To Wait or Not to Wait: Optimal Ordering Under Lead Time Uncertainty and Forecast Updating

Yimin Wang,[1] Brian Tomlin[2]

[1] *W. P. Carey School of Business, Arizona State University, Tempe, Arizona 85287*

[2] *Tuck School of Business at Dartmouth, Hanover, New Hampshire 03755*

Received 9 October 2007; revised 23 July 2009; accepted 9 August 2009
DOI 10.1002/nav.20381
Published online 7 October 2009 in Wiley InterScience (www.interscience.wiley.com).

Abstract: There has been a dramatic increase over the past decade in the number of firms that source finished product from overseas. Although this has reduced procurement costs, it has increased supply risk; procurement lead times are longer and are often unreliable. In deciding when and how much to order, firms must consider the lead time risk and the demand risk, i.e., the accuracy of their demand forecast. To improve the accuracy of its demand forecast, a firm may update its forecast as the selling season approaches. In this article we consider both forecast updating and lead time uncertainty. We characterize the firm's optimal procurement policy, and we prove that, with multiplicative forecast revisions, the firm's optimal procurement time is independent of the demand forecast evolution but that the optimal procurement quantity is not. This leads to a number of important managerial insights into the firm's planning process. We show that the firm becomes less sensitive to lead time variability as the forecast updating process becomes more efficient. Interestingly, a forecast-updating firm might procure earlier than a firm with no forecast updating.
 Naval Research Logistics 56: 766–779, 2009

Keywords: stochastic lead time; forecast updating; supply chain

1. INTRODUCTION

The pressure to reduce direct manufacturing costs has led many firms to outsource production to lower-cost countries. As a result, North American and European firms are faced with longer procurement lead times. Moreover, the procurement lead times are often uncertain. Delays can occur for many reasons, including transportation-infrastructure issues in rapidly-developing economies, congestion in foreign and domestic ports, customs inspections, and logistical issues involving export quotas. Long and uncertain lead times are especially problematic for firms with short selling seasons. One of the authors previously worked in the custom-design drawnwork (textile) industry. For reasons of cost, US drawnwork wholesalers source from Chinese suppliers. The typical lead time is on the order of 3 months but there is significant uncertainty around this for the reasons cited earlier. The US wholesalers sell the product to domestic customers in the fall but, because of the lead time, have to source the drawnwork before their customers place orders. In essence, the US wholesalers face a newsvendor-type problem with an uncertain lead time.

Firms from many industries face a similar problem. The issue of supply uncertainty is of growing managerial concern, as evidenced from the following excerpts from reports by the "The Economist and The Boston Consulting Group (BCG):"

> Last autumn some 80m items of clothing were impounded at European ports and borders because they exceeded the annual import limits that the European Union and China had agreed on only months earlier. Retailers had ordered their autumn stock well before that agreement was signed, and many were left scrambling. (When the Chain Breaks. The Economist. 2006.)

> In the run up to Christmas 2004, grid-lock hit the Los Angeles-Long Beach ports, the entry point for almost half the goods coming into the United States. Nearly 100 ships floated around, cooling their keels and waiting to be unloaded – a process that was taking up to twice as long as usual. The results of the dock jam were serious and far reaching.

Additional Supporting Information may be found in the online version of this article.

Correspondence to: Y. Wang (yimin_wang@asu.edu)

> The sharper image, for one, claimed that November sales had been adversely affected by reduced inventory resulting from congestion at the ports. Now companies are ordering earlier. (Avoiding Supply Chain Shipwrecks. BCG. 2005)

Ordering earlier reduces the "lateness" risk associated with uncertain lead times but it also increases the firms' demand risk, i.e., the potential mismatch between the quantity procured and the realized demand. In industries with short selling seasons, firms often use various sources of information to improve the accuracy of demand forecasts. Examples of such sources include pre-season orders, sales force interactions with customers, trade shows, and market trend reports. The demand forecast improves as more information is obtained over time and, therefore, the longer the firm waits before procuring the product, the more information it has and its demand risk decreases. Thus, the firm faces a trade-off: order earlier to reduce supply risk or order later to reduce demand risk.[1] A primary purpose of this research is to investigate this trade-off.

In this article, we characterize the firm's optimal procurement timing-and-quantity decision under supply risk and forecast updating. We also determine how supply and market attributes influence the firm's optimal procurement time and optimal expected cost. This leads to a number of interesting findings. We prove that, with multiplicative forecast revisions, the firm can determine its optimal procurement time at the start of its planning horizon but that its optimal procurement quantity will depend on the particular realization of the evolving demand forecast. One might expect that a forecast-updating firm would procure closer to the selling season to take advantage of increasing demand-forecast accuracy. We prove that this is not necessarily true. There are situations in which, all else being equal, a forecast-updating firm will place its order earlier than a firm with no forecast updating. We prove that this can only occur if the lead time is uncertain. Furthermore, we show that, as its forecast updating process becomes more efficient, a firm becomes less sensitive to lead-time variability.

The rest of this article is organized as follows. We discuss the relevant literature in Section 2. Section 3 describes the model. We characterize the optimal procurement policy in Section 4 and discuss managerial implications in Section 5. Extensions to our model are considered in Section 6 and concluding remarks are presented in Section 7. Additional materials (technical lemmas and heuristics) can be found in a supplementary appendix.

[1] This trade-off can also be seen as the interaction between two distinct but related trade-offs: between overage and underage cost in satisfying demand, and between earliness and lateness of order arrivals relative to the selling season. We thank the associate editor for pointing out this alternative interpretation.

2. LITERATURE

This article connects two important streams of literature in supply chain management, namely, the uncertain-supply and dynamic-forecasting streams. We first review the relevant literature on uncertain supply and then turn our attention to the dynamic-forecasting literature. In the end, we briefly discuss the literature on advanced demand information.

Generally speaking, the uncertain-supply literature can be divided into three different but related categories; supply-disruption models, random-yield models, and stochastic lead-time models. The disruption literature typically models a supplier as alternating between up and down phases e.g., [3, 13, 21, 23, 28, 33]. Orders placed when the supplier is up are received on time and in full. No order can be placed when the supplier is down. The random yield literature considers settings in which the quantity received varies in a random fashion from the quantity ordered. We refer the reader to [41] for a review of the random-yield literature. We note that the distinction between disruptions and random yield is less sharp from a manager's perspective, in that stochastically-proportional Bernoulli random-yield models are appropriate models for supply failures e.g., [2, 30, 34–36].

Our article is in the stochastic lead-time category. Stochastic lead times have typically been studied in a recurring-demand setting rather than in a newsvendor setting. For the recurring-demand setting, [4] and [26], for example, investigate the impact of lead time variability in a single-item inventory model. They focus on quantifying the impact of lead time uncertainty on metrics such as the risk of stock out and the optimal cost. A recent article [19] investigates contingency strategies under lead time uncertainty with a constant demand rate. We refer readers to [42] for a comprehensive discussion of stochastic lead times in the context of recurring demand.

A question arises in the newsvendor setting that does not arise in the recurring-demand setting; when should the order be placed? There is an extensive stream of literature in the optimal stopping (time) problem in economics and statistics about sequential decision making and sample testing. We refer the reader to [24] for a comprehensive treatment of the discrete-time Markov decision processes. From the inventory management perspective, this optimal timing decision has been studied in the context of a newsvendor that assembles multiple components into a finished product e.g., [18, 20, 25, 27, 40]. We note that all these articles resort to heuristics and they do not consider demand forecast updating. Our focus on a single-component product allows us to analytically characterize the optimal timing decision and generate important managerial insights.

We now turn our attention to dynamic forecast updating. There is a vast literature on inventory management with forecast updating, but to the best of our knowledge, ours is the first article to consider demand updating in an uncertain-supply setting. The forecast-updating literature can be broadly classified into three related approaches: Markovian, time series, and Bayesian. Our article most closely follows the Markovian forecast revision approach, as developed by [16] and later extended by [17] and [11]. The Markovian forecast revision approach has been commonly adopted in the operations literature e.g., [6, 7, 12, 22, 32].

There are two variants to the Markovian forecast revision approach: additive and multiplicative. The additive model assumes new information is additive and forecast revisions are independent and normally distributed. In contrast, the multiplicative model treats new information as multiplicative, implying that the ratios of successive forecast revisions are independent and lognormally distributed. In this article, we follow the multiplicative forecast updating model. In [17] it was noted that the multiplicative model provided a better fit to their empirical data, which exhibits the property that "the standard deviation of forecast error was roughly proportional to the size of the forecast." (p. 22) This property is consistent with the multiplicative model but not with the additive model.

A number of game-theoretic articles, e.g., [8, 9, 14], investigate forecast updating in a manufacturer-supplier setting. Their forecast-updating models do not fit naturally into the above categorization. In essence, they assume a second ordering opportunity at which time some of the initial forecast uncertainty has been resolved.

Our work is somewhat related to the literature on advanced demand information. One stream of that literature focuses on the coordination of operations and marketing decisions, i.e., the optimal ordering quantity and level of price discount, respectively. Refer [31, 38, 39]. Another stream e.g., [10, 15, 37], focuses on the optimality of base stock or (s,S) policies when customer orders do not require instantaneous fulfilment. Therefore, the firm, effectively, has advanced demand information.

Finally, we note that while there is an extensive literature on stochastic lead times and on demand forecast updating, to the best of our knowledge no existing article links these two literatures. A key contribution of our article is to provide one such link.

3. MODEL

In what follows, we first describe the time line, then the demand and supply models, and conclude with a description of the relevant cost and the problem formulation.

3.1. Time Line

We consider a firm that sells a seasonal product. We use a discrete time, finite horizon model and adopt the convention that 0 is the starting period, the selling season occurs in period T, and that $0 \leq t \leq T$ indicates the current period. Throughout the article we refer to the points of time where decisions are made as decision epochs, and we adopt the convention that a decision epoch corresponds to the beginning of a period. In addition, we refer to the time interval between two decision epochs (i.e., the length of a period) as the decision interval.[2]

The firm has a forecasting process in place: at the start of period 0 it has an initial forecast of demand, and, as time evolves, the firm updates its forecast in each period to incorporate new information. The firm's forecast becomes more accurate as the selling season approaches. By the start of the selling season, however, the firm's forecast may not resolve all demand uncertainty; some residual uncertainty may still exist. The amount of residual uncertainty depends on the forecast performance: a perfect forecasting process will leave no residual uncertainty whereas a less than perfect process will leave some residual uncertainty. In the extreme case of no forecast updating, no demand uncertainty is resolved before the season. The forecast model and the notion of residual uncertainty will be formalized in the next subsection.

The firm may procure in any period $0 \leq t \leq T$, but it is allowed to procure only once in the horizon.[3] This is a reasonable assumption for firms that source from emerging economies, especially for those who do not have a contingent domestic supplier that can provide quick responses. In deciding when to procure, the firm has to balance its supply and demand risk. Procuring early incurs a higher demand risk because the firm's forecast is less accurate and procuring late incurs a higher supply risk because the order may be late due to lead time uncertainty. If the order is late, i.e., arrives after T, similar to [27] the firm incurs a tardiness penalty cost but customer demand is not lost due to tardiness. Note that in Section 6 we explicitly model tardiness related lost sales.

3.2. Demand Updating

In what follows we first describe the forecast updating process, which is an application of the Multiplicative Markovian Forecast Revision (MMFR) approach [16]. We then derive the firm's demand distribution by extending the MMFR model to allow for residual uncertainty. Note that the MMFR model is supported by empirical evidence in actual forecasts (see, for example, p. 94 in [5], p. B-97 to B-102 in [16], and p. 22 in [17]). For ease of understanding, in what

[2] We note that T denotes the number of decision epochs, and therefore, the decision interval will depend on the length of the time horizon. When there is no possibility of ambiguity, we at times also use T to refer to the time horizon.

[3] We do not consider the case where the firm might order after the selling season already starts. This is a mild restriction and it can be easily relaxed at the cost of some additional notation.

follows we briefly state the properties of the discrete time MMFR and then derive the demand forecast distribution.

Let X_t and X_{t+1} denote successive forecasts made at time t and $t+1$ for the demand at time T. Then by MMFR,

PROPERTY 1: The ratios of successive forecasts, $\Delta_t = X_{t+1}/X_t, t = 0, \ldots, T-1$, form a series of independent, lognormally distributed random variables.

We further assume that the Δ_t ratios are identically distributed with parameters (μ, σ). In Section 6 we relax this assumption, but for expositional ease we focus on the independent-ratio model in this article. Note that, the parameters μ and σ are influenced by the decision intervals: if the decision intervals change, then these parameters will also change.

PROPERTY 2: The series of successive forecasts have the Markov property.

In other words, the future forecast evolution's dependence on the past is fully captured by the current realized forecast, x_t. Using Properties 1 and 2 and the fact that the product of lognormally distributed random variables is also lognormally distributed, one can establish that the forecast to be made j periods ahead, i.e., X_{t+j}, satisfies $\ln X_{t+j} \sim N(j\mu + \ln x_t, \sqrt{j}\sigma)$. It follows that if the realized forecast at current time t is x_t, then at time t, the distribution of the forecast made at T, i.e., X_T, satisfies $\ln X_T \sim N((T-t)\mu + \ln x_t, \sqrt{T-t}\sigma)$.

The distribution (at time t) of most interest to us is not that of X_T but rather that of X_D, i.e., the demand. As discussed earlier, we do not assume that demand uncertainty is fully resolved before the start of the selling season, that is, we do not assume that $X_D/X_T = 1$ with probability 1. Instead, we assume that $X_D/X_T = \widehat{\Delta}$, where $\widehat{\Delta}$ reflects residual uncertainty and is lognormally distributed with parameters $\widehat{\mu}$ and $\widehat{\sigma}$. We are now in a position to derive the firm's demand distribution at any time t. At current time t, let $f_t(\cdot|x_t)$ denote the conditional probability density function (PDF) of the demand X_D, given the realized forecast at t is $X_t = x_t$. From the above description, X_D is lognormally distributed and we have

$$f_t(x|x_t) = \frac{1}{\sqrt{2\pi\psi_\sigma(t)}x} \exp\left(\frac{-(\ln x - \psi_\mu(t) - \ln x_t)^2}{2\psi_\sigma(t)}\right),$$

where $\psi_\sigma(t) = (T-t)\sigma^2 + \widehat{\sigma}^2$ and $\psi_\mu(t) = (T-t)\mu + \widehat{\mu}$. Note that $\psi_\sigma(t)$ measures the remaining uncertainty in the demand that is not yet resolved by the forecasting process, and $\psi_\mu(t)$ measures the expected drift in the demand.

Finally, we introduce a useful measure to determine how "good" a forecasting process is. Intuitively, everything else being equal, a good process should have less residual demand uncertainty as compared with a poor process. Now, $\psi_\sigma(t)$ is a measure of the uncertainty remaining at time t. We therefore operationalize the concept of "goodness" of a forecasting process by defining $\lambda = 1 - \psi_\sigma(T)/\psi_\sigma(0)$ as the forecast efficiency. A higher λ reflects less residual uncertainty (relative to the initial uncertainty) and, therefore, the higher the λ, the more efficient is the forecasting process. At $\lambda = 1$, the forecasting process fully resolves demand uncertainty by time T. On the other hand, when $\lambda = 0$, there is no forecast updating, i.e., the residual uncertainty is equal to the initial demand uncertainty. Note that $\lambda = 1$ if and only if $\widehat{\sigma} = 0$. That is, as long as there exists residual uncertainty, the forecast cannot be perfect. Conversely, $\lambda = 0$ if and only if $\sigma = 0$. That is, the firm's belief about demand does not change over time, i.e., there is no forecast updating.

3.3. Supply

Oftentimes, there are two aspects to supply uncertainty. An order might be delayed or not, and if it is delayed, the length of delay may be uncertain. For example, a shipment might be selected for customs inspection and the resulting time spent in inspection might be uncertain. Similarly, a shipment might be delayed due to a port disruption, such as the 2002 US west-coast disruption, and the resulting delay might be uncertain. Let L denote the standard lead time and ω denote a stochastic, non-negative delay. Then, with probability θ, there is a delay and the lead time is $L+\omega$; with probability $1-\theta$, there is no delay and the lead time is simply L. Our model collapses to a constant lead time case when $\theta = 0$ and a pure stochastic lead time case when $\theta = 1$. Hereafter, we refer to θ as the delay probability and ω as the delay.

We assume that both the cumulative distribution function (CDF) and PDF of the delay exist and denote them by $G(\cdot)$ and $g(\cdot)$, respectively. Note that while decisions are made at discrete points in time, the delay is not restricted to decision epochs, and hence, the order can arrive at any time between decision epochs. It is therefore perfectly reasonable to account for the expected delay using the continuous distribution $G(\cdot)$. Let $x^+ = \max\{x, 0\}$ and $x^- = \max\{-x, 0\}$. Define expected earliness and tardiness as

$$A(t) = (1-\theta)(T-L-t)^+ + \theta \int_0^{(T-L-t)^+} ((T-L-t)^+ - \omega) g(\omega) d\omega,$$

$$B(t) = (T-L-t)^- + \theta \int_{(T-L-t)^+}^{\infty} (\omega - (T-L-t)^+) g(\omega) d\omega,$$

where t is the time when an order is placed. Note that $A(t)$ is the expected duration that the order spends in inventory before the season and $B(t)$ is the expected duration by which the order is late.

3.4. Costs and Problem Formulation

Let c, h, s, r, and p represent the unit purchasing cost, holding cost, salvage value, revenue, and tardiness penalty, respectively. We assume the same tardiness penalty structure as in [27], that is, the incurred tardiness cost is linear in the realized delay. The tardiness penalty can be a tangible cost, a proxy for good will or loss demand, or a combination of these factors.

If the firm procures at the beginning of period t, then the expected profit, given the realized forecast of the demand x_t and the procurement quantity y, can be expressed as

$$\tilde{v}(t,x_t) = \max_{y\geq 0}\Big\{ -cy + r\min\{y, \mathsf{E}[X_D|x_t]\} + s\,\mathsf{E}\big[(y-X_D)^+|x_t\big] - hA(t)y - pB(t)\,\mathsf{E}[X_D|x_t]\Big\}. \quad (1)$$

Note that (1) is a classic newsvendor formulation with two additional time-based terms that reflect the holding and tardiness costs[4]: the expected holding cost is $hA(t)y$ and the expected tardiness cost is $pB(t)\,\mathsf{E}[X_D|x_t]$. To follow the cost minimization approach commonly used in classical inventory literature, we rewrite (1) from a cost perspective, i.e.,

$$v(t,x_t) = \min_{y\geq 0}\Big\{cy + r\,\mathsf{E}\big[(X_D-y)^+|x_t\big] - s\,\mathsf{E}\big[(y-X_D)^+|x_t\big] + hA(t)y + (pB(t)-r)\,\mathsf{E}[X_D|x_t]\Big\}. \quad (2)$$

The firm's decision problem can be formulated as a binary discrete-time Markov decision process (MDP).[5] Let $u_t(x_t)$ be the minimum expected cost at the beginning of period t (if the firm has not yet ordered) when the forecast is x_t. Then, the optimality equation is

$$u_T(x_T) = v(T,x_T),$$
$$u_t(x_t) = \min\{v(t,x_t), \mathsf{E}\,u_{t+1}(\Delta_t \cdot x_t)\}, \quad (3)$$

where recall Δ_t is lognormally distributed with parameter (μ,σ). At each decision epoch, the firm must decide whether to place an order (and decides the order quantity) or wait for one more period. Recall that the last decision epoch in (3) is $T-1$, i.e., if the firm has not ordered by time $T-1$ then the firm will order at time T. The firm's objective is to find a policy for placing the order that minimizes the expected cost. Note that we ignore discounting so as to focus on the primary trade off between supply and demand risks.

We conclude this section by introducing the following assumptions. These are made solely for expositional clarity and are assumed to hold throughout the rest of the article (all analysis and results hold without these assumptions, details available upon request).

1. Without loss of generality, we scale the initial demand forecast $x_0 = 1$.
2. We restrict attention to $T > L$, i.e., the starting point is at least a standard lead time before the selling season.[6]

4. OPTIMAL PROCUREMENT POLICY

We first introduce a useful lemma for the optimal procurement quantity for any given procurement time t and realized forecast x_t.

LEMMA 1: For any given procurement time t and realized forecast of demand x_t, the optimal procurement quantity is given by

$$y^*(t,x_t) = F_t^{-1}\left(\frac{r-c-hA(t)}{r-s}|x_t\right), \quad (4)$$

where $F_t^{-1}(\cdot|x_t)$ denotes the conditional inverse distribution function of X_D at time t.

Under the conditions stated in Lemma 1, the optimal procurement quantity is a newsvendor-type expression with an appropriately defined distribution function. Note that $y^*(t,x_t)$ in (4) reduces to the standard newsvendor solution if $t \geq T-L$, i.e., if the firm does not procure until it is within a standard lead time of the selling season.

Given Lemma 1, finding the optimal procurement time t^* becomes the key to characterizing the optimal procurement policy. By (3), it is optimal to place an order (if the firm has not ordered yet) at time t if

$$v(t,x_t) \leq \mathsf{E}\,u_{t+1}(\Delta_t \cdot x_t), \quad (5)$$

and hence the firm's objective is to find $t^* = \inf\{t : v(t,x_t) \leq \mathsf{E}\,u_{t+1}(\Delta_t \cdot x_t)\}$. Note that (5) is defined for $t = 0,1,\ldots,T-1$. In what follows, we simplify (5) and establish an important theorem that considerably simplifies

[4] As in [27], the tardiness penalty is applied to the expected demand. However, all key results hold when the tardiness penalty depends on satisfied demand, i.e., $\mathsf{E}[\min\{X_D,y\}|x_t]$. This is further discussed in Section 6.

[5] We thank an anonymous referee who suggested using a discrete-time approach rather than our original continuous-time approach. This suggestion allows similar results to be established in a more concise and elegant fashion.

[6] The case of $T < L$ is quite straightforward (with the optimal solution being to either order immediately or wait until T.) This assumption allows us to focus on the more interesting case, i.e., $T > L$.

the determination of t^*. First, we establish the following lemma on the firm's minimum expected cost.

LEMMA 2: $v(t, x_t)$ is separable in x_t and t, i.e., there exists a real valued function $\widehat{v}(\cdot)$ such that $v(t, x_t) = x_t \cdot \widehat{v}(t)$.

The previous lemma utilizes the decomposability property of the lognormal distribution, where decomposability refers to the fact that a lognormal random variable can be expressed as the product of n independent random variables.[7] Using Lemma 2, we can establish by induction that the minimum expected cost function $u_t(\cdot)$ (refer (3)) is a homogeneous function of degree 1; that is, $u_t(\alpha x_t) = \alpha u_t(x_t)$. Therefore, we have

$$u_T(x_T) = x_T \widehat{v}(T),$$
$$u_t(x_t) = x_t \min\{\widehat{v}(t), \mathsf{E}\, u_{t+1}(\Delta_t)\}, \tag{6}$$

where $\widehat{v}(t)$ is independent of x_t. Hence (5) can be reduced to

$$\widehat{v}(t) \leq \mathsf{E}\, u_{t+1}(\Delta_t). \tag{7}$$

As a result, whether an order should be placed at time t depends on the sign of $\widehat{v}(t) - \mathsf{E}\, u_{t+1}(\Delta_t)$, which is independent of x_t. The following theorem ensues.

THEOREM 1: The optimal procurement time t^* is independent of the forecast updating process X_t.

Theorem 1 and Lemma 1 together characterize the structure of the optimal procurement policy: (i) the optimal time is independent of the forecast evolution X_t (Theorem 1), and (ii) the optimal quantity depends on the forecast evolution and is given by (4). Care should be taken when interpreting Theorem 1. Although it establishes that the current forecast level X_t does not influence the optimal procurement time, it does not state that the forecasting process has no impact. In fact, as will be seen later (as a consequence of Theorem 2), the forecasting process does influence the optimal procurement time through $\psi_\sigma(t)$, or in other words, through the forecast parameter σ. Although the current forecast level X_t does influence the magnitude of the expected total cost, it does not influence the timing-induced tradeoff amongst the expected procurement, overage, underage, earliness, and lateness costs.

While this timing independence result can be extended to allow for correlated forecast revisions and time-dependent parameters (refer Section 6), the result does depend on the decomposability attribute of the lognormal distribution. However, as discussed earlier, the lognormal assumption for forecast revisions is a reasonable one. This timing-independence result has an important managerial implication: the firm can determine its optimal procurement time in advance but it must wait until that time to determine its procurement quantity. The fact that the firm can determine its procurement time in advance is very beneficial from a planning perspective. For example, the firm may notify its supplier of the upcoming procurement plan (time) such that the supplier can better prepare for the future production request in advance.

[7] See [29] (p. 508) for a rigorous discussion of the decomposability of the lognormal distribution. We thank the AE for pointing out this decomposability approach that admits a more elegant proof for Lemma 2.

4.1. Optimal Procurement Time

Having characterized the structure of the optimal procurement policy, we now proceed to solve the firm's optimal procurement time t^*. By Theorem 1, when deciding whether to place an order now or wait for one more period, it is sufficient to study the behavior of $\widehat{v}(t) - \mathsf{E}\, u_{t+1}(\Delta_t)$ as a function of t. To simplify our analysis, we rewrite the optimality Eq. (6) in terms of $w(t)$, using the relationship $w(t) = u_t(x)/x$:

$$w(T) = \widehat{v}(T),$$
$$w(t) = \min\{\widehat{v}(t), \mathsf{E}[\Delta_t \cdot w(t+1)]\}. \tag{8}$$

We analyze (8) by backward induction in time t. First consider the last decision epoch $T-1$. We have

$$\begin{aligned} w(T-1) &= \min\left\{\widehat{v}(T-1), \mathsf{E}\left[\Delta_{T-1} \cdot w(T)\right]\right\} \\ &= \min\left\{\widehat{v}(T-1), \mathsf{E}\left[\Delta_{T-1}\right] \cdot \widehat{v}(T)\right\} \\ &= \min\left\{\widehat{v}(T-1), e^{\mu+\sigma^2/2} \cdot \widehat{v}(T)\right\}, \end{aligned} \tag{9}$$

where the last step follows from the fact that the forecast revision Δ_t's are lognormally distributed. Using Lemma 1 and 2, one can rewrite $\widehat{v}(t)$ as (refer technical Lemma A2)

$$\widehat{v}(t) = e^{(\psi_\mu(t)+\psi_\sigma(t)/2)} \cdot M(t), \tag{10}$$

where $M(t) = pB(t) - (r-s)(1+\mathrm{erf}(k(t) - \sqrt{\psi_\sigma(t)/2}))/2$, and $k(t) = \mathrm{erf}^{-1}(2(r-c-hA(t))/(r-s)-1)$. Note that $\mathrm{erf}(y) = \frac{2}{\sqrt{\pi}}\int_0^y e^{-x^2}dx$ is the "error" function associated with the normal distribution, and it has a nice property of being monotonic in its argument. Refer p. 297 in [1] for more details about the error function and the closely related complementary error function $\mathrm{erfc}(\cdot)$. Substituting (10) into (9), we obtain

$$\begin{aligned} w(T-1) &= \min\left\{e^{(\psi_\mu(T-1)+\psi_\sigma(T-1)/2)} \cdot M(T-1), e^{\mu+\sigma^2/2} \right. \\ &\qquad \left. \cdot\, e^{(\psi_\mu(T)+\psi_\sigma(T)/2)} \cdot M(T)\right\} \\ &= e^{(\psi_\mu(T-1)+\psi_\sigma(T-1)/2)} \cdot \min\{M(T-1), M(T)\}, \end{aligned} \tag{11}$$

where the second step follows from the definition of $\psi_\mu(t)$ and $\psi_\sigma(t)$. Therefore, if the firm has not yet placed an order, it can decide whether or not to place an order at time $T-1$ by evaluating the sign of $M(T-1) - M(T)$. In particular, the firm should place an order immediately if $M(T-1) \le M(T)$ and should wait until next period otherwise. More generally, one can establish by induction that

$$w(t) = e^{(\psi_\mu(t)+\psi_\sigma(t)/2)} \cdot \min\{M(t), \min\{M(t+1), M(t+2), \ldots, M(T)\}\}. \quad (12)$$

The following theorem ensues.

THEOREM 2: The optimal procurement time $t^* = \arg\min_{t\in\{0,1,\ldots,T\}} M(t)$.

The above theorem establishes that the optimal procurement time can be found simply by searching (restricting to decision epochs) over the $M(t)$ function for the minimum point.[8] In what follows, we consider the special case when there is no forecast updating (forecast efficiency $\lambda = 0$), i.e., the firm does not resolve any demand uncertainty before the start of the selling season T.

4.1.1. *The Case of No Forecast Updating ($\lambda = 0$)*

Note that while the $\lambda = 0$ case can be viewed a special case of our general model, the analysis involved differs significantly from the $\lambda > 0$ case. This is because there is no forecast updating in the $\lambda = 0$ case. Hereafter, we refer to the $\lambda = 0$ case as the no forecast updating case. The no forecast updating case itself is of interest because, to the best of our knowledge, the joint quantity and timing problem under lead time uncertainty has not been studied in the newsvendor model before.

We first note that, with no forecast updating, the firm's demand distribution is independent of time t. Because of this independence, we can simply write $v(t, x_t)$ as $v(t)$. Using (2) and recalling that we scale $x_0 = 1$, the minimum expected cost if the firm procures at time t is

$$v(t) = (hA(t) - (r-c))y^*(t) + (r-s)\,\mathsf{E}[(y^*(t) - X_D)^+] + pB(t)\,\mathsf{E}[X_D], \quad (13)$$

where, by Lemma 1, $y^*(t) = F_0^{-1}((r-c-hA(t))/(r-s))$. Hence, the optimal procurement time is given by $t^* = \arg\min_{t\in\{0,1,\ldots,T\}} v(t)$. Define $\underline{t} = A^{-1}((r-c)/h)$. It is straightforward to show that the optimal procurement time satisfies $\lfloor \underline{t} \rfloor \le t^* \le \lceil T-L \rceil$, where $\lfloor\cdot\rfloor$ and $\lceil\cdot\rceil$ represents floor and ceiling functions, respectively. Note that, for any $t < \underline{t}$, the implied procurement quantity is zero, i.e., the firm does not participate (in what follows we consider the more interesting case where the firm does participate). While we can restrict the range of the t^*, characterizing the optimal procurement time t^* is not straightforward because (13) is in general not unimodal in time t. For certain classes of the delay distribution function $G(\cdot)$, however, we can exploit structural properties of Eq. (13) and characterize the optimal procurement time.

For the rest of this subsection, we assume that the delay distribution satisfies $G''(\cdot) \le 0$. This is a reasonably mild assumption and is satisfied by a number of distributions, including the uniform, exponential, and certain classes of Weibull and Gamma distributions. The following theorem characterizes potential candidates for the optimal procurement time t^*. Define $\mathcal{I}(x) = \{\lfloor x \rfloor, \lceil x \rceil\}$.

THEOREM 3: (i) If $\theta = 0$, then $t^* \in \mathcal{I}(T-L)$. (ii) If $\theta > 0$, define

$$\widehat{H}(t) = y^*(t)\frac{(1-\theta) + \theta G(T-L-t)}{\theta(1-G(T-L-t))}.$$

(a) If $\sup\{\widehat{H}(t)\} < (p/h)\,\mathsf{E}[X_D]$, then $t^* \in \mathcal{I}(\underline{t})$.
(b) If $\widehat{H}(T-L) > (p/h)\,\mathsf{E}[X_D]$, then $t^* \in \mathcal{I}(T-L)$.
(c) Otherwise, let $\operatorname{sgn}(x) \doteq x/|x|$, define $\mathcal{T} = \{t : \operatorname{sgn}(v(t-1) - v(t)) \ne \operatorname{sgn}(v(t) - v(t+1))\}$, then $|\mathcal{T}| \le 2$. Let τ be the larger element in $\mathcal{T}$ (if $|\mathcal{T}| = 2$), then $t^* \in \{\tau\} \cup \mathcal{I}(T-L)$.

Part (i) of Theorem 3 states that, if the supply is perfectly reliable, then the optimal procurement time is around one standard lead time before the selling season. Part (ii) establishes the importance of the ratio of the tardiness penalty cost to the holding cost in the optimal procurement time. In particular, (ii)(a) says that if this ratio is very high, then the firm may not participate, i.e., $t^* \in \mathcal{I}(\underline{t})$; (ii)(b) says that if this ratio is very low, then the firm orders around a standard lead time L in advance of the selling season; for intermediate ratios (ii)(c), the firm uses a procurement time between these two extremes. Note that if the decision interval is infinitesimal, then the characterization of t^* is simpler: $\mathcal{T}$ has exactly two elements and $t^* = \tau$.

Given the delay distribution is concave, Theorem 3, combined with Lemma 1, completely characterizes the firm's optimal procurement policy. In closing, we note that we also investigate a number of heuristics in Appendix A5.

[8] The $M(t)$ function is in general not well behaved and can assume very complex shapes. If we do not restrict t to decision epochs, then (with some additional technical assumptions) we are able to characterize the procurement time that minimizes $M(t)$. This unrestricted procurement time provides a reasonable approximation to t^* when the decision intervals become very short. We refer the reader to Theorem A1 and Lemmas A7 and A8 in the Appendix for a complete characterization of this unrestricted procurement time.

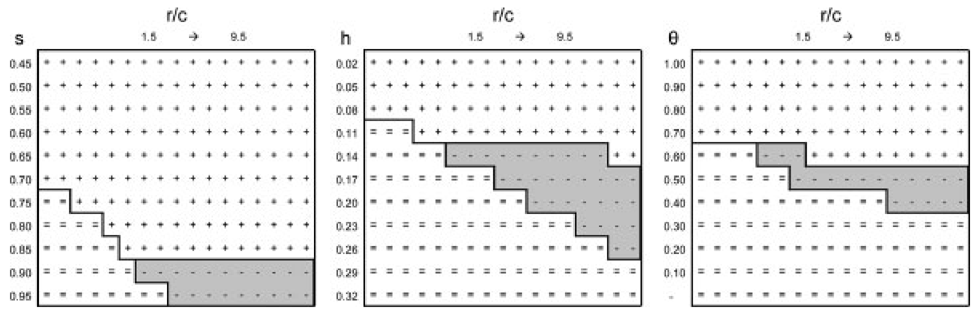

Figure 1. Optimal ordering time with and without forecast updating. "+" denotes a forecast updating firm orders closer to the selling season, "−" denotes a no updating firm orders closer to the selling season, and "=" denotes both ordering at exactly one standard lead time L before the selling season.

5. IMPLICATIONS OF OPTIMAL POLICY

In this section, we investigate how three important aspects of the system, i.e., the demand forecast, supply, and cost characteristics, influence the firm's optimal procurement policy and the resulting expected cost. At times we will refer to numeric studies we used to complement the analytical results developed in this section. The detailed description of the studies can be found in Appendix A2. In these numeric studies, all technical assumptions (e.g., $G''(\cdot) \leq 0$ assumed in Section 4.1.1) were relaxed. In what follows we refer to the expectation, variance, and coefficient of variation (CV) of the delay distribution $G(\cdot)$ as the mean delay, the delay variance, and the delay CV, respectively.

5.1. Demand Forecast Characteristics

We focus on an important characteristics of the forecast process: the forecast efficiency λ that measures the fraction of demand uncertainty resolved by the forecasting process. Intuitively, one might expect that given the same starting demand estimation, the optimal procurement time with forecast updating ($\lambda > 0$) would be later than that without forecast updating ($\lambda = 0$), because forecast updating enables the firm to reduce its demand risk by delaying its procurement. This intuition, however, is not true in general.

THEOREM 4: For identical initial forecasts of demand, let t_F^* and t_N^* denote the optimal order time with and without forecast updating, respectively.[9] (a) $\theta = 0 \Rightarrow t_F^* \geq t_N^*$. (b) $h = 0 \Rightarrow t_F^* \geq t_N^*$. (c) Otherwise, t_F^* can be less than t_N^*.

In the case of a deterministic lead time or no holding cost, i.e., (a) and (b) above, the intuition is correct: a forecast-updating firm will procure later than a firm that does not update its forecast. However, when neither of the above two conditions hold, i.e., case (c) above, the forecast-updating firm may in fact procure earlier than the "no-updating" firm. The reason lies in the interplay of the demand and supply risk. Starting with the same demand forecast, the forecast-updating firm's demand risk, i.e., demand uncertainty, reduces over time. For certain cost parameters, a lower demand variance results in a lower procurement quantity for a given procurement time. A lower quantity reduces the earliness component of the supply risk, i.e., the inventory-related cost of an early arrival is lower for a smaller order quantity. Because of this, it can be optimal under certain circumstances for the forecast-updating firm to procure earlier than the no updating firm.

We use a comprehensive numeric study (See Section A2.2 for details) to further investigate conditions under which case (c) above is more likely to occur. We observed that case (c) is more likely to occur when (1) the revenue to cost ratio (r/c) is higher, (2) the salvage value to cost ratio (s/c) is higher, (3) the unit holding cost h is higher, and (4) the delay probability θ is moderate. Conditions 1 to 3 push the no updating firm to order closer to the selling season because they all contribute to higher inventory holding cost: (1) prompts the firm to order more (and hence higher holding cost) because the product is more profitable; (2) also prompts the firm to order more because the cost of over-stocking is low; (3) directly contributes to higher holding cost. The following figure illustrates the region (shaded area) where a forecast updating firm (with $\lambda = 0.4$) orders earlier than a no updating firm ($\lambda = 0$). Note that Fig. 1 was obtained with delay probability $\theta = 0.5$, unit holding cost $h = 15\%c$, and unit salvage value $s = 90\%c$, except when they are varied in each subfigure. These parameter values are chosen to illustrate the likely region that case (c) above can occur, and therefore, do not represent typical scenarios.

In our numeric study, we observed instances, especially when the r/c and s/c ratios were high, for which the optimal procurement time decreased in the forecast efficiency λ. However, it was more typical for the optimal procurement time to increase in the forecast efficiency, that is, the more

[9] Note that F denotes "Forecast updating" and N denotes "No forecast updating."

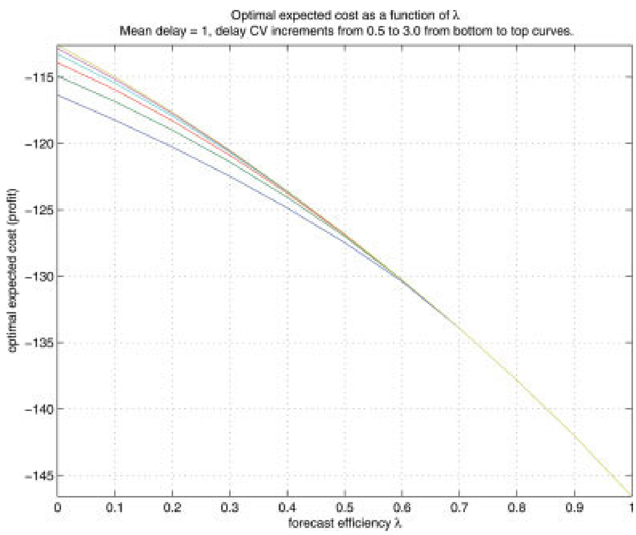

Figure 2. Optimal expected cost (profit) as forecast efficiency λ increases. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

the forecast process resolved uncertainty, the longer the firm waited to place its order.

We observed that the optimal expected procurement quantity typically increases in forecast efficiency. The intuition is that the optimal procurement quantity decreases in demand volatility when the critical newsvendor fractile is less than 0.5, which is the case for our base case settings (refer Section A2.1). As discussed earlier, however, when the revenue (salvage value) to cost ratio is relatively high, the firm's expected procurement quantity can be strictly decreasing in forecast efficiency. This observation is consistent with the fact that the firm orders earlier under these conditions.

In contrast to the procurement time (quantity), a more efficient forecast always leads to a lower optimal expected cost. In other words, forecast updating is valuable to the firm. To further understand the value of forecast efficiency, we conducted two numeric studies (refer Section A2.2 and Section A2.3 for details) to investigate the difference in the optimal expected cost as the forecast efficiency parameter λ changes. Refer Fig. 2.

We discuss two noteworthy findings. (1) The optimal expected cost decreases more rapidly as λ approaches to 1. Thus, a marginal increase in the relative strength of the forecast efficiency has a larger benefit if the forecast is already efficient. (2) The optimal expected cost becomes less sensitive to the lead time CV as λ increases. Therefore, the firm with more efficient forecast is more robust to variability in the lead-time delay.

On average (refer Section A2.3 for details), the cost with forecast updating at $\lambda = 1$ and $\lambda = 0.5$ was 29.7% and 8.5% lower than the cost with no forecast updating ($\lambda = 0$), respectively. We note that, in practice, the firm usually incur a positive forecasting cost, and therefore, the forecasting benefit discussed earlier would be need to be compared with the cost of forecasting to evaluate the net benefit.

5.2. Supply Characteristics

In this section, we study how lead time reliability and the standard lead time affect a firm's procurement policy and its expected cost.

5.2.1. Delay Probability

We observed that the optimal procurement time typically decreases in the delay probability θ. In other words, the firm places order earlier when the delay probability increases. The optimal expected procurement quantity follows a similar pattern as being typically decreasing in the delay probability. This is primarily driven by the fact that the remaining demand uncertainty is higher when the firm places order earlier.

Although one might expect that the optimal procurement time always decreases in the delay probability θ, this intuition is correct only when there is no forecast updating ($\lambda = 0$). Our numeric study indicates that the optimal procurement time can be increasing or decreasing in the delay probability when $\lambda > 0$. This effect is directly linked to the relationship between the delay probability and the delay CV, which is further discussed in Section 5.2.2.

The effect of the delay probability on the optimal expected cost is also nuanced. With deterministic demand, one can prove that a firm always (at least weakly) prefers a perfectly reliable lead time ($\theta = 0$) to an unreliable lead time ($\theta > 0$), but a marginal increase in the delay probability does not always hurt the firm: the firm's optimal expected cost is not necessarily increasing in the delay probability θ. Our numeric investigation revealed that even in the case of stochastic demand with or without forecast updating, the firm's optimal expected cost can decrease in the lead time delay probability. We note, however, the firm's optimal expected cost at $\theta > 0$ never falls below that at $\theta = 0$.

While surprising, the above observation can be explained as follows. An increase in the delay probability increases the mean lead time and this hurts the firm. However, an increase in the delay probability can decrease the lead time variance and this benefits the firm. In fact, one can show that the lead time variance increases in the delay probability if and only if the delay of CV is greater than or equal to $\sqrt{2\theta - 1}$. Therefore, an increase in the delay probability hurts the firm if the delay CV is high but may benefit the firm if the delay CV is low.

5.2.2. Delay CV

The effect of the delay CV is not straightforward. We observed the optimal procurement time to be increasing or decreasing in the delay CV, with it typically being decreasing for a low CV and increasing for a high CV. At a high delay CV, a marginal increase in the delay CV serves to increase the already-high probability of the delay being small, and this allows the firm to further postpone procurement. At a low delay CV, the probability of a small delay is low and the probability of a moderate or large delay is high. A marginal increase in the CV, while increasing the probability of a small delay, also increases the probability of a large delay, and this large-delay effect appears to dominate, with the result that the firm orders earlier.

The effect of the delay CV on the optimal expected procurement quantity follows a similar pattern, i.e., typically decreasing for a low CV and increasing for a high CV. The intuition is that the optimal procurement quantity is closely linked with the remaining demand uncertainty when the firm places the order. For typical parameter values, when the firm orders earlier the higher demand uncertainty leads to a lower procurement quantity. The case of the firm ordering later can be similarly reasoned.

In contrast to the effect of the delay probability, the optimal expected cost was always increasing in the delay CV. We note, however, that as the forecast becomes more efficient, the optimal expected cost becomes less sensitive to the delay CV.

5.2.3. Standard Lead Time

We first note that, as expected, the firm procures earlier as the standard lead time L increases. The firm's optimal procurement quantity and expected cost, however, is independent of the standard lead time L when there is no forecast updating ($\lambda = 0$). The intuition is that the standard lead time does not affect the demand uncertainty – the firm merely shifts its optimal procurement time. As there is no monetary discounting in our model, the shift in the optimal procurement time does not affect the expected procurement quantity and cost. This is not true in forecast updating case. With forecast updating ($\lambda > 0$), a reduction in the standard lead time allows the firm to order later, i.e., when its demand forecast becomes more accurate. Because of this, the firm's optimal expected procurement quantity (typically) decreases as the standard lead time decreases. In addition, we observed the optimal expected cost to be concave increasing in the standard lead time. Therefore, with forecast updating, lead time reduction efforts becomes increasingly valuable as the standard lead time decreases.

5.3. Cost Characteristics

In this section, we investigate the sensitivity of the firm's optimal procurement policy and the expected cost with respect to system cost parameters. Note that, for the $\lambda = 0$ case, most of the sensitive analysis can be analytically characterized. For the $\lambda > 0$ case, however, it is difficult to determine the sensitivity direction analytically and we therefore primarily present numeric observations. The following theorem presents the sensitivity analysis for the $\lambda = 0$ case, i.e., when there is no forecast updating.

THEOREM 5: (i) The optimal procurement time t^* is non-increasing in the unit tardiness cost p. (ii) The minimum expected procurement cost is non-decreasing in the unit procurement cost c, the unit holding cost h, and the unit tardiness cost p, and is non-increasing in the unit revenue r and unit salvage value s.

Part (i) of the above theorem tells us that the firm procures earlier as the tardiness penalty cost p increases. Note that the directional change in t^* as a function of the unit salvage value s or unit revenue r depends on the decision interval. Intuitively, one might expect that as the holding cost h increases, the firm would order closer to the season so as to reduce the expected duration over which inventory is held. Although we are unable to establish this analytically, our numeric studies indicate that this intuition is indeed correct. For typical parameter values, the above sensitivity analysis holds for the forecast updating case as well.

The optimal expected procurement quantity is closely related to the remaining demand uncertainty at the order time and it therefore follows a similar pattern as the optimal procurement time. In other words, the optimal expected procurement quantity decreases in the unit tardiness cost p. The effect of h is more nuanced: we observed the optimal procurement quantity to be initially decreasing in h and then increasing in h. The intuition lies in the fact that h has two opposing effects on the optimal procurement quantity. On one hand, a higher h results in the firm ordering later, which increases the optimal procurement quantity. On the other hand, an increase in h reduces the critical fractile, which decreases the optimal procurement quantity. At a smaller h, the latter effect seems to dominate but as h continues to increase, the former effect dominates with a result that the optimal procurement quantity being increase in h.

Part (ii) of the above theorem proves that the optimal expected cost is increasing in the unit cost c, in the holding cost h, and in the tardiness penalty cost p; is decreasing in the unit revenue r and in the salvage value s. These directional results are quite intuitive.

In summary, the firm's optimal procurement policy is influenced by a number of important factors, notably the forecast efficiency λ, the delay probability θ, the delay CV, and the system cost parameters, especially r/c ratio, s/c ratio, and the unit holding cost h. For typical parameter values, the sensitivity effect is as expected, e.g., the optimal procurement time is closer to the selling season when the forecast efficiency increases, when the delay probability decreases, or when the unit holding cost increases. Some sensitivity effects are less intuitive or parameter-value dependent. For example, the optimal procurement time may increase or decrease in the delay CV, depending on whether the current delay CV is low or high. Therefore, a firm should carefully analyze its system to determine its optimal procurement time. For the optimal procurement quantity, we note that for typical parameter values it follows a similar pattern to the optimal procurement time, with a notable exception of the unit holding cost, where the optimal procurement quantity can be increasing or decreasing in h.

6. ROBUSTNESS AND SOME GENERALIZATIONS

In this section, we discuss the robustness of Theorem 1 and, in doing so, illustrate how the order time/forecasting independence result developed in Section 4 can be generalized to alternative model formulations. In what follows we explore demand forecasting, the tardiness penalty cost, customer balking behavior, and time dependent system parameters. We relegate the detailed formulation and corresponding proofs to Section A3 in the Appendix.

- Tardiness Penalty Cost: The order time/forecasting independence result holds when the tardiness penalty cost is assessed not on the expected demand as in (2), but on the satisfied demand, i.e., $\mathsf{E}[\min\{X_D, y\}|x_t]$. It is straightforward to prove (refer Section A3 for details) that Lemma 2 holds under this alternative formulation and so does Theorem 1. Furthermore, one can establish that the independence result holds when the tardiness penalty p is time dependent, hence allowing the order time/forecasting independence result to be extended to models with a general nonlinear (in time) tardiness penalty cost.
- Customer Balking Behavior: We have used a tardiness penalty cost in (2) as a proxy for lost goodwill or demand due to tardiness. We now consider a model in which customer balking (tardiness induced lost sales) is directly modeled. Let $0 \leq \alpha(\tau) \leq 1$ denote the fraction of demand that remains after a delay of length $\tau \geq 0$. Then, the demand at time $T + \tau$ is $\alpha(\tau)X_D$. This tardiness induced lost sales replaces the tardiness penalty in our base model. Lemma 2 still holds for this model,[10] and therefore, so does Theorem 1. Refer Section A3 for formulation and further details.
- Demand Forecasting: From the forecasting perspective, Theorem 1 can be generalized to a forecasting process with autocorrelated successive revisions or non-stationary revisions where the Δ_t's are non-identical. In these cases, the expressions for $\psi_\mu(t)$ and $\psi_\sigma(t)$ are more complex, but it is evident from (4) that the homogeneity property of the optimality equation still holds, and hence, the Theorem 1. In fact, we have developed and fully characterized a model in which the Δ_t's are autocorrelated and the analytical results and insights obtained from this more complex model are similar (details available upon request). Note that we also describe and analyze an alternative demand forecasting model where the effect of successive forecast revisions are additive (as opposed to multiplicative). We prove that under this alternative model Theorem 1 does not hold. Refer Section A4 for details.
- Time-Dependent System Parameters. In our base model, the system parameters were time invariant. Our model can be extended to allow for a time-dependent unit procurement cost c, unit salvage value s, and unit penalty cost p. The analysis of the optimal procurement time t^* will be much more complex but the independence result still holds. Refer Section A3 for details.

We note that there are circumstances under which Lemma 2 does not hold, and hence, the independence result breaks down. One example is when a parameter, such as the unit procurement cost c, depends on the forecast realizations of x_t. This unit cost dependence might happen, for example, if the system depends on certain general economic environment states and a higher realization of x_t is a reflection of higher demand in the general economy. A second example when the independence result does not hold is when the procurement cost is nonlinear in the procurement quantity.

7. CONCLUSIONS

In this article, we investigate a key trade-off that will be faced by an increasing number of firms as more goods are sourced from distant suppliers. In particular, we study the optimal timing-and-quantity problem for a newsvendor-type firm facing supply and demand risk. The supply risk, i.e., lead time uncertainty, motivates the supplier to order earlier. In contrast, the opportunity to reduce its demand risk through forecast updating motivates the firm to order later. We study this timing-and-quantity problem in a quite general setting and establish a number of interesting technical and managerial results.

We prove an important timing-and-level separation result: under multiplicative forecast revisions the optimal procurement time is independent of the realization of forecast evolutions but the optimal quantity is not. We characterize the optimal procurement time and quantity, and analytically

[10] While it's reasonable to assume that $\alpha(0) = 1$ and that $\alpha(\tau)$ is convex decreasing in τ, such assumptions are not necessary for the separability result of Lemma 2 to hold.

establish the directional effect of many important supply and market attributes. Although, one might expect that, all else being equal, a firm with effective forecast would procure closer to the selling season (to take advantage of increasing demand-forecast accuracy) than would a firm with no forecast updating, this is not necessarily true. We prove that this intuition is correct in the case of a deterministic lead time. However, in a stochastic-lead time setting, a forecast-updating firm may procure earlier than the firm that does not. The reason lies in the interplay of the demand and supply risk.

We see future research opportunities along a number of dimensions. One could investigate the multi-component version of this problem or investigate a single-component problem with multiple ordering opportunities. In addition, it would be of interest to consider a risk-averse firm. We hope that future research, by ourselves and others, will address these dimensions.

8. PROOFS

PROOF OF LEMMA 1: Let $\mathrm{TC}(y) = (hA(t) - (r - c))y + (r-s)\,\mathsf{E}[(y-X_D)^+|x_t] + pB(t)\,\mathsf{E}[X_D|x_t]$. One can show that $\mathrm{TC}(y)$ is convex in y, and $\mathrm{TC}'(y) = hA(t) - (r - c) + (r-s)F_t(y|x_t)$. Setting $TC'(y) = 0$ obtains the desired result. □

PROOF OF LEMMA 2: By Property 1 of the MMFR process, the demand X_D follows a lognormal distribution, which exhibits the decomposability property. Specifically, define $\widetilde{\Delta}_t = \widehat{\Delta}\prod_{l=t}^{T-1}\Delta_l$, we have

$$X_D = x_t\widehat{\Delta}\prod_{l=t}^{T-1}\Delta_l = x_t\widetilde{\Delta}_t. \tag{14}$$

Substituting (14) into (2), we have

$$v(t,x_t) = x_t\min_{\widetilde{y}\geq 0}\left\{c\widetilde{y} + r\,\mathsf{E}\left[(\widetilde{\Delta}_t - \widetilde{y})^+\right]\right.$$
$$\left. -s\,\mathsf{E}\left[(\widetilde{y} - \widetilde{\Delta}_t)^+\right] + hA(t)\widetilde{y} + (pB(t) - r)\,\mathsf{E}[\widetilde{\Delta}_t]\right\},$$

where $\widetilde{y} = y/x_t$. The lemma statement then follows directly. □

PROOF OF THEOREM 1: Follow directly from Lemma 2 and (6). □

PROOF OF THEOREM 2: We prove (12) by induction on t. Analogous to (9) to (11), we have $w(T-2) = e^{(\psi_\mu(T-2)+\psi_\sigma(T-2)/2)}\cdot\min\{M(T-2),\min\{M(T-1),M(T)\}\}$. Suppose $w(t) = e^{(\psi_\mu(t)+\psi_\sigma(t)/2)}\cdot\min\{M(t),\min\{M(t+1),M(t+2),\ldots,M(T)\}\}$, we need to show that $w(t-1) = e^{(\psi_\mu(t-1)+\psi_\sigma(t-1)/2)}\cdot\min\{M(t-1),\min\{M(t),M(t+1),M(t+2),\ldots,M(T)\}\}$. Analogous to (9), we have

$$\begin{aligned}
w(t-1) &= \min\left\{\widehat{v}(t-1),\mathsf{E}\left[\Delta_{t-1}\cdot w(t)\right]\right\}\\
&= \min\Big\{\widehat{v}(t-1),\mathsf{E}\left[\Delta_{t-1}\right]\cdot e^{(\psi_\mu(t)+\psi_\sigma(t)/2)}\\
&\quad\cdot\min\{M(t),\min\{M(t+1),M(t+2),\ldots,M(T)\}\}\Big\}\\
&= \min\Big\{e^{(\psi_\mu(t-1)+\psi_\sigma(t-1)/2)}\cdot M(t-1), e^{\mu+\sigma^2/2}\\
&\quad\cdot e^{(\psi_\mu(t)+\psi_\sigma(t)/2)}\cdot\min\{M(t),M(t+1),\ldots,M(T)\}\Big\}\\
&= \min\Big\{e^{(\psi_\mu(t-1)+\psi_\sigma(t-1)/2)}\cdot M(t-1), e^{(\psi_\mu(t-1)+\psi_\sigma(t-1)/2)}\\
&\quad\cdot\min\{M(t),M(t+1),\ldots,M(T)\}\Big\}\\
&= e^{(\psi_\mu(t-1)+\psi_\sigma(t-1)/2)}\min\{M(t-1),\\
&\quad\min\{M(t),M(t+1),\ldots,M(T)\}\},
\end{aligned}\tag{15}$$

where the second step follows from the induction assumption. This completes the proof. □

PROOF OF THEOREM 3: (i) For $\theta = 0$, $A(t) = (T-L-t)^+$ and $B(t) = (T-L-t)^-$. It is straightforward to show that $t^* \leq \lceil T-L\rceil$. By (13), $v(t) - v(t+1) > 0 \Rightarrow t^* \in \{\lfloor T-L\rfloor, \lceil T-L\rceil\}$. (ii)(a) Follows from the fact that $\sup\widehat{H}(t) < (p/h)\,\mathsf{E}[X_D] \Rightarrow v(t+1) - v(t) \geq 0$ for any t (see Lemma A4). Note that the $\widehat{H}(\cdot)$ function defined in Theorem 3 is similar to the $H(\cdot)$ function defined in Lemma A3 and A4, with a change of variable $z = T-L-t$. (b) By Lemma A4, $v(\cdot)$ is minimized at corner solutions when $\widehat{H}(T-L) > (p/h)\,\mathsf{E}[X_D]$ (or equivalently $H(0) > (p/h)\,\mathsf{E}[X_D]$). The case of $v(\cdot)$ being minimized at $\underline{t}$ is trivial (because the firm does not participate). Hence, the optimal $t^* \in \{\lfloor T-L\rfloor, \lceil T-L\rceil\}$. (c) By Lemma A4, if t is not restricted to decision epochs, then $v(\cdot)$ has exactly one maximum and one minimum point. If $|\mathcal{T}| = 2$, then $v(\tau)$ must be the closest point to the minimum. Given the restriction of decision epochs, t^* must either equal to τ or equal to the corner solution $\mathcal{I}(T-L)$. □

PROOF OF THEOREM 4: (a) When $\theta = 0$, the supply lead time is constant and equals to L. Consequently, we have $t^*_{\mathrm{N}} \geq \lfloor T-L\rfloor$ and $t^*_{\mathrm{F}} \geq \lfloor T-L\rfloor$. By Lemma A1 (refer appendix), however, $t^*_{\mathrm{N}} \leq \lceil T-L\rceil$. Therefore, we have $t^*_{\mathrm{N}} \in \mathcal{I}(T-L)$. With forecast updating, we must have $t^*_{\mathrm{F}} \geq t^*_{\mathrm{N}}$ because otherwise the firm can do strictly better by adopting t^*_{N} in the forecast updating case. (b) When $h = 0$, we have $t^*_{\mathrm{N}} \leq t^*_{\mathrm{F}}$ because, with no forecast updating, *any* procurement time before $T-L$ is optimal. (c) We prove this by an example. Assume the delay ω follows a Weibull distribution with the shape parameter $\alpha = 0.85$ and expectation $\mathsf{E}[\omega] = 2$. Set unit cost $c = 2.1$. In addition, we set $r = 7$, $p = 0.7$, $s = 2$, $h = 0.14$, $L = 2$, $T = 6$, and $\theta = 0.5$. We set expected demand equal to 100 and the demand CV equal to 0.8. Note that $\alpha = 0.85 \Rightarrow G''(\cdot) \leq 0$. We assume the decision interval is 0.01. For the no forecast updating case, by definition in Lemma A3, $H(0) = y^*(0)\frac{1-\theta}{\theta} = 331.07$. Since $\frac{p}{h}\,\mathsf{E}[X_D] = 500 > H(0)$ and $\sup H(z) = 2{,}387.56 > \frac{p}{h}\,\mathsf{E}[X_D]$, by part (c) of Lemma A4, $z^* = z_1$. By evaluating equation (A-24), we obtain $z_1 = 0.41$, and hence, the optimal order time $t^*_{\mathrm{N}} = T - L - z_1 = 3.59$. For the forecast updating case, consider $M(t)$. For $0 \leq t < L$, $k(t) - \sqrt{\frac{1}{2}\psi_\sigma(t)}$ is strictly decreasing in t. Combining with the fact that $B(t)$ is linear in t, one can verify that $B(T) = 3$, $B(T - L) = 1$, so that $t > T - L$ cannot be optimal. For $t \leq T - L$, note $B(\cdot)$ is convex increasing in t and $\lim_{t\to-\infty} B(\cdot) = 0$, and $\mathrm{erf}(k(t) - \sqrt{\frac{1}{2}\psi_\sigma(t)})$ is monotonically increasing in t and $\lim_{t\to-\infty}\mathrm{erf}(\cdot) = -1$. As both functions are bounded, by an exhaustive numeric search we obtain $t^*_{\mathrm{F}} = 3.55 < t^*_{\mathrm{N}}$. □

PROOF OF THEOREM 5: Part (i). By the definition of t^*, we have $v(t^*) \leq \min\{v(t^*+1), v(t^*+2), \ldots, v(T)\}$. By (13), $v(t)-v(t+n)$ is decreasing in p for any $n \in \{1, 2, \ldots, T-t\}$. It follows that t^* is non-increasing in p. Part (ii). The theorem statement can be easily seen from the fact that for any fixed t, $v(t)$ is non-decreasing in c, h, and p and non-increasing in r and s. Because t^* minimizes $v(t)$, the theorem statement then follows directly. □

ACKNOWLEDGEMENTS

The authors are grateful for the invaluable comments and suggestions provided by the referees and the associate editor on earlier versions of this article. They also thank Serhan Ziya at the University of North Carolina and seminar participants at the University of Rochester and Purdue University for their helpful comments.

REFERENCES

[1] M. Abramowitz and I.A. Stegun, "Handbook of mathematical functions with formulas, graphs, and mathematical tables," Dover, 1972.
[2] R. Anupindi and R. Akella, Diversification under supply uncertainty, Management Sci 39 (1993), 944–963.
[3] A. Arreola-Risa and G.A. DeCroix, Inventory management under random supply disruptions and partial backorders, Nav Res Logistics 45 (1998), 687–703.
[4] U. Bagchi, J.C. Hayya, and C.-H. Chu, The effect of lead-time variability: The case of independent demand, J Oper Manage 6 (1986), 159–177.
[5] R.G. Brown, "Statistical forecasting for inventory control," McGraw-Hill, New York, 1959.
[6] J. Chod and N. Rudi, Strategic investments, trading and pricing under forecast updating, Management Sci 52 (2006), 1913–1929.
[7] J. Chod, N. Rudi, and J.A. Van Mieghem, Mix, time and volume flexibility: Valuation and corporate diversification, Work pap 2004.
[8] M.E. Ferguson, When to commit in a serial supply chain with forecast updating, Nav Res Logistics 50 (2003), 917–936.
[9] M.E. Ferguson, G.A. DeCroix, and P.H. Zipkin, Commitment decisions with partial information updating, Nav Res Logistics 52 (2005), 780–795.
[10] G. Gallego and Ö Özer, Integrating replenishment decisions with advance demand information, Management Sci 47 (2001), 1344–1360.
[11] S.C. Graves, D.B. Kletter, and W.B. Hetzel, A dynamic model for requirements planning with application to supply chain optimization, Oper Res 46 (1998), S35–S49.
[12] R. Güllü, On the value of information in dynamic production inventory problems under forecast evolution, Nav Res Logistics 43 (1996), 289–303.
[13] D. Gupta, The (Q,r) inventory system with an unreliable supplier, INFOR 34 (1996), 59–76.
[14] H. Gurnani and C.S. Tang, Optimal ordering decisions with uncertain cost and demand forecast updating, Management Sci 45 (1999), 1456–1462.
[15] R. Hariharan and P. Zipkin, Customer-order information, leadtimes, and inventories, Management Sci 41 (1995), 1599–1607.
[16] W.H. Hausman, Sequential decision problems: A model to exploit existing forecasters, Management Sci 16 (1969), B93–B111.
[17] D.C. Heath and P.L. Jackson, Modeling the evolution of demand forecasts with application to safety stock analysis in production/distribution systems, IIE Trans 26 (1994), 17–30.
[18] W.J. Hopp and M.L. Spearman, Setting safety leadtimes for purchased components in assembly systems, IIE Trans 25 (1993), 2–11.
[19] P. Kouvelis and J. Li, Contingency strategies in managing supply systems with uncertain lead times, Management Sci (2007).
[20] A. Kumar, Component inventory costs in an assembly problem with uncertain supplier leadtimes, IIE Trans 21 (1989), 112–121.
[21] R.R. Meyer, M.H. Rothkopf, and S.A. Smith, Reliability and inventory in a production-storage system, Management Sci 25 (1979), 799–807.
[22] J.M. Milner and P. Kouvelis, Order quantity and timing flexibility in supply chains: The role of demand characteristics, Management Sci 51 (2005), 970–985.
[23] M. Parlar and D. Berkin, Future supply uncertainty in EOQ models, Nav Res Logistics 38 (1991), 107–121.
[24] M.L. Puterman, "Markov decision processes: Discrete stochastic dynamic programming," Wiley-Interscience, New York, 2005.
[25] H. Shore, Setting safety leadtimes for purchased components in assembly systems: A general solution procedure, IIE Trans 27 (1995), 634–637.
[26] J.-S. Song, The effect of leadtime uncertainty in a simple stochastic inventory model, Management Sci 40 (1994), 603–613.
[27] J.-S. Song, C.A. Yano, and P. Lerssrisuriya, Contract assembly: Dealing with combined supply lead time and demand quantity uncertainty, Manuf Serv Oper Manage 2 (2000), 287–296.

[28] J.-S. Song and P.H. Zipkin, Inventory control with information about supply conditions, Management Sci 42 (1996), 1409–1419.
[29] F.W. Steutel and K. van Harn, Infinite divisibility of probability distributions on the real line, Marcel Dekker; 1 edition (September 1, 2003).
[30] J.M. Swaminathan and J.G. Shanthikumar, Supplier diversification: Effect of discrete demand, Oper Res Lett 24 (1999), 213–221.
[31] C.S. Tang, K. Rajaram, A. Alptekinoğlu, and J. Ou, The benefits of advance booking discount programs: Model and analysis, Management Sci 50 (2004), 465–478.
[32] L.B. Toktay and L.M. Wein, Analysis of a forecasting-production-inventory system with stationary demand, Management Sci 47 (2001), 1268–1281.
[33] B. Tomlin, On the value of mitigation and contingency strategies for managing supply-chain disruption risks, Management Sci 52 (2006), 639–657.
[34] B. Tomlin, Disruption-management strategies for short life-cycle products, Nav Res Logistics 56 (2009), 318–347.
[35] B. Tomlin, The impact of supply-learning when suppliers are unreliable, Manuf Serv Oper Manag 11 (2009), 192–209.
[36] B. Tomlin and Y. Wang, On the value of mix flexibility and dual sourcing in unreliable newsvendor networks, Manuf Serv Oper Manage 7 (2005), 37–57.
[37] T. Wang and B.L. Toktay, Inventory management with advance demand information and flexible delivery, Management Sci 54 (2008), 716–732.
[38] Z.K. Weng, Channel coordination and quantity discounts, Management Sci 41 (1995), 1509–1522.
[39] Z.K. Weng and M. Parlar, Integrating early sales with production decisions: Analysis and insights, IIE Trans 31 (1999), 1051–1060.
[40] C.A. Yano, Stochastic leadtime in two-level assembly systems, IIE Trans 19 (1987), 371–378.
[41] C.A. Yano and H.L. Lee, Lot sizing with random yields: A review, Oper Res 43 (1995), 311–334.
[42] P.H. Zipkin, Foundations of inventory management, McGraw-Hill/Irwin; 1 edition (January 24, 2000).