# Get Better Acquainted With Your Known Unknowns

By Dan Kim and Boyang Meng


Uncertainty of a predictive model is a fact of life that many insurers could be overlooking at their peril without a framework for assessing it.

Predictive analytics have become increasingly commonly used across the U.S. life insurance industry in areas such as mortality and policyholder experience analysis, automated pricing and underwriting, in-force management and fraud/claims analytics.

Predictive models are usually better at detecting signals from a large dataset and are more likely to be precise in making predictions compared to traditional approaches such as a tabular or one-way analysis. For example, U.S. life insurers often use actual-to-expected ratios in a tabular form to develop best estimate assumptions. Predictive models, like generalized linear models or tree models, may improve the traditional models, but the danger is that models become regarded as perfect and a silver bullet for decision making within the business.

For, as the renowned statistician George E. P. Box once said: "All models are wrong, but some are useful."

## ZERO ERROR IS A PIPEDREAM

What he was referring to is that while any predictive model will (or should) be built to minimize the generalized error, the error will never practically be zero. So, the question insurers need to think about is how much the future will emerge differently from their predictions. To answer this question, having a framework to determine a level of model uncertainty can be invaluable.

Such understanding matters, because it can be fundamental to things such as whether an insurance applicant that may be below the underwriting criteria is falsely approved, how much capital and reserves companies need to hold, the chances of fraudulent insurance claims making it through the vetting process, and decisions about risk transfer.

## FRAMEWORK GOALS

There are a few key issues to address in such a framework. Actuaries usually ask, "How credible is the data?" Instead, we can expand this into more specific, targeted questions. What is the confidence level of the model's average predictions? How significantly can reality differ from these predictions? The aim is to determine the degree to which your predictions may be uncertain so that you can augment your business strategy to minimize the impact from such uncertainty.

We can apply those questions to the simplest of predictive models—the outcomes of tossing a coin. Let's say 10 tosses of this coin have yielded six heads and four tails. Without knowledge of the fairness of this coin, what could the range of outcomes for 20 tosses be? What about 1,000 tosses?

First, we must develop an assumption or a predictive model on the fairness of the coin. Then, we need to quantify the uncertainty of the model. Finally, given all these, we can understand the range and probability distribution of possible outcomes from more tosses.

An assumption about the fairness of the coin can be illustrated as a probability of showing a head. So, a reasonable assumption, based on experience, is 60 percent. We can estimate the uncertainty of that probability using binomial distribution as shown below to give an estimated standard deviation of 15.5 percent.

$$Standard\ deviation\ of\ P(Head) = \sqrt{\frac{P(Head) \times \{1 - P(Head)\}}{Number\ of\ Observations}} = \sqrt{\frac{0.6\ \times 0.4}{10}} \approx 0.155$$

If we had more observations, for example, 60 heads from 100 observations, our belief about the model would be stronger. The more data that's available, the less the model uncertainty (see Figure 1).

Figure 1
Probability of a Head (P(Head)) Density Distribution of the Coin and Credibility

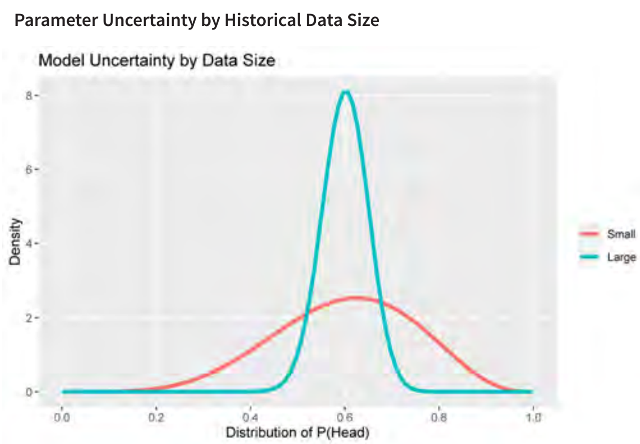


Some level of random noise would occur even if we have a model with a high degree of confidence in the underlying response. In our coin example, even if we are sure about the fairness of the coin, 20 tosses are expected to show a range of the number of heads due to random volatility. We may see 12 heads, but 11 or 13 would also be likely given a coin with 60 percent of probability of heads. More tosses would ensure the outcome will be close to 60 percent of heads (see Figure 2). This measure can be particularly relevant when making predictions for a small number of cases or exposures.

Figure 2
Impact of Random Noise

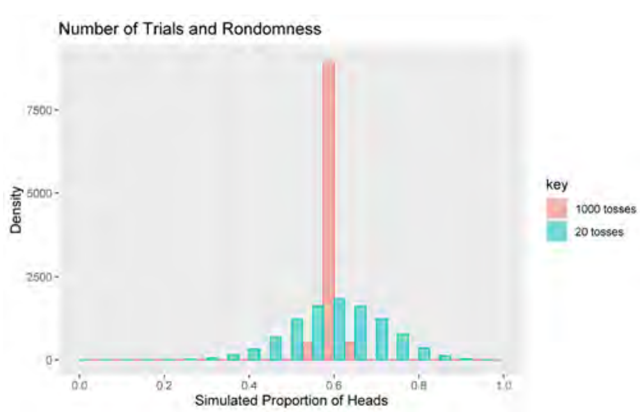


The overall uncertainty of the prediction is derived from both model uncertainty and randomness, as depicted in Figure 3. Even if the historical data indicated there is a 60 percent probability of heads with a high confidence, the future may surprise us (i.e., 10 percent or 95 percent heads).

Figure 3
The Prediction Uncertainty Combines Parameter Uncertainty and Randomness. (Based on 10,000 Simulations to Create the Density)

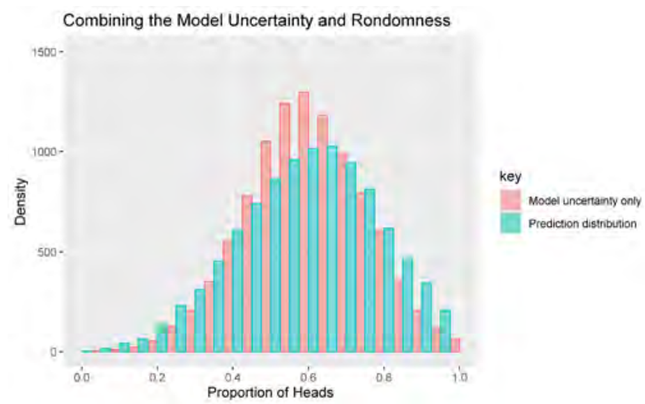


## BAYESIAN APPROACH

So far, we have seen a frequentist view by developing a model on uncertainty relying solely on the historical data. Another approach is to use a Bayesian method that combines historical data and judgment (or a belief in the prior distribution). Depending on the prior belief, the view of the model and prediction uncertainties would differ, as illustrated in four charts in Figure 4. If there had been no or little prior knowledge, the posterior distribution would be more dispersed and fit to data similar to the frequentist view as shown in the chart indicated as uninformative or weak prior belief. A stronger prior belief that the coin is fair would produce a posterior distribution closer to 50 percent with less dispersion, putting less weight to the historical data. The selection of the prior distribution relies on qualitative subject matter expertise and intuition. This is a great way to combine the insights and domain expertise with the historical data, especially when the data is scarce.

> Some level of random noise would occur even if we have a model with a high degree of confidence in the underlying response. In our coin example, even if we are sure about the fairness of the coin, 20 tosses are expected to show a range of the number of heads due to random volatility.

Figure 4
Model Uncertainty Combines Parameter Uncertainty and Randomness. Posterior and Prediction Distributions are Developed Based on 10,000 Simulations.

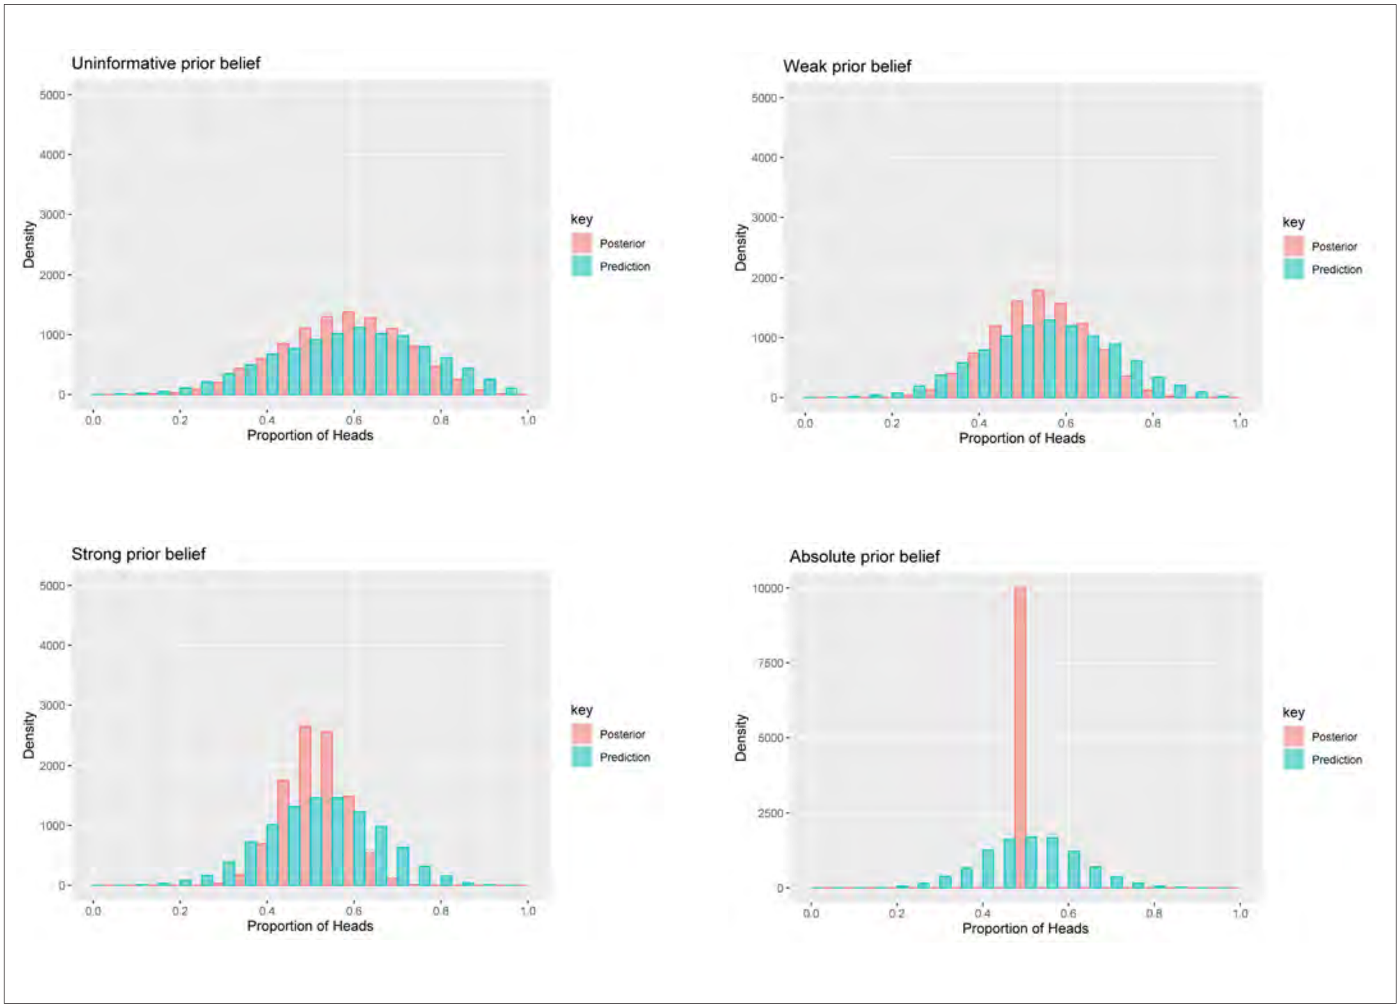


## CASE STUDY: MORTALITY ANALYSIS

With this theory in hand, we can now look to the operations of a life insurer, a much more complex problem than coin tosses. We will apply the same framework illustrated using coins, to a mortality study, which is one of key assumptions for life insurers.

We first developed actual-to-expected (A/E) ratios on historical experience based on a classical approach. The A/E ratios were developed in a tabular form by gender, smoking status and substandard. As a comparison, we performed a Poisson regression, which is one of the popular generalized linear models. We used the Bayesian approach, assuming we are quite confident with the base expected mortality table. The prior distribution is assumed to be 100 percent of the table with 5 percent of standard deviation.

We then developed predictions from a separate hold-out dataset, which was not used to develop the models. We then developed the mean and the 95 percent intervals of the predictions of the hold-out data. The predictions of the mortality rate (per 10,000 lives) of the group is shown in the chart in Figure 5. The actual mortality rate of the hold-out data set was 210 per 10,000 while the traditional model predicted 215. The Bayesian-Poisson model prediction of 212 was closer to the actual compared to the traditional model that overfit to the training data set. The lower bound of confidence interval of the traditional model was 211, which is still higher than the actual (the second chart). A better way to view this is to compare the actual to the prediction interval. The Bayesian-Poisson model expects the actual mortality would be between 206 and 219 with 95 percent confidence, which includes the actual mortality (the fourth chart). The actual data isn't a surprise given the Bayesian analysis.

Figure 5
Predictions (Mean, and Upper and Lower Bounds Based on 95 Percent Confidence Level)

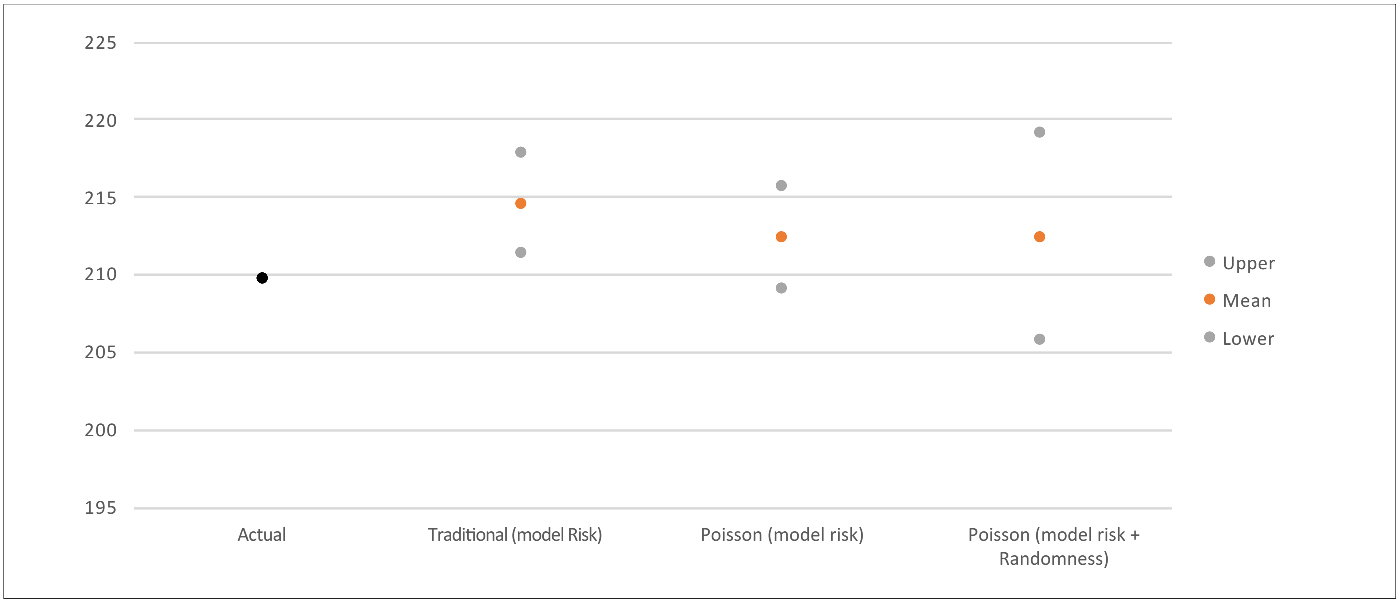


From this we can see that the Bayesian approach is less prone to overfitting, and this was also the case when we reviewed the results in a more granular subgroup level with less credible data. Additionally, the Bayesian framework allowed us to combine historical data and actuarial judgment and helps us directly address the question of how the model is uncertain through its posterior distribution. We could add randomness to the posterior distribution to create the prediction distribution. We believe the Bayesian approach is one of the most effective quantitative analysis tools to inform how the model can deviate from reality and support risk management strategy.

## THE SIGNAL AND THE NOISE

If insurers are truly going to get to grips with the known unknowns in their businesses, such as inevitable variances between predictions and actual experience, their use of predictive models needs to accommodate the inherent prediction uncertainty.

Figure 6 illustrates a framework for doing that and, consequently, avoiding the tendency to accept models without a critical thought. When we develop a prediction model, we try to remove the noise and capture the signal. Ultimately, any model is a generalization of the complex and seemingly chaotic reality; but still useful as an approximation. With that understanding of the model in hand, the business should not ignore the noise, which is and always will be part of the reality. The Bayesian framework provides a way to address the uncertainties associated both with determining the model used for capturing the signal and understanding the possible noise (randomness) that would undermine the accuracy of the prediction.

Figure 6
Illustrative Framework for Assessing Prediction Uncertainty

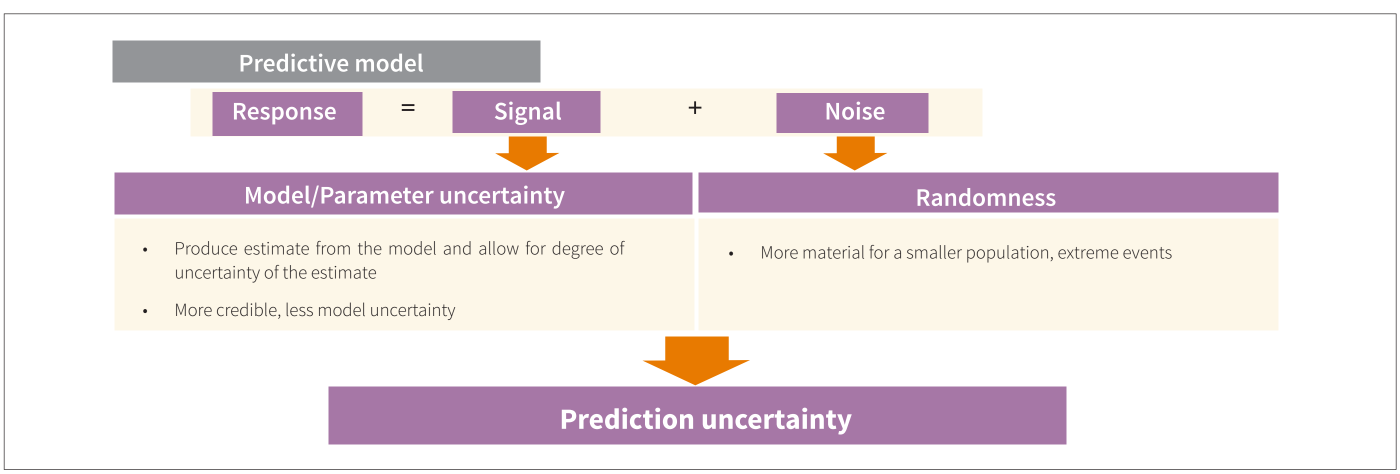

As much as predictive models have improved actuaries' abilities to make more accurate and precise projections and assumptions, our foresight will never be 20/20. That much we know, so building our knowledge of the prediction uncertainty in our models is an essential part of fully understanding them and making sound business decisions based on them. ■

Dan Kim is a director in Willis Towers Watson's Insurance Consulting and Technology business, based in Atlanta. He can be contacted at *dan.kim@willistowerswatson.com*

Boyang Meng is a senior consultant in Willis Towers Watson's Insurance Consulting and Technology business, based in Atlanta. He can be contacted at *boyang.meng@willistowerswatson.com*