

MAKING CLAIMS

APPLICATIONS OF PREDICTIVE ANALYTICS IN LONG-TERM CARE

BY ROBERT EATON AND MISSY GORDON

Predictive analytics has taken far too long in getting its foothold in the long-term care (LTC) insurance space, but it promises to have a lasting impact now that it has arrived. This article explores the current LTC predictive analytics landscape. We examine two case studies demonstrating the use of predictive analytics tools in assumption-setting for LTC claims. Finally, we take to the whiteboard and plot out ways in which predictive analytics may be used in the future.

BACKGROUND

2015 SOA Tables

In July 2015, the Society of Actuaries (SOA) completed its report “Long-Term Care Experience Basic Table Development”¹ (the SOA Tables). The SOA research team used predictive analytics to develop models for claim incidence, claim termination and benefit utilization. The model for LTC incidence and claim termination is a generalized linear model (GLM) using a log-link function and Poisson error structure (for benefit utilization, a Tweedie error structure is used). These types of models generally produce a friendly, multiplicative format that is easy to understand.

2016 Predictive Modeling Workshop

The Intercompany Long-Term Care Insurance (ILTCI) conference is widely attended among LTC actuarial professionals. This makes the conference a natural setting for hosting a workshop on predictive analytics. In March 2016, the researchers who developed the SOA Tables brought their software, models and expertise to San Antonio, Texas, to educate LTC professionals on predictive analytics. The workshop included a course laying out a theoretical framework for predictive analytics followed by a four-hour hands-on predictive modeling workshop. The workshop had attendees partitioning data between the model and validation, selecting and fitting models, interpreting results, validating assumptions and comparing the predictiveness of models.

CASE STUDY NO. 1

Following the workshop, we went in search of our own predictive analytics problems. LTC claim termination assumptions typically are expressed by monthly claim duration. Carriers may assume that monthly claim terminations vary by sex, site of care (i.e., a nursing home, assisted living facility or home care), benefit period (i.e., the number of years the policyholder may claim) and potentially other attributes. The claim termination assumption in an LTC model can be used to calculate an expected length of stay (LOS) on claim.

We wondered if other policyholder attributes may also help to describe the claimant LOS. In particular, would the quality of an LTC facility, or the quality of care provided in the claimant’s area, help us in understanding how long an insured would remain on claim?

We began with a data set for one carrier that included the claim LOS, along with relevant policyholder attributes and benefits. In addition to this claim data, we used the ZIP code associated with the insureds’ current residences at the time of claim to find both the average income in that ZIP code as well as an area factor from the Milliman Health Cost Guidelines (HCGs). We didn’t have provider quality ratings for all of the sites of care used by claimants, but we ventured (and for this analysis, assumed) that the average income by ZIP code and the HCGs area factor might serve as proxies. We also included an area factor-adjusted average income, which may remove some of the distortions from high cost of living areas. This creation of a new variable is a good example of feature engineering—an important step in predictive modeling. Our

ultimate goal was to understand which of these variables or features of the policy mattered most in determining the LOS. See **FIGURE 1** for more details on these policy attributes.

FIGURE 1: ATTRIBUTES USED TO ESTIMATE LTC CLAIM TERMINATION AND LENGTH OF STAY

Traditional	All Available
Age at claim	Average income
Benefit period	Average income (area factor adjusted)
Sex	Age at claim
Site of care	Area factor
Tax-qualified status	Benefit period
	Daily benefit amount
	Incurral year
	Inflation type
	Issue age
	Marital status
	Policy duration at claim
	Product generation
	Risk class
	Sex
	Site of care
	Tax-qualified status

To help estimate LOS, we applied the random forest and variable importance modeling techniques to two sets of variables: traditional descriptors of LTC length of stay, and a collection of all policy attributes we had on hand, whether or not they traditionally are used to model claim terminations and LOS. We chose

these predictive analytics techniques because they are robust, handling both quantitative and qualitative variables, and the results are straightforward to apply to future modeling efforts.

The table in **FIGURE 2** shows the rank of variable importance of the five attributes from our initial random forest, where the variable importance measures have been scaled to sum to 100 percent.

FIGURE 2: VARIABLE IMPORTANCE FOR TRADITIONAL ATTRIBUTES

Variable	Rank	Variable Importance
Site of care	1	52%
Age at claim	2	38%
Tax-qualified status	3	6%
Benefit period	4	2%
Sex	5	2%

It’s clear that site of care and age at claim are our most important predictors of LOS for the data in this case study. Our next step was to feed our random forest model each of our 16 prediction variables, to see which variables ranked the most important. The table in **FIGURE 3** shows the results of the variable importance algorithm for all attributes.

FIGURE 3: VARIABLE IMPORTANCE FOR ALL ATTRIBUTES

Variable	Rank	Variable Importance
Average income	1	15%
Age at claim	2	14%
Issue age	3	13%
Area factor	4	12%
Site of care	5	12%
Average income (area factor adjusted)	6	11%
Incurral year	7	8%

Daily benefit amount	8	4%
Policy duration at claim	9	4%
Inflation type	10	2%
Risk class	11	2%
Tax-qualified status	12	2%
Benefit period	13	<1%
Product generation	14	<1%
Sex	15	<1%
Marital status	16	<1%

The blue-shaded rows show where our initial five predictive variables ranked within the group. Right away, we see that certain variables, such as area factor and the average income in a claimant's area, may have similar predictive power as the age at claim in determining the LOS. Other variables that may be useful in this analysis, which weren't conducive to our data set, include diagnosis and coverage type (e.g., facility-only or comprehensive coverage).

There are many things to note about these results. To the extent that certain variables are correlated (e.g., issue age and age at claim), we should expect similar variable importance. However, this may not imply that we should develop models based on all variables with a high importance measure. The actuary should consider carefully the variable selection task so as to produce a model with meaningful predictors that also can be explained to management. We also should consider whether or not our proxy variables are doing a good job of representing the underlying fundamentals (in this case, estimating provider quality). Furthermore, the actuary also may consider which of the variables ultimately will be useful and available if the intent is to incorporate the assumption into other projection systems and models. Finally, the initial findings from this case study are relative to the data set we used. It's likely that other sets of claim data may produce different drivers of predictiveness for LOS.

A model of LOS that includes new predictors may be used, for example, to analyze claimants currently in the disabled life reserve. The actuary may be interested in understanding the impact of considering provider quality, or (proxy) affluence of a pool of disabled lives. This new view of the predictors of LOS may cause the actuary to consider, for example, high-level adjustments to disabled life reserves, particularly for companies with limited historical claims experience, and/or to revisit the claim termination assumption. The next case study outlines how a company may use predictive analytics to do just that.

CASE STUDY NO. 2

Traditionally, experience studies adjust a starting benchmark, whether it be an industry benchmark or prior assumption, for company-specific or updated data. We wondered if there was a way to use predictive analytics techniques to assist in making these adjustments, and so we researched applying such techniques to industry data.

One key question we had was: How can we incorporate the idea of credibility when using predictive analytics to update this assumption? We tip our cap to Brad Armstrong and Shea Parkes for shining light on a method for doing so in their article “Calibrating Risk Score Model with Partial Credibility,” which was published in the SOA’s *Forecasting and Futurism* section newsletter.² This article presents the idea of recalibrating a risk adjustment model for a company that had a limited amount of data to warrant full recalibration solely based on the company’s experience. Their approach to solve the limited data issues was to start with a benchmark model that was developed on a much larger data set. They then made adjustments to that model to better fit the company-specific experience, where adjustments were credible. This was done by using a penalized regression with the benchmark model as an offset in the regression.

A penalized regression is a GLM with an extra constraint that penalizes the coefficients in the model. This penalty can be thought of as a “credibility lever.” In this case, a large penalty would give essentially no weight to the company-specific data, leaving the benchmark unchanged. On the opposite side of the scale, a small penalty would give considerably more weight to the company data and potentially produce large adjustments to the benchmark. Hugh Miller discusses more details on this process in an article that shows the equivalence of penalized regression methods and credibility theory.³

After finding a way to incorporate credibility, our next step was to determine an approach to do so when developing a claim termination assumption. One common approach is to turn the claim termination assumption-setting process into a survival analysis problem. A nice trick we can perform under this setting is to use a GLM with a log-link and Poisson error structure to approximate the Cox proportional hazard rate model.⁴ Doing so allows for two important things. First, it allows us to incorporate an offset into the model, which serves as our benchmark starting assumption. Second, the nature of the Poisson assumption allows us to aggregate our data to the level of unique covariates in the model, which decreases the runtime to fit a model. The resulting coefficients from the model are then multiplicative adjustments that are applied to the benchmark rates.

We have been using this modeling framework (with the addition of penalization) in our research, and it has proven to be fruitful by producing reasonable and intuitive results. Our major takeaway is that the penalized regression technique is a useful tool that we encourage actuaries to add to their toolboxes. Incorporating the use of a benchmark as an offset in the penalized regression is a great approach to use in experience studies. It provides a robust way to give weight to experience compared to some other methods of applying credibility theory when updating an existing assumption. Giving the “right” amount of weight to the experience is important in producing an assumption that is reflective of the most recent experience, while not over-reacting to new data. This helps to avoid significant variation and generally applies well to other uses.

Additionally, we have found that by using predictive modeling techniques, we are able to update existing assumptions based on statistical concepts using an automated process. This leads to a robust and more transparent, reproducible process. Predictive modeling techniques also allow us to add new variables that we have not been able to use in the past, while simultaneously creating the adjustments. This in turn normalizes the effects of the other covariates, giving us a better understanding of the true relationships that drive the underlying experience. Predictive analytics is a powerful tool that requires great responsibility. We encourage actuaries to explore the application of predictive analytics, but to do so with the guidance of an experienced practitioner.

While our research involved industry data, we expect this modeling framework to be useful to a smaller data set as well, such as a company updating its assumptions. As a next step, we are performing a case study that compares and contrasts the development of a claim termination assumption for one company using traditional methods versus predictive analytics. We will explore the business case for this change in methodology and examine the impact of using predictive analytics. Full results from our case study will be published in early 2017.

PREDICTIVE ANALYTICS SUPPORTING LONG-TERM SERVICES IN THE FUTURE

Predictive analytics has the potential to reach beyond supporting today's actuarial conundrums. The future of the provision and financing of long-term services and support (LTSS) is evolving. Many states are looking to the decades ahead, and they perceive the heavy burden of Medicaid LTSS funding. They are looking for solutions from today's LTC professionals.

In response to that, and as part of a broader national dialogue on LTC, the SOA has sponsored the LTC Think Tank. The LTC Think Tank aims to provide new ideas⁵ to help people pay for long-term care, make care more accessible, reduce the cost of care and mitigate the need for care to start with. Two of the key ideas stemming from the LTC Think Tank are a Healthy Longevity App and an online Care Portal. Another idea in the area of care delivery is the "Uberification" of LTC services—in particular, home care services.

Each of these concepts lends itself to the adoption of predictive analytic techniques: to optimize recommendations from the Healthy Longevity App, to push salient recommendations to the front of the Care Portal, to anticipate when people may need certain home care services and more. Predictive analytics has a long and healthy future in the LTC space, and actuaries will be on the forefront of these advancements.

References

¹ "Long-Term Care Experience Basic Table Development." Society of Actuaries, April 2015. <https://www.soa.org/Files/Research/Exp-Study/2015-ltc-exp-basic-table-report.pdf>.

² Armstrong, Brad and Shea Parkes. "Calibrating Risk Score Model with Partial Credibility." *Forecasting and Futurism*, no. 11. July 2015. <https://www.soa.org/Library/Newsletters/Forecasting-Futurism/2015/July/ffn-2015-iss11-parkes-armstrong.aspx>.

³ Miller, Hugh. "A Discussion on Credibility and Penalised Regression, With Implications for Actuarial Work." August 2015. <http://actuaries.asn.au/Library/Events/ASTINAFIRERMColloquium/2015/MillerCredibilityPaper.pdf>.

⁴ Survival Models, <http://data.princeton.edu/wws509/notes/c7.pdf>.

⁵ "Long-Term Care Think Tank." February 22, 2016. <https://www.soa.org/Files/Sections/2016-03-long-term-care-think-tank.pdf>.

Robert Eaton, FSA, MAAA, is a consulting actuary at Milliman in Tampa, Florida. He can be reached at robert.eaton@milliman.com.

Missy Gordon, FSA, MAAA, is a principal and consulting actuary at Milliman in Minneapolis. She can be reached at missy.gordon@milliman.com.

Special thanks to Joe Long, an assistant actuary and data scientist at Milliman in Minneapolis, for his contribution to our research.