NSF GRANT #0905039/0800580
NSF PROGRAM NAME: Manufacturing Enterprise Systems

The Correlation of Supply Chain Risk and Financial Market Returns

Sridhar Seshadri
University of Texas at Austin

Vishal Gaur, Nikolay Osadchiy
Cornell University, New York University

Abstract: The retail sales are documented to be correlated with financial market returns. Through the placed orders, the effect propagates to the upstream members of the supply chain. In this paper we model and provide empirical evidence for the correlation between financial market returns and the sales in a supply chain. We observe that the effect of financial market is amplified for the upstream members of a supply chain and identify possible factors contributing to it. We conclude with discussing the applications to forecast updating, inventory and risk management.

1. Introduction: The retail sales are documented to be correlated with the state of the economy. The economic booms and recessions that happened over the last 20 years provided ample examples of retailers' growth and fall along with the economy, including recent bankruptcies of Circuit City, Linens 'N Things and store closings at many other retailers in response to the current economy downturn. The effect is not limited to retailing. In this paper we focus on the effects of the state of economy on retailers and their supply chains that include wholesalers, distributors and manufacturers. The goal is to understand the cross correlations between the sales in the supply chain and the state of the economy, and improve the sales forecast accuracy through incorporating the information on the state of the economy. The improved forecast accuracy reduces inventory cost and increases profits. Although the approach we present can be beneficial for all members of the supply chain, we observe that the upstream members of supply chains are more exposed to changes in the state of the economy, and also have substantially higher sales volatility, which makes them better suited for the implementation of the proposed forecast updating methodology.

The model and results presented in this paper are extensions of Gaur et al.(2008) [1], the paper written as a part of the work under the NSF grant. We consider the firm-level, dollar-denominated annual sales. Forecasts of annual sales of firms are generated by equity analysts up to three years in advance of their fiscal year-end dates, and are updated frequently. We model the joint evolution of the sales forecasts in a supply chain and the financial market using a continuous time stochastic process. We represent the financial market by one broad market indicator, such as the value weighted market index. Using a contemporaneous return on a tradable asset allows to apply our results for hedging and risk management purposes.

Our model has two key parameters which are estimated empirically. One is the correlation matrix of between sales forecast errors for the supply chain members and the financial market. The other is the vector of volatilities of the sales forecast errors. The effect of the state of economy is proportional to the product of the correlation coefficient and the demand volatility. We observe that the sales volatility increases as we move further upstream the supply chain, possibly due to the bullwhip effect. For example, our results show that the volatility of sales forecast errors at apparel manufacturers is 50% higher than at retailers. The increased volatility can also be due to the behavior of the downstream supply chain members. We develop a model that is aimed to isolate the effect of the downstream supply chain members from the state of the economy. In the model, the order fulfillment rates appear to play a critical role, as both correlation with the market (in absolute value) and volatility of the forecast errors increase with the service rate. This is consistent with the fact that the upstream members of a supply chain have more excess capacity, hence can accommodate more orders in the periods of booming economy.

The primary application of our results is forecast updating. We show that forecast updates from our model provide new information not contained in the forecast updates by equity analysts. Hence, a combination of these two inputs has higher average forecast accuracy than equity analysts. This result is most useful for firms with high correlation coefficients. Beyond forecast updating, our model can be applied for risk management, calculating the value of decision postponement, and planning other operational decisions. Several recent research papers have discussed financial hedging based on demand being correlated with the price of a traded asset (see, for example, [2] and [3]). The correlation information obtained from our model can be used to inform such hedging decisions

The paper is organized as follows. In Section 2 we present a generalized empirical model of supply chain. In Section 3 we present a summary of empirical results for retail industry, available at [1]. We analyze the results for the upstream members of the supply chain in Section 4. Section 5 presents a simple analytical model that explains the observed empirical results. We conclude the paper in Section 6 with discussing the applications to forecast updating, inventory and risk management.

2. A Generalized Model of Supply Chains: This model is an extension of the model of [1], aimed to accommodate the direct interaction between supply chain members (the bullwhip effect) in addition to the effect of the financial market.

For a two stage supply chain, consider the evolution of the following processes:

Information variable D_t^M: time t information about sales at manufacturers level for a certain fiscal year,

$$\ln D_t^M = \ln D_0^M + \int_0^t \sigma_D^M(u) \mathrm{d}B_D^M(u). \tag{1}$$

Information variable D_t^R: time t information about sales at retailers level for the same fiscal year,

$$\ln D_t^R = \ln D_0^R + \int_0^t \sigma_D^R(u) \mathrm{d}B_R^M(u). \tag{2}$$

Market index

$$\ln M_t = \ln M_0 + \mu_M t + \int_0^t \sigma_M(u) \mathrm{d}B_M(u), \tag{3}$$

where $\{B_D^M\}, \{B_D^R\}, \{B_M\}$ are the standard brownian motion processes with the correlation matrix

$$\Sigma = \begin{bmatrix} 1 & \rho_{12} & \rho_2 \\ \rho_{12} & 1 & \rho_1 \\ \rho_2 & \rho_1 & 1 \end{bmatrix}. \tag{4}$$

In the correlation matrix ρ_1 and ρ_2 represent the correlation between sales and the market for the retailer and manufacturer, respectively, and ρ_{12} is the correlation between retailer's and manufacturer's sales, after controlling for the market effect. All three stochastic processes are considered over the same time period $[0, T]$, where T is an end time of the fiscal year under consideration. Note that the set-up can be extended to an N-level supply chain.

To derive an expression suitable for data fitting, we suppose that at $t = 0$ the manufacturer has a perfect lookahead for: Case 1) $B_M(T)$, or Case 2) $B_M(T)$ and $B_D^R(T)$.

Then in Case 1), if $B_M(T) = \frac{1}{\sigma_M}\left(\ln(1 + r_{0T}) - \mu_M T + \frac{1}{2}\sigma_M^2 T\right) = a$,

$$(B_D^M(T), B_D^R(T)) \sim N(\bar{\mu}, \bar{\Sigma}),$$

where

$$\bar{\mu} = \begin{bmatrix} \rho_2 \\ \rho_1 \end{bmatrix} a \text{ and } \bar{\Sigma} = T \begin{bmatrix} 1 - \rho_2^2 & \rho_{12} - \rho_1\rho_2 \\ \rho_{12} - \rho_1\rho_2 & 1 - \rho_1^2 \end{bmatrix}$$

The Case 1) is mathematically identical to the model of [1], where they consider an interaction between a retailer and the market. The only difference between the retailer and the manufacturer is in values of parameters ρ_1, ρ_2 and σ_D^R, σ_D^M. To avoid duplication we consider only the retailer below.

Using the expression for the analysts' forecast $F_0^R = D_0^R \exp(\sigma_D^R T/2)$, the conditional expectation and variance of the sales can be written as:

$$\begin{aligned} \mathbb{E}[\ln D_T^R | \mathcal{F}_0, M_T] = \\ \ln F_0^R + \frac{\rho_1 \sigma_D^R}{\sigma_M} \ln(1 + r_{0T}) \\ + T\frac{\rho_1 \sigma_D^R}{\sigma_M}\left(-\mu_M + \frac{\sigma_M^2}{2} - \frac{\sigma_D^R \sigma_M}{2\rho_1}\right), \end{aligned} \tag{5}$$

$$\mathbb{V}\mathbf{ar}[\ln D_T^R | \mathcal{F}_0, M_T] = (\sigma_D^R)^2 T(1 - \rho_1^2). \tag{6}$$

At time $T_1 > t > 0$, where T_1 is the start of a selling season, the forecast can be updated with current financial information M_t, by the following formulae:

$$\begin{aligned} \mathbb{E}[\ln D_T^R | \mathcal{F}_0, M_t] = \ln F_0^R - \frac{1}{2}(\sigma_D^R)^2 T \\ + \frac{\rho_1 \sigma_D^R}{\sigma_M}\left(\ln(1 + r_{0t}) - \left(\mu_M - \frac{1}{2}\sigma_M^2\right) t\right), \end{aligned} \tag{7}$$

and

$$\mathbb{V}\mathbf{ar}[\ln D_T | \mathcal{F}_0, M_t] = (\sigma_D^R)^2 (T - t\rho_1^2). \tag{8}$$

In Case 2) manufacturer has perfect lookahead to the retail sales data, if $B_D^R(T) = b$,

$$B_D^M(T) \sim N(\bar{\mu}, \bar{\Sigma}),$$

where

$$\bar{\mu} = \begin{bmatrix} \rho_{12} & \rho_2 \end{bmatrix} \begin{bmatrix} 1 & \rho_1 \\ \rho_1 & 1 \end{bmatrix}^{-1} \begin{bmatrix} b \\ a \end{bmatrix},$$

and

$$\Sigma = T - \begin{bmatrix} \rho_{12} & \rho_2 \end{bmatrix} \begin{bmatrix} 1 & \rho_1 \\ \rho_1 & 1 \end{bmatrix}^{-1} \begin{bmatrix} \rho_{12} \\ \rho_2 \end{bmatrix} T$$

Substituting $\begin{bmatrix} 1 & \rho_1 \\ \rho_1 & 1 \end{bmatrix}^{-1} = \frac{1}{1-\rho_1^2} \begin{bmatrix} 1 & -\rho_1 \\ -\rho_1 & 1 \end{bmatrix}$:

$$\bar{\mu} = \frac{b(\rho_{12} - \rho_2\rho_1) + a(\rho_2 - \rho_{12}\rho_1)}{1 - \rho_1^2}, \tag{9}$$

$$\bar{\Sigma} = T\left(1 - \frac{\rho_{12}^2 - 2\rho_{12}\rho_1\rho_2 + \rho_2^2}{1-\rho_1^2}\right). \quad (10)$$

Substitute $a = \frac{1}{\sigma_M}\left(\ln(1+r_{0T}) - \mu_M T + \frac{1}{2}\sigma_M^2 T\right)$ and $b = \frac{1}{\sigma_D^R}\ln\frac{D_T^R}{D_0^R}$ and note the relationship between the time $t=0$ sales forecast by analysts F_0^M and the information variable D_0^M, $F_0^M = D_0^M \exp(\sigma_D^M T/2)$:

$$\begin{aligned}
&\mathbb{E}[\ln D_T^M | \mathcal{F}_0, M_T, D_T^R] = \\
&\quad \ln F_0^M - \frac{1}{2}(\sigma_D^M)^2 T \\
&\quad + \frac{\sigma_D^M}{\sigma_D^R}\frac{\rho_{12}-\rho_1\rho_2}{1-\rho_1^2}\ln\frac{D_T^R}{D_0^R} \\
&\quad + \frac{\sigma_D^M}{\sigma_M}\frac{\rho_2-\rho_1\rho_{12}}{1-\rho_1^2}\ln(1+r_{0T}) \\
&\quad + \frac{\sigma_D^M}{\sigma_M}\frac{\rho_2-\rho_1\rho_{12}}{1-\rho_1^2}(\frac{1}{2}\sigma_M^2 - \mu_M)T, \quad (11)
\end{aligned}$$

and

$$\begin{aligned}
&\mathbb{V}\mathbf{ar}[\ln D_T^M | \mathcal{F}_0, M_T, D_T^R] = \\
&\quad (\sigma_D^M)^2 T\left(1 - \frac{\rho_{12}^2 - 2\rho_{12}\rho_1\rho_2 + \rho_2^2}{1-\rho_1^2}\right). \quad (12)
\end{aligned}$$

We define the forecast error as a difference between the actual and expected sales. The equations (11)-(12) can be used to set up an MLE estimation of correlation matrix Σ, and parameters σ_D^M, and σ_D^R, under the assumption of the normally distributed forecast errors.

3. Summary of Results for Retailers: Formula (5) gives the expression for the forecast error as a function of the term of the forecast, financial market return over the term of the forecast, and various volatility parameters. We use this expression to test hypotheses about the correlation coefficient between sales forecast error and financial market return over the term of the forecast. Our data set is a panel of 4,698 observations across 97 US public retailers and fiscal years 1997-2007, each year containing multiple forecasts made at different times for each firm. The time period for these data is remarkable in that it includes two periods of high market returns, 1998-2000 and 2003-2005, and two periods of low or negative market returns, 2001-2002 and 2006-07. Another output of our model is a method to revise sales forecasts using up-to-date information on financial market prices (using (7)). We evaluate the accuracy of these forecasts and compare them to the forecast revisions issued by equity analysts at the same time epochs. Thus, we determine if experts incorporate publicly available historical financial return information in their forecast updates, and if our model adds new information augmenting their forecasts.

Our results show that the correlation between the sales forecast error and market return is significant. With an average value of 0.17 across our entire data set, it varies across retail segments from a high of 0.83 for office supplies stores to not significant or negative for shoe stores and auto parts and accessories stores (see Table 2 for the segment-by-segment results, and [1] for details). We classify retail segments into two categories, discretionary purchase and everyday use, analogous to a 'needs versus wants' distinction in economics. We find that retail segments in the discretionary purchase group, such as jewelry, home improvement, consumer electronics and apparel, have higher correlation coefficient estimates on average than retail segments in the everyday purchase group, such as discounters and wholesale clubs. We further classify retailers in each segment into high- and low-gross margin groups, and find that the average correlation coefficient for high-margin retailers is significantly higher than that for low-margin retailers. Our results are remarkable because they not only fit our expectations of retailers, but also show a wide range of variation across firms. For example, in the jewelry stores segment, Tiffany's, a high-margin retailer, has an estimated correlation coefficient of 0.87, whereas Zale Corp., a lower margin retailer has 0.29. In the discounters and wholesale clubs segment, BJ's Wholesale Club, a warehouse club selling durable discretionary purchase items has an estimated correlation coefficient of 0.91, whereas Walmart Stores, a discounter has 0.01, and Dollar General has -0.81. We further test the hypothesis whether the correlation coefficient increases with the term of the forecast. We find that its average value across the data set almost doubles as the term of the forecast increases from 12-16 months to 20-24 months.

4. Empirical Results for Manufacturers: The estimates in this section are based on the data for apparel and household appliances manufacturing segments. Our first set of estimates is based on the Case 1) model for the manufacturers, i.e., we derive the log-likelihood function from (5) and (6). Actual sales, analysts' sales forecasts and returns on value-weighted market index were obtained from I/B/E/S and CRSP databases respectively, accessed through the WRDS service. The data set summary is presented in Table 1.

Table 2 contains estimates of the correlation between sales forecast error and the market (ρ) and forecast error volatility (σ_D) on the pooled samples. For comparison, we include estimates of the same parameters for the apparel and consumer electronics retailers, also pooled. Notice the increase in volatility of forecast errors for the manufacturers. Estimates on the firm-by-firm basis support this conclusion. Pooled estimate of the correlation coefficient ρ also increases as we move upstream of the supply chain, however the result is not strongly supported by the firm-by-firm estimates. For example median correlation coefficient for the manufacturers of household ap-

Table 1: Summary statistics for the manufacturers data

Segment	NAI-CS	# Obs.	# Firms	# Firm years	Examples
Apparel manuf.	315	1553	26	150	Oxford Industries, LVMH
Househols appliances manuf.	3352	383	7	37	Maytag, Whirlpool

Table 2: WMLE estimates for the manufacturers and retailers

Segment	ρ	ρ_{median}	σ_D
	Manufacturers		
Appliances	.3066***	.1522	.2296***
Apparel	.2968***	.3720	.1261***
	Retailers		
Appliances	.0403	.4199	.0864***
Apparel	.2681***	.2938	.0971***

pliances is lower than for the electronics retailers. Overall, we observe an empirical support for amplification of the forecast errors and increased effect of the financial markets for upstream members of supply chains. In the next Section we develop a simple analytical model to explain the increase in volatility and the correlation coefficients for the upstream members of a supply chain.

5. Two Stage Supply Chain Model: Consider a two stage supply chain consisting of a retailer r and manufacturer m. The retailer facing a random customer's demand arriving at time T has to place an order to a manufacturer at time $t > 0$. Manufacturer, anticipating an order from retailer starts the production and/or ordering of raw materials at time $t = 0$. Suppose, that according to the time zero information, the consumer's demand is normally distributed with mean μ_0, variance $\sigma^2 T$ and CDF F_0. To model an effect of the state of the economy, we assume that the demand distribution evolves with time, and at time of the retailer's order its mean becomes μ_t, variance - $\sigma^2(T-t)$, and CDF - F_t. By the newsvendor model the retailer's order quantity is therefore $q_t^r = \mu_t + z_\alpha\sigma\sqrt{T-t}$, where α is the service level of the retailer. The value of μ_t is random at time 0, so we assume that $\mu_t \sim G = N(\mu_0, k^2 t)$. Accounting for the uncertainty of μ_t, order quantity of the manufacturer at time 0 is $q^m = \mu_0 + z_\alpha\sigma\sqrt{T-t} + z_\beta k\sqrt{t}$. Based on these results the forecast errors can be derived as follows.

Retailer's expected sales:

$$\mathbb{E}(S^r) = \int_0^{q_t^r} D\mathrm{d}F_t + (1-\alpha)q_t^r$$

and the forecast error is $\varepsilon^r \triangleq \mathbb{E}(S^r) - \min(q_t^r, q^m)$.

Similarly for the manufacturer:

$$\begin{aligned}\mathbb{E}(S^m) &= \int_0^{q^m} q_t^r \mathrm{d}G + q^m\mathbb{P}(q^m \le q_t^r) \\ &= \int_0^{q^m} \mu_t \mathrm{d}G + \beta z_\alpha\sigma\sqrt{T-t} + (1-\beta)q^m\end{aligned}$$

and the forecast error is $\varepsilon^m \triangleq \mathbb{E}(S^m) - \min(q_t^r, q^m)$.

Notice that $\mathbb{E}(\varepsilon^r) = \mathbb{E}(\varepsilon^m) = 0$. Consider the variance of the forecast errors and the conditions under which $\mathbb{V}\mathbf{ar}(\varepsilon^m) > \mathbb{V}\mathbf{ar}(\varepsilon^r)$.

$$\begin{aligned}\mathbb{V}\mathbf{ar}(\varepsilon^r) &= \mathbb{V}\mathbf{ar}(\min(D, q_t^r)) \\ &= \mathbb{V}\mathbf{ar}(\min(D, F_t^{-1}(\alpha)) \\ &= \mathbb{V}\mathbf{ar}(\min(D, \mu_t + z_\alpha\sigma\sqrt{T-t})); \\ \mathbb{V}\mathbf{ar}(\varepsilon^m) &= \mathbb{V}\mathbf{ar}(\min(q_t^r, q^m)) \\ &= \mathbb{V}\mathbf{ar}(\min(F_t^{-1}(\alpha), G^{-1}(\beta) + F_t^{-1}(\alpha) - \mu_t)) \\ &= \mathbb{V}\mathbf{ar}(\min(\mu_t, G^{-1}(\beta))) \\ &= \mathbb{V}\mathbf{ar}(\min(\mu_t, \mu_0 + z_\beta k\sqrt{t}))\end{aligned}$$

Note, that asymptotically, as $\alpha, \beta \to 1$, $\mathbb{V}\mathbf{ar}(\varepsilon^r) \to \sigma^2(T-t)$ and $\mathbb{V}\mathbf{ar}(\varepsilon^m) \to k^2 t$, so that asymptotic volatilities are σ and k respectively. Both expressions are the variances of censored normal distributions, with the censoring thresholds are set at the quantiles α and β respectively. We will use the following lemma (can be proved by taking derivative of the variance w.r.t. censoring limit):

Lemma 1 *For any random variable X that admits a density function, $\mathbb{V}\mathbf{ar}(\min(X, c))$ is increasing in c.*

Higher censoring thresholds correspond to greater variance, therefore, $\mathbb{V}\mathbf{ar}(\varepsilon^m)/t > \mathbb{V}\mathbf{ar}(\varepsilon^r)/(T-t)$ if $\beta > \alpha$.

We attribute changes in the mean demand μ_t to the evolution of the state of the economy which can be proxied by a return on a broad financial index. This translates into the stochastic process $\{\mu_t\}$ being correlated with the process describing evolution of the financial index $\{M_t\}$. Assume that $\mu_t = \mu_0 + aR_{(0,t]} + \epsilon$, where $R_{(0,t]} = M_t/M_0 - 1$ and $D = \mu_t + bR_{(t,T]} + \delta$, a, b are scalars and ϵ, δ are independent random variables with zero mean. Then the correlation of forecast errors with the market performance is:

$$
\begin{aligned}
&\mathbb{C}\mathbf{orr}(R_{(t,T]}, \varepsilon^r) = \\
&\quad = \mathbb{C}\mathbf{orr}(R_{(t,T]}, \min(q_t^r, D)) \\
&\quad = \mathbb{C}\mathbf{orr}(R_{(t,T]}, \min(F_t^{-1}(\alpha), \mu_t + bR_{(t,T]} + \delta)) \\
&\quad = \mathbb{C}\mathbf{orr}(R_{(t,T]}, \min(\mu_t + z_\alpha \sigma \sqrt{T-t}, \mu_t + bR_{(t,T]} + \delta)) \\
&\quad = \mathbb{C}\mathbf{orr}(R_{(t,T]}, \min(z_\alpha \sigma \sqrt{T-t}, bR_{(t,T]} + \delta)), \text{ and} \\
&\mathbb{C}\mathbf{orr}(R_{(0,t]}, \varepsilon^m) = \\
&\quad = \mathbb{C}\mathbf{orr}(R_{(0,t]}, \min(q_t^r, q^m)) \\
&\quad = \mathbb{C}\mathbf{orr}(R_{(0,t]}, \min(\mu_t, G^{-1}(\beta))) \\
&\quad = \mathbb{C}\mathbf{orr}(R_{(0,t]}, \min(\mu_0 + aR_{(0,t]} + \epsilon, \mu_0 + z_\beta k\sqrt{t})) \\
&\quad = \mathbb{C}\mathbf{orr}(R_{(0,t]}, \min(aR_{(0,t]} + \epsilon, z_\beta k\sqrt{t})).
\end{aligned}
$$

Again, we see that as the service level and leadtime increase the *absolute* value of the correlation coefficient increases. Typically, manufacturers have longer leadtime, i.e. $t \geq T - t$. Together with the condition on the service rates $\beta \geq \alpha$, this is sufficient to explain higher degree of correlation between the forecast errors and financial market performance at the upstream levels of supply chain.

6. Applications:

6.1. Forecast Updating: The main application of our model is to forecast updating. Firm-level sales forecasts are issued by retailers for planning and by equity analysts for estimating earnings. In both cases, substantial effort is required to generate and update forecasts. In contrast, financial data are readily available and highly reliable. Thus, our model, together with financial data, can be used to update sales forecasts in a cost-efficient way. In [1], we test the accuracy of the model for forecast updating by comparing errors of the forecasts generated by the model vis-a-vis forecasts issued by equity analysts at the same times for the panel of retailers. We find that, on average, analysts have lower forecast root mean squared error than our model. This is not surprising because analysts have access to many types of information for their forecast updates, whereas our model-based updates are based on only one variable. Indeed, we find only a weak evidence that analysts take financial market returns into account in their forecast updates. We conduct a regression of analyst forecast updates on financial returns over the period of update as per our model, and find that the financial returns are weakly significant even though they were public information at the times when analysts updated their forecasts. To further examine the usefulness of our model, we construct a combined model based on analysts' forecast updates as well as our model's forecast updates, and find that both variables are statistically significant in improving forecast accuracy. Thus, a combined forecast obtained from both updates perform better than the equity analysts. Furthermore, as expected, the improvement is higher for firms with high correlation coefficients and for longer term forecasts.

We give the following example of this model application. Suppose at time t a retail firm needs to make an operational decision based on the demand forecast made by experts at time t_0. The firm has three options: use the original forecast for the decision, hire experts to update the forecast, or use current financial information and our model to update the forecast. Parameters of the model ρ and σ_D can be either estimated from the historical data or taken as average values for the segment and margin group of the company (see Table 2). Then, formula (7) provides an expression for the forecast update incorporating the current financial information.

To demonstrate the benefits of this approach we perform out-of-sample testing. The observations are split into two, an estimation and testing subsamples, that include years 1996-2004 and 2005 respectively. The first subsample is then used to estimate the model parameters. Then, for every company in the 2005 sample, we take the forecast with the longest term as a starting point and perform forecast updating using the contemporary financial information according to (7). We choose the forecast update epochs to coincide with the times of forecast updates issued by experts. The performance of our model is then benchmarked to the performance of the experts using the root mean squared error, on the industry segment basis. Figure 1 shows how RMSE reduces when the forecast is updated by experts and the model, depending on the time between the original forecast and the update. The results demonstrate, that typically the forecast update generated by experts results in the smaller RMSE. However, we were able to find retail segments where our model comes close to the experts' performance and even surpasses it.

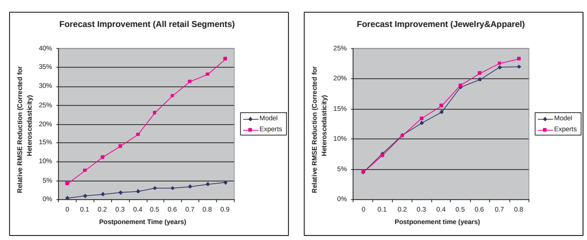

Figure 1: Forecast improvement by the experts and the model, characterized by the RMSE, over all data set and for the Apparel and Jewelry segment.

We conjecture that the model forecast updating performance depends on the initial forecast error. We would expect the model's performance (measured by RMSE) to be closer to the performance of the experts when the initial forecast error is small. To substantiate this conjecture, we subdivide the sample into the three time buckets with long, medium, and short term forecasts. For every company, we pick the forecast with the longest term from every time bucket and use it as a starting point. We consider all subsequent experts' forecasts for the retailer and compare them to the model generated forecasts. For the Durable Goods, Apparel and Jewelry, and Discount retail segments, errors by the model and the experts are within one standard deviation for all time buckets. On the other hand, for the Grocery, Everyday Clothing, Recreational Goods and Computer and Electronics stores experts' forecast errors are significantly less than the ones generated by the model for long and/or medium term forecast. The difference decreases for the short term forecasts (less than four month from the start of the selling season). Therefore, the model can be used for the forecast updating in the Durables, Apparel and Jewelry and Discount retail segments for all terms of the forecast (less than 24 months), where as in the other retail segments the effectiveness of the model is limited to the shorter term forecast updates.

6.2. Decision Postponement Model: The results of Section 6.1 show that the model can be used efficiently for forecast updating. Suppose a retail company has to make an operational decision, such as a procurement of a raw material for its private brand product. It has an option to purchase it from Supplier 1 with the leadtime l_1 (Case 1) or from Supplier 2 with the leadtime l_2 (Case 2). Assume $l_1 > l_2$. Assume also that the unit purchasing cost is greater for Supplier 2. If the company chooses Supplier 1, the order has to be placed at time T_0; for Supplier 2 it can be postponed for t periods where $t = l_1 - l_2$. Suppose the company has a demand forecast D_0 made at time T_0. The company can place an order to Supplier 1, based on the experts' forecast, or it can decide to order from Supplier 2 and use the updated forecast given the financial information available at time $T_0 + t$. We study the relative benefits of these two options.

We model the company as a profit maximizing newsvendor, and demand evolution as specified by (7) and (8). That is the optimal order quantity y is given by

$$y = F^{-1}\left(\frac{p-c}{p}\right),$$

where F is the c.d.f. of the demand, and p and c are respectively the selling price and the unit purchasing cost. If $\ln\mathcal{D} \sim N(\mu, \sigma^2)$, then $\ln y = \mu + \sigma\Phi^{-1}\left(\frac{p-c}{p}\right)$, or

$$y = \exp(\mu + \sigma z_\beta), \tag{13}$$

where $z_\beta = \Phi^{-1}(\frac{p-c}{p})$ and Φ is the standard normal c.d.f.

The optimal expected profit is given by

$$\mathbb{E}(\pi) = pe^{\mu+\sigma^2/2}\Phi(z_\beta - \sigma), \tag{14}$$

and the profit variance is

$$\begin{aligned}\mathbb{V}\mathbf{ar}(\pi) = p^2e^{2\mu}\Big[&e^{2\sigma^2}(1+\Phi(2\sigma - z_\beta)) - e^{\sigma^2} \\ &-e^{2\sigma z_\beta}\Phi(-z_\beta) - (e^{\sigma^2/2}\Phi(\sigma - z_\beta) - e^{\sigma z_\beta}\Phi(-z_\beta))^2\Big].\end{aligned} \tag{15}$$

If the retailer chooses to order from Supplier 1 (Case 1), the demand has the following distribution:

$$\ln\mathcal{D} \sim N\left(\ln D_0 - \frac{1}{2}\sigma_D^2 T,\ \sigma_D^2 T\right). \tag{16}$$

If the retailer chooses to postpone and order from Supplier 2 (Case 2), the conditional distribution of the demand given that return on the financial index realized over the period from T_0 to $T_0 + t$ equals r_t is:

$$\begin{aligned}\ln\mathcal{D} \quad\sim\quad & N\bigg(\ln D_0 - \frac{1}{2}\sigma_D^2 T + \frac{\rho\sigma_D}{\sigma_S}(\ln(1+r_t)\ , \\ & -(\mu_S - \sigma_S^2/2)t)\sigma_D^2(T - t\rho^2)\bigg).\end{aligned}$$

In Case 1, the expectation and the variance of the profit are given by (14) and (15) where parameters μ and σ of the demand distribution are defined by (16). That is,

$$\mathbb{E}(\pi_1) = pD_0\Phi(z_\beta - \sigma_D\sqrt{T}), \tag{17}$$

and

$$\begin{aligned}\mathbb{V}\mathbf{ar}(\pi_1) = p^2D_0^2\Big[&e^{\sigma^2}(1+\Phi(2\sigma - z_\beta)) - 1 - \\ &-e^{2\sigma z_\beta - \sigma^2}\Phi(-z_\beta) - (e^{\sigma^2/2-\sigma}\Phi(\sigma - z_\beta) - e^{z_\beta}\Phi(-z_\beta))^2\Big],\end{aligned}$$

where $\sigma = \sigma_D\sqrt{T}$.

In Case 2, conditional on r_t, the expectation and the variance of the profit are still given by (14) and (15), and parameters of the demand distribution are given by (17). By unconditioning, the expected profit can be written as

$$\mathbb{E}(\pi_2) = \mathbb{E}\left(\mathbb{E}(\pi_2|r_t)\right) = pD_0\Phi\left(z_\beta - \sigma_D\sqrt{T - t\rho^2}\right).$$

Profit variance equals

$$\mathbb{V}\mathbf{ar}(\pi_2) = \mathbb{V}\mathbf{ar}(\mathbb{E}(\pi_2|r_t)) + \mathbb{E}(\mathbb{V}\mathbf{ar}(\pi_2|r_t)),$$

where

$$\mathbb{V}\mathbf{ar}(\mathbb{E}(\pi_2|r_t)) = p^2D_0^2\Phi^2(z_\beta - \sigma)(e\rho^2\sigma_D^2 t - 1),$$

$$\begin{aligned}\mathbb{E}(\mathbb{V}\mathbf{ar}(\pi_t|r_t)) = p^2D_0^2e^{-\sigma_D^2(T-2\rho^2 t)}\Big\{&e^{2\sigma^2}(1+\Phi(2\sigma - z_\beta)) \\ &-e^{\sigma^2} - e^{2\sigma z_\beta}\Phi(-z_\beta) - (e^{\sigma_2/2}\Phi(\sigma - z_\beta) - e^{\sigma z_\beta\Phi(-z_\beta)^2}\Big\},\end{aligned}$$

and $\sigma = \sigma_D\sqrt{T - t\rho^2}$.

Note that expressions for the Case 1 and the Case 2 give identical results if $\rho = 0$. In that case the retailer does not benefit from the postponement. We compare the expected profits in cases 1 and 2 as well as the profit variance. Figure 2 shows the relative increase in the expected profit due to the decision postponement and the relative decrease in the standard deviation of the profit. Increase in the expected profit can be substantial for the companies with both high $|\rho|$ and σ_D, whereas only high $|\rho|$ is sufficient for the decrease in the profit variance. If the

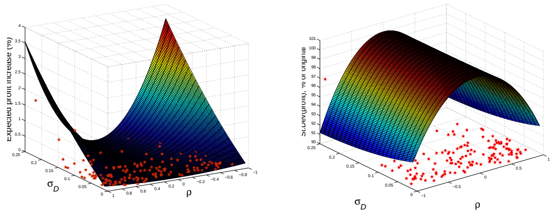

Figure 2: Relative profit increase and reduction of the standard deviation of the profit due to the postponement. t=6 months, $\mu_s = 0.2$, $\sigma_S = 0.18$. Estimates of ρ and σ_D are plotted on the horizontal plane.

retailer is risk neutral, the order should be postponed as long as profit increase is greater than the cost differential associated with purchasing from Supplier 2.

Consider another example. Suppose now that a retailer needs to purchase two different products, one from Supplier 1 and the other from Supplier 2. Assume that the lead times of the suppliers are $l_1 > l_2$. Demand forecasts for both products are generated by experts. Suppose both suppliers offer the company an option to decrease the lead time by the same period. Shorter lead time will allow the retailer to postpone the decision, update the forecast, therefore increase its expected profit as well as decrease its variance. Note that both the profit increase and the variance decrease are higher for the higher $|\rho|$. Now, the forecast for the Product 1 has a longer term, therefore the value of $|\rho|$ is higher. Hence the retailer should choose to decrease the longer lead time.

Conceptually this application of our model is similar to the paper of Fisher and Raman (1996)[4]. The difference is that Fisher and Raman observe the demand signal from the early sales data, whereas we observe the signal from the financial market.

For the companies with similar parameters σ_D, the value of postponement depends on the value of ρ. However, in order to update the forecast efficiently, the forecast error by the model should be comparable to the error by the experts. Table 3 gives the numbers of companies, their retail segments, and the average postponement value (as percent of the profit increase and variance reduction).

6.3. Risk Management: The dispersion in the correlation coefficients ρ across retailers presents an opportunity for risk pooling. Consider two retailers that have sales forecasts D_1 and D_2 with the term T. Let sales volatilities be σ_1 and σ_2 respectively and the correlation structure between $\ln(1+r)$, $\ln \mathcal{D}_1$ and $\ln \mathcal{D}_2$ be given by

$$\Sigma = \begin{pmatrix} 1 & \rho_1 & \rho_2 \\ \rho_1 & 1 & \rho_{12} \\ \rho_2 & \rho_{12} & 1 \end{pmatrix}. \quad (18)$$

Table 3: Value of decision postponement for the companies in the data set: number of companies, retail segments, average expected profit increase and profit variance reduction.

		Value of the correlation coefficient ρ	
		Below the median	Above the median
Model's performance relative to the experts	Poor	26 companies; no dominating segment. **Postponement on the case by case basis.**	13 companies, no dominating segment. **Postpone, Use experts**
	Good	35 companies; Grocery stores, Everyday Clothing. **Do not postpone.**	47 companies; apparel and jewelry, durables, discounters; $\bar{\rho} = 0.71$, $\bar{\sigma}_D = 0.07$, $\Delta\mathbb{E}[\pi] = 0.5\%$, $\Delta\mathbb{V}\mathbf{ar}[\pi] = 4\%$. **Postpone, use model.**

At time t_0, sales D_1 and D_2 are random variables with the following distribution

$$\ln \mathcal{D}_i \sim N(\ln D_i - \frac{1}{2}\sigma_i^2 T, \sigma_i^2 T),\ i = 1, 2.$$

Therefore

$$\begin{aligned} \mathbb{E}[\mathcal{D}_i] &= D_i \quad \text{and} \\ \mathbb{V}\mathbf{ar}[\mathcal{D}_i] &= (\exp(\sigma_i^2 T) - 1) D_i^2 \exp(\sigma_i^2 T),\ i = 1, 2. \end{aligned}$$

Consider a company with the pooled sales $\mathcal{D} = \mathcal{D}_1 + \mathcal{D}_2$. Then $\mathbb{E}[\mathcal{D}] = D_1 + D_2$ and

$$\mathbb{V}\mathbf{ar}[\mathcal{D}] = \mathbb{V}\mathbf{ar}[\mathcal{D}_1] + \mathbb{V}\mathbf{ar}[\mathcal{D}_2] + 2\mathbb{C}\mathbf{ov}[\mathcal{D}_1, \mathcal{D}_2]. \quad (19)$$

Note, $\mathbb{C}\mathbf{ov}[\mathcal{D}_1, \mathcal{D}_2] = \rho_{12}\sqrt{\mathbb{V}\mathbf{ar}[\mathcal{D}_1]\mathbb{V}\mathbf{ar}[\mathcal{D}_2]}$. Therefore the variance of the pooled sales decreases if and only if

$$\rho_{12} < 0. \quad (20)$$

Parameter ρ_{12} can be estimated from the data, by regressing residuals from the regression of $\ln D_1$ on $\ln(1+r)$ on the residuals from the regression of $\ln D_1$ on $\ln(1+r)$ (this will "detrend" the effect of the financial market). However, we can provide theoretical bounds on ρ_1 and ρ_2 that will guarantee that condition (20) is satisfied. Indeed, $\Sigma \succ 0$. Since $|\rho_1|, |\rho_2|$ and $|\rho_{12}|$ are less than unity, it is sufficient to check $\det \Sigma > 0$, which is equivalent to

$$(\rho_{12})^2 - 2\rho_1\rho_2\rho_{12} + \rho_1^2 + \rho_2^2 - 1 < 0.$$

Hence $\rho_{12} < 0$ if

$$\begin{aligned}(\rho_1\rho_2)^2 - \rho_1^2 - \rho_2^2 + 1 &> 0, & (21)\\ \rho_1\rho_2 &< 0, \text{ and} & (22)\\ \rho_1^2 + \rho_2^2 &> 1. & (23)\end{aligned}$$

Geometrically, the conditions above define the complement of the unit circle to the unit square in the second and fourth quadrants. Based on the empirical results we identify more than one hundred pairs of the companies for which the merger would result in the reduced demand variance. Similar idea can be extended to the assortment planning problem for a risk averse retailer. In addition to operational risk management, our results can be used for financial hedging and fine tuning of trading strategies.

7. Acknowledgements: The authors gratefully acknowledge the support of NSF Grants #0905039 and #0800580, NSF Division of Civil, Mechanical, and Manufacturing Innovation, Manufacturing Enterprise Systems NSF Program, Program Officer Cerry M. Klein.

8. References:

[1] V. Gaur, N. Osadchiy, and S. Seshadri, "Sales Forecasting with Financial Indicators and Experts' Input," *Working Paper*, 2008.

[2] V. Gaur and S. Seshadri, "Hedging inventory risk through market instruments," *Manufacturing & Service Operations Management*, vol. 7, no. 2, pp. 103–120, 2005.

[3] R. Caldentey and M. Haugh, "Optimal control and hedging of operations in the presence of financial markets," *Mathematics of Operations Research*, vol. 31, no. 2, p. 285, 2006.

[4] M. Fisher and A. Raman, "Reducing the cost of demand uncertainty through accurate response to early sales," *Operations Research*, vol. 44, no. 1, pp. 87–99, 1996.