SOCIETY OF ACTUARIES®

Article from

The Financial Reporter

September 2018
Issue 114

Using Predictive Modeling for UL Premium Assumptions

By Emily Cassidy and Frankie Logan

Life insurers face many challenges when valuing universal life (UL) business on a US GAAP basis due to the product's flexible nature to suit individual policyholder's needs. One common challenge is setting premium persistency assumptions on the in-force block of business. The presence of a secondary guarantee or a rider can greatly affect policyholders' future premium payment behavior. Additional factors like the insurer's methods for notifying policyholders of potential for their contract to lapse can also have an effect on payment behaviors. While the base US GAAP reserve on a UL contract is the account value, the deferred acquisition cost (DAC) asset will be influenced by the premium persistency assumption. Premiums are not a revenue item under ASC 944 (previously FAS 97), but other elements such as percent of premium loads and commission expense are dependent on the assumption and will flow through to the cash flows.

Insurers have historically struggled to perform experience studies on their premium persistency and to set a best estimate assumption with some degree of confidence. The difficulty of capturing some of these effects causes insurers' best estimate assumptions to be inaccurate which leads to earnings surprises and risk management problems. Predictive analytics is starting to gain traction as a solution for this challenge. Techniques such as generalized linear modeling (GLM), deep learning, Markov modeling, random forests, and clustering can help analyze insurers' data about policyholders and their prior behavior to better project short-term and long-term premium payments for each policyholder. Using predictive analytics, the premium assumption can vary by duration and be assigned at a seriatim level, providing a more accurate depiction of payment behavior.

We present in this article a case study in which machine learning was used to develop the premium persistency assumption for an in-force block of UL business.

DATA CONSIDERATIONS

Like all other modeling techniques, the results of a predictive analysis model are only as good as the underlying data. Selecting the data that will ultimately be used to create a model is crucial to building a good model. In addition, accuracy, ease of annual update, ease of implementation, and applicability are also key considerations when using predictive analytics.

Gathering, scrubbing and structuring the data to build a predictive model has a cost. Predictive power, availability, IT cost and implementation are some of the key factors to consider when collecting data. Internal data such as historical premium payments, product type and characteristics, funding intentions, funding level, historical grace letter/funding notices and payment methods (e.g., automatic withdrawal) are all good information to use when setting premium persistency assumptions. An insurer might also consider using external data such as interest rates and unemployment levels to supplement the internal data. Lastly, demographic and other personal information about the policyholder can be used to create a richer model. Once the data is gathered, a data model would need to be created to begin analysis. Documentation and validation are extremely important as the process will need to be repeated during the annual unlocking process.

The data model is often split into a training set, which is used for building the model, and a holdout set, which is used for validating the model. The purpose of splitting the data model is to avoid overfitting, i.e., matching the model parameters too closely to the data. Overfitting the model will cause the model to perform poorly as new data is introduced. For example, a 60-40 split can be applied—60 percent of the data is the training set and 40 percent is the holdout set. This ratio can be adjusted to find the best balance.

When splitting premium behavior data, users can consider splitting the data model by number of policies (e.g., 60 percent of the policies will be used as the training set and 40 percent will be used as the holdout set) or by calendar year (e.g., 2000–2014 as the training set and anything after as the holdout set). Both methods have advantages and disadvantages. Splitting the data model by number of policies can give the model an opportunity to learn the entirety of the policyholder's behavior to date. However, it can be difficult to apply economic factors into the model as economic factors might not affect every policy in the training set the same way (e.g., policies ended before the financial crisis versus policies during or after the financial crisis). Splitting the data model by calendar year makes it easier for the user to apply economic factors and cycles into the model. It is also a good way to validate the model since the goal is to predict future payment behavior. A disadvantage is that the mix of new and existing policies in a calendar year can influence the payment behavior, though this effect can be lessened by adding policy duration into the model.

Other data considerations include how to track the response variable in comparison to future experience, handling of "early" or "late" payments, segmentation of the model, deployment of assumptions for use in valuation, and IT infrastructure.

MODELING APPROACH

One of the first steps in modeling is defining a response variable. A response variable is something we want to predict, or measure, with the model. For this case study, the ratio of paid premium to target premium is selected as the response variable. Performing an exploratory analysis such as a one-way analysis can create a distribution profile for each of the potential variables and can serve as an indicator of which variables are strong predictors.

The historical premium payment information showed that large proportions of policyholders were either paying their target premium or making no payments so a two-part model can address the different behaviors in policyholder. A two-part model includes one part to indicate whether an event has occurred and the second part to indicate the size of the event. The advantages of taking a two-part approach are having the option to include predictive variables in either the first or second component, it is easier to communicate, and it provides a greater understanding of the business.

In this case study, the two-part model includes the policyholders' payment pattern (i.e., how policyholders behave in a given policy year) and their new planned premium.

We established four options for premium payment: paying target, steady payment not equal to target, change in payment and no payment. Paying target means the target premium is paid. Steady payment not equal to target means the premium paid is equal to the prior duration and is non-zero. Change in payment means premium paid in the current duration is not equal to the prior duration and is non-zero. No payment means zero premium is paid in the current duration. In this case study, the data showed over two-thirds of policies have no change in premium pattern. Most policies tend to stay in either the "paying target" or "steady payments" state for several durations. Less than 15 percent of policies showed zero payments in the previous duration and over half of those policies continue to pay no premium. Around 10 percent of the policies showed a change in payment amount.

The planned premium is the assumption for premium paid in the current year for policies that have made a change. This is expressed as a percentage of target. There are several ways to decide this amount. For example, the average and median of the policies' payment change in the different state transitions can be used in determining the percentage of target. Using median, and reducing the effect of large outliers, we see that those paying target will pay slightly less than target in the next period; steady, but not target, will pay about 15 percent over target in the next payment; changing in payment will pay around 5 percent over target in the next payment; and no payment will pay very close to target in the next period.

MODELING PROCESS

Now that the dataset was built and predictive variables were selected, we used a random forest to model premium behavior.

A random forest is a classification and an ensemble learning model. An average is taken from a number of decision trees. One property of random forests that users should be aware of is that the random forest decreases **variance** in the results but it doesn't decrease the **bias.** Using a random forest, we can directly predict which state a policyholder is in and which state they will move to in the next period. One of the challenges surrounding random forest modeling is that numeric variables like issue age need to be grouped into a range. Exhibit 1 shows an example of how individual decision trees work. Shaded boxes represent a policy characteristic while non-shaded boxes represent a classification.

Exhibit 1
Decision Trees Example

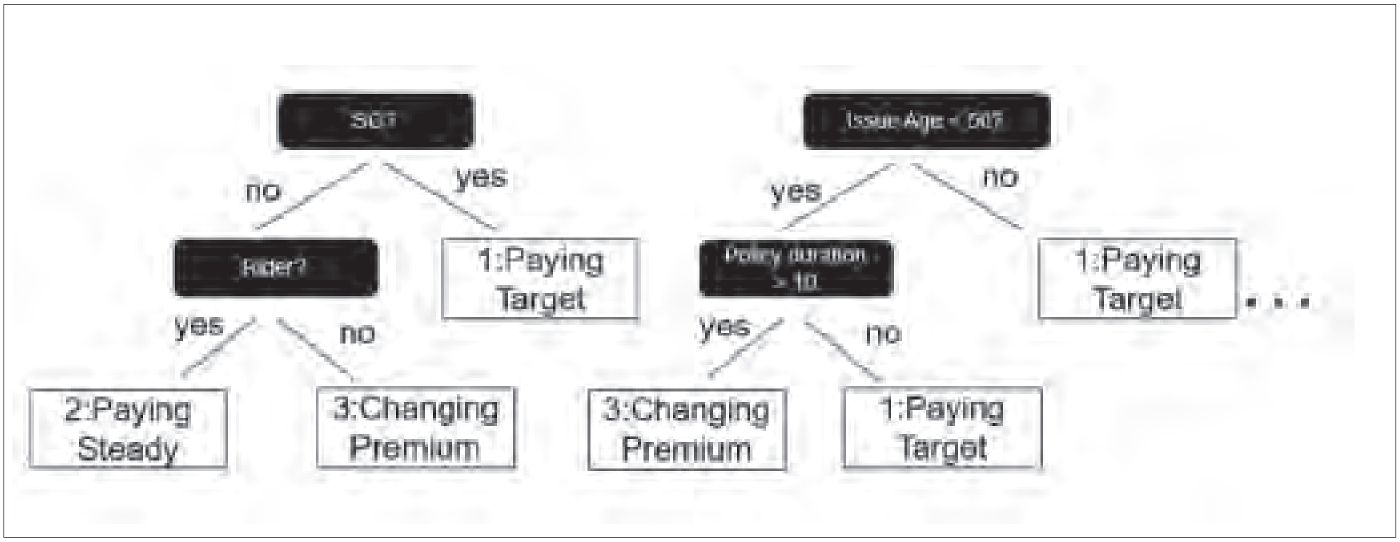

To create the model, we use the training set of the data model. Once the initial model is created, we can input our holdout data set to test how well the model is performing. There are many model performance matrices that we can use to check the performance of the model such as gains/lift chart, logloss chart, and receiver operating characteristic curve (ROC curve). Additional matrices such as confusion matrix and area under the curve (AUC) can be produced based on the ROC curve. A gains/lift chart measures the effectiveness of the predictive model compared to not using a model at all. Logloss is a metric that penalizes the model for having the wrong classification. An ROC chart is a graph that uses the true positive rate and false positive rate to test performance of classification models. We can then adjust the hyper-parameters (the number of decision trees in the forest and the number of features considered by each tree when splitting a node) of the random forest and/or the predictive variables to include in the model.

ASSUMPTION SETTING AND IMPLEMENTATION

Based on the information we gathered, a new premium assumption can be set for each of the states (paying target, paying steady, changing payment and no payment). Now that we have our model and assumptions, we can input seriatim policy data into our model and the output will be a seriatim assumption for future premium payments. With our model, we were typically able to predict future premium payments within 2 percent of the actual payments.

All of this work would be wasted if there was not an efficient way to implement the assumption. Some of the implementation options include deploying the premium assumption methodology dynamically within the valuation software or calculating the premium "upstream" (in a statistical software like R) and passing them to the valuation model. The first option allows the use of more sophisticated and up-to-date premium predictors while the second option provides more flexibility in analytical methods.

Additional considerations when using predictive analytics include the model capabilities and limitations, model size and processing time, the use of dynamic variables in the predictive model, the frequency of updates, and the assumption validation and quantification of financial impacts. ■

Emily Cassidy, FSA, MAAA, is a manager at KPMG LLP. She can be reached at *emcassidy@kpmg.com*.

Frankie Logan is an associate at KPMG LLP. He can be reached at *flogan@kpmg.com*.