INDIVIDUAL CLAIM LOSS RESERVING CONDITIONED BY CASE ESTIMATES*

Greg Taylor
Taylor Fry Consulting Actuaries
Level 8, 30 Clarence Street
Sydney NSW 2000
Australia

Phone: 61 2 9249 2901
Fax: 61 2 9249 2999
greg.taylor@taylorfry.com.au

Gráinne McGuire
Taylor Fry Consulting Actuaries
Level 8, 30 Clarence Street
Sydney NSW 2000
Australia

Phone: 61 2 9249 2906
Fax: 61 2 9249 2999
grainne.mcguire@taylorfry.com.au

James Sullivan
Taylor Fry Consulting Actuaries
Level 8, 30 Clarence Street
Sydney NSW 2000
Australia

Phone: 61 2 9249 2912
Fax: 61 2 9249 2999
james.sullivan@taylorfry.com.au

December 2006

* This work was carried out with the assistance of a research grant from the Actuarial Profession (UK), which is gratefully acknowledged.

Summary

This paper examines various forms of individual claim model for the purpose of loss reserving, with emphasis on predictive efficiency in the sense of prediction error in the reserve. Each form of model is calibrated against a single extensive data set, and then used to generate a forecast of loss reserve and an estimate of its prediction error.

The basis of this is a model of the "paids" type, in which the sizes of strictly positive individual finalised claims are expressed in terms of a small number of covariates, most of which are in some way functions of time. Models of this sort are capable of higher predictive efficiency than aggregate models, particularly those aggregate models that do not follow a strict statistical procedure in their calibration (see the example in Section 8.3.1).

The paids model is extended to include other static covariates. While this yields possibly interesting and useful information, it is found to degrade rather than improve predictive efficiency for the data set under consideration (Section 8.3.3). Doubtless, this conclusion would be different for some other data sets.

It is often possible to achieve high efficiency with a model of the "paids" type that has a small number of parameters. Further improvements are possible, but possibly with considerable effort. The construction of an individual claim "incurreds" model involves rather more modelling complexity. For the particular data set considered here, this did not yield any direct improvement in predictive efficiency (Section 8.3.4). Again, however, a different conclusion would be obtained for other data sets.

The paids and incurreds models were used to produce a blended (weighted average) estimate of greater efficiency than either one (Section 8.3.5).

Section 8.3.5 generalises this to a genuine unification of the paids and incurreds models in which forecasts of ultimate individual claim sizes are conditioned by their current case estimates in addition to other covariates. It produces high predictive efficiency, in relation to the present data set at least.

Keywords: case estimate, GLM, individual claim model, loss reserving, prediction error, statistical case estimation

1. Introduction

Conventional actuarial methods of loss reserving are commonly of the "triangulation" type. By this is meant the following.

Let *Y* denote some measure of claims experience, e.g. claim payments, incurred losses, etc. The observations on *X* are summarised into values X_{ij}, labelled according to i = origin period (often accident period) and j = development period. The array $\{X_{ij}\}$ will usually be triangular: $i = 1,2,\ldots,I$; $j = 0,1,\ldots, I\text{-}i$.

A model of the X_{ij} is chosen, fitted to the data array, and then used to forecast future values of X_{ij}, $j = I\text{-}i+1, I\text{-}i+2, \ldots, I$. The forecast claims experience is then manipulated into an estimate of liability for incurred but unpaid claims, and hence a loss reserve. Taylor (2000) provides a summary of this type of methodology.

This sort of approach was undoubtedly useful in the days of manual computation. It remains useful today as a simple approach to lines of business that are either of borderline materiality and/or highly stable statistical structure. It is evident, however, that the summary "triangulation data" discard a great deal of information, some, or even much, of which may be relevant to the forecast required.

From the viewpoint of a forecaster who is not laden with the baggage of loss reserving history and convention, there is no clear reason for using such summary data, and the loss of information provides an incentive (at least *prima facie*) for not doing so.

From this viewpoint, a more natural approach is to forecast future claims experience on the basis of the data in the fullness with which it is available. When the data are presented separately by claim, this will imply the modelling of individual claims. The model is used to forecast the ultimate outcome of individual claims or specified groups of them, again leading to an estimate of the liability for incurred but unpaid claims (henceforth referred to as the **loss reserve liability**), and hence a loss reserve.

A model of individual claims will be called an **individual claim model**, and loss reserving derived from it **individual claim loss reserving**. These are to be distinguished from **aggregate claim models,** which are those based on the aggregate data described in the opening paragraphs of this section, and **aggregate claim loss reserving**.

In view of the large volume of data and model complexity involved in individual claim modelling, it seems pointless for the model to be other than stochastic. Stochastic models are assumed throughout this paper. A stochastic model is capable of delivering the stochastic properties of its forecasts. It will also be assumed throughout that this is an objective.

The natural units of claims experience may not be claims as such. For example, it is common in Motor Bodily Injury insurance to distinguish between "claimants", the injured individuals, and "claims", the accident events that cause one or more of these injuries. In this case, the claimants might be the natural units of experience for analysis. Throughout this paper, "individual claim model" will be taken to include this wider interpretation.

It appears that Norberg (1986, 1993) and Jewell (1989, 1990) were the first to attempt to lay down a comprehensive architecture for individual claims loss modelling. This framework has recently been developed by Larsen (2006). Other specific individual claim models appearing in the literature are due to Hachemeister (1980) and Haastrup and Arjas (1996).

The papers of Norberg and Jewell are (quite properly) very general, whereas some of the others just mentioned are very specific applications. The present paper will attempt a compromise by discussing a very broad family of models, but without particular reference to the general frameworks of Norberg and Jewell.

These models will be formulated in terms of survival analysis and generalised linear models (GLMs). In this, the work will develop previous research by the same authors on individual claim modelling (Taylor and Campbell, 2002; Taylor and McGuire, 2004), and similar research from other authors (Brookes and Previtt, 2004).

The paper will also address a persistent problem faced in loss reserving, specifically whether to rely on paid loss or incurred loss data. The model will attempt to integrate both. This subject has received attention in relation to aggregate claim loss reserving (e.g. Taylor, 1985; Quarg and Mack 2004), but not previously in relation to individual claim loss reserving.

2. Individual claim models and loss reserving

2.1 Basic concepts and definitions

All loss reserving procedures consist of three major components:

- Model specification;
- Parameter estimation (model calibration);
- Forecast of liability.

The first and third of these components will be examined since an individual claim model may appear in the first, and individual claim loss reserving may appear in the third.

Individual claim model

Consider a data set consisting of n claims labelled $i = 1, 2,\ldots,n$. The claims carry some response variable Y_i of interest, and it is assumed that

$$Y_i \sim F\left(.; X_i, \theta\right) \qquad (2.1)$$

where F is a specific d.f., X_i is a vector of claim-dependent covariates, and θ is a vector of parameters that are **independent of the claims**.

Denote

$$\mu_i = E\left[Y_i\right] \qquad (2.2)$$

Explicitly

$$\mu_i = \mu\left(X_i, \theta\right) \qquad (2.3)$$

for some unique function μ.

According to the discussion in Section 1, (2.1) is an individual claim model. Models with other response variables, consisting of summaries of the Y_i, are aggregate claim models.

Individual claim loss reserving

Now consider a second set of *m* claims labelled *j* = 1, 2,…,*m* with covariate vectors X_j^*, but whose responses Y_j will be observable only in the future. A forecast $Y = \sum_{j=1}^{m} Y_j$ is required. If *Y* is forecast by

$$Y^* = \text{function}\left(X_1^*, ..., X_m^*, \hat{\theta}\right) \qquad (2.4)$$

where $\hat{\theta}$ is an estimate of θ, then, again by the discussion in Section 1, a loss reserve based on this forecast is an individual claim loss reserve.

A more specific form of (2.4) is

$$Y^* = \sum_{j=1}^{m} Y_j^* \qquad (2.5)$$

where

$$Y_j^* = \mu\left(X_j^*, \hat{\theta}\right) \qquad (2.6)$$

Here, Y_j^* is a forecast of the response of individual claim *j*, and is a special case of individual claim loss reserving that will be referred to as **statistical case estimation**, such as discussed by Brookes and Previtt (2004) and Taylor and Campbell (2002).

An example of individual claim loss reserving that is **not** statistical case estimation appears in Taylor and McGuire (2004), and the same type of model appears in Section 3 below. Briefly, in these models:

- the response variable is size of **completed** claim (open claims are excluded from the data);
- this enables the forecast of the ultimate claims cost of each accident period, since this is the total cost of all its claims, when completed;
- this forecast of ultimate cost does not require dissection between claims that are reported at the point of forecast and those that are unreported;
- hence there are no separate forecasts for these two groups of claims, and certainly no forecasts for individual reported claims.

2.2 Static covariates

Consider the covariate vector X_i. Certain of its components X_{ik} may assume values which do not change over the life of a claim. For example, the type of vehicle in which a motor accident claimant was an occupant would be such a covariate. These will be called **static covariates**.

2.3 Dynamic covariates

Covariates which are not static will be called **dynamic covariates**. By definition, they are liable to change over the lifetime of a claim. An example would be the case estimate (variously called the manual estimate or physical estimate) of ultimate incurred loss associated with a claim.

There are two fundamentally different categories of dynamic covariate, which are discussed in the following two sub-sections.

2.3.1 Time covariates

A number of covariates are likely to be time-related. Some may change with the passage of time. Development period would be such a covariate. However, the changes in their values are predictable. In the example just given, the value of development period is known for each future calendar period. If the development period last observed is d, then future periods will be $d+1$, $d+2$, etc.

These covariates will be called **time covariates**. They may be included among model covariates in such a way that model forecasts depend on their future values X^*_{jk} in the notation of Section 2.1, since these values are predictable.

2.3.2 Unpredictable dynamic covariates

Other dynamic covariates will change over time in a manner that is not predictable with certainty. The case estimate referred to above provides such an example. Such covariates will be called **unpredictable dynamic covariates**.

Inclusion of these among model covariates in such a way that model forecasts depend on their future values X^*_{jk} needs to be approached with much greater caution. If they are included, then a further model will be required to make predictions of these uncertain quantities.

Whether this is fruitful will depend on whether the formulation of the additional model, with its own prediction error, will improve or degrade the reliability of the ultimate forecasts. The more essential the covariate concerned is to an accurate model specification, and the smaller its prediction error, the greater the likelihood that it will improve the ultimate predictions.

Conversely, the inclusion of covariates that are of only moderate structural relevance and whose prediction is difficult is likely to degrade the performance of the overall model.

On the basis of Sections 2.2 and 2.3, the covariate vector X_i may in general be decomposed as $X_i = [X_i^{(S)}, X_i^{(T)}, X_i^{(U)}]$ where denote $X_i^{(S)}, X_i^{(T)}, X_i^{(U)}$ denote static, time and unpredictable dynamic covariate vectors respectively.

The meaning of the forecast (2.6) is straightforward if X^*_j contains only static covariates, but not otherwise. If it contains unpredictable dynamic covariates, their values will be unknown. Even time covariates, though predictable, may be unknown at the point of measurement of the response variable. For example, if X^*_{jk} is the time variable, development period of finalisation of the claim, its value will be unknown ahead of that time.

It will be reasonable to take a probability-weighted average of forecasts over these unknowns, so that (2.6) is extended to the following:

$$Y_j^* = \int \mu\left(X_j^*, \hat{\theta}\right) dP\left(X^{(T*)}, X^{(U)}\right) \tag{2.7}$$

where $X^{(T*)}$ is the subset of $X^{(T)}$ requiring prediction, and P is a measure on $X^{(T*)}, X^{(U)}$.

3. Individual claim loss reserving without case estimates

Henceforth in this section it will be supposed that the response variable Y_i is the ultimate cost of the i-th claim when finalised. If the vectors $X^*_1,\ldots,X^*_m$ appearing in (2.4) relate to the claims incurred in a portfolio prior to some valuation date but unfinalised at that date, then Y^* given by (2.4) is a forecast of the ultimate cost of those claims.

The unpaid cost of these claims at the valuation date is equal to Y^* less the amount paid to date in respect of the same claims. The loss reserve is based on this quantity, perhaps with discounting for investment return, addition of a safety margin, etc.

3.1 Individual claim models dependent on only time covariates

Consider initially a model based just on time covariates. In this case, (2.1) becomes

$$Y_i \sim F(.\,; X_i^{(T)},\theta) \tag{3.1}$$

One may be able to invent many components of $X_i^{(T)}$, but there are several obvious candidates, namely

a_i = accident period;
d_i = development period of finalisation;
p_i = a_i+d_i = experience period of finalisation (calendar period in which finalisation occurs);

where in each case the subscript i indicates that the value of the time variable concerned is that observed for the i-th claim.

To this small collection may be added **operational time**, denoted t_i. This concept from stochastic processes (see e.g. Feller, 1971) was introduced into the actuarial literature by Bühlmann (1970), and first applied to loss reserving by Reid (1978).

It is defined as follows. Suppose that, for a given but arbitrary accident period with a count N of claims incurred, a development time scale δ is defined to commence (δ=0) at the beginning of the accident period. Suppose that the i-th claim is finalised at $\delta=\tau_i$, i=1,2,…,N. Now define a new time scale

$$t_i(\delta) = \#\{ \tau_i \leq \delta \}/N \tag{3.2}$$

This is the operational time. It takes values in the interval [0,1], and maps development time δ to the proportion of the accident period's incurred claims that have been finalised at or before that time.

For many purposes, a_i, d_i, p_i, t_i form an adequate set of time variates for individual claim modelling.

3.2 Individual claim models dependent on only a_i and t_i

Consider the special case of (3.1) in which

$$Y_i \sim F(. \,; a_i, t_i \,;\theta) \tag{3.3}$$

Such models have very special properties. Consider a single accident period a with N claims incurred, some possibly unreported as yet, and suppose that F of these have been finalised at some chosen valuation date.

This means that operational time for the accident period stands at F/N and, over the future, will traverse the values $(F+1)/N$, $(F+2)/N$,…,1. Hence, if U_a denotes the ultimate cost of claims unfinalised at the valuation date, then

$$E[U_a] = \sum_{j=F+1}^{N} \mu(a, j/N\,;\theta) \tag{3.4}$$

which is estimated by

$$U^*_a = \sum_{j=F+1}^{N} \mu(a, j/N; \hat{\theta}) \tag{3.5}$$

Note that this quantity is fully determined by a and $\hat{\theta}$. These parameters determine the ultimate costs of the last (N-F) claims finalised, and in fact the cost of each of those claims taken in order. However, the order of finalisation of the claims open at the valuation date (including claims then unreported) is not specified, and so no individual forecasts is associated with these.

The forecast (3.5) is, therefore, **not** an example of (2.5) and (2.6), and so does not represent statistical case estimation.

3.3 Individual claim models dependent on time variates other than a_i and t_i

Consider now the case in which model (3.5) is enlarged to include either or both of d_i, p_i among its covariates. The case of p_i may be considered without loss of generality:

$$Y_i \sim F(. \,; a_i, t_i\,, p_i\,;\theta) \tag{3.6}$$

The forecast (3.5) now becomes (see (2.7)):

$$U^*_a = \sum_{i=F+1}^{N} \int \mu(a, j/N, p_j; \hat{\theta})\, dP(p) \qquad (3.7)$$

where p_j now denotes the experience period of finalisation of the j-th claim taken in order of finalisation, and $P(p)$ is the distribution associated with the vector $[p_{F+1},\ldots, p_N]$.

For example, the integrand in (3.7) might take the form:

$$\mu(a, j/N, p_j; \hat{\theta}) = \exp\,[\text{function}(a, j/N; \hat{\theta}) + \beta\, p_j] \qquad (3.8)$$

where the term involving p_j represents inflation at the rate exp β per period.

The forecast U^*_a still does not require a component forecast to be associated with each of the claims open at the valuation date, and still does not represent statistical case estimation. It is, however, fundamentally different from (3.5) in requiring a (stochastic) mapping of the operational times j/N to real times p_j.

This requires a further model of the connection between these two time scales. This is commonly achieved by modelling finalisation rates (with respect to real time), and is discussed further in Section 5.

3.4 Individual claim models dependent on time and static variates

If static variates are incorporated in the model, it becomes:

$$Y_i \sim F(.\,; X_i^{(S)}, X_i^{(T)}, \theta) \qquad (3.9)$$

For example, $X_i^{(S)}$ might denote gender of claimant i. This means that the forecast U^*_a depends on the composition of the (N-F) unfinalised claims with respect to $X_i^{(S)}$. A natural way of recognising this composition is to put U^*_a in statistical case estimation form. Other approaches tend to be cumbersome.

If this is done, the forecast is changed fundamentally from (3.5) and (3.7). It becomes (2.5) with

$$Y^*_k = \int \mu(a, t(p_k), p_k, X^{*(S)}_k; \hat{\theta})\, dP(p_k | X^{*(S)}_k) \qquad (3.10)$$

where

- μ(.) is the statistical case estimate in respect of the k-th unfinalised claim (taken in any order, as distinct from the j-th above, taken in chronological order of finalisation), conditional on a, t_k, p_k, $X^{*(S)}_k$;
- t is the mapping of real time p to operational time; and
- the measure $P(.)$ on p_k may now depend on $X^{*(S)}_k$.

3.5 Individual claim models dependent on unpredictable dynamic covariates

As mentioned in Section 2.3.2, the inclusion of these covariates generates a need for further models over which forecasts must be averaged, as in (2.7). This adds considerably to the complexity of the modelling and forecasting, and models of this type are not pursued further here.

3.6 IBNR claims

The forecasts in Sections 3.2 and 3.3 include allowance for IBNR claims at the valuation date as they incorporate all remaining finalisations. However, any forecast of the statistical case estimate type, as in Section 3.4, assigns a forecast just to each known claim, and therefore includes no allowance for IBNR claims

It will be supposed that the number of these claims has already been estimated. A model will be required to forecast their sizes. This will need to be a simplified version of (3.9), because none of the covariates $X^*{}_k^{(S)}$ will be available before a claim is reported.

The simplest approach to IBNR claims would be to forecast their sizes on the basis of a model formulated by simply deleting $X^*{}_k^{(S)}$ from (3.9), i.e. by re-fitting the data $\{Y_i\}$ with a model of the form (3.1).

4. Individual claim loss reserving conditioned by case estimates

4.1 Model structure

The models discussed in Section 3 fall generally within the family that actuaries think of as "paids" models. They depend on only paid losses, and not at all on the insurer's various estimates of those losses through the lifetimes of the claims.

A conventional alternative form of model forecasts ultimate losses on the basis of the insurer's estimates at the valuation date. This is usually referred to as an "incurreds" model. A statistical case estimation form of it will now be considered.

In the following, the term **case estimate** will be used to mean a subjective estimate of ultimate incurred loss placed by an insurer on an individual claim. The estimate will usually be made by a claims assessor, and is often referred to as a **physical estimate** or **manual estimate.**

The general thrust of a model that uses case estimates as predictors of ultimate incurred cost will be concerned with the factor relating these two quantities. This is the ratio usually referred to as the **age-to-ultimate ratio**, defined as:

$$\text{Age-to-ultimate ratio} = \frac{\text{Expected ultimate incurred loss}}{\text{Current case estimate}} \qquad (4.1)$$

where the "Expected" in the numerator is used in its statistical sense. Such a model will be referred to here as a **case estimates model**. It corresponds to the conventional "incurreds" model.

However, the possibility of zero numerator or denominator in (4.1) complicates the situation, with the result that the model needs to consist of a number of sub-models as illustrated in Figure 4.1, where I denotes the current case estimate, U the ultimate incurred loss, and R=U/I. F(.) denotes a distribution function.

Figure 4.1
Sub-model structure of a case estimate model

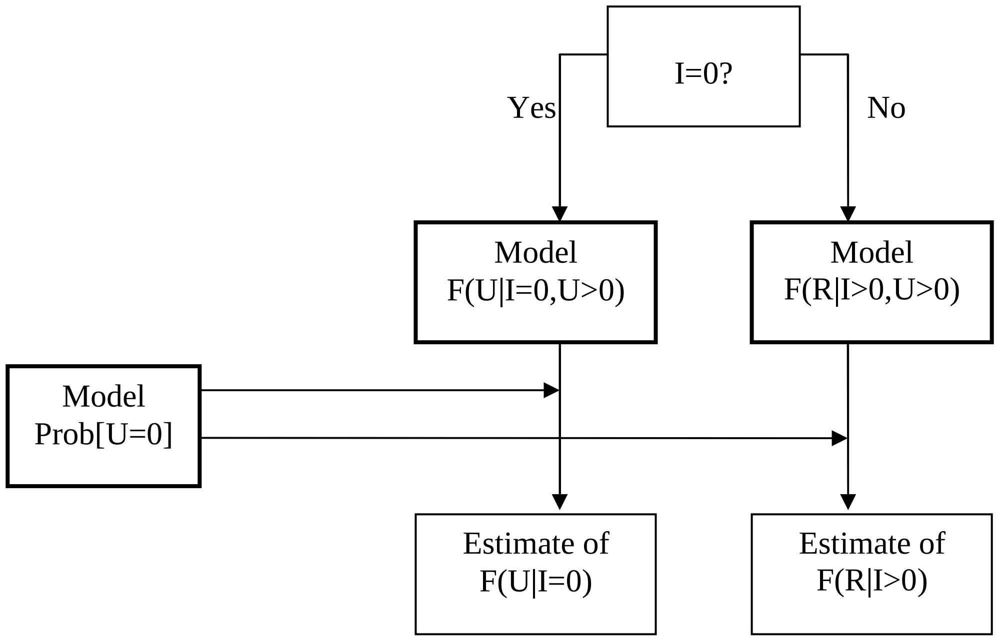

Note that the cases I=0 and I≠0 are dealt with separately in the figure, as are the cases U=0 and U≠0. This is because:

- In the case I=0, the ratio R does not exist, and so the size of U must be modelled directly, rather than as U=IxR.
- The distribution of U, with a discrete mass at U=0, is best modelled by separate recognition of the mass and the remainder of the distribution (assumed continuous).

This creates the need for three separate sub-models, indicated by the heavily outlined boxes in the figure. Further sub-models will be required for IBNR claims (Section 3.6) and finalisation rates (Section 5). Note, however, that only one model of Prob[U=0] is required. There is no difficulty in accommodating the cases of predictors I=0 and I≠0 within that one model.

The three sub-models are subject to the same considerations as discussed in Sections 3.2 to 3.4. In addition, the model of Prob[U=0] requires formulation as binomial.

4.2 Form of observations

The *i*-th observation was described at the start of Section 3 as comprising the ultimate cost of the *i*-th claim when finalised. The data set thus consists of one record per claim finalised. The data set for a case estimate model requires a different format because each claim involves a whole sequence of case estimates through its lifetime.

Consider a claim from accident period *a*, reported in development period *r*, and finalised in development period *f*. It will carry a case estimate I_d at the end of development period $d=r,r+1,\ldots,f-1$, and then the ultimate cost U. This implies a total of *f-r* records.

It would be possible to organise these as *f-r* records, each relating to a single development period, tracking the development of the case estimate from I_d to I_{d+1}. However, this would create difficulties of two types.

First, the case estimates would become unpredictable dynamic covariates of the sort discussed in Section 2.3.2, generating difficulties of the sort hinted at there. Second, any feasible means of forecasting future case estimate development would be likely to involve a highly dubious Markovian assumption.

It seems preferable to create multi-period data observations of the case estimate development from the end of development period $d=r,r+1,\ldots,f-1$ to finalisation in development period *f*. Thus, all observations periods would end at finalisation, and would be of an “age-to-ultimate” nature, as illustrated in Figure 4.2.

Figure 4.2
Structure of case estimate development observations

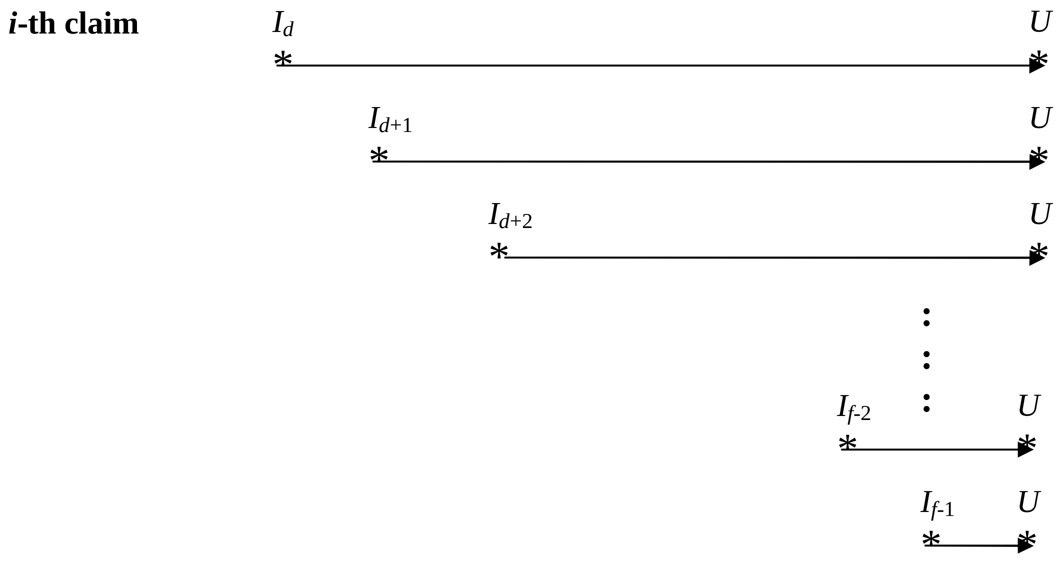

While this data structure solves the two problems mentioned above, it calls for comment on one other point. It is that the observations are unlikely to be stochastically independent. For example, in the data set that forms the subject of Section 8, case estimates typically remain constant for a number of consecutive quarters.

The GLM vehicle mentioned in Section 1 assumes independence of observations and so is not applicable to the data set. Accommodation of a dependency structure on observations would require the replacement of GLMs by Generalised Estimating Equations (GEEs) (Liang and Zeger, 1986; Zeger and Liang, 1986), but this has not been done here. Instead, one record from the multi-period observations corresponding to each claim has been sampled at random. This removes the dependency while retaining a selection of records over different period lengths.

5. Modelling claim finalisation rates

It is typical actuarial practice for the finalisation rates to be measured over single periods. Each rate is defined as the number of claims finalised in the period divided by some measure of exposure to finalisation, such as the average number open over the period. From these future operational times may be derived and thus finalisation of claims corresponds to the advancement of operational time.

Making use of this fact, a very simple model of future claim finalisation has been constructed in the form:

$$\Delta t_i(d) = t_i(d) - t_i(d\text{-}1) = \omega_i\ \delta_i(d) \qquad (5.1)$$

where $t_i(d)$ denotes the operational time at the end of development period d of origin period i, and $\Delta t_i(d)$ denotes the increment in operational time over that development period.

The model on the right of (5.1) consists essentially of selected increments $\delta_i(d)$, specific to origin periods, accompanied by an adjustment factor ω_i which is usually close to 1 and such as to ensure that

$$t_i(\infty) = t_i(d^*) + \sum_{d=d^*+1}^{\infty} \Delta t_i(d) = t_i(d^*) + \omega_i \sum_{d=d^*+1}^{\infty} \delta_i(d) = 1 \qquad (5.2)$$

where d^* is the value of d (for origin period i) at the valuation date. The effect of (5.2) is to ensure that the forecast numbers of finalisations for origin period i account for all claims unfinalised at the valuation date.

The model values $\delta_i(d)$ have been selected simply by inspection of historical increments. While they are in principle specific to individual origin periods, as a practical matter the same values are selected for **groups of origin periods**.

Details of such a model appear in Appendix C.1. It is an aggregate model, in the sense of Section 1, in that it does not seek to assign probabilities of finalisation to an individual claim that depend on the attributes of that claim. Such simple aggregate models often operate very satisfactorily in conjunction with aggregate models of claim sizes. However, the more individualised become the latter, the greater the need for individualisation of the former.

If, for example, the claim size model is capable of differentiating claim attributes sufficiently as to produce high and low forecasts of size for some claims, then it may not be reasonable to assume that all claims are subject to the same schedule of probabilities of finalisation. To do so may produce unreliable forecasts of liability.

It is then preferable to model such probabilities for each claim as dependent on the attributes of that claim. This is most naturally done by means of **survival analysis** (see e.g. Lee, 1992). This means that, for each i, the i-th claim, with vector X_i of covariates, just as in earlier sections, has a lifetime T_i, from reporting to finalisation, assumed subject to a **survival function** $S_i(.)$ such that

$$\text{Prob}[T_i > t] = S_i(t;X_i) \qquad (5.3)$$

The **hazard rate** associated with the i-th claim is $h_i(t) = -S'_i(t)/S_i(t)$. A convenient form for the present application is the **proportional hazards** form (Cox, 1972):

$$h_i(t) = \exp[X_i^T \beta] \qquad (5.4)$$

where β is a vector of parameters and the upper T denotes matrix transposition.

Note that, according to (5.4), the probability of a claim's finalisation in a specific period depends on the characteristics of the claim. For example, a claim carrying a larger case estimate might have a lower probability than one carrying a smaller case estimate, all other things equal.

The fitting of a survival model of this sort to survival data is called **Cox regression**. This has been applied to the survival of claims, with each claim regarded as:

- commencing its life when reported to the insurer;
- dying when finalised;
- having observation of its lifetime right censored (terminated other than by finalisation) if still open (alive) at the valuation date.

Details of the model are given in Appendix C.2.

6. Estimation of prediction error

Taylor (2000, Chapter 6) discusses a couple of components of the prediction error associated with a loss reserve. Generally this prediction error can be dissected into three components:

- **Model specification error** (self-explanatory);
- **Parameter error** (due to sampling error in the estimates of model parameters);
- **Process error** (random noise within the model).

The last two of these three components are usually regarded as quantifiable from the data, and they will be addressed here. Specification error is ignored, not because it is insubstantial but because it does not fit readily into the quantitative framework applied here.

For advice on how to estimate it, see O'Dowd, Smith and Hardy (2005). In any event, it should be remembered that specification error needs to be added to all of the estimates of prediction error reported in Section 8.

Prediction error can often be estimated by means of the **bootstrap** (Efron and Tibshirani, 1993). Here an abbreviated form of that procedure is applied to each

GLM and, for succinctness, it will be referred to here as the bootstrap. The framework for estimating prediction error is discussed below.

6.1 Prediction error of GLM forecasts

Consider a GLM of form:

$$E[Y_i] = h^{-1}(X_i^T \beta) + \varepsilon_i \tag{6.1}$$

where Y_i is the response variable, $X_i^T \beta$ has the same interpretation as in (5.2), and ε_i is a centred stochastic error term with d.f. $F(.)$.

Suppose that one is concerned with forecasting a quantity:

$$U_i = g(Z_i^T \beta) + \eta_i \tag{6.2}$$

where $g(.)$ is some function, Z_i is some new vector specific to U_i, and η_i is another centred stochastic error term with d.f. $F(.)$.

Let $\hat{\beta}$ and C denote the GLM estimate of the vector β (assumed unbiased) and the estimated covariance matrix of $\hat{\beta}$.

A forecast $\hat{U}_i$ of U_i is given by

$$\hat{U}_i = g(Z_i^T \hat{\beta}) \tag{6.3}$$

Diagonalise C:

$$C = P^T DP \tag{6.4}$$

for orthogonal P and diagonal D. If γ denotes $P\hat{\beta}$, then $E[\gamma] = P\beta$ and $Var[\gamma] = D$.

Assume a normal distribution for $\hat{\beta}$, so that

$$\gamma \sim N(P\beta, D) \tag{6.5}$$

Random realisations γ^* of γ are drawn in accordance with (6.5), and random realisations η^*_i of η_i drawn from $F(.)$. Then, by (6.2), a random realisation of U_i is given by

$$U^*_i = g(Z_i^T P^T \gamma^*) + \eta^*_i \tag{6.6}$$

The collection of replicates U^*_i may be considered as forming the empirical distribution of $\hat{U}_i$. This procedure is similar to that discussed by England and Verrall (2002).

6.2 Prediction error of Cox regression forecasts

Consider a survival model of the form (5.2), applied to claim finalisations in the way described at the end of Section 5. Again, let $\hat{\beta}$ and C denote the estimate of the vector β (again assumed unbiased) and the estimated covariance matrix of $\hat{\beta}$.

Let $\hat{S}_i(t;X_i)$ denote the survival function estimator obtained by replacing β by $\hat{\beta}$ in (5.2), and consider the discrete probabilities:

$$\hat{q}_{ir} = [\hat{S}_{ir}(t+r;X_i) - \hat{S}_{ir}(t+r+1;X_i)] / \hat{S}_{ir}(t;X_i) \tag{6.7}$$

which is an estimate of the probability that the i-th claim, currently open, is finalised in the (r+1)-th period ahead.

Random realisations $\hat{q}_r^*$ of $\hat{q}_r$ can be obtained by sampling the parameter vector $\hat{\beta}$ in the same way as described in Section 6.1. Process error (corresponding to η^*_i in (5.2)) is added by recognising that, for given $\hat{S}_i(t;X_i)$, finalisation period is a multinomial variate with probabilities $\hat{q}_{ir}$ and sampling it as such. This provides a set of replicates R^*_i of the number of quarters ahead that the i-th claim is finalised.

7. Combination of models

7.1 Blending of models based on paid and incurred losses

Sections 3 and 4 discuss various forms of model that may be fitted to the same data set. If forecasts are produced from more than one of these models, the question will arise as to how a single final forecast is obtained from them.

This question is discussed by Taylor (1985), who optimises convex combinations of the various forecasts. Specifically, suppose that there are m estimates of loss reserve, each by accident period a. Let $L_a^{(s)}$, s=1,2,…,m; a=1,2,…,A, denote the estimate for accident period a from the s-th model. The final estimate is

$$L_a = \sum_{s=1}^{m} w_a^{(s)} L_a^{(s)} \tag{7.1}$$

where the weights $w_a^{(s)}$ are chosen, with $\sum_{s=1}^{m} w_a^{(s)} = 1$ for each a, so as to minimise the following loss function that compromises between low prediction error and smoothness of the forecasts over periods of origin:

$$L^T ML + \lambda_1 (KR)^T(KR) + \lambda_2 \sum_{s=1}^{m} (K w^{(s)})^T(K w^{(s)}) \tag{7.2}$$

where

L=[L_1, L_2,…, L_A]^T;
M=Cov[L], the predictive covariance matrix of L;
R=[R_1, R_2,…, R_A]^T with $R_a = L_a / Q_a$, and Q_a denoting the case estimate of L_a;
K a matrix that takes d-th differences of its operand;
$w^{(s)}$ is the vector of weights $w_a^{(s)}$;

and λ_1, λ_2>0 are constants that set the compromise between low prediction error, smoothness of the ratios R_a over a, and smoothness of the weights $w_a^{(s)}$ over a.

Note that R_a has a meaning here slightly different from that assigned to it in Figure 4.1, where it represents the ratio of ultimate to case estimate of **total** incurred loss, not outstanding liability as here.

7.2 A unified model based on paid and incurred losses

Consider the case of a=1,2 in Section 7.1, where Models 1 and 2 are “paids” and “incurreds” models of strictly positive finalisations respectively, in the terminology of Section 4.1. Consider only claims with strictly positive current case estimates of incurred loss.

Suppose also that both models are GLMs with log link. Then Model 1 takes the form:

$$Y_i = \exp [X^{(P)}_i \beta^{(P)}] + \varepsilon^{(P)}_i \qquad (7.3)$$

and Model 2:

$$R_i = \exp [X^{(I)}_i \beta^{(I)}] + \varepsilon^{(I)}_i \qquad (7.4)$$

where R_i now has the meaning assigned in Section 4.1, the P and I suffixes denote “paids” and “incurreds” respectively, the $X^{(\cdot)}_i$ denote the covariates of the i-th finalised claim, the $\beta^{(\cdot)}$ the vectors of model parameters, and the $\varepsilon^{(\cdot)}_i$ centred stochastic error terms.

If I_i denotes the current case estimate of incurred loss for the i-th claim, then $Y_i = I_i R_i$, and (7.4) may be re-written as

$$Y_i = \exp [\ln I_i + X^{(I)}_i \beta^{(I)}] + \varepsilon^{(I)}_i \qquad (7.5)$$

The model blending procedure of Section 7.1 takes weighted averages of the model predictions:

$$Y^*_i = w \exp [X^{*(P)}_i \beta^{(P)}] + (1-w) \exp [\ln I_i + X^{*(I)}_i \beta^{(I)}] \qquad (7.6)$$

where the weight w is temporarily assumed the same for all claims.

The expression on the right side of (7.6) is an arithmetic weighted average of the model predictions. A geometric weighted average would be an equally valid choice, and yields the more convenient form:

$$Y^*_i = \exp \{w X^{*(P)}_i \beta^{(P)} + (1-w) [\ln I_i + X^{*(I)}_i \beta^{(I)}]\}$$

$$= \exp \{X^{*(P)}_i [w\, \beta^{(P)}] + (1\text{-}w) \ln I_i + X^{*(I)}_i [(1\text{-}w)\, \beta^{(I)}]\} \qquad (7.7)$$

The exponent on the right side is a very special form of linear combination of ln I_i, $X^{*(P)}_i$ and $X^{*(I)}_i$, which may now be generalised to the following:

$$Y^*_i = \exp [(1\text{-}w) \ln I_i + X^{*(U)}_i\, \beta^{(U)}] \qquad (7.8)$$

where the upper U denotes a **unified** paids and incurreds model with $X^{(U)}_I$ the union of $X^{(P)}_i$ and $X^{(I)}_I$, and w, $\beta^{(U)}$ are the model parameters.

This implies that the unified model takes the form:

$$Y_i = \exp [(1\text{-}w) \ln I_i + X^{(U)}_i\, \beta^{(U)}] + \varepsilon^{(U)}_{\mathrm{i}} \qquad (7.9)$$

where $\varepsilon^{(U)}_{\mathrm{i}}$ is a centred stochastic error term with distribution suitably chosen as in some sense "between" those of $\varepsilon^{(P)}_{\mathrm{i}}$ and $\varepsilon^{(I)}_{\mathrm{i}}$.

This model is more general than that underlying (7.7), and so certainly includes (7.7). It should therefore provide more efficient forecasts. Since it is more general than (7.7), it is no longer precisely a weighted average of paids and incurreds models. However, to the extent that one chooses to view it as such, (7.9) implies weights of *w* for paids and 1-*w* for incurreds.

As discussed in Section 3.4, it is necessary to deal with the dependency between the multi-period observations of the same claim. The same procedure is followed here as above; one multi-period observation is sampled at random for each claim.

It was assumed at the start of the present sub-section that both I_i and R_i are strictly positive. Other cases require separate treatment, just as in Figure 4.1.

8. Numerical examples

8.1 Data set

The data set is taken from a long tail portfolio from a liability line of business that includes risks underwritten on both Occurrence and Claims-Made bases over a period of about 30 years. It consists of roughly 23,000 unit records in respect of claims completed at strictly positive cost, and another 38,000 completed at nil cost.

There were no Claims-Made risks underwritten in the earliest 7 or 8 origin years. Subsequently, such risks constituted a steadily increasing proportion of those underwritten, though the proportion of Occurrence coverages has always remained substantial, and risks of this type still dominate the loss reserve.

Each claim record consists of:
- Certain time covariates;
- Certain static covariates;

- Final cost of claim;
- Quarterly history of case estimates of incurred loss.

Paid losses are heavily concentrated at the date of completion. Indeed, paid losses to date in respect of open claims are less than 3% of the estimated outstanding liability in respect of notified claims.

8.2 Conventional (Mack) stochastic model

Paid and incurred losses have been summarised by **origin quarter** and development quarter. Origin quarter means that in which the claim originates. For coverages of the Occurrence type, this may be thought of as accident year, but such terminology is not appropriate to Claims-Made coverages.

Development quarter is defined relative to origin quarter, and has its usual meaning.

This leads to the conventional triangular representation of the data set, which has been subjected to a chain ladder analysis. The resulting chain ladder estimate of loss reserve is taken as a baseline against which later estimates may be compared.

The Mack method (Mack, 1993) has been used to estimate the mean square error of prediction (MSEP) associated with the chain ladder estimate. In order that the chain ladder estimates be objective, the method has been applied blindly, in that age-to-age factors have been averaged over all origin quarters, and not subjected to any smoothing. This has three major consequences.

First, the paid chain ladder produces obviously false results. Its estimate of loss reserve, including IBNR, is about treble those produced by demonstrably reasonable methods. This appears to occur because of the shifting composition of Occurrence and Claims-Made exposures.

The latter increase relative to the former over origin periods, but are subject to a shorter payment pattern. This means that the correct age-to-age factors fall increasingly short of the historical averages as origin period increases, and the averages therefore cause dramatic over-estimation of loss reserve for the more recent origin periods.

This difficulty might have been overcome by separate chain ladder modelling of the two coverage types. If this were done, however, the absence of Claims-Made exposures in the earliest origin periods would leave this coverage type with no experience on which to estimate the payment tail.

In view of these difficulties the paid chain ladder has been discarded.

Second, the incurred chain ladder yields negative estimates of loss reserves for the earlier origin periods. The difficulties associated with the paid chain ladder do not necessarily affect the incurred chain ladder. However, age-to-age factors in this portfolio tend to be substantially less than unity, and even slight under-estimation of them can send ultimate incurred losses for an old origin period below losses paid to date.

Third, blind application of the chain ladder (or, for that matter, of any other model) can be expected to produce less efficient prediction than more considered modelling, and so can be expected to generate unnecessarily high prediction error.

Ultimately, the incurred chain ladder estimates of loss reserve have been retained here as baseline estimates, with all negative estimates of liability for individual origin years reset to zero. These estimates need to be read as subject to qualification on the above grounds.

An incurred chain ladder forecast includes allowance for IBNR claims. As noted in Section 3.6, forecasts of the type described in Sections 3.2 and 3.3 also naturally do this. However, it is remarked in Section 4.1 that forecasts based on models of case estimate development, such as discussed there, include no IBNR allowance.

It is necessary to place all forecasts on the same footing with respect to IBNR claims if they are to be compared. An "excluding IBNR basis" has been chosen for this purpose. All individual claim loss reserves estimated in Section 8.3 exclude IBNR allowance by the procedure outlined in Section 8.3.1. The chain ladder forecast has been placed on the same basis by subjecting the "including IBNR" estimate to the adjustment:

$$\text{Forecast liability excluding IBNR} = \text{Forecast liability including IBNR} \times \frac{\text{Section 8.3.2 forecast liability excluding IBNR}}{\text{Section 8.3.2 forecast liability including IBNR}} \qquad (8.1)$$

where the adjustment is applied separately to each origin quarter within coverage type, and each of the required numbers of IBNR claims has been separately estimated..

This adjustment invokes an implicit assumption that the open reported claims and the IBNR claims are sampled from the same size distribution. This may involve some approximation, and the comparisons between forecasts below may be slightly distorted as a result. A more complete treatment would incorporate an explicit forecast of the cost of IBNR claims, where needed, along the lines discussed in Section 3.6.

Table 8.1 reports this loss reserve estimate in total and its prediction error. Appendix D.1 gives the corresponding information on origin years, and this is represented in Figure 8.1. Origin years ("accident years" in this and subsequent figures) are numbered just 1, 2, etc, starting at the earliest.

For the sake of anonymity of the insurer providing the data set, all forecasts of loss reserve have been scaled in proportion so that the one in Section 8.3.2 totals $1,000M. The relativities between them are thus preserved.

Table 8.1
Chain ladder loss reserve

Estimated loss reserve	$888M
Predictive coefficient of variation ("CoV")	10.5%

Figure 8.1
Chain ladder loss reserve

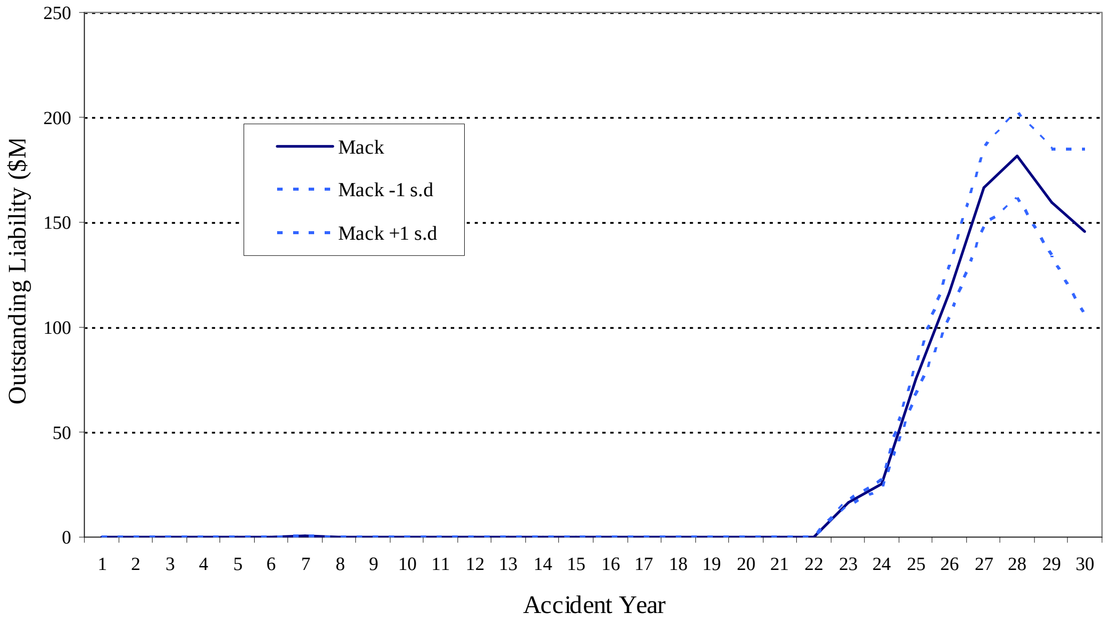

The figure refers to the chain ladder estimates as "Mack" because the Mack method has been used to estimate the prediction CoV. It shows the chain ladder forecasts with an envelope of one CoV in each direction. The envelope is seen to grow wide (27%) in relation to the most recent origin years.

8.3 Individual claim models

8.3.1 Forecast procedure

For each of the individual claim models discussed in Sections 8.3.2 to 8.3.5, the same broad outline of forecast procedure was followed. This consisted of the following steps. Let the valuation date be taken as the end of calendar quarter q.

Step 1: The simple aggregate model set out in Appendix C.1 was used to forecast, for each origin quarter a, **operational times** $t_{a:q+1}$, $t_{a:q+2}$,… at the ends of quarters q+1, q+2, etc. These forecast operational times **included** allowance for claims that were IBNR at the valuation date so that finalisation numbers derived from these forecasts will include IBNR claims. In other words, the total number of finalisations, past and forecast future, for an origin quarter was equal to the estimated total number incurred.

Step 2: For each claim i **open** at the valuation date, the probability of finalisation in quarters q+1, q+2, etc was estimated from one of the models of Section 5. The first model is that used in Step 1 (the future operational time increments being used to form a probability distribution in this case) but the second (the Cox regression model) produces probabilities that are claim-specific and relate to open claims, and so do not necessarily relate to the forecasts in Step 1 (which related to all unfinalised claims open or IBNR). More detail on the derivation of the probabilities appears in the sub-sections relating to the different models below. Let the probabilities be denoted $f_{i:q+1}$, $f_{i:q+2}$,… .

Step 3: For each of these claims, the **probability of a finalisation for zero cost**, given that finalisation occurs in quarters q+1, q+2, etc, was estimated from the model of Section 4.1 (see also Section 8.3.4 and Appendix B.3.1). These probabilities are also claim-specific. Let them be denoted $z_{i:q+1}$, $z_{i:q+2}$,… .As is evident from Appendix B.3.1, they depend on the $t_{a:q+1}$, $t_{a:q+2}$,…from Step 1.

Step 4: For each of these claims, the **claim size**, given that finalisation occurs in quarters q+1, q+2, etc, was estimated from the relevant model. The precise forms of the models are discussed in the sub-sections relating to the different models below. These claim sizes are also claim-specific. Let them be denoted $s_{i:q+1}$, $s_{i:q+2}$,… .As is evident from Appendices B.1 to B.4, they also depend on the $t_{a:q+1}$, $t_{a:q+2}$,…from Step 1.

Step 5: Combine the results of Steps 3 to 5 to produce forecast **expected claim cost** in respect of the i-th claim for experience quarter $q+r$, r=1,2, etc as $C_{i:q+r} = f_{i:q+r} (1-z_{i:q+r}) s_{i:q+r}$.

Step 6: Collate the $C_{i:q+r}$ over relevant values of i, r. For example, the loss reserve for open claims in respect of origin quarter a is obtained by summation over all r and all claims i belonging to origin quarter a.

8.3.2 Time covariates only

A model of the type discussed in Section 3.3 has been fitted to the data set of strictly positive finalisations. It takes the form (6.1) with log link, i.e. $h(.) = \ln(.)$. The error term has a quasi-likelihood from the exponential dispersion family (EDF), characterised by

$$\mathrm{Var}[\varepsilon_i] = \varphi \{E[\varepsilon_i]\}^{2.2} \qquad (8.2)$$

The design X_i is restricted to time covariates, and full details of the model are given in Appendix B.1. It is seen to be economical, involving just 9 parameters, including:

- An intercept;
- 3 describing the operational time effect;
- 4 giving rates of claims inflation over different past periods (coefficients of covariates defined in terms of finalisation year);
- 1 describing variation in the rate of claims inflation as operational time varies.

Claims inflation is seen in Appendix B.1 to be estimated at a constant rate of 3.9% per annum over the 15 years 1990 to 2004 for operational times above 70%. For lesser operational times, the rate of inflation (strictly the inflationary coefficient in the linear predictor) increases linearly with decreasing operational times. These rates have been extrapolated into the future for the purpose of forecasting loss reserve.

As noted in Section 3.3, the inclusion of claims inflation in the model necessitates a forecast of finalisation rates. This is done on the basis of the model described in Appendix C.1. Indeed, both loss reserve estimates described in this and the following section are based on the same forecasts of future numbers of finalisations by origin period and development period.

Note that the forecast of finalisation rates has only a second order effect on estimated loss reserve. If claims inflation were zero, it would have no effect, as explained in Section 3.2. The uncertainty induced in the loss reserve by the schedule of finalisation rates is therefore of the order of the uncertainty in the average term to finalisation multiplied by the claims inflation rate.

This influences the choice of model for forecast of numbers of finalisations in Step 1 of the protocol set out in Section 8.3.1. Two models are available from Section 5, the "simple aggregate model" (Appendix C.1) and the "Cox regression model" (Appendix C.2).

In principle, the second of these, being the more comprehensive, is preferable for any loss reserve forecast that requires a forecast of finalisations. In fact, however, the simple aggregate model has been used for the models of the present sub-section in view of the relative insensitivity of the loss reserve estimate to this choice.

The finalisation probabilities for the open claims of an origin quarter, introduced in Step 3 of Section 8.3.1, are taken as proportional to the numbers of finalisations in Step 1. This is equivalent to assuming that the distribution of finalisations over future quarters is the same for open and IBNR claims respectively.

This introduces an error but, on the basis of the insensitivity of the loss reserve estimate to the forecast of finalisations reasoned above, it has been ignored. Some control testing of the loss reserve estimates showed that they did, in fact, change little when the Cox regression model replaced the simple aggregate model in Step 3.

Table 8.2 reports this loss reserve estimate in total and its prediction error. Appendix D.2 gives the corresponding information on origin years, and this is represented in Figure 8.2. These results are compared with their chain ladder counterparts, reproduced from Table 8.1.

In the case of the individual claim model, the predictive CoV is dissected into a number of component parts. The CoV is shown as it accumulates various contributions, denoted as follows:

- CoV1: Claim size parameter error only;
- CoV2: Claim size process error added in;
- CoV3: Finalisations process error added in;
- CoV4: Probability of zero finalisation process error added in.

The final CoV is taken as CoV4, whence it can be seen that certain, apparently minor, contributions to CoV (e.g. parameter error in finalisations model) have been omitted on the grounds of materiality.

Table 8.2
Loss reserve based on time covariates only

	Chain ladder and Mack	Individual claim model based on time covariates only
Estimated loss reserve	$888M	$1,000M
Predictive CoV	10.5%	5.3%
CoV1		4.0%
CoV2		4.1%
CoV3		4.3%
CoV4		5.3%

The individual claim model estimates a quite different (12%) loss reserve from the chain ladder, but is virtually twice as efficient in terms of prediction error. As mentioned in Section 8.1, this comparison is somewhat unfair to the chain ladder, but it seems unlikely that the latter method could be refined to an extent where its efficiency would challenge that of the individual claim model.

Figure 8.2
Loss reserve based on time covariates only

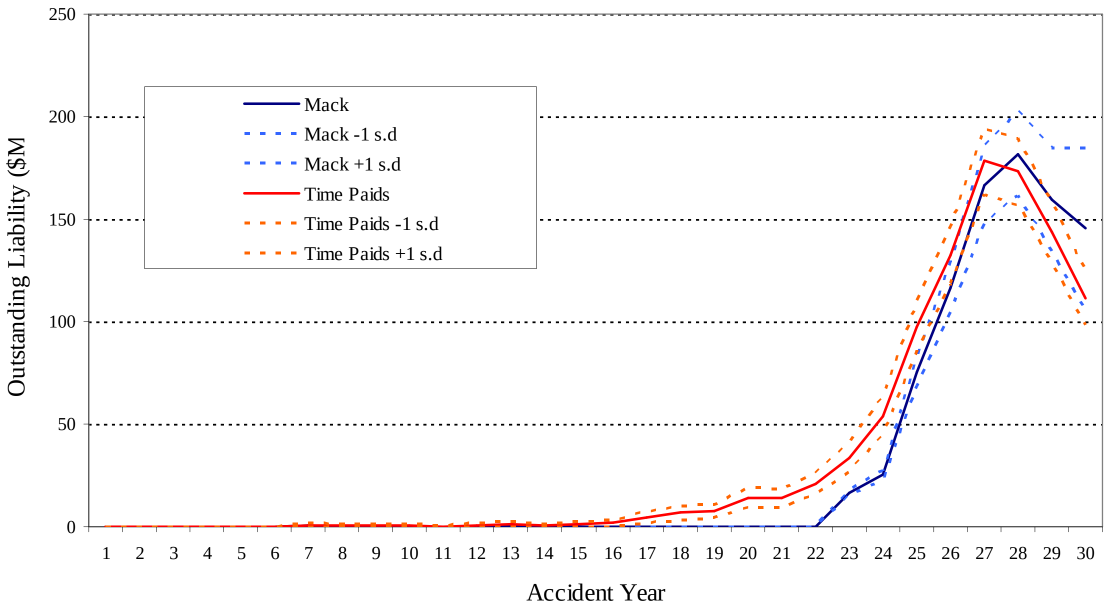

In Figure 8.2, "Mack" refers to the chain ladder and "Time Paids" to the individual claim model. It shows that the latter has a much narrower confidence envelope for the recent origin years. A plot of the individual claim model results on a log scale, displaying more detail for the early origin years, appears in Figure 8.3.

8.3.3 Static covariates included

The individual claim model of Section 8.3.2 has been extended to include several static covariates, generically described as follows:

- The risk class (6 of them) to which the policy generating the claim belongs;
- The geographic district (3 of them) in which the risk generating the claim is located;
- The type of entity covered (2 of them) by the policy generating the claim.

The model now has 20 parameters rather than the earlier 9. Claims inflation over recent years for higher operational times is estimated as 4.0%, rather than 3.9% per annum as formerly. Full detail appears in Appendix B.2. The simple model of finalisations set out in Appendix C.1 was retained here.

Table 8.3 reports this loss reserve estimate in total and its prediction error. Appendix D.3 gives the corresponding information on origin years, and this is represented in Figure 8.3. These results are compared with their counterparts from the individual claim model based on just time covariates, reproduced from Table 8.2.

Table 8.3
Loss reserve based on time and static covariates

	Individual claim model based on time covariates only	**Individual claim model based on time and static covariates**
Estimated loss reserve	$1,000M	$978M
Predictive CoV	5.3%	5.7%
CoV1	4.0%	4.3%
CoV2	4.1%	4.3%
CoV3	4.3%	4.7%
CoV4	5.3%	5.7%

The inclusion of the static covariates, despite their statistical significance, actually degrades the predictive efficiency of the model. Most of the loss of efficiency relates to parameter error, i.e. to the additional forecast variability arising from the additional parameters in the model.

This is not to deny the potential usefulness of the expanded model. It may well be the appropriate model for differentiating risks in the portfolio, such as in rate-making. It is simply not especially useful in a loss reserving context.

Figure 8.3
Loss reserve based on time and static covariates

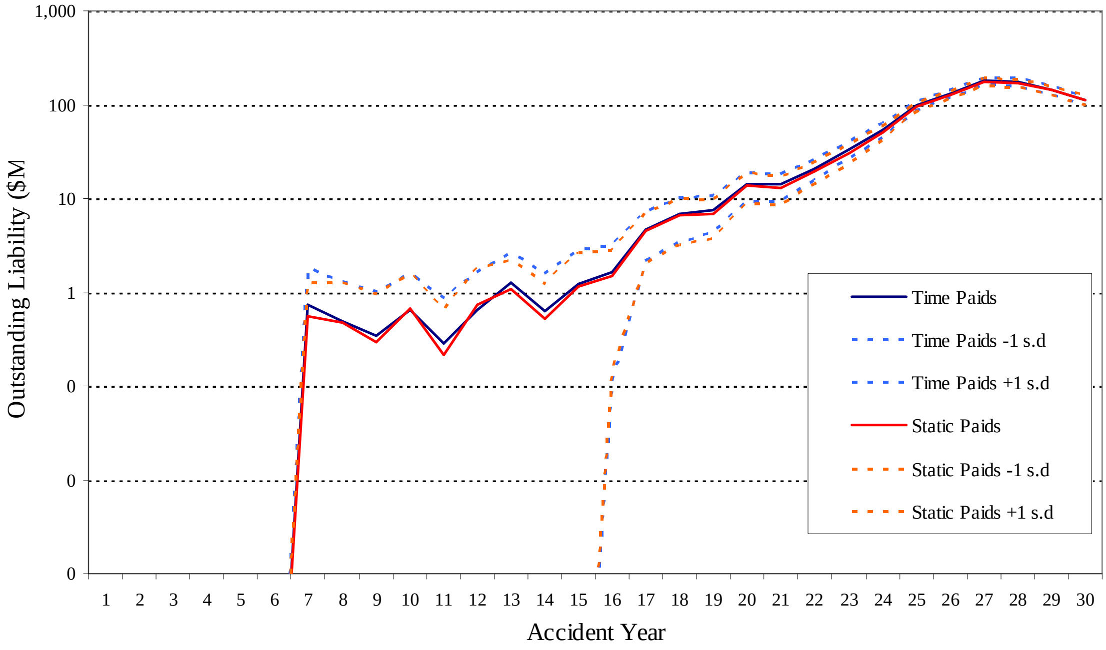

The vertical scale of Figure 8.3 is logarithmic in order to display detail of the earlier origin years. The results of the two models are seen to be quite similar.

8.3.4 Case estimates included in covariates

The general form of a model that includes case estimates of ultimate incurred loss in its covariates is discussed in Section 4.1. It is pointed out there that the model in fact consists of three sub-models, in addition to the additional models required to forecast future numbers of finalisations and the sizes of IBNR claims.

The three sub-models relate to:

- The probability that a claim settles for zero, given a current estimate of incurred loss (which can be zero or non-zero);
- The distribution of ultimate claim size given a zero current estimate of incurred loss;
- The distribution of incurred age-to-ultimate factor given a non-zero current estimate of incurred loss.

The last two of these could have been merged into a single model of ultimate incurred loss conditioned by current estimate. However, the separate sub-models were considered likely to lead to greater forecast efficiency since the range of ratios required to be forecast by the last of them (which accounts for the majority of claims) is considerably narrower than the corresponding range of claim sizes.

All three sub-models are GLMs. The first involves a binomial error with a logit link, while the other two have log links and EDF quasi-likelihoods of the form:

$$\mathrm{Var}[\varepsilon_i] = \varphi \{E[\varepsilon_i]\}^p \qquad (8.3)$$

with $p = 2.1$ and 1.85 respectively.

The details of the three sub-models are given in Appendices B.3.1 to B.3.3 respectively. The model of age-to-ultimate factors (B.3.3) is seen to involve many covariates and interactions, all of which are statistically significant and either decrease the Akaike Information Criterion or are approximately neutral with respect to it. The model has 72 parameters.

The modelling of future numbers of finalisations has been carried out in rough and ready manner in Sections 8.3.2 and 8.3.3. A more careful approach is required here, however, for the following reason.

According to the indications of the data, the model of age-to-ultimate factor recognises that this factor tends to increase with the delay in finalisation. Since large claims will tend to finalise later than small claims, it is necessary to make claim-specific forecasts of finalisation periods. The Cox regression model of Appendix C.2 does so, and has been used here.

Recall the form of data input to the third of the sub-models, discussed in Section 4.2. The response variable consists of observations on age-to-ultimate factors for closed claims. Applied to an open claim, the model produces a forecast of this factor for the claim concerned, and then the ultimate claim size of the claims is forecast as follows:

Forecast ultimate claim size = Current case estimate of incurred loss
X
Forecast age-to-ultimate factor.

The loss reserve in respect of open claims has been estimated as the aggregate of the individual ultimate costs, forecast in this manner, of all claims open at the valuation date.

The inclusion of finalisation period in model covariates is worthy of comment. While it is reasonable to interpret coefficients of finalisation period as representing claims inflation in the earlier models, this is not so obvious for the present one.

Age-to-ultimate factors already incorporate allowances for inflation, and so any trend in them over finalisation periods reflects change over time in the strength of case estimates relative to ultimate cost. Any such trend might reflect claims inflation in some way, but is more general than this.

The model of Appendix B.3.3 estimates that age-to-ultimate factors experienced a downward trend of about 0.8% per quarter over the most recent 7 or 8 years of finalisation. This has **not** been extrapolated in forecasts of ultimate claim sizes of open claims; a nil trend has been assumed for future finalisation periods.

Table 8.4 reports this loss reserve estimate in total and its prediction error. Appendix D.4 gives the corresponding information on origin years, and this is represented in Figure 8.4. These results are compared with their counterparts from the individual claim model based on just time covariates, reproduced from Table 8.2.

Table 8.4
Loss reserve based on model of case estimate development

	Individual claim model based on time covariates only	**Individual claim model based on case estimate development**
Estimated loss reserve	$1,000M	$1,062M
Predictive CoV	5.3%	6.0%
CoV1	4.0%	5.2%
CoV2	4.1%	5.2%
CoV3	4.3%	5.6%
CoV4	5.3%	6.0%

The comparison here is very interesting in a couple of respects. First, the difference between the two forecasts is roughly one standard deviation of either forecast, and so appears acceptable. The difference could be partially due to the chosen projection of past trends in to future such as discussed in:

- Section 8.3.2 in relation to claims inflation;
- The present sub-section in relation to age-to-ultimate factors.

Second, the inclusion of current case estimates in the predictors of ultimate claim size achieves no increase in prediction efficiency. At the level of individual origin periods, the standard errors of prediction are similar for both models, except in the early accident periods where the paids model predicts no outstanding liability, unlike the model based on case estimates. Overall, the model based on case estimates has somewhat lower predictive efficiency.

A further point to note is that, for the present model, CoV2 accounts for the prediction error of the "main" model, that of age-to-ultimate factors. The error contributed by the open cases carrying a zero case estimate has been omitted, as there are few such cases.

Figure 8.4 plots the loss reserves estimated by this model ("Incurreds") against year of origin, and compares them with the corresponding estimates from Section 8.3.2 ("Time paids").

Figure 8.4
Loss reserve based on model of case estimate development

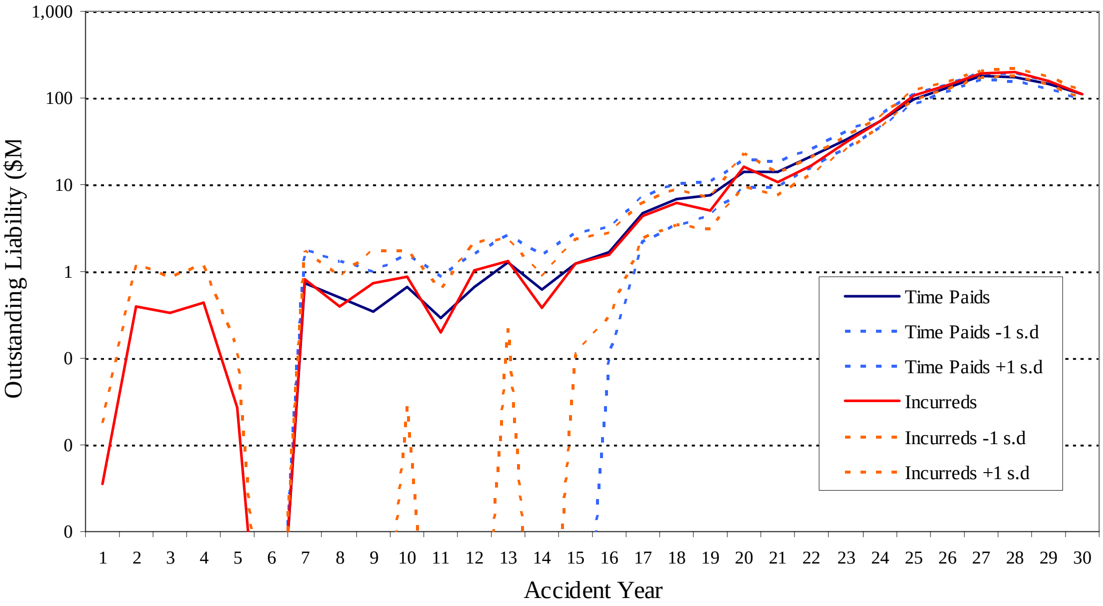

8.3.5 Combination of models of paid and incurred losses

Model blending

The procedure described in Section 7.1 has been applied to blend the models of Sections 8.3.2 and 8.3.4. After some experimentation, values of $\lambda_1=\lambda_2=10^{14}$ were chosen.

The predictive covariance *M* defined in Section 7.1 is estimated on the assumption that the results of the two models are stochastically independent. The matrix *M* then has block diagonal form. There are two blocks, representing the predictive covariance matrices associated with the two respective models. These matrices are obtained directly from the same bootstrap results as those used to produce the CoVs in Tables 8.2 and 8.4.

The weights calculated according to Section 7.1 are not constrained to non-negativity as such constraints would destroy the linearity of the problem set out there. However, where a negative weight is calculated, it has been set to zero here, and the weight for the opposing model set to 100%.

Figure 8.5 displays the resulting weights of the models of Sections 8.3.2 ("Time Paids") and 8.3.4 ("Incurreds"). The latter set commence at 100% for the earliest origin year.

Figure 8.5
Model blending weights

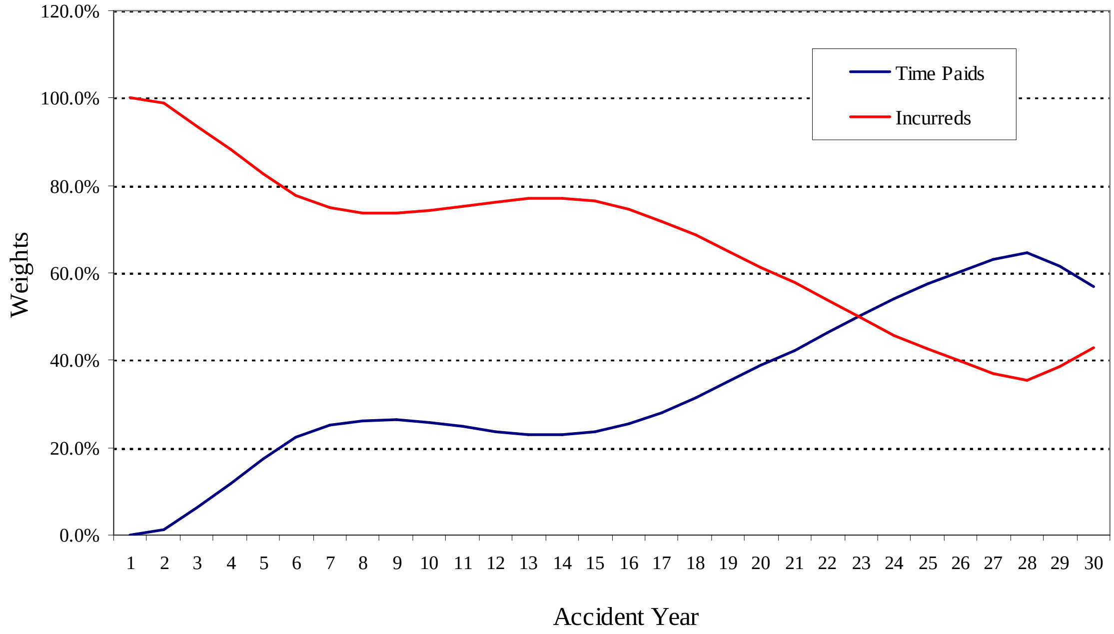

As one might expect, high weight is assigned for the early periods of origin to the model that relies on case estimates. Here, where there are relatively few claims remaining to be finalised, the average size associated with any origin year may differ from a typical value, in which case the case estimates will provide useful predictive information.

A little surprising, perhaps, is the result that the model based on case estimates continues to receive approximately 40% weight even for the most recent years of origin. One might reasonably expect the case estimates for these years to be of little predictive value, and it is indeed common to see models based on case estimates assigned low weight in respect of the most recent years. The results indicate that the data set analysed here is different in that its development of case estimates is less erratic in the early development periods than is typical.

Table 8.5 reports the blended loss reserve estimate in total and its prediction error. Appendix D.5 gives the corresponding information on origin years, and this is represented in Figure 8.6. These results are compared with their counterparts from the models of Sections 8.3.2 and 8.3.4, reproduced from Table 8.5.

Table 8.5
Loss reserve based on model of case estimate development

	Individual claim model based on time covariates only	Individual claim model based on case estimate development	Blended model
Estimated loss reserve	\$1,000M	\$1,062M	\$1,022M
Predictive CoV	5.3%	6.0%	4.0%

Figure 8.6
Loss reserve based on blending of models of paid and incurred losses

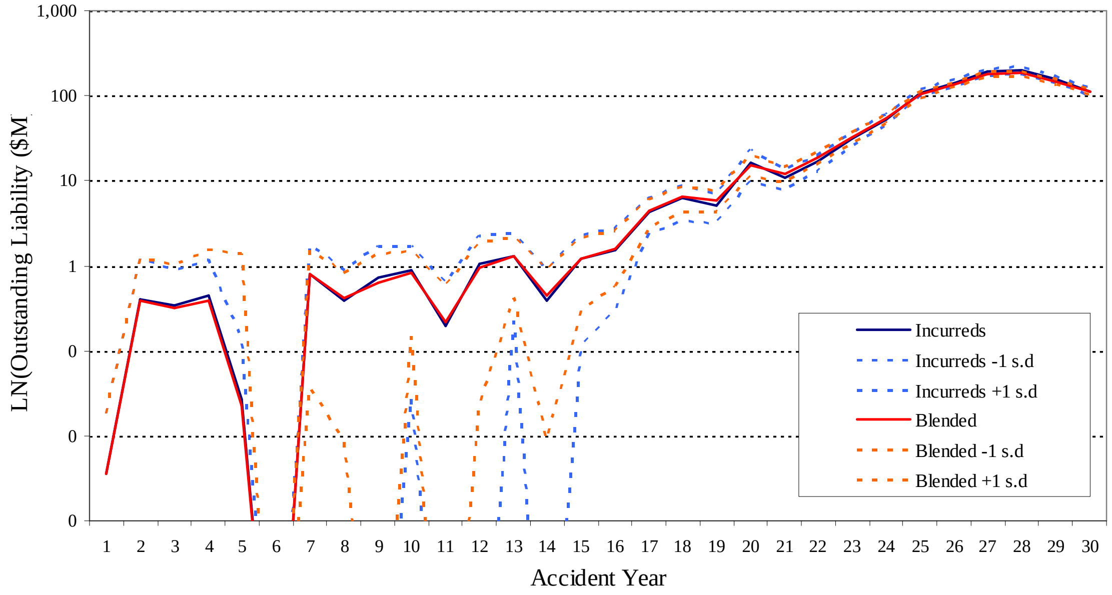

Figure 8.6 shows that the blended results are quite similar to those produced in Section 8.3.4 (similarly for those of Section 8.3.2). It also shows the narrower confidence envelope for the blended results, leading to the lower CoV that appears in Table 8.5.

Indeed, while the model based on case estimates did not, of itself, decrease the prediction error associated with estimated loss reserve, it does so when taken in conjunction with the model of Section 8.3.2.

Recall from Section 7.1 that one of the criteria accommodated in the loss function minimised by the blending process was smoothness over years of origin of the ratios of the blended to the case estimates of loss reserve (on open claims). Figure 8.7 plots these ratios against year of origin.

As is typical, one observes higher ratios for the more recent origin years. The grid lines in the plot are separated by intervals of 20%, so reasonable smoothness in the progression of ratios has been achieved. The break in the plot reflects one origin year for which there were no open claims at the valuation date.

Figure 8.7
Ratio of estimated liabilities to corresponding case estimates

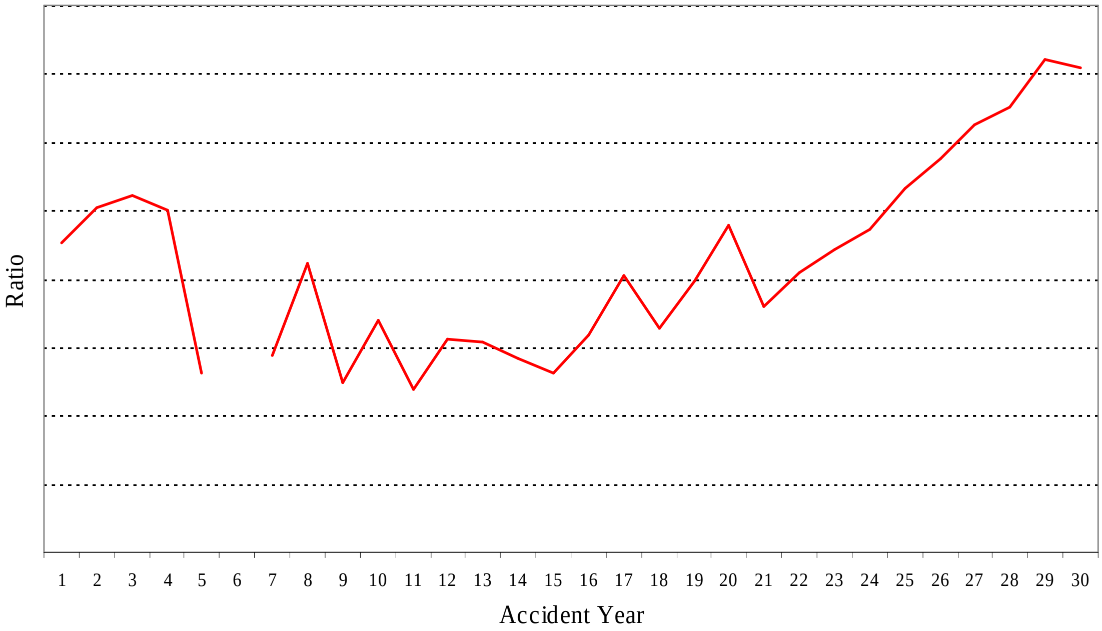

Unified model
Section 7.2 describes how a single model may be constructed, unifying (as opposed to blending) the models of Sections 8.3.2 and 8.3.4.

The error terms in those earlier sub-sections suggest a quasi-likelihood of the form (8.3) for the unified model, with p between 1.85 and 2.2. A convenient choice is p=2 (gamma distribution). Details of the model are set out in Appendix B.4.

As noted in Section 7.2, this model requires, in common with that of Section 8.3.4, the sub-models of Appendices B.3.1, B3.2 and C. The model of Appendix B.4 replaces that of B.3.3. The Cox regression model of claim finalisation in Appendix C.2 was used.

Although Section 7.2 remarks that the unified model is more general than a weighted average of the paids and incurreds models, it may nevertheless be useful sometimes to regard it in this way. This is particularly so in the selection of an allowance for future inflation for incorporation in forecasts.

Section 7.2 suggests that w may be interpreted as the weight assigned to the paids model in the unified one. Appendix B.4 indicates this weight to be roughly 20%, or equivalently an 80% weight on the incurred model.

Table 8.6 indicates that this interpretation is a reasonable one. Covariates are named here in the same way as in Appendices B.3.3 and B.4. It shows that, for a sample of influential covariates that occur in the incurreds but not the paids model, that their contribution to the unified model is not too different from 80% of that to the incurreds model. The comparison becomes erratic as one proceeds to covariates of more marginal influence.

Table 8.6
Comparison of linear predictors of incurreds and unified models

Covariate	Coefficient in linear predictor		
	Incurreds model	Unified model	Ratio: Unified/Incurred
			%
I(Fin qtr <60)	-0.042	-0.063	148
Log(max(ratio1-10,1))	0.053	0.054	101
Min(finqtr, 28)^2	0.001	0.000	44
Max(finqtr-27,0)	0.009	0.007	72
Max(8-dqdiff,0)	-0.039	-0.036	92
Min(dqdiff,50)	0.015	0.018	123
Max(10-optime*100, 0)	-0.098	-0.083	84

The paids model estimated claims inflation to have been 3.9% per annum over recent years (Section 8.3.2). Section 8.3.4, on the other hand, pointed out that temporal trends in the age-to-ultimate factors of the incurreds model generally represent something other than inflation. Age-to-ultimate factors were estimated to have exhibited a slight downward trend over recent years, but this was not extrapolated beyond the valuation date in forecasts.

Hence, the allowance for inflation over future years incorporated in the unified model was, in broad terms, 20% of that included in the paids model. More precisely, the contribution of finalisation period to the linear predictor:

- is scaled down by a factor of 20% at operational times above 70%, and
- its slope with respect to lower operational times is preserved from the paids model.

It is evident that this treatment of future inflation is heuristic, and it is identified as the major weakness of the unified model. Greater rigour would be a worthwhile objective of any future research.

Table 8.7 reports this loss reserve estimate in total and its prediction error. Appendix D.6 gives the corresponding information on origin years, and this is represented in Figure 8.8. These results are compared with their counterparts from the blended model, reproduced from Table 8.5.

Table 8.7
Loss reserve based on unified model of paid and incurred losses

	Individual claim model based on unified model	Individual claim model based on blended model
Estimated loss reserve	$1,091M	$1,022M
Predictive CoV	4.8%	4.0%
CoV1	3.7%	
CoV2	3.7%	
CoV3	4.1%	
CoV4	4.8%	

In view of the licence taken in the treatment of future inflation, the estimated loss reserve may be slightly suspect. The prediction error is similar to, though materially higher than, that based on the blended model. However, in creating the blended results, an assumption is made that the two models are independent; this assumption is not strictly true and may cause under-estimation of the prediction error.

Figure 8.8
Loss reserve based on unified model of paid and incurred losses

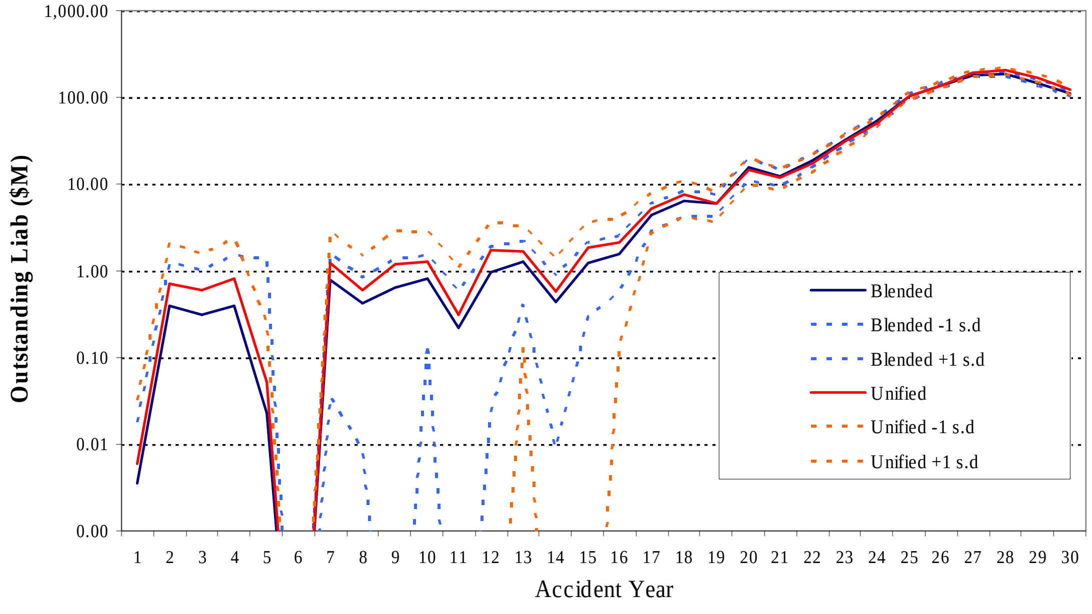

8.4 Summary

Table 8.8 summarises the estimates of loss reserve and predictive CoVs from Sections 8.3.

Table 8.8
Summary of loss reserves and estimated prediction errors

Model	Estimated loss reserve	Predictive CoV
	$M	%
Mack	888	10.5
Paids:		
Based on time covariates only	1,000	5.3
Including other static covariates	978	5.7
Incurreds	1,062	6.0
Blended	1,022	4.0
Unified	1,091	4.8

The bootstrap procedure for estimating prediction error (Section 6) yields not only the predictive CoVs reproduced in Section 8.3, but also estimates the entire distribution of the loss reserve forecast. Figure 8.9 plots this distribution for the unified model.

Figure 8.9
Estimated distribution of loss reserve forecast

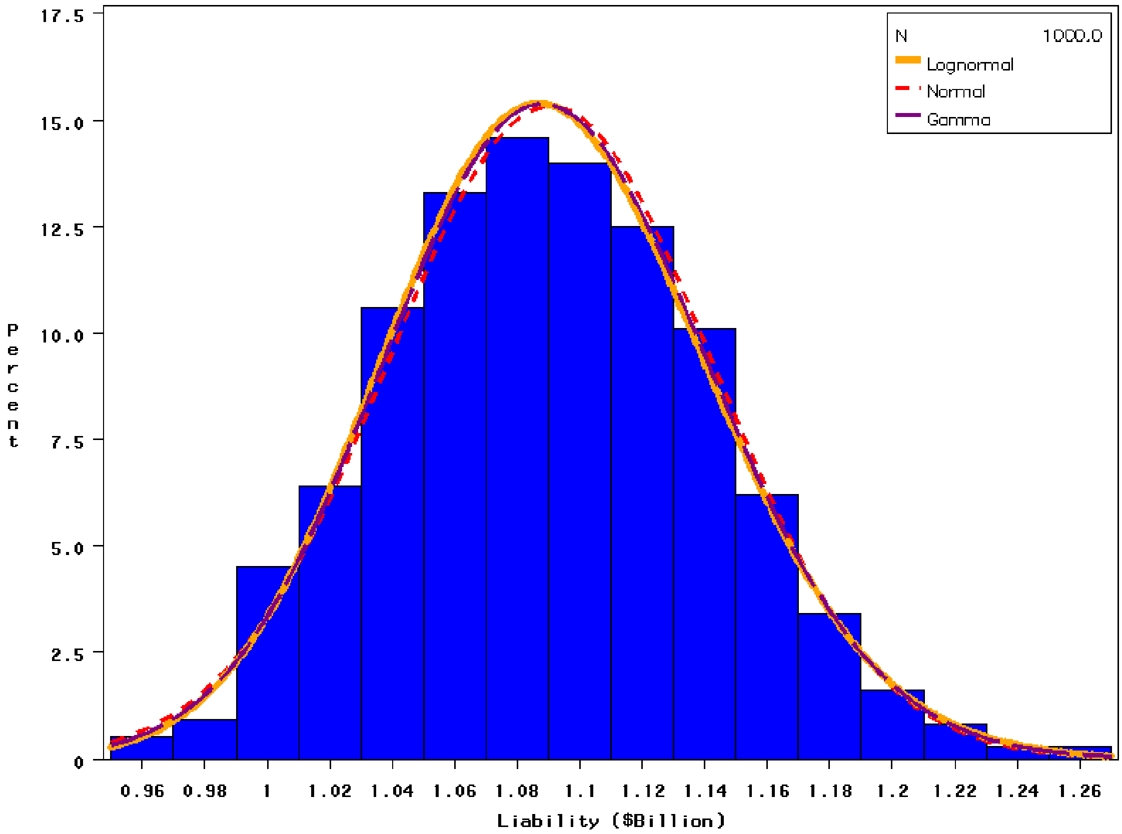

The bootstrap estimate of the distribution is based on 1,000 replications of the loss reserve forecast. The figure compares this empirical distribution with normal and log normal distributions having the same mean and variance. It is seen that:

- The skewness of the empirical distribution is low, despite its long tailed nature and the high skewness of the individual claim size distribution (see e.g. (8.2)), so that there is relatively little difference between the normal and log normal approximations to it.
- In fact all three distributions (normal – small dotted line, log normal – solid line and gamma – broken line) fitted to the curve are virtually indistinguishable. Even in the tail there are only small differences between each curve.

9. Conclusions

This paper examines various forms of individual claim model for the purpose of loss reserving, with emphasis on predictive efficiency in the sense of prediction error in the reserve.

Models of this sort are capable of higher predictive efficiency than aggregate models, particularly those aggregate models that do not follow a strict statistical procedure in their calibration (see the example in Section 8.3.1).

It is often possible to achieve high efficiency with a model of the "paids" type that has a small number of parameters (see the example in Section 8.3.2), certainly fewer than in most conventional actuarial models. The issues involved in the construction of such a model are relatively simple. Further improvements are possible, but possibly with considerable effort (Sections 8.3.3 and 8.3.4).

It is likely that any reasonable "paids" models will include time covariates, as defined in Section 2. Such models may or may not include other static covariates (as defined in the same section). It is interesting that, for the data set examined here, their inclusion, while yielding possibly interesting and useful information, degraded rather than improved predictive efficiency.

The construction of an individual claim "incurreds" model involves rather more modelling complexity. For the particular data set considered here, this did not yield any direct improvement in predictive efficiency (Section 8.3.4).

However, a different conclusion would be obtained for other data sets. If, for example, claim payment experience were highly erratic (unpredictable) but case estimates developed regularly (were predictable), it seems obvious that the loss reserve forecast conditioned by case estimate information would out-perform that which discarded this information.

Even though, in relation to the present data set, while the loss reserve forecast conditioned by case estimates did not, of itself, improve predictive efficiency, it did provide an alternative model that was largely stochastically independent of the paids model. The two models could then be used to produce a blended estimate of greater efficiency than either one (Section 8.3.5).

The blended model is a certain type of weighted average of its component models. Section 8.3.5 generalises this to a genuine unification of the paids and incurreds models in which forecasts of ultimate individual claim sizes are conditioned by their current case estimates in addition to other covariates.

The unified model, while more general than a weighted average, retains certain features related to it. It produces high predictive efficiency, in relation to the present data set at least.

The current weakness of this model is its lack of rigour in the interpretation of claims inflation, and the consequent lack of guidance as to the future values of this parameter. This area would merit further research.

Section 6 identifies the components of prediction error. It is to be emphasised that, as stated in that section, model specification error is not considered in this paper. This is not to minimise its significance. It is likely to be a substantial addition to the prediction errors quantified here, but needs to be addressed by different means, outside the scope of this paper.

10. Acknowledgment

Our thanks are due to our colleague Dr Richard Brookes for a number of valuable insights.

References

Brookes R and Prevett M (2004). Statistical case estimation modelling – an overview of the NSW WorkCover model. **Xth Accident Compensation Seminar**, Institute of Actuaries of Australia.

Bühlmann H (1970). **Mathematical methods in risk theory**. Springer-Verlag, Berlin.

Cox D R (1972). Regression models and life tables. **Journal of the Royal Statistical Society, B**, 34, 187-220.

Efron B and Tishirani R J (1993). **An introduction to the bootstrap.** Chapman and Hall, London

England P D and Verrall R J (2002). Stochastic claims reserving in general insurance. British Actuarial Journal 8: 443-544

Feller W (1971).. **An introduction to probability theory and its applications, Vol II**. John Wiley & Sons, Inc, New York.

Haastrup S and Arjas E (1996). Claims reserving in continuous time – a nonparametric Bayesian approach. **ASTIN Bulletin**, 26(2), 139-164.

Hachemeister C A (1980). A stochastic model for loss reserving. **Transactions of the 21st international Congress of Actuaries**, 185-194.

Jewell W S (1989). Predicting IBNR events and delays I: continuous time. **ASTIN Bulletin**, 19,25-55.

Jewell W S (1990). Predicting IBNR events and delays II: discrete time. **ASTIN Bulletin**, 20,93-111.

Larsen C R (2006). Paper included (though not presented) at the 2006 Stochastic Modeling symposium and investment seminar, Canadian Institute of Actuaries.

Lee E T (1992). **Statistical methods for survival data analysis**. John Wiley & Sons, Inc, New York.

Liang K Y and Zeger S L (1986). Longitudinal data analysis using generalized linear models. **Biometrika**, 73, 13-22.

Mack T (1993). Distribution-free calculation of the standard error of chain ladder reserve estimates. **ASTIN Bulletin**, 23, 213-221.

Norberg, R (1986). A contribution to modelling of IBNR claims. **Scandinavian Actuarial Journal**, 155-203.

Norberg (1993) Prediction of outstanding liabilities in non-life insurance. **Astin Bulletin**, 23(1): 95-115

O'Dowd C, Smith A and Hardy P (2005). A framework for estimating incertainty in insurance claims cost. **XVth General Insurance Seminar**, Institute of Actuaries of Australia.

Quarg G and Mack T. (2004). Munich chain ladder: a reserving method that reduces the gap between IBNR projections based on paid losses and IBNR projections based on incurred losses. Munich Re Group.

Reid D H (1978). Claim reserves in general insurance. **Journal of the Institute of Actuaries**, 105,211-296.

Taylor (1985) Combination of estimates of outstanding claims in non-life insurance. **Insurance: Mathematics and Economics**, 4(1985), pp. 321-438. Reprinted in **Actuarial Research Clearing House**, 185.2 issue, 335-371.

Taylor G (2000). **Loss reserving: an actuarial perspective**. Kluwer Academic Publishers.

Taylor G and Campbell M (2002). Statistical case estimation. Research Paper No 104 of the Centre for Actuarial Studies, University of Melbourne, at http://www.economics.unimelb.edu.au/actwww/html/no104.pdf.

Taylor G and McGuire H (2004). Loss reserving with GLMs: a case study. **Casualty Actuarial Society 2004 Discussion Paper Program,** 327-392. Paper presented to the CAS Spring 2004 Meeting, Colorado Springs, CO, May 16-19 2004. Also appears as Research Paper No 113 (2004) of the Centre for Actuarial Studies, University of Melbourne, at http://www.economics.unimelb.edu.au/actwww/wps2004/No113.pdf.

Zeger and Liang (1986) Longitudinal data analysis for discrete and continuous outcomes. **Biometrics,** 42, 121-130.