ISE Industrial and Systems Engineering

Modelling high-tech product life cycles with short-term demand information: A case study

S. David Wu
Lehigh University

Berrin Aytac
Lehigh University

Report: 10W-002

Modelling high-tech product life cycles with short-term demand information: A case study

Berrin Aytac and S. David Wu

Department of Industrial and Systems Engineering, Lehigh University, Bethlehem, Pennsylvania, USA

bea4@lehigh.edu, david.wu@lehigh.edu

Abstract

Increasing competition and volatile conditions in high-tech markets result in shortening product life cycles with non-cyclic demand patterns. This study illustrates the use of a demand-characterization approach that models the underlying shape of product demands in these markets. In the approach, a Bayesian-update procedure combines the demand projections obtained from historical data with the short-term demand information provided from demand leading indicators. The goal of the Bayesian procedure is to improve the accuracy and reduce the variation of historical-data-based demand projections. This paper discusses the implementation experience of the proposed approach at a semiconductor-manufacturing company; the key test results are presented using product families introduced over the last few years with a comparison to real-world benchmark demand forecasts.

Key Words: Bayesian forecasting; leading indicators; cumulative demand growth; short life-cycle products; high-tech industry.

Introduction

The high-tech industry, including semiconductor and computer manufacturers, had an accelerated growth in the mid-to-late 1990s; companies acquired advanced manufacturing and design capabilities to supply leading-edge systems at competitive prices. The successive releases of operating systems that required increasingly more hardware power fuelled the growth. However, the drop in demand with a downturn in economic conditions in 2001 resulted in an industry-wide excess capacity, thus deferring companies' investments in new resources. Moreover, the power of computing systems continually exceeded that of software systems, enabling computers to run effectively with existing software.

These market conditions led to the development of low-cost high-power systems. For instance, in order to decrease power consumption and increase features of microprocessors, core processors that included multiple processors of existing technologies were introduced. In order to stay competitive, companies committed to the development of a major new processor every two years as opposed to four or more years as in the 1990s. In addition, a minor architectural change began to occur in the year following the introduction of each new processor. This rapid adoption of technologies led to the shortening of product life cycles.

The rising complexity of the market, as well as the products, increases the challenge in demand forecasting. Current forecasting practices in the industry are based on (1) collecting field forecasts from all the regional markets that a company serves, then (2) aggregating the field forecasts to global demand by product lines and by adjusting with additional economic and marketing factors. The final forecast is approved and released by the senior management of the company and used as the basis for capacity expansion and operational decisions. However, it is reported that there is an industry-wide need for a systematic and repeatable approach to effective demand forecasting, utilising all the accessible data and leading to less volatile, more accurate and timely responses to market-significant events.

Aytac and Wu's (2008) approach, along the lines of these objectives, provides a methodology that characterises demand using multiple sources of information and quantifies the uncertainty in demand estimates. High-tech products typically demonstrate a single-modal life-cycle pattern with high volatility. In the proposed approach, the combined estimate of a number of life-cycle growth models describes the cumulative demand. A Bayesian procedure aggregates different sources of information that describe short-term demand characteristics. The aggregated information updates the estimate of each growth model projected from historical data. The procedure provides a distributional characterization of demand with the objective of improving the accuracy and reducing the variability in this characterization.

The current study contributes to the practice of operational planning by illustrating the use of this recently proposed demand-characterization methodology and by exploring a variety of sources of short-term demand information. The implementation takes place within one main product division of a major semiconductor-manufacturing company. However, the data is disguised due to the confidentiality agreement with the company. The following discussion summarises the methodology, introduces different sources of short-term demand information, demonstrates the implementation experience and testing results, and concludes with future research directions.

Life-cycle demand characterization

A typical demand pattern of short life-cycle products follows a single-modal curve with growth, maturity, and decline phases. The technological forecasting literature, which studies the diffusion of innovations in a population, provides a large number of growth models to describe life-cycle demand realizations observed in practice. A common approach in the literature is to characterise cumulative demand. This mitigates the impact of short-term demand fluctuations on the quality of parameter estimates and provides a useful means to separating the trend from the noise in the data.

Technological growth models are, in general, based on the following characterization of the diffusion rate of an innovation (Sultan *et al*, 1990):

$$\frac{d\hat{X}(t)}{dt} = g(t)[1 - \hat{X}(t)]$$

where $\hat{X}(t)$ is the estimated cumulative sales at time t, and $g(t)$ represents the fraction of the remaining expected life-cycle sales to realise at time t. Different functional forms for $g(t)$ lead to different characterizations of the diffusion process. For instance, for the Bass model (Bass 1969), a pioneer in the technological forecasting research and one of the most widely used growth models, $g(t) = p + q\hat{X}(t)$, where the constants p and q are known as the *coefficient of innovation* and *coefficient of imitation*, respectively. When the right-hand side of the equation is expanded, the diffusion rate at any time t is expressed as the sum of the fraction of the buyers that are not influenced by the previous buyers, $p[1 - \hat{X}(t)]$, and the fraction of the buyers that are influenced by the previous buyers, $q\hat{X}(t)[1 - \hat{X}(t)]$. The solution to the differential equation is an S-shaped cumulative curve.

Among the studies that illustrate the use of diffusion modelling in practice, Kurawarwala and Matsuo (1998) focus on life-cycle modelling combined with seasonality at a computer manufacturer. In addition, Meade and Islam (1998) provide an extensive survey of the 29 diffusion models proposed in the technological forecasting literature; the performances of the individual and combined model forecasts are compared using both simulated and real-life data sets for consumer durables. Meade and Islam (2006) provide a more recent review of the literature and further explore the issues related to multi-national and multi-generational diffusions of innovations.

Bayesian demand forecasting

In the proposed approach, a group of growth models characterise the baseline demand patterns. A Bayesian procedure describes the projection of each model in probability distributions; the combined projection of models is the final estimate of demand. This procedure improves the uncertainty in the characterization of demand using short-term demand information. For this purpose, the approach repeats the following procedure for each model.

Step 1 The projection of the growth model from historical data provides the *prior information* for the Bayesian procedure (Figure 1a).

Step 2 Demand leading indicators are the sources of short-term demand information. The information from a leading indicator extends the historical demand data, and the model is reestimated over the extended data. The collection of the projections from the data extended with each leading indicator forms the *sampling information* for the estimate of the model (Figure 1b).

Step 3 The Bayesian procedure updates the prior projection using the sampling information to generate the *posterior life-cycle demand projection* (Figure 1c).

For a more formal description of the procedure, let Θ_T represent the historical data of the cumulative fraction of total demand realised over time T, i.e., $\Theta_T = \{X(1), ..., X(T)\}$. Given Θ_T, in step 1, the Bayesian procedure produces the prior projection of each model ($k = 1, ..., K$) to predict the cumulative demand at time $T + M$, i.e., $X(T + M)$, as

$$\widehat{X}_k(T + M|\Theta_T),\ k = 1, ..., K,\ M > 0.$$

Total variance of the prior estimate, σ_k^2, is equal to the sum of the variance of the model's estimate and variance of its estimation error. The former is calculated based on the explicit density approach (Meade and Islam, 1995), where random sets of parameters are drawn from the probability distributions of the parameter estimates; sample model estimates derived from each set of parameters determine the variance. On the other hand, the latter component, variance of estimation error, is approximated using the mean square error of estimation.

In the second step, each source of short-term demand information ($j = 1, ..., m$) provides an estimate for the L-period future demand, i.e., $\{X(T+1), ..., X(T+L)\}$, and extends the available

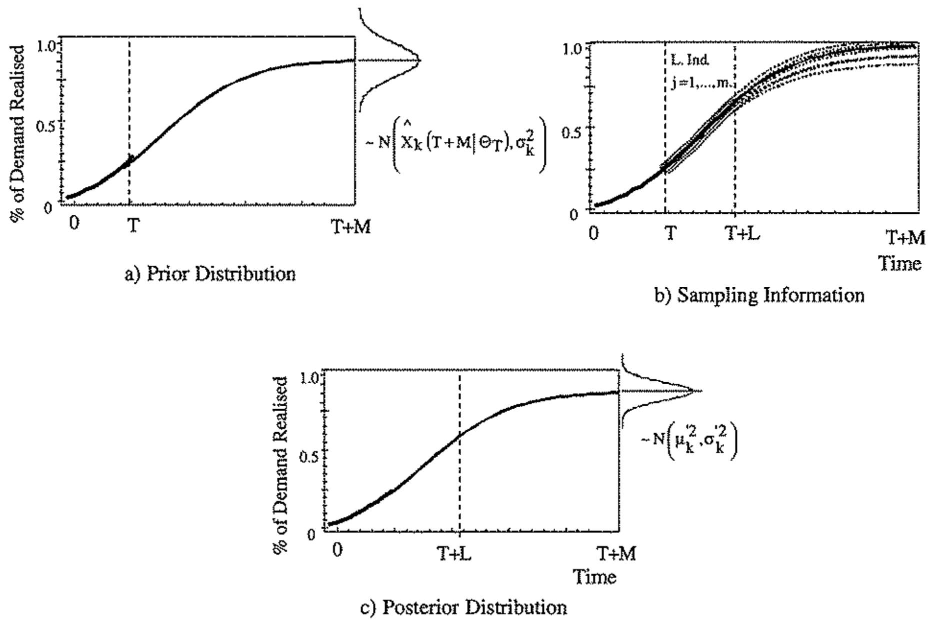

Figure 1: Bayesian update of the probability distribution of a growth model's estimate

data to $T+L$, i.e., Θ^j_{T+L}. The estimate of a growth model obtained from the extended data,

$$\widehat{X}_{kj}(T+M|\Theta^j_{T+L}),\ k=1,...,K,\ j=1,...,m,$$

is a sample estimate for the model's projection to be obtained with the actual data, Θ_{T+L}. The collection of the estimates with all the leading indicators forms the sampling information. Furthermore, the sum of the variance of leading-indicator-based estimates and variance of estimation error determines the variance of sampling information, i.e., τ_k^2.

In the third step, assuming that normal distribution characterises the sampling and estimation errors, the conjugate prior distribution is normally distributed (Press, 2003). Thereby, in the Bayesian procedure with m leading-indicator-based sampling estimates that follow a Normal distribution with mean $X(T+M)$ and variance τ_k^2, where $X(T+M)$ has prior normal distribution with mean $\widehat{X}_k(T+M|\Theta_T)$ and variance σ_k^2, the posterior estimate of model k is normally distributed as

$$\widetilde{X}_k(T+M) \sim N(\mu'_k, \sigma'^2_k), \qquad k=1,...,K$$

where

$$\begin{aligned}\mu_k' &= \frac{1/\sigma_k^2}{1/\sigma_k^2+m/\tau_k^2}\,\widehat{X}_k(T+M|\Theta_T) + \frac{m/\tau_k^2}{1/\sigma_k^2+m/\tau_k^2}\,\frac{1}{m}\sum_{j=1}^{m}\widehat{X}_{kj}(T+M|\Theta_{T+L}^j),\\ \sigma_k'^2 &= \frac{\sigma_k^2\tau_k^2}{m\sigma_k^2+\tau_k^2}.\end{aligned}$$

The updated life-cycle demand projections are the posterior means, which are the weighted averages of the prior life-cycle projections and sampling means, with the weights being inversely proportional to their variances. As a measure of the quality of the additional demand information, relative estimates of leading indicators determine the variance of sampling information and also the extent of the update on the historical-data-based projections. Hence, depending on the quality of the additional information, this weighting scheme improves the accuracy and variability of the estimates. Aytac and Wu (2008) further explore the impact of additional information on the reduction of variability in the demand characterization.

Combined demand modelling

While it is difficult to select a single growth model that best describes the historical demand and that generates the most accurate forecast, a final step that combines the projections of multiple life-cycle growth models legitimises the approach. In the combination, model k is assigned the weight φ_k that is inversely proportional to its variance:

$$\varphi_k = \frac{1/\sigma_k'^2}{\sum_{i=1}^{K} 1/\sigma_i'^2}, \quad k = 1, ..., K.$$

to obtain the forecast

$$\mu_c' = \sum_{k=1}^{K}\varphi_k\mu_k',$$

This is a weighting scheme that minimises the variance of combined forecast (Dickinson, 1973). Its effectiveness over individual model forecasts has been shown through several empirical studies (c.f., Winkler and Makridakis, 1983; Timmermann, 2006). For an outline of the entire demand-characterization approach, one can refer to Figure 2.

Important to note are the studies that suggest Bayesian approaches to forecasting demand over product life cycles. For instance, in Zhu and Thonemann (2004), the Bass model characterises demand, and a Bayesian procedure updates the distribution of its parameters using parameter estimates from earlier product generations. In another study, Bewley and Griffiths (2001) compare

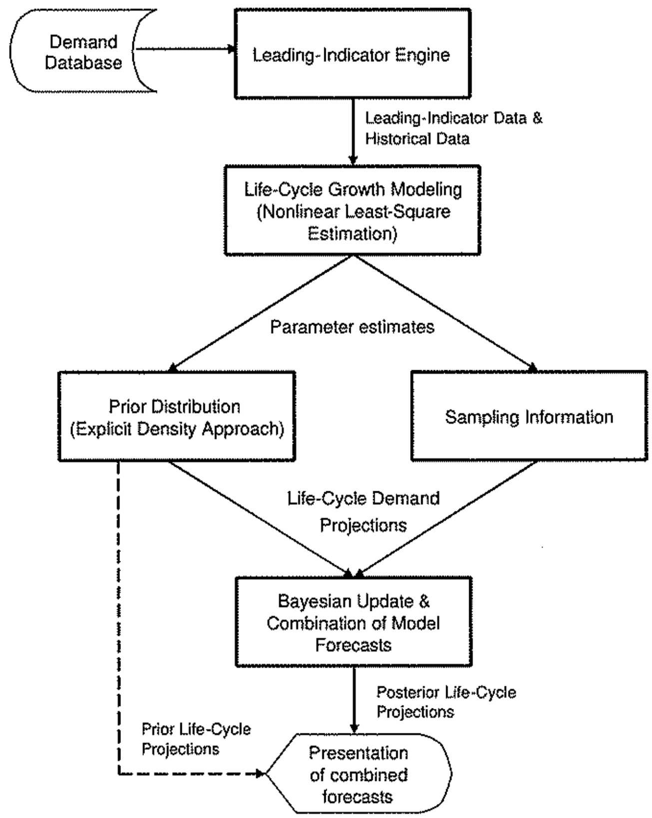

Figure 2: Flow diagram outlining the demand-characterization approach

the Bayesian and classical forecasts for the logistic model using the diffusion of compact disks in several countries. The main difference in the current study is the use of leading indicators to stretch the window of information that is available to the demand planner.

Sources of short-term demand information

In different industries, a variety of information sources are available as early demand signals. For instance, pre-season sales and initial point-of-sales data are informative in the fashion industry (Fisher and Raman, 1996; Eppen and Iyer, 1997; Kim, 2003). Similarly, advanced bookings data induced by the discounts offered for early commitments can signal the end-market demand in the retail industry (Tang *et al*, 2004).

Human judgement is another source of information. In the apparel industry, for instance, Fisher *et al* (1994) illustrate a forecasting process that gathers the opinions of a group of individuals with expertise in business, forecasting, and marketing. The aim of this study is to provide accurate demand forecasts and eliminate bias by aggregating input from different sources. Similarly, there is a forecast-improvement project that takes place in the semiconductor-manufacturing company based on the anonymity of its participants and providing incentives, with the opinion of each expert being a candidate for a short-term demand scenario in the proposed modelling framework.

Of the products following similar demand patterns, those that are introduced earlier into the market are potential additional sources of information. In Meixell and Wu (2001) and Wu *et al* (2006), identification of these products is based on a cross-correlation analysis of their time-series data. A comparable concept occurs with macroeconomic leading indicators, for which statistics that measure the similarity of the turning points in the data series (Quinn and Mawdsley, 1996) identify leading indicators.

Designing the implementation framework

This study tests the proposed demand-modelling approach using data from a major semiconductor manufacturer. The analysis includes five successive generations of product families from one of the company's main product categories. The product families are referred to as A, B, C, D, E, and they contain $34, 82, 24, 19$, and 26 products, respectively. The data is the monthly billing quantity over a five-and-half-year time period.

For each product family, forecasting performance of the Bayesian-updated projection is evaluated against that of the historical-data-based projection and a benchmark demand (BD) forecast; specifically, BD is known as the best-exercised forecast in the company and is based on the data-collection procedure described in the introduction. To test the approach, we designed a study according to the selection of estimation period, forecast-validation period, life-cycle growth models, and sources of short-term demand information as outlined below.

First, comparison is performed at different stages of the demand life cycle as the estimation period, in particular, when a product family is 6, 9, and 12 months into the market (i.e., $T = 6, 9, 12$). These time periods are chosen to represent the early, mid, and late life-cycle stages. The data shows that on the average less than 35% of the total sales is met when a product family is 9 months into its life cycle. Second, the total length of the forecast-validation period is the maximum forecast horizon available for the BD forecast, typically covering 9 to 12 months. Moreover, a number of validation periods that change in increments of three months (i.e., $M = 3, 6, 9$, and up to 12 months) are tested in order to gain insights into the differences between the short-term and long-term performances of the forecasts.

Third, seven models suggested by Meade and Islam (1998) as the best-performing models characterise demand. The models have been selected among the 29 models from the literature based on their fitting and forecasting performances across 3000 simulated and real data sets. Among these models, the Simple Logistic and Mansfield are symmetric with respect to the fraction of total sales realised when the peak sales rate is reached; the Gompertz and Floyd are non-symmetric; the Weibull, Extended Logistic, and Cumulative Log-normal are flexible models.

Last, two groups of leading indicators are considered. The first group includes the earlier products with similar demand patterns. In collaboration with the company's demand-planning team, the second group has been determined as the externally-given data from sources such as the design-win data, data from other key product categories, and the company's BD forecast. The next section elaborates on the implementation experience along these two streams of information.

Comparable products as demand leading indicators

Earlier generations of products with similar demand patterns are potential leading indicators. Their identification is based on the correlation analysis between the time-series data of the product family and that of the potential leading indicator (Wu *et al*, 2006); the regressed time-series data of a selected leading indicator extends the time-series data of the product family by two months ($L = 2$). This analysis is for the product families B, C, and D. Forecast errors are reported in terms of mean absolute percentage error ($MAPE$), that is, the average value of the absolute percentage deviation of the forecasts, $F(T+h)$, from the actual data $X(T+h)$ over the validation period ($h = 1, ..., M$):

$$MAPE = \frac{\left|\sum_{h=1}^{h=M} F(T+h) - X(T+h)\right|}{X(T+h)}.$$

Table 1 summarises the average forecasting errors of the updated life-cycle projections (i.e., posterior LC) against the benchmark demand forecasts (i.e., BD forecast) and historical-data-based projections (i.e., prior LC) over the three product families. The results show that it is possible to obtain significant improvement in the accuracy of forecasts through the Bayesian procedure, with the change in the performances being in the range of $[-0.77\%, 35.05\%]$. On the average, the most significant improvement is achieved at the early life-cycle stage when smaller amount of data is available; however, there is no notable difference in the short-term and long-term forecasting performances. In addition, as opposed to the prior projections, the posterior projections produce forecasts comparable to the BD forecast, in particular, when forecasting longer into the future.

Table 1: Accuracy of forecasts in terms of $MAPE$ (Product families B, C, D)

Life-Cycle	Forecast	Forecast Horizon			
Stage	Type	$M = 3$	$M = 6$	$M = 9$	$M = 10, 11, 12$
$T = 6$	BD forecast	11.92%	11.79%	12.14%	12.34%
	Prior LC	17.50%	20.73%	21.24%	20.60%
	Posterior LC	11.35%	9.76%	8.46%	9.03%
$T = 9$	BD forecast	1.43%	3.61%	5.68%	6.96%
	Prior LC	4.26%	6.67%	6.40%	5.52%
	Posterior LC	2.09%	3.55%	3.39%	3.38%
$T = 12$	BD forecast	2.35%	4.07%	5.43%	6.56%
	Prior LC	3.95%	3.63%	3.34%	3.60%
	Posterior LC	2.25%	2.49%	2.93%	3.02%

The variability in sampling information determines the extent of change in the prior estimates. As the results indicate, in case that a consistent set of leading indicators is used, forecast improvement occurs; otherwise, there is no significant change. This confirms the claim that the proposed approach aims to improve on the historical-data-based estimates.

Sensitivity of the estimates to total market size

The proposed procedure projects the cumulative proportion of demand that is realised over time. Hence, the raw data of the billing quantity needs to be expressed in terms of cumulative proportions. This requires an estimate of the total billing quantity (μ), which is also referred to as the total market size. The above results assume that μ is known a priori; hence, the following analysis purports to examine the sensitivity of the results to μ.

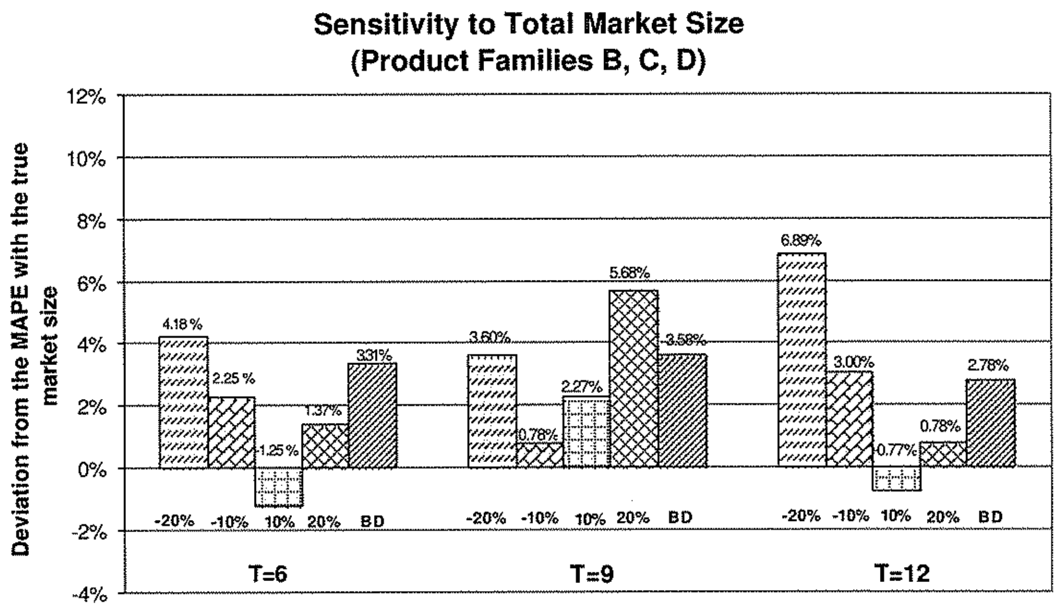

Figure 3: Sensitivity of the Bayesian estimates to total market size as compared to the benchmark (BD) forecast, the estimate of market size deviating from its true value by [-20%,+20%].

Figure 3 summarises the change in $MAPE$ of the posterior forecasts when the estimate of μ deviates from its true value by an amount within the range of $[-20\%, 20\%]$, in increments of 10%. The results are the deviations of the forecast errors from those obtained with the true value of μ over the entire validation horizon. For instance, at the early life-cycle stage, underforecasting μ by 20% results in an average forecast error that is 4.18% more than the forecast error with the true value of μ.

Across all the product families, with a 10% deviation in μ, the change in forecasting performances is less than 4%; with a 20% deviation, it is less than 10%. The change is larger as a product family progresses into its life cycle; however, a more accurate estimate of μ becomes available as more historical data accumulates, reducing the risk associated with misforecasting μ. In addition, when compared to the benchmark forecast, the posterior projections produce comparable estimates, particularly at the early and mid life-cycle stages.

Variability in demand projections

Another dimension of forecast evaluation is the measure of its variability in demand characterization. Accordingly, Table 2 presents the change in the standard deviations of the combined forecasts with the Bayesian update, where there is an improvement of at least 25% over the entire validation horizon.

Table 2: Improvement in the variability of the combined estimates with the Bayesian update

Product	Life-Cycle Stage		
Family	$T = 6$	$T = 9$	$T = 12$
B, C, D	38.10%	31.58%	44.44%
E	28.57%	55.56%	50.00%

External sources of information as demand leading indicators

The projection of a product family's design-win data, sales data from another main product category, and the BD forecast are the sources of short-term demand information. The first source, *design-win* data, occurs when a customer agrees to integrate one of the company's products at the design stage prior to the product's launch; and the customer submits an initial estimate of total order quantity in addition to the timing of market entry. The total quantity is then projected into monthly figures by the company's demand-planning team. In the present study, a regression analysis between the monthly projections of the design-win data and the actual demand data realised by the time of forecasting provides a means to debiasing the design-win data. In this section, the analysis is for the product family E since its design-win data is known to be reliable due to a recent improvement in the company's data-management system.

The second source of information is a related product within another major product category of the company. It is less costly and shipped a few months ahead by means of cheaper modes of transportation. Through historical data analysis, it has been confirmed that a debiasing procedure similar to that of the design-win data provides unbiased estimates of short-term demand information. Finally, the last source of information is the company's BD forecast, which aggregates the knowledge of a group of individuals in demand planning.

Forecasting Performance

With the inclusion of external sources of information, at the early life-cycle stage, the posterior projection performs better than the BD forecast and the prior projection by about 8% and 3%, respectively, with an improving long-term forecasting performance (Table 3). On the other hand, at the later stages, the performances of all the forecasts are comparable, with a value of $MAPE$ of less than 5%.

Table 3: Accuracy of forecasts in terms of $MAPE$ (Product family E)

Life-Cycle	Forecast	Forecast Horizon			
Stage	Type	$M = 3$	$M = 6$	$M = 9$	$M = 12$
$T = 6$	BD forecast	13.76%	16.37%	20.46%	—
	Prior LC	14.05%	12.14%	9.24%	—
	Posterior LC	11.36%	9.53%	6.96%	—
$T = 9$	BD forecast	5.07%	5.64%	4.66%	4.05%
	Prior LC	1.29%	2.44%	2.25%	1.87%
	Posterior LC	2.00%	3.19%	2.91%	2.43%
$T = 12$	BD forecast	0.60%	1.17%	1.56%	1.41%
	Prior LC	1.68%	1.08%	1.04%	0.96%
	Posterior LC	0.46%	0.84%	1.15%	1.05%

Similar to the other product families, at $T = 6$, the projections with inaccurate estimates of μ by 10% and 20% remain within the neighbourhood that deviates by less than 4% and 7% in $MAPE$ from the cumulative projections obtained with the accurate estimate of μ, respectively; the outperformance over the benchmark forecast is preserved (Figures 4). However, as the product family progresses into its life cycle, the sensitivity of the results is comparatively higher. In particular, at $T = 12$, an inaccurate estimation of μ by 20% results in an $MAPE$ that is approximately 12% higher for the cumulative projections. The most likely reason is that for this product family a larger volume of total sales occurs prior to 12 months into its life cycle, however, also reducing the possibility of a large deviation from the true market size.

Conclusion

The Bayesian modelling framework characterises demand by combining information from historical data and different sources of short-term demand information. The model updates the long-term demand projections using short-term demand information with the objective of improving

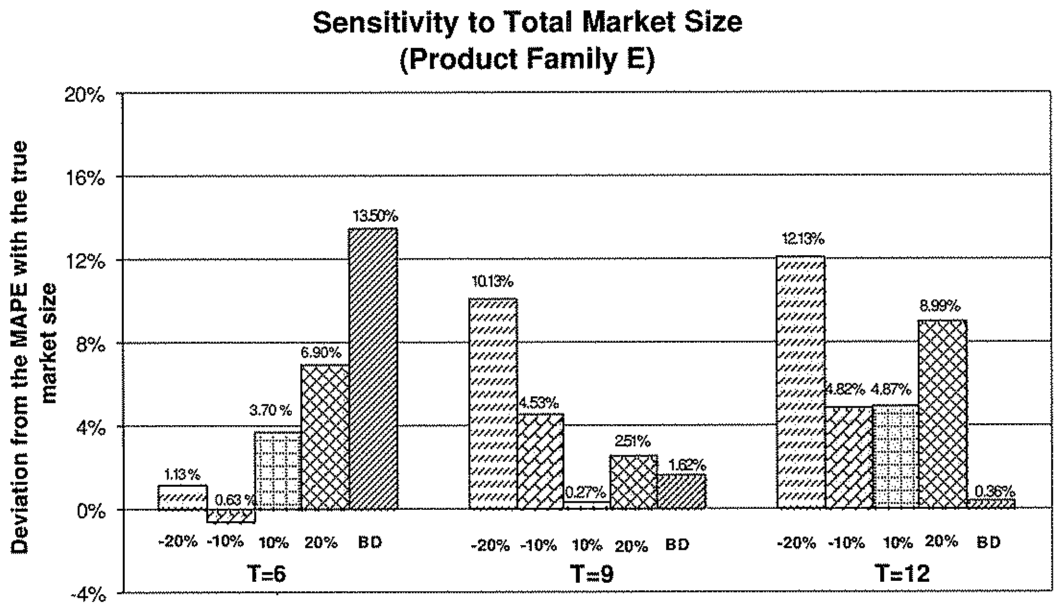

Figure 4: Sensitivity of the Bayesian estimates to total market size as compared to the benchmark (BD) forecast, the estimate of market size deviating from its true value by [-20%,+20%].

the accuracy and uncertainty in the characterization of demand. The methodology serves as an information-aggregation mechanism by allowing the integration of different information sources, including human judgement. This case study validates its applicability using real-world demand data over multiple years.

The results of the study indicate that it is possible to obtain significant improvement over the historical-data-based estimates. In addition, the outperformance of the updated forecasts over the benchmark forecasts is prominent, specifically at the early and mid life-cycle stages. On the other hand, the sensitivity of the projections to the estimate of total market size is higher at the later stages; however, a more accurate estimate becomes available as the historical data accrues. The overall results indicate that the proposed demand-characterization approach will provide an invaluable input to the forecasting processes for short life-cycle products, particularly in the high-tech industry.

This study makes a contribution by illustrating the use of a new forecasting methodology for short life-cycle products. Improved demand uncertainty, as important as forecast accuracy, has implications for operational decision-making in a capital-intensive business environment. However, when the proposed methodology is to be implemented as an integrated part of a forecasting system, its use in the first few months of product introduction will require different sources of prior information. For this purpose, a multi-generational analysis of product transitions can be performed to study whether the ramp rates of successive generations of technologies provide a beneficial source of information.

Acknowledgment

This research is supported by the Semiconductor Research Corporation (SRC) grant 2004-OJ-1223.

References

Aytac B and Wu S D (2008). Characterization of demand for short life-cycle technology products. Technical report 08T-001, Lehigh University.

Bass F M (1969). A new product growth for model consumer durables. *Mgmt Sci* **15**: 215–226.

Bewley R and Griffiths W E (2001). A forecasting comparison of classical and Bayesian methods for modelling logistic diffusion. *J Forecast* **20**: 231–247.

Dickinson J P (1973). Some statistical results in the combination of forecasts. *Opl Res Q* **24**: 253–260.

Eppen G D and Iyer A V (1997). Improved fashion buying with Bayesian updates. *Ops Res* **45**: 805–819.

Fisher R, Hammond J H, Obermeyer W R and Raman A (1994). Making supply meet demand in an uncertain world. *Harv Busin Rev* **72**: 83–93.

Fisher R and Raman A (1996). Reducing the cost of demand uncertainty through accurate response to early sales. *Ops Res* **44**: 87–99.

Kim H (2003). A Bayesian analysis on the effect of multiple supply options in a quick response environment. *Nav Res Logist* **50**: 937–952.

Kurawarwala A A and Matsuo H (1998). Product growth models for medium-term forecasting of short life cycle products. *Technol Forecast Social Change* **57**: 169–196.

Meade N and Islam T (1995). Prediction intervals for growth curve forecasts. *J Forecast* **14**: 413–430.

Meade N and Islam T (1998). Technological forecasting– model selection, model stability, and combining models. *Mgmt Sci* **44**: 1115–1130.

Meade N and Islam T (2006). Modelling and forecasting the diffusion of innovation– a 25-year review. *Int J Forecast* **22**: 519–545.

Meixell M J and Wu S D (2001). Scenario analysis of demand in a technology market using leading indicators. *IEEE Trans Semicond Mfg* **14**: 65–75.

Press S J (2003). Subjective and Objective Bayesian Statistics: Principles, Models, and Applications. J. Wiley and Sons, Inc.: New Jersey.

Quinn T and Mawdsley A (1996). Forecasting Irish inflation: A composite leading indicator. Technical report 4/RT/96, Central Bank of Ireland.

Sultan F, Farley J U and Lehmann D R (1990). A meta-analysis of diffusion models. *J Marketing Res* **27**: 70–77.

Tang C S, Rajaram K, Alptekinoğlu A and Ou J (2004). The benefits of advance booking discount programs: Models and analysis. *Mgmt Sci* **50**: 465–478.

Timmermann A (2006). Forecast Combinations. In: Elliott G, Granger C W J and Timmermann A (eds). Handbook of Economic Forecasting 2006. North Holland: Amsterdam, pp 135–196.

Winkler R L and Makridakis S (1983). The combination of forecasts. *J R Statist Soc* **146**: 150–157.

Wu S D, Aytac B, Berger R T and Armbruster C A (2006). Managing short life-cycle technology products for Agere Systems. *Interfaces* **36**: 234–247.

Zhu K and Thonemann U (2004). An adaptive forecasting algorithm and inventory policy for products with short life cycles. *Nav Res Logist* **51**: 633–653.