

Managing Short-Lifecycle Technology Products for Agere Systems

S. David Wu

Department of Industrial and Systems Engineering, Lehigh University, Bethlehem, Pennsylvania
email: david.wu@lehigh.edu

Berrin Aytac

Department of Industrial and Systems Engineering, Lehigh University, Bethlehem, Pennsylvania
email: bea4@lehigh.edu

Rosemary T. Berger

Department of Industrial and Systems Engineering, Lehigh University, Bethlehem, Pennsylvania
email: rosemary.berger@lehigh.edu

Chris A. Armbruster

Agere Systems, Allentown, Pennsylvania
email: carmbruster@agere.com

Over the past decade, the high-tech industry has been experiencing an unprecedented acceleration of technology innovations and rapid product introduction cycles. At the same time, the industry has been moving rapidly away from regional vertical integration toward a supply chain structure that is fundamentally global. In this paper, we describe a critical operational problem that arises in this environment. The research stems from a project started in the fall of 2002 partnering Agere Systems and the Center for Value Chain Research at Lehigh University. The main focus of the project is to study means of characterizing demands for short-lifecycle technology products. Agere is particularly interested in using demand characterization tools for capacity planning and capacity negotiation with their global supply partners. The shortened technological lifecycle poses major challenges to the time series forecasting methods fundamental to all commercial demand planning systems. As part of this project, we propose a new demand characterization method based on the notion of demand “leading indicators.” Given a cluster of products, our method identifies one or more products that provide advanced indication of demand behaviors for the rest of the cluster. Based on extensive analysis of a data set from Agere that covers twenty-six months from December 2001 to January 2004 and includes some 3,500 semiconductor products, we have discovered that we can consistently find leading indicators that predict the cluster demand pattern two to eight months ahead of time with correlation values ranging from 0.51 to 0.95. These findings have significant implications to capacity management in an increasingly global high-tech supply chain. We discuss opportunities for applying the leading indicator approach to various planning functions.

Key words: Short-Lifecycle Technology Products; Supply-Demand Planning; Semiconductor Manufacturing;

MSC2000 Subject Classification:

OR/MS subject classification:

1. An Overview of Short-Lifecycle Technology Markets In the mid to late 1990’s, high-tech industries such as consumer electronics, telecommunications equipment, and semiconductors were experiencing unprecedented growth and expansion. During that time, many firms developed and deployed supply chain management systems to integrate and optimize their operations. With goals of reducing costs and cycle times, companies focused on internal integration but continued to rely on a traditional model of demand planning, in which marketing adjusts customer projections to produce a unit forecast against which operations executes. Against the backdrop of rapid demand growth fueled by the dot com build out, planning to customer-driven marketing forecasts was adequate, because companies were more concerned with keeping pace with demand and ensuring availability of products than with the accuracy of the data being provided by customers. However, this approach to planning prohibited many companies from reacting more quickly to the industry decline when it began in 2001. With the decline initially predicted to be short-lived, many customers were reluctant, or slow, to revise their forecasts, and many suppliers were left in the difficult position of trying to reconcile optimistic forecasts with increasingly negative economic indicators. By the time the industry acknowledged the depth and potential duration of the decline, many companies were left to assume financial responsibility for large buffers of inventory and underutilized capital equipment, further depleting already limited cash reserves.

Even in good economic times, the demand for high-tech products is volatile and challenging to manage; the rapid rate of innovation causes short product lifecycles, while the long production lead time hampers a firm’s ability to respond. Uncertain economic times, however, increase the challenge significantly.

Whereas in an environment of sustained demand growth, supply chain partners might build inventory or hold excess capacity to buffer against demand variability, many are reluctant, or unable, to assume such financial risks in a downside market. Firms recognize, however, that they must provide both innovative products and exceptional service in order to retain their customer base and to gain new revenue opportunities. To do so, they must structure their supply chains to respond to upside demand and to absorb downside risks without creating excessive inventory or capacity. It is for this reason that the high-tech industry has gone through a profound transformation during the past decade, starting with growth and expansion in the mid 1990's and continuing through contraction in the early 2000's.

As part of this transformation, major corporations are moving aggressively away from vertical integration; instead of owning and operating the entire process of product realization, they are focusing on those aspects where they hold the strongest value proposition. Contract manufacturing emerges as a component of this supply chain restructuring. By consolidating demands across multiple customers and investing in and developing highly flexible processes, contract manufacturers are able to achieve high utilization on their equipment, thereby reducing unit costs. In addition, by consolidating the component procurement for multiple customers, contract manufacturers are able to take advantage of economies of scale from their suppliers. Thus, contract manufacturers can offer their customers a greater variety of products at significantly lower cost, while offering the component suppliers greater stability in their demands. Contract manufacturing has grown from a few billion dollar industry in the early 1990's to over \$180 billion in 2001 (Kador, 2001). This rate is expected to accelerate in the next few years with the share of manufacturing done on a contract basis expected to be well over 50% (Gartner, 2003). In the semiconductor industry, contract manufacturing is no longer limited to a few fabless companies. Even fully integrated firms are beginning to contract their back-end (packaging) operations to assembly and test facilities or contract their front-end (wafer fabrication) operations through partnerships with major foundries.

The shortening product lifecycle and the emergence of contract manufacturing reflect broader trends in the global economy toward rapid product innovation cycles and increasingly complex manufacturing and supply chain partnerships. High-tech contract manufacturers, in particular, have a significant presence in the Asia-Pacific region. They represent a dominating force and a significant economic driver for China, Taiwan, Korea, and Malaysia. In addition, major ports such as Hong Kong and Singapore have become logistics consolidation points for many of these operations as well as a major sources of investment capital.

This paper addresses a critical operational dimension of the supply chain transformation described above. Specifically, we study a new approach for characterizing the demand of short-lifecycle technology products with the specific purpose of facilitating capacity management in a contract manufacturing environment. The research is part of a project started in September 2002 that partners Agere Systems and the Center for Value Chain Research at Lehigh University. Specifically, Agere seeks recommendations for a comprehensive decision tool that would characterize short-lifecycle product demands in the context of supply-demand planning. Agere is particularly interested in using the demand characterization tools for capacity planning and capacity negotiation with their global supply partners.

2. Agere's Products and Business Environment Originally the Microelectronics Division of Lucent Technologies, Agere Systems was spun off from Lucent in March 2001, in what was the fifth largest IPO in Wall Street history at the time. Agere specializes in providing semiconductor products for wireless data, high-density storage, and multiservice networking markets. Its wireless data products include GPRS chips that offer Internet connectivity for cellular phones and wireless voice over IP. In the high-density disk drive market, the company is the leading provider of chips for hard disk drives, including read channel, preamplifiers, and "systems-on-a chip" that integrate several functions into a single device. Agere remains a strong presence in the telecommunications infrastructure market; the company provides custom and standard integrated circuits (ICs) for multiservice networking equipment that moves information across wired, wireless, and enterprise networks.

Agere's business is organized by key product lines into four business units: Enterprise and Networking, Mobility, Storage, and Telecommunications. Each business unit is subdivided into business entities based on key product technologies. Each business entity serves different technology markets, each of which tends to follow a particular pattern and rate of technological evolution, i.e., a technology improves in successive generations, each more advanced than its predecessor. For instance, the technology lifecycle

for custom IC's on cellular phones is quite different from that for the motor controller IC's for hard drives. New technologies are accepted more readily in the former market and replaced more frequently, driving significantly shorter product lifecycles.

As is typical in semiconductor manufacturing, Agere's operations consist of two main stages. In the "front-end" operation, silicon wafers are fabricated in clean room facilities (fabs), and in the "back-end" operation, wafers are cut, packaged into IC chips, and tested. The front-end operation typically is the bottleneck since the process involves a manufacturing lead time of six to twelve weeks, while the back-end operations requires two to four days. The wafer fabs are capital intensive and require significant lead time to build; a typical fab costs \$1 billion to \$4 billion and takes 12 to 18 months to build. The high capacity cost and long production lead time, compounded by short product lifecycles (some are less than 18 months), push many semiconductor manufacturers to go "fabless." While Agere retains a portion of its fab capabilities in house, a substantial portion of the front-end operation is now handled by foundry partners such as Chartered Semiconductor and TSMC. The back-end operations typically are performed at Agere's facilities in Asia.

Agere is part of a high-tech supply chain that is truly global and increasingly complex. Agere's key customers include large PC manufacturers (e.g., Apple Computer), wireless handset providers (e.g., Samsung Electronics), network equipment suppliers (e.g., Lucent, Cisco Systems), and high-density storage device manufacturers (e.g., Maxtor). During the past decade, many of these firms have moved from vertical integrated operations to horizontally integrated operations that involve multiple contract manufacturers. An immediate consequence of the restructuring is that instead of receiving a single demand feed from each customer, Agere now receives a demand feed from each of the customer's manufacturing facilities and each of their contract manufacturers. The multiple feeds arise because each customer splits its demands across multiple subcontractors, who in turn pull demand separately from the component suppliers (Armbruster, 2002). The multiple demand feeds for the supplier lead to multiple inventory buffer locations and more complex demand characteristics. Note that in a separate research project, we have studied with a key customer inventory strategies to consolidate multiple demand feeds, to increase overall inventory turns and to increase transparency. Although contract manufacturing is not the focus of this paper, we reiterate an earlier point that the fast changing environment due to the increasing use of contract manufacturing has a significant impact to supply-demand planning in the high-tech supply chain.

Key motivations for this research are the increasingly complex high-tech demand structure due to the compressed technology evolution cycles and the emergence of contract manufacturing. One operations manager at Agere summarizes his perspectives as follows:

The planning and coordination environment for our industry is extremely complex and difficult to manage due to the exceptionally volatile nature of product demands and the complex manufacturing processes. In addition, the semiconductor supply chain has to constantly battle with short product lifecycles and capital intensive capacity that requires long lead time for expansion. Our objective is to more accurately anticipate our market conditions, better estimate production capacity needs in order to procure the appropriate manufacturing capability, and make optimum use of our capital assets. Improving how we plan and make decisions in this industry will also contribute to the success of our customers. For example, achieving early production ramp for custom logic IC for our customers in the multimedia computing or communications equipment markets will help them to strengthen their competitive advantage...

3. Exploring the Research Questions The high-tech manufacturing environment is primarily driven by time-based competition, where a manufacturer's ability to provide responsive and flexible supply to a customer defines its competitive advantage. To this end, our project focuses on demand characterization tools that will allow Agere to handle demand signals proactively such that capacity can be aligned for the right time at the right level. This is known in the industry as *supply-demand planning*. The project addresses the following specific research questions:

1. What is the extent of high-tech demand volatility experienced by planners and decision makers? How well can the standard time series forecasting tools handle this level of demand volatility?
2. Are there discernable patterns that can be derived from historical or current demand data that

would enhance our understanding of high-tech demands? Is it possible to identify “leading indicators” that provide advanced warning of demand changes? Are there effective ways to identify and monitor these leading indicators?

3. If leading indicators do exist, are they capable of producing reliable demand forecasts? Is it possible to develop general-purpose analysis tools building on the concept of leading indicators? Is it possible to test whether a particular product of interest is a strong leading indicator for some set of products?

We address these research questions in the remainder of the paper. In Section 3.1, we summarize the findings of a comprehensive study of Agere’s demand data covering a fourteen-month time period during the early 2000’s. In Section 3.2, we evaluate the potential effectiveness of the time series forecasting methods that are widely available in commercial software systems. In Section 4, we describe an innovative demand analysis process, which we developed into a “leading indicator engine” that can be used to identify a subset of products that predicts the demand trends for a larger product group. In Section 5, we illustrate the use of the leading indicator analysis on data provided to us by Agere. In Section 6, we discuss the implications of the leading indicator approach to capacity planning and capacity negotiation. In Section 7, we conclude with some future directions.

3.1 Understanding High-Tech Demand Volatility Agere and other high-tech manufacturing firms are cognizant that the compressed technology lifecycle and the increasingly complex supply structure (due to contract manufacturing) have stretched existing supply-demand planning systems to the limit. To understand the extent of this phenomenon as experienced by planners and decision makers alike, we initiated our study with a close examination of the demand information available to the decision makers via a sophisticated order management system. The idea is to understand demand volatility as experienced by the decision makers by reconstructing from historical data the demand views as they unfold over time. The analysis was conducted using a representative sample of Agere’s products in telecommunications, personal computing, and storage over a fourteen-month period between 2001 and 2002. For each product, we reconstructed the progression of customer-placed orders (known as the order board or backlog data) over time that culminated in an actual shipment. Weekly views of the backlog and shipment data were used in the analysis. Specifically, for a particular product in a given shipment week, the time series of order board data for the week leading to the final shipment was studied. This series represented the weekly updated views of an anticipated shipment in a particular week. To demonstrate demand volatility, we computed the mean value of this stream of order board data for each product and compared the value against the actual shipment quantity. The histogram in Figure 1 summarizes the results of this analysis by

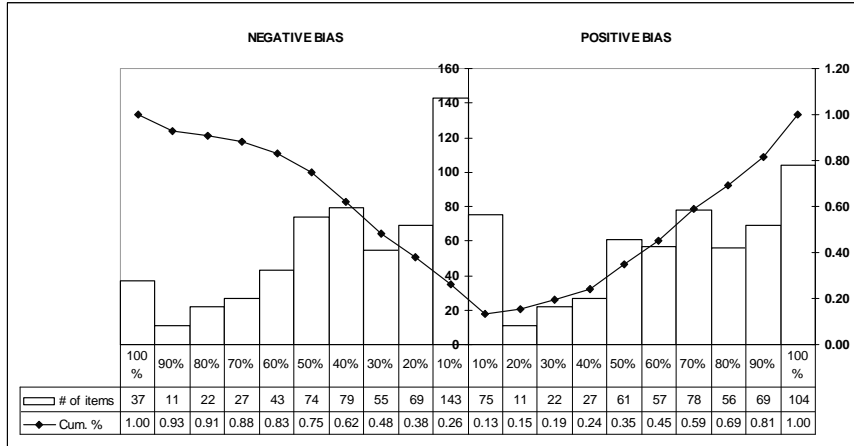


Figure 1: Deviation (%) between Order Board and Actual Shipment(2001-2002).

tracking the order and shipment quantity of 560 products over a 14-month period. The histogram shows the percentage of deviation between the order board mean and the actual shipment for each product and the frequency of occurrence of that value. The left half of the histogram represents the case of a “negative” deviation where the actual shipment is *lower* than the mean backlog quantity. The right half of the histogram represents “positive” deviation where the actual shipment is *higher* than the mean backlog

Forecasting Methods	PC 1	PC 1	PC 2	PC 2	TC 1	TC 1	TC 2	TC2
	EP	VP	EP	VP	EP	VP	EP	VP
Random Walk	55	7	281	100	179	173	37	154
Linear Trend	58	17	153	47	226	328	27	50
Quadratic Trend	40	120	152	39	229	324	27	35
Exponential Trend	52	60	109	48	99	323	25	84
Moving Avg (3 mon.)	75	24	224	49	154	372	65	134
Exponential Smoothing	56	11	163	56	450	140	34	154
Brown's Linear	56	12	166	59	488	137	33	120
Holt's Linear	59	10	197	65	254	326	30	38
Brown's Quadratic	56	12	161	59	493	136	31	105

Table 1: Representative Forecast Errors as Measured by MAPE (in %)

quantity. As many demand managers might expect, both the percentage of deviation and the number of occurrences are alarmingly high; the results suggest a highly volatile market for which even order board data is a poor indicator of actual shipments. Not surprisingly, when we replace the order board data with the forecast data over the same period, the deviation is *significantly* greater. These results motivate two questions. First, are high-tech demands simply too volatile for meaningful forecasting? Second, to what extent can standard forecasting methods help in this situation? We explore these questions in the next section.

3.2 Applying Time Series Forecasting In many high-tech manufacturing environments, operations managers have resigned themselves to the fact that demand is too volatile to forecast. A common belief is that timely information updates, reduced lead times, and well-controlled operations would enable production to be driven completely by orders from the order board. However, as illustrated in Figure 1, even the order board data may be highly unreliable. For the purpose of long-term planning, we need a comprehensive and in-depth characterization of demand, despite the inherent difficulty in constructing one. For specialized semiconductors, there are significant technological barriers to reducing production lead time, and there is no meaningful way to build finished goods inventory, since most IC chips are customized for special functionality. Therefore, capacity planning plays a crucial role in combating demand uncertainty both for in-house capacity expansion and outsourced capacity negotiation and capacity reservation. Many high-tech firms including Agere are interested in investigating forecasting methods for the purpose of capacity planning.

We know from the literature that time series forecasting methods generally are not appropriate for high-tech products such as semiconductor and telecommunications (c.f., Skiadas 1986; Mahajan et al. 1990; Sharma et al. 1993; Islam and Meade, 1997). Technology products tend to have short product lifecycles as a result of continued innovations, and the data available early in the lifecycle typically are insufficient for time series analysis. Since we were interested to know to what extent Agere's demands corroborate the thesis, we analyzed the shipment data described in the previous section to assess the general effectiveness of time series methods in this context.

Recall that the data set includes shipment data for 560 products over a fourteen-month time period. We selected the first twelve months as the estimation period (EP) in which we estimated trends using different time series methods. We reserved the remaining two months as the validation period (VP), and we measured the forecast error. More details of the methodology are given in Appendix A. Table 1 shows the analysis for four representative IC products from the personal computer (PC) and telecommunications (TC) markets using nine different forecasting methods. The mean absolute percentage error (MAPE) associated with an EP represents the trend-fitting error, while the MAPE associated with a VP represents the forecast error.

As can be seen from the Table 1, the forecast errors (see the VP columns) are enormously high. The time series forecasting methods applied here are admittedly simple and straightforward, however, the results are sufficient to demonstrate the difficulty of judging the predicting strength of a forecasting method given its trend-fitting performance on historical data. Traditional time series forecasting methods are designed for situations where the demand trend is stable or cyclic. This is not characteristic of high-




Figure 2: Examples of Demand Leading Indicators

tech products, whose demand can vary tremendously going through the different stages of its lifecycle. As predicted by the technology forecasting literature, time series forecasting methods that rely on a product’s historical demands do not yield satisfactory results. These results motivated us to investigate and develop a fundamentally different approach.

4. The Leading Indicator Analysis In this section, we address the second group of research questions posed in Section 3. We are interested in determining if any discernable patterns can be derived from historical or current demand data that would enhance our understanding of Agere’s demands. More specifically, we are interested in determining if there exist certain demand “leading indicators” that provide advanced warning of major demand changes. These questions were motivated by our experience in 1997 to 1999 of analyzing demand data of Agere, which was Lucent Technologies at the time. As documented in Meixell and Wu (2001), after analyzing demand data for some 3,500 products, we found that the products followed approximately six lifecycle patterns and that the products could be grouped according to these patterns using statistical cluster analysis. More importantly, after performing correlation analysis on historical shipment data, we found that in each cluster there exists a subset of “leading indicator” products that provide advanced indication of changes in demand trends.

A leading indicator of a product group can be characterized by the correlation of its demand pattern in relation to the group and the *time lag* by which the demand pattern leads the rest of the group. Clearly, there is a trade-off between the two. For example, the first chart in Figure 2 shows a leading indicator that predicts the demand pattern of a larger group three months ahead of time with a correlation of 0.96; the second chart shows a six-month time lag with a correlation of 0.82. In each example, the leading indicator’s demand is less than 2% of the cluster demand and is excluded from the cluster demand calculation. This prevents a product from being identified as a leading indicator simply because of its large volume.

The Lifecycle Effect Technology lifecycles for high-tech products are known to follow a general demand cycle that starts with an initial *growth* (ramp up) followed by a period of *stability* and then a *decline* in sales when a new generation of products is introduced. The lifecycle of a product is driven in part by technological innovation as well as market competition. As discussed previously and in the technology forecasting literature, one reason that the traditional time series forecasting approach is ineffective for high-tech products is due to the short technological lifecycle demonstrated by these products, i.e., there is no reason to believe that the demand trend demonstrated in historical data is going to continue in the future. The basic hypothesis behind the leading indicator analysis is that there exists a subset of products (the leading indicators) that captures the lifecycle effect of a larger product group.

To test this hypothesis, we performed extensive analysis on a recent Agere data set using the spreadsheet-based “leading indicator engine” that we developed. In Appendix B, we describe the core analysis procedure of the leading indicator engine but omit details of the complex data management functions. The data set covers 26-month demand including some 3,500 semiconductor products. The most important result from our analysis is that given a product group of interest, the leading indicator engine can often find one or more indicator(s) that predicts the group demand pattern 2-8 months ahead of time with a correlation ranging from 0.51 to 0.95. More importantly, these leading indicators are capable of producing reliable forecasts for the larger product group.

In following section, we provide a few highlights of the above analysis, and we use examples to illustrate

MAPE				One-Month				11-Month
	lag 1,2,3	lag 4,5	lag 6,7	Total	lag 1,2,3	lag 4,5	lag 6,7	Total
0-20%	5	3	26	34	0	0	0	0
20-40%	6	2	2	10	4	2	22	28
40-60%	1	9	9	19	4	5	8	17
60-80%	1	2	7	10	3	7	3	13
80-100%	0	1	9	10	2	5	11	18
> 100%	1	8	8	17	1	6	17	24
Total:	14	25	61	100	14	25	61	100

Table 2: Summary of Top 100 Leading Indicators

the different aspects of the leading indicator analysis.

5. Empirical analysis The leading indicator engine analyzes the data associated with a specified group of products and systematically search for a set of leading indicator products for the group and generates demand forecasts based on the leading indicator identified. The tool can be also used in a scenario analysis mode to test whether a particular product is a strong leading indicator for some group of products, which is a question of great managerial interest. Our experiments were conducted using monthly demand data from Agere that covered the (twenty-six month) period from December 2001 to January 2004. The data set included 3,500 semiconductor (IC) products across eight business entities.

5.1 Identifying Leading Indicators To begin the process of identifying leading indicators, the supply-demand planner must specify a set of products. The tool then can be used to identify leading indicators for any product group or for a subset of the product group that shares the same capacity. In this section, we illustrate both situations for the set of 643 products within a particular business entity. In Example 5.1, we consider the entire group of 643 products as one cluster, and in Example 5.2, we restrict our analysis to a subset of the products that require capacity from a particular fab. The detailed experimental procedure is given in Appendix C.

EXAMPLE 5.1 In this example, we group all 643 products within a particular business entity into one cluster, and we allow each product within the cluster to be a candidate leading indicator. We perform the leading indicator analysis over three different time intervals in order to gain insight into how the length of the time interval and the age of the historical data impact which leading indicators are selected.

Correlation Value and Forecasting Performance

To begin, we consider an estimation period (EP) covering months 1 through 15 and a validation period (VP) covering months 16 through 26. We perform the leading indicator analysis for each of the 643 products in the cluster and each product is evaluated with time lag from 1 to 15 months. The correlation is calculated between the (time lag-offset) product demand series and the cluster demand series (excluding the product under consideration). We then rank all of the product-time lag pairs by their absolute correlation over the EP. For the top 100 product-time lag pairs (leading indicators), we produce a leading indicator based forecast using the procedure described in Appendix C, and we evaluate the forecasting error (in MAPE) using actual shipment data available from the validation period. The basic procedure here is similar to that of Section 3.2, except that we use the data in EP as a basis to identify the leading indicator, while using the data in VP to evaluate the forecast performance of each leading indicator.

Table 2 summarizes the results of the above analysis for the top 100 leading indicators. Note that all of them show a correlation above 0.6. In the table, we consider the one-month and eleven-month forecasting performance of the leading indicators and show the distribution of indicators according to MAPE value and time lag. For example, the entry with value 26 in the row labeled “0 – 20%” and the column labelled “lag 6,7” indicates that there are 26 leading indicators with absolute correlation value above 0.6 with time lags of 6 or 7 months that have a one-month forecasting MAPE value in the range of 0% to 20%.

For the current example, Table 2 indicates that there are 34 products with MAPE values of at most 20% for the one-month forecast and 28 products with MAPE values of at most 40% for the eleven-month




Figure 3: Top 100 Demand Leading Indicators.

forecast. These results suggest that there exists a strong pool of leading indicators for products in this business entity. To help visualize the characteristics of this pool of leading indicators, Figure 3 plots the eleven-month forecasting performance of the top 100 leading indicators against time lag and absolute value of correlation. The figure reveals that a large number of leading indicators have time lags longer than four months, suggesting that they are capable of providing warnings for demand changes sufficiently far in advance. Notice however that some of the products with the longer time lags perform well in forecasting whereas others perform poorly. One reason why products with the longer time lags may perform poorly is that fewer data points are available for the correlation analysis after the data has been shifted to account for the time lag. Figure 3 also reveals that a strong correlation value alone is not a sufficient measure in determining a leading indicator. The results show that there are several instances in which products with relatively low absolute correlation values have relatively good forecasting performance.

Incorporating New Information

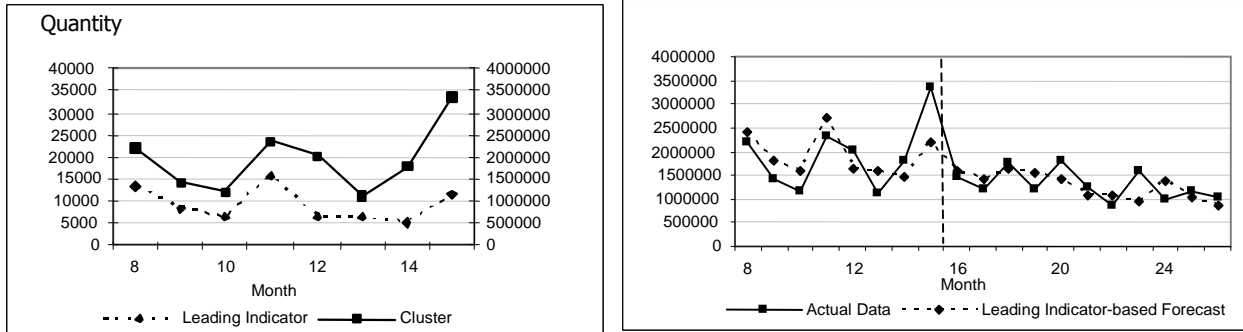
As new information becomes available over time, the correlation value and the forecasting performance of a leading indicator are likely to change. Therefore, we need mechanisms by which we update a previously selected leading indicator product and determine the amount of historical data to be used. To gain insight into the issue of updating, we repeated the leading indicator analysis for Example 5.1 but set EP to months 1 through 20 and set VP to months 21 through 26, and compared the set of leading indicator products identified for the two different EPs. Of the top 100 leading indicators identified using an EP from months 1 through 15 (in Example 5.1), forty of them appear on the list of the top 50 leading indicators under the new setting. This result indicates that the set of 100 candidate leading indicators includes both leading indicators that remain strong with the new information, and some that are misleading and should be disregarded.

To gain insight into the issue of the amount of historical data to use, we repeated the leading indicator analysis for the above example using an EP of months 6 through 20 and a VP of months 21 through 26 and again compared the set of leading indicators identified for the respective EP. Of the top 50 leading indicators that were identified using the EP of months 1 through 20, twenty-five of them appear on the list of the top 50 leading indicator products for the EP of months 6 through 20. We cannot conclude that more recent data leads to better performance of the leading indicator. Using the longer estimation period (months 1 through 20) requires more data but identifies leading indicators that perform well over a longer time horizon. Using the shorter estimation period allows for the possibility that products with an initial period of poor performance but with a high predictive value with respect to the more recent data are likely to be identified as candidate leading indicators.


Examples of Leading Indicator Forecasts

Figure 4 illustrates the forecast performance of three leading indicators, one corresponding to each of the three cases of EP and VP specified in above example. Each chart on the left shows the actual data of the selected leading indicator product and of the rest of the cluster over the given estimation period. In the figure, the dashed line shows the time series data of the leading indicator product as measured by the scale given on the left axis. The solid line shows the time series data of the cluster as measured by the scale given on the right axis. Note that the time series of the cluster is shifted ahead by the appropriate time lag so that the chart shows the mapping between the demand pattern of the indicator product and the cluster. Each chart on the right shows the actual demand of the cluster (solid lines) plotted against

a) Estimation Period: Month 1 to 15, Validation Period: Month 16 to 26
Leading Indicator: Time Lag= 7, Correlation=0.625, 11-Month Forecast: MAPE=20.11%



b) Estimation Period: Month 1 to 20, Validation Period: Month 21 to 26
Leading Indicator: Time Lag= 6, Correlation=0.696, 6-Month Forecast: MAPE=20.18%



c) Estimation Period: Month 6 to 20, Validation Period: Month 21 to 26
Leading Indicator: Time Lag= 5, Correlation=0.575, 6-Month Forecast: MAPE=30.72%




Figure 4: Leading Indicators from Example 5.1

the forecast (dashed line) generated from the leading indicator product. The vertical line separates the estimation period from the validation period.

The first set of charts illustrates the performance of a leading indicator for an EP from months 1 through 15. This leading indicator provides a signal for the demand pattern of the cluster seven months ahead of time with a correlation of 0.6251. The forecast, generated from a regression model fit within the EP (see Appendix C), results in a 20.11% MAPE over the eleven-month VP. The second set of charts show a leading indicator for an EP of months 1 through 20. This leading indicator predicts the demand pattern of the cluster six months ahead of time with a correlation of 0.6955. The leading indicator forecast results in a 20.18% MAPE over a six-month VP. Similarly, the third set of charts illustrates a leading indicator generated from an EP from months 6 through 20, which predicts the cluster demand pattern five months ahead of time with a correlation of 0.5953. The leading indicator forecast results in a 30.72% MAPE over an VP of six months.

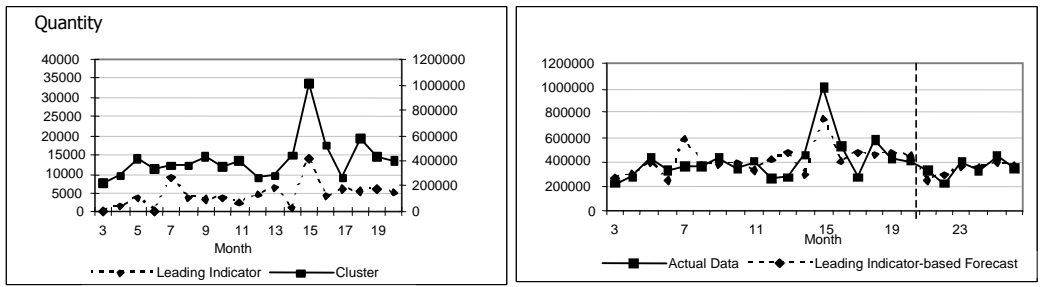
EXAMPLE 5.2 In this example, we consider products in the same business entity as in Example 5.1 but restricted to a subset of the products that use a particular wafer fab process. Future demands for this particular subset of products will have direct implications on the capacity of this wafer fab. We allow each of the 120 products within the subset to be a candidate leading indicator, and we perform the leading indicator analysis for an EP of months 1 through 20 and a VP of months 21 through 26.

For this smaller cluster, there are ten candidate leading indicators with absolute correlation values above 0.5. The average MAPE value for these candidates is 25% for the one-month forecast horizon and 40% MAPE for the six-month forecast horizon. In Figure 5, the top chart shows the results for the product with the highest absolute correlation value, 0.6680, which provides a signal for the demand pattern of the cluster two months ahead of time. This candidate leading indicator predicts the demand pattern of the cluster during the six-month VP with only 13.76% MAPE. The bottom chart shows the results for a product that exhibits similar performance to the first with respect to the subcluster. While the first leading indicator also appeared among the top 50 indicators with respect to the entire cluster (in Example 5.1), the second indicator does not. There are managers who believe that there are only a small number of leading indicator products that drive the demand for all product groups of similar characteristics. The above result seems to suggest otherwise. There is no reason to believe that a strong leading indicator for a subgroup is necessarily going to be a good indicator for the wider group.

5.2 Evaluating Candidate Leading Indicators During the process of implementing the leading indicator engine, we maintained close interactions with the supply-demand planning group at Agere. Most planners believe that there are “intuitive leading indicators,” i.e., products with characteristics that suggest they might naturally be strong leading indicators. This is also relevant from a business perspective, where managers may want to keep track of a high-volume, revenue-driving product from an important customer, and it is useful to know if this product is in fact a leading indicator for a group of related products. For this purpose, we designed the leading indicator engine such that it can be used to test whether a product in question is in fact a strong leading indicator for a specified set of products. Due to generally short product lifecycles, however, the candidate leading indicator may not correspond to an individual item but instead may be composed of successive generations of one particular technology. To handle this situation, we create a composite leading indicator to represent the progression of the technology over time. In this section, we consider a leading indicator composed of twelve products from a business entity that includes short lifecycle products. The twelve products account for about 15% of the total volume of the products within the business entity over the 26-month time horizon. Figure 6 shows the time series data associated with these products. The dotted line shows the total volume of the twelve products, while the individual curves show a rather complex pattern of technology migration over the 26-month period. In Example 5.3, we analyze whether the composite of the twelve products (consisting of multiple technology generations and modifications) is a strong leading indicator for other products in the same business entity. In Example 5.4, we determine if the composite product can serve as an indicator for other products (in the same business entity) that also share the same fab capacity. Note that since the composite product accounts for a large portion of the overall volume of the cluster, we perform the analysis in two ways – both including and excluding the leading indicator products from the cluster.

Estimation Period: Month 1 to 20, Validation Period: Month 21 to 26

a) Leading Indicator: Time Lag= 2, Correlation=0.668, 6-Month Forecast: MAPE=13.76%



b) Leading Indicator: Time Lag= 5, Correlation=0.651, 6-Month Forecast: MAPE=19.50%




Figure 5: Leading Indicators from Example 5.2

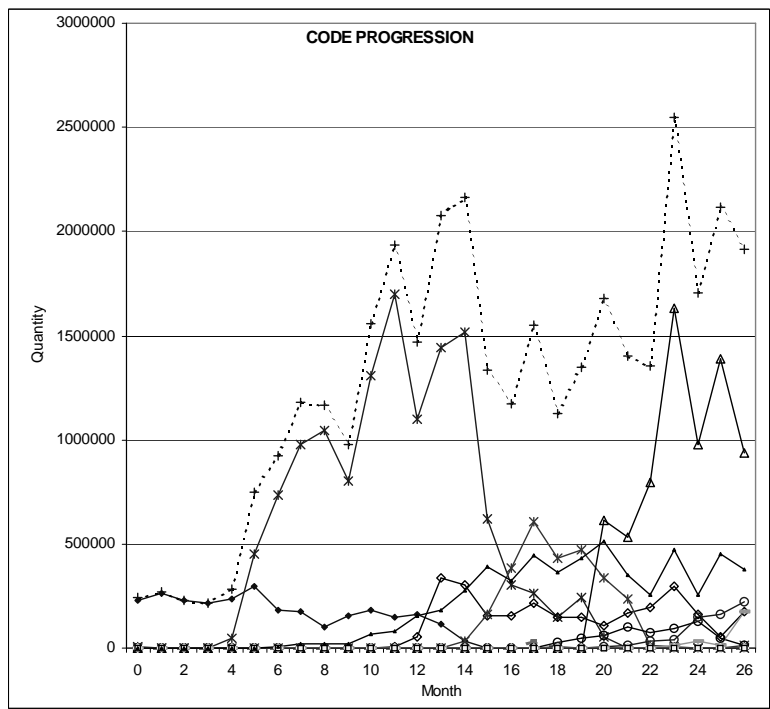


Figure 6: Progression of a Product Over Multiple Generations and Overlapping Lifecycles




Figure 7: Demand Series of the Composite Product (CP) vs. the Cluster

Time lag	MAPE	MAPE	MAPE	MAPE	MAPE	MAPE
	EP:9-24 (%)	VP:25 (%)	VP:25,26 (%)	EP:14-24 (%)	VP:25 (%)	VP=25,26 (%)
0	8.25	7.39	2.42	7.42	3.83	2.20
1	11.51	9.25	6.01	9.44	14.99	11.87
2	11.33	20.19	12.31	7.29	31.79	22.33
3	9.78	13.01	7.70	6.99	16.68	13.70
4	10.62	12.20	8.46	8.23	15.11	11.37
5	9.75	15.11	9.59	10.43	16.52	12.16
6	8.88	15.73	11.17	6.96	17.69	10.88
7	8.92	16.26	12.68	10.32	10.66	8.72

Table 3: Forecast Errors when Using the Composite Product as the Leading Indicator

EXAMPLE 5.3 In this example, we group all products within the business entity into one cluster. We perform the leading indicator analysis over two different time horizons, both of which start after the initial transient phase of the progression. The first time horizon has an EP from months 9 through 24 and the second has an EP from months 14 through 24. We use months 25 and 26 as the VP.

Figure 7 shows the actual shipment data for the composite product (CP) and for the cluster both excluding the CP and including the CP. Table 3 shows the results for the two time horizons. In this example, a time lag of zero has been considered to compare the concurrent similarity of the demand pattern of the composite leading indicator to the demand pattern of the cluster. The similarity in the two patterns can be seen both in Figure 7 and in results in Table 3. For time lags greater than zero, the results indicate that the forecast error for the validation periods are generally low. In other words, the composite product is indeed a strong leading indicator for the cluster.

EXAMPLE 5.4 In this example, we consider a cluster comprised of products that share the same fab capacity with the composite product above. To keep the example simple, we will focus on the subset of products that require the same fab process and thus share the capacity; seven of the twelve products in the composite fall in this group, and we will call the cluster CP2. There are 74 products within the business entity that share the same fab capacity with CP2. CP2 constitutes approximately 22% of the volume of the subcluster.

Since the shipment data for CP2 appears only in months 14 through 26, the analysis is performed over a thirteen-month time horizon. The EP covers months 14 to 24 and the VP covers months 25 and 26. Since large time lag values result in a small number of data points for the correlation calculations, we restrict the time lags to values from one through four to avoid misleading correlation values. Table 4

Time lag	CP2 Correlation	MAPE EP:14-24 (%)	MAPE VP:25 (%)	MAPE VP:25,26 (%)	CP Correlation	MAPE EP:14-24 (%)	MAPE VP:25 (%)	MAPE VP:25,26 (%)
0	0.9192	9.51	11.42	10.93	0.5721	20.80	26.91	25.53
1	0.7023	16.57	22.65	13.92	0.0338	25.54	37.93	34.25
2	0.8882	10.20	20.74	14.68	-0.3582	21.08	57.32	45.73
3	0.9186	7.21	15.46	24.83	0.0253	18.52	32.42	28.00
4	0.5434	12.85	19.19	14.43	-0.5053	13.94	28.42	24.03

Table 4: Forecast Errors when Using CP2 or CP as the Leading Indicator

shows the results of the analysis when the time series data of the cluster excludes CP2. (We obtain similar results when the time series data of the cluster includes the leading indicator products.)

As shown in the table, CP2 performs very well as an leading indicator for the cluster sharing the same capacity with respect to the correlation values and the MAPE values. However, due to the data available, we are only able to examine the time lag up to four months. This example demonstrates that if a correct leading indicator is selected, we will be able to use this leading indicator to provide advanced demand signal needed for capacity planning. We should point out, however, this does not happen by accident. For instance, if we use the original CP as the leading indicator for this cluster, it would perform rather poorly. This is demonstrated in the right half of the table.

6. Implications to Capacity Planning and Capacity Negotiation Agere recognizes that not only does the leading indicator engine provide a new perspective to demand forecasting, it also provides a tool to support capacity planning and capacity negotiation with supply partners. More specifically, the leading indicators provide a time-lagged model that predicts the demand pattern of a broader demand group. Suppose that the broader demand group is about to experience a shortage in the following quarter and the capacity planners have this information ahead of time, then they could renegotiate capacity levels with the partner foundries. In this context, clustering products by technology or by manufacturing resources may make sense, since the predicted aggregate demand corresponds directly to future capacity requirements.

For example, a leading indicator for a certain technology group might suggest a demand surge a few months from now. While this prediction may be highly variable and unreliable at the individual product level, the prediction for the group as a whole tend to be more robust. Moreover, the strength of the (leading indicator) prediction is quantified by the coefficient of correlation and the fitting error (in MAPE), which provides a measure for the quality of the information.

While capacity configuration and allocation are important decisions for any manufacturing firm, a few factors make this problem especially crucial to semiconductor firms such as Agere. The first factor is that there are high costs and long lead times associated with equipment procurement and clean room construction. Although a significant portion of the capacity is owned by outside foundries, state-of-the-art manufacturing equipment often costs millions of dollars and must be ordered months in advance. The clean rooms cost several hundred million dollars to a few billion dollars and take one to two years to construct. During a market upside, there may be a shortage of capacity, which means that the foundry will not be able to react to a sudden surge in demand. In this environment, an advanced signal of demand changes (e.g., from the leading indicator) is a significant advantage at the negotiation table. Specifically, if reliable demand information is available on aggregated technology groups, more favorable terms on capacity level may be negotiated a few months ahead of the competition. This could result in major savings in capacity costs, while avoiding detrimental capacity shortages during market upside.

A second factor that complicates capacity planning is the rapid advancement of fab technologies and the pace of transition from old technologies to new. Typically, fab technologies are defined by line width (the space between features on a semiconductor die) and wafer size. With each improvement in photolithography technology, new and more expensive equipment must be purchased so that features with smaller line widths can be produced. At the same time, wafer sizes are increasing, which increases the number of chips to be made at once and produces higher yields. This in turn reduces the unit cost of manufacturing. As semiconductor technologies improve, foundries must migrate their manufacturing

capability to the newer technologies. However, they are cautious with decisions on technology transition; transitions take time, and they must be anticipated correctly. A premature transition could lead to costly underutilization of equipment or necessary production of older technologies on newer, more expensive equipment. A delayed transition could lead to missed market opportunities and a lower ROI for the capital investment. The leading indicator approach could play an important role here. For instance, a leading indicator for a particular technology group could provide advanced notice on demand changes and thus signal the need for a technology migration. Since the technology migration is likely to involve contract negotiation with outside foundries, the advanced notice provided by the leading indicator may shorten the lead time for a major technology migration, enabling more favorable terms with the foundry suppliers.

A third factor that complicates capacity planning is that actual execution of the capacity plan is subject to much uncertainty, requiring frequent adjustments and reconfigurations. The “effective capacity” required to manufacture the same technology may be different in each location, depending upon the technology mix (capacity configuration), the wafer sizes made at a facility, the skill level of the labor and myriad other factors. The leading indicator approach may play a significant role in capacity reconfiguration during execution. For instance, a certain number of wafer-starts (production units) are allocated for a particular product that requires a certain technology; suppose the leading indicator projects that the demand for this product will be postponed for a few months. In this case, an operations manager may decide to act on the leading indicator information and reallocate this capacity to a mature product requiring an older technology. This can be done since newer equipment can typically be used to manufacture older technologies, albeit at a lower cost efficiency. Nonetheless, it may be cost effective to reconfigure the capacity proactively rather than reacting to the demand changes later on. Similar to the previous situations, the leading indicator could provide significant advantage by providing earlier warnings of an undesirable situation.

7. Conclusions and Future Directions We have described a project at Agere Systems in which we studied demand characterization methods for short-lifecycle technology products. We have developed and implemented a spreadsheet based “leading indicator engine” that systematically searches among a specified group of products for leading indicator(s). The leading indicator engine provides a multi-purpose decision support tool that has significant implications to capacity planners, supply-demand planners and others.

In addition to capacity planning and capacity negotiation, the leading indicator analysis has important implications to other planning functions such as financial forecasting and inventory forecasting.

Financial Forecasting Although financial forecasting was not the motivation for this research, the leading indicator approach could be a useful tool for projecting revenue and inventory for a fiscal period. In this context, a leading indicator could be used to drive and adjust revenue projections based on the trends of main revenue streams in the near future. For financially critical product groups, leading indicators could be developed to provide advanced notice of potential revenue short falls or new business opportunities. One major difference between capacity planning and revenue forecasting is the form of the data that drives the planning process. Capacity planning is concerned with expected unit volume requirements of specific resources, whereas revenue forecasting is concerned with estimated sales for a specific market segment, business entity or customer. To reflect this difference in the leading indicator approach, demand could be characterized in terms of sales rather than unit volume.

Inventory Forecasting The goal of inventory forecasting is to project inventory cost and/or inventory velocity for a given future period. Inventory is perhaps one of the most difficult phenomena to project in the high-tech industry, because it is a product of many highly volatile factors, including sales, product mix, product cost, manufacturing yield, cycle time variation and supply volatility. As such, a methodology that could simplify the process of forecasting inventory would be very valuable. The leading indicators studied in this research are based on a characterization of “demand”. Since inventory is a phenomenon driven by more than just demand, the specific analysis in this research has limited applicability for predicting inventory. However, one might ask whether a similar leading indicator analysis based on identifying leading indicators for inventory cost rather than demand could be developed and applied as a useable inventory model. Another approach that warrants consideration is to develop leading indicators for each of the factors that significantly influences inventory and combine these leading indicators to derive a

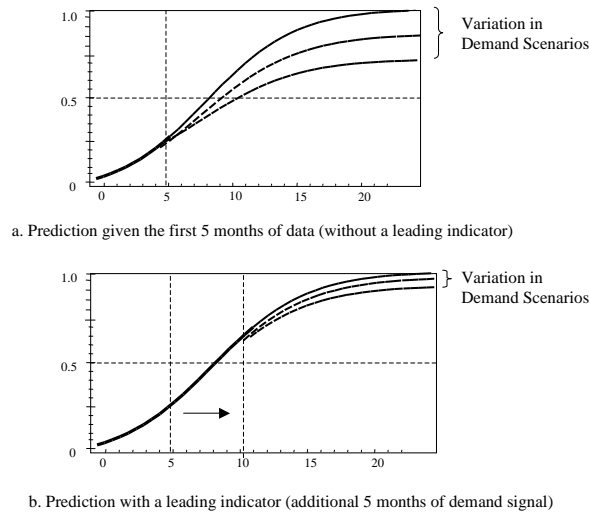


Figure 8: Predicting Demand Growth with or without the Leading Indicator.

leading indicator for inventory.

Predicting Demand Growth An important realization concerning semiconductor products is that a particular product only goes through one lifecycle of *growth*, *stability*, and *decline*, i.e., a single modal lifecycle curve. Therefore, the cumulative demand of a product over its lifecycle can be expressed as an *S*-shaped function. The shape of this function specifies the precise pattern of demand growth over time. More specifically, the demand growth pattern can be characterized by the point of inflection of the *S*-shaped function, which represents the most drastic change. We are currently examining statistical methods that use the leading indicator as a means to streamline the projection of demand growth patterns for a product group of interest. The main idea is illustrated in Figure 8. For a product group of interest, it is possible to project probabilistically a number of different demand growth patterns from the current point in time to the end of the demand lifecycle (see Figure 8-a). However, the variance associated with such projections could be too high to be useful. Using the leading indicator, it is possible to reduce the variance of the projected demand growth patterns. This can be accomplished by monitoring the demand of the leading indicator product and use its advanced demand signal to (Bayesian) update the initial demand projection, thereby reducing its variance. As illustrated in Figure 8-b, the reduction in variance can be significant; this is due to the fact that a small movement on the time axis might correspond to a drastic change on the demand curve, especially when the point of inflection is included in the movement.

Technology Substitutions In the context of technology forecasting and demand characterization, an additional complication is the replacement effect demonstrated by subsequent generations of a technology. For example, at some time during the lifecycle of a particular chip designed for a cell phone model, a next-generation chip is in the process of being designed and developed, perhaps for a new cell phone model. The demands for the new product will begin to replace the demands for the old product during its lifecycle. As illustrated in Figure 6, in reality, the migration of technology innovation over multiple generations may not be “clean-cut” and could resemble significant overlaps, driven by a complex replacement relationship (e.g., several existing chips maybe replaced by one new chip). Researchers (c.f., Kumar and Kumar, 1992; Islam and Meade, 1997; Sharif and Kabir, 1976) have proposed simplified technological substitution models as a means to capture the successive generations of technology products. We are currently working on extensions that examine the implications of technology substitution to the leading indicator analysis.

Acknowledgments. We are grateful to Kumar Abhishek, Christopher Cassin, Nian Hua, and Douglas Coburn who contributed to the development of the leading indicator engine. We acknowledge the generous support from Agere Systems. This research is also supported, in part, by the National Science Foundation Grants DMI-0121395, DMI-0432422, and the Semiconductor Research Corporation (SRC) grant 2004-OJ-1223.

Appendix A. Measuring Forecast Error using Available Shipment Data To measure forecast errors as summarized in Table 1, we use the following process. For a given part ID, consider the time series of monthly shipment quantities. Let $T = m + n$ be the number of periods (months) in which we have actual shipment data. Given a forecasting method, we use the first m -month shipment quantities as the estimation period (EP), where a trend line is generated to fit the data. We then use the remaining n months as the validation period (VP), where the forecast predicted by the trend lines is validated using actual shipment data. To measure how well each trend line fits the data, the mean absolute percentage error (MAPE) was calculated as follows:

$$MAPE(\xi) = \frac{1}{\xi} \sum_{i=1}^{\xi} \frac{|y_i - \hat{y}_i|}{y_i} \quad (1)$$

where y_i is the actual quantity shipped during period i and \hat{y}_i is the trend line estimated quantity shipped during period i . MAPE measures forecast error as a percentage of the actual (shipment) quantity. Thus $MAPE(m)$ measures how well a particular trend line fits the data during the estimation period and $MAPE(n)$ measures how well the trend line predicts the demand.

Appendix B. The Leading Indicator Search Procedure

1. The user identifies a product group of interest and sets a threshold specifying the minimum *time lag* and *correlation* required. To initialize, we put all products in the group into one common *cluster*.
2. Finding Leading indicators. Within each *cluster*, we find all leading indicators above the required threshold as follows:
 - (a) Initialization. Given a cluster C of products, select a product i from the cluster, set time lag $k = 1$.
 - (b) Main Step. Compute the correlation between (1) the demand time series associated with product i where the time series is offset by $(t - k)$ and (2) the demand time series associated with the cluster excluding i (set $C \setminus \{i\}$)
 - (c) Set $k = k + 1$. Repeat the Main Step and record the correlation number ρ_{ik} computed for each product i with time lag k .
 - (d) Repeat steps (b) and (c) for each product $i \in C$.
3. Examine all the correlation numbers ρ_{ik} computed. If at least one of the correlation ρ_{ik} and its corresponding time lag k satisfy the specified thresholds, go to Step 4. Otherwise, perform re-clustering as follows:
 - (a) Re-clustering. Using statistical cluster analysis, subdivide the product group into clusters based on statistical patterns demonstrated by each product's historical demand; a variety of attributes may be used for clustering, e.g., mean shipment quantity, shipment frequency, volatility, skewness, etc.
 - (b) Repeat Steps 2 and 3 for each cluster.
4. Return the leading indicator(s) and the corresponding product cluster(s).

Appendix C. Experimental Settings and Procedures In the following, we describe the general settings and the procedures used for all experiments reported in Section 5.

The Estimation-Validation Procedure

In all experiments, the 26-month data set is split into an estimation period (EP) and a validation period (VP) as described in Appendix A. Let $[1, T]$ be the time period (in months) in which the shipment data is available, and let the subperiod $[t_0, t_1]$ be the EP in which the leading indicators are identified and the parameters for the forecast are determined. The remaining time period $[t_1 + 1, T]$ is used as the VP over which the forecasting performance of a candidate leading indicator is tested. As such, any h -months forecast can be validated by the data set by comparing the forecast to actual shipment and $h \in [1, T - t_1]$.

The Coefficient of Correlation

Over the estimation period $[t_0, t_1]$ the degree of the linear relationship between the time series of cluster C and product i at time lag k is quantified by using the following correlation coefficient:

$$\rho_{ik} = \frac{\sum_{t=t_0+k}^{t_1} (x_{i,t-k} - \bar{x})(y_t - \bar{y})}{\sqrt{\sum_{t=t_0+k}^{t_1} (x_{i,t-k} - \bar{x})^2 \sum_{t=t_0+k}^{t_1} (y_t - \bar{y})^2}} \quad (2)$$

where $x_{i,t}$ and y_t denote the actual shipment quantities of a candidate leading indicator i and the rest of the cluster in month t , and \bar{x}_i and \bar{y} are the average shipment quantities over the corresponding time horizons in which correlation is calculated. Thus, the correlation coefficient ρ_{ik} measures how well the demand of item i over time period $[t_0, t_1 - k]$ predicts the demand of the cluster over $[t_0 + k, t_1]$.

Note that the correlation coefficient is determined by comparing the time series of the item against that of the rest of the cluster. The total shipment quantity of the cluster is adjusted by removing the item's quantity from each month's shipment quantity. In this way, the bias that might be introduced from a (high-volume) dominating item is eliminated.

The Leading Indicator Based Forecast

After a leading indicator i is identified from a cluster C based on time lag k and the coefficient of correlation ρ_{ik} , we construct a forecast for cluster C based on the time series of the leading indicator using the following procedure.

1. Regressing the time series of cluster C over $[t_0 + k, t_1]$ against the time series of the leading indicator over $[t_0, t_1 - k]$. Determine the corresponding regression parameters $\hat{\beta}_0$ and $\hat{\beta}_1$.
2. For a given month t , generate the forecast for the cluster, \tilde{y}_t , using k -month earlier time series data of leading indicator i as follows:

$$\tilde{y}_t = \hat{\beta}_0 + \hat{\beta}_1 x_{i,t-k} \quad (3)$$

3. Calculate the forecast error for leading indicator i over the validation period $[t_1 + 1, T]$: for an h -month forecast during VP, calculate $MAPE(h)$ based on (1).
4. Calculate the overall fitting error over the estimation period $[t_0, t_1 + h]$: calculate $MAPE(m)$, $m = t_{1+h} - t_0$.

References

- [1] Chris A. Armbruster, *Case study: Optimizing the supply chain for horizontal business models*, Tech. report, Agere Systems Technical Report, Allentown, Pennsylvania, 2002.
- [2] Gartner Group, *Semiconductor manufacturing and design predictions for 2004*, Research Report COM-21-7334, Gartner Group, 2003.
- [3] T. Islam and N. Meade, *The diffusion of successive generation of a technology: A more general model*, Technological Forecasting and Social Change **56** (1997), 49–60.
- [4] U. Kumar and V. Kumar, *Technological innovation diffusion: The proliferation of*, Journal of Economic Perspectives **10** (1996), no. 3, 135–152.
- [5] V. Mahajan, E. Muller, and F.M. Bass, *New product diffusion models in marketing: A review and directions for research*, Journal of Marketing **54** (1990), 1–26.
- [6] Mary J. Meixell and S. David Wu, *Scenario analysis of demand in a technology market using leading indicators*, IEEE Transactions on Semiconductor Manufacturing **14** (2001), 65–75.
- [7] N.M. Sharif and C. Kabir, *A generalized model for forecasting technological substitution*, Technological Forecasting and Social Change **18** (1976), 353–364.
- [8] L. Sharma, A. Basu, and C. Bhargava, *A new model of innovation diffusion*, Journal of Scientific and Industrial Research **52** (1993), 151–158.
- [9] C. Skiadas, *Innovation diffusion models expressing asymmetry and/or positively or negatively influencing forces*, Technological Forecasting and Social Change **30** (1986), 313–330.