

WHEN CAN ACCIDENT YEARS BE REGARDED AS DEVELOPMENT YEARS?

GLEN BARNETT, BEN ZEHNWIRTH AND EUGENE DUBOSSARSKY

Abstract

The chain ladder (volume-weighted average development factor) is perhaps the most widely used of the link ratio (age-to-age development factor) techniques, being popular among actuaries in many countries. The chain ladder technique has a number of interesting properties. We present one such property, which indicates that the chain ladder doesn't distinguish between accident years and development years. While we have not seen a proof of this property in English language journals, it appears in Dannenberg, Kaas and Usman [2]. The result is also discussed in Kaas et al [3]. We give a simple proof that the chain ladder possesses this property and discuss its other implications for the chain ladder technique. It becomes clear that the chain ladder does not capture the structure of real triangles.

1. INTRODUCTION

Link ratio (loss development factor) methods are widely used for reserving. The chain ladder technique is one such method applied to cumulatives (or sometimes incurreds). The development factor is a weighted average of the individual link ratios, weighted by the previous cumulative (volume-weighted average). The chain ladder is normally applied to cumulated paid loss arrays, incurred loss arrays, or sometimes to cumulated claim numbers, such as claims incurred, claims notified or claims closed. This 'formal' chain ladder is described by Mack in [5], but we give a detailed description of it below. We present the chain ladder for a paid loss array with annual data; the exposition is essentially the same for other kinds of data. We assume that the reader is familiar with the usual triangular development layout.




Figure 1 Incremental paid loss array

Consider the incremental array $P = [p_{ij}]$, $i = 1, \dots, s$; $j = 0, \dots, s-1$; $i+j \leq s$ (the array of incremental payments – the actual amounts paid in each development year in respect of each accident year – contains the fundamental observed quantities).

The chain ladder is usually presented in something like the following fashion. Let us take an array of paid losses (incremental amounts paid), p_{ij} , and cumulate along the accident years, $c_{ij} = p_{i0} + p_{i1} + \dots + p_{ij}$, so that $c_{ij} = \sum_{k=0}^j p_{ik}$ are the corresponding cumulatives. Then compute ratios $r_j = \sum_i c_{ij} / \sum_i c_{i,j-1}$, where the sum is over all available terms that are present in both the j^{th} and $(j-1)^{\text{th}}$ columns. Forecasts are produced by projecting elements on the last diagonal $c_{i,s-i}$ to the next development by multiplying by the development ratio r_{s-i+1} , and recursively projecting those forecasts in turn by multiplying by the next ratio.

Now imagine a version of the chain ladder working the other direction (‘down’ rather than ‘across’) – where you cumulate downward, take ratios running down (accident-year-to-accident-year ratios), project down into the future, and difference back to incrementals, as in the diagram below. It turns out that the incremental forecasts for both the usual chain ladder (the version that runs across) and this new ‘down’ version of the chain ladder are the same.




Figure 2 Two incremental arrays to which an “across” (the standard chain ladder) and a “down” version of the chain ladder are applied.

2. THE INCREMENTAL CHAIN LADDER

To see this, we will first write the chain ladder purely in terms of incrementals (which we call *the incremental chain ladder*).

Consider that we are attempting to forecast a cumulative, c_{ij} , in the next calendar year. Let $A_{ij} = \sum_k c_{kj}$, that is, A_{ij} is the sum of all the cumulatives in the column above c_{ij} . Then the j^{th} ratio, $r_j = A_{ij}/A_{i,j-1}$. Note that A_{ij} is also the sum of all the incrementals above the (i,j) cell (p_{kj} , $k=1, 2, \dots, i-1$) plus all the incrementals to the left of those. That is, A_{ij} and $A_{i,j-1}$ are the sums of all the incrementals in the regions marked below. If the values of the incrementals are represented by heights of square prisms in each (i,j) cell, values represented by A , B and C may be thought of as the ‘total volume’ in the marked regions.

Note further that (since the forecasts of the cumulatives are in the same ratio), the formula $r_j = A_{ij}/A_{i,j-1}$ as a ratio of sums of incrementals as defined above applies to observations in later (further into the future) calendar years as well, as long as any unobserved incrementals in the sum are replaced with their predicted values.

In the usual form of the chain ladder, you compute forecasts $\hat{c}_{ij} = r_j c_{i,j-1}$ (where $c_{i,j-1}$ is again replaced by its forecast when it is unavailable). That is, compute the forecasts $\hat{c}_{ij} = A_{ij}/A_{i,j-1} c_{i,j-1}$. Predicted incrementals may be formed by differencing predicted cumulatives. Computation of incremental forecasts is essential for incorporating future inflation and discounting, (where relevant) and for computation of annual claim cash flows.

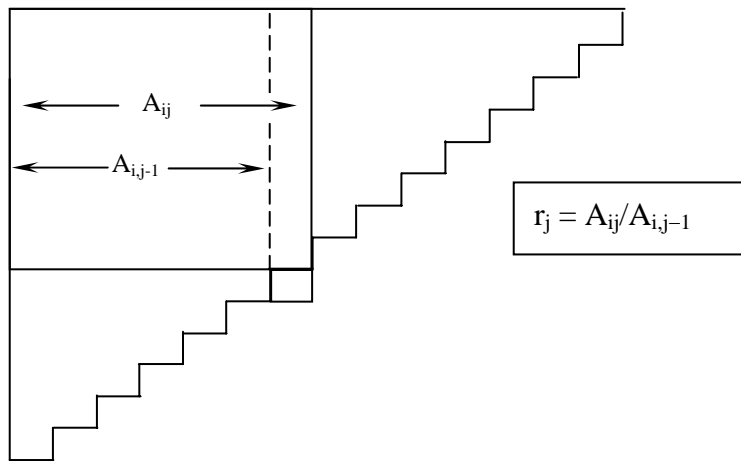


Figure 3 Depiction of the incrementals involved in the calculation of the ratio.

Now let $b_{ij} = A_{ij} - A_{i,j-1}$, which is the sum of the incrementals above p_{ij} . For simplicity, in the diagram below this is just called B. Similarly, let $C = c_{i,j-1}$, and let $A = A_{i,j-1}$.

Then the forecast may be written

$$\begin{aligned} \hat{c}_{ij} &= A_{ij}/A_{i,j-1} c_{i,j-1} \\ &= (A_{i,j-1} + b_{ij})/A_{i,j-1} c_{i,j-1} \\ &= [(A+B)/A] C \\ &= [1 + B/A] C. \end{aligned}$$

Similarly, the incremental forecast is

$$\begin{aligned} \hat{p}_{ij} &= \hat{c}_{ij} - c_{i,j-1} \\ &= [1 + B/A] C - c_{i,j-1} \\ &= [1 + B/A] C - C \\ &= B.C/A \end{aligned}$$

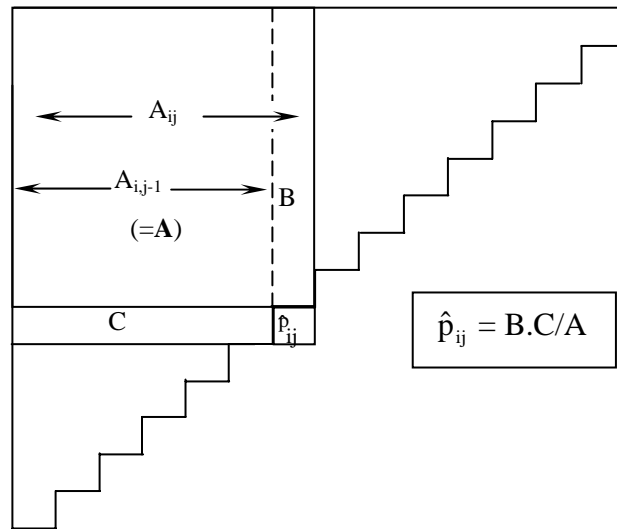


Figure 4 Calculation of cumulative and incremental forecasts in terms of incrementals. Labels represent the sums of the incrementals in their region.

That is, the forecast of the incremental observation is the product of (the sum of the incrementals above it) by (the sum of the incrementals to its left) divided by (the sum of all the incrementals that are both above and to the left). Note that this is symmetric in B and C (and also A) – interchanging i and j merely changes the role of B and C. Thus we see that the chain ladder may be neatly defined directly in terms of the incrementals. See the appendix for a brief but somewhat more formal proof of the above symmetry.

It may help to give an example. Imagine we have an incremental paid array as follows, and we wish to predict the incremental cell labeled P:

12	26	19	⇒	12	38	57	⇒	12	38	57	P = 63-42 = 21
18	24	P		18	42	C		18	42	63	
15				15				15			

cumulate $r_{2:1} = 57/38$ $C = 57/38 \times 42$

Note what the cumulative forecast, C, consists of in terms of the incrementals:
 $C = 57 \times 42 / 38 = (12+26+19) \times (18+24) / (12+26)$.

Hence, we have

$$\begin{aligned}
 P &= 57 \times 42 / 38 - 42 = (12+26+19) \times (18+24) / (12+26) - (18+24) \\
 &= (12+26+19) \times (18+24) / (12+26) - (18+24) \times (12+26) / (12+26) \\
 &= (12+26+19 - \{12+26\}) \times (18+24) / (12+26) \\
 &= 19 \times (18+24) / (12+26) = B.C/A .
 \end{aligned}$$

Every incremental forecast turns out to work the same way (recall that you must replace unobserved values in A, B and C by their forecasts in this formulation). Consequently,

look at what happens when we interchange accident and development years in (i.e. transpose) the original array and apply the chain ladder:

12	18	15	⇒	12	30	45	⇒	12	30	45	P = 40-19 = 21	
26	24			26	50			$r_{2:1} = 80/38$	26	50		
19	P			cumulate	19	C			$C = 80/38 \times 19$	19		40

Considered in terms of cumulatives, it is not immediately clear that the chain ladder prediction of P will always be the same. However, if you consider it directly in the incremental form, while B and C have interchanged, their product is obviously the same, and A is unchanged, so the prediction P is unchanged. In each case, we have $P = 19 \times (18+24) / (12+26)$.

This applies generally – interchanging accident and development years (transposing) in the incremental array and applying the chain ladder, when considered in terms of the incremental formula, merely interchanges the values of B and C, and leaves A unchanged, so the incremental forecasts are unchanged.

One advantage of this incremental version of the chain ladder is that it is substantially more convenient to implement in a spreadsheet. This is because it can be implemented in terms of formulas that can be successfully cut and pasted without the effort involved in computing the ratios first. The usual ratios (and cumulative forecasts if needed) are then easily computed from the completed array.

3. BRIEF DISCUSSION OF SOME RELATED WORK

Kremer [4] recognizes the connection between a ratio model (which he calls a multiplicative model) and two-way ANOVA within missing values on the logarithms. He uses this to derive an approach to forecasting outstanding claims (he derives the expectations of the final outstanding). Kremer explicitly points out the strong connection to the chain ladder method.

Mack [5] derives standard error calculations (including process and parameter error) for a mean-variance model whose forecasts reproduce the standard chain ladder technique in a recursive fashion. Mack makes use of the Gauss-Markov theorem to avoid specific distributional assumptions for the losses. He compares results for a particular case study with results from similar models for which computation of exact or approximate standard errors are available.

In his later paper, Mack [6] argues that while several different stochastic ratio models had previously been referred to as the stochastic chain ladder, that the model he discusses in [5] reproduces the classical chain ladder forecasts and that models that don't shouldn't be referred to as chain ladder models.

Murphy [7] explicitly writes several loss development factor methods as stochastic models and derives forecast variances, working in a least-squares framework. He argues that it is often necessary to extend ratio models to include intercepts.

Barnett and Zehnwirth [1] develop a statistical framework extending Murphy's approach to include some adjustment for common accident and calendar period trends as a general diagnostic tool for testing the suitability of ratio models to data. Multiple examples point toward some common deficiencies of ratio models, including the need for an intercept, and the lack of predictive power of ratios after incorporating obvious predictors.

Renshaw and Verall [8] derive another model that reproduces the chain ladder. Their formulation makes the number of parameters describing the mean process in the chain ladder explicit. The model is initially presented as a Poisson model, which extends to a quasi-likelihood framework as a model with variance proportional to the mean.

Even though we started with a form of the chain ladder that looked something like the stochastic form presented by Mack [5, 6] and Murphy [7], by the end of section 2 there are strong similarities to the stochastic form presented by Renshaw and Verrall [8]. Despite arguments in the literature, the two approaches differ mainly in the data they appear to condition on when describing the past, and in how many variance parameters they employ. They are identical in the way they describe the mean predictions for the future, which is why they both reproduce the chain ladder forecasts. Given a quasi-likelihood approach, differences in forecast standard errors appear to be largely due to two factors – the number of variance parameters, and the number of degrees of freedom to fit the data (i.e. parameters) for which the parameter uncertainty is ignored.

Note that Kremer [4] describes the relationship between the ratios of the chain ladder and the column parameters in the two-way loglinear model. This is akin to the relationship between the parameters in the Mack formulation and those of Renshaw and Verrall.

4. INTERCHANGEABILITY OF ACCIDENT AND DEVELOPMENT YEARS

We immediately see from the abovementioned symmetry that the incremental predictions from the chain ladder of the incremental array with accident years and development years interchanged (with the array transposed) are simply the corresponding predictions from the original array, with accident years and development years interchanged (transposed). That is, the chain ladder has the property that its incremental forecasts are the *same* whether the chain ladder is applied to an incremental array running across (as is usual) or down – where you cumulate down, take accident-year-to-accident-year ratios, project down into the future, and difference back to incrementals!

Note that this property must hold for the forecasts of all models that reproduce the chain ladder forecasts. Such a property might most accurately be called the “transpose-forecast commutativity property of the chain ladder”. However, in the interest of brevity we simply call it *transpose-invariance*.

This property implies that any fact that applies to the accident years applies to development years, and vice versa, and that any asymmetry of directions in our description of the chain ladder is an artifact of our description, and is not an inherent part of the chain ladder itself. That is, the chain ladder doesn't differentiate between accident and development periods - it treats them in identical fashion, in spite of the actual structure in the two directions being completely different. This result obviously applies to forecasts for all the stochastic chain-ladder-reproducing models as well.

If accident years are treated like they were development years, imagine you have homogeneous accident years (a not uncommon occurrence, especially after you adjust for changes in exposure and inflation, assuming no superimposed inflation). You wouldn't predict the level of the next accident year using ratios – it would be far more sensible and informative to take some kind of average. But the chain ladder in fact *does* use ratios.

If this way of looking at the chain ladder seems a little nonsensical, it is because we are inferring additional meaning in the usual form of the chain ladder that it doesn't really possess. The two descriptions (the across version and the down version) are in reality the same description of the data.

Note also that we can now see that there are in fact parameters in both directions in the chain ladder. This is not a consequence of any particular formulation of the chain ladder - every chain-ladder-reproducing model has *degrees-of-freedom to fit the data* (i.e. parameters) that run both across and down. Some formulations make the existence of both kinds of parameters explicit (as in Renshaw and Verrall [8]); some other formulations do not (such as Mack [6]) – the row parameters become hidden by the fact that the model is conditioned on the first column. The chain ladder itself still unavoidably has degrees-of-freedom to fit changes in accident level, so the parameters remain, even where not explicitly represented in the formulation. All formulations of the chain ladder have $2s-1$ parameters for the mean, though the number of variance parameters and distributional assumptions may vary.

We note that so many parameters makes the forecasts quite sensitive to relatively small changes in a few values, making the chain ladder unsuitable for forecasting. Yet even with so many parameters the chain ladder is still unable to model changing superimposed inflation.

A further important consequence of this property is that parameters in the two directions can take the roles of both a level and a ratio.

We know within ourselves that the two directions are fundamentally different, both in general appearance of their trends and in spirit. The development year direction tends to have a smooth "run-off" shape, where the incrementals tend to increase initially to a peak somewhere in the first few developments and then smoothly decrease in the tail, while the accident years tend to have quite a different pattern. Yet the model itself makes no such

distinction – it does not contain important information we already know about claims payments (i.e. the structure in loss data).

Indeed, the chain ladder model, being a two-way cross classification model (as has been recognized by numerous authors), not only fails to distinguish between accident years and development years, it ignores the relationships between years within either category.

Consider the two plots of logs of paid data against development year in figure 5. Can you tell which one is the real data?




Figure 5 Plot of exposure adjusted log paid incremental data against development year. One of the plots has had its development period labels randomly allocated.

Most practitioners will instantly (and correctly) guess that the lower plot is the real one. We *know* that paid data often has a strong pattern to its runoff – that nearby development

periods tend to be more alike than ones further away, and further, we usually observe smooth trends relating them.

Clearly the accident and development labels mean something, and you can't arbitrarily relabel them without affecting the information in the data. The observations in development year 3 do not just tend to be closer to each other than to observations from other columns, they also tend to be more like the observations in development years 2 and 4 than they are like observations from columns further away. In a two-way cross-classification model, we can arbitrarily rearrange the group labels within both factors, and even interchange the factors *without changing the fit*. If the labels do carry information over and above being arbitrary identifiers of a category the observation falls in (as they do in claims runoff), the chain ladder model is inappropriate.

With the lower plot of figure 5, one could omit *all* of the data for any development period between 1 and 8 (i.e. replace the observations with missing values) and still be able to get a good estimate of the values in that development period. Nearby periods carry a great deal of information about each development period's level. If we consider only the first plot, if we omitted a development period, what would we know about it? The information was there, but we threw it away when we threw away the ordering in the development year. There is also information in the accident year direction – nearby accident years often tend to be more alike than ones further away. The chain ladder ignores that information in both directions. This loss of information causes the predictive distributions of chain ladder forecasts to be very wide, much wider than they should be if the model used what we know about losses.

The top plot of figure 5 actually looks a bit more like a plot against accident years (though nearby accident years in practice are often closer together than those further away, and so they tend to be smoother than the top plot, even though they don't normally exhibit the smooth curves of the development direction).

That is not to say that a plot against development years looking something like the top one could *never* arise, but it is quite rare – and if it *does* arise, an ANOVA-style model is not all that helpful in forming good forecasts, particularly in the tail and for future developments. It has parameters where it has little data, and that makes for poor forecast prediction errors. An underparameterized model is often substantially better for forecasting in that circumstance. If we were in the rare circumstance that the means for each development didn't have any strong trend to them, we'd want to quantify the extent to which the means tend to shift around their overall average, and use as much information as possible in identifying what little trend there might be. When there is less information in the data, it is even more crucial not to waste it.

Many of the problems discussed in respect of the chain ladder apply to other link ratio methods. The exact transpose-invariance property no longer applies (since different weights are involved), but basic link ratio methods are still two-way cross classification models (with different assumptions about variance), so they generally share the problem of overparameterization in the development and accident year direction, and ignore the

relationships between adjacent year levels. Further, although the correspondence isn't exact, there is generally a strong similarity between forecasts (on the incremental scale) and the transposed-forecasted-transposed forecasts. This is hardly surprising, since other ratios may be written as weighted versions of the chain-ladder; the transposing merely results in a differently weighted version of a thing that *is* transpose-invariant.

REFERENCES

- [1] Barnett, Glen, and Zehnwirth, Ben (2000) "Best Estimates for Reserves," *PCAS*, LXXXVII No. 166-167, pp. 245-303.
- [2] Dannenburg, D.R., Kaas, R, and Usman, L.N. (1998) "Gegeneraliseerde Lineaire Modellen voor IBNR-driehoeken," *Het Verzekerings Archief*, Jaargang 75, 4e kwartaal 149-158
- [3] Kaas, R., Goovaerts, M., Dhaene, J., and Denuit, M. (2001) *Modern Actuarial Risk Theory*, Kluwer Academic Publishers, Dordrecht, p208-209.
- [4] Kremer, E. (1982). IBNR claims and the two way model of ANOVA. *Scandinavian Actuarial Journal*, 47-55.
- [5] Mack, Thomas (1993), "Distribution-Free Calculation of the Standard Error of Chain Ladder Reserve Estimates," *ASTIN Bulletin*, Vol 23 No. 2, pp. 213-225.
- [6] Mack, Thomas (1994), "Which stochastic model is underlying the chain ladder method?" *Insurance Mathematics and Economics*, Vol 15 No. 2/3, pp. 133-138.
- [7] Murphy, Daniel M. (1994), "Unbiased Loss Development Factors," *PCAS*, LXXXI No. 154-155, pp. 154-222.
- [8] Renshaw, A. and Verrall, R. (1998), "A Stochastic Model underlying the Chain-Ladder Technique," *British Actuarial Journal*, Volume 4, Part IV, pp. 903-923.

APPENDIX

Derivation of the incremental chain ladder

In the following, when a term appears on the right hand side of an equation, if it is unavailable, it is replaced by its predicted value. The usual form of the chain ladder predictions is given by:

$$\begin{aligned}
 \hat{c}_{ij} &= \hat{\beta}_j c_{i,j-1} \\
 &= [\sum_{h=1}^{s-j} c_{hj}] / [\sum_{h=1}^{s-j} c_{h,j-1}] \cdot c_{i,j-1} \\
 &= (1 + c_{s-j,j} / [\sum_{h=1}^{s-j} c_{h,j-1}]) \cdot c_{i,j-1} \\
 &= (1 + [\sum_{h=1}^{s-j} p_{hj}] / [\sum_{h=1}^{s-j} \sum_{k=0}^{j-1} p_{hk}]) \cdot \sum_{k=0}^{j-1} p_{ik}
 \end{aligned}$$

$$\begin{aligned}
 \text{Hence } \hat{p}_{ij} &= (1 + [\sum_{h=1}^{s-j} p_{hj}] / [\sum_{h=1}^{s-j} \sum_{k=0}^{j-1} p_{hk}]) \cdot \sum_{k=0}^{j-1} p_{ik} - \sum_{k=0}^{j-1} p_{ik} \\
 &= ([\sum_{h=1}^{s-j} p_{hj}] / [\sum_{h=1}^{s-j} \sum_{k=0}^{j-1} p_{hk}]) \cdot \sum_{k=0}^{j-1} p_{ik} \\
 &= (B/A) \cdot C, \quad (\text{see figure 4}) \\
 &= B \cdot C / A
 \end{aligned}$$

where $A = \sum_{h=1}^{s-j} \sum_{k=0}^{j-1} p_{hk}$, $B = \sum_{h=1}^{s-j} p_{hj}$ and $C = \sum_{k=0}^{j-1} p_{ik}$. Note the symmetry in the subscripts.

The transpose invariance property

The symmetry immediately establishes the transpose invariance property. Equivalently, refer to figure 4, and note that the numerator of the equation for \hat{p}_{ij} is the product of the total of the values above it and the total of the values to its left. Consequently, if the array were transposed (rows and columns interchanged), the numerator for \hat{p}_{ji} would be unchanged (and of course the denominator is also unchanged) from that for \hat{p}_{ij} .