An Alternative Method for Vintage Forecasting Using SAS®

Table of Contents

Abstract

The success of financial institutions is predicated – in large part – on the ability to manage the composition and expected performance of their customer bases. Sophisticated techniques for customer segmentation, target marketing, account ranking and performance analysis have been incorporated into banks' customer lending and retention processes and are critically important for short- and long-term profitability and solvency. In particular, the ability to anticipate, track and control the behavior of a group of accounts established during the same time period (i.e., vintage) is at the heart of marketing activity and risk management policy establishment. Many widely accepted statistical techniques for vintage analysis use a standard, parametric curve when defining the month-to-month performance (e.g., monthly revenue, balances, delinquency rates, etc.) for a group of accounts. This curve is then modified using traditional time series techniques based on factors such as seasonality, bank policy, economic conditions, etc.

This approach, while generally effective, assumes the irrefutable soundness of the basic curve shape – initially based on historical vintages – and incorporates continuous time covariates ex post facto. This paper presents a technique that reverses the analytical sequence; namely, treating the components of the vintage as a set of predictions and forecasts from a cross-vintage data stream. This methodology allows for the inclusion of a variety of analytical techniques, including time series analysis, dynamic segmentation and clustering, vintage profiling, and forecast reconciliation.

The objective is to employ an integrated approach to vintage curve modeling that unifies internal bank drivers, external economic factors and past performance into a cohesive strategy free from internal biases and more closely aligned with market reality. It is believed that this methodology can provide critical insight that supplements an institution's sales and operations planning process.

Background

At the risk of oversimplification, this section presents the basic business conceptual background that will assist the reader in understanding the proposed process.

The importance of a sound revenue and loss forecast for a financial institution is underscored by three major objectives:

1. Management of earnings expectations for Wall Street.
2. Management and evaluation of an effective acquisition, marketing and risk policy.
3. Anticipation of economic and competitive conditions that promote volatility with revenue and loss streams.

The interdependency of these objectives is defined by its link to a bank's customer segmentation and profiling strategy. Accordingly, understanding customer behavior throughout volatile economic times is at the core of a bank's analytical strategy.

All banks use some form of segmentation strategy that provides a road map for targeting, managing and evaluating prospective customers, preferably those with high revenue potential, while mitigating risk. By doing so, banks can establish a logical, manageable set of portfolios that partition the broad customer base into groups with similar attributes and products. Logistically, this segmentation promotes the establishment of a management structure that enables specialization due to the perceived homogeneity of the customer groups (see Way, 2009).

Segmentation can be classified broadly into two groups: products and profitability/risk assessment.

A sample of product lines includes:

- Personal loans.
- Mortgages.
- Credit cards.
- Investment accounts.
- Lines of credit.
- Merchant card processing.
- Payroll services.
- Commercial loans.

Profitability and risk assessment categorization is more complex and – in many cases – unique to a bank's preferred methodology for customer evaluation. This segment includes factors such as:

- Account credit history (balance and payment).
- Delinquency and charge-off history.
- Fair Isaac Credit Score (FICO) from Experian, Equifax or TransUnion.
- Internally developed risk scores.

Moreover, there is a degree of interaction between product line and risk assessment that should be considered. For example, an account's payment history for a home mortgage could be dramatically different than its credit card pattern.

Most analytical methodologies have been formulated to support this segmentation (and subsegmentation) strategy. Typically they involve some form of account performance tracking for a specified period of time, usually for the duration of 36 to 60 months. The metrics defining performance vary. In the case of loans, for example, the metric could be delinquency of payment, usually classified as a days past due (DPD) number – e.g., DPD60 means a payment that is 60 days past due. Revenue could be measured, for example, by accumulated monthly interest or fees. Account purchasing behavior could be measured by the active balance in a given month over time.

Delinquency measures are usually short-term. DPD is tracked to charge-off status, usually at the 150-180 days past due mark. Revenue measures are tracked long-term. Additionally, the sheer number of accounts and wide variability of monthly behavior between accounts necessitate some form of aggregate analysis. In this case, the segmentation strategy is crucial in maximizing the degree of portfolio homogeneity and minimizing the in-segment account variation. The usual assumption is that each portfolio subsegment established in a given month will have some uniformity of behavior with a subsegment established in a prior month. This has given rise to the traditional form of techniques known collectively as vintage analysis.

In its most literal sense, the term "vintage" is borrowed from the wine industry denoting a yield from a crop base bottled in a given time period. The yield is tracked and evaluated for a specified period and progresses through various stages of maturity. In banking terms, new accounts opened in a given month represent the base and the yield could reflect any number of metrics (both positive and negative), each of which defines account credit management and loan behavior over time. Each vintage has its own set of characteristics and can be evaluated qualitatively while in progress or retrospectively following closeout. The premise of vintage analysis is based on the following assumptions:

- The initial account base composition is defined by a combination of bank policy, target account profiling and marketing campaigns.

- Vintage metrics are defined in the aggregate as a summary or average value at a given point in time (e.g., total mortgage loan balance, average card account balance, total accounts in delinquency, total account fees, etc.).

- Metrics are typically measured every month that the vintage is active relative to its starting month – most often referred to as a month on book (MOB) time period. In the case of loss analysis, a subcomponent of time refers to months (or days) since the last minimum payment was due (DPD).

- The initial set of metrics (the first few months or so) establishes a baseline that can help with predictions for the remainder of the vintage life cycle. It is assumed that the characteristic (or quality) of the vintage is reflected numerically during this period and thus should permeate the remaining months.

- The effect of outlier accounts on vintage performance is usually minimal, although it can be magnified should the account base be small to begin with or if there is excessive volatility in the account base.

- Internal bank policies are usually developed and modified over time to constrain vintage performance to manageable ranges.

- When tracking vintage performance over the MOB period, a particular curve shape is noted that reflects the maturity cycle. These curves share similarities across vintages and thus are often used as historical proxies for predicting performance of young or future vintages.

It would appear then that a reasonable vintage curve projection could be developed from a review of historical vintage performance combined with association analysis (intuitive or mathematical) on metric values with internal policies, account characteristics and perhaps market conditions. However, a number of issues need to be addressed by this approach – as will be shown in the next section.

Traditional Vintage Analysis

Although there are variations, the basic setup is as follows:

1. The metric to be analyzed is determined (e.g., revenue, active balance, delinquency, etc.).

2. Monthly information from historical vintages is captured.

3. Each data point is tagged with a MOBx designator (x = number of months since the vintage was established).

4. Vintages are segmented by predetermined classifications. The most typical classification involves a combination of FICO score band (e.g., 700-759) and product segmentation (e.g., mortgage loan, fees, credit card, etc.).

5. If possible, covariate information is associated with each vintage. If it is not time-based, then it could be considered as a classification covariate. If it is time-based, it is associated with a particular MOBx value.

6. A projection for an initial MOB1 value for a future vintage is determined, perhaps using some form of regression technique.

7. A linear or nonlinear regression technique is used to construct a projected vintage curve for the new vintage. The MOB1 value from the previous step is used as a constraining factor.

8. The curve is adjusted based the anticipated impact of external factors, internal decisions (e.g., APR adjustment, promotions, etc.) and seasonal influences.

9. As MOBx values for the vintage develop, they are used to adjust the curve going forward.

There are a number of issues that are difficult to address with this approach:

1. This approach assumes complete independence of the vintages from one another. Although each vintage consists of a distinctly partitioned customer base, they are conjoined by the progression of shaping factors at given points in time.

2. This approach does not rigorously identify a lagged effect from MOBx to MOB(x+1).

3. This approach does not factor in trend, seasonality and business cycles.

4. This approach may disproportionately weigh older vintages more heavily in the computations. More specifically, these older vintages may have been shaped by economic conditions and business decisions that are no longer relevant.

5. Covariate effects are embedded within the curve and difficult to isolate.

6. There is typically wide variation in the account base that makes up each vintage; there is little consideration for weighing the in-vintage variation.

7. The segmentation strategy makes assumptions about the homogeneity of performance across vintages within the segment.
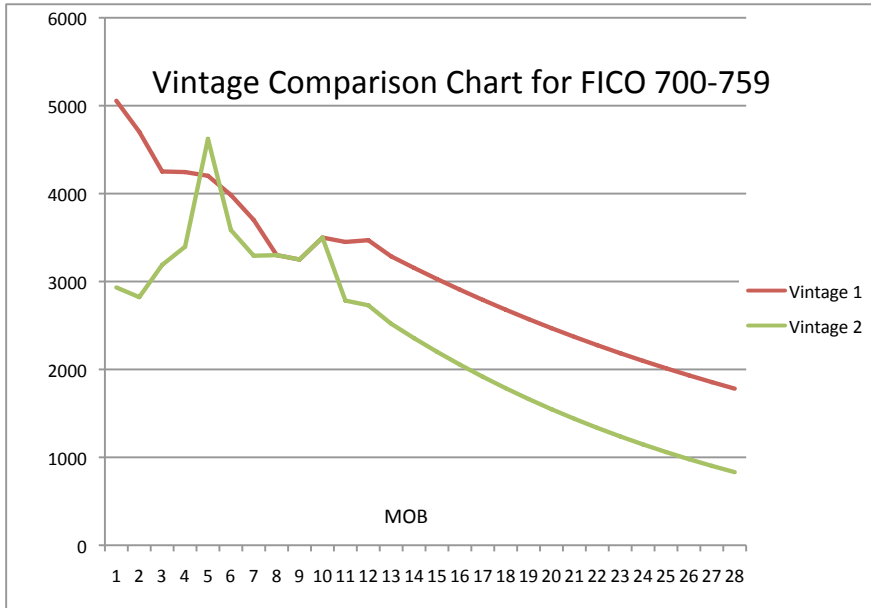
Figure 1 illustrates some of the problems.

Figure 1: Vintage curves for average active balance.

Vintage 1 and Vintage 2 come from the same segmentation strategy (e.g., FICO band 700-759, similar internal account score, average monthly card balance), yet their patterns are clearly different. Vintage 1 peaks early in its life cycle, shows a small upward spike in MOB5-MOB7, then progressively tapers. Vintage 2 peaks abruptly at two points, then drops. One possible explanation could be:

a) Vintage 1 was established just prior to the Christmas holiday season with a promotional campaign starting four months in.

b) Vintage 2 was established during the period just prior to the 2008 economic downturn and peaked at that point before dropping abruptly over the next few months due to many accounts going into charge-off status or forced repayment.

This is one of several plausible explanations. Intuition is often employed at this point to adjust the curve, with considerable likelihood of misinterpretation. From a practical standpoint, the true interpretation is that some (or all) of the problem issues identified above have not been addressed.

Formulating an Alternative Methodology

The objective is to craft a methodology that augments *intra*vintage performance analysis with *inter*vintage analysis – recognizing it as a continuum in which time series techniques play a critical role. This objective expands the statistical tool base to include a variety of analyses that can be linked in a sequential fashion and unified through reconciliation techniques. Given this foundation, we can approach each of the aforementioned problems as follows:

Issue 1: Vintages are not independent.

Approach:

- Examine performance across vintages, looking at the progression of key metrics (e.g., first month's balance, DPD90 rate, etc.) from month to month.
- Align MOBx values with the calendar along with internal/external covariates that influence performance across vintages.
- Provide a unique time series forecast for each MOBx value, including covariates. This effectively transposes the single (MOB1 – MOBend) view into multiple views (MOB1, MOB2, …MOBend).

Issue 2: Incorporating lagged effects from month to month.

Approach:

- Time series methodologies allow for the inclusion of lagged variables, transfer functions and autoregressive effects.
- Consider the use of relevant covariance structures for linear and nonlinear mixed modeling within a vintage. This helps to refine the shape of the curve during maturation and provides a measure of constraint during subsequent reconciliation processes.

Issue 3: Factoring in trend, seasonality and business cycles.

Approach:

- Time series methodologies inherently allow for identification of these factors.
- Business cycles can be accounted for with the use of covariate proxies (internal/external) that can be projected into the future and used to shape the vintage forecast.

Issue 4: Avoiding disproportionate weighting of older vintages.

Approach:

- Cross-vintage time series analysis assigns primacy to more recent vintages.
- Vintage sensitivity to changing business conditions can be modeled.

Issue 5: Covariate effects are embedded within the curve and are difficult to isolate.

Approach:

- Time series decomposition techniques can be employed to assess component contributions to the predictive model.
- Variable reduction techniques can be used to identify critical factors of influence and to reduce multiple collinearity.

Issue 6: Controlling for in-vintage variation.

Approach:

- Establishing the profile of the vintage early in its life cycle and carrying that forward within the model allows for natural constraints on remaining periods.
- Advanced time series clustering techniques can be used to identify highly irregular patterns as potential outliers that can skew forecasts.
- Additional measures of variability can be used as attributes during vintage profiling.

Issue 7: Controlling for non-homogenous segmentation.

Approach:

- Dynamic clustering techniques can be used.
- Advanced vintage profiling techniques can be used to quantitatively identify vintage characteristics. This approach could involve, for example, forecasting key metrics (e.g., projected account base, first month balance, month-to-month variation, etc.) as an initial step in forming the curve.

Methodology Flow

The general approach is a five-phase process:

| Prepare Vintage Data | → | Forecast Drivers, Trends and Seasonality | → | Build Vintage Clusters | → | Model Current and Future Vintage Curves | → | Reconcile Curves |

Figure 2: The five-phase methodology.

1. **Prepare Vintage Data:** All relevant *in-vintage* information (historical vintage performance, internal/external drivers and events) is captured, segmented using business rules, aligned with the calendar, assessed for relevant influencing variables and transposed into a hierarchy of time series.
2. **Forecast Drivers, Trends and Seasonality:** Internal/external drivers (covariates) are isolated and forecasted along with trends and seasonality across *all vintages* within a segment.

3. **Build Vintage Clusters:** All vintage curves within a segment are grouped for similarity of curve shape and augmented by similar vintage characteristics. These curve clusters form the basis for modeling.

4. **Model Current and Future Vintage Curves:** Driver forecasts are merged with the detrended/deasonalized time series and predictions for all vintage MOBx values are produced. This takes into account values of future drivers (either for upcoming or existing vintages) developed via forecasts or preplanned activity (e.g., APR changes, promotional offers, etc.).

5. **Reconcile Curves:** The final MOBx forecasts are transposed back to vintage curve format and reconciled with bank-defined constraints for vintage performance.

Phase 1: Prepare Vintage Data

The objective of this step is to explore individual account-level data within each vintage, isolate relevant variables, establish a segmentation strategy, and prepare for time series, cross-vintage analysis. An example of a data outline follows:

Metric to be forecast:

- Active balance (combination of purchase, payments, adjustments and fees) at month's end.

Possible segmentation identifiers:

- Primary channel (online vs. in bank).
- Acquisition channel (direct mail vs. Internet).
- Behavior profile (retail transaction vs. cash advance).
- Sector (target audience group: sports, college, professional, AAA, etc.).
- Product (HELOC, credit card, etc.).
- FICO (credit score).

Account-level identifiers:

- Average daily balance.
- Active status.
- Delinquency status.
- Average payment.
- Total available credit.
- Promotional offer in place.
- Payment rate.
- Delinquency rate.
- Homeowner status.
- Initial balance.
- Gender.
- Age group.

- Fees.

- Income.

- Marital status.

Step 1: Account-level data mining.

1. This step is designed to extract relevant account-level information that could be related to performance of the metric. SAS® Enterprise Miner™ can be used to impute missing values and conduct primarily regression/correlation analysis on classification and continuous variables within the account. This can be done using a high-level partitioning structure (e.g., primary channel, product and FICO) initially.

Step 2: Collapse account-level data.

1. This step isolates the selected identifiers and establishes subsegments (classification variables), summarizes continuous variables and establishes percentages for selected class variables (e.g., percentage homeowner) by MOBx value.

2. A necessary preliminary step is to decide on the lowest level of granular segmentation needed for business analysis. This step is required to avoid obvious inaccuracies in subsequent grouping (clustering) steps (e.g., different products, different promotional strategies, etc.). Note that the process will produce predictions for each desired business segment; however, it is not necessary that this same level of segmentation be adhered to at this point.

Step 3: Align vintages with the calendar.

1. This step transposes the vintage-level view (MOBx) into a date-specific view based on vintage start date. As an example, MOB1 for the August 2012 vintage would be coded as Aug2012. MOB2 for the same vintage would be coded as Sep2012, MOB4 for the September 2012 vintage would be coded as Dec2012, etc.

2. This step allows for cross-vintage time series analysis and can easily be performed using SAS DATA step coding and PROC TRANSPOSE.

Phase 2: Forecast Drivers, Trends and Seasonality

The objective of this forecasting phase is twofold:

1. Forecast cross-vintage drivers (internal/external) that determine short- and long-range business trends and activity.

2. Standardize the comparison of vintage curves in subsequent clustering steps.

Step 1: Prepare drivers for forecasting.

1. This step develops time series data sets (primarily via PROC TIMESERIES) for subsequent hierarchical forecasting of the drivers using SAS Forecast Server. Anticipated macroeconomic forecasts and expected bank policy drivers can be included in advance as future covariates.

Step 2: Forecast drivers.

1. This step executes SAS Forecast Server using the prepared data. The expansive functionality of the SAS Forecast Server engine is used to develop scenarios and model combinations. Forecasts are performed within the specified hierarchical structure.

Step 3: Merge forecasts into aligned vintage data.

1. This step incorporates outputs from SAS Forecast Server into the vintage-level data that has been aligned with calendar. Some of these forecasts may be considered as time invariant (i.e., fixed throughout the life cycle of the vintage) or time varying (changing with the MOBx value). This determination is based primarily on whether the covariate is considered a profile variable or continuous variable.

Step 4: Develop trend and seasonal forecasts.

1. This step builds a weighted (in the case of ratio-based metrics) cross-vintage forecast for purposes of determining trend and seasonal factor adjustments that will be applied to individual vintages prior to subsequent curve-clustering activities.
2. SAS Forecast Server is used in conjunction with PROC X11 to build these factors, which are stored for later use.

Phase 3: Build Vintage Clusters

The objective of this phase is to establish a modeling foundation for new or evolving vintages. Because vintages are autonomous, those that have yet to be initiated require a structured method (see Leonard, Trovero, et al.) for establishing a historical proxy that can serve as a starting point. The outcome of this phase will be a series of segmented curve clusters that represent assigned points for vintages to be modeled.

Step 1: Detrend and deseasonalize curves.

1. Trend and seasonal factors from the prior phase are applied to each curve.

Step 2: Conduct curve similarity analysis.

1. PROC SIMILARITY is used to produce a matrix of curve shape similarity measures between detrended/deseasonalized vintages.
2. Quantitative profile measures (based on the account-level identifiers determined in Phase 1) are used to augment the similarity matrix.
3. PROC DISTANCE is run on the augmented matrix, thus producing a distance matrix that can be used as input into PROC CLUSTER. The analyst can opt for the desired number of clusters using PROC TREE. Typically, this is constrained to a manageable number.
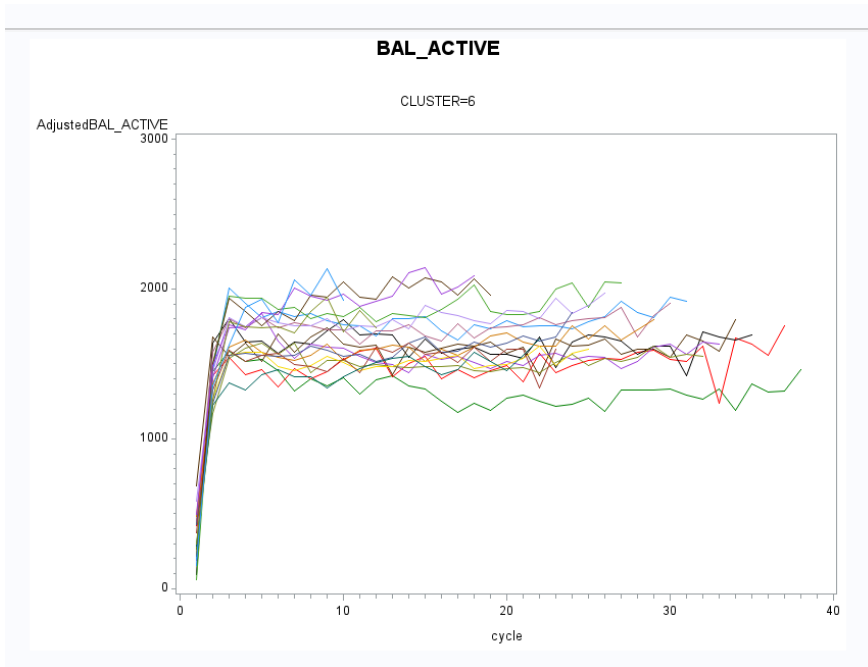
Figure 3: Sample curve cluster. This figure superimposes approximately 15 standardized vintage curves that collectively define "Cluster 6." The cycle (x-axis) equates to the MOBx value.

Step 3: Assign vintages to clusters.

1. Each vintage to be modeled must have an assigned cluster. This assignment can be accomplished via a combination of manual selection and automated techniques that incorporate business rules and judgment.

2. In the case of the illustrative data (documented in Phase 1), the "hard" segmentation incorporates primary channel, product and FICO. This is the BY layer of the analysis. Vintages to be modeled include this layer plus acquisition channel, behavior and sector, which can be assigned to different clusters within the BY group. After building clusters, the analyst can determine the level of homogeneity within each cluster (e.g., whether the cluster includes several overlapping segments) and use that as the basis for assignment.

3. SAS DATA step and SQL procedures can be used to associate clusters with vintage key identifiers.

Phase 4: Model Current and Future vintages

The objective of this phase is to build predictive models for each existing or planned vintage. The data prepared in Phase 1 has been augmented with time variant (MOBx) and invariant covariate predictions from Phase 2 along with deterministic business information (e.g., planned promotions) and clusters from Phase 3. The next step is to model the critical period for the vintage.

Step 1: Model critical vintage period.

1. The exact length of the critical period (e.g., first MOBx values) can vary based on the metric; however, it is assumed to be no longer than 18 months. After this point, time activity enters a more steady state period that can be modeled with basic time series techniques.

2. The critical period must overlap the period where covariate forecasts were developed.

3. The nature of the data allows for the use of a mixed modeling technique that uses fixed effects (e.g., MOB), covariates (e.g., initial balance, promo, etc.) and random effects (e.g., each vintage) in a repeated measures study with a user-defined covariance structure, typically AR(1) or MA(1).

4. The MIXED procedure is used to estimate the model parameters for each BY group/cluster combination (i.e., all vintages within the same cluster are used in the model). Those that are new or incomplete will have data points with missing metric values for future MOBs; however, they will contain the necessary classification and future covariate values. The following code presents an example with subsequent output:

```
Proc mixed data=LIB.MIXEDMODELCLUSTER;
by PrimaryChannel FICO cluster;
class keyname MOB;
model Adjusted&driver = MOB VintageExpectedBalance AvgAPR AvgIncome
 / alpha =.05 outpredLIB.MIXEDMODELPREDICTION  singular = 1e-6;
repeated / subject=keyname type=ar(1) r rcorr;
ods output Solution=LIB.MIXEDMODELSOLUTION covparms=cov rcorr=corr;
run;
```

Figure 4: Mixed model sample code. The above code generates mixed model predictions for each cluster. Each vintage is uniquely identified (keyname) and treated as a random subject. The AR(1) covariance structure is one of several that can be used to define month-to-month covariance within the vintage.

Solution for Fixed Effects									
Effect	Months on Book - based on CRACCT_DW open_dt	Estimate	Standard Error	DF	t Value	Pr > |t|	Alpha	Lower	Upper
Intercept		-7.1969	191.19	56	-0.04	0.9701	0.05	-390.19	375.80
MOB	1	-764.22	36.2378	56	-21.09	<.0001	0.05	-836.81	-691.62
MOB	2	-259.07	36.2378	56	-7.15	<.0001	0.05	-331.66	-186.48
MOB	3	-115.16	36.2378	56	-3.18	0.0024	0.05	-187.75	-42.5620
MOB	4	-112.26	36.2378	56	-3.10	0.0030	0.05	-184.85	-39.6669
MOB	5	-85.4164	36.2378	56	-2.36	0.0219	0.05	-158.01	-12.8233
MOB	6	-75.2135	36.2378	56	-2.08	0.0425	0.05	-147.81	-2.6205
MOB	7	-51.7716	36.2378	56	-1.43	0.1587	0.05	-124.36	20.8214
MOB	8	-46.4645	36.4915	56	-1.27	0.2082	0.05	-119.57	26.6368
MOB	9	-39.6369	36.9058	56	-1.07	0.2874	0.05	-113.57	34.2943
MOB	10	-43.6121	37.5585	56	-1.16	0.2505	0.05	-118.85	31.6266
MOB	11	-35.9367	38.8261	56	-0.93	0.3586	0.05	-113.71	41.8413
MOB	12	-62.0533	41.0577	56	-1.51	0.1363	0.05	-144.30	20.1951
MOB	13	0	.	.	.	.	.	.	.
VintageExpectedBalan		1.3265	0.2417	56	5.49	<.0001	0.05	0.8423	1.8106

Figure 5: Mixed model parameter estimation. The above figure shows an example of parametric coefficients for each of the MOBx values in addition to the fixed covariate "VintageExpectedBalance." These estimates can be used to score new vintages.

Step 2: Model remainder of vintage period.

1. After the mixed model generates predictions for the first "x" MOB values, several options for modeling the remaining period can be considered:

   a. The mixed model can be extended to cover the remaining MOBs. However, the lack of historical data may be problematic.

   b. A simple smoothing factor based on adjacent historical period values can be used.

   c. Time series forecasting techniques can be used on the individual vintage under the assumption that the critical period has adequate history (typically 18 months) to determine the pattern for remaining months. In this case, ARIMAX modeling in SAS Forecast Server can be used under the assumption that intervention variables would be limited to those found in the time series (e.g., promotions).

Phase 5: Reconcile Curves

The objective of this phase is to assemble all vintage curves (historical and projected values) into a cross-vintage hierarchical display that allows for hierarchical reconciliation, cross-vintage comparison and adjustments.

Step 1: Establish a visual display area.

1. This step uses SAS Forecast Server to establish a baseline forecast project that arranges the vintages in a user-specified hierarchical structure. The technique developed by Burness and Roehl (2013) is used to overlay forecast outputs from the baseline projects with predictions from the mixed model. This technique allows for display of predictions and confidence intervals for future vintage even though no history exists.

Step 2: Establish a hierarchical reconciliation method.

1. Most metrics to be modeled are in the form of a ratio (e.g., balance per account, delinquency rate, etc.). Current reconciliation techniques within SAS Forecast Studio provide the means to reconcile using simple averaging. However, this imposes an unrealistic constraint of equally sized groups during vertical reconciliation. Ratio reconciliation techniques (see Leonard, Trovero, Dillman, Elsheimer, 2013) external to the application can be developed and applied to each level of the forecast, most likely using indirect techniques that assign weights to the forecasts based on the size of the account base.
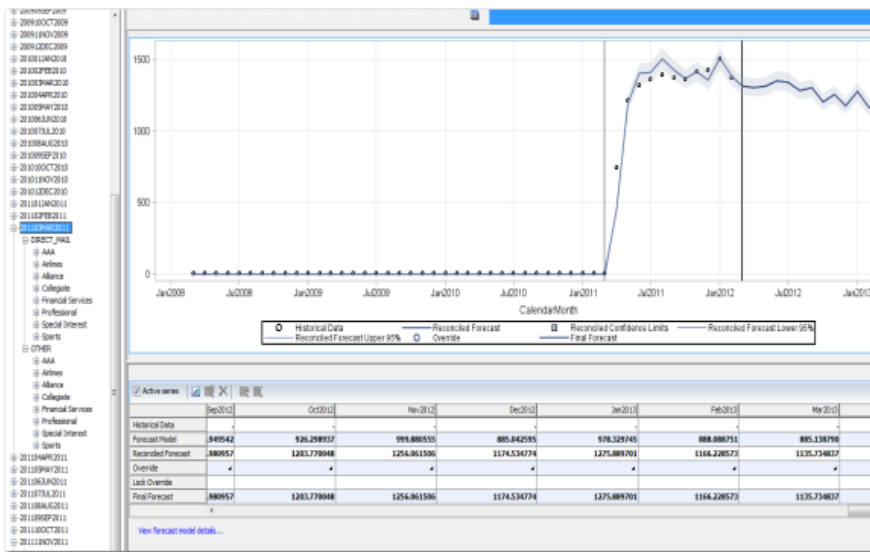

Figure 6: Example of reconciled forecast. The above figure shows an example of a SAS Forecast Studio project. The vintage starts in March 2011 with the first 12 months' predictions (solid line in between the two vertical lines) compared against the actual data points (circles). This reflects a weighted reconciled estimate based on the size of the other vintages below it in the hierarchy.

Step 3: Incorporate adjustments.

1. SAS Forecast Server allows for the incorporation of overrides that can be propagated throughout the hierarchy using various user-driven constraints for application. For example, should there be a need to increase projections for a given segment by 5 percent over the course of the forecast period, the metric for all vintages within that segment would be proportionately scaled to reflect this change. Those that are exempt from this adjustment (determined by business decisions) can be excluded from this process.

Caveats

1. The methodology assumes the generation of vintages at regular intervals.

2. The methodology does not assume full automation. Some degree of manual intervention is required for cluster assignment as well as initial segmentation analysis.

3. The mixed modeling methodology requires a rigorous assessment of the covariance term in the REPEATED statement.

4. Reconciliation weighting procedures will require eventual alignment across the SAS product line.

Conclusion

This paper has presented a methodology that treats the components of a vintage as a set of predictions and forecasts from a *cross-vintage data stream*. It allows for the use of a variety of SAS techniques, including time series analysis, clustering, vintage profiling and reconciliation, thus employing an integrated approach to vintage curve modeling that blends internal drivers and external factors with past performance.

References

Leonard, Michael, Trovero, Michele, Dickey, Thomas, Guseman, Samuel and Celestino, Robert. *A New Product Forecasting Process: Combining Statistics and Judgment.* SAS white paper, 2007.

Leonard, Michael, Trovero, Michele, Dillman, Peter and Elsheimer, Bruce. *Devices for Forecasting Ratios in Hierarchies.* Patent Pending. US Patent Client Ref No: 53649, Attorney Docket No. 94926/860712, 2013.

Burness, Geoffrey C. and Roehl, Nitzi M. *An Automated Approach for Integrating Externally Developed Forecasts into SAS Forecast Studio.* SAS white paper, 2013.

Way, Robin. *Credit Card Industry Best Practices for Revenue and Loss Forecasting.* SAS white paper, 2009.

About SAS

SAS is the leader in business analytics software and services, and the largest independent vendor in the business intelligence market. Through innovative solutions, SAS helps customers at more than 60,000 sites improve performance and deliver value by making better decisions faster. Since 1976 SAS has been giving customers around the world THE POWER TO KNOW® For more information on SAS® Business Analytics software and services, visit **sas.com**.