

Retail Loan Portfolio Dynamics

Becoming a Better Vintner

by Joseph L. Breeden

Maturing portfolios should result in fine wines, not bad whines. To help ensure they do, some organizations are using new techniques for quantifying the drivers of portfolio performance. These new techniques analyze retail loan “vintages,” groups of accounts originated in a given time period.

Surprise is the enemy of any business. Stock value, securitization, lending capital, strategic planning, and government oversight all rely on portfolio visibility. Unexpected shocks to receivables or losses are always punished by investors and regulators.

Yet, eliminating surprises is not simply a matter of publishing more data. The difficulty is in extracting cause-and-effect from performance data to explain why the portfolio is performing as observed. Is it the economy? A new originations score? The natural maturing of the portfolio?

The lack of such information not only worries investors and regulators, it also forestalls effective portfolio management. In the fall of 2000, many risk managers were observing degradation in their portfolios and wondering “Is it us, or is it the economy?” Months were lost before the bankruptcy shock became apparent as a real industry-wide event. That time should have been used for action rather than self-doubt.

At the onset of the Asian economic crisis of 1997, many financial institutions suffered from terminal indecision. Retail loan portfolios in Asia had grown rapidly with aggressive expansion into lower-income consumer

segments and new products. As such, they were untested in an economic crisis.

When the Thai interest rate shock began, officials denied that it could spread through the economy. On May 10, 1997, IMF Managing Director Michel Camdessus said, “I don’t see any reason for this crisis to develop further.”

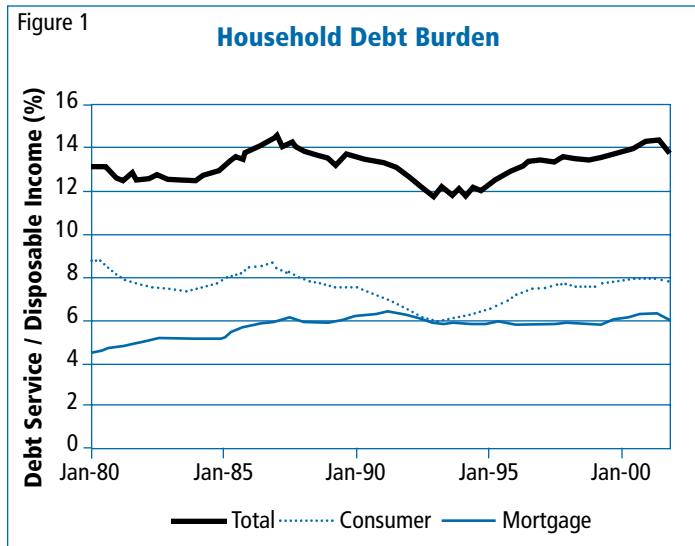
In the spring of 1997, banks saw unusual activity in their portfolios, but could not tell whether the Chinese New Year, lax originations, or economic impacts were at fault. The result was management committee meetings consisting mostly of finger-pointing and no concrete action. At some organizations, five months passed before any concrete actions were taken—far too late to be effective.

Even as the Federal Reserve declares the U.S. recession of 2001 over, bankers are aware that consumer debt burden is at all-time highs (Figure 1). Banks navigating these rocky shoals need tools that will make the dangers clear with sufficient time for correcting course.

Many different causes operating on different time scales drive retail loan portfolios. Originations

quality, account maturation, seasonality, management actions, competition, and the macroeconomic environment are some of the factors affecting portfolio

performance. This article reviews the most important effects, common ways of looking for them, and some new approaches to understanding portfolios.



Originations Quality



Credit bureau scores are the dominant method for assessing quality at origination. However, many factors, such as negative selection and application systems failures, can create surprises. To monitor the quality of new originations, the most common approach is called *vintage analysis*.

A retail loan vintage is a group of accounts originated in a given time period. For example, all new customers from 2001 make up the 2001 vintage. Vintage groupings can be monthly, quarterly, or annual, depending on the application and the amount of data available. The OCC's ALLL guidelines call for lenders to consider vintage analysis at least intuitively when studying homogenous loan pools.

Like a bottle of wine, consumer vintages mature with age. By plotting multiple vintages together against months-on-books, we can make comparisons for similar maturities.

The difficulty with vintage plots is that effects other than originations quality can influence the result. Changes in the economic environment, product policies, collections systems, and more can mask the result being sought. In the hypothetical example in Figure 2, can we conclude that the 2001 vintage is intrinsically higher risk and lower quality?

Recent changes in the economic environment could be at least partly responsible for the worsened performance. The sudden rises seen in the 2000 and 1999 vintages suggest environmental impacts could be affecting all vintages. With such ambiguity, quantifying cause-and-effect is a daunting task.



Vintage Maturation

Vintage analysis as just described is intended to provide a visual method for adjusting vintage performance for the maturation process. The maturation process itself is often valuable information. Figure 3 shows the maturation curves for two consumer segments. Aside from exhibiting different

overall delinquency risk, the segments differ in when the risk is greatest. Such differences have implications for timing loss reserves and collections staffing.

The shapes of the maturation curves can vary by consumer segment and product. Policies specific to an institution can introduce unique features in the curves. For example, credit cards that carry an annual fee will show a spike every 12 months in the maturation curve for attrition. Quantifying the maturation curve is an essential starting point for understanding a portfolio.

The standard approach for measuring a maturation curve is to align the historical vintages by months-on-books and compute the average at each point. This will provide a rough estimate of the maturation process.

When Originations Change

The difficulty with averaging algorithms is that they cannot adapt to changes in originations or the environment. Consider a case where originations quality steadily improved over a four-year period. An average across vintages would not represent the true maturation process because vintages of different risk are being observed at different months-on-books.


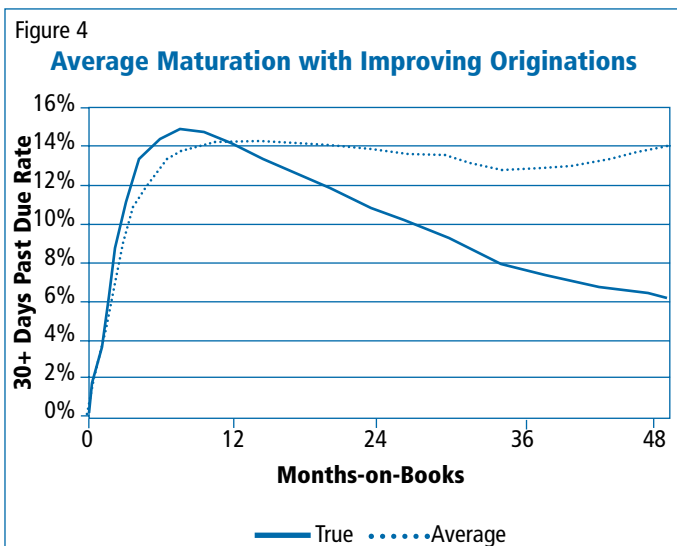
Figure 4 shows a hypothetical example built with an improving originations scenario, from which we can see that the averaged curve does not decrease as it should. Relying upon a distorted curve such as this could lead management to a false prediction of future loss levels and overly pessimistic loss reserves. If the maturation curve in Figure 4 were used to generate an 18-month delinquency forecast, vintage delinquency rates would be overestimated by an average of 49%. The total portfolio loss forecast would be overestimated by 66%. In practice, few portfolios are so seriously wrong, but this is due to the use of intuition and hand-crafting of the forecast. As an automated tool, this approach falls far short. Actual performance will follow the true curve, independent of the average.

The maturation curve estimation process is analogous to a vintner sampling old bottles of wine to see what his new vintage will taste like when it matures. If the older vintages were made from sour grapes, he might sell the vineyard before realizing the quality of his new wines.

When the Environment Changes

Changes in the environment can distort the averaged maturation curve as well. Using an environmental scenario similar to what has occurred through the recent recession, the averaged maturation curve would again be distorted similarly to what was shown in Figure 4. With a scenario that the environment will improve over the next few years, the averaged maturation curve would be as shown in Figure 5.

The practical consequence of Figure 5 is that data obtained in the years following the bottom of a recession (2002 to ?) can mislead the analyst into an overly optimistic view of how the portfolio will mature. Conversely, studying data from 1999 through 2001 could give an overly pessimistic view of how the portfolio will mature.



These examples demonstrate that a maturation curve obtained through a months-on-book average of vintages is useful for understanding the portfolio, but has questionable accuracy in detailed portfolio forecasting.

Macroeconomic Environment

All bankers understand that the economy drives their businesses. If the economy grew at a steady rate with no booms or busts, performance could be predicted with precision. The variability of the vineyard would be replaced with the reproducibility of the brewing tank.

Measuring the impact of the environment is arguably the most challenging of portfolio modeling tasks. The two most common approaches are *total portfolio analysis* and *vintage level analysis*. A few organizations go a step further to computing *maturation residuals*.

Total portfolio analysis is the simplest and most widely used. This once-standard approach of directly comparing total portfolio performance to macroeconomic variables by way of econometric models has been widely discredited. Changes in segment mix, changes in credit policies, and the maturation of the portfolio all contribute to confound the estimation of macroeconomic impacts.

A more effective approach to measure economic impact is to compute a vintage-level average. By aligning the vintages by calendar month rather than months-on-books and computing the average, we can easily eliminate biases from portfolio growth.

Some organizations take this analysis one step

further. The vintage average is still vulnerable to biases due to portfolio maturation. By computing the vintage residuals relative to the maturation curve (maturation residuals) and averaging those for each month, we can estimate a time series of environmental impacts.

$$\text{Average residual for month } t = \text{average across all vintages at month } t \text{ of } (\text{vintage performance} - \text{expected maturation value})$$

Few organizations actually compute the vintage residuals because of the possible instabilities in the estimate. In a portfolio with changing originations criteria or product policies, distortions will appear that may mask economic impacts.

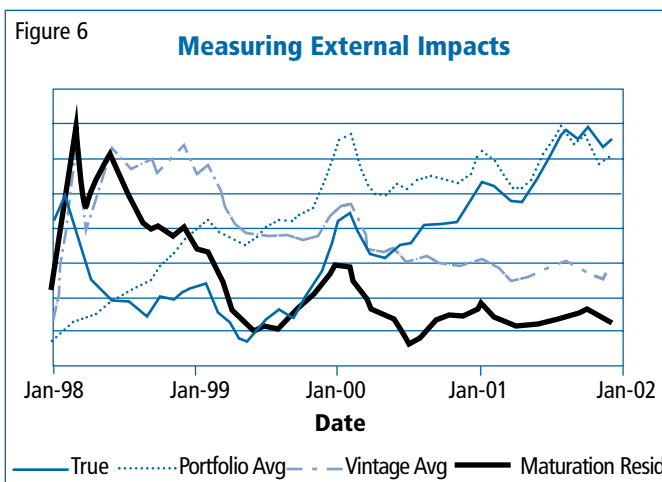
Figure 6 shows the results of these three approaches for a hypothetical portfolio. It shows the growth of a product launched in January 1998 and measured with quarterly vintages. A real-world example of environmental impact was included to show the rise in risk through 2001 and the onset of the recession (solid line). As the environment degraded, new bookings were curtailed and originations criteria were tightened. This scenario, while hypothetical, rings true for a number of portfolios. Comparing the three approaches, which does the best job of finding the external impact?

The total portfolio approach looks best until you realize that this is just a lucky accident. The example was of a new and growing portfolio; thus, losses just happened to coincide with a worsening environment.

Neither of the other two approaches would have even noticed the recession. Although the maturation residuals seem like a good idea, the poor quality of the maturation estimation (Figure 4) makes the maturation residuals one of the worst results. Since long-term trends are crucial for macroeconomic comparisons, none of these approaches reproduces the environmental impacts sufficiently to support econometric modeling.

Even going outside the portfolio to industry averages can leave management in doubt. Has the industry shift to subprime lending or aggressive pricing for prime lending contributed to an observed rise in delinquency?

During the 1995 Mexican peso crisis, loan rewrites were extensive in Argentina—so much so that it became difficult to determine when



the crisis might be over. Was the dramatic drop in slow debt just due to the rewrites, or was the economy improving? In retrospect, it became clearer that rewrites had created an artificial lull, but not in time to assist policy decisions. Figure 7 shows the actual slow debt rate for consumer lending in Argentina. The dashed line represents what might have happened without the rewrite programs.

Recognizing the rewrite wave in the Argentina data is crucial to making econometric models for that market. In the absence of that industry insight, the double peak in slow debt leads to misleading correlations to macroeconomic data. Industry data alone is not a complete answer for understanding external impacts on a portfolio.

Seasonality and Other Shocks

Portfolio analysis cannot be complete without a detailed study of seasonality. Intuitively, it is the simplest component of portfolio performance to understand. Credit card use always rises at Christmas time. Delinquency usually falls during tax refund season.



The most common way to view seasonality is to make a cyclic or polar plot versus month of the year. Figure 8 shows an example of delinquency rate in a cyclic plot. This example shows the growth of a new portfolio from 1997 through 2001. The natural maturation of that portfolio introduces an overall trend in the data. Further, events in 2001 create a dramatic rise. Figure 8 does show the customary tax-time drop in delinquency, but the end-of-year drop is suspicious. Creating an average measure of seasonality is somewhat helpful qualitatively, but not quantitatively.

For this example, where portfolio maturation transitions smoothly into a recession, the worst distortions can be removed by first *detrending* the data. Detrending is always a good idea in computing seasonality, but it is not always completely effective. The complexity driving many portfolios introduces more instability than a polynomial or moving average detrending can correct.

The final step in understanding portfolio performance is to analyze consumer response to credit policy changes, systems changes, and other management actions. Ideally, these effects should be quan-

tified by maintaining control groups. During any change, some subset of the portfolio should be managed under the old policies so that response specific to the new policy can be observed separately from maturation, seasonality, and the environment.

Unfortunately, accidents happen. Failures in payment processing or collections systems can cause unintended shocks to the portfolio. Some policies are enacted with a “no-going-back” attitude, making the impact universal. In such cases, the only available approach is to understand in detail all the fundamental drivers of the portfolio first. Shocks due to unique events will become apparent only after stripping away the known effects of maturation, origination changes, seasonality, and environmental impacts.



Solutions

Is there no better way to predict the quality of a bottle of wine? Well, with wine, perhaps not, but the situation in consumer lending is not hopeless.

Data. The simplest way to get better answers is to collect more data. Vintage-level data is essential, but more of it is needed. The industry standard seems to be to save only two to four years of performance data. When archiving account-level data, these limitations are natural. However, analyzing portfolio performance does not necessarily require account-level data. All of the preceding analysis can be done from vintage-level aggregated data, preferably monthly vintages. For a typical portfolio, even 10 years of monthly vintage cohorts with monthly performance reporting take on the order of tens of megabytes of data.

Data storage in relational databases at the account level is a leading reason that institutions do not do more vintage-level portfolio analysis. Portfolio-wide analysis requires touching every account in the database, which is the worst-case scenario for relational database performance. For portfolio analysis, a separate database with a time-series orientation is preferable. Recent developments suggest that this vintage data is even sufficient to support IRB models under the proposed Basel II Accord.

Additional data is useful in portfolio analysis because the portfolio can be observed under a broad range of conditions. The simple approaches described earlier fail most dramatically for portfolios measured only through part of a cycle—an economic downturn or a trend to tighter originations.


Observing the same portfolio through several changes in originations policy or several economic cycles allows the trends to be averaged away.

New technologies. Additional data, while a long-term goal for many organizations, is unavailable in the short term. In such cases, more sophisticated analysis is required. The techniques described in the first half of this article are basic. Practitioners have experimented with numerous refinements and can do better than the estimates shown. Nevertheless, the underlying weaknesses persist.

Some organizations are starting to use new techniques that have been developed recently for quantifying the drivers of portfolio performance. *Decomposition approaches* attempt to simultaneously separate maturation and exogenous effects, such that the components are independent and unbiased. While analogous to principal component analysis in linear regression, decomposition uses an approach that is adapted to the nonlinearity of the components and the different time dimensions of months-on-books and calendar time.

Such nonlinear decomposition techniques leverage research in scientific fields involving nonlinear dynamical systems and nonlinear partial differential equations. Rather than treating each vintage as an independent forecasting problem or just looking at the total portfolio, the vintages are used as probes. Each vintage measures how consumers mature with time and how consumers respond to their environment. The added complication is that the probes are uncalibrated, because each vintage varies in composition.

These new techniques analyze the data from the vintages as if it represented a surface, as, say, the surface of a pond. Changes in the environment ripple across the vintages. The maturation process appears as a large wave running down the vintages in a different direction. Nonlinear partial differential equations provide a natural framework for analyzing such overlapping waves. The major components of portfolio performance (originations quality, maturation, and environmental impacts) are extracted in a mutually independent fashion. The result is a view of the maturation process without the distortions of originations or the environment, quantification of originations quality independent of the environment, and a



better measure of environmental impacts than is available from industry-wide data. Most importantly, this analysis can be performed on the short, noisy datasets common among retail lending institutions.

Analysis of commercial loan portfolios has been revolutionized over the last decade through the introduction of sophisticated mathematical techniques. Because of the fundamental differences between commercial lending and retail lending, commercial portfolio management techniques have a poor track record when applied to retail loans. Using techniques from nonlinear time-series analysis and partial differential equations, retail portfolios can now be analyzed with the fidelity prevalent in commercial lending.

Figure 9 shows a re-analysis of the same data shown in Figure 6 extracting the external impacts. This result shows how external effects are extracted independently of maturation and originations quality to produce a much higher fidelity result than any of the standard algorithms. The error bars are largest in the first portion of the curve, where there are very few vintages.

There are numerous, substantial benefits to be realized from this new technology. Since the real, underlying maturation curves can be isolated for specific portfolio segments, portfolio caps and score cutoffs can be set so that the long-term profitabili-

ty of the total portfolio can be maximized. By isolating the response of the portfolio to macroeconomic and other secular developments, we can make accurate estimates of portfolio response to stress scenarios. Since the impact of past management actions can now be estimated accurately, effective contingency plans can be developed and the identification of specific situations can be made earlier in the process as these situations unfold.

This technology also improves the accuracy of loss forecasting, which, in turn, helps ensure that adequate reserves are maintained and that corrective actions are taken well in advance of deteriorating conditions. Improved loss forecasting also facilitates the allocation of capital to loan portfolios and can either free excess capital for use elsewhere within the organization or can identify undercapitalized situations. Finally, it can serve as an independent check on business forecasts and alert senior management to situations where managers may be either too optimistic or too pessimistic regarding their expectations.

Wouldn't a vintner love a recipe for how the grape, the weather, and the cask combine to produce a bottle of wine? In retail lending, that is becoming a reality. □

Contact Breeden at
breeden@strategicanalytics.com



Be in Good Company!

JOIN Mark Hillis, Jack Goldstein, Marc Intrater, Garry Barnes, Robert Davidson, Harry Glenos, David McLeod, JAMES ZEIGER, Annette Larabee, Carl Horsch, Darcy Walker, Dev Strischek, Sean Keenan, Konrad Alt—bankers who write for the *Journal*. Share best practices, lessons learned, risk management innovations, practical advice, opinions, and more with other readers.

E-mail bfoster@rmahq.org or call 215-446-4096