Two Symmetric Families of Loss Reserving Methods

Andy Staudt, FCAS, MAAA

Abstract.
In this paper, we introduce two families of loss reserving methods – the Actual vs. Expected family and the Mean-Reverting family. The Actual vs. Expected family can be used to credibly adjust prior expectations, either in terms of a fixed initial estimate or just a prior period's estimate, for deviations between actual and expected experience in the same direction as the deviation. In this regard, methods within this family are useful as an alternative to a fixed *a priori* expectation and when rolling-forward estimates of ultimate loss. Conversely, the Mean-Reverting family can be used to credibly adjust *a posteriori* estimates for deviations between actual and expected experience in the opposite direction of the deviation. In this regard, methods within this family are useful in situations where either the occurrence (or absence) of events decreases (or increases) the likelihood of similar events in the future.

Keywords.
Reserving; Bornhuetter-Ferguson; Chain-Ladder; Benktander; Actual vs. Expected; Mean-Reversion.

1. INTRODUCTION

In this paper, we introduce two families of loss reserving methods – the Actual vs. Expected family and the Mean-Reverting family. The Actual vs. Expected family can be used to credibly adjust prior expectations, either in terms of a fixed initial estimate or just a prior period's estimate, for deviations between actual and expected experience in the same direction as the deviation. In this regard, methods within this family are useful as an alternative to a fixed *a priori* expectation and when rolling-forward estimates of ultimate loss. Conversely, the Mean-Reverting family can be used to credibly adjust *a posteriori* estimates for deviations between actual and expected experience in the opposite direction of the deviation. In this regard, methods within this family are useful in situations where either the occurrence (or absence) of events decreases (or increases) the likelihood of similar events in the future.

Although the primary characterization and purpose of these families are different, they can be expressed generally using the symmetric formulations shown in Table 1.

Table 1. General formulations of the Actual vs. Expected and Mean-Reverting families of loss reserving methods.

Family	Formulation
Actual vs. Expected (AE) Family	$U_{AEi} = U_0 + w_i\left(C_k - p_k U_0\right)$
Mean-Reverting (MR) Family	$U_{MRi} = U_i - w_i\left(C_k - p_k U_0\right)$

Here p_k is the percentage of ultimate loss developed at time k, C_k is the actual loss at time k and w_i is a weighting function.[1] We use U_i as a generic estimate of ultimate loss using method i where

[1] For the purposes of this paper, we take the development pattern and the selection of percentage to ultimate figures p_k as a given rather than discuss the computation or updating of such patterns based on experience.

U_0 is our initial expectation of ultimate and U_{AEi} and U_{MRi} represent the Actual vs. Expected and Mean Reverting variants of projection method i, respectively. When referring to projections of ultimate loss, we drop the time k subscript for simplicity.

To roughly understand these families, note that $C_k - p_k U_0$ is an actual vs. expected adjustment as C_k is the actual loss at time k and $p_k U_0$ is the amount of loss expected at time k based on our initial expectation and the loss development pattern. So, where the Actual vs. Expected family takes as its starting point our *a priori* expectation and credibly adjusts this amount *upward* for the difference between actual and expected experience to date, the Mean-Reverting family takes as its starting point our *a posteriori* estimate of ultimate loss and adjusts this amount *downward* for the difference between actual and expected experience to date.

In this regard, and considering Table 1 in detail, the symmetry of the methods is somewhat obvious. We should note, however, that this symmetry is primarily a mathematical nicety which proves useful in later sections as we derive key members and properties for each of the individual families, rather than a characteristic which intrinsically links these two families. And indeed, each of these families can be considered and used independently of one another. However, as will be discussed in Section 3.3, the Actual vs. Expected family can be used to solve a key shortcoming of the Mean-Reverting family.

1.1 Notation, Abbreviations and a Recap of Common Loss Reserving Methods

Notation and abbreviations will play an important role in this paper, both to understand the methods presented and to reflect their commonalities and lineage. For ease of reading and clarity then, it is useful to include a short but comprehensive discussion on the notation and abbreviations which will be subsequently used.

The basic notation is taken from Mack [2] with the key elements already defined above. But to recap, we define p_k as the percentage of ultimate loss developed at time k, C_k as the actual loss at time k, U_i is the estimate of ultimate loss using loss reserving method i at time k (recall that we have dropped the k subscript for simplicity) where U_0 represents the fixed *a priori* expectation of ultimate; and w_i is a weighting function.

For the remainder of the paper, the subscript i in the term U_i will be replaced with the initials of the loss reserving method used. So, for the Chain-Ladder Method we use CL, for the Bornhuetter-Ferguson we use BF and for the Gunnar-Benktander Method we use GB. We will also use the abbreviation IE, standing for Initial Expected method, and notation U_{IE}, as well as U_0, to refer to our fixed *a priori* expectation of ultimate loss. The former usage is practical as it formalizes our initial expectation as a loss projection method comparable to the CL or BF method. And the latter usage is to be consistent with Mack [2] which defines the *a priori* expectation as the estimate at time $k = 0$. But also, the term U_0, indexed by the time k subscript, will become useful in later sections where we discuss rolling forward a prior period's estimates of ultimate loss (not to be confused with the initial *a priori* estimate). In these cases we use the notation U_k to reflect our current estimate of ultimate loss at time k and U_{k-1} to reflect our prior estimate of ultimate loss at time $k-1$ regardless of the method selected.

For each basic loss reserving method described in the preceding paragraph, the following paper will define an Actual vs. Expected variant and a Mean-Reverting variant. To differentiate the basic loss reserving methods from their variants, we will precede the subscript i in U_i with AE for members of the Actual vs. Expected family and MR for members of the Mean-Reverting family. For instance, we will use the abbreviation AEBF and the notation U_{AEBF} to refer to the Actual vs. Expected Bornhuetter-Ferguson method and the abbreviation MRBF and the notation U_{MRBF} to refer to the Mean-Reverting Bornhuetter-Ferguson method.

1.2 Common Loss Reserving Methods and the Experience Adjusted Method

As a refresher, Table 2 below shows the calculations underlying each of the basic loss reserving methods used in this paper with both the traditional as well as the credibility formulations shown to highlight the relationships between these methods.

Table 2. Notation and formulations of common loss reserving methods as well as the Experienced Adjusted method.

Abbrev.	Name	Traditional Formulation	Credibility Formulation
IE	Initial Expected	$U_{IE} = U_0$	N/A
EA	Experience Adjusted	$U_{EA} = U_0 + p_k\left(C_k - p_k U_0\right)$	$= p_k U_{BF} + \left(1 - p_k\right)U_0$
BF	Bornhuetter-Ferguson	$U_{BF} = C_k + (1 - p_k)U_0$	$= p_k U_{CL} + \left(1 - p_k\right)U_0$
GB	Gunnar Benktander	$U_{GB} = C_k + \left(1 - p_k\right)U_{BF}$	$= p_k U_{CL} + (1 - p_k)U_{BF}$
CL	Chain-Ladder	$U_{CL} = {C_k}/{p_k}$	N/A

While the IE, CL, BF, and GB methods should be familiar to most actuaries, this paper introduces a new method which we call the Experience Adjusted method, denoted using EA. The EA method, although to the best of our knowledge not defined in the actuarial literature, is useful for presenting an evenly-spaced spectrum of potential members of the Actual vs. Expected and

Mean-Reverting families; and is defined as

$$U_{EA} = U_0 + p_k\left(C_k - p_k U_0\right). \quad (1)$$

Where the CL and IE methods are polar opposites, and the BF method is the credibility-weighted average of these two methods, the EA method is the polar opposite of the GB method. To understand this note that the GB method can be expressed as the credibility weighted average of the BF and CL methods as shown in Table 2 or Mack [2], whereas the EA method can be expressed as the credibility-weighted average of the BF and IE methods. Working backwards, we shows this as

$$\begin{aligned} p_k U_{BF} + (1-p_k)U_0 &= p_k\left[C_k + (1-p_k)U_0\right] + (1-p_k)U_0 \quad (2) \\ &= p_k C_k + p_k(1-p_k)U_0 + (1-p_k)U_0 \\ &= p_k C_k + (1+p_k)(1-p_k)U_0 \\ &= p_k C_k + \left(1-p_k^2\right)U_0 \\ &= p_k C_k + U_0 - p_k^2 U_0 \\ &= U_0 + p_k\left(C_k - p_k U_0\right) \\ &= U_{EA} \end{aligned}$$

And thus, these five methods – the IE, EA, BF, GB, and CL – form a spectrum from no credibility to full credibility with respect to current experience.

1.3 Outline

The remainder of the paper is structured as follows. In Section 2 we present the Actual vs. Expected family and in Section 3 we present the Mean-Reverting family. Generally these sections follow the same outline where we first present the general formulation of the family and then discuss the family's "generator function" which is used to derive specific members of that family. We then discuss considerations when selecting a specific member of each family before focusing on the practical uses of each family and note any contra-indications. We will also introduce extensions to the basic versions of these families as defined in Table 1. In Section 2, we introduce the Generalized Actual vs. Expected family which can be used to roll forward prior estimates of ultimate loss. And in Section 3, we introduce the Adjusted Mean-Reverting family which corrects for a flaw in the basic version of the Mean-Reverting family. Finally, in Section 3, we also comment upon the relative accuracy of the Mean-Reverting family using hindsight testing. The conclusion of this paper highlights the four members of these families which might prove most useful to the actuary in a practical setting. We also include an Appendix which discusses the motivation for this paper (Appendix A) and attach an Excel file which shows how to implement these methods (Appendix B).

2. THE ACTUAL VS. EXPECTED FAMILY

In this section, we focus on the Actual vs. Expected family. This family can be used to credibly adjust prior expectations, either in terms of a fixed initial estimate or just a prior period's estimate, for deviations between actual and expected experience in the same direction as the deviation. In this regard, methods within this family are useful as an alternative to a fixed *a priori* expectation and when rolling forward estimates of ultimate loss.

2.1 General Formulations

2.1.1 The actual vs. expected formulation

As discussed above, the Actual vs. Expected family is defined as

$$U_{AEi} = U_0 + w_i\left(C_k - p_k U_0\right). \tag{3}$$

Without any loss of generality, Equation (3) can be used to develop any data triangle (i.e., paid or incurred losses as well as reported or closed claim counts). For the purpose of understanding this family, note that $C_k - p_k U_0$ is an actual vs. expected adjustment as C_k is the actual loss at time k and $p_k U_0$ is the amount of loss expected at time k based on our initial expectation and the loss development pattern. For example, if actual losses are more than expected, this family would adjust the initial expectation upward allowing for some portion, w_i, of this deviation, and *vice versa*.

From this interpretation, it is obvious that the critical factor is the weighting function w_i which determines the amount of reliance we place on the actual vs. expected adjustment relative to initial expectations. If we were to set $w_i = 1$, then we would adjust the *a priori* expectation fully for the deviation between actual and expected experience. On the other hand, if we were to set $w_i = 0$, then Equation (3) would reduce to the *a priori* expectation, ignoring actual experience. This loosely suggests that the weighting function w_i can be viewed as the credibility of the actual vs. expected adjustment and that an acceptable constraint is $w_i \in [0,1]$.

Consider the following example. Suppose that the historical percentage of loss developed at time k is 25%, the initial expectation of ultimate is $200, and the current loss amount is $150. In Equation (3), suppose that we set the weighting factor equal to the percentage of loss developed at time k (i.e., $w_i = p_k$). As will be shown in the next section, this is actually a special case of the Actual vs. Expected family – namely the Actual vs. Expected Bornhuetter-Ferguson (AEBF) method. Table 3 below compares our fixed initial expectation against the AEBF method. From this comparison, we see that as actual losses were $100 more than expected ($150 less 25% of $200), but only 25% credible according to the weighting function defined above, we only adjust our initial expectation upward by $25 (25% of 100).

Table 3. Simple example comparing IE method with the AEBF method.

IE Method			AEBF Method		
U_{IE}	=	U_0	U_{AEBF}	=	$U_0 + p_k(C_k - p_k U_0)$
	=	200		=	$200 + 25\% \times (150 - 25\% \times 200)$
				=	$200 + 25\% \times 100$
				=	225

2.1.2 The credibility formulation

Even restricting $w_i \in [0,1]$, there are still an infinite number of members of the Actual vs. Expected family, which, practically, isn't a very useful result. Rather, it is more constructive to limit ourselves to a finite subset of the family. One or two methods which could be used regularly during a reserve review. To this end, consider the following credibility formulation:

$$U_{AEi} = p_k U_i + (1 - p_k) U_0 . \tag{4}$$

Equation (4) takes the standard form of credibility-weighted averages defined in the actuarial literature (see Mahler and Dean [3]) where we use p_k to weight together our "observation" U_i based on experience with our initial estimate U_0 based on "other information." Although it is not immediately obvious, this credibility equation defines a subset of members of the Actual vs. Expected family. Similar to the moment or probability generating functions in statistics, this equation can be used as a "generator function" for the Actual vs. Expected family, where, by inserting common loss reserving methods into the U_i term, the resulting formula (rearranged) returns a member of this family.

For example, suppose we were to insert the BF method into Equation (4). Then we can derive what we will call the AEBF method as

$$\begin{aligned} U_{AEi} &= p_k U_i + (1 - p_k) U_0 \\ U_{AEBF} &= p_k U_{BF} + (1 - p_k) U_0 \\ &= p_k \left[C_k + (1 - p_k) U_0 \right] + (1 - p_k) U_0 \\ &= p_k C_k + p_k U_0 - p_k^2 U_0 + U_0 - p_k U_0 \\ &= p_k C_k - p_k^2 U_0 + U_0 \\ &= U_0 + p_k (C_k - p_k U_0) \end{aligned} \tag{5}$$

Or, to show another derivation, consider inserting the EA method into Equation (4) to derive what we will call the AEEA method as

$$\begin{aligned}
U_{AEi} &= p_k U_i + (1-p_k)U_0 \\
U_{AEEA} &= p_k U_{EA} + (1-p_k)U_0 \\
&= p_k\left[U_0 + p_k\left(C_k - p_k U_0\right)\right] + (1-p_k)U_0 \\
&= p_k U_0 + p_k^2 C_k - p_k^3 U_0 + U_0 - p_k U_0 \\
&= p_k^2 C_k - p_k^3 U_0 + U_0 \\
&= U_0 + p_k^2\left(C_k - p_k U_0\right)
\end{aligned} \tag{6}$$

Table 4 below presents three other distinct members of this family, along with the AEBF and AEEA methods, which were all derived in a similar manner – by inserting the named loss reserving method into the generator function in Equation (4) and rearranging. Table 4 also explicitly presents the weight function which defines each of these methods within the actual vs. expected formulation. This function reflects weight or credibility each method gives the actual vs. expected adjustment, with the AEIE method placing no weight on the adjustment and the AEEA, AEBF, AEGB, and AECL methods placing increasing weight on actual relative to expected experience.

Table 4. Members of the Actual vs. Expected family of loss reserving methods differentiated by their weight function.

Method	Credibility Formulation	Actual vs. Expected Formulation	Weight Function
AEIE	$U_{AEIE} = p_k U_{IE} + (1-p_k)U_0$	$U_{AEIE} = U_0 + 0\left(C_k - p_k U_0\right) \Rightarrow U_{IE}$	0
AEEA	$U_{AEEA} = p_k U_{EA} + (1-p_k)U_0$	$U_{AEEA} = U_0 + p_k^2\left(C_k - p_k U_0\right)$	p_k^2
AEBF	$U_{AEBF} = p_k U_{BF} + (1-p_k)U_0$	$U_{AEBF} = U_0 + p_k\left(C_k - p_k U_0\right) \Rightarrow U_{EA}$	p_k
AEGB	$U_{GB} = p_k U_{GB} + (1-p_k)U_0$	$U_{AEGB} = U_0 + \left(2p_k - p_k^2\right)\left(C_k - p_k U_0\right)$	$2p_k - p_k^2$
AECL	$U_{AECL} = p_k U_{CL} + (1-p_k)U_0$	$U_{AECL} = U_0 + 1\left(C_k - p_k U_0\right) \Rightarrow U_{BF}$	1

The table above begins to indicate an important relationship. Namely, the Actual vs. Expected family is fundamentally a generalization of the BF method. Table 5 below illustrates this by placing the credibility formulation of the Actual vs. Expected family as defined in Equation (4) next to the credibility formulation of the BF method. Rather than restricting our "observation" within the credibility formula to the CL method, the Actual vs. Expected family lets us use any alternative method. In fact, we note that if we use the BF method as our plug-in estimator, than the resultant Actual vs. Expected variant is the EA method discussed in Section 1.2.

Table 5. Comparison of the Actual vs. Expected family and the BF method.

Family	Formulation
Credibility Formulation of Actual vs. Expected Family	$U_{AEi} = p_k U_i + (1-p_k)U_0$
Bornhuetter-Ferguson Method	$U_{BF} = p_k U_{CL} + \left(1-p_k\right)U_0$

As a result of this exercise, we have gone from an infinite set of members defined by the weighting function in Equation (3), to an infinite subset defined by the credibility formulation in Equation (4), to a finite subset of five members which can each be expressed as a variant of a common loss reserving method. Figure 1 illustrates this progression graphically.

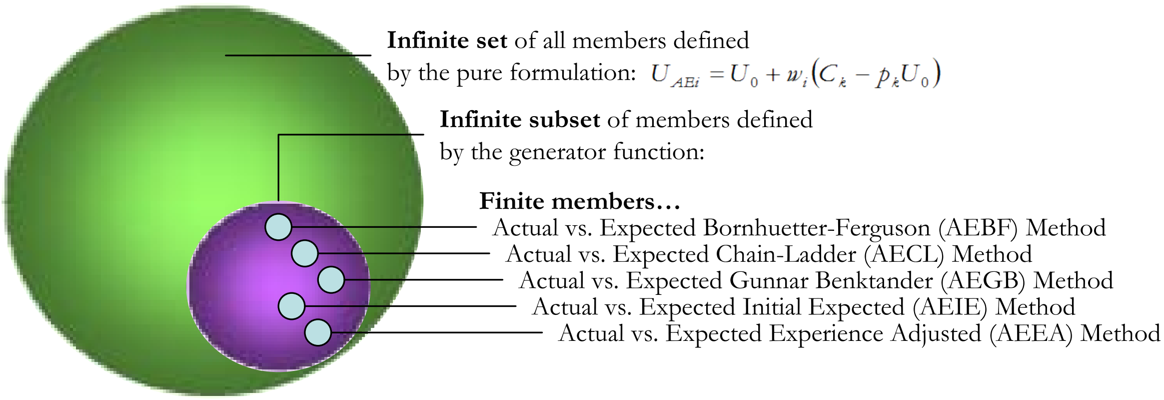

Figure 1. Graphical representation of possible members of the Actual vs. Expected family.

Although the methods shown in Table 4 are by no means the optimal members of the Actual vs. Expected family, the fact that we can express each of these methods as a credibility-weighted average of common loss reserving methods, with our fixed *a priori* expectation as the complement of credibility, makes them a sensible first choice.

2.2 Selecting a Specific Member

The primary motivation for the Actual vs. Expected family is the obvious inability of a fixed *a priori* expectation to learn with experience updating expectations with new information. Because the Actual vs. Expected family is effectively our *a priori* expectation U_0 with an adjustment for experience $w_i(C_k - p_k U_0)$, this family is quite useful as an alternative seed to the BF method with the weight function controlling the degree of responsiveness relative to stability when updating initial expectations, and thus providing us with a natural heuristic to choose between alternative members of the Actual vs. Expected family.

2.2.1 Using the weight function to select a method – in general

Considering the five distinct members shown in Table 4, the first and most obvious way to choose a member of this family is with reference to the weight function w_i, which describes the reliance we place on deviations between actual and expected experience. Consider the AEBF method as defined by Equation (5). Here, as losses develop to ultimate (i.e., $p_k \to 100\%$), the AEBF method tends toward the ultimate loss amount rather than staying fixed at the initial expectation. Now, in this instance, we do not apply the full actual vs. expected adjustment, rather the rate at which we adjust the *a priori* expectation for actual experience is commensurate with the percentage of loss developed $w_i = p_k$ at time k.

Figure 2 below illustrates the weight each of the defined Actual vs. Expected methods place on the actual vs. expected adjustment as a function of the amount of developed experience.

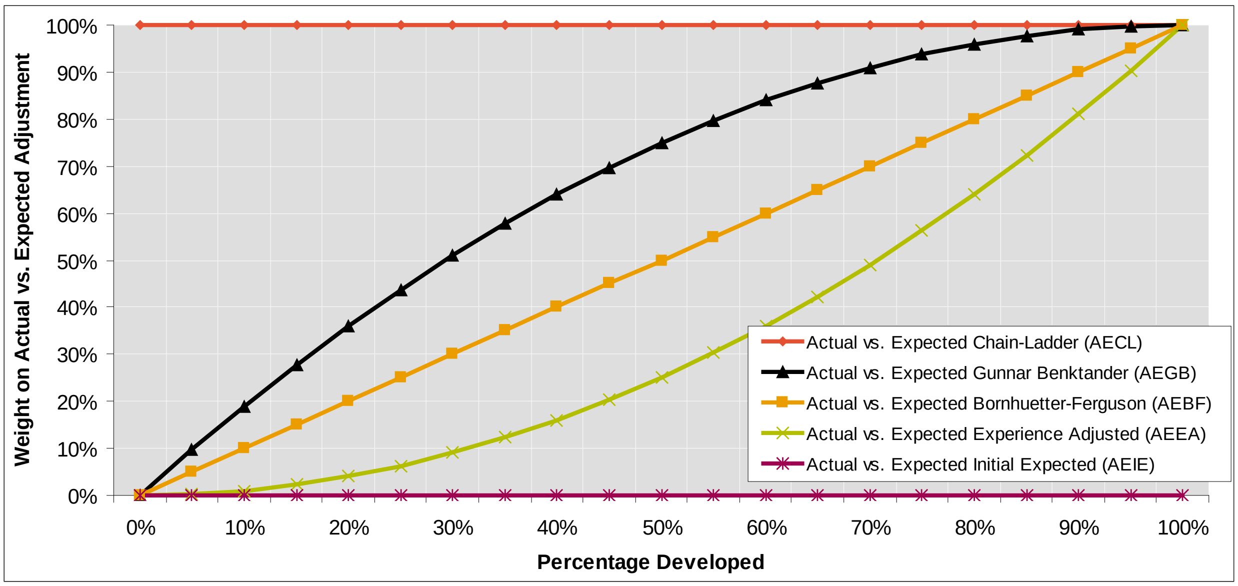

Figure 2. The weight these members of the Actual vs. Expected family give the actual vs. expected adjustment.

From this illustration, we can see that the AEGB method is more responsive than the AEEA method. As was discussed in Section 1.1, this makes sense, given that the GB method places more weight on developed experience than the EA method which places more weight on the initial expectation. And the AEBF method, as with its namesake, takes the middle ground and places "equal" weight on the actual vs. expected adjustment as on the *a priori* expectation. In contrast, the AECL method makes a full allowance for the actual vs. expected adjustment.

Or, presented another way, return for a moment to the previous example (i.e., the percentage developed at time k is 25%, the initial expectation is \$200, and the current loss amount is \$150). Table 6 shows the estimate of ultimate loss in this example using each of the five defined Actual vs. Expected methods. Note that we have also explicitly specified the weight given to developed experience as a means of indicating the relative responsiveness / stability of each method.

Table 6. Projections of ultimate loss using various members of the Actual vs. Expected family.

Method	Ultimate Loss Projection			Weight Function		
AEIE	$U_{AEIE} = U_0 + 0(C_k - p_k U_0)$	=	\$200.00	$w_{IE} = 0$	=	0.00%
AEEA	$U_{AEEA} = U_0 + p_k^2(C_k - p_k U_0)$	=	\$206.25	$w_{EA} = p_k^2$	=	6.25%
AEBF	$U_{AEBF} = U_0 + p_k(C_k - p_k U_0)$	=	\$225.00	$w_{BF} = p_k$	=	25.00%
AEGB	$U_{AEGB} = U_0 + (2p_k - p_k^2)(C_k - p_k U_0)$	=	\$243.75	$w_{GB} = 2p_k - p_k^2$	=	43.75%
AECL	$U_{AECL} = U_0 + 1(C_k - p_k U_0)$	=	\$300.00	$w_{CL} = 1$	=	100.00%

2.2.2 Using the weight function to select a method – more specifically

While there is not a most accurate member of the Actual vs. Expected family, it is useful to take a position. For the moment, let us consider the AEBF, as the weight it gives to the difference between actual and expected losses is directly proportionate to the amount of experience, so that neither prior expectations nor actual relative to expected losses are unduly favored.

Or, put another way, actuaries will be familiar with actual vs. expected diagnostics which compare the change in ultimate to actual less expected experience. In these diagnostics, if the actuary is using the BF method, the change in ultimate will perfectly mirror the actual vs. expected statistic. Or if the actuary is pegging loss to a prior or initial estimate, then the change in ultimate will be zero regardless of actual experience. However, as is very often the case when reviewing these diagnostics, the change in ultimate generally lies somewhere between zero and the actual vs. expected statistic indicating that partial credibility has been given to actual vs. expected experience in the period. This makes sense, given that actuaries will often select an estimate of ultimate loss based on not just one projection method but a variety of methods utilizing averaging, rounding, and potentially manual adjustments to the methods where necessary. To this end, the Actual vs. Expected family is useful for formalizing the results of these diagnostics into a projection method, where at one extreme, the diagnostic is ignored, and, at the other extreme, the diagnostic is believed. And in between, the diagnostic is given a degree of credibility proportional to the amount of experience in the period.

In the case of the AEBF method, using the percentage of expected loss development in the period p_k as the degree of credibility seems like a natural and sensible choice. Note that this construction allows the credibility we place on actual loss experience to be linearly proportional to our expectation of loss emergence over the same period.

2.3 The Generalized Actual vs. Expected Family

Alternatively, a natural use of the Actual vs. Expected family arises when rolling forward prior actuarial work. Consider the situation where full reserve reviews are done periodically (perhaps annually or quarterly) and actual vs. expected diagnostics are used in the interim to adjust for experience over the period. Most often, actuaries tend to one extreme or the other and either allow for 100% of the experience in the period (as updating estimates of ultimate loss using the BF method would) or make no adjustment for the experience in the period (as fixing estimates of ultimate loss at prior selections would). In the case of the former instance, the change in ultimate would exactly mirror the actual vs. expected statistic, and in the case of the latter, the change in ultimate would be zero regardless of the actual vs. expected statistic.

In these situations, given the all-or-nothing nature of movements over what are potentially short and not fully credible time intervals, it is perhaps more useful to first assess the credibility of experience in the period and then adjust our estimates of ultimate as such. Hopefully, this process more appropriately balances the need for responsiveness with the need for stability, or at least provides a formalized means of doing so. Equation (7) generalizes the AEBF method for exactly this purpose – to string together estimates of ultimate loss in subsequent development periods controlling for the random volatility vs. credibility of loss emergence within relatively short intervals. We call this method the Generalized Actual vs. Expected Bornhuetter-Ferguson method or the GAEBF method.

$$U_{GAEBF} = U_k = U_{k-1} + \left(\frac{p_k - p_{k-1}}{1 - p_{k-1}}\right)\left(\left(C_k - C_{k-1}\right) - \left(\frac{p_k - p_{k-1}}{1 - p_{k-1}}\right)\left(U_{k-1} - C_{k-1}\right)\right) \tag{7}$$

Here the subscript k refers to the current period and the subscript $k-1$ refers to the prior period. Note that for development from time $k=0$ where $p_{k-1}=0$, $U_{k-1}=U_0$ and $C_{k-1}=0$, Equation (7) reduces to the AEBF method described in Equation (5). And as with Equation (5), this projection method adheres to the general principle that the longer the period over which actual experience is measured, the more weight given to actual experience relative to prior expectations, either with regard to a prior estimate or an initial expectation.

To help understand how the AEBF method works in this situation, it is useful to further develop the simple example of the previous section. Suppose that one month has elapsed since our previous actuarial review where we ended up selecting $225 (i.e., the amount as projected under the AEBF method) and the incurred loss amount is now $195 (i.e., actual incurred in the period of $45) and the percentage developed is now 40% (i.e., expected percentage developed in the period of 15%). From Equation (7), we can roll forward our prior estimate of ultimate loss of $225 as

$$
\begin{aligned}
U_k &= U_{k-1} + \left(\frac{p_k - p_{k-1}}{1 - p_{k-1}}\right)\left(\left(C_k - C_{k-1}\right) - \left(\frac{p_k - p_{k-1}}{1 - p_{k-1}}\right)\left(U_{k-1} - C_{k-1}\right)\right) \\
&= 225 + \left(\frac{40\% - 25\%}{1 - 25\%}\right)\left(\left(195 - 150\right) - \left(\frac{40\% - 25\%}{1 - 25\%}\right)\left(225 - 150\right)\right) \\
&= 225 + 20\%\left(45 - 20\% \times 75\right) \\
&= 225 + 20\%\left(45 - 15\right) \\
&= 225 + 6 \\
&= 231
\end{aligned}
\tag{8}
$$

Actual development in the period was $45 and expected development was 20% of the unreported amount of $75, or $15; thus, the actual vs. expected adjustment is $30. However, as this was a short period with only 20% expected development on unreported, we only adjust the ultimate loss amount by 20% of total implied adjustment, or $6, for an ultimate of $231. Note that the degree of credibility depends both on the length of the period as well as the shape of the paid or incurred development patterns.

2.4 Contra-Indications

With regard to using a member of the Actual vs. Expected family in either of the situations listed above, there aren't necessarily any obvious contra-indications.

In the former instance, when using a non-trivial member of the Actual vs. Expected family as an alternative to a fixed *a priori* expectation, this family will often be preferable to fixing an initial expectation and failing to update this expectation as more evidence becomes available. Additionally useful, this family allows the actuary to determine the extent to which they wish to peg their *a priori* estimate to initial expectations with the AEEA being the most sticky and the AEGB being the most aggressive (ignoring the trivial case of the AEIE method).

Similarly, in the latter instance, when using a non-trivial member of the Generalized Actual vs. Expected family to roll-forward prior estimates, it is perhaps more a judgment call (rather than a case of selecting a "most accurate" method) when deciding between allowing for 0% of the actual vs. expected experience as is true of fixing estimates of ultimate loss at prior selections, 100% of actual vs. expected experience as is the case with updating BF projections or somewhere in the middle taking into consideration the credibility of experience over the time interval. In any event, this approach should provide the actuary with more freedom when it comes to balancing responsiveness and stability.

3. THE MEAN-REVERTING FAMILY

This section introduces the Mean-Reverting family of loss reserving methods. As the doppelganger of the Actual vs. Expected family, the Mean-Reverting family can be used to credibly adjust *a posteriori* estimates for deviations between actual and expected experience in the opposite direction of the deviation. In this regard, methods within this family are useful in situations where the occurrence of events *decreases* the likelihood of similar events in the future, or likewise, when the absence of events *increases* the likelihood of similar events in the future.

Put another way, this family effectively relaxes the independence assumption of the BF method and the positive dependence assumption of the CL method and allows for the potential of some negative dependence between current and future losses.

3.1 General Formulations

3.1.1 The actual vs. expected formulation

Similar to the Actual vs. Expected family, members of the Mean-Reverting family are grounded in an actual vs. expected adjustment. However, where in respect of the Actual vs. Expected family, the adjustment is used to fine-tune *a priori* expectations for actual experience, in respect of the Mean-Reverting family, the adjustment is used to bring *a posteriori* projections back toward some long-run estimate of the mean. Mathematically, this is formulated as

$$U_{MRi} = U_i - w_i\left(C_k - p_k U_0\right). \tag{9}$$

To understand the name and purpose of the Mean-Reverting family, note that through some simple manipulations (adding and subtracting C_k) we can rearrange Equation (9) as

$$U_{MRi} = C_k + \left[\left(U_i - C_k\right) - w_i\left(C_k - p_k U_0\right)\right]. \tag{10}$$

This arrangement is useful as it isolates both the unadjusted and adjusted outstanding reserve, $\left(U_i - C_k\right)$ and $\left[\left(U_i - C_k\right) - w_i\left(C_k - p_k U_0\right)\right]$, respectively, from the current amount of loss C_k. In doing so, it becomes clear that the Mean-Reverting family offsets the unadjusted outstanding reserve for the amount by which actual losses deviated from expected losses (subject to some weight w_i). For example, if losses to date were more than expected, the outstanding reserve would be decreased to reflect the propensity for future losses to be less than expected, and *vice versa*. It is this type of "mean-reversion" from which the name of the family is derived.

To understand the mechanics of this family, we return again to the simple example from the previous section. Remember that the percentage developed is 25%, current incurred loss is $150, and our initial expectation of ultimate is \$200. In Equation (9), assume that we are using the BF method as our unadjusted estimate of ultimate loss (i.e., $U_i = U_{BF}$ and $U_{MRi} = U_{MRBF}$) and set the weighting factor equal to the percentage developed at time k (i.e., $w_k = p_k$). As will be shown in the next section, this is actually a special case of the Mean-Reverting family – namely the Mean-Reverting Bornhuetter-Ferguson (MRBF) method. From Table 7, which compares the BF method with its Mean-Reverting variant, we see that as actual losses were more than expected, the MRBF method adjusts the BF outstanding reserve/ultimate liability downward back toward initial expectations allowing for a degree of mean-reversion over the future experience period.

Table 7. Simple example comparing BF method with the MRBF method.

BF Method			MRBF Method		
U_{BF}	=	$C_k + (1-p_k)U_0$	U_{MRBF}	=	$U_{BF} - p_k\left(C_k - p_kU_0\right)$
	=	$150 + (1-25\%) \times 200$		=	$300 - 25\% \times \left(150 - 25\% \times 200\right)$
	=	300		=	$300 - 25\% \times 100$
				=	275

Here, although the difference between actual and expected experience is \$100, we only adjust the BF projection by 25% of this amount representing the credibility we assign the degree of mean-reversion. It is this weight, as well as the basis of the *a posteriori* projection, which distinguishes members of the Mean-Reverting family. We explore the link between these two components in the next section.

3.1.2 The credibility formulation[2]

Similar to as was done with the Actual vs. Expected family, we can define a generator function for the Mean-Reverting family that isolates a subset of this family and expresses these members as the credibility-weighted average of the fixed *a priori* expectation and the unadjusted loss reserving method. This generator function is shown below in Equation (11). What is immediately obvious, and to some extent reasonable, given the relationship between the Actual vs. Expected and Mean-Reverting families, is that this formulation is the mirror opposite of the credibility formulation of the Actual vs. Expected family.

$$U_{MRi} = p_k U_0 + (1-p_k)U_i \tag{11}$$

[2] Technically, the generator function formulation presented in Equation (11) is the opposite of a credibility-weighted projection where the estimate of ultimate loss tends toward the complement of credibility U_0 and away from experience as losses develop to ultimate. This is obviously not ideal and will be addressed Section 3.3, where we present an adjusted version of this family which tends toward experience and away from prior expectations as losses develop to ultimate.

Using the same basic loss reserving methods as above, Table 8 shows both the actual vs. expected and credibility formulations of the Mean-Reverting variants for the IE, EA, BF, GB, and CL methods. As in Section 2, each of these methods were derived by plugging the unadjusted method into the generator function in Equation (11) and solving for the weight in the pure formulation shown in Equation (9).

Table 8. Five members of the Mean-Reverting family of loss reserving methods.

Method	Credibility Formulation	Actual vs. Expected Formulation	Weight Function
MRIE	$U_{MRIE} = p_k U_0 + (1-p_k)U_{IE}$	$U_{MRIE} = U_{IE} - 0(C_k - p_k U_0)$	0
MREA	$U_{MREA} = p_k U_0 + (1-p_k)U_{EA}$	$U_{MREA} = U_{AE} - p_k^2(C_k - p_k U_0)$	p_k^2
MRBF	$U_{MRBF} = p_k U_0 + (1-p_k)U_{BF}$	$U_{MRBF} = U_{BF} - p_k(C_k - p_k U_0)$	p_k
MRGB	$U_{MRGB} = p_k U_0 + (1-p_k)U_{GB}$	$U_{MRGB} = U_{GB} - (2p_k - p_k^2)(C_k - p_k U_0)$	$2p_k - p_k^2$
MRCL	$U_{MRCL} = p_k U_0 + (1-p_k)U_{CL}$	$U_{MRCL} = U_{CL} - 1(C_k - p_k U_0)$	1

There are two items of interest concerning Table 8. The first is that the method which is plugged into the generator function in Equation (11) is the same method which is used as the *a posteriori* projection in the actual vs. expected formulation. This is useful, as it reduces the complexity of this family from two free parameters (the *a posteriori* projection and the weight given the actual vs. expected adjustment) to a single free parameter (the *a posteriori* projection) which fully defines members of this family. The second is that the weights given to the actual vs. expected (or mean-reverting) adjustments are identical to the weights given the adjustments in the Actual vs. Expected family; however, the starting points, the *a priori* expectation in terms of the Actual vs. Expected family and the *a posteriori* estimates in terms of the Mean-Reverting family, are different. We explore this symmetry in the next section.

3.2 Selecting a Specific Member

3.2.1 The notion of relative mean-reversion

For completeness, in Table 8 we also show the trivial case of the MRIE, noting that this method reduces to the *a priori* expectation. This is useful, as it highlights that members of the Mean-Reverting family, as with the Actual vs. Expected family, form a spectrum from 0 to 100% weight on the actual vs. expected adjustment.

With that said, it is important when interpreting these methods that the weight given the actual vs. expected adjustment is not mistaken for the degree of mean-reversion. In contrast to the Actual vs. Expected family, where each method applies a different adjustment to the same starting point (i.e., the *a priori* expectation), each member of the Mean-Reverting family applies a different adjustment to a different starting point (i.e., the chosen *a posteriori* estimator). In this regard, each member of the Mean-Reverting family should primarily be considered in relation to its unadjusted variant, rather than with respect to other actuarial methods or some absolute reference.

To explore this concept further, we define the coefficient of mean-reversion as

$$C_{MRi} = \frac{U_i - U_{MRi}}{U_i - U_0}. \tag{12}$$

To understand this equation, note that the denominator expresses the amount by which the unadjusted method deviates from our *a priori* expectation and the numerator expresses the amount by which the Mean-Reverting variant of the unadjusted method pulls the answer back toward the mean or initial expectation.

We derive the coefficient of mean-reversion for the MRBF and MRCL methods in Table 9.

Table 9. Derivation of coefficient of mean-reversion for MRBF and MRCL methods.

MRBF Coefficient of Mean-Reversion			**MRCL Coefficient of Mean-Reversion**		
C_{MRBF}	$=$	$\frac{U_{BF} - U_{MRBF}}{U_{BF} - U_0}$	C_{MRCL}	$=$	$\frac{U_{CL} - U_{MRCL}}{U_{CL} - U_0}$
	$=$	$\frac{U_{BF} - [U_{BF} - p_k(C_k - p_k U_0)]}{C_k + U_0(1 - p_k) - U_0}$		$=$	$\frac{U_{CL} - [U_{CL} - 1(C_k - p_k U_0)]}{C_k / p_k - U_0}$
	$=$	$\frac{p_k(C_k - p_k U_0)}{C_k - p_k U_0}$		$=$	$\frac{(C_k - p_k U_0)}{(C_k - p_k U_0)/p_k}$
	$=$	p_k		$=$	p_k

Note that the coefficient of mean-reversion is the same for both these methods. And indeed, using this definition, we can easily demonstrate that the relative mean-reversion for each of the Mean-Reverting methods shown in Table 8 (or derived via Equation (11)) will always be equivalent and equal to the percentage developed p_k at time k. This makes intuitive sense, given that the mean to which the ultimate loss reverts in Equation (11) is our a priori expectation and the credibility assigned to this initial expectation is p_k.

3.2.2 The notion of absolute mean-reversion

Although each member of the Mean-Reverting family introduces the same degree of mean-reversion relative to its unadjusted variant, this does not necessarily imply that each method has the same absolute mean-reversion. Rather, the absolute mean-reversion of the family (i.e., the degree of negative dependence between current and future losses) depends not just on the mean-reverting adjustment, but also on the degree of dependence between current and future losses in the underlying method. For instance, consider the BF and CL methods. Where the BF method assumes that future losses are independent of losses to date, the CL method assumes a large degree of positive dependence between current and future losses with the unearned reserve leveraged for experience to date.

We can demonstrate this roughly by considering a slightly unusual version of the outstanding reserve $\left(U_i - C_k\right)$ for each of these methods as shown in Table 10.

Table 10. Comparison of the CL and BF estimates of the outstanding reserve.

Method	Outstanding Reserve	Dependence
CL Method	$U_0\left(1-p_k\right)+\left(\frac{1-p_k}{p_k}\right)\left(C_k - p_k U_0\right)$	$p_k \in (0,1] \Rightarrow$ *Positive*
BF Method	$U_0\left(1-p_k\right)+\left(0\right)\left(C_k - p_k U_0\right)$	$p_k \in (0,1] \Rightarrow$ *Independent*

Here we split the outstanding reserve into two components – the "independent" reserve $U_0\left(1-p_k\right)$ which bears no relationship to loss experience C_k and the "dependent" reserve $\left(C_k - p_k U_0\right)$ which is the actual vs. expected adjustment. This is a useful formulation, as we can easily assess the dependence of the outstanding reserve on experience to date. For the BF method, as the dependent reserve is zero, this method assumes future loss experience is fully independent of current loss experience. But for the CL method, as the independent reserve adjustment factor $\left(1-p_k\right)/\,p_k$ is always positive, the CL method assumes that future loss experience is positively dependent on current loss experience adjusting the independent reserve upward.

Now consider Table 11 which shows a similar comparison for the MRCL and MRBF methods.

Table 11. Comparison of the MRCL and MRBF estimates of the outstanding reserve.

Method	Outstanding Reserve	Dependence
MRCL Method	$U_0\left(1-p_k\right)+\left(\frac{1-2p_k}{p_k}\right)\left(C_k - p_k U_0\right)$	$p_k \in \begin{cases} (0,0.5) & \Rightarrow \textit{Positive} \\ 0.5 & \Rightarrow \textit{Independent} \\ (0.5,1] & \Rightarrow \textit{Negative} \end{cases}$
MRBF Method	$U_0\left(1-p_k\right)-\left(p_k\right)\left(C_k - p_k U_0\right)$	$p_k \in \{(0,1] \quad \Rightarrow$ *Negative*

Here it becomes evident that while each Mean-Reverting method does introduce some negative dependence or mean-reversion into its unadjusted variant, the final absolute dependence between future and current loss experience is not necessarily negative. Rather it depends on the interaction between the dependence of future and current loss experience in the underlying method and the strength of the mean-reversion adjustment.

For the MRBF method, because the underlying BF method assumes that current and future loss experience is independent, applying a mean-reverting adjustment to this method will obviously produce estimates of the outstanding reserve which are negatively dependent on experience to date. And from Table 11 we can see that, at all stages of loss development, this is indeed the case. For the MRCL method, however, the result is a little bit trickier, as the CL estimate of the outstanding reserve is positively dependent on losses to date. Thus, the absolute degree of mean-reversion depends on the interaction between the credibility given the negative dependence in the mean-reverting adjustment and the credibility given the positive dependence in the CL method. Specifically, when $p_k < 50\%$, the positive dependence of the CL method dwarfs the mean-reversion adjustment and the estimate of the outstanding liability is still positively dependent on experience to date (however, less so than with the CL method). When $p_k > 50\%$, the mean-reversion adjustment is more influential than the leveraged effect of the CL projection and the estimate of the outstanding liability is negatively dependent on experience to date. And when $p_k = 50\%$, there is balance and the estimate of the outstanding liability and experience to date are largely independent.

The summation of these two sections implies that there are two layers of interpretation regarding the Mean-Reverting family. The first is that each member of the Mean-Reverting family introduces a relative degree of mean-reversion into its underlying variant. The second is that the absolute mean-reversion in the final result depends on the relationship between the underlying method chosen and the credibility given the mean-reversion adjustment. This is a useful result, as these two interpretations begin to hint at a two-step procedure for selecting a member of the Mean-Reverting family. First, select a best unadjusted method, and then, if the situation warrants, adjust that method for some degree of mean-reversion. This is discussed in the next section and will be illustrated using actual data in Section 3.4.

3.2.3 Putting it all together

As mentioned above, the Mean-Reverting family is most useful in situations where either the occurrence or the absence of an event has the opposite impact on the likelihood of similar events in the future. These situations arise when reserving for a variety of lines characterized by total / near-total losses or some notion of risk aging or mortality. Such examples might include marine, crop, credit disability, construction defect, and extended warranty.

For instance, consider an extended warranty policy. As the policy ages, the loss potential generally increases with product wear and tear. And, although we may try to take this into account through our earning pattern, if losses to date are less than expected, it may become necessary to make an adjustment for the increased future loss propensity given the weighted aging of the account. Conversely, if losses are more than expected, then it may become necessary to reduce the future possibility of losses, as policies either exit the portfolio or the products are replaced with newer versions and the weighted age of the account decreases.

When deciding whether to use, say, the BF or CL method or their mean-reverting variants, it is useful to query why actual losses were more or less than expected. Continuing with the example of extended warranty insurance, suppose that there is no discernable reason why losses were more than expected. In this case, we are implying that, although losses were more than expected, we do not expect such trends to continue. Here the BF method is potentially more useful than the CL method, as the estimate of the outstanding liability is independent of current experience to date. However, although losses were more than expected, we might reasonably expect some degree of mean-reversion associated with the replacement of products, and the MRBF method is potentially more useful than the BF method.

On the other hand, suppose that losses were more than expected due to a "catastrophe" event such as a substantial product defect. In this case, the CL method is probably more useful than the BF method, as we should probably expect a higher number of future losses because of the defect. However, in this situation there is still potentially a degree of mean-reversion associated with the policy exit or product replacement decreasing the future propensity to claim on at least that portfolio of the book which has had a loss. In this situation then, the MRCL method is potentially more useful than the CL method.

Generalizing this exposition, selecting a member of the Mean-Reverting family is effectively a two-step process. In situations where losses are more (or less) than expected, we first select the best unadjusted method, given our understanding of the situation and loss drivers. Then, in situations also involving a degree of mean-reversion, we make an adjustment to this method to allow for the decreased (increased) loss potential associated with losses to date being more (or less) than expected.

3.3 Contra-Indications (or the Adjusted Mean-Reverting Family)

Unlike the Actual vs. Expected family, there is one near-fatal flaw to the Mean-Reverting family – namely that as $p_k \to 100\%$, the Mean-Reverting estimate of ultimate loss approaches the initial expectation U_0. In many regards, this is not as significant a problem at younger maturities when losses are not yet fully developed and the *a priori* expectation is an *as*-reasonable if not *more*-reasonable estimate of ultimate than actual losses to date. However, at later maturities, where actual losses approach ultimate losses, this becomes an undesirable characteristic. In fact, it practically becomes a nuisance, because it forces the actuary to define some rule-of-thumb regarding the percentage developed at time k, above which the actuary should not typically rely on a Mean-Reverting method, but below which it is reasonable (that is, if the situation involves some degree of mean-reversion).

This, however, is not necessarily a shortcoming of the Mean-Reverting family. Rather it is a limitation of using a fixed *a priori* expectation which ignores actual experience. And as such, the mirror image of the Mean-Reverting family, the Actual vs. Expected family, offers a simple solution. Rather than using a fixed initial expectation as the mean to which this family reverts, instead use a member of the Actual vs. Expected family of loss reserving methods in place of U_0 in Equation (3) or (4). By doing so, note that as $p_k \to 100\%$, except in the trivial case where $w_i = 0$, the estimate of ultimate loss will tend toward actual.

So far, we have discussed five distinct members of the Actual vs. Expected family and five distinct members of the Mean-Reverting family, and so the above combinations potentially give us twenty-five methods, which is – admittedly – a bit much to digest. Instead, we propose considering just two combinations – the Adjusted MRBF (AMRBF) method, which uses the AEBF method in place of the fixed *a priori* expectation, and the Adjusted MRCL (AMRCL) method, which uses the AECL method in place of the fixed *a priori* expectation. We derive the AMRBF method as

$$\begin{aligned}
U_{MRBF} &= U_{BF} - p_k\left(C_k - p_k U_0\right) \\
U_{AMRBF} &= U_{BF} - p_k\left(C_k - p_k U_{AEBF}\right) \\
&= U_{BF} - p_k\left(C_k - p_k\left[U_0 + p_k\left(C_k - p_k U_0\right)\right]\right) \\
&= U_{BF} - p_k C_k + p_k^2 U_0 + p_k^3 C_k - p_k^4 U_0 \\
&= U_{BF} - \left[C_k\left(p_k - p_k^3\right) - p_k U_0\left(p_k - p_k^3\right)\right] \\
&= U_{BF} - \left(p_k - p_k^3\right)\left(C_k - p_k U_0\right)
\end{aligned} \tag{13}$$

Importantly, from Equation (13), we can see that as $p_k \to 100\%$, $p_k - p_k^3 \to 0$, and thus this method approaches the BF method (which in turn approaches actual as losses develop). However, more interestingly, note that the weight this method places on the mean-reversion relative to the unadjusted projection is given as $p_k - p_k^3$, whereas the amount of weight the MRBF method places on the mean reversion is p_k. In this regard, the AMRBF method not only tends toward actual ultimate losses, but acts as a mechanical rule of thumb, determining the point at which the fixed *a* *priori* estimate of ultimate loss becomes less relevant relative to developed experience. Figure 3 shows this graphically, comparing the amount of weight the AMRBF method gives to the initial expectation relative to the MRBF method.

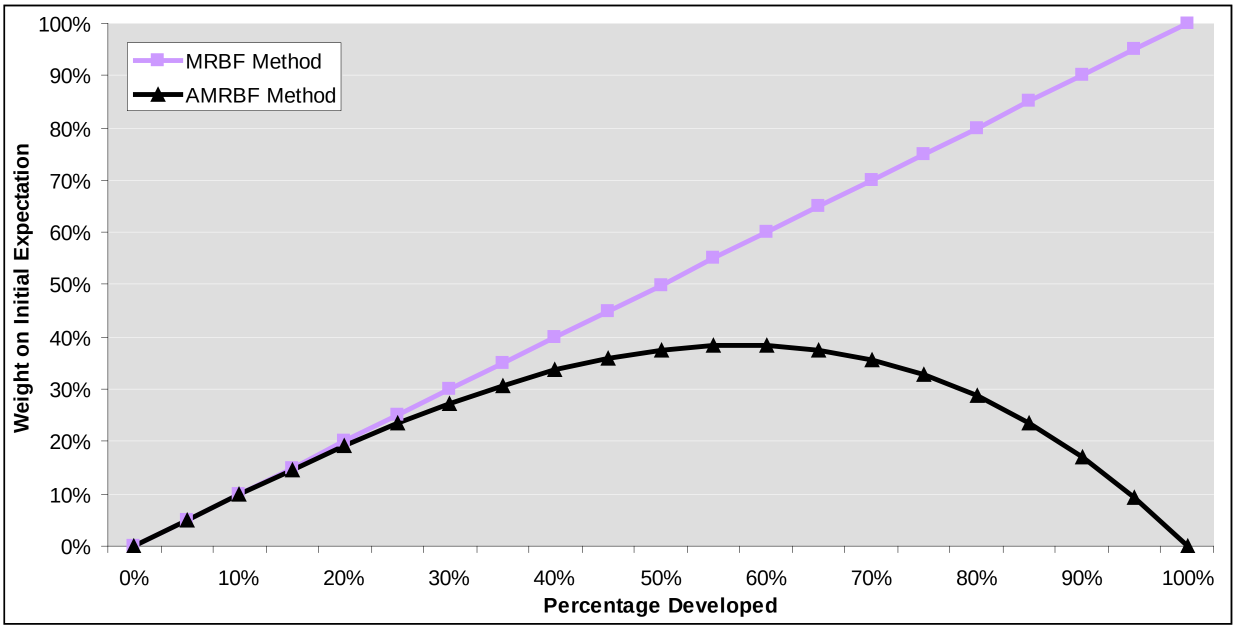

Figure 3. Amount of weight the MRBF and AMRBF methods assign initial expectations.

Similarly, we can derive the AMRCL method as

$$\begin{aligned}
U_{MRCL} &= U_{CL} - 1(C_k - p_k U_0) \\
U_{AMRCL} &= U_{CL} - 1(C_k - p_k U_{AECL}) \\
&= U_{CL} - 1(C_k - p_k[U_0 + 1(C_k - p_k U_0)]) \\
&= U_{CL} - C_k + p_k U_0 + p_k C_k - p_k^2 U_0 \\
&= U_{CL} - [C_k(1 - p_k) - p_k U_0 (1 - p_k)] \\
&= U_{CL} - (1 - p_k)(C_k - p_k U_0)
\end{aligned} \qquad (14)$$

Here, as with the AMRBF method, the AMRCL method tends to the CL method and thus actual losses as the percentage developed tends to 100%.

3.4 Hindsight Testing

To further explore the relevance of the Mean-Reverting family, it is useful to consider the performance of this family in a real-world situation. To do so, we test how the AMRCL and AMRBF methods (as described in the previous section) would have performed historically relative to their bases – the CL and BF methods, respectively. Using crop insurance as an example, we consider data from the Risk Management Agency (RMA) of the United States Department of

Agriculture (USDA), which administrates the Federal Crop Insurance Program (FCIP). Specifically, we consider claims frequency (defined as policies indemnified to total policies) over the ten-year period from 2001 to 2010 in Texas. For reference, the data used is summarized in Table 12.

Table 12. Total policies vs. policies indemnified by month / year from 2001 through 2010 for Texas.

	Total	Policies		Cumulative Policies Indemnified by Month									
Year	Policies	Indemnified	Frequency	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
2001	232	95	41%	7	12	18	31	53	71	82	88	91	95
2002	225	86	38%	10	19	28	39	58	69	78	81	84	86
2003	228	87	38%	3	9	17	25	56	66	78	84	86	87
2004	207	38	19%	5	8	12	16	26	28	30	33	35	38
2005	194	36	18%	1	3	7	13	22	26	29	31	33	36
2006	196	104	53%	15	24	37	55	78	90	95	98	101	104
2007	226	37	17%	7	12	17	21	29	33	35	35	36	37
2008	245	115	47%	15	27	35	51	83	97	102	104	111	115
2009	237	98	42%	17	33	50	60	80	90	94	96	97	98
2010	203	21	10%	1	2	4	6	11	15	17	18	20	21
Total	**2,194**	**718**	**33%**	**81**	**149**	**225**	**318**	**497**	**586**	**641**	**669**	**696**	**718**

Hindsight testing most typically involves projecting historical amounts on an as-if basis and then using the benefit of hindsight to evaluate the performance of these projections. Here, specifically, we projected ultimate claim amounts, using each of the BF, CL, AMRBF and AMRCL methods, for each year at each evaluation month March through December. We then computed the mean-squared error (MSE) by month as the squared difference between projected and actual claims normalized by the actual number of ultimate claims averaged over all years.

Of course, as we didn't actually project ultimate loss amounts at each of the historic points in time, this hindsight test is on a somewhat artificial basis and of course dependent on our selection of the initial expected frequency and frequency development pattern. To these ends, we used 35% as our initial frequency for all years, which appears fairly reasonable given the above ultimate frequencies, and we estimated the development pattern as the volume-weighted average of all years. However, we sensitivity tested the following results based on several different sets of reasonable assumptions and, while the exact estimates of error change, the same key results hold.

Given the two-step process for selecting a member of the Mean-Reverting family, it is useful to first compare the performance of the BF method relative to the CL method to select the best unadjusted method. Figure 4 below plots the normalized MSE for each month averaged across all years for the BF and CL methods. Note that the error is largest when the year is most immature (i.e., March), but as the years age to ultimate (i.e., December), the error tends to zero.

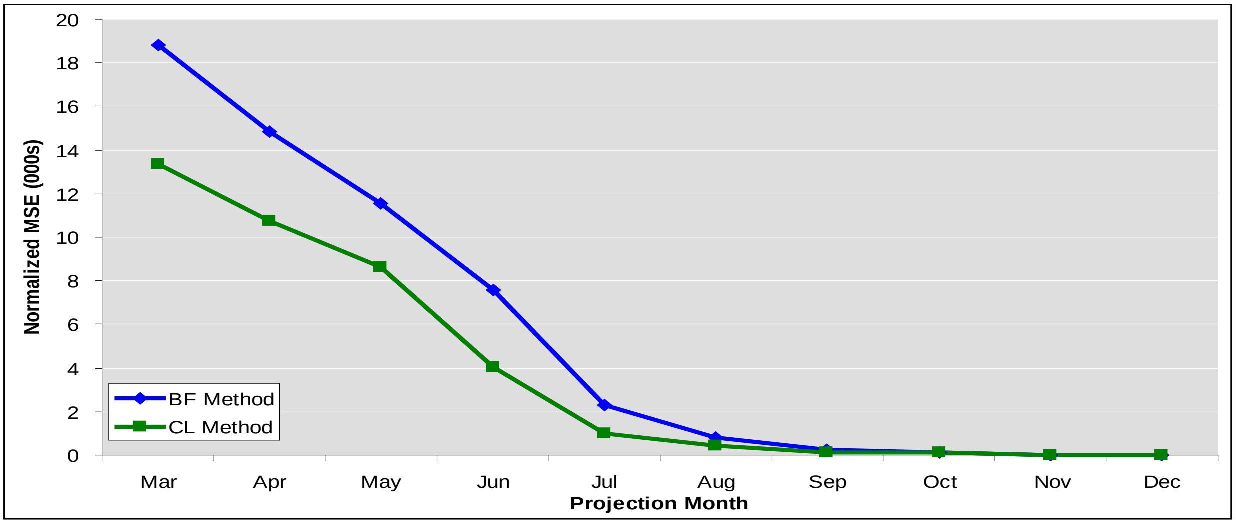

Figure 4. Hindsight testing of the BF and CL methods.

Here we see that the CL method performs better than the BF method. As discussed above, this comparison is actually quite useful as it indicates a key loss driver here – namely, that losses beget losses. For instance, a heavy rainfall in April or a drought in May will certainly cause losses during those months, but they will probably also cause losses in subsequent months due to late reporting or knock-on effects which became more apparent toward harvest. The CL method performs better than the BF method, as it assumes a degree of positive dependence between current and future losses gearing-up future losses to be more than would have initially been expected as experience to date was more than expected.

Now, this analysis may seem slightly at odds with the fundamental message of the Mean-Reverting family, but it isn't. Remember, the Mean-Reverting family does not produce in all situations an absolute level of mean-reversion; rather it applies a mean-reverting adjustment to an unadjusted projection of loss (i.e., the CL or BF method). So in this case, although a dominant loss driver appears to be that losses beget losses, we can now evaluate the AMRCL and the AMRBF methods to assess whether in addition to this market force there is also a degree of mean-reversion at work. The results of this analysis are shown in Figure 5, where Panel (a) shows the AMRCL method relative to the CL method and Panel (b) shows the AMRBF method relative to the BF method.

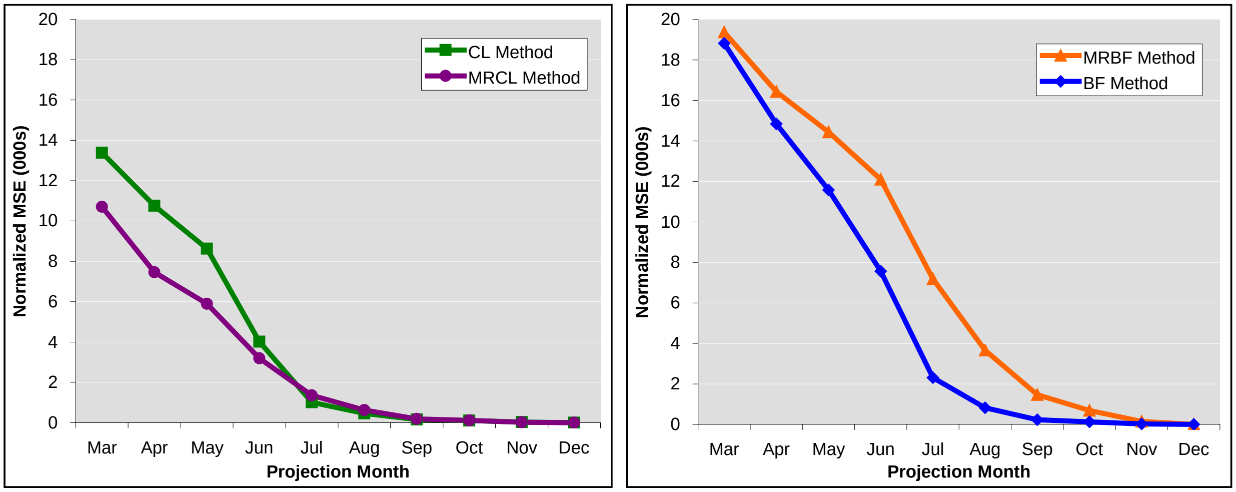

Panel (a) **Panel (b)**

Figure 5. Hindsight testing of the CL method vs. the AMRCL method; and the BF method vs. the AMRBF method.

Considering Panel (a) first, note that the AMRCL method performs substantially better than the CL method. This result indicates that although there is a degree of positive dependence between current and future losses, there is also a degree of mean-reversion where the occurrence (absence) of losses now decreases (increases) the potential of similar losses later. Because of this mean-reversion, the AMRCL method is more accurate than the CL method as well as the BF method by transitivity. This is similar to the example of extended warranty insurance discussed above, where a hypothetical product defect caused both an increase in losses to date as well as a potential uptick in future loss experience, but there was also a degree of mean-reversion associated with policy exit and product replacement.

Panel (b) compares the AMRBF method with the BF method. Here, the BF method is more accurate than the AMRBF method. This is an interesting result, as it indicates in this particular situation that the losses beget losses force is stronger than the mean-reversion force. In order to understand this, note that the BF method assumes that future losses are fully independent of current losses, whereas the MRBF method assumes negative dependence between future and current losses. However, if the mean-reversion force was stronger in this instance, the BF method would be more accurate than the CL method and the AMRBF method would be the most accurate.

4. CONCLUSION

In this paper, we introduced two families of loss reserving methods – the Actual vs. Expected family and the Mean-Reverting family. We showed that the Actual vs. Expected family is useful as an alternative to a fixed *a priori* expectation and when rolling forward prior estimates of ultimate loss. And we showed that Mean-Reverting family is useful in situations where either the occurrence (or absence) of an event decreases (or increases) the likelihood of similar events in the future.

To distill the above into something which is most useful for the practicing actuary, the four key methods to take to take away from this paper are the Actual vs. Expected Bornhuetter-Ferguson (AEBF) method, the Generalized Actual vs. Expected Bornhuetter-Ferguson (GAEBF) method, the Adjusted Mean-Reverting Bornhuetter-Ferguson (AMRBF) method, and the Adjusted Mean-Reverting Chain-Ladder (AMRCL) method. These methods are shown in Table 13.

Table 13. Key methods to take away from this paper.

Method	**Formula**
AEBF	$U_{AEBF} = U_0 + p_k\left(C_k - p_k U_0\right)$
GAEBF	$U_{GAEBF} = U_k = U_{k-1} + \left(\frac{p_k - p_{k-1}}{1-p_{k-1}}\right)\left(\left(C_k - C_{k-1}\right) - \left(\frac{p_k - p_{k-1}}{1-p_{k-1}}\right)\left(U_{k-1} - C_{k-1}\right)\right)$
AMRBF	$U_{AMRBF} = U_{BF} - \left(p_k - p_k^3\right)\left(C_k - p_k U_0\right)$
AMRCL	$U_{AMRCL} = U_{CL} - \left(1 - p_k\right)\left(C_k - p_k U_0\right)$

The AEBF method is a solid alternative to a fixed *a priori* expectation in that the AEBF method credibly updates initial expectations for actual experience balancing responsiveness with stability. Furthermore, the AEBF method can easily be generalized (i.e., the GAEBF) in order to credibly roll forward prior estimates of ultimate loss while balancing responsiveness with stability. The AMRBF method introduces an absolute degree of mean-reversion into projections of ultimate loss and is particularly useful in situations which involve some degree of mean-reversion, but the occurrence (or absence) of losses to date are roughly independent of one another. The AMRCL method introduces a relative, but not always absolute, degree of mean-reversion into projections of ultimate loss and is useful in situations which involve some degree of mean-reversion and the occurrence (or absence) of losses to date are predicated on some underlying force which is expected to effect future events as well.

5. REFERENCES

[1] Bornhuetter, Ronald L. and Ronald E. Ferguson, "The Actuary and IBNR," *PCAS*, **1972**, Vol. LIX, 181-195.

[2] Mack, Thomas, "Credible Claims Reserve: The Benktander Method," ASTIN *Bulletin*, **2000**, Vol. 30, No. 2, 333-347.

[3] Mahler, Howard C. and Curtis Gary Dean, "Credibility," Chapter 8 in *Foundations of Casualty Actuarial Science*, **2001**, 4th Edition.

Appendix A. Author's Note

To better understand the substance of this paper, it is useful to understand the motivation for writing it. Although only briefly alluded to in the text, the Actual vs. Expected Bornhuetter-Ferguson (AEBF) method is equivalent to the Experience Adjustment (EA) method. The original motivation for this paper was to present the EA method as an artificially intelligent version of a fixed *a priori* expectation. However, it soon became apparent that the EA method as well as the IE and BF method could be generalized as members of the same family indexed using the weight each method assigns to an actual vs. expected adjustment. This seemed to be a more useful and pliable result as it not only defines the EA method, but also presents an entire spectrum of methods which take as their seed a fixed *a priori* expectation and update that expectation for experience with varying degrees of responsiveness.

Then, while writing this paper, the question of reserving for a crop insurance program arose. Specifically a situation where floods had knocked out a large portion of crops and an adjustment was needed to make an allowance for the reduced future potential of losses within the loss projections. Although the obvious solution involves making an adjustment to the unearned exposure, given the importance of mechanizing loss reserving techniques as well as the importance of mean-reversion in many traditional actuarial time-series models, the Mean-Reverting Chain-Ladder (MRCL) and Mean-Reverting Bornhuetter-Ferguson (MRBF) methods were born. Again, similar to the Actual vs. Expected family, it soon became apparent that these methods could be generalized into a family of loss reserving methods which interestingly enough bore a striking resemblance to the Actual vs. Expected family. Hence this paper, and the title: *Two Symmetric Families of Loss Reserving Methods.*

Appendix B. Supporting Excel File

Included with this paper is an Excel file showing how to program these various methods.

The first tab – Projection – compares the unadjusted projections (i.e., the BF method) with their Actual vs. Expected variants (i.e., the AEBF method), Mean-Reverting variants (i.e., the MRBF method) and Adjusted Mean-Reverting Variants (i.e., the AMRBF method).

The second tab – Roll-Forward – shows how to extend the AEBF method in order to roll-forward prior estimates of ultimate loss for development during interim periods (i.e., the GAEBF method). A comparison is also done to roll-forwards using the IE method which gives 0% credibility to actual experience in the period and the BF method which gives 100% credibility to actual experience in the period. In contrast, the AEBF gives partial credibility to the experience in the period proportionate to the expected percentage of developed loss in the period.

The third and fourth tabs – Example_BFvsMRBF and Example_CLvsMRCL – contain the calculations underlying the hindsight testing performed in Section 3.4.

Acknowledgment

I would like to thank Ernest Wilson, Brad Andrekus and Joth Tupper for their thorough reviews. This paper is much stronger because of their helpful suggestions and commentary.

Biography of the Author

Andy Staudt is a consultant with Towers Watson in London. He is a Fellow of the Casualty Actuarial Society (CAS), a member of the American Academy of Actuaries (AAA) and has a Masters in Statistics from the University of California at Berkeley. He is also a Fellow in the Institute of Actuaries (FIA) through mutual recognition. He can be reached at: andy.staudt@towerswatson.com or +44 (0) 207 886 5146.