

Reserving on a contract-by-contract basis

Joe Monk, Peter Brinck, Kathryn Hudson,
Katherine Peet, Leith Raw

Introduction

This paper has the basic aim of applying familiar projection methods to a slightly wider set of data in order to add more value from the reserving process. It also neatly avoids some of the complications involved in providing regulators with the information they now require.

There are two simple ideas in this paper.

The first is to use development patterns measured from the inception date of a policy and not from the beginning of the calendar year.

The second is to use projection methods for individual contracts to derive estimated ultimate claims on a contract-by-contract basis. This allows the reserves and ULR projections to be split any which way that is desirable.

For example, with estimated ultimate claims on a contract-by-contract basis you could show the estimated ULRs by exposure type, by broker or by geographic area.

You could even analyse the data in homogeneous reserving groups and then re-allocate into other slightly arbitrary reserving groups, geographic areas and individual currencies that would not be possible to analyse on their own due to a lack of data.

There are two rules required for reading this paper.

First, you cannot criticise the methods here for failing to address something that the standard methods do not address.

Secondly, you must acknowledge that this is a work in progress. The main aim of writing this paper is to air the approach, generate discussion and gather feedback so the methodology can be refined.

Inception date development

The majority of insurance contracts are for one year. Under the traditional methods the development factor model starts at the beginning of the calendar year. Yet for a contract which starts in the last half of the year it is not possible to have claims for the first half of the year.

While in a steady state with uniform earning patterns this should not make any difference. The reality tends to be different. Earning patterns do move from year to year and this in itself creates uncertainty for fitting a pattern to the underlying development factor triangle.

By moving to an inception date basis, for losses occurring during contracts all business is earned during the first year of the development pattern.

This not only accelerates the development pattern relative to traditional methods but also creates a more uniform pattern as the noise from shifting earning patterns over time is removed.

This in turn reduces the uncertainty in the tail. With uncertainty most actuaries understandably add additional loadings.

Thus, the ultimate claims position becomes clearer sooner. For an individual contract this means that the ultimate claims position is clearer on average six months sooner than under the traditional methods.

We note that the same principles outlined above also apply to risks attaching business with a 24 month earning period.

At this stage, all well and good, but bonus marks to those who have read this and worked out a rather large flaw.

By moving everything to inception date cohorts the last portion of the triangle partly disappears. This is an important point, so forgive me if the following labours this too much.

Let us take a cohort of an underwriting year with quarterly development.

For the 2010 underwriting year, the first quarter's incurred snapshot is taken as at 31st March 2010, the second quarter is as at 30th June 2010, the third quarter is as at 30th September and, well, you get the idea.

At the 2011 year-end the incurred snapshot is taken as at 31st December 2011, which is the 8th development quarter in the triangle,

However, for the inception date basis it does not work that way. Let us take two contracts which incept on 1st January 2010 and 1st July 2010. We shall assume that both of these contracts have claims.

For the contract incepting on 1st January the triangle works in the same way as the traditional method.

However, for the contract starting on 1st July 2010 the first quarter's incurred snapshot is taken as at 30th September 2010, the second as at 31st December 2010, the third as at 31st March 2011 and so on.

At the 2011 year-end the incurred snapshot is taken as at 31st December 2011, which is the 6th development quarter in the triangle.

This means that for the 2010 inception year cohort one of these contracts has claim development up to the 8th quarter, whereas the other has claim development only up to the 6th quarter.

The incurred claims drop off in the last four quarters of the triangle.

This is potentially quite a problem, especially as this is the most recent data on which to make the projections.

After discussions amongst the group, we considered that there were four main options.

First, ignore the four most recent quarters of the triangle and go ahead with the projections anyway.

Secondly, fill in the missing data using standard development factors based on the data. In this example, we have the development factors from the 6th to 7th period for contract 1, which can simply be applied to the incurred claims from contract 2. Clearly, there is more data in the previous portions of the triangle to make this work.

Thirdly, we could apply an exposure base to put the triangle on a consistent basis for each period. This means that the claims would be weighted by the premium from the contract that they arise from. This means that when the premiums drop off for the 7th and 8th quarters, the incurred claims are grossed-up to the full premium value.

Fourth, abandon the idea as it is getting too complicated.

We decided to keep things simple and go for the first option.

This means that we simply ignore the last four data points of the triangle and use the remaining data to estimate the full development.

This conclusion was reached by a process of elimination. In reverse order:

Option four seemed like a waste of everyone's time and we really did think we had something here.

Option three had some intuitive appeal, but does assume a strong connection between premiums and claims. In practice, this didn't appear to be as strong as we had hoped. By breaking up the cohorts we would also be introducing additional volatility with smaller claim samples.

Option two would probably work. However, the filling in of the missing data is based on the development factors of the data we already have so it does not really add that much. We cannot create something out of nothing.

In a triangle with a reasonable number of cohorts and stable data, filling in the missing data is actually not that important for the overall projections. If a development factor is based on six cohorts then adding a seventh, derived from the previous six adds relatively little.

This led us to the pleasing conclusion that simply ignoring the past year's worth of data in the triangle was the most appropriate and easiest choice.

As you will no doubt have realised, this has the fundamental problem that we are throwing away data, which really is not part of the actuarial training. Not only that but it's the most recent data, so are we really advocating doing projections a year in arrears, we may end up missing something rather important.

The answer to this is two-fold.

First, for the claim triangle we are removing the development factors for the most recent periods. ie the leading diagonal and three diagonals before. The question here is how much influence this really has on the development pattern chosen.

Under the traditional methods a purist view would be that the patterns need to be re-projected each quarter based on the latest data. Given the amount of judgement required in deriving patterns, I am not sure I agree with that approach.

Re-projecting once a year is more than adequate and, even then, the patterns should be reasonably consistent from year-to-year.

Therefore, under the inception date method not using the data on the leading diagonal is not so bad. It would be concerning if those data points really were driving the results.

Of course, for new business with only a few cohorts then excluding the leading diagonal would cause problems. However, difficulties exist due to a lack of claims data, regardless of the projection method chosen.

Secondly, in reality we are not throwing away this data as it is still used further on in the method, it is simply not used for deriving the development factor model. This is a small price to pay for the benefits later on.

In practice, we also look at the development on a normal underwriting year basis, so any unusual jumps in the most recent quarters' data would be picked up.

So now we have an inception date triangle that we can project using normal development factor methods to derive a pattern.

After using this method for several lines of business (both Property and Casualty) we found that the acceleration in the development was slightly more than six months.

Contract-by-contract projection

Now that we have a development pattern based on the inception date of a contract, we can apply this pattern to each contract.

This can be done using paid or incurred development factor modelling. We can incorporate Bornhuetter-Ferguson methods as well, potentially using the contract pricing initial estimated ULR as the apriori.

However, there is a further flaw in making contract-by-contract projections that is not apparent when projecting a whole cohort of aggregated contracts.

The incurred development pattern is applied to both open and closed claims, so when applying patterns to individual contracts then a closed claim would be projected.

The solution to this is relatively straightforward although does require some tinkering. We shall present a rather naive view without the tinkering.

For a given cohort of claims, ignoring pure IBNR, it is the open claims that develop, while the closed claims do not. Therefore, for a development factor of, say, 1.2 for all incurred claims with a value £X, this must be transformed into a development factor to apply to all open incurred claims with a value £Y.

If we are to believe the development factor model then the answer to these calculations must be the same. Thus, the adjusted development factor to apply only to open claims would be:

As an alternative, instead of the open claim value being the total reported claims from this cohort; it could be the outstanding value of the claims. It is a matter of personal preference which to go for, although we tend to favour the latter.

Beyond this, in practice there are a number of relatively minor issues here.

First, the value of the open claims may be particularly small and the projection of these claims unrealistically high. However, this probably means that the standard projections are also over-stated.

Secondly, it does assume that all development comes from IBNER and ignores pure IBNR. This does not appear to be such a concern in practice as early on in the development Bornhuetter-Ferguson would generally be used which is premium related. Pure IBNR itself would generally be premium related.

Thirdly, the adjusted cumulative development pattern subsequently becomes rather less smooth as it is based on the value of the open claims. It is a matter of personal preference here. Sometimes we would smooth this pattern, other times we would not.

The advantage of smoothing the pattern is that it is this pattern that would subsequently be used. For a development factor model, is the future development related to the value of open claims or to the potentially significant number of closed claims? We would argue the former rather than the latter.

The Pay-off

Having overcome the two issues above, we now have estimated ultimate claims on a contract-by-contract basis.

We have found two major uses for this.

First, we have performed analyses on the sub-lines of business that are too small to have credible volumes of data on their own. Although, when we say 'performed analyses' what we actually mean is run a pivot-table on the results.

For example, for a professional indemnity account we have premiums, estimated ultimate claims and estimated ULRs by exposure type. We can observe the difference between solicitors, architects and engineers. We could also look at the practice area by solicitor. Two-dimensional analysis is feasible so we could look at solicitors by type of practice and geographic area.

This type of analysis can feed straight back into pricing and highlights the areas of the business we should be writing less premium and, just as importantly, more premium.

Secondly, we have taken the output of our analysis for our homogenous reserving groups and run a pivot-table for the SII reserving groups. This ensures that results are entirely consistent with our GAAP reserves.

It also means that subsequent calculations for moving from GAAP to TP are much clearer and to a large extent can be automated from the GAAP projections.

Conclusion

The two ideas from this paper are first, to derive more stable and faster development patterns and secondly, to derive contract-by-contract estimates of ultimate claims. It works within the confines of the standard methods of claim reserving.

The paper does not attempt to solve all the secondary issues that arise from applying these ideas. A few are relatively simple to solve, some are rather trickier, while for others we haven't yet realised that the problem even exists.

For example, for catastrophe events the claims arising would appear at different quarters in the inception date development pattern. This is because the event and subsequent claims are measured relative to the inception date of the policies involved. For such circumstances we would project catastrophe events separately as they would distort the underlying patterns in any case. However, it is possible that other seasonal events could have a distorting effect on patterns.

The ultimate aim here is to extract more value out of the information that is provided as part of the reserving exercise.

Feedback, discussion and debate are part of the process of moving the actuarial profession's methods forward and we welcome all three for this paper.

With grateful thanks to Peter Minton for his actuarial insights.

August 2012