds on and contributes to two distinct bodies of literature: postponement of inventory decisions and real options. The benefits of postponement come from two sources (Whang and Lee 1998): uncertainty resolution and forecast improvement. The more recent information can be used to better infer future demand (Lee and Tang 1997, Aviv and Federgruen 2001), and better forecasts can be used to adjust the safety stock (Fisher and Raman 1996, Kaminsky and Swaminathan 2004, Kouvelis and Tian 2014). The paper by Xiao et al. (2015) uses the uncertainty in purchasing cost to adjust inventory decisions over time. These papers focus on demand uncertainty and treat price as deterministic or as a decision variable. Reversing the uncertainty pattern, Li and Kouvelis (1999) study the value of flexible order timing in an environment with deterministic demand but uncertain procurement costs, and analyze benefits of timing flexibility via a cost-minimizing dynamic program.

The problem of optimal order timing under price uncertainty has been studied in the context of spot market trading. Guo et al. (2011) model a firm that, in anticipation of demand arriving at some random future time, can buy or sell inventory at the spot market at any moment before demand arrives. The random demand timing allows transforming the problem into infinite horizon expected discounted profit maximization. Assuming a fixed markup over the spot price and positive transaction costs associated with buying or selling on the spot market, the paper shows that the optimal inventory policy is of the price-dependent threshold type. Secomandi and Kekre (2014) study a two-period commodity procurement problem, with the first period representing the forward market and the second the spot market. They allow spot price and demand to be correlated, and show that the value of forward procurement increases with the degree of correlation.

The real options literature provides a framework for valuation of real assets using methods initially developed for valuation of financial options. Flexible production or ordering capabilities are examples of such assets. An important work by McDonald and Siegel (1986) studies the value of delaying an investment into a project with uncertain value and cost. In their setting, both project value and costs evolve as geometric Brownian motions, investors are risk averse, and the investment opportunity exists over some finite time interval. They show that there exists some time-dependent factor, such that it is optimal to invest once the project benefits exceed investment costs by that factor.

Our model can be viewed as a refinement of the above mentioned models, with continuous time evolution of demand and price uncertainty, time varying costs, and endogenous decision on the amount of invested capital (i.e., purchased inventory quantity). As such, it combines the logic of demand uncertainty resolution pertaining to the postponement literature and the logic of value of an option, pertaining to the real options literature.

Our modeling of price builds on the prior research in economics and finance studying evolution of commodity prices (Schwartz 1997), as well as on inventory management for traded commodities (see Haksöz and Seshadri 2007, for a review). We further allow the price and the demand for the product to be correlated with asset prices in the financial market. The correlation between product price and financial market has been observed in commodity markets (Caballero et al. 2008), and the demand-side correlation has been documented in retailing by Osadchiy et al. (2013) and in wholesale trade and manufacturing by Osadchiy et al. (2015). This correlation is important for risk adjustment by risk-averse investors, as well as for financial hedging, which can be a substitute to timing flexibility (Chod et al. 2010).

The adjustment for the risk taken by risk averse investors in our model is based on Constantinides (1978) and McDonald and Siegel (1985). In this framework, the market risk adjustment is based on the 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 approach is, therefore, "descriptive of value-maximizing publicly-owned firms, and is widely used in the finance literature" as argued by McDonald and Siegel (1985). In the operations management literature, it has been used by

Kouvelis (1999), Birge (2000), and more recently by Kouvelis and Tian (2014).

Finally, the option we consider bears some resemblance to an American-style Quanto option (Piros 1998). Similarly to a Quanto option, the payoff of the option to invest in a newsvendor depends on two variables, price and demand, in a non-linear fashion. However, 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 inventory decision, and the role of the information state variable. Consequently, the comparative statics analysis of our option price is more complex.

3 Model

We set up a single-period model of a firm that sells a product with stochastic demand and stochastic selling price. Time is indexed from 0 to T. The demand D_T and the price P_T are realized at time T, and the stocking decision can be taken at any 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. The firm seeks to determine when to procure inventory, and in what quantity, in order to maximize its expected profit.

Stochastic Prices. The price of the product at the future time instant T is stochastic due to environmental factors that are not fully known before time T. There can be many such factors, including the extent of competition, the quality of the product, changes in customer preferences, etc. We assume that the firm is a price-taker in its market. It has a forecast of price at time t, which it uses to make its stocking decision. The forecast is constructed from currently available information. It becomes increasingly accurate with time due to the gradual revelation of information, with the forecast at time T being equal to the actual realization P_T.

We model this price forecast evolution as a standard mean-reverting continuous time stochastic process:

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

This specification for price forecast is similar to models of convergence of futures prices of commodities to cash as constructed by Schwartz (1997), and the investment project value as in Dixit and Pindyck (1994). 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. These parameters can be estimated by fitting the model on historical data. As h tends to zero, the price forecast process (1) reduces to the special case of Geometric Brownian Motion (GBM).

Stochastic Demand. We model the demand to be isoelastic in price with a random intercept. Let $D_T = aP_T^{-\eta}\epsilon_T$, 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 is a state variable representing the information about the scale of demand at time t, α_ϵ is the rate of drift in the scale of demand, s_ϵ is the instantaneous volatility of ϵ_t, and z_ϵ is a standard Brownian motion. We call this process as the demand forecast process. Such a multiplicative demand model has been commonly used in the literature, see Petruzzi and Dada (1999), Federgruen and Heching (1999), and Bernstein and Federgruen (2005), with the difference that those papers consider price to be a decision variable, whereas we consider a price-taking firm.

Demand changes with price in two ways in our model, the first is along the demand curve due to the price elasticity of demand η, and the second due to a scaling of the demand curve caused by the conditional distribution of ϵ_T changing with price. For this to occur, we allow z_ϵ to be correlated with z_P with a correlation coefficient ρ, i.e., $\rho dt = dz_P dz_\epsilon$. Both positive and negative correlations between z_ϵ and z_P are plausible in practice. For example, changes in competition due to entry or exit of firms or emergence of a new technology would cause the scale of demand and prices to move in the same direction, i.e., the scale of demand and price could rise or fall simultaneously, manifesting as a positive correlation between z_ϵ and z_P. Large-scale economic downturns can also depress prices and demand simultaneously. On the other hand, technology learning curves or a maturing of the product lifecycle may expand the market and lower prices simultaneously, introducing a negative correlation between z_ϵ and z_P.

Market price of risk. In general, the firm can be risk averse and adjust the value of the option to invest according to the market price of risk. To account for the market price of risk we will use the equivalent risk-neutral valuation. The approach adjusts the expected rate of return for the systematic, economy-wide risk, the only risk for which investors are compensated. Our adjustment is a direct application of the approach of McDonald and Siegel (1985), Section 3, and is standard in the finance literature. 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. (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.) 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). The difference between our setup and McDonald and Siegel (1985) is two-fold: the option value depends on price and demand variability, and the price is mean-reverting. The first issue is addressed by using the bi-variate Ito's Lemma. The second issue results in the drift adjustment of the price process that depends on the mean reversion speed. The derivation is omitted for brevity. Hereafter, to account for the market price of risk, we use the following adjustments to the drift rates: in (1), replace m by $m - \lambda\rho_{Pm}s_P/h$, and in (2), replace α_ϵ by $\alpha_\epsilon - \lambda\rho_{\epsilon m}s_\epsilon$.

Stocking Decision. The firm can make its stocking decision at any time $t \in [0, T]$. The firms' objective is to maximize profit and the inventory decision may be based on price and demand information revealed at time t. Given the uncertainty in price and demand, the ordering time t is a random variable. 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 unit cost of purchase is denoted as c_t. It includes not only production and transportation costs, but also the cost incurred for holding one unit of the finished product from time t to T. The cost trajectory c_t is known to the buyer for the entire period $[0, T]$. For simplicity, we assume that the firm can place its order 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 (a) the time at which to place the order, and (b) the quantity of product to purchase. We assume zero salvage value of inventory left over at time T for parsimony.

Note that this 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. First, it can be extended to accommodate a lower bound on lead-time by modifying the set of feasible dates for inventory procurement to some interval $[0, T - \bar{L}]$, where $\bar{L} < T$ is a constant. Second, a finite non-zero production time can be accommodated in the model. Third, a 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. Finally, the formulation can be extended to accommodate multiple opportunites for inventory ordering. The resulting problem will be similar to a considerably more complex optimal multiple stopping time problem (Kobylanski et al. 2011) with endogenous inventory decisions. We exclude these considerations from the model in order to focus on the implications of price and demand volatility on the timing and quantity of the stocking decision.

4 Optimal Timing of Inventory Decisions

At each time instant, as the price and demand forecast evolve, the firm has to decide whether to purchase inventory at that instant or to delay the purchase. Thus, the firm is endowed with an option on the newsvendor because its time of purchase is flexible. The option derives its value from the randomness of price and demand, the ability of the firm to optimize the purchase quantity, and the ability to shut down. The cost arises from the time value of money and the increase in procurement cost over time. The firm exercises this option when it decides not to delay the purchase any longer.

We first set up the firm's optimization problem. Let $V(t)$ denote the value of the option to invest in inventory at time $t \in [0, T]$ given information (P_t, ϵ_t) available at time t. The firm has full flexibility to order inventory at any moment $\tau \in [t, T]$ or not at all. In general, τ could depend on the realizations of the processes $\{P_t\}$ and $\{\epsilon_t\}$, thus τ is a random variable, a stopping time. Let $\mathcal{T}[t, T]$ be the set of stopping times taking values between t and T. Also let $\pi(\tau, q)$ denote the time t value of expected profit if quantity q is purchased at time τ; $\pi(\tau, q) = e^{-r(T-t)} E[P_T \min\{q, D_T\} | P_\tau, \epsilon_\tau] - c_\tau q$. Let $q^*(\tau)$ denote the optimal stocking quantity if the purchase decision is made at time τ, and $\pi^*(\tau, P_\tau, \epsilon_\tau)$ denote the corresponding optimal expected profit. P_τ and ϵ_τ are added as arguments of π^* to emphasize the dependence of $q^*(\tau)$ on the information $(P_\tau,\ \epsilon_\tau)$ and the dynamic nature of the problem.

We solve the following optimization problem to determine $V(t)$:

$$V(t) = \sup_{\tau \in \mathcal{T}[t,T], q \in [0,\infty)} E[\pi(\tau, q)|P_t, \epsilon_t] \tag{3}$$
$$= \sup_{\tau \in \mathcal{T}[t,T]} E[\pi^*(\tau, P_\tau, \epsilon_\tau)|P_t, \epsilon_t].$$

In other words, $V(t)$ is the supremum of the expected profit over all feasible time instants at which the option to purchase inventory can be exercised. We proceed as follows. First, we solve for the optimal order quantity and derive the expected profit function if the inventory is purchased at a fixed time. Second, we derive the value function of the firm that takes an optimal inventory decision at some prespecified time. Third, we show that the inventory purchase timing follows a threshold policy in price, i.e., it is optimal to buy inventory if and only if price exceeds a certain time-dependent threshold. We conclude this section by defining the optimal static postponement policy, which we use as a benchmark to assess the value of order timing flexibility.

4.1 Preliminaries

To derive the expected profit function, we need to obtain the conditional distribution of price and demand given the information available at any time $t \in [0, T]$. Given that the price process is mean-reverting whereas the demand process is a GBM, the distribution of demand that would result from the combination of these two processes can be non-intuitive.

Let x_t denote $\log P_t$ and y_t denote $\log \epsilon_t$. Applying Ito's Lemma to (1) and (2), we get

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

for $t \in [0, T]$. Thus, the conditional distribution of (x_τ, y_τ) given information (x_t, y_t) at time $0 \le t \le \tau \le T$ is bivariate normal. Let $\mu_x(\tau, t)$ and $\mu_y(\tau, t)$ denote the means of x_τ and y_τ, $\sigma_x^2(\tau, t)$ and $\sigma_y^2(\tau, t)$ denote the variances of x_τ and y_τ, and $\sigma_{xy}(\tau, t)$ denote the covariance of x_τ and y_τ, given the time t information. When $\tau = T$, which will be clear from the context, we are going to suppress the (T, t) notation, for example, μ_x means $\mu_x(T, t)$. We obtain the following lemma.

Lemma 1. *The conditional distribution of (x_τ, y_τ) at time t, $0 \le t \le \tau \le T$ is bivariate normal with parameters*

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

All proofs are given in the Appendix. Note that $\sigma_{xy}(\tau, t)$ tends to a constant at a slower rate than σ_x. This implies 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 constant.

Lemma 1 yields the conditional expectations of price and demand at time t, $E[P_T|P_t, \epsilon_t]$ and $E[D_T|P_t, \epsilon_t]$, respectively, as well as $E[P_T D_T|P_t, \epsilon_t]$, the conditional expectation of the revenue at time t, if the firm carried infinite inventory and no stockout occurred.

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

Note that $E[D_T|P_t, \epsilon_t]$ and $E[P_T D_T|P_t, \epsilon_t]$ are multiplicative functions of e^{μ_y}, and, therefore, are proportional to ϵ_t, because $\mu_u = y_t + (\alpha_\epsilon - s_\epsilon^2/2)(T-t)$ and $y_t = \log \epsilon_t$. This property will be important for the optimal ordering policy characterization.

4.2 Optimal Order Timing and Quantity

To derive the expected profit, suppose the firm buys inventory at $\tau = t$ using time t available information. Using Lemma 1, the expected profit at time t, $\pi(t, q)$, can be written as a function of the order quantity q by taking conditional expectation over the joint distribution of (P_T, D_T) 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. It is a function of both q and price P_T, and is defined by the condition $D_T \leq q$, i.e., $aP_T^{-\eta} e^{y_T} \leq q$, which implies that $y_T \leq \eta x_T + \log(q/a)$, and thus, $y_q = \eta x_T + \log(q/a)$.

Solving for $\pi(t, q)$, we obtain a closed-form expression 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.

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

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

where

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

$$d_2 = d_1 - \sigma_z,$$

$$\sigma_z = \sqrt{\eta^2\sigma_x^2 + \sigma_y^2 - 2\eta\sigma_{xy}}.$$

The above expression for $\pi(t, q)$ is analogous to the familiar Black and Scholes (1973) option pricing formula. The parameter σ_z is the standard deviation of $y_T - \eta x_T$ 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. The profit function $\pi(t, q)$ is concave in the order quantity q. Thus, the optimal inventory level and optimal expected profit are given by the following proposition.

Proposition 2. *If the firm purchases inventory at time t, then the optimal inventory level is given by*

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

where

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

and the optimal expected profit is given by

$$\pi^*(t) = e^{-r(T-t)}E[P_T D_T|P_t, \epsilon_t]\Phi\left(d_1^* - \sigma_z\right). \tag{4}$$

Here, d_1^* is the inverse of the critical fractile similar to a newsvendor formula. Interestingly, 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.

Let $Y(t,\tau) \equiv e^{-r(\tau-t)}E[\pi^*(\tau)|P_t,\epsilon_t]$ denote the value of the firm at time t if the optimal inventory decision is taken at a *fixed* future date τ, and the information at time t is (P_t,ϵ_t). We now compute $Y(t,\tau)$ for $\tau \geq t$ by taking the expectation of $\pi^*(\tau)$ with respect to x_τ and y_τ given the information at time t, (x_t,y_t). The function Y will help us determine the value of postponing the stocking decision from time t to a later time τ.

Proposition 3. *The time t value of the expected profit if optimal inventory decision is taken at time $\tau \geq t$ is given by:*

$$Y(t,\tau) = e^{-r(T-t)}E[P_TD_T|P_t,\epsilon_t]\int_{\underline{x}_\tau}^{\infty}\Phi(d_1^{**}-\sigma_z)\frac{1}{\sqrt{2\pi}\sigma_x(t,\tau)}\exp\left[-\frac{\left(x_\tau-\mu_x(t,\tau)-k\sigma_x^2(t,\tau)\right)^2}{2\sigma_x^2(t,\tau)}\right]dx_\tau.$$

Here, d_1^{**} is the critical fractile for inventory decision taken at time τ, k is a constant as defined below, $\underline{x}_\tau$ is the smallest value of the logarithm of price at which it becomes profitable to sell:

$$\begin{aligned} d_1^{**} &= \Phi^{-1}\left(\sup\left[0, 1-\frac{c_\tau}{e^{-r(T-\tau)}E[P_T|P_\tau,\epsilon_\tau]}\right]\right), \\ k &= (1-\eta)e^{-h(T-\tau)}+\rho\frac{\sigma_y}{\sigma_x}, \\ \underline{x}_\tau &= \left[\log c_\tau + r(T-\tau)-\frac{\sigma_x^2}{2}-(m-\frac{s_P^2}{2h})(1-e^{-h(T-\tau)})\right]e^{h(T-\tau)}. \end{aligned}$$

Parameters $\mu_x(t,\tau), \mu_y(t,\tau), \sigma_x(t,\tau)$ and $\sigma_y(t,\tau)$ are as in Lemma 1.

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

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

$$Y(t,\tau) = e^{-r(T-t)}E[P_TD_T|P_t,\epsilon_t]\int_{\underline{\xi}}^{\infty}\Phi(d_1^{**}-\sigma_z)\phi(\xi)d\xi, \tag{5}$$

where

$$\underline{\xi} = \left[\left\{\log c_\tau + r(T-\tau)-\frac{\sigma_x^2}{2}-(m-\frac{s_P^2}{2h})(1-e^{-h(T-\tau)})\right\}e^{h(T-\tau)}-\mu_x(t,\tau)-k\sigma_x(t,\tau)^2\right]\frac{1}{\sigma_x(t,\tau)},$$

and

$$d_1^{**} = \Phi^{-1}\left(\max\left[0, 1-\frac{c_\tau}{\exp\{(\mu_x(t,\tau)+\sigma_x(t,\tau)\xi)e^{-h(T-\tau)}-r(T-\tau)+(m-\frac{s_P^2}{2})(1-e^{-h(T-\tau)})+\frac{\sigma_x^2}{2}\}}\right]\right),$$

the newsvendor critical fractile corresponding to the ordering time τ.

Given expressions for the optimal expected profit and the value of the firm, we can now present the structure of the optimal inventory ordering policy. The main insight is that the optimal time to purchase inventory depends on the current state of the price forecast process but is independent of the demand forecast process. The timing decision follows a threshold policy: once the price exceeds a certain (time dependent) level, it is optimal to purchase inventory. The following theorem formalizes the result.

Theorem 1. *There exists a threshold value, P_t^*, of the price forecast process such that it is optimal to purchase inventory if the price forecast at time t exceeds that threshold, and wait otherwise. That is, the optimal time to purchase inventory is given by*

$$\tau = \min\{t \in \mathcal{T}[0,T] : P_t \geq P_t^*\}.$$

Therefore, the optimal time τ to purchase inventory depends on the current state of the price forecast process but is independent of the current state of the demand forecast process.

The result of Theorem 1 is intuitive once we observe that even though the payoff from the investment in inventory is a non-linear function of price and demand, the expected value of that investment $Y(t,\tau)$ is proportional to ϵ_t. Thus, the value of the option to invest scales perfectly with the demand forecast, and decision to invest is driven by the margin and safety stock considerations, but is independent of the demand forecast.

Having characterized the optimal order policy, we can now define the value of order timing flexibility. We define it as the percentage difference in the firm's value under the two scenarios: The first scenario is based on our model, i.e., the firm is free to order inventory at any moment $\tau \in [0,T]$ taking the most recent price and demand information. In the second scenario, the firm solves for the best procurement time τ at time 0 using the available information at that time and commits to making the order at time τ. That is, in the second scenario, the firm follows the optimal static policy. Note that it still uses the optimal ordering quantity in each state (P_τ, ϵ_τ), and still retains the option of not buying inventory at all. Thus, the only difference in the two scenarios is in timing flexibility—the procurement time τ is dynamically determined in scenario one, and is determined in advance in scenario two.

Denote the value of timing flexibility by V_F:

$$\begin{aligned} V_F &= V(0) \Big/ \max_{\tau \in [0,T]} E[\pi^*(\tau)|P_0, \epsilon_0] - 1 \qquad (6) \\ &= \sup_{\tau \in \mathcal{T}[0,T]} E[\pi^*(\tau, P_\tau, \epsilon_\tau)|P_0, \epsilon_0] \Big/ \max_{\tau \in [0,T]} E[\pi^*(\tau)|P_0, \epsilon_0] - 1. \end{aligned}$$

We conclude this section by noting that Theorem 1 is key for a practical implementation of the ordering policy. It reduces the dimensionality of the problem by eliminating the dependence on the demand process. In the following section, we assess the value of timing flexibility.

5 Comparative Statics

Our model answers questions that are useful for buyers and suppliers. Using this model, a buying firm can determine when and whether to purchase inventory and in what quantity. A supplier can determine the lead-times and corresponding procurement costs to offer to retailers so that they are economically viable. In this section, we first analytically show the sources of the value of timing flexibility, and then turn to a numerical analysis of the value of optimally timing inventory orders, its mediating factors, and ordering times.

5.1 Sources of Value of Timing Flexibility

We begin with a stylized example of an option to postpone inventory purchasing from a fixed time t to a fixed time $\tau > t$ and two specific cases of the price process, constant price and GBM price forecast. The option to postpone to a fixed time has limited timing flexibility. Nonetheless, this stylized example yields a lower bound on the value of flexibility. It also helps us illustrates three sources of the value of flexibility, namely, (i) a volatility effect due to the reduction in the variance of price and demand forecasts, (ii) an optionality effect in ordering due to the option to shut down, and (iii) a newsvendor effect due to the ability to fine tune purchase quantities in response to the revealed price and demand information.

In the case of constant prices, equations (4) and (5) reduce to $\pi^*(t) = e^{-r(T-t)} D_t e^{\alpha_\epsilon(T-t)+m}$ $\Phi\left(d_1^* - s_\epsilon\sqrt{T-t}\right)$ and $Y(t,\tau) = e^{-r(T-t)} D_t e^{\alpha_\epsilon(T-t)+m} \Phi\left(d_1^{**} - s_\epsilon\sqrt{T-\tau}\right)$, where d_1^* and d_1^{**}, respectively, capture the profit margin for ordering at times t and τ . Comparing the expressions for $\pi^*(t)$ and $Y(t,\tau)$, it follows that the firm should postpone the inventory decision from time t to τ if and only if

$$\Phi\left(d_1^{**} - s_\epsilon\sqrt{T-\tau}\right) \geq \Phi\left(d_1^* - s_\epsilon\sqrt{T-t}\right).$$

Due to the monotonicity of the function Φ, the comparison is between the arguments. Thus, postponement is optimal if and only if:

$$d_1^{**} - s_\epsilon\sqrt{T-\tau} \geq d_1^* - s_\epsilon\sqrt{T-t}. \tag{7}$$

Equation (7) gives all the factors that determine the postponement decision under constant price, namely, the costs c_t and c_τ, and the demand volatility s_ϵ.

When the price forecast follows a GBM, a sufficient condition for the convergence of (1) to a GBM is $h \to 0$ s.t. $hm = \alpha_P = const.$[1] Then, the expected profit of the firm if ordering at time t and the time t value if the decision is postponed until τ are given by:

$$\pi^*(t) = e^{-r(T-t)} P_t D_t e^{[(1-\eta)\alpha_P - \eta(1-\eta)s_p^2/2 + (1-\eta)\rho s_P s_\epsilon + \alpha_\epsilon](T-t)} \Phi\left(d_1^* - \sigma_z\sqrt{T-t}\right)$$
$$Y(t,\tau) = e^{-r(T-t)} P_t D_t e^{[(1-\eta)\alpha_P - \eta(1-\eta)s_p^2/2 + (1-\eta)\rho s_P s_\epsilon + \alpha_\epsilon](T-t)} \int_{\underline{\xi}}^{\infty} \Phi(d_1^{**} - \sigma_z\sqrt{T-\tau})\phi(\xi)d\xi$$

where

$$\sigma_z^2 = \eta^2 s_P^2 + s_\epsilon^2 - 2\eta\rho s_P s_\epsilon$$
$$d_1^* = \Phi^{-1}\left(\max\left[0, 1 - \frac{c_t}{e^{(-r+\alpha_P)(T-t)} P_t}\right]\right)$$
$$d_1^{**} = \Phi^{-1}\left(\max\left[0, 1 - \frac{c_\tau}{\exp[(-r+\alpha_P)(T-t) + \{r + s_P^2/2 + \rho s_P s_\epsilon - \eta s_P^2\}(\tau - t) + \xi s_P\sqrt{\tau - t}] P_t}\right]\right)$$
$$\underline{\xi} = \frac{1}{s_P\sqrt{\tau - t}}\left[\log(c_\tau/P_t) + (r - \alpha_P^*)(T-t) - \{r + s_P^2/2 + \rho s_P s_\epsilon - \eta s_P^2\}(\tau - t)\right].$$

Thus, there is value in postponing the stocking decision from time t to time τ if the following inequality holds:

$$\int_{\underline{\xi}}^{\infty} \Phi(d_1^{**} - \sigma_z\sqrt{T-\tau})\phi(\xi)d\xi \;\geq\; \Phi\left(d_1^* - \sigma_z\sqrt{T-t}\right) \tag{8}$$

In this case, the postponement decision depends not only on the costs and s_ϵ, but also the state of the price P_t and the price volatility parameters, s_P, ρ, and η. We derive a lower bound on the value of the integral in (8). Let $a = \sigma_z\sqrt{T-\tau}$ and consider the function $\Phi(d_1^{**} - a) - (\Phi(d_1^{**}) - (\Phi(0.5a) - \Phi(-0.5a)))$. This function equals zero if $d_1^{**} = 0.5a$, and tends to a positive quantity equal to $\Phi(0.5a) - \Phi(-0.5a)$ if d_1^{**} tends to positive or negative infinity. Also, note that $|d_1^{**} - a| \geq |d_1^{**}|$ for $d_1^{**} \leq 0.5a$ and $|d_1^{**} - a| \leq |d_1^{**}|$ for $d_1^{**} \geq 0.5a$. Thus, the function decreases until $d_1^{**} = 0.5a$ and increases thereafter. Therefore, we obtain the lower bound: $\Phi(d_1^{**} - a) \geq \Phi(d_1^{**}) - (\Phi(0.5a) - \Phi(-0.5a))$. Integrating,

$$\int_{\underline{\xi}}^{\infty} \Phi(d_1^{**} - a)\phi(\xi)d\xi \geq \Phi^c(\underline{\xi})\left(1 - (\Phi(0.5a) - \Phi(-0.5a))\right) - \frac{c_\tau \Phi^c(\underline{\xi} + s_P\sqrt{t-\tau})}{P_t e^{(\alpha_P - r)(T-t) + (r + \rho s_P s_\epsilon - \eta s_P^2)(\tau - t)}}.$$

The bound illustrates three sources of value loss and gain: volatility, optionality, and the newsvendor effect. The effect of volatility on profits is captured by the term $1-(\Phi(0.5a)-\Phi(-0.5a))$,

[1] This condition 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, $m = const$ as $h \to 0$, we get a GBM with zero drift.

which decreases with σ_z because $a = \sigma_z\sqrt{T-\tau}$. The effect of the option to shut down due to the variability of price is captured by the terms $\Phi^c(\underline{\xi})$ and $\Phi^c(\underline{\xi} + s_P\sqrt{t-\tau})$. As price volatility s_P increases, $\underline{\xi}$ decreases, and the option to shut down becomes more valuable. Finally, as the price state variable P_t increases, the profit margin increases, which thus enhances the value of the option to procure inventory. Computationally, the accuracy of this bound improves as the profit margin increases, but the approximation is poor near zero profit margin.

Equations (7) and (8) also show that the drift rate α_ϵ and the information about ϵ_t have no effect on the postponement decision at time t. Under constant prices, (7) implies that the optimal timing decision is path independent. Thus, the optimal time to procure inventory is independent of time t and can be fixed in advance. Under stochastic prices, (8) implies that the optimal postponement decision at time t depends on the value of P_t, although it is independent of the sample path of ϵ_t. Thus, the optimal time to procure inventory cannot be fixed in advance, but rather obeys the threshold policy formulated by Theorem 1.

5.2 Ordering Policy Computation

This section describes our computational experiment and presents insights regarding the threshold value of price, the value of timing flexibility relative to static postponement, and the implications of the option being in-the-money or out-of-the-money. Our computation of the price thresholds P_t^*, distributions of ordering time τ, and value of timing flexibility V_F are based on a discretized version of problem (3). We simulate the processes $\{x_t\}$ and $\{\epsilon_t\}$ with binary trees, and check whether or not to buy inventory in each state (P_t, ϵ_t), by comparing $\pi^*(t, P_t, \epsilon_t)$ and $V(t)$. We compute $P_t^* = \min\{P_t : \pi^*(t, P_t, \epsilon_t) \geq V(t)\}$ and verify that the purchasing policy is indeed a threshold type in P_t.

To generate price and demand forecast trajectories, we set $T = 0.1$ and use a time discretization step equal to 0.005 resulting in a binary tree with 20 levels, representing periods over the interval $[0, T]$. We use a GBM price process with the initial price forecast $P_0 = 1$ and three levels of price volatility $s_P = \{0.2, 0.5, 1\}$, corresponding to low, medium, and high price volatility scenarios. Similarly, we consider three levels of the initial cost $c_0 = \{0.8, 1, 1.2\}$, corresponding to profitable, marginally profitable, and non-profitable conditions at time $t = 0$. Scenarios with $c_0 = \{1, 1.2\}$ highlight the value of timing flexibility and price volatility. *Ex-ante*, it is not profitable to procure at costs $c_0 = \{1, 1.2\}$ at $t = 0$, i.e., the option to invest in the newsvendor is out-of-the-money at time 0 for $c_0 = \{1, 1.2\}$, yet, similarly to financial options, it derives positive value in the

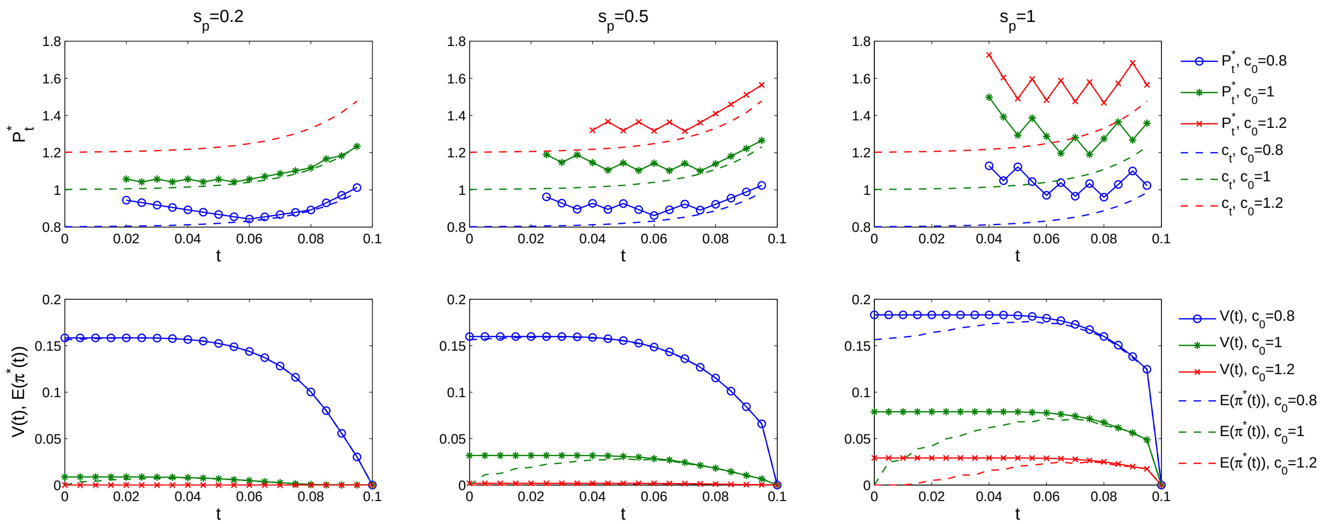

Figure 1: *Top row*: Optimal price thresholds P_t^* for the initial costs $c_0 = \{0.8, 1, 1.2\}$, and price volatilities $s_p = \{0.2, .5, 1\}$. Costs increase exponentially over time, the trajectories are shown in dashed lines.
Bottom row: Value of the option to invest $V(t)$ under the optimal ordering policy (solid lines) and expected profit $E(\pi^*(t))$ given that an optimal inventory decision is taken at time t (dashed lines) for the initial costs $c_0 = \{0.8, 1, 1.2\}$, and price volatilities $s_p = \{0.2, .5, 1\}$.
Parameter values: $T = 0.1,\ r = 0.01,\ h = 0,\ m = \ln P_0 = 0,\ \rho = 0,\ \alpha_\epsilon = 0,\ s_\epsilon = 1,\ \eta = 0,\ \lambda = 0.$

presence of price volatility. We allow unit purchasing cost to increase exponentially over time with $c_t = c_0(1 + e^{kt}/K)$, where $k = 50$ and $K = 500$ for $t \in [0, 0.95T]$. If $t > 0.95T$, we set $c_t = \infty$ to prevent perfect matching between supply and demand. We set the initial demand signal $D_0 = 1$. The time value of money $r = 0.01$.

The above parameters give 9 instances of the problem covering all combinations of s_P and c_0. For each instance, we plot the price threshold P_t^*, the value of the option to invest at time t, $V(t)$, and, for comparison, the expected profit $E[\pi^*(t)|P_0, \epsilon_0]$ if the procurement decision is made at time t. Both the value to invest and the expected profit are discounted to time 0.

Figure 1 shows a summary of the results plotted against time t. In the upper panel, we observe that the price thresholds P_t^* are non-monotone: they decrease at first despite increasing costs. Note that P_t^* does not exist if t is small. That shows a strong postponement effect in which it is optimal to postpone the decision in all states of the price process in our binary tree. When price volatility is low, the postponement effect is observed when the option is out-of-the-money, but when price volatility is high, the postponement effect is observed regardless of whether the option is in- or out-of-the-money. In the lower panel, the gap between $V(t)$ and $E[\pi^*(t)|P_0, \epsilon_0]$ shows the value of

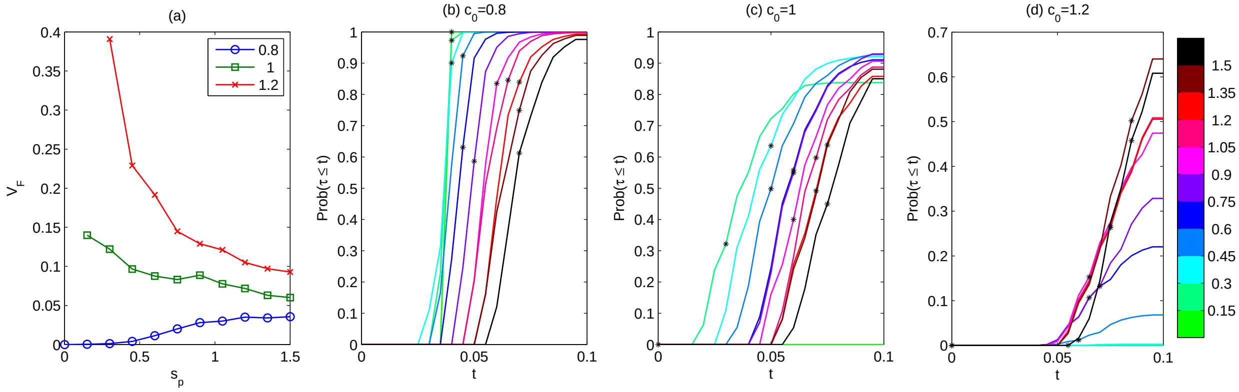

Figure 2: (a): Value of timing flexibility as a function of s_p for the initial costs $c_0 = \{0.8, 1, 1.2\}$.
(b)-(d): CDF of inventory order timing for $s_p = \{0, 0.15, 0.3, ..., 0.15\}$ and the initial costs $c_0 = \{0.8, 1, 1.2\}$. Markers show the optimal static ordering times for each s_P. The colorbar represents s_P.
Parameter values: $T = 0.1$, $r = 0.01$, $h = 0$, $m = \ln P_0 = 0$, $\rho = 0$, $\alpha_\epsilon = 0$, $s_\epsilon = 1$, $\eta = 0$, $\lambda = 0$, $c_t = c_0(1 + e^{50t}/500)$, if $t \in [0, 0.95T], c_t = \infty$ if $t > 0.95T$.

timing flexibility in absolute terms. This gap decreases with t and increases with s_p. Thus, the value of timing flexibility is higher for longer time horizons and more volatile prices. Furthermore, observe that the gap between $V(t)$ and $E[\pi^*(t)|P_0, \epsilon_0]$ is the largest for $c_0 = 1$. This shows that the value of timing flexibility is large if the option is neither in- nor out-of-the-money, and is smaller for deep in- or out-of-the-money options. A comparison across figures shows that the value of the option to invest $V(t)$ increases with s_p.

Finally, note that in all instances, the rate of cost increase is substantially higher than the risk free rate. This suggest that offering timing flexibility can be economically attractive for suppliers. We tested several alternative cost trajectories. If costs are constant, then waiting until the end is optimal. If costs increase linearly, then the optimal strategy is to either execute immediately or wait till the end.

5.3 Effect of Price Volatility on Optimal Timing

In this section, we study the value of timing flexibility V_F and estimate the distribution of optimal ordering times as functions of price volatility s_P. We vary s_P from 0 to 1.5 and keep the rest of the parameters the same as in Section 5.2.

Figure 2 shows the results obtained. If $s_P = 0$, i.e., if prices are constant, then the value of timing flexibility is zero because it is optimal to postpone ordering to a fixed time in the future if $c_0 < P_0$, or not order at all if $c_0 \geq P_0$. The latter is the option to shutdown. As s_P increases,

V_F increases if the option to invest is in the money (i.e., if $c_0 = 0.8$) and decreases otherwise (i.e., if $c_0 = \{1, 1.2\}$, see Figure 2(a)). This happens because the numerator of (6) increases with s_P at about the same rate regardless of whether the option is in- or out-of-the-money, whereas the denominator for out-of-the-money options is initially very small, but increases quickly with s_p (see Figure 1 for an illustration of the effect of s_P on $E(\pi^*(t)|P_0, \epsilon_0)$). We find that, for the in-the-money option, V_F ranges between $0 - 4\%$ increasing with s_p, and for the out-of-the-money option, it ranges between $6 - 35\%$ decreasing with s_P. This shows another evidence that timing flexibility is particularly valuable when the profit *ex-ante* is close to zero or even negative.

We plot cumulative distributions of the optimal order time on the panels (b)-(d) of Figure 2. Each line represents a cumulative distribution function (CDF) color-coded according to the value of s_P. Dots on the lines represent best static ordering times for the respective values of s_P. The ordering of the CDFs shows that, as s_P increases, it is optimal to postpone the ordering decision more. This effect persists for all levels of c_0. Interestingly, we find that the average order time under timing flexibility is smaller than the optimal static order time. That is, offering flexibility induces earlier orders on average. Instances where order timing under flexibility is earlier than the optimal static order time correspond to the sample paths with higher prices. On the contrary, if prices along a sample path are low, it may be optimal to postpone the purchase more, or not order at all. The distributions of ordering times for $c_0 = \{1, 1.2\}$ illustrate this pattern. The optimal order times for these costs follow a defective distribution in which there is a non-zero probability that no inventory will be procured at all. We call these distributions defective because they have a point mass at infinity corresponding to an infinite lead time when no order is placed; we compute the means and variances of these distributions conditionally. Finally, we find that order times become more variable with increase in s_p if the *ex-ante* profit margin is positive. If the *ex-ante* profit margin is negative, then the order time variability is greater, and can be a non-monotonic, inverted U-shaped, function of s_P.

To summarize, we find that the value of timing flexibility can be substantial when costs are high, and the optimal order lead time exhibits a large variation.

5.4 Effect of Demand Volatility on Optimal Timing

Similarly to the effect of price volatility, in this section, we investigate the effect of demand volatility s_ϵ on the value of timing flexibility. We vary s_ϵ from 0 to 1.5, set $s_P = 1$, and keep the rest of the parameters the same as in Section 5.2. Intuitively, when s_ϵ is high, there is more incentive to

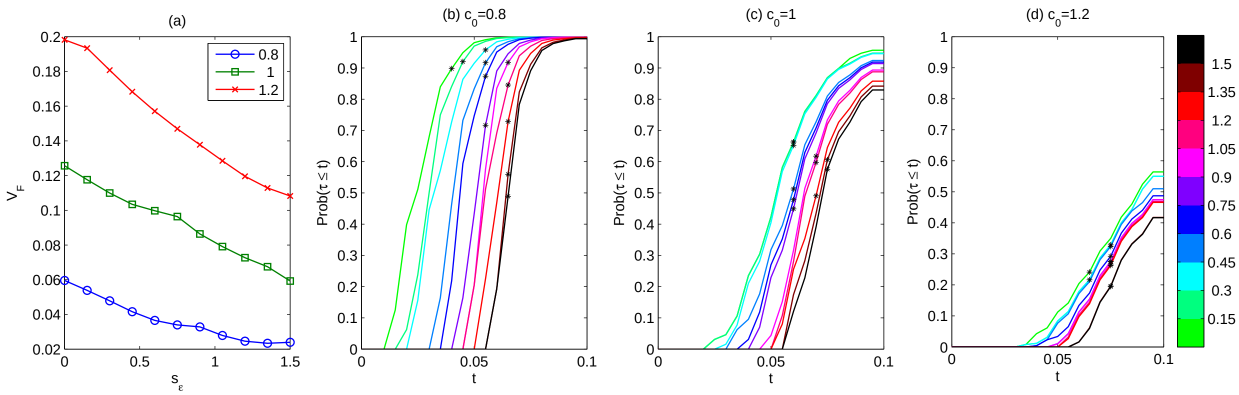

Figure 3: (a): Value of timing flexibility as a function of s_ϵ for the initial costs $c_0 = \{0.8, 1, 1.2\}$.
(b)-(d): CDF of inventory order timing for $s_\epsilon = \{0, 0.15, ..., 1.5\}$ and the initial costs $c_0 = \{0.8, 1, 1.2\}$. Markers show the optimal static ordering times for each s_ϵ. The colorbar represents s_ϵ.
Parameter values: $T = 0.1,\ r = 0.01,\ h = 0,\ m = \ln P_0 = 0,\ s_p = 1,\ \alpha_\epsilon = 0,\ \rho = 0,\ \eta = 0,\ \lambda = 0,$ $c_t = c_0(1 + e^{50t}/500),$ if $t \in [0, 0.95T], c_t = \infty$ if $t > 0.95T$.

postpone ordering decision. However, the effect on value of timing flexibility is less obvious.

We find that V_F decreases with s_ϵ. The intuition for this result is that when s_ϵ is large compared to s_P, V_F is driven by the reduction of demand uncertainty, which is the source of value for static postponement. Therefore, the additional value derived from order timing flexibility decreases, and as the result, V_F decreases with s_ϵ. We also find that the order lead time distributions are stochastically increasing in s_ϵ. This is similar to the effect of s_P. Increasing s_ϵ provides more incentives to postpone the procurement decision. However, unlike s_P, the variance of order lead times decreases as s_ϵ increases because static postponement becomes progressively more valuable as s_ϵ increases. We also observe that V_F increases with the initial procurement cost c_0. Thus, demand volatility increases the value of timing flexibility for the out-of-the-money options. This is similar to the effect of price volatility. In addition, order timings become more variable and the probability of no order placement increases when c_0 is high.

Overall, the combined effect of price and demand volatility on V_F is the following: whereas s_P can amplify or decrease the value of timing flexibility, depending on the intrinsic value of the option, s_ϵ decreases the value of timing flexibility regardless of the intrinsic value of the option to invest in a newsvendor contract.

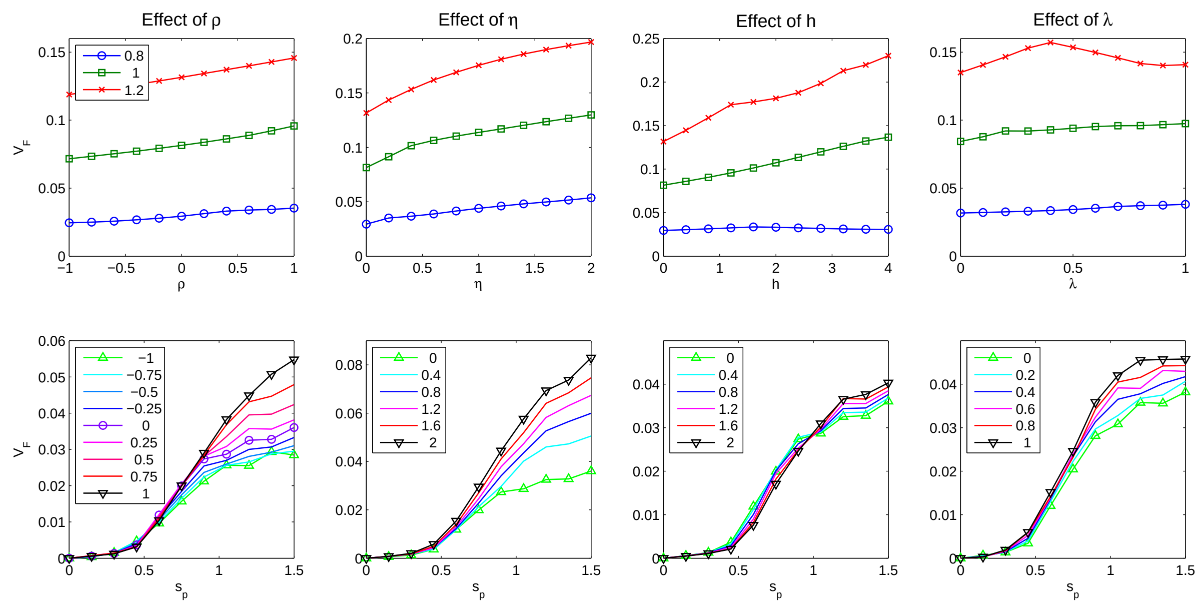

Figure 4: *Top row*: Value of timing flexibility as a function of ρ, η, h, and λ for the initial costs $c_0 = \{0.8, 1, 1.2\}$. *Bottom row*: Value of timing flexibility as a function of s_P for $c_0 = 0.8$ and sets of possible values of ρ, η, h, and λ. Triangular markers indicate minimum and maximum values of ρ, η, h, or λ.
Baseline parameter values: $T = 0.1,\ r = 0.01,\ h = m = \ln P_0 = 0,\ s_P = s_\epsilon = 1,\ \alpha_\epsilon = \rho = \eta = \lambda = 0,$ $c_t = c_0(1 + e^{50t}/500)$, if $t \in [0, 0.95T], c_t = \infty$ if $t > 0.95T$. For the effect of λ: $\rho_{Pm} = \rho_{\epsilon m} = 0.5, \rho = 0.25$.

5.5 Mediating Effects

The value of timing flexibility is further affected by the price-demand elasticity η, the correlation between price and demand forecasts ρ, the speed of price mean reversion h, and market price of risk λ. In this section, we study the sensitivity of V_F to these parameters and assess how these parameters mediate the effect of s_p on V_F. Figure 4 presents four sets of plots, two plots per parameter.

We find that the value of timing flexibility increases with ρ. A positive correlation between price and demand forecasts increases the value of timing flexibility, and a negative correlation decreases that value. Negative correlation between price and demand forecasts reduces the variance of the expected profit, thus, the value of timing flexibility is also reduced. The opposite happens for positive ρ. The effect of ρ is amplified when s_P is large.

A higher price demand elasticity η also amplifies the value of timing flexibility. Although the result is similar to the effect of ρ, the mechanism though which η affects V_F is different. Equation (8) may clarify. There, the effect of ρ depends on its sign, whereas the effect of η depends on the sign of $\eta s_P - 2\rho s_\epsilon$. In our example, $\rho = 0$, so that the value of timing flexibility increases with η. However, if ρ is a large positive number and price volatility is low, then a high value of η creates a natural hedge on revenue and decreases the value of timing flexibility.

The direction of the effect of mean reversion h on V_F depends on the price volatility s_P. The value of timing flexibility decreases in h if s_P is small, but increases in h if s_P is large. This effect occurs because a higher price volatility enables a broader range of opportunities, but also increases the risk of low prices. A high mean reversion ensures that price forecasts do not stay low for too long. When price forecasts revert to higher values, the firm may find it profitable to invest, thus the value of timing flexibility increases.

A higher market risk premium λ increases the value of timing flexibility, but the effect is non-monotone for out-of-the-money options. For such options, the probability of not buying inventory at all increases with λ, which leads to non-monotonicity. Increasing price volatility reduces the probability of foregoing inventory purchase and shifts the maximum of $V_F(\lambda)$ to the right.

The increasing value of timing flexibility with λ suggests that timing flexibility would be more valuable when the price is more positively correlated with the market. It is useful to relate the correlation of price with the market with the role of financial hedging 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 timing flexibility 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 flexibility is enhanced since the investors are willing to take a lower price risk.

Commonly to all parameters, V_F is greater for out-of-the-money options. These options derive value from price volatility, and increase in value is greater if one has complete time flexibility to execute them. Also the effects of ρ, η, and λ are amplified when price volatility s_P is large.

To investigate the effect of ρ, η, h, and λ on optimal order timing, we compute cumulative distribution functions of optimal order times. The plots are presented for $c_0 = 0.8$, i.e., in-the-money option (Figure 5) with *ex-ante* profit margin of 20%. The results are similar for the out-of-the-money options with the caveat that the probability of not buying inventory at all increases, resulting in defective probability distributions. We find that high positive ρ and high h shift orders to earlier times, high η and low h increase variability of order times, and high λ shifts orders towards later times. The result for ρ does not contradict that $V_F(\rho)$ is increasing. The price threshold P_t^*

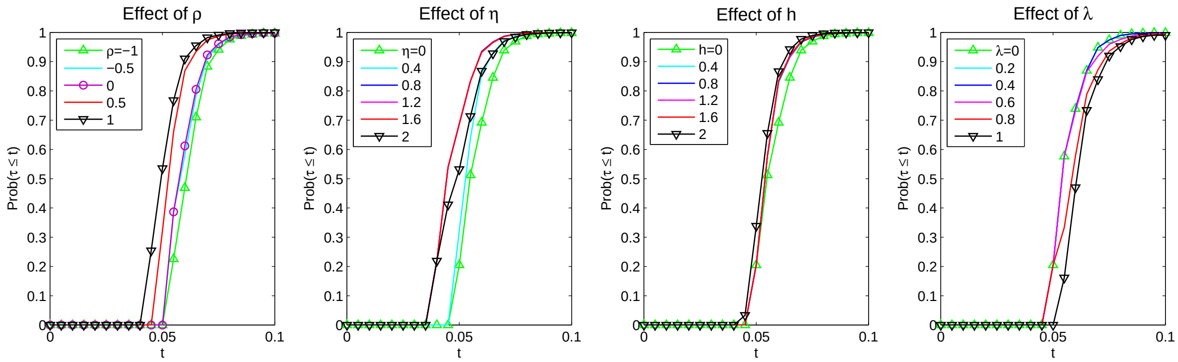

Figure 5: Cumulative distributions of optimal order times for $c_0 = 0.8$ and sets of $\rho, \eta, h,$ and λ. Triangular markers indicate minimum and maximum values of $\rho, \eta, h,$ or λ.
Baseline parameter values: $T = 0.1,\ r = 0.01,\ h = m = \ln P_0 = 0,\ s_P = s_\epsilon = 1,\ \alpha_\epsilon = \rho = \eta = \lambda = 0,$ $c_t = c_0(1 + e^{50t}/500)$, if $t \in [0, 0.95T], c_t = \infty$ if $t > 0.95T$. For the effect of λ: $\rho_{Pm} = \rho_{\epsilon m} = 0.5, \rho = 0.25$.

for early execution is high. Thus it shows that it is optimal to forego benefits of updated forecasts of price and demand, and buy early if forecasted prices are high enough.

Finally, note that the risk premium has no effect on the timing decision if the selling price is constant. From (7), note 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 when prices are constant but different decisions if prices are stochastic. This result follows from the fact that the optimal timing decision is independent of the demand forecast sample path. Thus, when the risk premium affects the discount rate from time T to time t, it has the same proportional effect on $Y(t, \tau)$ and $\pi^*(t)$, and therefore, has no bearing on the postponement decision.

6 Conclusion

This paper studies the problem of optimally timing an inventory purchasing decision when demand and price are both stochastic. The demand and price forecasts get progressively more accurate with time, but the unit purchasing cost also increases with time. The model presented in this paper contributes to the two distinct streams of literature on inventory timing decisions and real options. The former literature has recognized that demand uncertainty affects the optimal inventory timing decision. We contribute to this literature by showing the consequences of price uncertainty as well as the correlation between the demand and price variables. With respect to the real options literature, we develop a model of a complex American-style option with an endogenized newsvendor profit function. We identify sources of the value of the option and numerically analyze its comparative statics.

We find that the optimal order policy requires flexibility in timing, i.e., it is optimal to place an order if and only if the price variable exceeds a certain time dependent threshold. Thus, the optimal order timing is random. This highlights the distinction between order timing flexibility and postponement of ordering decision to a fixed time. Order timing flexibility adds value through demand volatility reduction, adjusting the order quantity in response to price fluctuations and increased costs, and through the option of postponing ordering decision and possibly not ordering at all.

Benchmarking the performance of the time flexible ordering policy with respect to the optimal static postponement policy, we find that the time flexible ordering can substantially increase the value of the firm. If business conditions are *ex-ante* profitable, timing flexibility can increase the value of the firm by up to 4%, but if the *ex-ante* margin is close to zero or negative, the increase can be much larger (up to 30-40%). In these conditions, volatile prices provide opportunity to substantially increase profit margin, if the firm has the order timing flexibility. We further find that the relative value of timing flexibility increases in price volatility, but decreases in demand volatility. In other words, high demand volatility reduces the benefit of reacting to price changes over time.

Comparing the optimal timings of inventory decisions, we find that time flexible ordering policy can be beneficial for suppliers in two ways. First, it can induce earlier purchases compared to the optimal static postponement policy depending on the cost profile and the evolution of the price forecast. The effect is more pronounced when the business is *ex-ante* profitable. Second, suppliers may charge substantially higher premiums for quick deliveries by providing the option of flexible procurement to buyer firms. Thus, price volatile environments can be substatially beneficial for both suppliers and buyers adopting time flexible ordering policies.

References

Aviv, Yossi, Awi Federgruen. 2001. Design for postponement: a comprehensive characterization of its benefits under unknown demand distributions. *Operations Research* **49**(4) 578–598.

Bernstein, Fernando, Awi Federgruen. 2005. Decentralized supply chains with competing retailers under demand uncertainty. *Management Science* **51**(1) 18–29.

Birge, John R. 2000. Option methods for incorporating risk into linear capacity planning models. *Manufacturing & Service Operations Management* **2**(1) 19–31.

Black, Fischer, Myron Scholes. 1973. The pricing of options and corporate liabilities. *The journal of political economy* 637–654.

Caballero, Ricardo J, Emmanuel Farhi, Pierre-Olivier Gourinchas. 2008. Financial crash, commodity prices, and global imbalances. *Brookings Papers on Economic Activity* (2).

Chod, Jiri, Nils Rudi, Jan A Van Mieghem. 2010. Operational flexibility and financial hedging: Complements or substitutes? *Management Science* **56**(6) 1030–1045.

Constantinides, George M. 1978. Market risk adjustment in project valuation. *The Journal of Finance* **33**(2) 603–616.

Dixit, Avinash K, Robert S Pindyck. 1994. *Investment under uncertainty*. Princeton university press.

Federgruen, Awi, Aliza Heching. 1999. Combined pricing and inventory control under uncertainty. *Operations research* **47**(3) 454–475.

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

Guo, Xin, Philip Kaminsky, Pascal Tomecek, M Yuen. 2011. Optimal spot market inventory strategies in the presence of cost and price risk. *Mathematical Methods of Operations Research* **73**(1) 109–137.

Haksöz, Ç, Sridhar Seshadri. 2007. Supply chain operations in the presence of a spot market: a review with discussion. *Journal of the Operational Research Society* **58**(11) 1412–1429.

Kaminsky, Philip, Jayashankar M Swaminathan. 2004. Effective heuristics for capacitated production planning with multiperiod production and demand with forecast band refinement. *Manufacturing & Service Operations Management* **6**(2) 184–194.

Kobylanski, Magdalena, Marie-Claire Quenez, Elisabeth Rouy-Mironescu. 2011. Optimal multiple stopping time problem. *The Annals of Applied Probability* 1365–1399.

Kouvelis, Panos. 1999. Global sourcing strategies under exchange rate uncertainty. *Quantitative Models for Supply Chain Management*. Springer, 625–667.

Kouvelis, Panos, Zhongjun Tian. 2014. Flexible capacity investments and product mix: Optimal decisions and value of postponement options. *Production and Operations Management* **23**(5) 861–876.

Lee, Hau L, Christopher S Tang. 1997. Modelling the costs and benefits of delayed product differentiation. *Management science* **43**(1) 40–53.

Li, Chung-Lun, Panos Kouvelis. 1999. Flexible and risk-sharing supply contracts under price uncertainty. *Management Science* **45**(10) 1378–1398.

McDonald, Robert L, Daniel R Siegel. 1985. Investment and the valuation of firms when there is an option to shut down. *International economic review* 331–349.

McDonald, Robert L, Daniel R Siegel. 1986. The value of waiting to invest. *The Quarterly Journal of Economics* **101**(4) 707–728.

Merton, Robert C. 1973. An intertemporal capital asset pricing model. *Econometrica: Journal of the Econometric Society* 867–887.

Oksendal, Bernt. 2013. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media.

Osadchiy, Nikolay, Vishal Gaur, Sridhar Seshadri. 2013. Sales forecasting with financial indicators and experts' input. *Production and Operations Management* **22**(5) 1056–1076.

Osadchiy, Nikolay, Vishal Gaur, Sridhar Seshadri. 2015. Systematic risk in supply chain networks. *Management Science, forthcoming* .

Petruzzi, Nicholas C, Maqbool Dada. 1999. Pricing and the newsvendor problem: A review with extensions. *Operations Research* **47**(2) 183–194.

Piros, Christopher D. 1998. The perfect hedge: Quanto or not to quanto. *Currency derivative* 340–354.

Schwartz, Eduardo S. 1997. The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance* **52**(3) 923–973.

Secomandi, Nicola, Sunder Kekre. 2014. Optimal energy procurement in spot and forward markets. *Manufacturing & Service Operations Management* **16**(2) 270–282.

Whang, S, H L Lee. 1998. *Value of Postponement*, vol. 10. In Product variety management: research advances, T.H.Ho, C.S.Tang (Eds), Springer Science & Business Media.

Xiao, Guang, Nan Yang, Renyu Zhang. 2015. Dynamic pricing and inventory management under fluctuating procurement costs. *Manufacturing & Service Operations Management* .

A Proofs

Proof of Lemma 1. The expressions for $\mu_x(\tau,t)$, $\sigma_x^2(\tau,t)$, $\mu_y(\tau,t)$ and $\sigma_y^2(\tau,t)$ are straightforward (for example, see Oksendal 2013, Chapter 5). The expression for $\sigma_{xy}(\tau,t)$ can be computed directly using the integral representation of processes $\{x_t\}$ and $\{y_t\}$:

$$x_\tau = x_t e^{-h(\tau-t)} + m\left(1 - e^{-h(\tau-t)}\right) + e^{-h\tau}\int_t^\tau s_P e^{hu}(dz_P)_u$$
$$y_\tau = y_t + \alpha_\epsilon(\tau - t) + s_\epsilon \int_t^\tau (dz_\epsilon)_u.$$

Using Ito isometry:

$$\begin{aligned}
\sigma_{xy}(\tau,t) &= Cov(x_\tau, y_\tau | x_t, y_t) \\
&= E_t\left[(x_\tau - E_t(x_\tau))(y_\tau - E_t(y_\tau))\right] \\
&= E_t\left[e^{-h\tau}\int_t^\tau s_P e^{hu}(dz_P)_u \cdot s_\epsilon \int_t^\tau (dz_\epsilon)_u\right] \\
&= s_P s_\epsilon e^{-h\tau}\int_t^\tau e^{hu}\rho du = \frac{\rho s_P s_\epsilon}{h}\left(1 - e^{-h(\tau-t)}\right).
\end{aligned}$$

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

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

where $\rho_{xy} = \sigma_{xy}/(\sigma_x\sigma_y)$. ■

Proof of Proposition 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:

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

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

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

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:

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

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:

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

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. ■

Proof of Proposition 2. 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. ■

Proof of Proposition 3. Given information (x_t,y_t) at time t, the distribution of (x_τ,y_τ) is bivariate normal. The conditional means $\mu_x(\tau,t)$ and $\mu_y(\tau,t)$ variances $\sigma_x^2(\tau,t)$ and $\sigma_y^2(\tau,t)$, and covariance $\sigma^2_{xy}(\tau,t)$. are given by Lemma 1. We integrate $\pi^*(t)$ with respect to the conditional distribution of (x_t, y_t). This gives the expression for $Y(t,\tau)$ given in Proposition 2. Unlike Lemma 1 and Proposition 1, a closed form expression for $Y(t,\tau)$ cannot be obtained due to the finite lower limit $\underline{x}_\tau$ on the integral over values of x_τ. ■

Proof of Theorem 1. Recall, $V(t)$ is the value of the option to invest in the newsvendor firm at time $t \in [0,T]$ given information (P_t, ϵ_t). We have

$$V(t) \equiv V(t, P_t, \epsilon_t) = \sup_{\tau \in \mathcal{T}[t,T]} \{E[\pi^*(\tau, P_\tau, \epsilon_\tau)|P_t, \epsilon_t]\}.$$

Here, $\pi^*(\tau, P_\tau, \epsilon_\tau)$ is the expected profit at time τ given information (P_τ, ϵ_τ) if the firm exercises the option to invest in the newsvendor at time τ. Using expression (5) and Proposition 2:

$$\begin{aligned}
\pi^*(\tau, P_\tau, \epsilon_\tau) &= e^{-r(T-\tau)} E[P_T D_T | P_\tau, \epsilon_\tau] \Phi\left(d_1^* - \sigma_z\right) \\
Y(t, P_t, \epsilon_t, \tau) &= E[\pi^*(\tau, P_\tau, \epsilon_\tau)|P_t, \epsilon_t] = e^{-r(T-t)} E[P_T D_T | P_t, \epsilon_t] \int_{\underline{\xi}}^{\infty} \Phi(d_1^{**} - \sigma_z)\phi(\xi) d\xi \\
&= e^{-r(T-t)} a e^{\mu_y + \sigma_y^2/2 + (1-\eta)\mu_x + (1-\eta)^2\sigma_x^2/2 + (1-\eta)\sigma_{xy}} \int_{\underline{\xi}}^{\infty} \Phi(d_1^{**} - \sigma_z)\phi(\xi) d\xi \\
&= \epsilon_t \psi(t, P_t, \tau),
\end{aligned}$$

where $\psi(t, P_t, \tau)$ is a function independent of ϵ_t. Indeed, per Lemma 1 $E[P_T D_T | P_t, \epsilon_t]$ is proportional to ϵ_t, and the integral does not depend on ϵ_t. Note that

$$\pi^*(t, P_t, \epsilon_t) = Y(t, P_t, \epsilon_t, t) = \epsilon_t \psi(t, P_t, t) \text{ and } V(t, P_t, \epsilon_t) = \sup_{\tau \in \mathcal{T}[t,T]} \{Y(t, P_t, \epsilon_t, \tau)\} = \epsilon_t \sup_{\tau \in \mathcal{T}[t,T]} \{\psi(t, P_t, \tau)\}.$$

Thus, $\pi^*(t)$ and $V(t)$ scale linearly in ϵ_t. Now, consider two distinct states at time t, (P_t^*, ϵ_t) and (P_t^*, ϵ_t'), such that the optimal policy is to purchase inventory in state (P_t^*, ϵ_t) and to wait in state (P_t^*, ϵ_t'). Thus, $V(t, P_t^*, \epsilon_t) = \pi^*(t, P_t^*, \epsilon_t)$ and $V(t, P_t^*, \epsilon_t') > \pi^*(t, P_t^*, \epsilon_t')$. But $V(t, P_t^*, \epsilon_t) = \epsilon_t \sup_{\tau \in \mathcal{T}[t,T]} \{\psi(t, P_t, \tau)\} = \epsilon_t V(t, P_t^*, \epsilon_t')/\epsilon_t' > \epsilon_t \pi^*(t, P_t^*, \epsilon_t')/\epsilon_t' = \pi^*(t, P_t^*, \epsilon_t)$. This gives a contradiction. Hence, proved. ■