Article from

Health Watch

December 2015
Issue 79

Simple and Effective Reserve Practices: Approaches to Combine Your Reserve Estimates for Better Prediction

By Dale Cap, Chris Coulter and Kevin McCoy

> The idea is simple—take the various predictions you are already making and weight them in a way that minimizes variance and increases accuracy.

Actuarial judgment is pervasive in our work. In many cases, judgment is a necessary element to our modeling and analysis. Over the past four decades behavioral research has shown that simple linear models can do much better than a human practitioner in many cases (Kahneman & Tversky, 2011; Wacek, 2007).

We present a couple of simple but effective reserving techniques that an actuary can add to his or her current reserving practices to produce significant reductions in reserve bias as well as reductions to reserve variance. Aggregating reserve estimates using only actuarial judgment can result in high variance and biased results, which can have consequences in many other areas of your company.

According to the Washington State Office of the Insurance Commissioner's data, the range of reserve error reported on financial statements for the largest insurance entities for the years 2008 through 2014, was –10 percent to 40 percent (Company Annual Statements, n.d.). More importantly, the standard deviation of these errors is 11 percent. This data supports the possibility of biases that actuaries generally believe to exist. Biases in reserve estimates include overcompensation (when you've reserved low one year, you overcompensate the next year by reserving way too high); or keeping too much weight on the prior estimates when new information is available; and more. It also indicates that the reserving techniques that are being employed are not very precise. With an 11 percent margin and an 11 percent swing, companies can easily see reserve estimates exceeding the final paid claims by up to 40 percent. This leaves capital in the prior year that could be used to benefit this year. This could impact the bottom line, distort the company's profitability over time, adversely affect ratings in the following year, trigger regulatory action, or impact pricing and forecasting models. Under-reserving can have similar effects. In addition to pricing and forecasting impacts, accruals may be set aside assuming a medical loss ratio (MLR) or other rebates are due, causing inappropriate payments on performance bonuses and bringing additional scrutiny to your department and deteriorating your credibility as the reserving actuary.
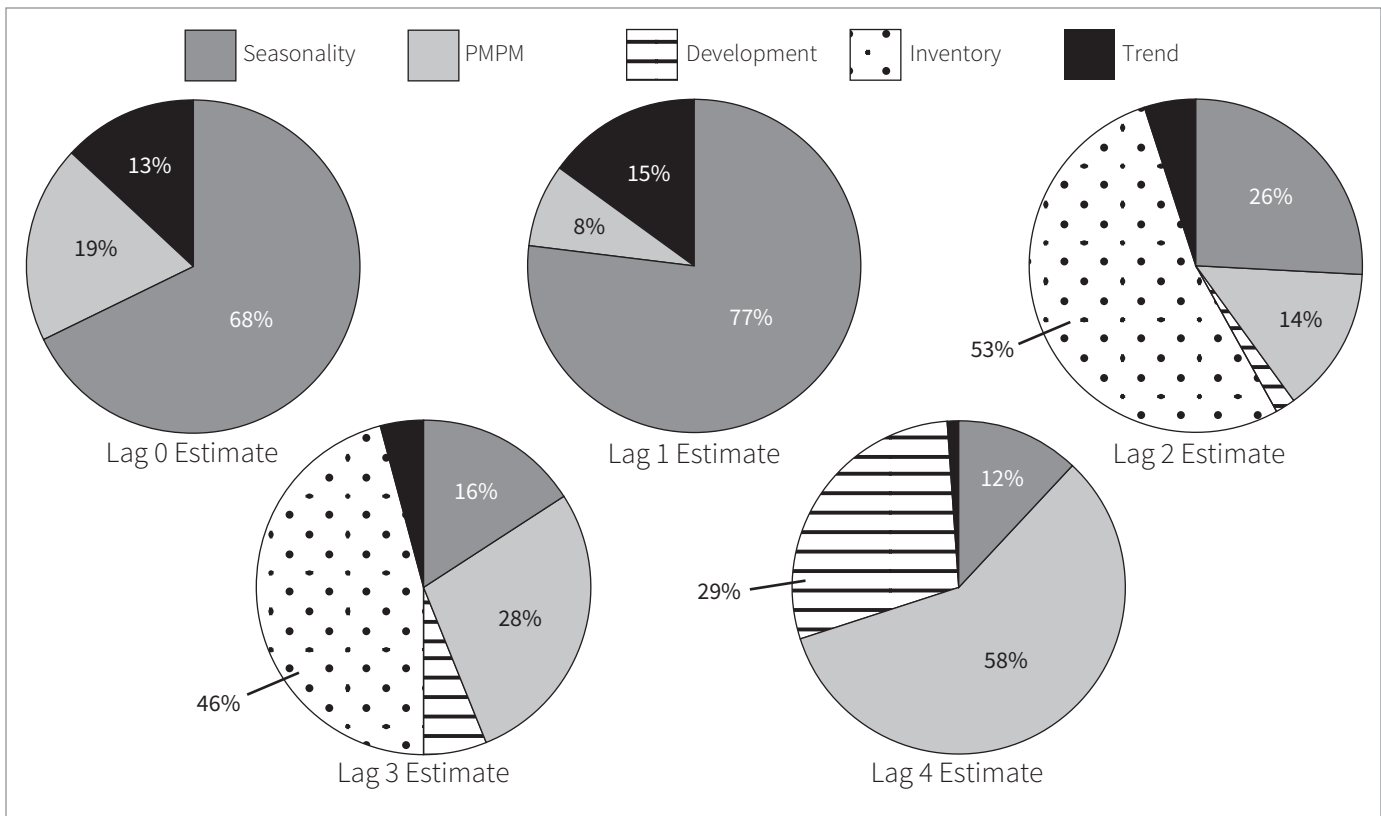
The following results are based on a simulation study with 8,000 simulations of claims run-out. The simulations took into account a seasonality component, a benefit change component, and a large claim component. Each of these components was developed with some randomness in each simulation. These simulations show a reduction of 5 percent variance to the reserve estimates. Unless estimators are completely correlated, these techniques should produce a reduction in variance and a more consistent estimate of the mean. With reduced variance and more accurate predictions, the margins needed could be reduced, resulting in a better estimate of each year's results.

The remainder of this article will outline the proposed techniques, followed by a high-level summary of the simulated data used to illustrate the results. Note: Although we illustrated the results by way of simulation, these techniques have been used in real practice and have shown a significant impact.

WEIGHTING TECHNIQUES

The idea is simple—take the various predictions you are already making and weight them in a way that minimizes variance and increases accuracy. This paper will discuss two weighting techniques you can use. However, there are many different ways to calculate the weights. Every reserving actuary is inherently doing this weighting in some fashion, whether it is via a mental algorithm or a more formalized approach. We advocate using a formalized approach that is testable and avoids potential human biases. In addition, the proposed formalized approach will tend to discredit reserving methods that perform poorly, focusing on those methods that are more reliable and consistent. If nothing else, this will give you a better baseline in which to apply judgment.

The following is an example illustrating the outcome from a weighting technique over multiple reserve methods by lag month.

Legend: Seasonality | PMPM | Development | Inventory | Trend

Lag 0 Estimate: Trend 13%, PMPM 19%, Seasonality 68%

Lag 1 Estimate: Trend 15%, PMPM 8%, Seasonality 77%

Lag 2 Estimate: Seasonality 26%, PMPM 14%, Inventory 53%

Lag 3 Estimate: Seasonality 16%, PMPM 28%, Inventory 46%

Lag 4 Estimate: Seasonality 12%, PMPM 58%, Development 29%

In this example, we used the weighting technique to combine the seasonality, paid per member per Month (PMPM), development, inventory, and trend methods. As you can see each lag differs in the weights applied to each method. In Lag 0, the seasonality method had the highest weight, indicating that it was the "best" model for that lag. However, the seasonality method alone is not the best method. Rather, the weighting given in the above panel minimizes the variance of the estimate, so we would use that weighting for our predictions of Lag 0 claims.

We recommend ongoing monitoring and measurement of any approach used to ensure the intended outcomes and expectations are being met. One of the pitfalls of this more data-driven weighting approach is over-fitting. This is a common pitfall in any estimation or prediction procedure.

TECHNIQUE 1: INVERSE VARIANCE

Inverse variance weights each of the reserve methods based on the inverse proportion of error variance when comparing to actuals. Therefore, lower weights are applied to those methods that have historically produced a larger variation of errors.
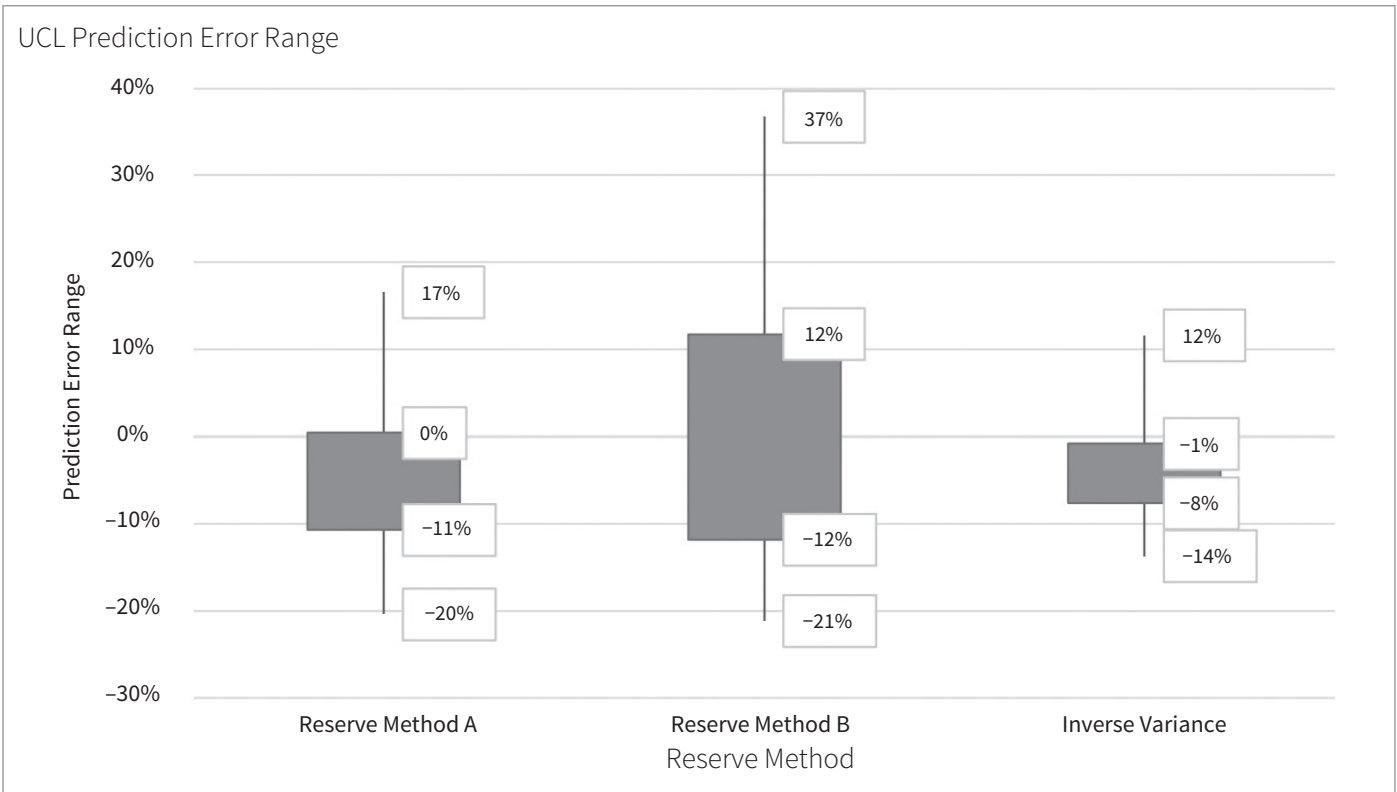
This approach is straightforward and simple to implement without having to add any additional features to one's existing reserve model. It also avoids any complex calculations, making it easy to explain to others. On the other hand, this type of approach ignores the correlations between the reserve methods being used and their distance from the target, which could be used to help lower the variance even further. This is why we offer two approaches.

Example

Suppose you have two methods for reserving: A and B. Each of these methods has a historical monthly reserve error associated with it (variance of 10 and 20, respectively). Based on the inverse variance technique, the proposed future weights when developing a projection could be 86 percent A and 14 percent B. This type of back-test has established that A is a better predictor; however, the mix of the two methods is still preferred. This technique provides a systematic approach to choosing a good mix and possibly better starting point prior to applying judgment in your reserve picks going forward.

Historical Experience	Method A	Method B	Actuals
Month 1	150.00	155.00	151.10
Month 2	160.00	145.00	155.20
Month 3	170.00	180.00	172.30
Variance of Monthly Errors	14.44	88.94	
Inverse Variance	0.07	0.01	
Proposed Future Weights	**0.86**	**0.14**	

After applying the inverse variance against our simulated claims database, using two of the more common reserving methods, we captured the unpaid claim liability (UCL) estimates for each incurred month. These estimates were then compared to the actual known liability, and their range of error is illustrated below. As seen below, the range of error using the inverse variance approach reduces the overall range of error when compared to each reserve method independently. However, you can also see that the technique doesn't improve accuracy significantly.

UCL Prediction Error Range

Chart: Prediction Error Range by Reserve Method

Reserve Method A: 17%, 0%, −11%, −20%
Reserve Method B: 37%, 12%, −12%, −21%
Inverse Variance: 12%, −1%, −8%, −14%

TECHNIQUE 2: LINEAR REGRESSION

The linear regression approach should produce more accurate weightings than the inverse variance approach, but it is far more computationally intensive. To ensure accuracy, the linear regression technique minimizes the sum of squared prediction errors for all points, penalizing larger errors disproportionately. On the other hand, the inverse variance focuses on reducing the dispersion of the estimates instead of the size of the error. In other words, the inverse variance method tends to enhance the precision of the estimate, but not necessarily the accuracy.

Example

Suppose you have two methods used for reserving: A and B. Each of these methods produced a historical estimate for the month. If we define A and B as X (a 2 x 3 matrix with A being column 1 and B column 2) and Y being the actuals, we could use the normal equation to solve for the proposed weights (assuming the matrix is invertible). Below is an example of the equation, where T is the transpose of the matrix and –1 is the inverse.

$$\text{Weights} = (X^T X)^{-1} X^T y$$

Applying this to the table below, the proposed future weights for these methods would be 71 percent A and 29 percent B (for this particular lag).

Historical Experience	Method A	Method B	Actuals
Month 1	150.00	155.00	151.10
Month 2	160.00	145.00	155.20
Month 3	170.00	180.00	172.30
Proposed Future Weights	**0.71**	**0.29**	

This type of back-test has established that A is a better predictor; however, the mix of the two methods is still preferable. This technique provides a systematic approach to choosing a good mix and possibly better starting point prior to applying judgment in your reserve picks going forward.

A similar illustration using linear regression against our simulated claims database can be found below. As discussed above, accuracy is what sets linear regression apart from the inverse variance approach. Unlike the previous results, the results here tend to center themselves on zero.

UCL Prediction Error Range
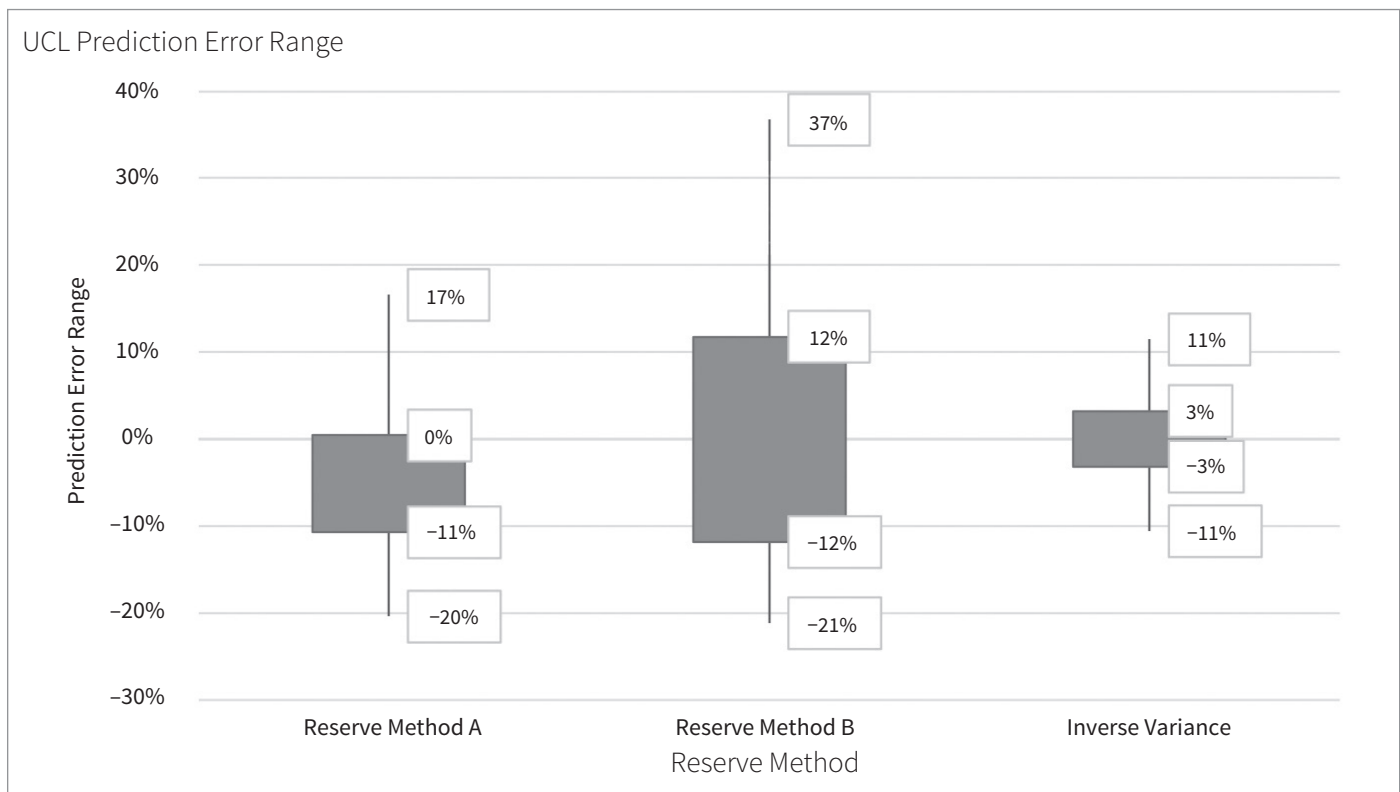
Although we provided an actual example where only two predictors are used, you can include more. Typically, actuaries may have many methods at their disposal like the development method, the paid PMPM method, loss ratio methods, trend-based methods, seasonality-based methods, etc. You can also integrate other variables into the analysis, such as the size of the current claims inventory. For whatever methods are ultimately chosen, we encourage you to pick methods that are diverse and not well-correlated with one another. We also encourage the methods to be consistent and stable over time. At the same time, you should be careful not to over-fit your data.

SUMMARY

In the examples outlined above, we presented two high-level techniques to weight existing reserve estimates. We showed how these techniques can improve your already defined reserving process with little extra work. In addition to the improvement to your estimates, there are two other benefits: the techniques will help the reserving actuary more precisely quantify where and when each reserving method works; and linear regression allows the actuary to integrate stochastic techniques in the calculation of reserve margins. However, there are limitations, and you should be aware of these and use judgment where necessary.

Predictive analytics is the practice of extracting information from existing data to determine patterns and predict future outcomes and trends (Predictive analytics, n.d.). If you don't use a weighting algorithm to combine your reserve estimates, you probably have a pretty good sense of which of your models performs the best for each lag month. But, the question is by how much. A weighting algorithm trained on real data can give you more precision around which models work better and when.

"Predictive analytics" is the new catch phrase, but not long ago stochastic analysis was a hot topic. Reserving is certainly a place where more stochastic models can prove beneficial. A Society of Actuaries sponsored report gives a definition of what margin is for incurred but not reported (IBNR). In math, it is written as:

$$\text{Probability}(\text{Estimate} + \text{Margin} > 95\%) > 85\%$$

The report also gives the reader a couple of ideas on how to obtain this estimate (Chadick, Campbell & Knox-Seith, 2009). In this report, the authors also point you to another Society of Actuaries published report, *Statistical Methods for Health Actuaries IBNR Estimates: An Introduction*, which outlines some more sophisticated ways to statistically approximate your IBNR (Gamage, Linfield, Ostaszewski & Siegel, 2007). Using Technique 2 is a great first step in integrating the stochastics into your already defined reserving system.

The idea of combining two or more estimates for better prediction or lower variance is used in many other contexts; it's called meta-analysis in statistics and ensemble methods in data science, while in finance the capital asset pricing model (CAPM) uses an optimal weighting structure. In any case, they work and can help to reduce the biases that exist in your reserving process.

DATA AND SIMULATIONS

Although these techniques have been shown to be successful in practice, the results included in this paper were developed using data from our simulated claim database to avoid the use of actual data in this paper. The ultimate incurred claims were developed by lag month and include adjustments for changes in claim processing patterns, number of weekly paid claims in a month, benefit design, workday factors, random large claim shocks, seasonality, leveraging, and other factors (which include random noise within each component and overall).

Consistent with actual experience, our simulated examples have shown improved performance when compared to using a single method for reserving. Although we are not able to simulate judgment, we have seen actual improvement when comparing to our final picks (adjusting for margin and implicit conservatism), but we will leave it to the readers to test their own historical performance and whether these techniques add value (or just a better baseline from which to build their estimates).

In the end, we believe if employed correctly—using various reliable and stable methods—these techniques (particularly regression) can help reduce both the bias and variance in the estimates.

Below are the results obtained from applying these techniques to our claims database. Roughly 8,000 simulations were generated estimating the ultimate claim liability for a given month.

Summary Statistics				
Statistics	Reserve Method A	Reserve Method B	Inverse Variance	Linear Regression
Mean Error	−5.2%	0.0%	−4.2%	0.0%
Std Error	5.6%	11.8%	3.4%	3.2%
Kurtosis	−1.9%	−14.8%	24.1%	−0.8%
VAR95%	13.5%	23.4%	9.36%	6.3%
Skew	43.2%	68.3%	30.8%	2.4%
Worst Error	**20.4**%	**36.7**%	13.7%	11.5%

VAR95% represents the point at which 95 percent of the errors (in absolute terms) fall below.

Simulated Distribution of Prediction Errors
IBNR Estimates
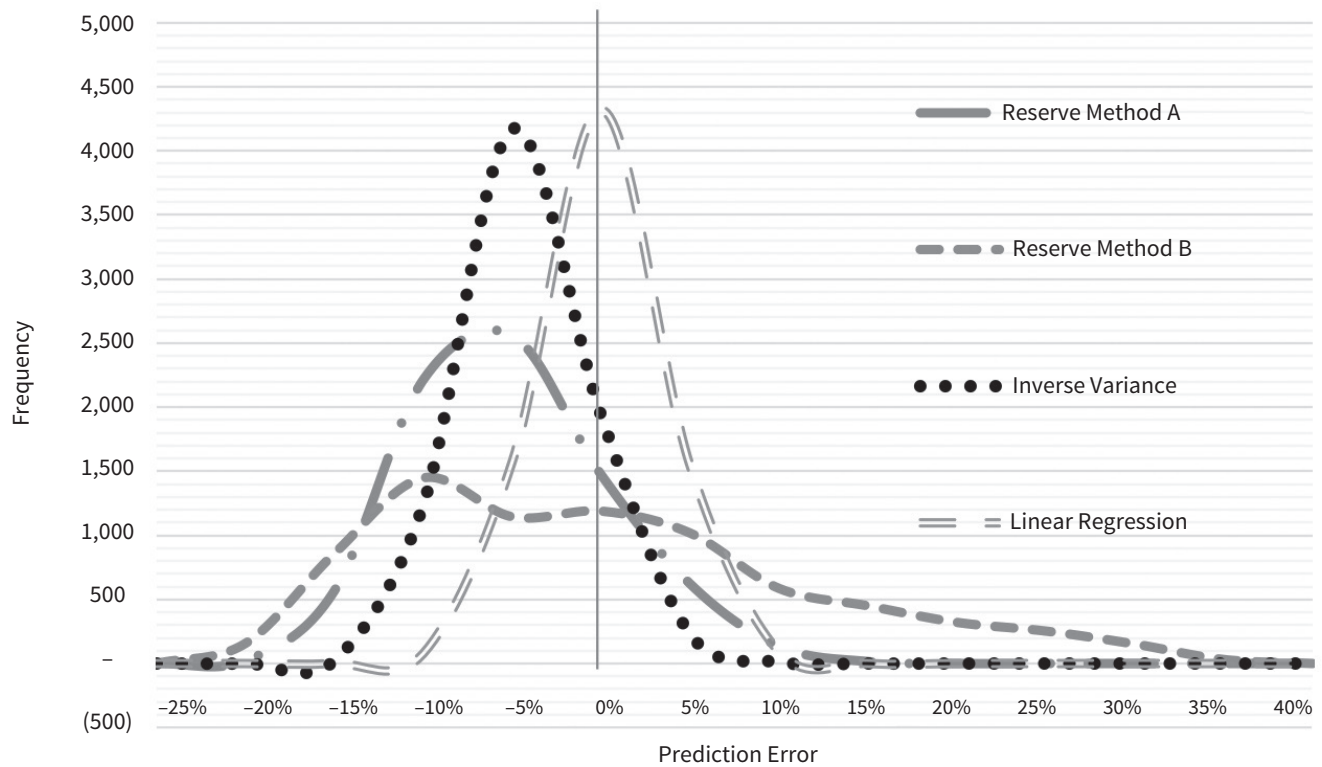
REFERENCES

Chadick, C., W. Campbell and F. Knox-Seith. (2009). *Comparison of Incurred but Not Reported (IBNR) Methods*. Society of Actuaries and Lewis and Ellis.

Company Annual Statements. (n.d.). Retrieved from Washington State Office of the Insurance Commissioner: http://www.insurance.wa.gov/for-insurers/look-up-companies-agents/ins-company-annual-statements/.

Gamage, J., J. Linfield, K. Ostaszewski and S. Siegel. (2007, Oct. 31). *Statistical Methods for Health Actuaries—IBNR Estimates: An Introduction.* Retrieved from Society of Actuaries: https://www.soa.org/research/research-projects/health/research-stats-hlth-act.aspx.

Kahneman, D., and A. Tversky. (2011). *Thinking Fast and Slow*. Farrar, Straus and Giroux.

Predictive analytics. (n.d.). Retrieved from Wikipedia: https://en.wikipedia.org/wiki/Predictive_analytics.

Wacek, M. (2007). *A Test of Clinical Judgement vs. Statistical Prediction in Loss Reserving for Commercial Auto Liability.* Casualty Actuarial Society. ◼

Dale Cap, ASA, MAAA, is a director at Aetna in Seattle and can be reached at *capd@aetna.com.*

Chris Coulter, FSA, MAAA is an actuary at Cambia Health Solutions in Seattle and leader of the Behavioral Finance Subgroup of the SOA's Health Section. He can be reached at *actuarychris@gmail.com.*

Kevin McCoy is a senior actuarial analyst at Cambia Health Solutions in Seattle. He can be reached at *Kevin.Mccoy@regence.com.*