Dynamic Modeling and Forecasting on Enterprise Revenue with Derived Granularities

Jerry Z. Shan, Hsiu-Khuern Tang, Ren Wu, Fereydoon Safai
Intelligent Enterprise Technologies Laboratory
HP Laboratories Palo Alto
HPL-2005-90
May 12, 2005*

Bayesian inference, data granularity, modeling and forecasting, seasonal ARIMA models

Timely and accurate forecasts are crucial in decision-making processes and have significant impacts on many business aspects. We at HP Labs have developed a complete set of quantitative forecasting methods that can enable the establishment of a reliable predictive reporting system, so that executives can discern as early as possible where the company is heading financially. This paper reports some of our technical developments in building such a predictive reporting system.

* Internal Accession Date Only
To be published in and presented at IEEE International Conference on Granular Computing (IEEE GrC 2005), 25-27 July 2005, Beijing, China
Approved for External Publication

Dynamic Modeling and Forecasting on Enterprise Revenue with Derived Granularities

Jerry Z. Shan, Hsiu-Khuern Tang, Ren Wu, Fereydoon Safai

***Abstract*—Timely and accurate forecasts are crucial in decision-making processes and have significant impacts on many business aspects. We at HP Labs have developed a complete set of quantitative forecasting methods that can enable the establishment of a reliable predictive reporting system, so that executives can discern as early as possible where the company is heading financially. This paper reports some of our technical developments in building such a predictive reporting system.**

***Index Terms*—Bayesian inference, data granularity, modeling and forecasting, seasonal ARIMA models.**

I. Introduction

Commercial enterprises compete for customers by promising, among other things, low prices and fast delivery. Successful competition often requires careful monitoring of profit margins and deadlines. One key to success in this environment is a system that provides accurate and timely business information, which includes forecasts of major financial metrics. Specifically, businesses often wish to use observed data to forecast some outcome (e.g., end-of-quarter revenue) or to monitor the probability of achieving some goal to support current business decisions. Given the fact that a large enterprise's ongoing transactions are complex and difficult to model, this task may be quite challenging. One alternative to constructing transaction-based models is to employ stochastic modeling techniques for forecasting, as we are demonstrating in this paper.

II. Problem Statement and Objectives

Some examples of the specific problems that we addressed successfully are as follows.

- What are the revenue and margin forecasts for the whole HP Company, say, for this month and this quarter?
- What is the probability of achieving the ASPIRE goal that the company has set at the beginning of a planning period?
- How can we expand the dimensions to include many of the business entities?
- How can we expand the time horizon so that specific operational guidance on a quarterly basis can be derived?
- In what granularities and at what aggregation levels, in the time dimension, in the geography dimension and in the business operation dimension, should we monitor, model and forecast the relevant metrics of our interest?

A. Background Information and Business Needs

HP executives need a system that can provide reliable and accurate predictions for the company's total revenue and other financial metrics for any given fiscal period. Without such a system in place, HP was under enormous pressures for quite some time in our recent history. The most noticeable quarter was FY01Q1. HP had a nearly disastrous financial reporting experience for that quarter, with a forecast error about 12% when the company was maintaining to Wall Street that we were still on track with the target growth of 15% as late as in the middle of December 2000. The difficulty was due to such complicating factors as the Agilent spin-off, and the slow down of both the tech industry and the overall economy in that period. In 2002, the merge with Compaq proved to be another significant compounding factor later.

In response to this urgent and business-critical need, the RMAP (Revenue Modeling and Prediction) team was created in HP Labs in December 2000, with the strong support and collaboration from Corporate Finance and Corporate IT.

B. Data Flow and Structure

The data surrounding the overall revenue stream has the following flow and structure.

First, there is the daily revenue part. It starts with a customer order event, followed by the product shipment event for the order, and then the customer invoice event. After the invoice event, the revenue is recognized following various Security Exchange Commission (SEC) rules. The order and shipment transaction data go to the order and shipment data warehouse, and the recognized revenue goes to the revenue data warehouse. In combination, the order and shipment data warehouse and the revenue data warehouse are linked to the online executive reporting system. From this part, we see three main data streams: order, shipment, and revenue.

Second, there is the part for the Journal Vouchers, monthly, and non-standard transactions. One example is currency adjustment. These month-end events happen at the end of a

Manuscript received March 17, 2005.

All the four authors are with HP Labs, Hewlett-Packard Company, CA 94304, USA (corresponding author email: jerry.shan@hp.com; phone: 650-857-5779; fax: 650-857-6278).

fiscal month, and after the closure of the previous calendar month. With various reconciliations between the observed daily revenue stream and the month end adjustments, the final monthly revenue data is ready to go to the financial statement generation process, which is usually performed in the third month of a fiscal quarter. On a quarterly basis, with some additional quarterly reconciliation, we generate the quarterly financial statements that executives report to the financial community.

The above is a generic description for the data flow and structure surrounding the revenue data. For HP, the flow and structure are similar to this generic description but there are also some additional components that are unique to the company.

C. *Objectives*

To help understand the problems, Figure 1 provides a simplified description that illustrates the monthly revenue prediction. Suppose we are entering into a new fiscal month, and we are interested in predicting the revenue for the whole month. How can we make a prediction as accurate as possible and as early as possible? The curve that starts at the lower-left corner and extends to the upper-right corner represents the MTD (month-to-date) amounts, with solid curve between days 1 and 20 to indicate that the MTD amounts have already been observed, and dotted curve afterwards to indicate they are not observed yet. The dotted horizontal line is the actual total amount for the fiscal month, which is unknown beforehand. The green curve represents the daily point predictions. Note that a fiscal month extends beyond a calendar month, usually about one week to 10 days.

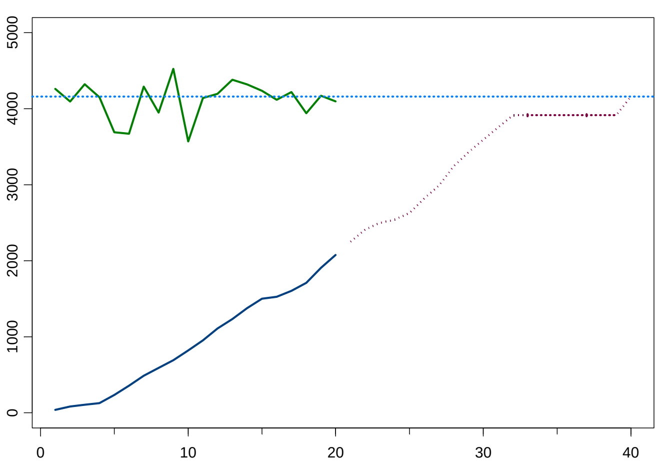

Figure 1. Illustration for the Monthly Revenue Prediction

D. *Technical Challenges*

As it turned out, the problem of forecasting financial metrics for a big enterprise such as HP differs significantly from those in the academic setting, as many technical assumptions are no longer valid with the data.

From Subsection II-B on the data flow and structure, we see three main data streams: order, shipment, and revenue. Conceptually, they're closely related, presumably with various time delays or lags. It is very tempting to explore those lag structures in order to forecast the revenue. However, as it turned out that the lag structures were too complex to model, as HP is such a complex enterprise.

One principle in forecasting is to use the observed data that is as close as possible to the metric that one is forecasting. In this case, the observed revenue data is a natural choice. However, with this thinking, we limit our explorations virtually to the univariate time series modeling approaches. How good can it be and what challenges are we facing ahead? Let us first take at look at some of the unique characteristics of the revenue data.

- A fiscal month goes beyond a calendar month, usually with a variable ending day in the following calendar month.
- The daily revenue amount can change significantly from work days and weekend days, and the delayed revenue and the month end adjustment can be quite different from the daily revenue flow.
- The changes in revenue recognition rules mandated by the SEC can bring many undesirable changes in the historical monthly profiles and patterns.

The non-homogeneous time series structure, (especially at low granularity levels), the open-ended period length, and the monthly profile change, all present significant difficulties to the modeling part, and hence to the forecasting task. The usual textbook efforts with stationary or derived stationary time series modeling turns out to be not that useful in solving this problem. In addition, the varying forecasting time horizons, the various data granularities and multiple layers of dimensions all add complexity and difficulty to the problem, as exemplified by the fact that different business entities can have different cyclic and trend effects.

III. Model Development

In this section, we will present the various development components in relation to the revenue prediction with statistical approach. As to be seen shortly, many of the development pieces originated from, and are still centered at, the revenue prediction problem; nevertheless we studied them with some broader perspectives in mind, which resulted in a suite of technological advances in the modeling and forecasting research and development area.

A. *Methods Review*

Most existing prediction methods in the statistical literature are based on the various time series models [1, 2]. These models can be classified as unstructured models and structured models. Unstructured models such as the ARIMA models usually assume that the observed data, or the differenced data of proper multiple orders, is a stationary stochastic process. The differencing part intends to take care of identifying and quantifying the trends. To further catch any seasonal effects, seasonal differencing is also applied. Their difficulties are as follows. To successfully quantify a trend, we need to correctly identify the order of differencing. This can only be achieved by an expert's interactive analysis. Similarly, to successfully quantify a seasonal effect, we need to identify the seasonal cycle, which can be very difficult in the constantly changing business environment. In addition, there can be multiple layers of trends and seasonality such a monthly, quarterly and yearly. The fact that the fiscal month, quarter and year can be different from their calendar counter-parts, produces additional complexity. No matter what orders of differencing are performed, the resulting data can still be highly non-stationary for data like daily revenue loaded from financial reporting systems. The reasons are due to many factors such as the weekend loading and weekday non-loading, both of which can be very irregular, and interruption and delay of loading services.

A structured model overcomes the difficulty of having to perform the differencing task, as it specifies explicitly the many effects, including trend and seasonality, in the model. However, if a particular effect is not explicitly specified in the model and it turns out it indeed exists, there is no way to extract it from the fitted model later on. In addition, the omission of an existing effect from the model can seriously compromise the identification and estimation of other effects. For the daily revenue data, identifying all major effects beforehand is near impossible. One example is the effect by the quick slowing-down of the high tech industry during the past few years. For structured models, there are still many significant open issues such as estimation initialization for trend and seasonality in the Holt-Winters prediction method.

B. Prediction with Profile Extrapolation

In the early stage of our engagement, we developed a monthly revenue prediction method, which is based on linear extrapolation and modeling on profiles. Due to the scope limitation of this paper, we will not get to the development details. The major steps are listed as follows.

(1) From financial reporting systems, extract the historical daily revenue data, which include those in the past months and in the current forecasting month.

(2) Identify a set of past months, called training months, by using recent months to reflect the market trends.

(3) For each training month identified in Step 2, compute the cumulative sum (i.e., month to date amount) on each day. Compute the total revenue for the fiscal month, which includes the total cumulative sum from the beginning to the last calendar day, plus any amount that is recognized in the following months but is for the fiscal month. Compute the percentages of the daily cumulative sums relative to the monthly total. This percentage time series is mathematically a function defined on all calendar days of the given training month, and we call it the percentage function.

(4) For each training month identified in Step 2, normalize its length to that of the forecast month. (A month length is defined as the number of calendar days in the month.) The normalization of lengths can be expanding, say from 30 days to 31 days, and can also be shrinking, say from 30 days to 29 days. Any month before this normalization is called an original base, and after the normalization, is called normalized base.

(5) For each percentage function obtained in Step 3, define by linearly interpolating the function values for all the points in the normalized base obtained in Step 4, from those values in the original base. Do this for each of the training months. Call the resulting function the generalized percentage function. Note that the generalized percentage functions for all the training months now have the same definition domain due to the aforementioned normalization in Step 4, regardless their original lengths, which can be different. This common definition domain is the set of all days in the forecast month by design in Step 4.

(6) At each point in the normalized base, get the generalized percentage function values for all training months, as obtained in Step 5, and compute a robust average value, with the statistical median as a preferred choice. By connecting all these robust average values defined on the normalized base, we get a robust average monthly behavior time series for the cumulative realized percentages.

(7) Model the expected monthly behavior for the forecast month with the robust average monthly behavior of the training months.

(8) For the forecast month, compute the daily cumulative sums from the very beginning to the last day when month to date revenue amount is available.

(9) Synthesis the month to date revenue amount for the forecast month obtained in Step 8 with the robust average monthly behavior time series obtained in Step 7, by dividing the former by the latter. Take proper care for the zero value denominator cases.

(10) The ratios derived in Step 9 are the forecasts for the fiscal month's total revenue, made for days starting from the very beginning to the last day that we have observed. In particular, the last value is the most updated forecast.

(11) Connect the ever-updating forecasts obtained in Step 10 to an interface such as a web service, and make it available to the users to support their decision-making processes.

C. Dynamic Modeling and Prediction

In this subsection, we present some of the major development that we have created after our earlier development phase. The objectives are to address some of the difficult issues that the earlier solutions can not resolve. We will explain the issues, along with our solution, in the following.

Given the complexity of the data flow, our solution goes as follows. First, decompose the data into three streams based on the characteristics of the monthly data flow: the pre-month stream, the within-month stream, and the post-month stream, which consists of the month-end delay part and the month-end adjustment part. Second, analyze the historical data for each stream and model them separately. Finally, make a forecast for each stream based on its developed model. With this frame, the task now is to focus on modeling each individual stream.

We exemplify the modeling for the within-month stream in the following. For the forecasting month, we model the whole within-month profile (normalized revenue accumulation vs. normalized time) with the average profile of a properly chosen set of historical months. Based on the model and the month-to-date amount for the forecasting month, a prediction can be made, which includes a point prediction and a confidence interval prediction. This profile based method is relatively straightforward. However, the resulting performance, especially the prediction intervals, can be unsatisfactory in the early days of a month as there is usually too much volatility in the derived model.

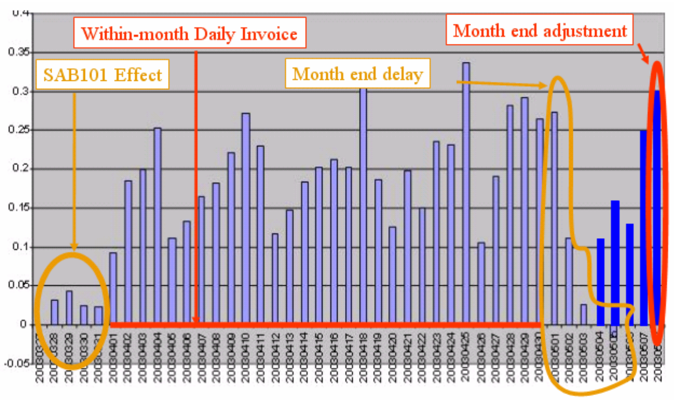

Figure 2. Decomposition of Recognized Revenue Flow on a Monthly Basis

Our solution next is to explore the Bayesian modeling approach. *Bayesian inference* is a statistical inference in which probabilities are interpreted not as frequencies or proportions, but rather as degrees of belief. With the Bayesian modeling approach, one starts with an initial set of beliefs about the relative plausibility of various hypotheses, estimates, or predictions in our case, collects new information (for example, the new daily revenue), and adjusts the original set of beliefs in the light of the new information to produce a more refined set of beliefs of the plausibility of the hypotheses or predictions.

Now we explain our Bayesian modeling development. First, we note that very early in a month the month's total will not exceed either bounds of the unnecessary-wide interval that was derived with the profile based method. Second, as the days go by, we have the month-to-date information in terms of accumulated revenue and the covered time, both we did not make full use of. Our main development for this part is the innovative modeling for all the components that are needed in deriving for the posterior distribution of the total amount in the whole prediction period, given what we have observed so far such as the MTD amount. The explanations are as follows.

(a) First, we quantify and incorporate all those pieces of information into the proper models, using the standard Bayesian inference and modeling techniques. With this, the out-of-bound problem is solved. This is expressed in the first part of the equation.

(b) Next, we go beyond the standard techniques by figuring out a way in linking the posterior distribution, which is difficult to model directly, to a stochastic invariant (something that is stochastically stable from period to period), which has good prediction power and can be readily modeled. That is the density function for the normalized ratio at any given time.

We find a multiplicative separation in modeling the stochastic invariant and the monthly total part. For the stochastic invariant, we use proper nonlinear smoothing techniques for the within-month behaviors. For the monthly total part, we have developed a methodology that can identify and quantify the major seasonal effects, and with that, we model the density. In combination with (b), we can now derive a generally excellent Bayesian prior and start the prediction and automatic updating process.

For complete details on our generic method, we refer the reader to our references [5, 6, 7].

D. Summaries

Modeling and forecasting the within-month stream is only one of the four parts for the monthly prediction work. Due to the scope limitation, we can only highlight some of our overall key development and innovations in this paper.

(1) Capturing. We built models that can interactively analyze and capture the unique characteristics of different business entities (business segments and/or geographic regions)

(2) Quantification. We developed methodologies and tools that can automatically identify and quantify various cyclic and trend effects in any temporal data.

(3) Initial Prediction. We developed methodologies for combining traditional SARIMA (seasonal autoregressive integrated moving average) approaches in deriving the a-priori (initial) period-end forecasts for each business entity. Those initial forecasts are generally very good ones.

(4) Dynamic Updating. We enhanced the standard Bayesian modeling techniques, and built updating models that can dynamically update forecasts as the most recent observation and/or business information comes available.

(5) Seasonal Adaptive Selection. We developed a method that adaptively select the historical reference period in building the model profile for the prediction period.

Figure 3 illustrates and reports the prediction accuracy at the HP enterprise level, with a MAPE at only about 1.3% for the FY2004Q1 quarter. The prediction results at various business entity levels are also very good but are in general not as good as at the HP enterprise level, which is not unexpected due to the loss of data aggregation from high level to lower levels [4].

With the automatic nature of our prediction solution, it can identify the proper data aggregation levels and granularities, in the time dimension, the geographic dimension and the business operation dimension, to model and forecast the metrics of interest.

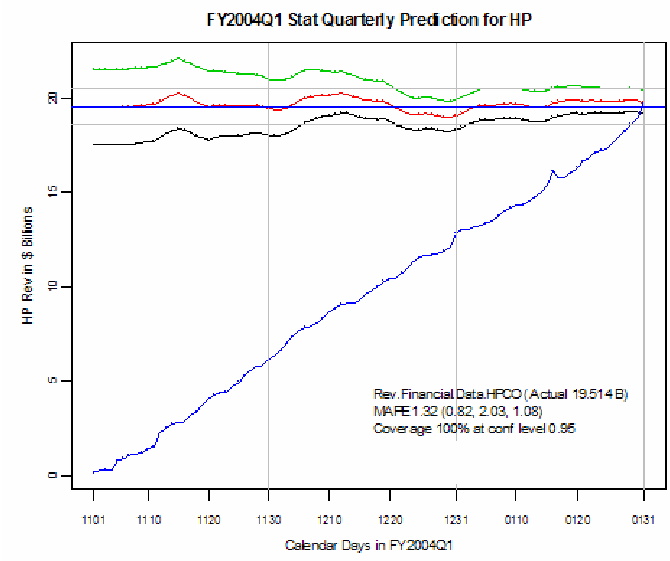

Figure 3. Statistical Quarterly Revenue Prediction

IV. A Truly Successful Collaboration

We compared various forecast metrics between our methods and other commercially available software solutions. Those forecast metrics include the forecast error, confidence interval width, forecast bias, and coverage. In general, we found our methods readily outperform the competitors. In fact, most existing prediction methods in the literature and in implementation of commercially available software are based on the various assumptions on the input data, and we found those assumptions could be seriously invalid for our data.

For judgment-based forecast, for example, there are methods that are derived on such information as sales quota data. As we now realize, the sales quota information can be very misleading. For instance, if a sales manager sees the likelihood of going moderately or severely below the original quota, he or she would have no incentive to generate more sales toward the end of this period. In fact, he or she would have every incentive to delay the sales in this period to the next so that he or she can meet his or her next target more comfortably and get an inflated bonus. The prediction accuracy based on sales quotas for this period could thus be seriously compromised.

Other forecasting methods are based on collective judgmental calls such as the ASPIRE process that HP has been using (now in combination with our quantitative methods). We believe, as it has been proven consistently, the combination of good collective judgment based methods and our quantitative methods can deliver systematically the best forecast.

Our solution has been proved to work very satisfactorily. The monthly statistical revenue forecast method has been in use as for the company's revenue prediction since April 2001, and the quarterly statistical revenue forecast method has been in use since FY2003 Q2. Top executives have benefited significantly from the increased visibility. The tool resides at the hp portal Dashboard, and about 3000 users so far (mostly executives and senior managers) have access to and are effectively using it. For example, HP CFO Bob Wayman is a frequent user of the tool, and he has developed a great deal of confidence in it over years. To quote Bob Wayman, “It is reassuring to have a solid projection algorithm, it’s crisper and cleaner, it has rigor and methodology as opposed to my own algorithm.” In the past seven quarters, we have consistently achieved a forecast error of less than 2%, targeted by the Corporate Finance. Figure 4 is a snapshot of the revenue forecast tool on “My DashBoard” that HP executives and senior management use, which is integrated as a web service at the @hp portal.

In summary, timely and accurate forecasts are crucial in decision-making processes and have significant impacts on many business aspects. Our experiences at HP and the evidences with our innovative forecasting solutions have shown that with improved forecasting models in place, decision makers at various levels can be in a much better informed position.

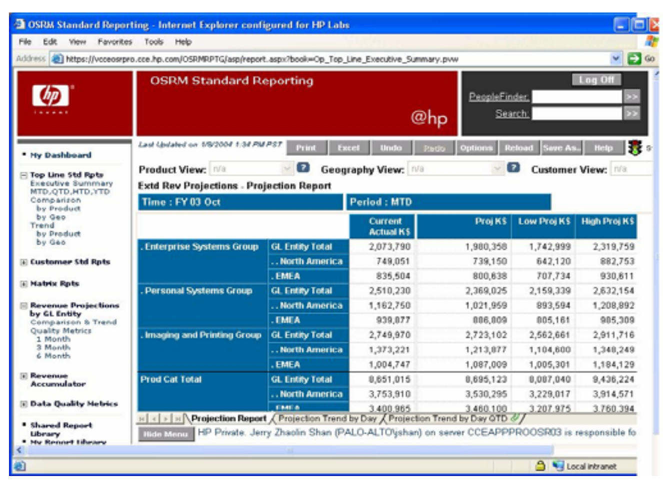

Figure 4. A Snapshot of the Revenue Forecasting Tool on “My DashBoard” at the @hp portal

Acknowledgment

The authers wish to thank all the other team members in the RMAP team: Peter Wright and Phil Blocher for our joint effort. We would also like to thank the following people: CFO

Bob Wayman and Jon Flaxman for their frequent use and insightful feedback, Craig Vent for his strong management support and stimulating conversations; Umeshwar Dayal for bringing the original problem to our attention; Shailendra Jain, Haydee Rinehart, Mary Bach, Martin Lahiff and all other managers for their strong management support. Finally, we would like express our gratutute to all others who have helped the project at various time points: Lauren Turner, Mark Misroch, Betty Cant and others.

REFERENCES

[1] George E. P. Box, & G. M. Jenkins (1994) Time Series Analysis, Forecasting and Control (3rd ed.), Englewood Cliffs, NJ: Prentice-Hall.

[2] Chris Chatfield (1999) The Analysis of Time Series (5th ed.), Chapman & Hall/CRC.

[3] Joa S. Lim and Marcus O'Connor (1996) "*Judgmental Forecasting with Time Series and Causal Information*", International Journal of Forecasting, 12, 139-153.

[4] Jerry Z. Shan (1997) "*A Theoretical Derivation of Relationships between Forecast Errors*", Hewlett-Packard Journal. (Also appeared in HP Tech Report, HPL-97-45)

[5] Jerry Z. Shan, Hsiu-Khuern Tang, Ren Wu, Fereydoon Safai, Craig Vent, Peter Wright, Phil Blocher (2004) "*Quantitative Modeling and Forecasting with Financial Applications*", HP TechCon 2004.

[6] Jerry Z. Shan, Hsiu-Khuern Tang (2004) "*Method and Systems for Cumulative Attribute Forecasting Using a PDF of a Current-to-future value Ratio*", US Patent Application 200406989-1.

[7] Jerry Z. Shan (2002) "*Method and Systems for Predicting Revenue Based on Historical Pattern Identification and Modeling*", US Patent Application 200209271-1.

[8] Jerry Z. Shan (2002) "*Method and Systems for Predicting Revenue Based on Historical Pattern Identification and Modeling*", US Patent Application 200209271-1.