Enhancements to the Shane-Morelli Method to Provide Technical Guidance in Implementation and Proposed Solutions for Challenges Encountered in the Application to a Workers Compensation Tail

Shon Yim, PhD, ACAS

Dolph Zielinski, FCAS

Dawn (Morelli) Fowle, FCAS

Abstract

Motivation. Application of the Shane-Morelli method in practice for multiple reserve reviews revealed potential areas of refinement.

Method. A theoretical examination of curves best used to develop workers' compensation tail factors resulted in a proposed enhancement to this part of the original methodology. Further, a melding of theoretical and practical considerations gave rise to an approach for gradual introduction of mortality to the curve for determination of a more realistic tail. Finally, the determination of a process for updating the proposed model was developed for use by a practicing actuary.

Results. The paper presents a new curve for tail fits, the gradual introduction of mortality into the tail calculation and a process for regularly updating the model.

Conclusions. The updated approach provides an enhanced way of incorporating mortality into traditional aggregate reserving methods in a manner that can be readily explained to business partners.

Keywords. Workers' Compensation; Reserving Methods; Tail; Mortality; Gamma Curve; Shane-Morelli

1. INTRODUCTION

1.1 Background

Modeling of the tail development in worker compensation is well known for being challenging, but it is an important task and thus remains an active area of interest [1]. One approach is to fit a parametric curve to the loss development factors (i.e., LDFs or link ratios) and extrapolate the curve beyond the triangle in order to effectively achieve an extension of the chain ladder method. The limitation with this method is that many parametric curves, such as the popular inverse power, do not converge, so the extrapolation must be truncated at some selected point in the future. Unfortunately, the choice of the truncation point often drives significant swings in the estimate of reserves and can be highly subjective when selected solely based on judgment.

A recent paper by Shane and Morelli provides a practical solution to the truncation problem, hereafter referred to as the Shane-Fowle method [2]. Their method is based on the curve fit approach using the inverse power curve to fit the LDFs. Their key contribution is the use of the life expectancy of the actual claimants to determine the duration of development in each accident year so that the extrapolation of the curve can be truncated using relevant information rather than simple judgment. The Shane-Fowle method is intuitive and can be applied readily. As with any model, however, application to new problems reveals challenges and potential for improvement. The purpose of this paper is to offer enhancements to the application of the Shane-Fowle method in the form of additional discussion around the selection of an appropriate curve, the gradual application of mortality to replace the truncation, and additional guidance around the updating of the underlying assumptions over time.

1.2 Scope

The analysis in this paper is limited to the development of workers' compensation claims in the form of the familiar aggregate runoff triangles. Both paid and incurred data are in scope for medical and indemnity, as are related defense and cost containment expenses. Our model is concerned only with the development of "lifetime claims," by which we mean claims expected to be paid regularly for the remainder of a claimant's life. Based on our experience, ten years is a reasonable threshold for the development age prior to which most shorter duration claims will have closed. This means that our analysis and model are not intended to be used on the most recent ten accident years.

1.3 Outline

The remainder of the paper proceeds as follows. In Section 2, we propose enhancements to the Shane-Fowle method. In particular, we first focus here on the selection of an appropriate curve, with discussion of the commonly known inverse power and exponential curves and the introduction of the gamma curve as a potential improvement. We then move on to a discussion of a method to gradually apply the effect of mortality, which we believe is an improvement to the truncation described in the Shane-Fowle method. In Section 3 we discuss business considerations, including the frequency of model and parameter updates, where we propose a "locked in" approach to the parameters for a pre-selected period of time to avoid overreactions to noise in the underlying data and a framework for understanding and explaining the changes in reserve estimates produced by the model. Section 4 is the conclusion, followed by appendices to support the main body of the paper.

2. ENHANCEMENTS TO THE SHANE-FOWLE METHOD

In this section, we present two enhancements to the original Shane-Fowle method. The first enhancement involves providing guidance to the selection of the curve fit; the second enhancement offers a modification to the application of mortality.

2.1 Notation

The mathematical notation in this section closely follows the 2013 paper by the CAS Tail Factor Working Party [1].

d development age in years.

$c(d)$ cumulative loss (paid or incurred) as of development age d.

$q(d)$ incremental loss from age $d-1$ to d.

$f(d)$ age-to-age LDF (or link ratio) such that $c(d+1) = c(d)f(d)$.

$\nu(d)$ "development portion" of the link ratio, such that $f(d) = 1 + \nu(d)$. For brevity, this is referred to as the development ratio in this paper.

β annual decay rate of incremental losses such that $\beta = [q(d) - q(d+1)]/q(d)$.

2.2 Enhancement to the Curve Fit Selection

Citing the analysis by Sherman [3], Shane-Fowle chose the inverse power curve for fitting the age-to-age LDFs,

$$\hat{\nu}(d) = ad^b, \tag{2.1}$$

where $\hat{\nu}(d)$ denotes the fitted development ratio at age d, and a and b are the fit parameters. The inverse power curve is a widely accepted choice for fitting LDFs [1], but justification for using it appears to be based on the observation that it often provides a reasonably good fit rather than providing sufficient information to allow the user to make an informed selection of the appropriate curve for the data being fit. Even if mortality is arguably the most important factor for a tail model, mortality doesn't start to dominate until much later in the tail, at claimant age 70 or so based on a review of the current life tables. Until then, the projection of loss is determined largely by the curve fit. Thus we believe it makes sense to invest more analysis into the selection of the curve fit.

2.2.1 Analysis of constant decay loss model

To gain some insights that will help guide our selection of the curve fit, we first conduct a simple analysis of a theoretical loss development model. This model will focus on only one accident year. We assume that the incremental loss $q(d)$ decays with age d at a constant rate of decay β such that the sequence of incremental losses starting at $d = 1$ are $q(1)$, $q(1)(1-\beta)$, $q(1)(1-\beta)^2$, and so on. The decay rate β is restricted to the range [0, 1], which means that negative development and increasing incremental losses are not considered. The latter omission might be concerning if we were working with severity data subject to inflation, for example, but the Shane-Fowle method is based on aggregate runoff triangles where losses are expected to decay over time, making the scenario where $\beta < 0$ unlikely. This constant decay model was previously analyzed by McClenahan [4], but his effort was focused more on the formulation of the reserves. In contrast, our purpose is to understand how this particular development pattern relates to the curve fit of the development ratio.

Given a constant value of β, the development ratio can be expressed as

$$\nu(d) = \frac{\beta(1-\beta)^d}{1-(1-\beta)^d}\,. \tag{2.2}$$

The derivation of Equation 2.2 is provided in Appendix A. We will next evaluate this expression for the scenarios of no decay ($\beta = 0$) and significant decay ($\beta \to 1$).

<u>No decay</u>

At $\beta = 0$, Equation 2.2 yields the indeterminate expression $0/0$, so we use L'Hôpital's rule to evaluate the limit,

$$\lim_{\beta \to 0} \nu(d) = \lim_{\beta \to 0} \left\{ \frac{-\beta d(1-\beta)^{d-1} + (1-\beta)^d}{d(1-\beta)^{d-1}} \right\} = d^{-1}\,. \tag{2.3}$$

This is the inverse power curve of Equation 2.1 with parameters $a = 1$ and $b = -1$.

<u>Significant decay</u>

At $\beta = 1$, Equation 2.2 collapses to the trivial solution of $\nu(d) = 0$, which is not meaningful since it equates to no development. So instead of evaluating at $\beta = 1$, we can see what happens as β approaches one. Looking at the denominator of Equation 2.2, the unity term will dominate over the $(1-\beta)^d$ term as β gets closer to one, and consequently Equation 2.2 will tend toward the expression $\nu(d) = \beta(1-\beta)^d$. This is the familiar form of the geometric distribution, which is the discrete analogue of the exponential curve.

To summarize, within the construct of the constant decay model, it is the decay rate that determines which curve fit is optimal. At one extreme, where losses do not decay, the inverse power curve provides the ideal fit. At the other extreme of significant decay, the exponential curve is ideal. In between, the appropriate fit is some blend of the two forms, with a gradual transition from the inverse power to the exponential as β goes from 0 to 1. Appendix B illustrates this transition with a hypothetical example. As expected, the inverse power gives the perfect fit when there is no decay. As long as the decay is low ($\beta < 0.07$), the inverse power fits better than the exponential. But as β increases, the exponential curve becomes superior, and by the time $\beta = 0.2$ the exponential is already practically an ideal fit with an R^2 value of 1.000.

2.2.2 Curve fit selection

The insights from the constant decay loss model analysis can be used to guide the selection of an appropriate curve fit. Below, we discuss the inverse power and exponential forms, finishing with a recommendation for the gamma form.

Inverse power curve

The inverse power curve is ideal when incremental losses have zero or very low decay with development age. For aggregate triangles, this is most likely to happen when the claimant count and the loss run rate (i.e., dollar amount per claimant per year) both remain steady or decrease very slowly over time. Slowly decreasing claimant count is certainly possible for lifetime claims beyond ten years of development and prior to very mature ages, where mortality takes over. Steady run rates can be expected in indemnity paid data with fixed wage replacement payments and without cost of living adjustments. Medical paid data could also exhibit level run rates if most claims are in a steady state of routine treatments and if medical inflation is unremarkable. Finally, paid expense triangles often settle into regular increments in the lifetime development phase. The conclusion here is that the inverse power curve should always be considered for paid data.

Exponential curve

Even paid triangles can exhibit fast decaying incremental losses if the claimant attrition rate is high, which would likely be true for an aged population. Therefore, when the inverse power curve is clearly struggling to fit due to excessive decay, the exponential curve should be considered. With incurred triangles, significant decay may be expected at early ages in cases where a sizable initial case reserve is followed by smaller reserve adjustments in subsequent years. Decay could continue even into mature years to the extent that information about remaining liabilities generally improves over time and results in continually smaller reserve adjustments. In fact, it would be unusual for incremental incurred losses to have little or no decay so exponential curves may often be chosen over inverse power when it comes to incurred data.

Equation 2.2

At first, Equation 2.2 seems like an appealing choice since it effectively represents both the inverse power and the exponential curves as well as the spectrum in between. But there are two reasons against using it. The first is that the premise of this formula is the assumption of a constant decay rate, which is unlikely to be observed in practice. The second reason is that Equation 2.2 is a single parameter curve. As any practitioner can attest, an attempt to fit data as nonlinear as workers' compensation losses to any single parameter curve would likely be a futile exercise. Real data is messy with decay rates and other trends that often change unpredictably over time. For practical reasons, we would want a curve form with more than one parameter to be able to accommodate such data.

Gamma curve

A single curve that can fit both slow and fast decaying losses would be beneficial. We therefore recommend the following form[1]:

$$v(d) = Ad^b e^{rd} \, . \tag{2.4}$$

This equation has the form of the gamma distribution. It consists of three parameters: the scale parameter A, the inverse power decay rate parameter b, and the exponential decay parameter r. This curve form has the following benefits and drawbacks:

- It contains both the inverse power and the exponential forms, so it can fit both low and high decaying loss patterns. The data itself will determine which curve form dominates, or the final form could be a blend of the two, but we do not have to manually choose one or the other.
- It is versatile in that it can be forced to be purely inverse power by constraining r to be zero, or it can be forced to be purely exponential by constraining b to be zero.
- It has three parameters, which allows for more flexibility than the inverse power or exponential alone, but also limits the number of parameters to avoid overfitting.
- Unlike the exponential and the inverse power, curve fitting for the gamma curve is not readily available through standard plotting applications such as Excel. While the choice of the curve fitting method is left to the readers' discretion, one option is to use an iterative algorithm for

[1] We credit Robert Ballmer, FCAS for suggesting this curve form.

minimizing the sum of squared error using Excel's Solver function.

- The gamma curve is not commonly used for this purpose, and as a result, additional communication with business partners regarding the justification of its application may need to take place.

2.3 Enhancement to the Application of Mortality

The second enhancement to the Shane-Fowle method has to do with the application of mortality in the loss projection. Shane and Fowle applied mortality by estimating the average remaining life expectancy of the underlying claimants and calculating a weighted average of life expectancy by accident year. They additionally introduced a useful method of evaluating life expectancy at desired percentiles, which allows the practitioner to judgmentally use percentiles in lieu of the expected value. The selected life expectancy for an accident year was then used to truncate the extrapolated curve fit to obtain the ultimate loss and the corresponding tail factor. Their approach has several merits, including being intuitively appealing and having high practical value. However, we see opportunities for improvement in view of the following considerations.

- The average of the individual life expectancies in a cohort will underestimate the life expectancy of the cohort. If we say that a cohort "dies" with its last remaining member, then the life expectancy of the cohort should exceed the average of the individual life expectancies. Consider the extreme example of a cohort of two people, one with a life expectancy of one year and the other with nineteen years. The average of their life expectancies is ten years. But the life expectancy of the cohort would be greater than ten years because the life expectancy of the youngest member is greater than ten years.
- A cohort in runoff dies off gradually, so it follows that the loss development should also be affected gradually by the force of mortality. Under Shane and Fowle's method of truncation, mortality is introduced at a single point in time – the point considered to be the ultimate. Prior to then, losses are assumed to develop without contribution from mortality other than the mortality which is embedded in the observed triangle. But mortality is nonlinear and observed mortality is a not a good predictor of future mortality.

2.3.1 Methodology for gradual application of mortality

To address the above considerations, we propose a gradual application of mortality[2]. This methodology requires the distribution of claimant ages as well as life tables for the claimant population. If gender is to be considered, and we believe it should, separate life tables for males and females are required along with the gender of each claimant. The number of years of projection is somewhat arbitrary as long as it is enough to cover the remaining life of all claimants. For illustrative purposes, one hundred years is likely a reasonable selection.

Prior to the projection, we assume that the link ratios have already been obtained using standard chain ladder methods and that a curve has already been fit to the development ratios. The steps below are then used to project losses for an individual accident year. For an illustration of the calculations for an example accident year, refer to Appendix C.

1. Group the claimants for the accident year and compute the group's current average age as well as its average age for every year of projection. The projected average group age is estimated using life contingencies whereby the group age at a future year is a weighted average of the member ages with the weights being equal to the probability of survival.

$$\text{age}_{\text{G}}(t) = \frac{\sum_{k=1}^{N} \text{age}_k(t) S_k(t)}{\sum_{k=1}^{N} S_k(t)} \tag{2.5}$$

   In Equation 2.5, N is the number of members in the accident year and t is a time variable representing the number of years into the projection. $\text{age}_{\text{G}}(t)$ is the group average age at time t. $\text{age}_k(t)$ is the age of member k at time t, which is equal to member k's current age plus t. $S_k(t)$, the probability that member k will survive the next t years, can be calculated from the life tables. For example, if hypothetical member k is a 60-year-old male, $S_k(3) = p_{60}p_{61}p_{62}$, where p_x denotes the probability that a male at age x will survive to age $x + 1$. Note that the group ages more slowly than an individual because the oldest members of the group are more likely to leave the group in the following year.

2. Estimate the group mortality rate for each year of projection by using the group average age calculated in the previous step to look up the mortality rate from the selected life table. If using gender-specific life tables, we can obtain a weighted average mortality rate using weights that reflect the gender split for the group.

---

[2] We again credit Robert Ballmer, FCAS for his significant role in the development of this methodology.

3. Use the curve fit for the development ratios (Equation 2.4) to project annual incremental future losses. First, cumulative losses can be projected year after year using $c(d+1) = c(d)(1 + \nu(d))$, starting with the current diagonal as the initial value. Then the incremental losses can be obtained by $q(d) = c(d) - c(d-1)$.

4. Use the incremental projected losses to compute an implied loss decay rate for each projected year.

5. The implied loss decay rate obtained in the previous step includes mortality observed in the triangle data. We want to remove this observed mortality in order to avoid double counting when we apply mortality later in a separate step. A precise evaluation of the observed mortality for each development age is impractically cumbersome, so we rely instead on the most recently observed mortality rate. This observed mortality rate is subtracted from the implied loss decay rate at every projected year. See Appendix C for additional detail.

6. Next we add the group mortality rate estimated in Step 2 to the restated implied loss decay rate from Step 5. This gives us the implied loss decay rate adjusted to include mortality. The subtle assumption made in this step is that loss decay and life mortality rates are additive. This assumption implies that the loss decay comes entirely from the loss of claimants and not from changes in the per claimant severity. This is a reasonable assumption for lifetime paid indemnity since many claimants receive a fixed wage replacement amount. Medical payments are also often steady due to routine treatments and medications although there is more variability here for a number of reasons, including the potential for medical technology and treatments to change over time, and the impact of claim handling practices around closed/reopened claims for routine infrequent treatments. The use of aggregate data should help smooth this volatility to some degree.

7. The final step is to apply the loss decay rate adjusted for mortality in Step 6 to estimate the future incremental losses. The future losses are then summed to produce the unpaid for the accident year.

2.3.2 Comparison of the gradual and truncated application of mortality

Using the projected incremental losses from the example in Appendix C, we can compare the gradual application of mortality described above against a truncated application. Figure 2.1 provides a visual representation to help understand the differences between the two methods. In the Appendix C example, the average life expectancy of the accident year cohort was computed to be 26 years, and that is where the truncated method stops projecting. On the other hand, the gradual application method projects incremental losses that taper down to zero over a period of 50 years. We can also see that losses from the gradual method decrease at a faster rate than losses from the truncation method due to the mortality adjustment applied every year. Comparing the two loss patterns, the gradual method also probably provides a more realistic cash flow projection. The total unpaid estimates for the gradual and truncated application are 1,047 and 964, respectively. It is not surprising that the gradual produced a higher unpaid estimate since it projects a longer life expectancy for the group (though also recall that the original Shane-Fowle method would have used an inverse power which, in practice, would have yielded a higher unpaid estimate than shown here, so this is not a representation of the difference between the original method and our proposed enhancements). In practice, the actuary may choose to use either or both of the methods due to any number of considerations. For example, consider a scenario where the case reserves are presumed to be set adequately to cover all future payments at some point prior to the ultimate life expectancy of the cohort. In this case, the actuary may opt to use the truncated method for the incurred development, while the gradual method may be appropriate for paid development.

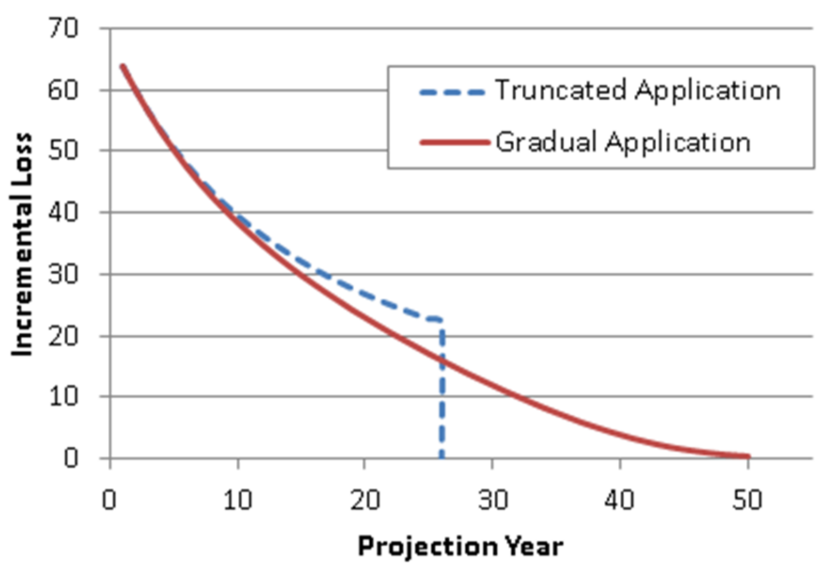

Figure 2.1: Projected incremental losses using the gradual and truncated application of mortality. Results are from the example presented in Appendix C. For clarity, curves are shown as continuous lines, but incremental losses are discrete for each projection year.

3. PRACTICAL CONSIDERATIONS RELATED TO APPLYING THE MODEL FOR SUBSEQUENT RESERVE REVIEW CYCLES

3.1 UPDATING THE MODEL

Since the "tail" analysis may often be one of the most leveraged assumptions in the projection to ultimate on long tail lines, it can also drive some of the most significant financial impacts. Subsequent utilizations of the model described above will likely generate ultimate indications that differ from prior model results. This introduces the age-old debate around stability versus responsiveness. To address this concern, we will explore the concept of "locking in" parameters for an extended period of time with contemporaneous reasonability checks of those parameters.[3]

We propose that when the model is updated, two contemporaneous exercises take place:

- Exercise 1: Updating the Data

  Underlying data is updated, incorporating a new diagonal, or several new diagonals, of paid and incurred information as well as updated information on the distribution of claimant ages and genders on open claim. For this exercise, we lock down the parameters that generated the original loss development curve (specifically parameters *a, b,* and *r,* from equation 2.4 above). We then run through each step of section 2.3.1. with updated claims data. The resulting indications will then be driven entirely by underlying data changes and not changes to the model's parameters.

- Exercise 2: Re-Fitting the Parameters

  In a separate analysis, after updating the underlying data per Exercise 1, we go through the entire process as outlined in Section 2, above, from selecting a curve to the application of mortality. This should essentially be starting from scratch to determine the best model for the data without looking back at the prior results. In this way, the practitioner will have the impacts from updating the underlying data as well as the impacts from updating the model parameters.

For purposes of reducing bias, the practicing actuary should determine ahead of time how long

[3] We credit Michael Shane, FCAS for this suggestion.

the parameter assumptions will be locked. We believe a time period of three to five years is a reasonable starting point. This locking of the parameters insulates against responding to the year-to-year process noise introduced when the underlying data is updated and the curve fit parameters are refit.

Of course, locking in parameters without monitoring changes in the underlying data would expose one to missing changes in underlying trends or material shifts in development patterns. For this reason a determination should be made on a threshold to be used to measure whether or not parameters should be re-fit earlier than the predetermined timeframe. This threshold would be measured against the parameter refit impact, calculated as the total variance from Exercise 2, above, less the process variance from Exercise 1 above. Suggested benchmarks might include a percentage of capital or surplus, or perhaps a percentage of the underlying reserves being modeled. Additionally, if the model has been used for multiple years, this parameter refit impact should be monitored over time, since consistent directional impacts might indicate the need for a refit sooner than anticipated. In theory, one would expect these parameter refit impacts to oscillate around zero if no systemic change is taking place.

Alternatively, one could use a goodness of fit test with the new data and the locked parameters. One other option may be to use a Bayesian framework to assess the new parameters; given a prior assumption, test whether the new information suggests that the prior is no longer valid. The specific test to be used is left to the reader. Again, the determination should be made beforehand around the tolerance around the results of this test and the level at which a decision needs to be made on whether or not to refit the parameters. When the decision is made to refit the parameters, the model fit guidelines presented in Section 2 should again be utilized. The parameter locking process should also be essentially reset at the same time.

The authors would like to clarify that this approach assumes a single and consistent mortality table. In practice, there is likely more than one mortality table that may be used. We leave it to the reader to select the most appropriate mortality table for the purposes of this model. Discussion around which mortality table(s) to utilize are beyond scope of this paper. Further, note that if the underlying mortality table that the practitioner had decided to use for the prior iteration of this model was updated or changed, the practitioner should consider utilizing the most up-to-date mortality table in the next iteration of this model.

3.2 DISCUSSION OF RESULTS WITH BUSINESS PARTNERS

When the data is updated and new indications are generated, the discussion of results with business partners should be fairly intuitive. In theory, the drivers of changes to indications should be directly related to changes in open claim attributes, such as the distribution of ages and genders, or changes to the run rate of payments (e.g., escalating medical severity). This is one of the key benefits of the original Shane-Fowle model, since the resulting changes to indicated ultimates are often fairly obvious and reasonably easy to explain.

When the data and the parameters are both updated, the discussions could be more difficult. The completion of both Exercises 1 and 2, above, allow for the same insights as noted above that relate to changes in the open claims, leaving the actuary to explain the remaining changes as related to the updating of the parameters. If the parameters were updated due to shifts in the underlying data leading to parameter refit variances greater than the selected threshold, it is likely that specific internal or external drivers could be identified that explain the need to refit the parameters and the directional change to indications. These drivers might include:

- Internal – changes to the claim settlement practices, changes to case reserving standards, new cost or expense mitigation efforts, etc.
- External – changes to medical inflation, changes due to state specific reforms, changes to the legal environment due to court rulings, changes to life expectancies, etc.

When parameters are updated due to the predetermined passage of time, such specific drivers may be more difficult to find, but the impact of the parameter update on the overall indications is also likely to be less significant in this case.

4. CONCLUSIONS

The Shane-Fowle model described an intuitive method for incorporating mortality assumptions into otherwise standard actuarial methods to improve the resulting reserve estimates and offer insights for discussion with management and other business partners. This model extends and enhances the Shane-Fowle method in two ways. We introduce the use of the gamma curve, which results in a tail that is fit based on the underlying data with less subjectivity than the more traditional inverse power or exponential curves. Additionally, we propose a gradual application of mortality to the curve fit, which allows for the projection of incremental incurred or paid losses to converge to zero as expected, which is not achieved by either a straight curve fit or the truncation method originally proposed.

This paper has also included a proposal for the regular updating of the model that should facilitate discussion with business partners about the underlying cause of changes to estimates. This recommended approach cautions against modifying parameters overly often, proposing instead that parameters are updated either at predetermined intervals or when there are truly significant changes to the data that suggest the previous model might no longer be appropriate.

Acknowledgment

The authors gratefully acknowledge the generous instruction and support by Michael Shane, FCAS. We also thank Robert Ballmer, FCAS for his significant contributions to this work.

Supplementary Material

Included with this paper is a supplemental Excel tool titled WC Tail Model Template.

5. REFERENCES

[1] CAS Tail Factor Working Party, "The Estimation of Loss Development Tail Factors: A Summary Report," *Casualty Actuarial Society E-Forum*, **Fall 2013**, 1-111.

[2] Shane, Michael, and Dawn (Morelli) Fowle, "Using Life Expectancy to Inform the Estimate of Tail Factors for Workers' Compensation Liabilities," *Casualty Actuarial Society E-Forum,* **Fall 2013**, 1-12.

[3] Sherman, Richard E., "Extrapolating, Smoothing, and Interpolating Development Factors," *Proceedings of the Casualty Actuarial Society*, **1984**, Vol. LXXI, 122-155.

[4] McClenahan, Charles L, "A Mathematical Model for Loss Reserve Analysis," *Proceedings of the Casualty Actuarial Society*, **1975**, Vol. LXII, 134-153.

Biographies of the Authors

Shon Yim, ACAS is an actuarial consultant at CAN, where he has worked on reserving analytics and Long Term Care Insurance pricing. He received a doctoral degree in Mechanical Engineering from MIT. Prior to joining CNA, he worked for a commodity trading advisor, focusing on research and development of technical trading systems.

Dolph Zielinski, FCAS is an AVP and Actuary at CNA with over eight years of reserving experience across multiple lines, including Workers' Compensation. He received his master's degree in Applied Mathematics at Roosevelt University.

Dawn Fowle, FCAS is a senior actuarial consultant at Ernst & Young LLP (EY), where she provides reserving support for numerous clients. Prior to joining EY, Dawn worked for both ongoing and runoff insurers, with significant exposure to Workers' Compensation.

Disclaimers

The views expressed in this article are solely the views of the authors and do not necessarily represent the views of their respective employers. The information presented has not been verified for accuracy or completeness by their employers and should not be construed as legal, tax or accounting advice. Readers should seek the advice of their own professional advisors when evaluating the information.

Appendix A

In this appendix we derive a closed form expression for $\nu(d)$, the development portion of the link ratio. Link ratios for aggregate triangles are typically evaluated over multiple accident years, but this derivation will focus on the development in only a single accident year. Our analysis will impose only one condition: the incremental losses will decay with development age at a constant rate, β. We will derive the formula here on a discrete basis, but the continuous analogue is straightforward.

We denote the cumulative loss as of age d as $c(d)$. The incremental loss from age $d-1$ to d is $q(d)$. In equation form, $q(d) = c(d) - c(d-1)$. The link ratio, also known as the age-to-age loss development factor, is defined as $f(d) = c(d+1)/c(d)$. We are interested in the development ratio $\nu(d)$, which is equal to $f(d) - 1$. It can be shown easily that

$$\nu(d) = q(d+1)/c(d) \tag{A1}$$

Next we define the incremental loss decay rate β in equation form:

$$\beta = \frac{q(d) - q(d+1)}{q(d)} = 1 - \frac{q(d+1)}{q(d)} \tag{A2}$$

Recall that the decay rate is constant, which means that the incremental loss decreases by the same percentage of the previous incremental loss, regardless of the age. Equation A2 can be rearranged to yield the following expression:

$$q(d+1) = (1-\beta)q(d) = (1-\beta)^d q(1), \tag{A3}$$

where the last equality comes from a recursive relation made possible by the assumption that β is constant. Next we recognize that $c(d)$ is simply the sum of q's:

$$c(d) = \sum_{k=1}^{d} q(k) = \sum_{k=0}^{d-1} q(k+1) = \sum_{k=0}^{d-1} (1-\beta)^k q(1), \tag{A4}$$

where the last equality utilizes Equation A3. The right side of Equation A4 is the geometric series with the well-known solution

$$c(d) = q(1)\left[\frac{1-(1-\beta)^d}{\beta}\right]. \tag{A5}$$

We next combine Equation A1 and A3 to get

$$v(d) = \frac{(1-\beta)^d q(1)}{c(d)} \ , \tag{A6}$$

and substituting Equation A5 for $c(d)$,

$$v(d) = \frac{(1-\beta)^d q(1)}{q(1)\left[\frac{1-(1-\beta)^d}{\beta}\right]} \ . \tag{A7}$$

Finally, cancelling the $q(1)$ terms and rearranging yields the formula

$$v(d) = \frac{\beta(1-\beta)^d}{1-(1-\beta)^d} \ . \tag{A8}$$

Appendix B

In this appendix, we examine the inverse power and exponential curve fits for the development ratio of a loss pattern where the incremental losses are subject to a constant rate of decay. We are specifically interested in how the quality of fit responds as the decay rate changes. We conduct the analysis using a simple hypothetical loss development for a single accident year in which incremental losses start at 1,000 at age one and decrease annually by a decay rate of β. For example, a decay rate of 10% produces the following incremental and cumulative loss patterns. Note that we are only looking at development past the first ten years in keeping with the lifetime scope as identified in Section 1.2.

Development Age, d	Incremental Loss, $q(d)$	Cumulative Loss, $c(d)$	Development Ratio, $\nu(d)$
10	387	6,513	0.054
11	349	6,862	0.046
12	314	7,176	0.039
13	282	7,458	0.034
14	254	7,712	0.030
15	229	7,941	0.026

Decay rate β = 0.1

Next, the development ratio as a function of the development age was fit to the inverse power curve,

$$\hat{\nu}(d) = ad^b, \tag{B1}$$

where a is the scale parameter. The more interesting parameter is b, which governs the rate of decay of the inverse power curve. The fitting was done by regression in log space using only data points for development years 10-29. The same data was also fit to the exponential curve,

$$\hat{\nu}(d) = \alpha e^{rd}. \tag{B2}$$

Parameter r controls the decay of the exponential curve. The fits to the inverse power and exponential curve were conducted for a range of decay rates. The quality of fit was measured with the R^2 metric in log space, but it was also confirmed that measuring R^2 in the original dimensions yields similar conclusions. The results are summarized below.

Decay	Inverse Power			Exponential		
Rate, β	a	b	R^2	α	r	R^2
0.00	1.000	-1.000	1.000	0.155	-0.054	0.980
0.05	2.716	-1.537	0.997	0.157	-0.084	0.993
0.07	4.519	-1.801	0.995	0.161	-0.099	0.995
0.10	10.849	-2.248	0.991	0.170	-0.124	0.998
0.15	60.884	-3.116	0.986	0.193	-0.172	0.999
0.20	447.933	-4.119	0.983	0.226	-0.228	1.000

At zero decay rate, the inverse power curve provides a perfect fit with $b = -1$. The response to an increasing decay rate is that b becomes more negative in an attempt to keep up with the faster decay, but this comes at the cost of the goodness of fit, as indicated by the decreasing R^2 value. On the other side of the table, it can be seen that the exponential is an inferior fit to the inverse power at a decay rate of zero, judging by the R^2 value. But the exponential's fit improves as the decay rate increases to the extent that above a decay rate of 7%, the exponential is better than the inverse power.

Appendix C

This appendix illustrates the projection of future losses using the gradual application of mortality as described in Section 2.3. This example uses hypothetical but realistic paid loss and claimant data. Suppose that we have the following information for an accident year:

- Cumulative loss at latest diagonal is 10,000.
- Development age at latest diagonal is 20 years.
- The development ratio has been fit to the curve in Equation 2.4, and the parameters have been estimated to be: $A = 0.4,\ b = -1.37,\ r = -0.00165$.

The table below shows the calculations for this accident year. Note that some figures are rounded for display.

Proj Year t [A]	Avg Age $\text{age}_G(t)$ [B]	Mort Rate [C]	Dev Age d [D]	Fitted LDF $\hat{f}(d)$ [E]	Cumul Loss $c(d)$ [F]	Increm Loss [G]	Loss Decay Rate [H]	Adj Loss Decay [K]	Adj Increm Loss [L]
1	60.6	0.0056	20	1.0064	10,000	64			64
2	61.5	0.0062	21	1.0060	10,064	60	0.060	0.061	60
3	62.4	0.0068	22	1.0056	10,124	57	0.058	0.059	56
4	63.4	0.0075	23	1.0052	10,180	53	0.055	0.057	53
5	64.3	0.0082	24	1.0049	10,234	51	0.053	0.056	50
6	65.2	0.0088	25	1.0047	10,284	48	0.051	0.055	48
7	66.1	0.0094	26	1.0044	10,332	46	0.049	0.053	45
45	94.0	0.1695	64	1.0012	11,275	14	0.022	0.186	2
46	94.6	0.1872	65	1.0012	11,288	13	0.022	0.203	1
47	95.3	0.2039	66	1.0012	11,302	13	0.021	0.219	1
48	95.9	0.2039	67	1.0011	11,315	13	0.021	0.219	1
49	96.5	0.2204	68	1.0011	11,328	13	0.021	0.235	1
50	97.1	0.2394	69	1.0011	11,340	12	0.020	0.254	0

[A] Projection period in years.

[B] Projected average age of the group at the start of year t, calculated with Equation 2.5. This calculation requires a selected life table and the distribution of claimant ages.

[C] Mortality rate looked up from the life table using the age from column [B]. This represents the probability that a person will die within the next year.

[D] Development age at the start of year t in years, starting with age at the latest diagonal.

[E] $= 1 + 0.4[D]^{-1.37}\exp(-0.00165[D])$ is the fitted link ratio using the given fit parameters with Equation 2.4.

[F] = [F previous]*[E previous] is the cumulative loss at the start of year t, obtained from the curve fit. First row is the latest diagonal.

[G] = [F next] – [F] is the incremental loss predicted for projection year t.

[H] = 1 – [G]/[G previous] is the implied loss decay rate between year $t-1$ and year t. Note that there will be discrepancies due to rounding in the table above.

[K] = [H] + ([C] – 0.0056) is the implied loss decay rate adjusted for mortality. The 0.0056 value is the latest observed mortality rate (first row of [C]), which is removed from the adjustment to avoid double counting mortality.

[L] = [L previous]*(1 – [K]) is the adjusted incremental loss.

The sum of column [L] is equal to 1,047. This is the total unpaid estimate for the accident year.

Appendix D

This appendix provides a comparison of reserving estimates using three different techniques: the enhanced Shane-Fowle method presented in this paper, the original Shane-Fowle method and a traditional method. The three methods were applied to example workers' compensation medical paid data from industry sources. All three methods utilized the same LDFs, which were five-year volume weighted averages excluding the first ten development years. The modeling distinctions among the three methods are as follows.

Enhanced Shane-Fowle: the gamma curve (Equation 2.4) was fit to the LDFs, and mortality was gradually applied as described in Section 2.3.

Original Shane-Fowle: the inverse power curve was used to fit the LDFs, and the curve was truncated at the life expectancy according to the average age of the accident year cohort.

Traditional: LDFs were used without curve fitting, and a tail factor was selected based on the incurred/paid ratio at the end of the triangle.

The table below shows the cumulative development factors (CDFs) and unpaid estimates using the three methods.

Acc Year	Cumul. Paid ($M)	Enhanced Shane-Fowle		Shane-Fowle		Traditional	
		CDF	Unpaid ($M)	CDF	Unpaid ($M)	CDF	Unpaid ($M)
1985	984	1.031	31	1.031	30	1.030	30
1986	1,133	1.035	40	1.034	38	1.033	37
1987	1,326	1.039	51	1.037	49	1.037	49
1988	1,530	1.041	63	1.041	62	1.040	61
1989	1,784	1.045	80	1.045	79	1.044	78
1990	2,029	1.048	98	1.047	96	1.048	97
1991	2,180	1.053	115	1.053	116	1.051	112
1992	1,741	1.058	101	1.057	99	1.055	97
1993	1,491	1.062	93	1.062	93	1.060	89
1994	1,449	1.068	99	1.068	98	1.065	95
1995	1,591	1.077	122	1.074	118	1.072	115
1996	1,681	1.081	136	1.077	130	1.081	136
1997	1,975	1.091	180	1.087	171	1.091	179
1998	2,585	1.100	258	1.094	243	1.101	261
1999	2,963	1.110	325	1.106	313	1.113	334
2000	3,486	1.121	422	1.114	396	1.127	442
2001	5,225	1.132	691	1.125	652	1.142	741
2002	5,342	1.149	798	1.141	756	1.158	842
2003	4,900	1.166	815	1.159	778	1.175	856
2004	3,915	1.184	720	1.177	695	1.194	760
2005	3,505	1.207	727	1.204	715	1.217	762
Total			5,967		5,727		6,171

In comparing the enhanced and the original Shane-Fowle estimates, it is interesting to recognize that the two enhancements made to the Shane-Fowle method had opposite effects on the unpaid estimate. The first enhancement of using the gamma curve in lieu of the inverse power lowered the estimate of future development because the exponential component of the gamma made the tail thinner. The second enhancement of the gradual application of mortality increased the estimate of development by projecting a longer life expectancy for the group.

Here we see an example where the three methods produce results that are not materially different from each other. This will not always be the case. For example if the gamma fit is very close to the inverse power, the first enhancement would not have much impact. If at the same time there is a wide distribution of claimant ages, the second enhancement would make a significant difference. The combined effect would be an overall material difference.