Article from:

Risk Management

August 2011 – Issue 22

Is Predictive Modeling the Answer?

By David J. Weinsier and Guillaume Briere-Giroux

LIFE INSURERS THAT WRITE VARIABLE ANNUITY (VA) BUSINESS WITH GUARANTEES FACE A VARIETY OF SIGNIFICANT CHALLENGES.

The recent financial crisis put hedging programs and, in certain cases, the statutory solvency of VA writers to the test and demonstrated that future financial success for VA writers will rely on sustainable product pricing, accurate hedging and robust risk management. Each of these actions depends on the insurer's ability to study, forecast and properly manage policyholder behavior risk.

In this Insights article, we explore how important tools used in property & casualty (P&C) insurance—predictive modeling and data mining—can be applied to more effectively model policyholder behavior risks in VA contracts. Traditional modeling approaches have attempted to reflect VA policyholder behavior patterns based on product design, policy characteristics and policy performance. However, in practice, policyholder behavior is driven by numerous interrelated factors.

David J. Weinsier, FSA, MAAA, is a consultant in the Towers Watson Atlanta office. He specializes in mergers and acquisitions, reserve financing solutions, indexed products and life insurance mortality. He leads the firm's Predictive Modeling for Life Insurers initiative.

Guillaume Briere-Giroux, FSA, CFA, MAAA, is a consultant in the Towers Watson Hartford office. He specializes in the pricing, risk management and valuation of products with capital market-based guarantees, such as variable and indexed annuities.

Many of these factors are difficult to account for under traditional approaches, which typically consider only a limited number of variables and fail to adequately capture certain correlations and interactions among them.

In this article, a case study is used to demonstrate how a predictive modeling approach can improve upon traditional methods used to model VA lapse behavior.

TRADITIONAL APPROACHES TO MODELING VA LAPSE BEHAVIOR

Figure 1 on page 11 describes the primary factors that drive VA lapse behavior and indicates whether traditional modeling approaches typically reflect each factor. We have categorized the factors into four groups:

- Product and guarantee design
- Policy characteristics
- Policy performance
- Distribution.

What is Predictive Modeling?

Predictive modeling is the application of certain algorithms and statistical techniques to a data set to better understand the behavior of a target variable based on the co-relationships of several explanatory variables.

Rather than relying on a simple understanding of basic risk elements, predictive modeling enables the user to consider many confounding factors simultaneously by mining across a set of scenarios. This analysis permits more informed decisions and limits the use of subjective judgment.

Predictive modeling techniques have primarily been used in the P&C space to enhance understanding of current and/or future insured risks.

This knowledge has led to improved risk segmentation, underwriting, pricing and marketing decisions. For example, auto insurance premiums reflect the fact that younger drivers are poorer risks than middle-aged and older drivers, and males are poorer risks than females. However, data also show a clear interaction between age and gender (i.e., the difference in relative risk between male and female drivers is much less pronounced at older ages than at younger ages). Traditional pricing techniques typically do not quantify this interaction between risk parameters, but a predictive model will recognize this and other interactions, enabling the insurer to develop premiums that accurately reflect the relative risk characteristics of the pool of underlying policyholders.

Figure 1. Factors that Drive VA Lapse Behavior

Category	Factor	Traditional Industry Modeling Practices
Product and guarantee design	Surrender charge length and strength	Reflected via grading up base rate, shock, shock + 1 and ultimate lapse rates
	Share class (A-share, B, C, L)	Reflected in accordance with specific surrender charge schedule
	Presence and nature of living benefits	Reflected, but approximate and somewhat speculative
Policy characteristics	Policy size	Typically not reflected
	Policyholder age and sex	Typically not reflected
	Life stage (i.e., accumulation vs. income)	Typically not reflected
	Qualified vs. nonqualified	Typically not reflected
Policy performance	Guaranteed benefits in-the-moneyness	Reflected via deterministic formulas applied uniformly to base lapse rates
	Recent fund performance	Typically not reflected
Distribution	Commission structures (heaped vs. trail)	Typically not reflected for in-force modeling (often reflected in pricing, however)
	Distribution channel/target market	Typically not reflected (beyond what is captured in aggregate experience)

Traditional approaches to modeling VA lapse behavior can have the following shortcomings:

- **Inability to distinguish between base and dynamic behavior.** Historical data will show a single lapse rate, which is a function of both base behavior and dynamic behavior. However, the use of traditional approaches makes it challenging to identify which component of the single aggregate rate is base and which is dynamic. When attempts are made to separate these impacts, the credibility of the resulting groups decreases. Thus the impact of these separate pieces cannot be precisely validated.
- **Suboptimal use of historical experience data.** In a typical experience study, the data are categorized, aggregated and analyzed. By splitting the data into categories, the exposure bases available to analyze a given relationship (e.g., policy year effect for a particular product) become smaller, which results in a loss of credibility. Aggregating the experience for a given variable does not control for the contribution of other variables influencing the experience for that group. This creates "noise" that increases the amount of data required to extract a credible relationship when analyzing a single variable at a time.
- **Traditional approaches typically consider a limited number of explanatory variables to account for a complex behavior.** This is often done to maintain the credibility of the results. In fact, many of these variables are readily available (e.g., age, gender, asset allocation, past withdrawals), but others could be categorized as "exotic" variables that could also be collected and analyzed to help predict VA lapse behavior (e.g., indicators of financial sophistication such as credit score, education levels, profession/industry).
- **Interactions between variables, where the impact of one variable is affected by a second variable, are typically not captured.** Consequently, these methods fail to consider explanatory variables and their impact on the target variable.
- **Correlations between explanatory variables are not fully accounted for,** which can result in double-counting effects or not attributing an effect to the correct variable.

CONTINUED ON PAGE 12

Is Predictive Modeling the Answer? | from Page 11

CASE STUDY: APPLICATION OF PREDICTIVE MODELING TECHNIQUES TO VA LAPSE BEHAVIOR

A predictive model can address many of these shortcomings by permitting consideration of all risk factors simultaneously, in addition to reflecting interactions between variables, without significantly reducing the credibility of results. This allows for both a macro view and a focus on the subtle, micro-interactions between risk factors. Specifically, predictive modeling enables management to better understand the factors influencing policyholder behavior, the interaction of such factors and the potential impact on profitability and risk.
The results of this case study show how the application of a predictive model to modeling VA lapse rates can improve on traditional approaches. The underlying analysis was performed on a large sample of hypothetical but representative data. The data were developed based on actual industry experience, with certain adjustments, resulting in an exposure base and product mix representative of a typical VA writer. The resulting data set features a typical age, share class, fund allocation, commission type and rider mix by year of issue (and includes more than 10 issue years).

The in-the-moneyness (ITM) for the living benefit riders (e.g., GMWB, GMIB) is representative of actual historical market conditions, including actual experience in the tumultuous years of 2008 and 2009.

The **traditional model** employs a typical industry approach to modeling VA lapse rates, reflecting the following factors:

- Base lapse rate varying by policy year
- Surrender charge length and strength
- Shock lapse at the end of the surrender charge period
- Commission structure
- Presence and nature of living benefits
- ITM of living benefits, defined as: 1 – (account value/benefit base).

The **predictive model**, derived as a generalized linear model (GLM), is based on the following variables present in the case study data:

- Base rate varying by policy year
- Surrender charge length and strength
- Proximity to end-of-surrender charge
- Commission structure
- Presence and nature of living benefits
- ITM of living benefits
- Premium (i.e., policy size)
- Fund value
- Portfolio mix (aggressive, balanced, conservative, cash)
- Attained age.

MODEL VALIDATION

The data set was randomly split into two distinct groups in order to facilitate an objective model validation. The first group, made up of 70 percent of the aggregate data set, was used to set the model parameters. The second group, the remaining 30 percent of the aggregate data set, was then used to test how effectively the model predicted actual lapse behavior. That is, the first group of data was used to fit the models. These models then projected an expected set of lapse rates for the policies in the second group (the "E" in an actual-to-expected study). The actual lapse experience in the second group was then designated as the "A" to see how well the models predicted actual results.

CASE STUDY RESULTS

Figure 2 shows actual-to-expected results by policy year, while Figure 3 shows results by ITM bands.

The predictive model shows an appreciably better fit than the traditional model when considering actual-to expected ratios by policy year and ITM bands. This result is primarily driven by correlations between policy year and ITM that are captured in the predictive model but ignored by the traditional model.

The comparisons of actual-to-expected lapse rates on an aggregate basis shown in Figures 2 and 3 are useful; however, additional comparisons and analysis should be performed to verify this result. Figure 4 on page 14 compares expected lapse rates emerging from the traditional model to the predictive model. This allows for a comparison and validation of the fit of the two models at more granular levels. The x-axis is the ratio of the predictive model expected lapse rate to the traditional model expected lapse rate.

A ratio of 1.0 indicates that the two models produce the same lapse rate for a given policy. A ratio less than 1.0 indicates that the predictive modeling approach produces a lower lapse rate than the traditional model, whereas a ratio greater than 1.0 indicates that the predictive model produces a higher rate.

This comparison tells us that, for a significant proportion of the policies, the two models produce very different expected lapse rates. The absolute difference in the ratio is greater than or equal to 20 percent for 65 percent of the policies and greater than or equal to 60 percent for 23 percent of the policies. As depicted on the far right side of Figure 4, this analysis also shows that for roughly 3 percent of policies, the predictive model produces a rate greater than or equal to 3.0 times the traditional rate, suggesting that the traditional model may have limitations in capturing the tails.

Figure 4 shows that expected lapse rates differ significantly between the models at the policy level, and further analysis is needed to test the viability of the predictive model at a granular level. For this purpose, we developed a typical graph commonly referred to in the P&C space as a "gains chart," as portrayed in Figure 5 on page 14. A gains chart sorts the policies by expected lapse rate in descending order. The cumulative lapse rate is then recorded as the data are stepped through policy by policy.

Figure 2. Actual Versus Expected Lapse Rates by Policy Year

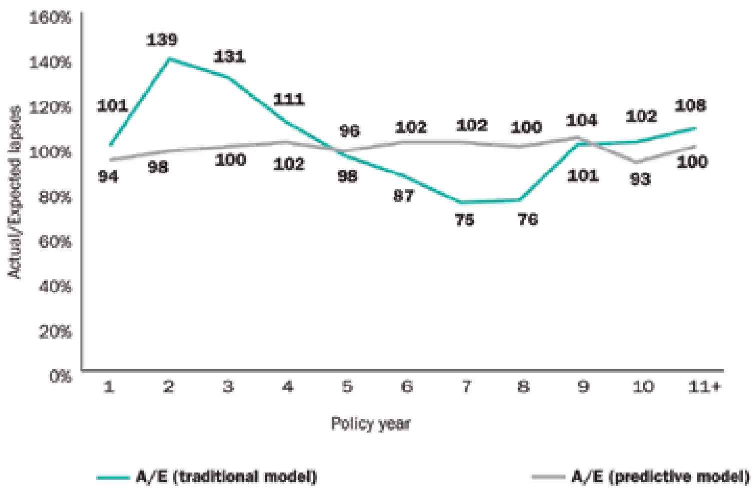

Figure 3. Actual Versus Expected Lapse Rates by ITM Bands

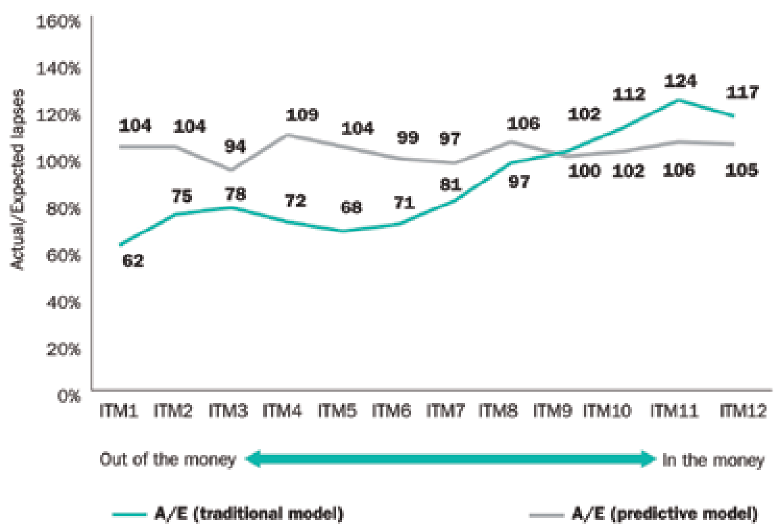

CONTINUED ON PAGE 14

Is Predictive Modeling the Answer? | from Page 13

Figure 4. Comparison of Predictive Model to Traditional Model Expected Lapse Rates

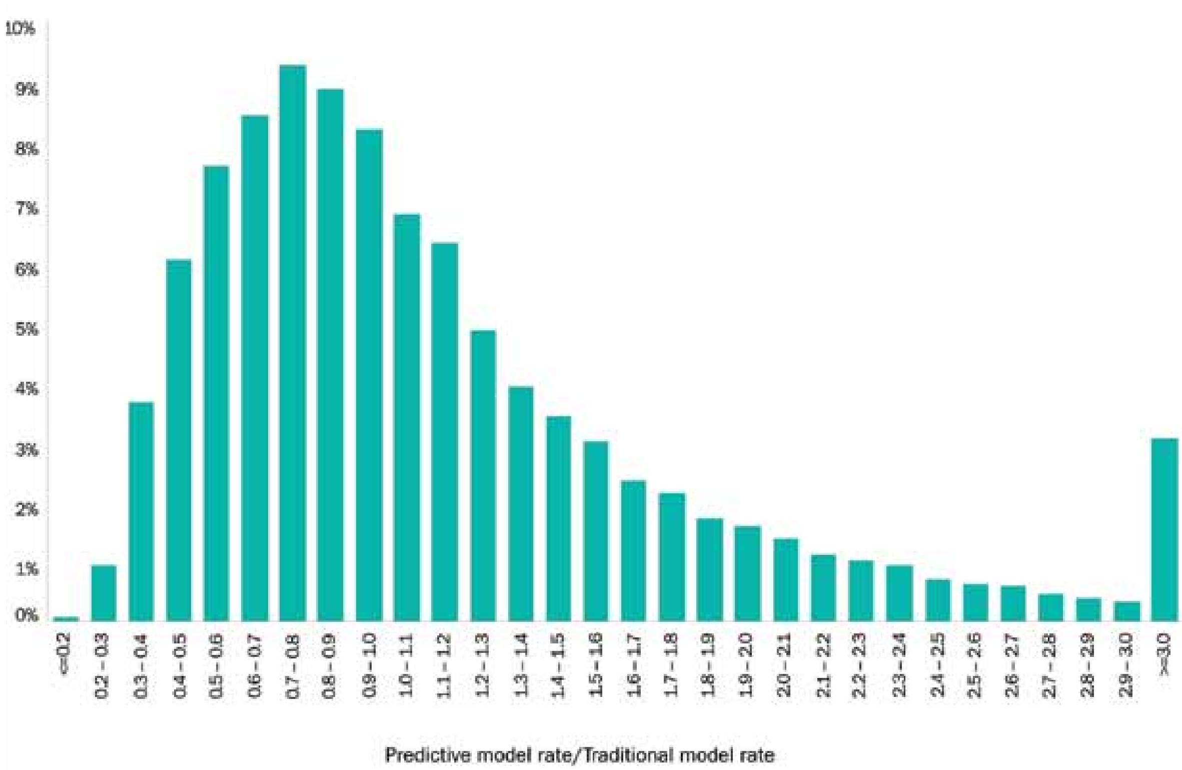

Figure 5.
Comparison of Traditional and Predictive Models Using a Gains Chart

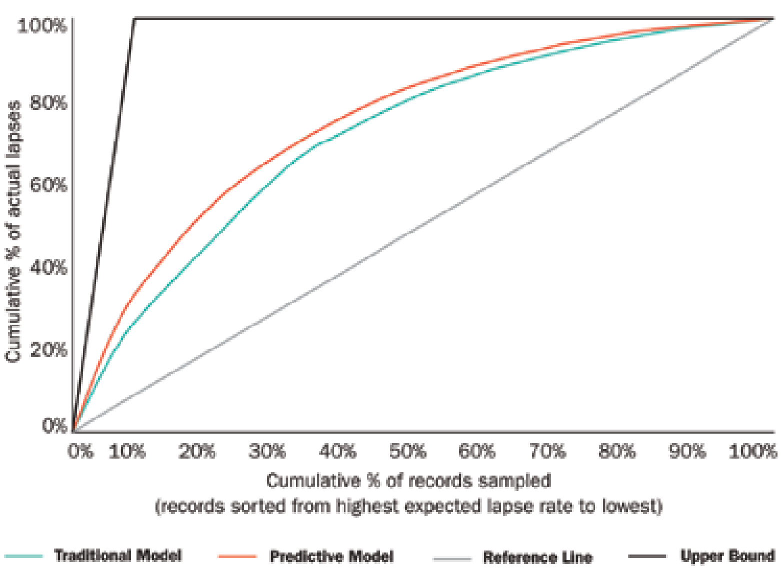

By definition, before the first record, the cumulative percentage of the total number of lapses will be 0 percent. At the end of the projection, it will be 100 percent. If the model is no better than a random sort of the data, then we would expect a straight diagonal line that we label the reference line (gray line in Figure 5). In this case, 50 percent of the lapses have been found (y-axis) after sampling 50 percent of the records (x-axis). At the other extreme, a perfect model would have predicted 100 percent of the lapses in roughly the first 8 percent of records (8 percent is the average annual lapse rate). This is labeled as the upper bound (black line).

Since the model is better than a random sort, we expect the cumulative percentage of lapses to increase more quickly than the cumulative percentage of records counted, and the line produced on the graph to be bowed to the left. The greater the area under the model line, the better the model is able to differentiate policies by risk of lapsing. The graph shows, for example, that if the first 20 percent of policies are targeted, the predictive model (red line) would have predicted roughly 55 percent of actual lapses, as compared to 45 percent for the traditional model (blue line), indicating a stronger model.

CONCLUSIONS

Predictive modeling and data-mining techniques commonly used in the P&C space can be applied to effectively measure, analyze and forecast complex VA lapse rate behavior. The results of the case study showed that, as compared to the traditional model, the predictive model achieved an appreciably better fit under a typical actual-to-expected analysis, produced a more granular fit, and better differentiated between policies with a low and high risk of lapsing.

The overall assessment is that, compared to traditional approaches, the predictive model can improve modeling of VA lapse behavior because it can:

- Capture a greater number of risk factors (or variables) that drive VA lapse behavior
- Account for correlations between explanatory variables; in the case study, the predictive model was able to obtain a better fit due to its ability to disentangle the effect of ITM and policy year
- Make optimal use of the data available by avoiding segmenting and grouping, which can result in a loss of credibility; the predictive model uses less data to achieve convergence
- Capture interactions between variables, where the impact of one variable is affected by another.

AUTHORS

David J. Weinsier, FSA, MAAA
David Weinsier is a consultant in the Towers Watson Atlanta office. He specializes in mergers and acquisitions, reserve financing solutions, indexed products and life insurance mortality. He leads the firm's Predictive Modeling for Life Insurers initiative. Mr. Weinsier is a Fellow of the Society of Actuaries and a Member of the American Academy of Actuaries.

Guillaume Briere-Giroux, FSA, CFA, MAAA
Guillaume Briere-Giroux is a consultant in the Towers Watson Hartford office. He specializes in the pricing, risk management and valuation of products with capital market-based guarantees, such as variable and indexed annuities. Mr. Briere-Giroux is a Fellow of the Society of Actuaries, a Member of the American Academy of Actuaries and a CFA charterholder.

Why use Predictive Modeling?

The use of predictive modeling by life insurers can lead to the following business and strategic benefits:

- Identification of more profitable segments, distribution and target markets
- More reliable pricing assumptions, less subjectivity and reduced assumption risk
- Product development based on more accurate estimates of policyholder behavior (e.g., surrender rates, withdrawal/annuitization utilization, asset allocation/rebalancing)
- Improved risk mitigation (e.g., hedging, asset/liability management) by reducing policyholder behavior variances
- More accurate modeling of policyholder behavior in the tail, resulting in more accurate reserve and capital estimates
- More streamlined models and better controlled model implementation by replacing multiple tables and dynamic formulas with a single parameterized predictive model
- Easing compliance with certain regulatory, rating agency and reporting requirements (e.g., Actuarial Guideline 43, Solvency II, MCEV principles).

For comments or questions, contact David Weinsier at +1 404 365 1781, david.weinsier@towerswatson.com or Guillaume Briere-Giroux at +1 860 843 7083, guillaume.briere-giroux@towerswatson.com. Recognition is owed to Jean-Felix Huet, Rob Spaul and Andy Staudt for their assistance in developing this article.

ADDITIONAL RESOURCES

Towers Watson Society of Actuaries Research paper Predictive Modeling for Life Insurers—Application of Predictive Modeling Techniques in Measuring Policyholder
Behavior in Variable Annuity Contracts.

NA-2010-18812