Reserving with Incomplete Exposure Information

David R. Clark, FCAS, MAAA

Abstract. This paper outlines a reserving method that allows the actuary to use exposure information, such as onlevel premium, even if that information is only available for a limited number of years. The method is a simple blending of methods already in wide use, but can be shown to be based on a common underlying statistical model. The paper provides an overview of the Over-Dispersed Poisson model, and how it relates to Multiplicative LDF, Cape Cod, and Bornhuetter-Ferguson methods.
Motivation. The reserving actuary may have reliable exposure information (e.g., onlevel premium) for only a few recent years of data, rather than for the full historical period for which reserves need to be set.
Method. This incomplete exposure information can still be used, by implementing a hybrid reserving method equivalent to the Cape Cod method for the recent years and the Multiplicative LDF method for older years.
Results. We show how common reserving methods can be derived from a single statistical model, and then show how these methods are best combined when partial information is available.
Conclusions. This is a practical solution to the problem of stabilizing loss projections for recent accident years, incorporating available rate change information, and being responsive to actual loss emergence.

Keywords. Reserving, GLM, Chain ladder, Cape Cod, Bornhuetter-Ferguson.

1. INTRODUCTION

The purpose of this paper is to outline a method for estimating a stable reserve for immature years on long-tailed lines of business.

In order to bring more stability to these reserve estimates, it is helpful to bring in an exposure base that is proportional to expected loss by year. Optimally, this exposure base would be something like payroll or sales, but more commonly only historical premium is available. Historical premium is not directly applicable because of significant changes in rate adequacy over time – a phenomenon called the "insurance cycle." Instead we need to adjust the historical premium to an "onlevel premium" basis that is truly proportional to the expected losses by year. Unfortunately, the rate level indices required to make this adjustment may only be available for a limited number of years.

We will propose that even this limited information can be used in the reserve review in a straightforward way, with a method that is a combination of the Multiplicative LDF (a.k.a. Chain ladder) and Cape Cod (a.k.a. Stanard-Bühlmann) methods to form a single unified method. This unified method can be shown to be a best use of the available data and to be consistent with the other methods because they are all relying on the same underlying statistical model.

1.1 Research Context

There have been several past papers surveying statistical models applied to insurance loss development (payment or reporting) patterns. Recent examples are the CAS Working Party on Reserve Variability (2005), and the classification paper by Schmidt (2006). This prior research has aided greatly in viewing the loss development phenomenon from a statistical viewpoint; and showing connections between various models.

1.2 Objective

We will not intend to break new ground from a theoretical standpoint. Instead, we will build on the theory already established and draw some important practical implications. Specifically, we will show how best to incorporate limited exposure information into a reserve review in a consistent manner. By grounding this method in sound theory, we can show how it is consistent with current models and how it is an improvement over some popular techniques such as the Bornhuetter-Ferguson method.

What is new in this paper is the demonstration that a single unified method, which combines a Multiplicative LDF for older years and Cape Cod for more recent years, is built upon a single statistical model. The result is that limited exposure information can be incorporated for the years in which it is available.

1.3 Outline

The remainder of the paper proceeds as follows.

Section 2.1 will provide a description of the reserving problem faced for long-tailed business. We will introduce a numerical example to illustrate this problem.

Section 2.2 will give some basic definitions to set the groundwork for addressing the problem.

Section 3.1 will describe the Over-Dispersed Poisson (ODP) model as the basic structure underlying all of the methods to be discussed.

Section 3.2 will look at three methods in common use; and how they relate to the ODP model.

Section 3.2.1 The Multiplicative LDF method (a.k.a., Chain ladder)

Section 3.2.2 The Cape Cod method (a.k.a. Stanard-Bühlmann)

Section 3.2.3 The Bornhuetter-Ferguson method

Section 3.3 will look at a unified method that combines the Multiplicative LDF and Cape Cod methods to incorporate limited exposure information.

Section 4 gives further discussion of practical issues of the Unified method, including issues in creating an appropriate exposure index.

2. PRELIMINARIES: THE RESERVING PROBLEM

We now proceed to give a more detailed description of the reserving problem to be addressed.

2.1 A Realistic Example

You are a reserving actuary reviewing the medical malpractice line of business. You will be working with an eight-year development triangle of cumulative paid loss data as shown below.[1]

Cumulative Paid Loss Triangle

AY	12	24	36	48	60	72	84	96
1999	257	1,143	2,402	3,478	4,456	5,080	5,284	5,481
2000	266	1,167	2,604	3,897	4,522	5,299	5,464	
2001	347	1,400	2,839	3,984	5,131	5,427		
2002	279	1,186	2,450	3,858	4,417			
2003	245	992	2,508	3,047				
2004	220	1,269	1,714					
2005	214	829						
2006	215							

This data shows a development pattern in which relatively little loss is paid in the first year. As a benchmark, you calculate standard chain ladder development factors, which confirm that only about 5.4% (=1/18.520 as shown below) of the loss would be paid as of the first twelve months—

assuming that there is no tail beyond the eighth year. Based on this, accident year 2006 seems too immature to expect the loss development method to yield a reliable result.

Development Factors (age-to-age link ratios)

AY	12-24	24-36	36-48	48-60	60-72	72-84	84-96	96-Ult
1999	4.447	2.101	1.448	1.281	1.140	1.040	1.037	
2000	4.387	2.231	1.497	1.160	1.172	1.031		
2001	4.035	2.028	1.403	1.288	1.058			
2002	4.251	2.066	1.575	1.145				
2003	4.049	2.528	1.215					
2004	5.768	1.351						
2005	3.874							
Wtd Avg	4.369	2.028	1.427	1.217	1.120	1.036	1.037	
LDF	18.520	4.239	2.090	1.465	1.203	1.074	1.037	1.000

In the past, reserves for immature years were often set using the Bornhuetter-Ferguson method, with a plan loss ratio used as the a priori expected value. However, in researching old reserve reviews, you have found that the plan loss ratio has consistently been set at about a 60% ELR, plus or minus a few points. By contrast, the actual experience has displayed a long-term cyclical pattern with a much wider range of loss ratios.

AY	Earned Premium	Latest Diagonal	LDF	Ultimate Loss	Loss Ratio
1999	5,400	5,481	1.000	5,481	101.5%
2000	5,900	5,464	1.037	5,668	96.1%
2001	6,500	5,427	1.074	5,829	89.7%
2002	8,500	4,417	1.203	5,315	62.5%
2003	10,200	3,047	1.465	4,464	43.8%
2004	11,000	1,714	2.090	3,582	32.6%
2005	11,300	829	4.239	3,514	31.1%
2006	11,500	215	18.520	3,982	34.6%
Total	70,300	26,594		37,835	53.8%

The 60% ELR might have been right for some periods (as Lewis Carroll observed: even a stopped clock is right twice a day…), but in general it has not proved to be an accurate number.

[1] This triangle is based on a section of industrywide medical malpractice data, but has been modified. The example is intended to be realistic, if somewhat better behaved than most accounts, but should not be used for any purpose other

Instead we have some evidence that the improving loss ratios from 1999 to 2006 were due in large part to significant rate increases. We know that this information should be used in the analysis, but unfortunately we only have a reliable monitor for rate changes starting in 2002.

What do we do?

2.2 Laying the Groundwork for a Solution

Before giving a detailed explanation of the models available to us and a proposed solution to the example above, it is worth carefully defining some key concepts.[2]

Model = A mathematical or empirical representation of a specified phenomenon

Method = A systematic procedure for estimating the unpaid claims

The "Model" is a mathematical description of the form of the world that we are analyzing, though with simplifying assumptions, such as the assumption that all accident years have the same expected loss development pattern.

The "Method" is the step-by-step procedure, or algorithm, that a person will follow to get from the original data to a final numerical result. In our insurance example above, we applied the chain ladder method to calculate our ultimate loss ratios.

Some may ask: why bother defining a model at all? Why not just select a method that seems reasonable and leave it there? There are three reasons:

1) A model gives criteria for deciding which of several possible methods is the "best" one (e.g., criteria of unbiasedness and minimum variance).
2) A model forces us to make all of our assumptions explicit so that they can be tested (e.g., with residual plots and goodness-of-fit criteria).
3) A model provides the theory for creating ranges around our reserve estimate (either

---

than illustrating the ideas in this paper.

[2] These two definitions come from Actuarial Standard of Practice No. 43; see Shapland (2007) for a more rigorous definition of these terms.

standard deviation or percentile distributions).

The 2005 CAS Working Party on Reserve Variability gives a more complete explanation of these reasons for creating a model. For the present paper, the primary purpose of introducing the mathematical model will be to show the "family relationship" of the methods presented.

Two more concepts need to be introduced before we proceed with our model.

Over-Parameterization = when we have too few data points relative to the number of model parameters

Model Constraints = user-supplied information that sets parameters, or relationships between parameters, rather than having them estimated from the data

The concept of over-parameterization is sometimes referred to as over-fitting or responding to the noise in the data rather than the signal. This can be a significant problem in the loss reserving context where we are working with data summarized into the triangle format. Constraining the model parameters is one way of reducing the instability from over-parameterizing and will be key to understanding the differences in the methods that we discuss below.

3. A FAMILY OF RESERVING MODELS AND METHODS

We now turn to a model that provides a framework for all of the familiar reserving methods. It points to a useful solution to our particular problem.

3.1 The Over-Dispersed Poisson (ODP) Model

The model presented here is derived from the theory of generalized linear models (GLM). GLM theory is an expansion of the theory of linear regression that allows for a broader category of error distributions beyond the normal Gaussian distribution, and also allows for the linear relationship of independent variables to be transformed by a "link function" in predicting the dependent variable.[3]

The structure of our model will be a multiplicative combination of accident year (y) and development period (d) factors. The dependent variable that we are attempting to fit will be the incremental loss for a given accident year in a given development period, and will be denoted $c_{y,d}$. For our example, this will be referred to as incremental paid, but the theory could be equally applied to reported data.

$$E(c_{y,d}) = \mu_{y,d} = v_y \cdot ELR \cdot \beta_d \tag{3.1.1}$$

Within this formula, the parameter v_y is an exposure or volume measure by accident year that is proportional to ultimate loss. This can be thought of as onlevel premium, though Section 4 of this paper will give a more detailed discussion as to how to create the measure. The ELR is an expected loss ratio, which represents the ratio of expected ultimate loss to the exposure measure. Because the exposures v_y already vary by accident year in proportion to expected loss, we only need a single value for ELR. The last parameter β_d is the development period relativity and may be thought of as the percent paid during a given calendar year.

This type of multiplicative combination of independent parameters indicates a log-link within GLM. That is, we would need to take logarithms of each side of the equation in order to transform the problem into a linear form.

Next, we will assume that the expected variance of an actual point from the expected value is in proportion to the expected value. The variance-to-mean ratio is represented as a dispersion parameter ϕ.

$$\mathrm{Var}(c_{y,d}) = \phi \cdot E(c_{y,d}) = \phi \cdot \mu_{y,d} \tag{3.1.2}$$

The GLM framework makes use of distributions within the exponential family for the error

[3] See Mildenhall (1999) for a good introduction to GLM in general, or Renshaw & Verrall (1998) for the GLM directly corresponding to this reserving application.

function. The assumption that the variance is proportional to the mean uniquely identifies the distribution as Poisson. The Poisson distribution is defined on the positive integers, $\{0,1,2,3,\cdots\}$, with variance equal to its mean, but this is generalized to the over-dispersed Poisson (ODP) model to be defined on multiples of the dispersion parameter, $\{0\phi,1\phi,2\phi,3\phi,\cdots\}$.[4]

With this model defined, the maximum likelihood estimates for the parameters can be found. We can actually do this by maximizing the quasi-log-likelihood (QLL) function,[5] a simplified version of the log-likelihood that does not depend on the dispersion parameter ϕ.

$$QLL \quad = \quad \sum_{y=1}^{n}\sum_{d=1}^{n+1-y}\left\{\ln\left(\mu_{y,d}\right)\cdot c_{y,d} - \mu_{y,d}\right\} \quad = \quad \sum_{y=1}^{n}\sum_{d=1}^{n+1-y}\ln\left(v_y \cdot \text{ELR} \cdot \beta_d\right)\cdot c_{y,d} - v_y \cdot \text{ELR} \cdot \beta_d \tag{3.1.3}$$

We maximize the quasi-log-likelihood by solving for the parameters that set all of the derivatives equal to zero. For example:

$$\frac{\partial\, QLL}{\partial\, \beta_d} \quad = \quad 0 \qquad \forall\, d \tag{3.1.4}$$

Taking these derivatives guarantees that totals of the fitted losses in each column (development age) are equal to the actual losses. The model may therefore be described as unbiased.[6]

$$\sum_{y=1}^{n+1-d} c_{y,d} \quad = \quad \sum_{y=1}^{n+1-d} v_y \cdot \text{ELR} \cdot \beta_d \qquad \forall\, d \tag{3.1.5}$$

---

[4] Venter (2007) prefers to call this the Poisson-constant-severity (PCS) model rather than ODP, because it can be interpreted as a collective risk model in which the number of claims follows a Poisson distribution, and every claim amount is the same value. However, there is no need to force this interpretation; we can simply view it as a discretized aggregate loss model for a given mean and variance.

[5] See Renshaw and Verrall (1998) for the full detail on this. They also note "We find it easiest to retain the assumption that the data have a Poisson distribution at the moment, although in all that follows in this section it is only the form of the likelihood which is important."

[6] The unbiasedness of row and column parameters as seen in "balancing" their totals may be familiar from the problem of classification ratemaking as described in Mildenhall (1999). More rigorously, we define unbiasedness as a characteristic of an estimator whose expected value is equal to the expected value of the random variable. That is, $E\left(\sum_{y=1}^{n+1-d} c_{y,d}\right) = E\left(\sum_{y=1}^{n+1-d} v_y \cdot \hat{\text{ELR}} \cdot \hat{\beta}_d\right)$. This means that the total of the fitted values corresponding to the observed payments will be unbiased; this does not mean that the estimated reserve for the future periods will also be unbiased (cf. Taylor 2003).

Depending on the further constraints on v_y and ELR, certain row totals will also have fitted values that equal the actual values. Our choice of reserving method will depend on how we define these constraints.

3.2 Common Methods – Based on Constraining the ODP Model

Having defined the basic ODP model, we proceed to show how it is related to three familiar reserving methods.

3.2.1 The Multiplicative LDF Method

We begin with a fully unconstrained model, for which we assume that the vector of exposure measures is not available and must be estimated from the data in the development triangle. The exposure values v_y and ELR are therefore considered parameters to be estimated by the model. We start by defining:

$$\alpha_y = v_y \cdot \text{ELR} \tag{3.2.1}$$

Then we need to add a fitting criterion that the derivative of the QLL with respect to each α_y is set equal to zero.

$$\frac{\partial\, \text{QLL}}{\partial\, \alpha_y} = 0 \qquad \forall\, y \tag{3.2.2}$$

Taking these derivatives guarantees that the row totals of fitted and actual values are equal for every accident year.

$$\sum_{d=1}^{n+1-y} c_{y,d} = \sum_{d=1}^{n+1-y} \alpha_y \cdot \beta_d \qquad \forall\, y \tag{3.2.3}$$

An easy way of estimating the α and β parameters for this model is to use the chain-ladder method of loss development factors. The parameters α_y represent the ultimate loss by year; the parameters β_d are a function of the weighted average LDFs.

AY *y*	Ult. Loss (alpha)	LDF	% of Ult	Incremental % (beta)
1999	5,481	1.000	100.00%	3.59%
2000	5,668	1.037	96.41%	3.31%
2001	5,829	1.074	93.10%	10.00%
2002	5,315	1.203	83.10%	14.84%
2003	4,464	1.465	68.26%	20.41%
2004	3,582	2.090	47.85%	24.26%
2005	3,514	4.239	23.59%	18.19%
2006	3,982	18.520	5.40%	5.40%
			Total of Betas:	100.00%

A simple inspection of the actual and fitted incremental triangles will confirm that both the row and column totals are equal.

Actual Incremental Payments

AY	0-12	12-24	24-36	36-48	48-60	60-72	72-84	84-96
1999	257	886	1,259	1,076	978	624	204	197
2000	266	901	1,437	1,293	625	777	165	
2001	347	1,053	1,439	1,145	1,147	296		
2002	279	907	1,264	1,408	559			
2003	245	747	1,516	539				
2004	220	1,049	445					
2005	214	615						
2006	215							
Total:	2,043	6,158	7,360	5,461	3,309	1,697	369	197

Fitted Incremental Payments

AY	0-12	12-24	24-36	36-48	48-60	60-72	72-84	84-96
1999	296	997	1,330	1,119	814	548	181	197
2000	306	1,031	1,375	1,157	841	566	188	
2001	315	1,060	1,414	1,190	865	583		
2002	287	967	1,289	1,085	789			
2003	241	812	1,083	911				
2004	193	652	869					
2005	190	639						
2006	215							
Total:	2,043	6,158	7,360	5,461	3,309	1,697	369	197

An important observation from this exercise is that we have set the tail factor at age 96 months equal to 1.000. That is, we are assuming that there is no further development beyond the eighth year. In fact, this is merely done by convention – we can include any tail factor that we would like beyond the eighth year. We include a tail factor by dividing all of our β parameters by the selected 96-ultimate LDF, and then also multiplying all of the α parameters by the same amount. The cross-product will produce fitted values equal to the model above.

What this tells us is that the fully unconstrained model provides us with no information about development beyond the periods in the historical data.[7]

A second observation from this unconstrained model is that, while we usually think of it in multiplicative terms, it can equivalently be considered an additive model:

Ultimate Loss	=	(Paid Loss) × LDF
Ultimate Loss	=	(Paid Loss) + (Expected Ultimate)×(1-1/LDF)
	where	Expected Ultimate = (Paid Loss)×LDF

A final observation is that our example includes 36 actual data points, but those 36 data points are estimating 15 parameters (eight accident year factors plus seven development factors). This gives us few data points per parameter and, therefore, should be described as an over-parameterized model.

3.2.2 The Cape Cod Method

As noted above, the fully unconstrained model that produces the chain-ladder method has a problem with over-parameterization. We therefore move to a model that adds more constraints, by introducing an exposure measure that forces a relationship between the accident year ultimates.

$$\frac{\text{Expected Ultimate Loss in Year } i}{\text{Expected Ultimate Loss in Year } j} = \frac{\alpha_i}{\alpha_j} = \frac{v_i}{v_j} \qquad \forall\, i, j \tag{3.2.4}$$

Because the exposure or volume measures are supplied by the user, we only need to estimate the parameter ELR instead of the full vector of α_y. The Maximum Likelihood Estimator (MLE) for

[7] One way to fit a tail factor to the data is to constrain the model by assuming that all of the β s follow a known development pattern form. This is the model outlined in Clark (2003), but will not be addressed here.

the ELR is found by setting the derivative of the quasi-log-likelihood function (QLL) equal to zero.

$$\frac{\partial \,\text{QLL}}{\partial \,\text{ELR}} = \frac{\partial \sum_{y=1}^{n} \sum_{d=1}^{n+1-y} \left\{ \ln\left(v_y \cdot \text{ELR} \cdot \beta_d\right) \cdot c_{y,d} - v_y \cdot \text{ELR} \cdot \beta_d \right\}}{\partial \,\text{ELR}} = 0 \tag{3.2.5}$$

This criterion results in a requirement that the sum of all the losses in the entire triangle is the same for fitted and actual values.

$$\sum_{y=1}^{n} \sum_{d=1}^{n+1-y} c_{y,d} = \sum_{y=1}^{n} \sum_{d=1}^{n+1-y} v_y \cdot \text{ELR} \cdot \beta_d \tag{3.2.6}$$

This does not add anything to our MLE criteria, since we had already required that column totals would be equal.

The method for estimating model parameters is:

1) Estimate an incremental loss ratio IncrLR_d for each development period:

$$\text{IncrLR}_d = \frac{\sum_{y=1}^{n+1-d} c_{y,d}}{\sum_{y=1}^{n+1-d} v_y} \qquad \forall d \tag{3.2.7}$$

2) Set the ELR as the sum of the incremental loss ratios:

$$\text{ELR} = \sum_{d=1}^{n} \text{IncrLR}_d \tag{3.2.8}$$

3) Set the development pattern parameters such that $\sum_{d=1}^{n} \beta_d = 1$:

$$\beta_d = \frac{\text{IncrLR}_d}{\text{ELR}} = \frac{\sum_{y=1}^{n+1-d} c_{y,d}}{\sum_{y=1}^{n+1-d} v_y \cdot \text{ELR}} \qquad \forall d \tag{3.2.9}$$

With this procedure, we accomplish the goal of having all of the column totals for the fitted triangle match those of the actual triangle; therefore the results are the maximum likelihood estimates.

We have been again assuming that there is no "tail" beyond the last age represented in the triangle. As with the Multiplicative LDF method, this is only by convention, and we can introduce any tail factor we wish by re-scaling the β and ELR parameters.

$$\beta_d \to \frac{\beta_d}{\mathrm{LDF}_n} \qquad \text{so that} \qquad \mathrm{LDF}_n = \frac{1}{\sum_{d=1}^{n} \beta_d} \tag{3.2.10}$$

The original ELR is then multiplied by the selected tail LDF_n to produce a final ELR.

$$\mathrm{ELR} = \left(\mathit{Original}\ \mathrm{ELR}\right) \cdot \mathrm{LDF}_n \tag{3.2.11}$$

The key concept to note is that the ELR and tail LDF_n are interdependent. If we change one of them, then the other will also need to change. This concept will be critical when we examine the Bornhuetter-Ferguson method.

In order to perform these calculations, we must first create an exposure index covering all of the accident years in the experience period. We saw above that the ultimate loss ratios were not constant by year, and so we cannot assume that historical premium is a good measure of exposure. We will instead make use of an onlevel factor to adjust for changes in rate adequacy. This way we can create a surrogate exposure base.

AY y	Earned Premium	Onlevel Factor	Exposures v_y
1999	5,400	2.200	11,880
2000	5,900	2.050	12,095
2001	6,500	1.850	12,025
2002	8,500	1.400	11,900
2003	10,200	1.200	12,240
2004	11,000	1.100	12,100
2005	11,300	1.050	11,865
2006	11,500	1.050	12,075
Total	70,300		96,180

The exposures v_y are estimated as the historical earned premium times the onlevel factor. These exposures are now assumed to be proportional to the ultimate expected losses by accident year and can be used in formula 3.2.7 to estimate the preliminary development parameters.

Actual Incremental Payments divided by Exposure

AY	0-12	12-24	24-36	36-48	48-60	60-72	72-84	84-96
1999	2.16%	7.46%	10.60%	9.06%	8.23%	5.25%	1.72%	1.66%
2000	2.20%	7.45%	11.88%	10.69%	5.17%	6.42%	1.36%	
2001	2.89%	8.76%	11.97%	9.52%	9.54%	2.46%		
2002	2.34%	7.62%	10.62%	11.83%	4.70%			
2003	2.00%	6.10%	12.39%	4.40%				
2004	1.82%	8.67%	3.68%					
2005	1.80%	5.18%						
2006	1.78%							
IncrLR:	2.12%	7.32%	10.19%	9.08%	6.91%	4.71%	1.54%	1.66%
Cumul:	2.12%	9.45%	19.63%	28.71%	35.62%	40.34%	41.88%	43.53%
Beta	**4.88%**	**16.82%**	**23.40%**	**20.86%**	**15.87%**	**10.83%**	**3.54%**	**3.81%**
Cumul:	4.88%	21.70%	45.10%	65.96%	81.83%	92.66%	96.19%	100.00%
LDF	20.495	4.609	2.217	1.516	1.222	1.079	1.040	1.000

These numbers are calculated additively rather than via chain ladder link ratios but the calculations are still very straightforward. The ELR to onlevel premium is calculated directly as 43.53% by summing the preliminary incremental loss ratios.

We can also calculate an LDF from the β s. However, this development pattern is not exactly equal to that produced by the chain ladder method. The key reason for this difference is that we are now making use of more information. For example, the 2006 year has loss as of 12 months of $215, which would not affect the chain ladder calculation (no link ratio is calculated from the 2006 year), whereas it does affect the result for the constrained model.

The next step is to use these parameters to project the ultimate losses by year. This is done with an additive formula.

$$\text{Ultimate Loss} \quad = \quad (\text{Paid Loss}) + (\text{Expected Ultimate}) \times (1\text{-}1/\text{LDF})$$

$$\text{where} \quad \text{Expected Ultimate} = \text{Exposure} \times \text{ELR}$$

AY y	Exposures v_y	ELR	Expected Ultimate	LDF	IBNR% 1-1/LDF	Latest Diagonal	Final Ultimate	Final L / R
1999	11,880	43.53%	5,172	1.000	0.00%	5,481	5,481	46.14%
2000	12,095	43.53%	5,265	1.040	3.81%	5,464	5,665	46.83%
2001	12,025	43.53%	5,235	1.079	7.34%	5,427	5,811	48.33%
2002	11,900	43.53%	5,181	1.222	18.17%	4,417	5,358	45.03%
2003	12,240	43.53%	5,329	1.516	34.04%	3,047	4,861	39.71%
2004	12,100	43.53%	5,268	2.217	54.90%	1,714	4,606	38.07%
2005	11,865	43.53%	5,165	4.609	78.30%	829	4,874	41.08%
2006	12,075	43.53%	5,257	20.495	95.12%	215	5,215	43.19%
Total	96,180	43.53%	41,871			26,594	41,871	43.53%

We may note that this is the same calculation that is often thought of as the Bornhuetter-Ferguson method, except that the ELR has been estimated from the data rather than from some a priori input.

This method can be equivalently applied by showing the ELR as the ratio of the latest diagonal of loss divided by the exposure corresponding to the expected loss-to-date. This is the format typically seen in the Cape Cod method, as shown below.

AY y	Exposures v_y	LDF	Expos / LDF	Latest Diagonal	Ultimate L / R	
1999	11,880	1.000	11,880	5,481	46.14%	
2000	12,095	1.040	11,634	5,464	46.96%	
2001	12,025	1.079	11,142	5,427	48.71%	
2002	11,900	1.222	9,737	4,417	45.36%	
2003	12,240	1.516	8,073	3,047	37.74%	
2004	12,100	2.217	5,457	1,714	31.41%	
2005	11,865	4.609	2,574	829	32.20%	
2006	12,075	20.495	589	215	36.49%	
Total	96,180		61,088	26,594	43.53%	= 26,594 / 61,088

This result is significant because it derives from the same underlying ODP model as we used for the Multiplicative LDF method. The only difference is that we have added a constraint that forces a certain behavior in the expected ultimate losses.

As with the Multiplicative LDF method, this Cape Cod method tells us nothing about development beyond the eight years in the historical data. We can again introduce a tail factor to change all of our β parameters, with an exact offsetting change to the ELR.

3.2.3 The Bornhuetter-Ferguson (BF) Method

As noted in the previous section, the Cape Cod method looks very much like a traditional Bornhuetter-Ferguson (BF) method, except that in the Cape Cod method the ELR is estimated from the data itself instead of being supplied by the analyst.

The BF method was originally created as a means of enforcing stability in the IBNR loss reserve estimate. As was stated in the original 1972 paper:

> The decision as to whether to develop the reserve as a direct function of case incurred losses or as a function of expected losses turns on the expected volatility of the data. If the data are extremely thin, the presence or absence of several large losses will impact greatly on the IBNR reserves if the reserve is a function of the case incurred.

This original quote implies an either/or decision: the IBNR reserve is either a function of case incurred losses or a function of expected losses. The GLM framework allows us to incorporate both sources of information in a single consistent model. We will start with the more general model, which incorporates the ELR into a GLM, and then move on to how the BF method is traditionally applied in practice.

For our example, let us suppose that the analyst has selected a 50% ELR for use in the BF method. To calculate the β parameters in this constrained model, we perform the same calculation as we used in the Cape Cod method, except that the denominator is the exposures times our selected 50% ELR.

$$\beta_d = \frac{\sum_{y=1}^{n+1-d} c_{y,d}}{\sum_{y=1}^{n+1-d} v_y \cdot \text{ELR}} \qquad \forall\, d \tag{3.2.12}$$

The form shown in formula 3.2.12 is the same as the pattern recommended in Mack (2006) as most consistent with the BF method.

Actual Incremental Payments divided by Exposure times ELR of 50%

AY	0-12	12-24	24-36	36-48	48-60	60-72	72-84	84-96
1999	4.33%	14.92%	21.20%	18.11%	16.46%	10.51%	3.43%	3.32%
2000	4.40%	14.90%	23.76%	21.38%	10.33%	12.85%	2.73%	
2001	5.77%	17.51%	23.93%	19.04%	19.08%	4.92%		
2002	4.69%	15.24%	21.24%	23.66%	9.39%			
2003	4.00%	12.21%	24.77%	8.81%				
2004	3.64%	17.34%	7.36%					
2005	3.61%	10.37%						
2006	3.56%							
Beta	**4.25%**	**14.64%**	**20.38%**	**18.16%**	**13.82%**	**9.43%**	**3.08%**	**3.32%**
Cumul:	4.25%	18.89%	39.27%	57.43%	71.25%	80.67%	83.75%	87.07%
LDF	23.539	5.293	2.547	1.741	1.404	1.240	1.194	1.149

The β parameters will all be in the same proportion to those estimated for the Cape Cod method. However, we no longer have the freedom to introduce a tail factor to go from the 96-month age to ultimate. Instead, the data and our selected ELR have forced a tail factor upon us (again formula 3.2.10).

$$\text{LDF}_n = \frac{1}{\sum_{d=1}^{n} \beta_d} \tag{3.2.10}$$

AY y	Exposures v_y	ELR	Expected Ultimate	LDF	IBNR% 1-1/LDF	Latest Diagonal	Final Ultimate	Final L / R
1999	11,880	50.00%	5,940	1.149	12.93%	5,481	6,249	52.60%
2000	12,095	50.00%	6,048	1.194	16.25%	5,464	6,447	53.30%
2001	12,025	50.00%	6,013	1.240	19.33%	5,427	6,589	54.79%
2002	11,900	50.00%	5,950	1.404	28.75%	4,417	6,128	51.49%
2003	12,240	50.00%	6,120	1.741	42.57%	3,047	5,652	46.18%
2004	12,100	50.00%	6,050	2.547	60.73%	1,714	5,388	44.53%
2005	11,865	50.00%	5,933	5.293	81.11%	829	5,641	47.54%
2006	12,075	50.00%	6,038	23.539	95.75%	215	5,996	49.66%
Total	96,180	50.00%	48,090			26,594	48,090	50.00%

This format is the same as for the Cape Cod method, except that the ELR has been fixed by the model user. We may again note that the final ultimate loss ratio (relative to onlevel premium) is equal to the selected ELR.

In this BF example, the selection of the 50% ELR results in an implied tail factor of 1.149. We could have used the Cape Cod method instead, including a 1.149 tail factor, and produced the same results as the BF method. The two methods are algebraically equivalent: either the ELR determines the tail factor or the tail factor determines the ELR.

Method	Values Supplied by User	Estimated Parameters
Multiplicative LDF	LDF_n	$v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8$ $\beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7, \beta_8$ ELR
Cape Cod	$v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8$ LDF_n	$\beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7, \beta_8$ ELR
Bornhuetter-Ferguson as a GLM	$v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8$ ELR	$\beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7, \beta_8$ LDF_n

There may be objections at this point that we are not presenting the traditional BF method as found in the original 1972 paper. In that paper, the development pattern (including the tail factor) is selected prior to and independent of the ELR; the ELR implied by the data is ignored and implicitly overwritten by user.

When model parameters are overwritten by the user, bias is introduced: the fitted values for the triangle will no longer balance to the actual values. This bias may remain unrecognized because the model assumptions underlying the BF selections are never made explicit and are therefore left untested.

For this reason, we seek a method that keeps the stability of the traditional BF method, but is more responsive to the loss experience by balancing to the historical paid loss values.

3.3 A Unified Method

Having reviewed the three traditional methods used in a reserve review, we may note some limitations in each.

- The Multiplicative LDF method is clearly over-parameterized.
- The Cape Cod method is attractive but requires an exposure for every AY.
- The traditional Bornhuetter-Ferguson method involves user-intervention, making it less responsive to the actual loss experience.

Given these limitations, the most attractive option would be the Cape Cod method. Unfortunately, we may not have the full data to implement it. This is where a combination or unified method becomes most useful.

We begin by slightly modifying our original model to have the ELR apply to a subset of years. For example, the most recent four years may be grouped together under the same ELR, with the older years being estimated separately. We begin again with the general model.

$$E\left(c_{y,d}\right) \quad = \quad \mu_{y,d} \quad = \quad v_y \cdot \mathrm{ELR} \cdot \beta_d \tag{3.3.1}$$

The key concept is that the exposure values, v_y, are not available for the older years and so must be estimated in the model just as was done for the Multiplicative LDF method. We define a group of years, g, in which the exposures are available as containing the indices for the more recent years 2003-2006: $g = \{5, 6, 7, 8\}$.

If all of the years are part of the group, $g = \{1, 2, \cdots, n\}$, then the Unified method is equivalent to the Cape Cod method. On the other extreme, if only the most recent year is included in the group, $g = \{n\}$, then the Unified method is equivalent to the Multiplicative LDF method.

To solve for the Maximum Likelihood Estimates (MLE) of this model, we again have the condition that the fitted column totals must equal the actual column totals. We also have a condition that the sum of all the rows in the subset of years, *g*, must balance between fitted and actual values. This can be written using an indicator function, $\delta(y \in g)$, which is equal to unity for years in the group and zero otherwise.

$$\sum_{y=1}^{n}\sum_{d=1}^{n+1-y} c_{y,d} \cdot \delta(y \in g) \quad = \quad \sum_{y=1}^{n}\sum_{d=1}^{n+1-y} v_y \cdot \mathrm{ELR} \cdot \beta_d \cdot \delta(y \in g) \tag{3.3.2}$$

For our eight-year example, this implies:

$$\sum_{d=1}^{9-y} c_{y,d} \quad = \quad \sum_{d=1}^{9-y} \hat{v}_y \cdot \mathrm{ELR} \cdot \beta_d \qquad \text{for } y = \{1, 2, 3, 4\} \tag{3.3.3}$$

$$\sum_{y=5}^{8}\sum_{d=1}^{9-y} c_{y,d} \quad = \quad \sum_{y=5}^{8}\sum_{d=1}^{9-y} v_y \cdot \hat{\mathrm{ELR}} \cdot \beta_d$$

This method requires an iteration to solve for the maximum likelihood values, but it is not difficult. The iteration finds the values for the β_d and ELR parameters such that the column total and the grouped-row totals match the actual values.

The result is the "best" model in that it uses all of the available information, produces an unbiased fit, and satisfies the maximum likelihood criteria.

The concept of the Unified method may sound abstract at first, but a numerical example will show that the application is actually quite simple.[8]

We first assume that the rate adequacy index is only available for the second half of the experience period. The exposures for the earlier years are just placeholders and do not affect the

---

[8] For an alternative discussion of this approach to reducing the number of parameters, see Venter (2007). While he is reducing the number of development period parameters rather than the number of accident year parameters, the technique is the same.

final result.

AY y	Earned Premium	Onlevel Factor	Exposures v_y	
1999	5,400	na	5,400	Separate Years
2000	5,900	na	5,900	
2001	6,500	na	6,500	
2002	8,500	na	8,500	
2003	10,200	1.200	12,240	Grouped Years
2004	11,000	1.100	12,100	
2005	11,300	1.050	11,865	
2006	11,500	1.050	12,075	
Total	70,300			

The results of the Unified method can be displayed in the same format as was used for the other methods.[9] The difference in this final version is that the ELR is the same for the recent years but different for the earlier years. The expected loss for the earlier years is simply the result from the Multiplicative LDF method.

AY y	Exposures v_y	ELR	Expected Loss	LDF	IBNR% 1-1/LDF	Latest Diagonal	Final Ultimate	Final L / R
1999	5,400			1.000	0.00%	5,481	5,481	101.50%
2000	5,900			1.037	3.59%	5,464	5,668	96.06%
2001	6,500			1.074	6.90%	5,427	5,829	89.68%
2002	8,500			1.203	16.90%	4,417	5,315	62.53%
2003	12,240	33.14%	4,057	1.465	31.74%	3,047	4,335	35.41%
2004	12,100	33.14%	4,011	2.104	52.47%	1,714	3,818	31.56%
2005	11,865	33.14%	3,933	4.293	76.71%	829	3,846	32.41%
2006	12,075	33.14%	4,002	18.745	94.67%	215	4,004	33.16%
Total	74,580	na	na			26,594	38,296	51.35%
2003-2006	48,280	33.14%	16,002	2.757		5,805	16,002	33.14%

As can be seen in this example, the Unified method is the Multiplicative LDF applied to the old years and the Cape Cod applied to the more recent years. In order for this to be the maximum likelihood estimate, the development pattern and the ELR should be calculated simultaneously. This

[9] The results shown require a numerical iteration to find the MLE parameters, so the reader can verify that the numbers satisfy the balance for row and column totals but cannot easily re-derive the parameters. A practical compromise is given in section 4.1.

requirement can be relaxed in practice, by having the analyst separately select the loss development pattern (see Section 4.1).

This example also shows that the group of years for which the exposure base is available can be treated as a unit with an average LDF applied multiplicatively. The average LDF is calculated as a harmonic average using the exposures as weights.

$$2.757 \quad = \quad \frac{48{,}280}{\left(\dfrac{12{,}240}{1.465}+\dfrac{12{,}100}{2.104}+\dfrac{11{,}865}{4.293}+\dfrac{12{,}075}{18.745}\right)}$$

We may summarize the relationship of this Unified method to the Multiplicative LDF and Cape Cod cases in the following chart.

Method	**Values Supplied by User**	**Estimated Parameters**
Multiplicative LDF	LDF_n	$v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8$ $\beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7, \beta_8$ ELR
Cape Cod	$v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8$ LDF_n	$\beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7, \beta_8$ ELR
Unified	v_5, v_6, v_7, v_8 LDF_n	v_1, v_2, v_3, v_4 $\beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7, \beta_8$ ELR

4. PRACTICAL ISSUES FOR THE "UNIFIED" METHOD

Having outlined the general approach for applying a Unified method that combines Multiplicative LDF and Cape Cod methods, we now wish to address two practical issues.

4.1 Separating the Selection of the Development Pattern

As was noted in the description of the theory underlying the Unified method, it is necessary that the development pattern (viewed either as β s or LDFs) must be estimated simultaneously with the other parameters in order to have the maximum likelihood estimate for the reserves. This may be unrealistic in practice, because the reserving actuary will often choose to smooth out the development pattern by removing outlier points or giving more weight to more recent diagonals.

All of these methods allow this step to be done separately. What results is a model that is further constrained by the selection of the β parameters. For example, the Multiplicative LDF method now seeks to find the "best" (MLE) α parameters, representing ultimate losses by accident year, given a selected development pattern. Within the ODP model, the maximum likelihood estimate is found by applying the selected LDF to the latest diagonal of the cumulative loss triangle. Likewise, for the Unified method, we simply apply the same method as outlined in section 3.3, using the selected LDFs.

Having selected a loss pattern of β parameters, either from the triangle or from external information, we apply this to the latest diagonal of account data: multiplicatively for the old years and additively (via Cape Cod) for the more recent years. This is equivalent to a GLM with the β parameters constrained by the user and the ELR fit via MLE.

Method	**Values Supplied by User**	**Estimated Parameters**
Multiplicative LDF	LDF_n	$v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8$ $\beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7, \beta_8$ ELR
Cape Cod	$v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8$ LDF_n	$\beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7, \beta_8$ ELR
Unified	v_5, v_6, v_7, v_8 LDF_n	v_1, v_2, v_3, v_4 $\beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7, \beta_8$ *ELR*
Unified method with development pattern selected by user.	v_5, v_6, v_7, v_8 $\beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7, \beta_8$ LDF_n	v_1, v_2, v_3, v_4 ELR

Because of this result, we can also interpret the Unified method as a purely multiplicative approach. In our example, we have grouped the latest four years together to apply the Cape Cod method. The ultimate for that group of years can also be calculated by applying a single average LDF to the four-year block.

$$\mathrm{LDF}_{\mathrm{Group}} = \frac{v_{2003} + v_{2004} + v_{2005} + v_{2006}}{\dfrac{v_{2003}}{\mathrm{LDF}_{48}} + \dfrac{v_{2004}}{\mathrm{LDF}_{36}} + \dfrac{v_{2005}}{\mathrm{LDF}_{24}} + \dfrac{v_{2006}}{\mathrm{LDF}_{12}}} \tag{4.1.1}$$

In other words, we take a weighted harmonic average of the development patterns for each year in the block, using the exposures as the weights. This average factor then is applied to the four-year block itself. The IBNR can be allocated back down to the individual years using the same Cape Cod method.

The averaging approach accomplishes the same stabilizing goal that is the reason that many people now use the Bornhuetter-Ferguson method, but it better responds to the actual experience.

We should also note that this concept is not original with this paper. This averaging method is the same as would be used if you had a development pattern from accident quarters (AQ) and needed to estimate an accident year (AY) development factor. You would perform an average as below.

$$\mathrm{LDF}_{AY\_12} = \frac{4}{\dfrac{1}{\mathrm{LDF}_{AQ\_3}} + \dfrac{1}{\mathrm{LDF}_{AQ\_6}} + \dfrac{1}{\mathrm{LDF}_{AQ\_9}} + \dfrac{1}{\mathrm{LDF}_{AQ\_12}}} \tag{4.1.2}$$

This is the same as our Unified group, with the assumption that exposures are uniform across quarters.

4.2 Creating the Exposure Index

A second practical problem is the need to create an appropriate onlevel factor. As stated previously, the resulting exposure measure v_y should be proportional to the expected loss for accident year "*y*."

The starting point for this calculation should be changes in the underlying pricing, including the key components:

- Changes to base rates and increased limits factors
- Changes to discretionary pricing modifications

- Changes to terms and conditions (e.g., removal of exclusions)
- Enforcement of underwriting standards (e.g., correct classifications, audits)

We want to adjust for these components so as to remove the effects of the "insurance cycle."

The second component for the onlevel factor is an adjustment for inflation trend. That is, we want to have each year's premium adjusted to a common rate level, but at the loss cost level for the specific year. We do this by adjusting the premium to a projected future level, reflecting rate changes and increases due to exposure inflation. That adjusted premium is then de-trended based on loss inflation.

AY y	Earned Premium	Rate Onlevel Factor	Exposure Trend at 3.0%	Onlevel Premium	Loss Trend at 6.0%	Exposures v_y	Final Onlevel Factor
A	B	C	D	E=B*C*D	F	G=E/F	H=G/B
1999	5,400	2.690	1.230	17,863	1.504	11,880	2.200
2000	5,900	2.435	1.194	17,157	1.419	12,095	2.050
2001	6,500	2.136	1.159	16,092	1.338	12,025	1.850
2002	8,500	1.570	1.126	15,023	1.262	11,900	1.400
2003	10,200	1.308	1.093	14,578	1.191	12,240	1.200
2004	11,000	1.165	1.061	13,596	1.124	12,100	1.100
2005	11,300	1.081	1.030	12,577	1.060	11,865	1.050
2006	11,500	1.050	1.000	12,075	1.000	12,075	1.050

Because this index involves estimates of inflation, as well as components of price adequacy that may be difficult to quantify, it is not an easy task to estimate it reliably for a long historical period. This is a practical argument for the Unified method to be applied rather than Cape Cod method.

5. CONCLUSIONS

This paper has outlined a "Unified" reserving method that is a combination of familiar Multiplicative LDF and Cape Cod methods. This Unified method allows the reserving actuary to make use of exposure information even if it can only be compiled for a few recent periods. This Unified method is based on the same statistical model that is common to both of these other methods.

This Unified method achieves the goal of stabilizing the reserves for immature periods, while also being more responsive to the actual loss payments than the traditional Bornhuetter-Ferguson method.

Acknowledgment

The author gratefully acknowledges the helpful reviews by Denise Ambrogio, Bobby Hancock, Gerhard Quarg, Simone Rechtsteiner, Ira Robbin, Jim Sandor, and Chuck Thayer. They offered valuable and constructive criticism even where they did not fully agree with all of the conclusions of this paper.

5. REFERENCES

[1] Bornhuetter, Ronald L., and R.E. Ferguson, "The Actuary and IBNR," *PCAS* LIX, 1972, 181-195.
[2] Clark, David R., "LDF Curve-Fitting and Stochastic Reserving: A Maximum Likelihood Approach," *CAS Forum*, Fall 2003, 41-91.
[3] Mack, Thomas and Gary Venter, "A Comparison of Models that Reproduce Chain Ladder Reserve Estimates," *ASTIN Colloquium International Actuarial Association - Brussels, Belgium* 1999: Toyko.
[4] Mack, Thomas, "Parameter Estimation for Bornhuetter/Ferguson," *CAS Forum*, Fall 2006, 141-157.
[5] Mildenhall, Stephen J., "A Systematic Relationship Between Minimum Bias and Generalized Linear Models," *PCAS* LXXXVI, 1999, 393-487.
[6] Renshaw, Arthur E., and Richard Verrall, "A Stochastic Model Underlying the Chain-Ladder Technique," *British Actuarial Journal* 4, 1998, 903-923.
[7] Schmidt, Klaus D., "Methods and Models of Loss Reserving Based on Run-Off Triangles: A Unifying Survey," *CAS Forum*, Fall 2006, 269-317.
[8] Shapland, Mark, "Loss Reserve Estimates: A Statistical Approach for Determining 'Reasonableness'," *Variance*, Vol. 1, No. 1, Spring 2007, 120-148.
[9] Taylor, Greg, "Chain Ladder Bias," *ASTIN Bulletin*, Vol. 33, 2003, 313-330.
[10] Venter, Gary, "Refining Reserve Runoff Ranges," *CAS E-Forum*, Summer 2007, http://www.casact.org/pubs/forum/07sforum/07s-venter2.pdf.
[11] Verrall, R.J. and P.D. England, "Comments on 'A comparison of stochastic models that reproduce chain ladder reserve estimates'," *Insurance: Mathematics and Economics* 26, 2000, 109-111.
[12] CAS Research Working Party on Quantifying Variability in Reserve Estimates, "The Analysis and Estimation of Loss & ALAE Variability: A Summary Report." *CAS Forum*, Fall 2005, 29-146.

Abbreviations:

BF	Bornhuetter-Ferguson
CL	Chain ladder
GLM	Generalized Linear Model
IBNR	Incurred But Not Reported loss (all loss beyond the amounts in the historical triangle)
LDF	Loss Development Factor, also known as an "age-to-ultimate" factor
ODP	Over-Dispersed Poisson

Reserving Model Notation:

$c_{y,d}$	Actual incremental losses in accident year "y" and development period "d"
β_d	Parameter for development period "d"; can be thought of as the percent of ultimate loss paid during a given development period
ELR	Expected Loss Ratio
v_y	Exposures or volume measure for accident year "y"; can be thought of as onlevel premium
ϕ	Dispersion parameter for ODP model, ϕ = ratio of variance to mean

Biography of the Author

Dave Clark is part of the Actuarial Research and Modeling team for Munich Reinsurance America, Inc. His paper "LDF Curve-Fitting and Stochastic Reserving: A Maximum Likelihood Approach" received the 2003 Reserves Call Paper Prize.