Article from:

The Financial Reporter

December 2005 – Issue 63

Estimating Incurred But Not Reported Life Claims

by Alfred Raws

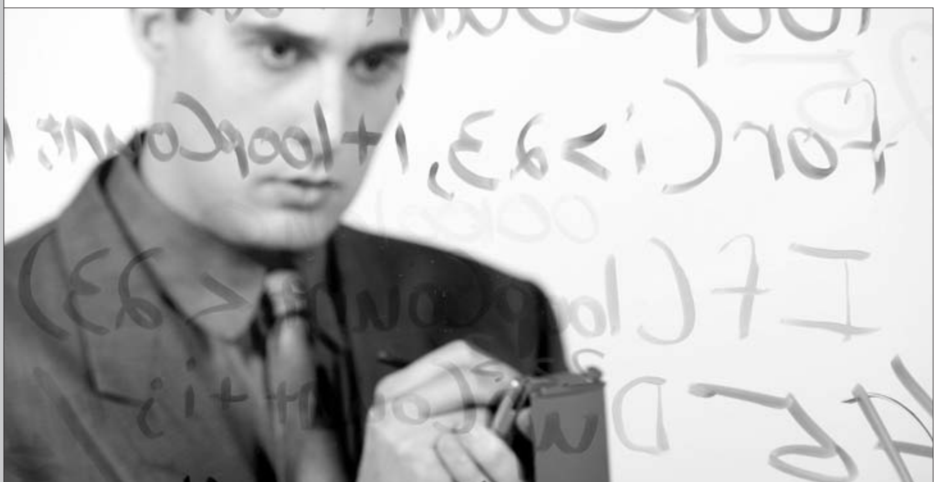

This article is designed to provide a statistical method for reviewing the accuracy of the incurred but not reported (IBNR) life claim liability estimates from prior points in time. The basic concept of this article is to treat the amount of IBNR claims as a random variable. The use of standard statistics related to the distribution of IBNR will provide information to answer questions with respect to the reasonableness of prior estimates of the IBNR claim liability. The first step will be to determine how to estimate *Var(IBNR)*. The second step will be to determine a reasonable function of this amount that will constitute an acceptable difference in our prior IBNR estimates.

The Formula for *Var(IBNR)*

The IBNR amount is actually a combination of two other random variables. The first of these, X, is associated with the net amount at risk for the company's portfolio. The second, N, is associated with the number of claims that are IBNR at any point of time. Then

$$IBNR = \sum_{i=1}^{n} NetAmountAtRisk(i)$$

Under the assumption that the random variables X_i = NetAmountAtRisk(i) associated with claim i are independent identically distributed variables, we have the following:

$$E(IBNR \mid N = n) = E(\sum_{i}^{n} X_i) = nE(X)$$

$$E(IBNR) = E(E(IBNR \mid N = n)) = E(nE(x)) = E(X)E(N)$$

$$Var(IBNR \mid N = n) = Var(\sum_{i}^{n} X_i) = nVar(X)$$

$$Var(IBNR) = E(Var(IBNR \mid N = n)) + Var(E(IBNR \mid N = n))$$

So we have

(I) $Var(IBNR) = Var(X)E(N) + E(X)^2 Var(N)$

The objective is to create an estimate of the right hand side of (I) as a means to estimate *Var(IBNR)*. Our first challenge is to develop an estimate of the components of (I).

A Component of *Var(IBNR): E(X)*

We are looking for an approach that combines ease of calculation with a reasonable degree of accuracy. If you assume that the business is homogeneous, then the calculation of $E(X)$ is straightforward: take the total face in force, subtract the total reserve and divide by the policy count. However, for most companies the block of life business will not be homogeneous. More recent issues tend to have larger face amounts than policies that have been in force for a long time. Different blocks of business have different underwriting standards. The real problem is that the average net amount at risk of all policies in force may not be the average net amount at risk of the policies that die. One possible solution is to segment the business into more homogeneous blocks and then apply the formula above to each block separately. Another choice, and here is where the balance between ease of calculation and accuracy comes into play, is to do a seriatim calculation as follows:

$$E(X) = \sum_{i=1}^{k} (X_i * q(i)) / \sum_{i=1}^{k} q(i),$$

where k is the number of policies in force and q(i) is the expected mortality rate associated with policy i. If you are going to do a seriatim calculation for $E(X)$, you might as well do a seriatim calculation of $Var(X)$ by calculating $E(X^2)$ at the same time.

An Approximation to *Var(IBNR)*

Alternatively, we could assume that the first term in (I) on the previous page is not significant relative to the magnitude of the second term. So a lower bound to (I), and a first approximation to the full calculation, would be

(II) $Var(IBNR) = E(X)^2 * Var(N)$

It would only be necessary to spend the additional effort to calculate the exact value of *Var*(*IBNR*) from (I) if the approximation in (II) does not provide adequate coverage of the observed variation between the historically calculated IBNR and the corresponding actual retrospective amounts.

The Distribution of *X*

It should be noted that the X_i are not, strictly speaking, identically distributed. This is essentially because the "sample" of the in force that is in the IBNR claims has been drawn without replacement. However, given the large number of policies in force in any company, the inclusion or exclusion of one policy cannot affect the parameters very much, so the use of the independence assumption appears reasonable.

We know the following about the distribution of *X*:

- $X \geq 0$
- The values of *X* are heaped at the lower values.
- *X* has a long but comparatively thin tail to the right.

The Weibull distribution, with parameters α (the shape parameter) and β (the scale parameter), meets these requirements, especially for $1 \leq \alpha \leq 2$. In general terms, for fixed β as α gets smaller, the distribution becomes more narrowly heaped to the left, and σ gets larger. In both examples below, μ, the mean of the distribution, has been taken to be 300,000. The following table shows some representative sets of values:

α	β	σ
2.0	338,000	156,379
1.7	336,000	181,627
1.3	325,000	232,985
1.0	311,000	300,000

For smaller values of α the Weibull appeared to be too narrowly distributed around the mean to be reasonable for our purposes.

The Distribution of *N*

In addition, some assumption will need to be made with respect to the distribution of *N*. One possibility would be a uniform distribution over some reasonable range. Considerations that apply to other choices for *N* include the fact that negative values are to be excluded, any tail should be skewed to the right, and the bulk of the distribution should occur in a reasonable interval. Some tracking of the number of historical IBNR claims will provide some insight into this assumption.

If the number of claims is assumed to be uniformly distributed from a to b, then

$E(N) = (a + b) / 2$ and

$Var(N) = [(b - a)*(2*(b - a) + 1)]/6 - [(b - a) /2]^2$

Example 1

Consider the following example:

$E(X) = 300,000$

N uniform on (90,110)

Then the lower bound from (II) is
$Var(IBNR) = 3.3 * 10^{12}$.

We also have: $\sigma(IBNR) \approx 1,816,000$

$E(IBNR) = E(X) * E(N) = (3*10^5) * (100) = 30,000,000$

In the event that $Var(X) = 3.3*10^{10}$, the first term in the right hand side of (I) equals the second term, so now doubling the right hand side of (II) would be a better approximation to *Var*(*IBNR*). What this means is that if $\sigma(X) \approx 180,000$, then $Var(X) \approx 3.3 * 10^{10}$ and (I) gives

$Var(IBNR) = 2 * E(X)^2 * Var(N) = 6.6 * 10^{12}$

$\sigma(IBNR) \approx 2,570,000$

The assumption about $\sigma(X)$ also means that the ratio $\sigma(X)/E(X) = 0.6$. Using (II) implies that $Var(X) << 3.3*10^{10}$, so $\sigma(X)/E(X) << 0.6$. Given the table above it appears that the ratio of 0.6 is much more likely than a substantially lower ratio. So (II) appears to give too low an answer.

Example 2

Modify Example 1 only by assuming that *N* is uniform on (125,175). Then the approximation from (II) is

$Var(IBNR) = E(X)^2 * Var(N) = (3*10^5)^2 *$
$(50* 101/6 - 25^2) = 1.95*10^{13}$

continued on page 30 >>

In the event that the first term in (I) is equal to the second term, we get

$Var(IBNR) = 2*(1.95*10^{13}) = 3.9*10^{13}$

$Var(X) = Var(IBNR)/(2*E(N)) = 3.9*10^{13} /300 = 1.3*10^{11}$

$\sigma(X) \approx 360{,}000$ and

$\sigma(X)/E(X) \approx 1.2$

Based on the table above, it would appear unlikely that this last ratio could be met in practice. However, similar to Example 1, the use of (II) implies that $\sigma(X)/E(X) << 1.2$, which also appears unlikely. So for this example it seems reasonable that

(II) $< Var(X) <$ (I), or $4{,}400{,}000 < \sigma(IBNR) < 6{,}245{,}000$, and $E(IBNR) = 45{,}000{,}000$.

Determining a Decision Rule for Evaluating IBNR Estimates

The object of the second step is to find a function, f(z), which will be used to assess the adequacy of the current estimation process for the reported IBNR. The test to be applied is as follows: if

(III) |actual IBNR – reported IBNR|< f(Var(IBNR)),

then the current estimation process is supported by the statistical test. Otherwise, the current method is not supported by the test.

The challenge here is to choose f(z) in a manner that will be acceptable to someone who reviews the results. Reviewers would include internal management as well as auditors and examiners.

If IBNR is considered to be approximately normally distributed, then use of standard tables permits the determination of the value w such that $f(z) = w*\sqrt{z}$ will produce a boundary in (III) which will support the statistical test p% of the time, for any choice of p. This reduces the issue to choosing p, which is largely a matter of professional judgment.

Alfred Raws III, FSA, MAA, ACAS, is a consulting actuary with KPMG, LLP in Malvern, Penn. He can be reached at arawsiii@kpmg.com.

Consideration of Reinsurance

Applying this approach to the ceded IBNR requires a decision that may affect the outcome. Should the statistics above be based on all claims and in-force policies, or only those on which reinsurance exists? If all policies are considered, then you may have a large number of policies (and claims) with a zero ceded amount. This will have an effect on the expected value and the variance of the ceded amounts, and may, therefore, affect the conclusions reached by the retrospective review. On the other hand, the second choice will have a smaller number of claims and a smaller range of amounts at risk, both of which should reduce the indicated variance described above.

The considerations above can be applied to either the direct, ceded or net liability individually. However, it will be necessary to calculate the variance separately for direct, ceded and net. The point is that variances do not add linearly, so the variance of the net losses does not relate to the variance of the direct and ceded losses. This is doubly true if the choice made with respect to how to count ceded losses means that the number of ceded losses does not agree with the number of direct losses.

An Alternative Approach to Determining the Parameters

When all is said and done, it is always reasonable to ask whether or not a potentially time consuming, theoretically correct process as described here is worth the effort. Why not just take a retrospective look at an adequate number of fully developed dates in the past and calculate the four values that are needed to apply (I): $E(X)$, $Var(X)$, $E(N)$ and $Var(N)$? It may turn out that this approach gives reasonably accurate results. One would expect it to give better results when the two variables X and N are very stable. But if X and N are stable, one would not expect much variation in IBNR, therefore possibly removing the need for the approach in its entirety.

Applying This Approach

Using the above analysis for any given block of business, it should be possible to arrive at an estimate of the variance of the amount of IBNR that will ultimately develop. This calculation should be done periodically, perhaps once a year, until a reasonable pattern is established, and less frequently thereafter. Then as long as the retrospective review of the IBNR at the end of a previous financial period results in a difference between carried and actual that is consistently less than the calculated limit in (III), it would appear reasonable to conclude that the estimation process is producing results that are within an acceptable variance from the actual IBNR. If the retrospective review produces differences that are consistently larger than the calculated standard deviation, then consideration should be given to revising the estimation process.

This approach provides the valuation actuary with a statistical demonstration of the reasonableness of past IBNR estimates, as opposed to relying on more judgmental approaches as to what constitutes a large difference between carried and actual IBNR. S