Extending Bass for Improved New Product Forecasting

S. David Wu
P.C. Rossin College of Engineering and Applied Science, Lehigh University, Bethlehem, PA, 18015, david.wu@lehigh.edu

Karl G. Kempf
Decision Technologies Group, Intel Corporation, Chandler, AZ, 85226, karl.g.kempf@intel.com

Mehmet O. Atan, Berrin Aytac
Department of Industrial and Systems Engineering, Lehigh University, Bethlehem, PA, 18015
mehmet@lehigh.edu,bea4@lehigh.edu

Shamin A. Shirodkar, Asima Mishra
Customer Planning and Logistics Group, Intel Corporation, Chandler, AZ, 85226
shamin.a.shirodkar@intel.com,asima.mishra@intel.com

Forecasting demand for new products is increasingly difficult as the technology treadmill drives product lifecycles shorter and shorter. The task is even more challenging for electronic goods, where product lifecycles are measured in quarters, manufacturing processes in months, while the market volatility takes place on a day-to-day basis. We present a model that perpetually reduces forecast variance as new market information is acquired over time. Our model extends Bass's original idea of product diffusion to a more comprehensive theoretical setting. We first describe how forecast variances can be reduced when combining predictive information from multiple diffusion models. We then introduce the notion of demand-leading indicators in a Bayesian framework that reduces forecast variance by incorporating a wide variety of information that emerges during the product lifecycle. We describe a successful implementation of this model at Intel, where one-third of the microprocessor products were tested. When compared to the current forecasting method, our model reduces forecasting time/effort from 3 days to 2 hours while decreasing forecast errors by 33%, which leads to $11.8M in cost savings over four months of demand realization.

In today's fast moving, ultra competitive markets, companies are introducing new products at a higher frequency with shorter and shorter product lifecycles. Electronic products such as personal computers, mobile phones, and video games are familiar examples. In these dynamic market environments, a company's competitive advantage and capability to achieve and sustain profitability come from its ability to manage frequent product entries and market exits. To hone this competency, the company must develop capabilities to utilize diverse and fast-changing market information so that their demand views sharpen perpetually, and their demand forecasts improve over time.

Upstream, in the supply chain, the challenges only intensify. A component manufacturer may need to introduce a growing variety of new products for multiple main market segments, all at a fast pace. For example, at Intel, the Microprocessor Units (MPUs) have a prominent presence at three major vertical markets (server, desktop, and mobile devices). This leads to the introduction of more distinct products with shorter lifespans, generating multiple product releases/transitions per year. To stay competitive, it is critical to release each new product to the right market(s), at the right time, with the right volume, paced over its entire lifecycle. There is a diffusion process for any new product introduced into the market; thus, appropriate timing and volume are critical to its acceptance, adoption, and ultimate financial success. As such, an in-depth understanding of the demand process over a product's lifecycle and the ability to forecast lifecycle demand diffusion are vital to a company's ability to manage product transitions and to maintain its competitive edge.

Lifecycle forecasting is critical not only for *demand* management, but also for *operations* management. After its introduction into the market, production shortage of a new product can seriously retard its adoption and negatively impact revenue. Excess inventory erodes profits and uses capacity that could have been better employed for other products. Production capacity can be extremely expensive; re-allocating the right amount of capacity at the right time (e.g., from old to new products) is critical to avoid both stock-out and inventory build-up. As there are many products sharing a common capacity across multiple market segments, the ability to quickly generate a more accurate forecast that can be mapped onto appropriate capacity will lead to significant operational savings.

Fortunately, during the lifecycle of a new product, a wealth of information can be potentially utilized for demand forecasting. The product lifecycle typically involves stages starting with pre-launch/introduction, continuing with ramp-up, maturity/saturation, ramp-down, and followed by the end-of-life when the product is being replaced by a new generation. During each stage, different sources of information become available, which can be utilized to provide projections, advanced indications, or status updates for future demands. With this massive amount of information, the challenge is how to systematically extract relevant information that will help the planners to comprehend future demands in the context of operations.

In this paper, we will address two prevailing business problems during new product introduction. These problems are:

1. The need to capture complex product diffusion patterns across diverse and multi-faceted vertical markets.
2. The need to utilize dynamically evolving market information and business intelligence during the product diffusion process.

The ultimate goal is to develop a new approach to lifecycle-demand forecasting that perpetually enhances forecast accuracy as more and different information becomes available, leading to significantly enhanced performance and consistency. The motivation of the work described here has been to generalize Bass's original idea (Bass 1969) of *product-lifecycle diffusion* to a more comprehensive theoretical setting and then to apply the theory to realize the potential benefits. Significant progress has been made toward this end. We have tested our model at a handful of technology companies (Wu et al. 2006) and successfully implemented the model at Intel, which has provided significant improvement in forecasting accuracy with corresponding savings in operational costs.

Previous Work in Technological Forecasting

Technological forecasting literature uses diffusion models to characterize product demand lifecycles. To predict patterns of future demand realization, researchers have proposed a variety of models that mainly differ in their cumulative diffusion profiles throughout the product lifecycle. Meade and Islam (1998) and Kumar and Kumar (1992) provide extended surveys of the various diffusion models used in technological forecasting. The most well-known and widely-used diffusion model was introduced by Bass (1969). When initially tested on consumer durable goods, the Bass model was shown to provide accurate predictions on both timing and magnitude of sales throughout the product lifecycle. Since then, the Bass model has been extended to incorporate additional features of technology diffusion and widely used to forecast diffusion in markets such as retail, education, pharmaceuticals, and agriculture (Mahajan et al. 2001). However, demand characteristics for technology products are different from most traditional markets due to the rapid innovation cycle that leads to significantly shorter lifecycles and higher volatility. Studies that consider technology product diffusion include Norton and Bass (1987), who built upon the Bass model to forecast successive generations of products in the semiconductor industry. Kurawarwala and Matsuo (1996) incorporated a seasonal influence parameter to the Bass model to predict demands for a personal computer manufacturer. Modis and Debecker (1988) also analyzed the demand for computer manufacturers using an S-shaped logistic diffusion curve. More recently, Wu et al. (2006) used diffusion models to generate forecasts while reducing forecast variation in a custom semiconductor manufacturing setting. Aytac and Wu (2008) are the first to introduce a demand characterization framework based on multiple diffusion models and a Bayesian updating procedure that uses advanced demand signals (leading indicators) to systematically reduce forecast variation. This latter paper provides the theoretical underpinnings of the work described here.

Demand Characterization and Forecast Analysis

Technology products typically have a single-modal demand lifecycle that goes through (once) the phases of pre-launch, ramp-up, ramp-down, and end-of-life. This lifecycle demand can be expressed as a bell-shaped time-series curve (e.g., billings over time) or as a cumulative curve, in which each point on the curve represents the percentage of lifecycle demand satisfied up to that time. Note that the cumulative curve is S-shaped; since it is in second order, some of the short-term fluctuations in the time series are ignored. Researchers (Meade and Islam 1998) have proposed various S-shaped diffusion models (including the well-known Bass model) to forecast cumulative demand during a product's lifecycle. Each of these models differs in the rate of adoption achieved and the peak diffusion rate reached, as well as the steepness of growth or decline of the diffusion rate. In other words, each diffusion model projects, in a unique way, how a particular product's lifecycle unfolds over time.

Given up-to-date information about realized demand (e.g., early sales) and an estimation of total market volume, lifecycle forecasting (1) finds a diffusion model and determines its parameter values that provide a strong goodness-of-fit, and (2) generates demand forecast by projecting the fitted diffusion model over the entire product lifecycle.

Characterizing Demand Diffusion Across Multiple Markets

A company may need to introduce a variety of new products for multiple vertical markets. Each of these verticals has its own market drivers and dynamics, which overlap and interact. Demands in different verticals are likely to demonstrate a distinctly different goodness-of-fit for particular diffusion models. Moreover, many of the products share the same capacity during the manufacturing process, thus a cohesive understanding of their diffusion in the markets is critical. The goal of our model is to capture demand characteristics across diverse and multi-faceted market environments by systematically combining the projections from multiple diffusion models.

While it has been suggested that combining multiple forecasts outperforms forecasts generated from any single model (Bates and Granger 1969), and a variety of techniques were suggested to combine the forecasts of individual models and to estimate model parameters (Mahajan and Muller 1979, Sultan et al. 1990, Timmermann 2005), it is less clear if combining diffusion models derived from different vertical markets would necessarily help in characterizing the overall, multi-faceted market demand. More importantly, when combining multiple diffusion models, does one risk introducing additional variances and biases into the forecast? Is it better to find a particular diffusion model that performs the best across all markets? In the following, we will summarize some of the key theoretical insights that form the basis for our forecasting model.

In forecasting lifecycle product demand, it is important to find the cumulative percentage of total market demand that has been observed by time $T+\tau$, denoted by $X\left(T+\tau\right)$, given that actual demand observations up to time T, $\Theta(T) = X\left(1\right), X\left(2\right), .., X\left(T\right)$, are available. Let $\hat{X}_k\left(T+\tau|\Theta(T)\right)$ denote an estimate of cumulative demand percentage observed by time $T+\tau$, projected by diffusion model k. Then,

$$X\left(T+\tau\right) = \hat{X}_k\left(T+\tau|\Theta(T)\right) + \epsilon\left(T+\tau|\Theta(T)\right)$$

where the estimate $\hat{X}_k\left(T+\tau|\Theta(T)\right)$ is provided by $F_k(T+\tau)$, which is the cumulative percentage of total demand observed by time $T+\tau$, according to diffusion model k, and $\epsilon\left(T+\tau|\Theta(T)\right)$ is the estimation error.

Suppose the forecast generated from a diffusion model represents an unbiased estimation for the actual demand and the estimation error is normally distributed with mean zero and a known, fixed variance σ_ϵ^2; then, the actual cumulative demand at $T+\tau$ can be represented by a normal random variable:

$$\tilde{X}_k\left(T+\tau\right) \sim N(\hat{X}_k\left(T+\tau|\Theta(T)\right), \sigma_k^2)$$

Note that the mean of the random variable is defined by diffusion model k ($F_k(T+\tau)$) while the variance σ_k^2 is the sum of variances of the forecast and the estimation error (σ_ϵ^2), assuming that they are independent. Although point estimates are used widely in practice, there is an uncertainty inherent in forecasts obtained by diffusion models. This uncertainty originates from the nonlinearity of model fitting and the errors in parameter estimation. The uncertainty in the estimate of a future realization of the random variable is described by a prediction interval. The prediction intervals for the forecast obtained by a diffusion model are illustrated in Figure 1.

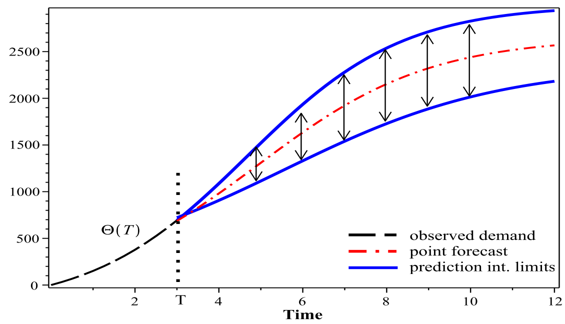

Figure 1 Prediction Intervals for Lifecycle Forecast by a Diffusion Model

We are interested to know if combining diffusion models derived from different vertical markets would help in improving the overall market forecast. Specifically, does one introduce additional variances and biases into the forecast, and how does this compare to finding a diffusion model (e.g., Bass) that performs well across all markets? Some of these questions can be answered using the setting above. Given a particular diffusion model, the actual cumulative demand at $T+\tau$, $\tilde{X}_k\left(T+\tau\right)$, can be represented as a normally distributed random variable (assuming normally distibuted fitting errors). With the combination of multiple diffusion models to forecast demands τ-period ahead, $\tilde{X}\left(T+\tau\right)$, the combined forecast is also normally distributed. Thus, as long as the combination of the diffusion models using weights inversely proportional to each model's *forecast variances* occurs, the variance of the combined forecast is always smaller than the variance of any individual diffusion model (see Proposition 1 in Appendix A). Note that the prediction intervals for the lifecycle forecast will shrink with the decrease in the forecast variance (see Figure 2).

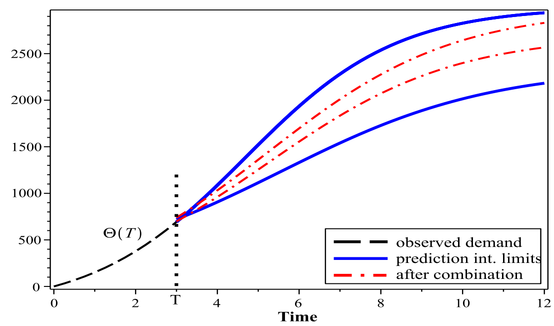

Figure 2 **Prediction Intervals after Combining Multiple Diffusion Models**

Incorporating Dynamically-Evolving Information

The intent of our model is to effectively utilize diverse and fast-changing market information to improve forecast accuracy. The goal is to perpetually reduce forecast variance as new market information is acquired over time. Given the inherent diversity and complexity of market information, we propose a unifying view that considers market information as a leading indicator for product demands. We adopt the Bayesian statistical framework, as described in (Aytac and Wu 2008), and extend it to process information provided by a wide variety of demand-leading indicators.

We define a leading indicator as a demand series, typically in the form of a time-series, that predicts the pattern of one or more new demand series before they materialize. The idea of demand-leading indicators was first proposed by Meixell and Wu (2001) and later verified and tested in an industry setting by Wu et al. (2006). Multiple leading indicators can be used at the same time or over time. We generalize the notion of leading indicators to include any information indicative of future demand patterns so long as verifiable connections can be drawn in a consistent manner. For instance, a leading indicator can be the historic demand series from an older generation product, with a sales pattern demonstrating high correlation with that of the new product; it can be derived from pre-horizon market research results or particular market or business cycle indexes that show strong connections to future demands of interest.

As shown in the first part of Figure 3, at time T one can fit a diffusion model to the observed demand data, $\Theta(T)$, and then project the adopted model over the product lifecycle, which provides a *prior distribution*, i.e., $\hat{X}_k(T+\tau)$. Now, suppose that time series data is available from a group of m leading indicators that provide unbiased estimates for the actual demand from $T+1$ through $T+L$. The data corresponding to each leading indicator is then projected over the entire product lifecycle (using a diffusion model). The collection of m leading indicator projections form the *sampling distribution*, as illustrated in the second part of Figure 3. The sampling distribution can be summarized as:

$$\tilde{X}_k(T+\tau) \sim N\left(\frac{1}{m}\sum_{j=1}^{m} \hat{X}_{kj}\left(T+\tau|\Theta^j(T+L)\right), \frac{\hat{\sigma}_k^2}{m}\right),$$

where the mean is provided by the diffusion curve k's projection, using the data extended by m leading indicators for L time periods.

Using a Bayesian-update procedure, we combine the prior and sampling distributions to generate a *posterior distribution*, which provides a distribution of lifecycle forecasts taking into consideration the new information provided by the leading indicator. An important theoretical insight provided by Aytac and Wu (2008) is that so long as each leading indicator represents an unbiased estimate of the actual demands, and this information is incorporated using the Bayesian procedure outlined above, then the variance of the lifecycle forecast will *always* decrease. This is because the variance of the posterior distribution is smaller than or equal to the variance of the prior distribution (Figure 3). Further, the variance asymptotically approaches zero as the number of leading indicators (m) increases, provided that each leading indicator is an unbiased estimate of actual data (see Proposition 2 and the proof in Appendix A).

Note that the Bayesian updates may take place multiple times throughout the unfolding of the product lifecycle; as new information becomes available, model parameters and combination weights for each leading-indicator-generated sample path are re-estimated. Moreover, this procedure can be used at any stage during the product diffusion (pre-launch, ramp-up, ramp-down, etc.). The only difference in a stage would be the choice of leading indicators, since the prediction quality of an indicator may vary at different time points throughout the planning horizon, discussed in detail in the implementation section of this paper.

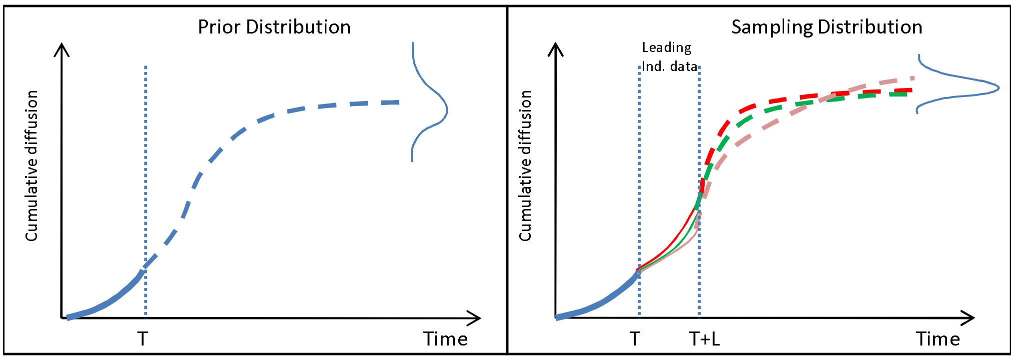

Figure 3 **Incorporating Leading Indicators in the Bayesian Update**

The Integrated Forecasting Model

We now describe an integrated model for new product forecasting using the theoretical results developed above. The model is to (1) capture demand diffusion across multiple vertical markets, and (2) incorporate dynamically evolving market information using leading indicators. Implemented using a Bayesian statistical framework, the model provides continuous improvement of forecast accuracy as verifiable new information (leading indicators) is introduced as the lifecycle unfolds. When multiple diffusion models are used, this procedure is repeated for each diffusion model, and a forecast is generated by combining them as described earlier. In Appendix B, we outline the specific algorithm in detail.

Our model generalizes the concept of Bass diffusion to broader dimensions, recognizing the richness of diffusion patterns across multiple markets and the importance of utilizing dynamically evolving market information using various indicators. The intent of the model is to perpetually reduce forecast variance, rather than generating the best possible point forecast. Note that the forecast variance is guaranteed to reduce if the leading-indicator-generated *sampling distribution* represents an unbiased estimate of the means of the actual data. In practice, it is possible to verify *ad hoc* through standard hypothesis testing if the indicators are indeed unbiased estimates. When systematic biases are identified, efforts should be made to adjust the bias for future uses. In Appendix C, we outline a simple procedure that helps to un-bias the leading indicator through learning.

Improve Forecasting at Intel

We now describe our experience of implementing the new product forecasting model at Intel. Intel Corporation is a supplier of microprocessor units (MPUs), boards, systems, and software for the computing and communications industries. Founded in 1968, Intel has emerged as the world's largest semiconductor company over the past 40 years, with 2008 revenues reaching $37.6 billion. The majority of the revenues are generated from three key market segments known as the server, desktop, and mobile markets.

To improve new product forecasting, a team involving three key groups was assembled and has been collaborating over the past few years. Researchers in supply-chain management and operations research at Lehigh University constitute the first group. The Lehigh group has been developing and integrating the notion of demand leading indicators and lifecycle diffusion models as a means to reduce forecast variations (Wu et al. 2006, Wu 2008, Aytac and Wu 2008). The second group is the Microprocessor Marketing and Business Planning (MMBP) team at Intel, who has been historically responsible for the MPU forecast. The third group, Decision Technologies (DT) at Intel, has acted as the bi-directional conduit between the Lehigh and MMBP groups. DT at Intel is chartered to identify critical business problems and supply effective decision support tools to the appropriate decision makers.

Collaboration between the three groups led to the development of a new set of decision support tools. The new tools support demand-lifecycle analysis utilizing extensive Intel business data sets. The data tested includes some 60 different Intel products with lifecycles completed by the end of year 2008; the data set spanned three product verticals, including 17 mobile, 17 desktop, and 26 server products. The tools and the extensive data set allow the team rigorous validation of theoretical insights using real-world data. The decision support software has been implemented with special attention given to the ease of use for demand planners, including a graphical interface. The system has been in use for an extended period of time (10 monthly forecasting cycles) on a large segment (about one-third) of Intel's MPU products, which makes possible the in-depth quantitative and qualitative assessment of its performance. This section documents the implementation details and provides a summary of the performance assessment. While the improvements in forecast quality (over existing methods) and the overall impact on streamlining Intel's business processes are both overwhelming and positive, evaluation of sustained performance using Intel's data-driven continuous improvement process is still ongoing.

Combining Demand Diffusion Models

Based on the analysis described earlier, the team implemented a procedure to combine a selected number of diffusion models, which collectively characterize the multifaceted aspects of Intel's three vertical markets. To begin with, the team selected 10 distinct diffusion models (i.e., Bass, Cumulative Lognormal, Extended Riccati, Simple Logistics, Extended Logistics, Gompertz, Skiadas, Mansfield, Floyd, and Weibull). The intent was to start with a wide variety of models that encompassed a good mix of differing symmetry and points of inflection characteristics in the S-Shaped diffusion curves. The 10 diffusion models were further tested to see which of them best fit Intel's historical MPU demand data. These models were fit to demand data described in the previous section. The top five models that minimized the sum of squares error (SSE) over the lifecycle of the products were selected. The five models that consistently performed well across the three verticals were the Skiadas, Extended Logistics, Bass, Weibull, and Simple Logistic.

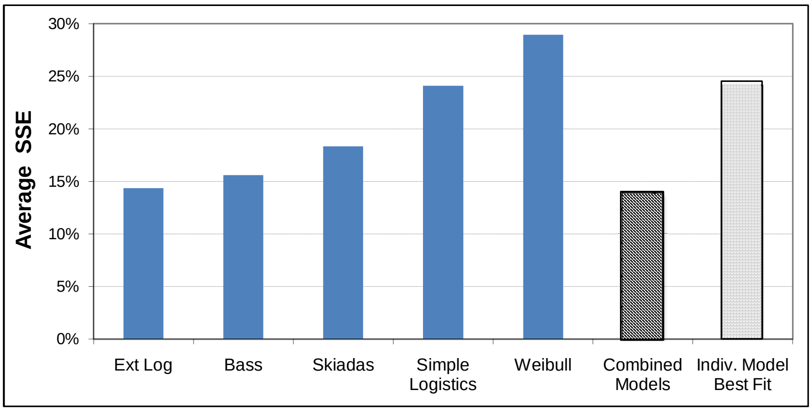

Figure 4 **Average Forecasting Error across all Mobile Families**

The average forecast error across all Mobile families (6 months into the product lifecycle) are presented in the bar chart in Figure 4. The first five bars represent the average forecast errors for the five individual models. The sixth bar shows the error for the combined forecast, while the last bar depicts the error for the diffusion model that demonstrates the best goodness-of-fit 6 months into the lifecycle. As shown in the figure, the combined forecast outperforms all individual model forecasts. However, note that the performance of an individual diffusion model is not known *a priori*. If one is to select a diffusion model that demonstrates the best goodness-of-fit 6 months into the lifecycle to project the rest of the lifecycle demand, the model presented by the last bar will be chosen. Hence, model combination not only performs better than every individual model on average, but also avoids the risk of choosing the "wrong" model given limited information.

Incorporating New Information at Different Stages of the Lifecycle

Before the introduction of a new Intel product through its end of life, a wealth of information is available from a variety of sources that can serve as demand leading indicators. To utilize different sources of information during the product lifecycle, we divide the lifecycle into pre-launch, ramp-up, ramp-down and end-of-life (Figure 5), where these stages are defined formally based on *time* (percentage of the estimated life realized) and *volume* (total percentage of estimated market size realized). A product is defined to be in the ramp-up phase from the beginning of production until 40% of the estimated lifecycle or 40% of the estimated market size for the product is realized. Similarly, a product is defined to be in end-of-life phase, when at least 90% of its estimated lifecycle or 90% of its estimated market size is realized. The rest of the active products are classified to be in ramp-down phase. For each product, the percentage of total life realized is defined as the ratio $(number\ of\ months\ into\ lifecycle/T)$ and the percentage total market size realized is given by $(volume\ realized\ till\ now/M)$, where T is the estimated total length of the product life-cycle and M is the estimated total volume of sales (market size) over the product lifecycle. For each product, T is a forecast that is published by the Long Range Planning group at Intel and M is calculated as the sum of past demand realized and the estimated future sales that is obtained from forecasts by the MMBP and Long-Range-Planning groups at Intel. Since at least 3 data points are necessary to fit the lifecycle diffusion models, these models are used once a product is 3 months into its lifecycle.

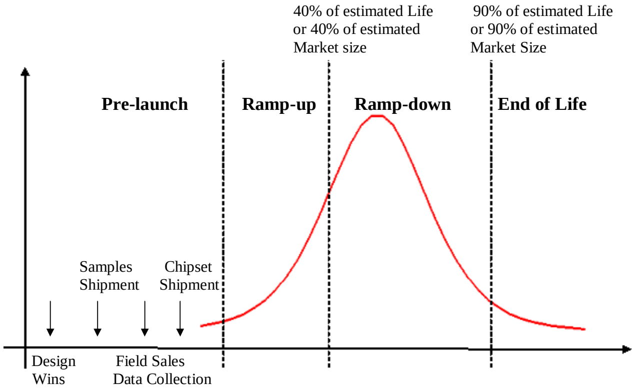

Figure 5 **Intel's Timing of Leading-Indicator-Data Collection and Different Phases of a Product Lifecycle**

Depending on the stage of the lifecycle of the product, different forms of leading indicators are used to generate the forecast. Leading indicators serve a dual role. They add business intelligence to the diffusion models by capturing changing dynamics in the market as early as possible. They also help by extending the demand data set needed to fit the models, which is very valuable early in the product lifecycle. In the following, specific leading indicators implemented at Intel are described.

Design Wins. Design Wins are early leading indicators that are available from the pre-launch stage. Designs are declared as Wins when a customer takes pre-production MPU samples and conducts preliminary evaluations. The customer then produces preliminary designs for the printed circuit board for a new end-product. The customer supplies an initial forecast of the expected order quantity in addition to the timing of the entry into the market. Figure 6 (left chart) shows the Design Wins for the same product collected prior to its launch. In this case, the correlation between actual demand and the Design Wins leading indicator is 0.81. This is quite typical for MPU products.

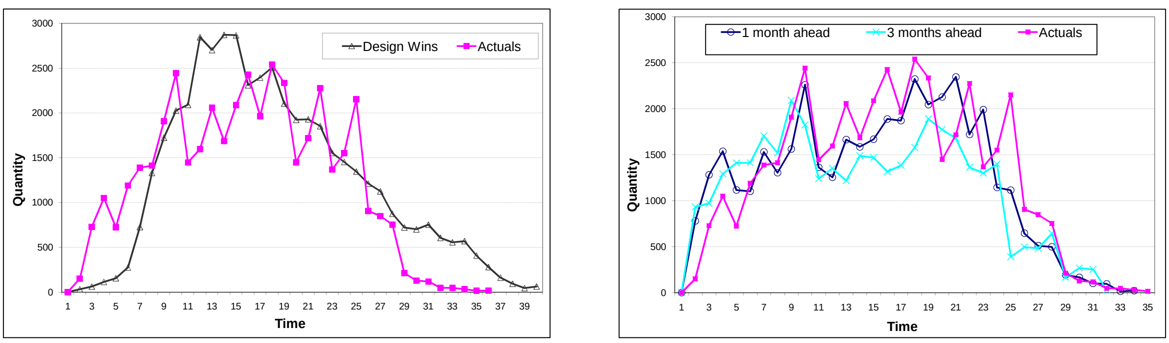

Figure 6 **Design Win (Left) and Forecasts from Field Sales (Right) as Leading Indicators for an Intel Product**

Field Sales Intelligence. Field sales personnel who are in daily contact with major customers often develop dynamic forecasts at a customer and product granularity level. This type of field sales intelligence is available from pre-launch, and continues throughout ramp-up and ramp-down, which has proven to be a good leading indicator. Figure 6 (right chart) shows examples of “x month ahead” forecasts from field sales for an Intel product in comparison to the actual demand realized for the product. Each point in the “x month ahead” curve was collected x months ahead of the realization of the actual demand. For example, the forecast on the “one month ahead curve” corresponding to time period 10 was forecasted in time period 9. Notice that, as expected, the quality of the leading indicator improves as one moves closer to the actual realization of demand, also reflected in the correlation between the actual demand and the leading indicators.

Chipset Sales. Chipsets are microprocessor chips that cover functionalities ranging from input/output to memory to graphics and are typically paired with the MPUs in the customers' end products. Chipsets are historically much less expensive than MPUs and are shipped weeks ahead of the corresponding MPUs to the customer to facilitate circuit board testing. As such, the Chipset “ship-aheads” serves as strong leading indicators for the MPUs.

As cautioned above, each type of leading indicator could have some degree of bias due to a variety of reasons. The bias in Design Wins is primarily because the information is collected much before product launch and the customers may not have visibility very far into the future. Compounding this is the situation where the customer's product design group (who typically provide the Design Win estimates) may be psychologically inclined to believe its next product will be a winner; this translates into positive biases. There could be bias in the field sales intelligence in that the gaming behavior of some customers could exaggerate field forecasted quantity. Hence, the quality of these leading indicators can be further improved by correcting for bias. The theoretical results established earlier asserted that if the projection made by the combined diffusion model and/or the leading indicator represent an *unbiased estimate* for the mean of the actual demand, then the forecast variance will reduce. We have implemented the un-biasing, or “learning” mechanism described in Appendix C, which allows the forecasting performance to improve over time.

Business Results and Conclusions

The Microprocessor Marketing and Business Planning (MMBP) team at Intel is a strategic group responsible for supply-demand matching and pricing for all MPU products. One of the main tasks of the group is to forecast customer demand for MPU products in the desktop, mobile, and server markets. Each month MMBP generates a new 12-month demand forecast for each active product relying on historical data systems, collective mental models, and the latest news from the market. The forecast is communicated first to the senior management team for final approval and then to supply chain operations for execution. The monthly forecast has at least three operational uses: Given Intel's 3-month production cycle, the first few months of the forecast serve to re-prioritize work in progress, re-position existing inventory, and finalize logistics arrangements. The fourth month of the forecast is the most critical since it is used to release raw materials into the fabrication facilities at the beginning of the production process, and to trigger the placement of orders for materials used in assembly factories. The remaining months of the forecast are used by the operational team as an input for production, materials, inventory, and logistics planning activities.

The Integrated Forecast Model has been incorporated into the forecasting process over 10 forecasting cycles, beginning with a trial run in December 2008. It has been used to generate forecasts for 10 desktop MPU products that form 32% of Intel's active MPU products. The team has been tracking the performance of the new forecast model as the actual demand volumes are realized every month. Since there is a 3-month production cycle for forecasts generated (and used to release materials into fabrication) in January, February, March, and April, we collect actual shipment data in April, May, June, and July, respectively. This provides matching pairs of time-series data that allows us to compare the forecast and actual demands in detail. Since there is a one-month accounting cycle at Intel, at the conclusion of this analysis in September 2009, we are able to report four months of demand realization.

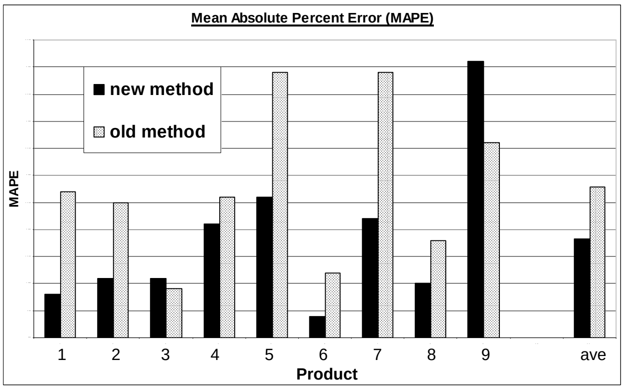

Figure 7 **Comparison of the Mean Absolute Percentage Error of the Forecasting Methods for 9 Products over 4 Demand Cycles**

Figure 7 summarizes the overall performance of the Integrated Forecast Model (new method) compared to the original MMBP methodology (old method) for 9 of the 10 desktop MPUs. On average the new method shows a 9.7% reduction in the 12-month forecast error (measured in Mean Absolute Percent Error, or MAPE) per product; however, if we focus our comparison on the critical fourth month forecast, the new method shows a 33% improvement in forecast accuracy. One of the 10 desktop products has been excluded due to the multi-modal nature of its lifecycle: rising to maximum volume in the first 6 months, falling to 60% in the next 6 months, rising back to 95% in the following half year, and then sitting at 35% for the following year. The single-modal diffusion models are not appropriate for such cyclic demand patterns. We are studying models that will allow us to extend the Integrated Forecast Model to handle products with multi-modal life-cycles. Products 3 and 9 are also under study to determine what information was included in the MMBP forecast that can used to improve the Integrated Forecast Model.

Impact on Operational Costs

To estimate the financial impact of the Integrated Forecast Model, the fourth month forecasts for the old method and the new method were compared to actual shipments realized for April, May, June, and July of 2009. The analysis proceeded through the following steps:

- It was assumed that there was no inventory on hand at the beginning of April.
- For each forecast, if production in a given period resulted in *more* product than actually shipped in that period, then (1) the excess product was held for use in future periods when production was insufficient to meet actual demand, and (2) an *inventory holding cost* was charged for each unit for each time period held.

• For each forecast, if production in a given period supplied *less* product than actually shipped in that period, and there was not enough inventory to cover the shortfall, an *underage cost* was charged for each unit for each time period that the shortage would persist.

The inventory holding costs and the underage cost were both supplied by MMBP based on financial data. Figure 8 summarizes the financial performance of the Integrated Forecast Model (new method) compared to the original MMBP methodology (old method). On average the new method shows over \$1.3M revenue enhancement per product, translating to an \$11.8M gain for the 9 products over the 4-month analysis period. In a broader sense, the *underage costs* should reflect not only potential loss in revenue but also increase in costs due to supply-demand mismatch. It has been shown that when Intel misses on-time shipment, some 7.5% of the time the customer turns to the competitor or the open market. As such, it may be necessary for Intel to shift (more expensive) capacity from other products to increase production, so as to avoid losing market share and other long-term consequences. This could cause significant and compounded increases in operating cost. Our current financial analysis does not take into account the compounded effect of under-forecasting, thus the above cost-saving is likely to be understated.

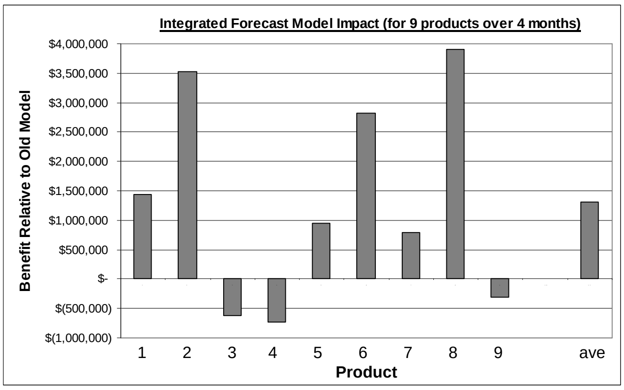

Figure 8 **Financial Comparison of the Forecasting Methods for 9 Products over 4 Demand Cycles**

Comparison of Figures 7 and 8 shows that the forecast errors and financial metrics capture related but different assessments of the performance of the methods. Considering the new method, Products 1, 5, and 7 showed the largest improvement in forecast accuracy (Figure 7), but Products 2, 6, and 8 showed the largest financial improvement (Figure 8). The new method performed worst for Product 9 by the MAPE metric, but the impact on the financial metric was minimal. Part of the reason is that each product has different inventory holding and underage costs (e.g., Product 9 has among the lowest inventory holding and underage costs). Moreover, the financial metric captures dynamics of the system that the MAPE metric does not. For instance, over-forecasting in an earlier period can build inventory (with a relatively low cost penalty) that covers under-forecasting in a later period (with a relatively high cost penalty if not covered). Conversely, under-forecasting followed by over-forecasting results in much more severe financial consequences.

In addition to improvements in forecast accuracy and revenue, the new forecast tools provide an opportunity to decrease the time and effort required to generate the forecast. The old process takes roughly 3 days, while the new tools can produce an initial base forecast in 2 hours and facilitates the evaluation of a number of additional business scenarios with small additional investments of time and energy. The integrated approach also helps to standardize the forecasting methodology and make the forecasting process both systematic and repeatable compared to the old methodology. Considering the high attrition rate seen in this profession, this is especially useful during knowledge transfer from forecaster to forecaster.

Diverse plans have been set for continuously improving forecasting of product transitions. Realizing and measuring the theoretical predictions for manufacturing cost saving is also of great interest. Expansion and refinement of the theory as well as the use of the New Product Forecast Model across broader sets of Intel products will continue to occur to generate improved forecasts and to streamline the overall business processes.

References

Aytac, B., S. D. Wu. 2008. Characterization of demand for short-lifecycle technology products. *Technical Report 08T-001* ISE Dept. Lehigh University, Bethlehem, PA.

Bass, F. M. 1969. A new product growth model for consumer durables. *Management Sci.* **15**(5) 215–227.

Bates, J. M., W. J. Granger. 1969. The combination of forecasts. *Oper. Res. Quart.* **20**(4) 451–468.

Kumar, U., V. Kumar. 1992. Technological innovation diffusion: The proliferation of substitution models and easing the user's dilemma. *IEEE T. Eng. Manage.* **39**(2) 158–168.

Kurawarwala, A. A., H. Matsuo. 1996. Forecasting and Inventory Management of Short Life-cycle Products. *Oper. Res.* **44**(1) 131–150.

Mahajan, V., E. Muller. 1979. Innovation diffusion and new product growth models in marketing. *J. Marketing* **43** 55–68.

Mahajan, V., E. Muller, F. M. Bass. 2001. New-product diffusion models: A review and directions for research. *J. Marketing* **54** 1–26.

Meade, N., T. Islam. 1998. Technological forecasting – Model selection, model stability, and combining models. *Management Sci.* **44**(8) 1115–1130.

Meixell, M.J., S. D. Wu. 2001. Scenario analysis of demand in a technlogy market using leading indicators. *IEEE T. Semiconduct. M.* **14** 65–75.

Modis, T., A. Debecker. 1988. Innovation in the computer industry. *Technol. Forecast Soc.* **33** 267–278.

Norton, J. A., F. M. Bass. 1987. A Diffusion Theory Model of Adoption and Substitution for Successive Generations of High-Technology Products. *Management Sci.* **33**(9) 1069–1086.

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

Timmermann, A. 2005. *Handbook of Economic Forecasting*, chap. Forecast Combinations. North Holland.

Wu, S. D. 2008. Prediction method and system. *U.S. Patent Application US-2008-0147486-A1* .

Wu, S. D., B. Aytac, R. T. Berger, C. A. Armbruster. 2006. Managing Short-Lifecycle Technology Products for Agere Systems. *Interfaces* **36**(3) 234–247.

Appendix A:

Proposition 1. *Combining forecasts of different diffusion models by using weights that are inversely proportional to their forecast variances yields a combined forecast variance that is smaller than forecast variance of each individual diffusion model.*

Proof. Let K be the set of different diffusion models that are used in forecasting ($k \in K$). Since the combined forecast is a linear combination of independent normal random variables $\tilde{X}_k\,(T+\tau)$, it is also normally distributed with mean $\sum_{k\in K} w_k \cdot \hat{X}_k\,(T+\tau|\Theta(T))$ and variance $\sum_{k\in K} w_k^2 \cdot \sigma_k^2$, where w_k is the weight assigned to model k's forecast by the combination method. Note that combined forecast's variance is:

$$\sigma_c^2 = \sum_{k\in K}\left(\frac{1/\sigma_k^2}{\sum_{i\in K} 1/\sigma_i^2}\right)^2 \cdot \sigma_k^2 = \sum_{k\in K}\frac{1}{\sigma_k^2\cdot\left(\sum_{i\in K} 1/\sigma_i^2\right)^2} = \frac{1}{\sum_{i\in K} 1/\sigma_i^2} < \sigma_k^2, \forall k \in K$$

□

Proposition 2. *The variance of the (estimates for) posterior distribution is smaller than or equal to the variance of the (estimates for) prior distribution. Further, the variance (of the estimates) asymptotically approaches zero as the number of leading indicators (m) increases, provided each leading indicator is an unbiased estimate of actual data.*

Proof. The *prior distribution* for $X_k(T+\tau)$ is $N(\hat{X}_k(T+\tau|\Theta(T)), \sigma_k^2)$; the *sampling distribution* is obtained from leading indicator based projections, which can be viewed as observations that are independently and identically distributed with respect to $N(\hat{X}_k(T+\tau|\Theta^j(T+L)), \hat{\sigma}_k^2)$. If we simplify $\hat{X}_{kj}(T+\tau|\Theta^j(T+L))$ by $\hat{X}_{kj}$, and $X_k(T+\tau)$ by X_k, according to the Bayes' theorem, the posterior probability density function for $X_k(T+\tau)$ can be obtained from the following formula:

$$p\left(X_k|\hat{X}_{k1},...,\hat{X}_{km}\right) = \frac{p\left(\hat{X}_{k1},...,\hat{X}_{km}|X_k\right)p\left(X_k\right)}{\int p\left(\hat{X}_{k1},...,\hat{X}_{km}|X_k\right)p\left(X_k\right)dX_k}$$

After the probability density functions of *prior* and *sampling distributions* are substituted in the above formula, the variance of the *posterior distribution* is found as:

$$\tilde{\sigma}_k^2 = \frac{\sigma_k^2\hat{\sigma}_k^2}{m\sigma_k^2 + \hat{\sigma}_k^2}$$

□

(More detailed proofs are given in Aytac and Wu (2008)).

Appendix B:

algorithm: *Integrated New Product Forecasting*;
Input at time T:

- Actual demand observations $\Theta(T) = \{X(1),..,X(T)\}$.
- Leading indicators $l_1,..,l_m$.
- Diffusion models $k \in K$

begin
For each diffusion model $k \in K$, **do**:
begin $\{|K|$ *passes*$\}$
  Estimate parameters for diffusion model k by fitting demand observations $\Theta(T)$.
  Project a demand series from $(T+1)$ to $(T+\tau)$ using parameters fitted for model k;
  add the demand series to the *prior distribution*.
  For each leading indicator $l_i, i \in \{1,..,m\}$, **do**:
  begin $\{|m|$ *passes*$\}$
    Use leading indicator l_i to extend $\Theta(T)$ by L_i periods;
    Estimate parameters for diffusion model k by fitting data $\Theta(T+L_i)$.
    Project a demand series from $(T+L_i)$ to $(T+\tau)$ using parameters
    fitted for model k; add the demand series to the *sampling distribution*.
  end; $\{|m|$ *passes*$\}$
  Perform Bayesian updates using the *prior distribution* and the *sampling distribution*
  from above to obtain the *posterior distribution* for model k.
end; $\{|K|$ *passes*$\}$
Combine the $|K|$ *posterior distributions* to obtain final forecast.
end;

Appendix C:

procedure: *Un-bias Leading Indicator*;

1. Regress the time series of past actual demand onto the time series of past leading indicator to obtain the extent of bias in the leading indicator data. At the beginning of time $t+1$, let $X_1, X_2, ..., X_t$ be the time series of the demand data realized and let $I_1, I_2, ..., I_t, ..., I_{t+k}, ..., I_{t+L}$ be the time series of leading indicators observed at $t+L$, where L is the time lag between the realization of product demand and the collection of leading indicators for each point t. Linearly regress the t actual demand onto the t leading indicators to get the following relationship: $X_t = \alpha + \beta \cdot I_t$.

2. The relationship established in Step 1 is then used to "un-bias" the leading indicator data for future time periods. Let $I^{'}_{t+k}$ be the unbiased leading indicator for future period $t+k$, then $I^{'}_{t+k} = \alpha + \beta \cdot I_{t+k}$.