

**PERFORMANCE TESTING AGGREGATE AND
STRUCTURAL RESERVING METHODS:
A SIMULATION APPROACH**

By John Rollins, ACAS, MAAA

**Submitted in Response to the Call for Papers for the 1997
Casualty Loss Reserve Seminar
November 2, 1997**

ABSTRACT

Aggregate financial data histories are extensively used by actuaries in projecting ultimate liabilities, but the claim occurrence, reporting, and settlement process which generates these data is not perfectly understood, rarely modeled directly, and not incorporated into the structure of most popular reserving methods. This paper utilizes today's computer applications to create an automated simulation tool. This module allows the user to choose statistical assumptions for each element of the claims process and specify the structure of the simulation experiment. It then generates random paid loss and claim count histories based on these inputs.

The module is used to perform an experimental test of the performance of aggregate versus structural reserving methods over time. The methods chosen are the paid loss development method and a new "closed claim cost" structural method presented in the paper. In each trial, a database of ten accident years at ten annual evaluations is simulated. Then both methods are run at five successive calendar year evaluations of the simulated data. Several error functions are tabulated at each valuation date and the speed of the approach of each method's indication to the true value of the aggregate costs is examined. Suggestions for enhancement of the simulation model and performance tests follow.

INTRODUCTION

Here is one way to think about the problem of loss reserving and the reason we actuaries are highly qualified to do it. As actuaries, we are expected to predict financial outcomes which are results (or values) of a relatively well-defined dynamic process or “objective function”, a term used by analysts in some other fields for the mathematical description of the process whose result they are trying to predict. This is a function of a number of inputs (or variables), whose values are in turn determined by externally imposed characteristics of the stochastic process (or parameters). Under this technical description, it is easy to see that the job in itself does not distinguish us from many other professionals, notably statisticians, economists, and engineers. What does set us apart is that, in addition to having the imagination, modeling abilities and the general set of quantitative skills required to perform any estimation job, we are presumed to have expert knowledge of the insurance function which generates financial outcomes for risk-bearing entities and the variables on which it depends, as well as extraordinary prior judgment about the parameters of the process. As a result, we spend much of our time trying to understand which issues drive the occurrence, reporting, development, and settlement of insurance claims and the magnitude and direction of their effects on ultimate costs.

One of the most powerful tools for gaining this understanding of the objective function and improving our parameter judgments is historical empirical data. The trouble is, data collection is typically financial in nature and focuses on the effects of the process - aggregate costs - rather than the causes, namely exposure to loss, accident frequency, delays in reporting claims, inflation of liabilities for open claims, and delays in claim settlement. Therefore, this branch of actuarial science has evolved as such: given prior values of the objective function (historical cost outcomes), we must (explicitly or implicitly) identify the form of the underlying objective function and the key variables on which it depends, and estimate the values of the parameters governing each variable. Predictive success for a given financial or insurance product is usually measured by the proximity of our predicted values to those which do emerge, and by how quickly our predicted values improve when we get additional information and update our parameter estimates.

A manifestation of this evolution is that most of our canonical reserving methods depend heavily on availability, validity and consistency of regularly collected historical data. This dependence necessitates the assumption that the “rules of the game” (the objective function, its variables, and their parameter values) have not changed substantially over our experience period. We are all familiar with the forces operating on the players in the game and the low likelihood that this is the case over a long period of time. Both operational and financial managers tend to make big changes based on short-term considerations which overshadow their desire for integrity of the actuarial database. When we are confronted with the knowledge of interim changes, we usually turn to a collection of piecemeal, though sometimes ingenious, methods for adjusting the canonical model, rather than looking for an alternative approach.

The problem of making projections from aggregate data which do not reflect a constant process is accompanied by the dual problem of testing loss reserves. The same databases which distort reserve projection methods also wreak havoc on traditional runoff tests - just ask anyone who tries to make a firm statement based on results from Schedule P of the statutory Annual Statement. It is difficult from these retrospective tests on historical data to solidly determine which reserving methods actually are less biased or volatile over time than others.

Another manifestation is that most popular reserving methods are designed to input aggregate financial data rather than claim frequency, reporting, settlement, and severity information. This is certainly understandable when no such data is collected responsibly, but we should be designing methods that will work with comprehensive data sets and improve over the performance of the old methods, anticipating the glorious day when technological and management advances will allow us to capture all the information we need.

These observations about loss reserving problems are hardly new, but our professional literature has historically been long on identifying theoretical issues and short on empirical solutions. In this paper, I try to make a contribution to addressing the issues by creating and doing a simple experiment. In particular, I design a flexible algorithm and a software tool for simulating the structure of the claims process, set up an aggregate and a new structural reserving method to make projections of ultimate losses, run the simulation algorithm under a basic scenario, and test the accuracy and volatility of the projections of each method as the experience period runs off over time. I also ignore several phenomena which certainly enter the structure of the process in the real world, but hopefully leave the architecture of the experiment open enough so that we can flesh it out in a modular fashion as research progresses.

OBJECTIVES

The objectives of this paper are several.

Fill Gaps in Historical Literature

Most of the reserving literature concerns projection methods which rely exclusively on historical data. A notable exception is Stanard's 1985 paper on simulation testing [1], to which this experiment may be compared. His paper is an excellent starting point for designing simulation techniques and choosing performance criteria. This paper is somewhat dissimilar in that Stanard used a different algorithm for generating his database, did not consider structural reserving methods in his experiment, and did not track the bias of the selected methods as the accident periods matured over calendar time.

A moderate amount has been written on testing the performance of loss projections. Khury [2] emphasized the idea of a time-dependent radius of confidence about point estimates of reserves and considered the form of the decay of that radius function $R(t)$ to zero as reserves are run off. He also focused on structural knowledge of the "pressure points" underlying the current process as the key to the prospective reserve computations, rather than historical knowledge of the adequacy of past reserves. Unfortunately, when new reserving methods are proposed in our literature, regardless of the data sources, the authors rarely undertake performance tests which track the accuracy of the method's projections or the radius of the confidence interval over several ensuing calendar years.

Tool Development

With today's technology, we can create executable computer applications to run even sophisticated or computationally cumbersome simulations, and generate output which is compatible with commercial spreadsheet or database systems, allowing further analysis or presentation of results. I chose to write a set of macro modules for Microsoft Excel in the Visual Basic framework. The software platforms available today also encourage modularity of code and therefore easy addition of enhancements to the model in the future.

Model Building and Refinement

An advantage of simulation techniques is that we can specify an algorithm which models the claims process, identify the variables (drivers of the claims process) in the algorithm, and make statistical assumptions about them individually, whether or not those assumptions facilitate easy computations or closed-form aggregate distributions. Running the simulation will still give us valuable information about ultimate costs and, perhaps more importantly, the sensitivity of those costs to changes in the type of variables in the model and the parameters which govern how they are generated. For instance, if we learn the claim department is three months behind in logging new reports, we can estimate the effect of the operational change directly and immediately rather than mumbling something about the "robustness of the aggregate loss distribution" to management or the client when our model is not flexible enough.

All of this would be no good if we had to reconstruct the model every time a major strategic or operational change occurred. Fortunately, simulation models are built like a house; namely, the foundation and walls are completed before the trim is painted and the azaleas are planted out front. We add enhancements to the algorithm as we learn more about the process we are modeling. In fact, the model described in this paper is intentionally kept very simple, and I will call upon the creativity and keen senses of the rest of the profession to add details. Even supposing the simulation algorithm is adequate, the parameters of each of the variables will require calibration in order to customize the model to any particular company. This is the work of keeping up with a constantly changing business environment, but I think it beats adjusting historical data.

Better Structural Reserving Methods

Structural methods which rely on claim reporting, claim closure, and other information about the loss process have historically been underdeveloped in the literature, probably because they require more data and computations to apply. The results of my experiment may prove me wrong, but my *a priori* belief is that estimates produced by structural methods have the potential to be more accurate predictors of ultimate losses, especially at immature ages, than those from aggregate methods. Another hypothesis of mine is that structural methods will respond with less volatility and bias to specific changes in the operating environment, like slowdowns in claim reporting, improvements in processing in the claim department, or commutation efforts on claims which would otherwise generate periodic partial payments. Without the albatross of insufficient or unsuitable historical data dangling from our necks, development of better structural reserving methods will provide a different and possibly more lucid actuarial perspective in loss reserving.

I have set up a simple structural method which projects ultimate claim counts from reported count data, projects closed claims incrementally from closed count data, projects average severity from paid loss data, and computes the required reserve for unsettled claims. It does not duplicate any algorithm of which I know, though it borrows some mechanics from the closure model for bodily injury claims in Adler and Kline's 1988 paper [3]. The method is intended to be a reasonable archetype, not universal or suited to all areas of actuarial practice. By way of comparison, it is applied to the simulated databases in parallel with an aggregate paid loss development method and the performance of the two methods is compared over time.

Time-Dependent Performance Standards

Khury elucidated an important concept when he graphed the radius function of a reserving method against time. Since this paper is not predicated on an aggregate statistical model, confidence radius functions are not calculated, but error functions are. A good performance test should consider the speed and consistency of the approach to ultimate of projected losses at successive valuation dates. Unfortunately, both actuaries and lay users of financial statements tend to "turn it around" and think in terms of the favorable or adverse development from the original estimate over time, not the development toward the target answer.

Moving Beyond Historical Data

This theme pervades the entire discussion, so I won't belabor it here. It should be apparent that the more intellectual capital we as a profession invest in models and technologies which free us from the rigid historical data triangle, the more flexible and valuable our skills will be. Economists, for example, have built many useful models predicting the effects of policy changes on aggregate costs (or profits) which are not as dependent on either sheer volume or consistency of historical data.

THE SIMULATION MODEL

The Environment

The simulation algorithm, its inputs and outputs, the reserving methods used in the experiment, and the test results are all stored in Microsoft Excel spreadsheets. The simulation model and its interface with the analyst are written as macro modules and dialog boxes in Visual Basic for Applications. This platform is nice because it allows seamless transitions among an aesthetically pleasing user interface, a structured but flexible programming environment, and a powerful number-crunching and exhibit-making tool, all three of which are desired to accomplish the (perhaps overly ambitious!) objectives set out for this paper.

How it Works

The algorithm is intended simply to approximate the reality of the claims process rather than conform to any nice closed-form solution or coordinate with any prior actuarial literature. It generates a claims and actuarial (triangular format) database in a few simple steps. While the following quick description of the algorithm will gloss over several assumptions about the variables and their relationships, we will discuss them in detail later.

We should first structure the output database by choosing values for the number of accident periods to simulate, the length of each evaluation interval (e.g. annual, quarterly, etc.), the last evaluation age (where the tail starts), and the maximum number of claims in a period. These inputs tell the modeling software the size of the output matrix. The input box used in my application is shown in [Exhibit 1, page 1](#). I will speak of data matrices rather than triangles - matrices only become triangles when we analyze them at a particular calendar point in time, before the values of all the cells in the matrix are known. We are simulating the database, so we know the answers before we make and test estimates. The maximum number of claims is necessitated by computing rather than theoretical considerations.

Parameter values for the statistical distributions used to generate the random variables in the model must be chosen. The change in exposure volume (from one accident period to the next), claim frequency (per unit of exposure) for each accident period, report lag for each claim, settlement lag for each claim, and the ultimate severity (in today's dollars) for each claim are all drawn randomly from sampling distributions. Parameter values specify the scale and shape of those distributions. Finally, the rate of inflation applicable to claim payments and the base exposure volume (a scale factor) must be specified. The inflation rate is nonstochastic in the model as constructed. The input box for the statistical assumptions is shown on [Page 2 of Exhibit 1](#).

The algorithm then goes to work. Using the selected structural and statistical parameters, we iteratively compute the exposure amount for each accident period as last year's volume multiplied by a percent change which is randomly drawn from a normal distribution. (For the first accident period, the base exposure volume is used.) We draw each period's claim frequency (independently from the exposure volume and prior frequency values) from a normal distribution. The ultimate number of claims for the period is the product of the exposure level and the frequency per unit.

For each claim indicated within each accident period, we draw its accident date from a uniform distribution over the accident period, its report lag from a gamma distribution, its settlement lag from a gamma distribution, and its ultimate severity (in today's dollars) from a gamma distribution. These random variables are all drawn independently and contemporaneously. Along with the predetermined parameters above such as the inflation rate, they give us enough information to allocate claim reports, claim counts, and loss payments to the appropriate cells in the simulated data matrices. (The payout amount for each claim is allocated entirely to the interval of settlement.) Computers do this kind of thing very efficiently, so we let them do all the sorting, data-handling, and arithmetic operations necessary to leave us with a finished actuarial (and, as a by-product, detailed claims) database. This data is ready to use as input to any reserving model, excepting those which use outstanding case reserves. Methods that do will not be considered in this paper.

Structural Assumptions

The model is flexible in terms of how the final loss matrices and triangles will look. A typical set of output matrices is shown in [Exhibit 3, page 2](#). A maximum number of claims is specified because theoretically it is desirable to draw the claim frequency from a distribution which is not truncated from above, such as the normal or lognormal, but in practice the spreadsheet which is used as the backbone of the application would run out of space if the number of simulated claims in one trial exceeds about 8,000. Since the number of claims depends directly on the exposure and frequency parameter assumptions, the effect of this limit can be made insignificant if the mean and standard deviation of the claim frequency and the exposure growth variables, as well as the base exposure amount, are chosen accordingly.

The other limitation concerns the tail. The algorithm produces matrices for reported counts, closed counts, and loss payments which are of a finite number of columns. Like the claim frequency, the report and settlement lag for each claim are also drawn from distributions with no upper limit. Rather than artificially truncating the distributions, it is theoretically purer to handle the upper bound by specifying a final evaluation of the data and let any payments, reports or settlements made at more advanced ages be allocated to a tail interval.

Outside of limitations of these and other structural parameters to finite values, the accident periods and evaluation intervals can be any length. The software measures all lengths in months, so six-month policy periods, quarterly or monthly data collection, and other common situations can be handled effectively. Note that the model assumes an exposure base measured in real terms - an "inflated exposure" column showing nominal exposure values inflated at a constant 5% per annum is shown on some exhibits, to satisfy the analyst who is modeling Workers Compensation or some other line with an inflation-sensitive exposure base.

Statistical Assumptions

The outcomes of the simulation model depend on some stochastic assumptions about the variables. I hypothesize that most of the observed results and conclusions drawn from the upcoming simulation experiment will be highly robust to reasonable changes in the statistical assumptions. In fact, given the time and inclination, sensitivity tests on this model are easy to run. Yet we all know actuaries love to quibble about mathematical assumptions, so I will discuss them here. For a quick summary of the assumptions used in the scenario tested in this paper, turn to [Exhibit 2](#).

Exposure Changes

Assuming some base exposure amount, the real change in exposure from period to period should depend on operational and strategic considerations, as well as the economic constraints facing the client company. Ignoring underwriting cycles and macroeconomic fluctuations for a moment, the periodic change can be approximated by some kind of random walk with a drift. The classic assumption for random walk time series models is that unsystematic changes are normally distributed. The variability of the random walk is represented by the standard deviation parameter, and the drift represented by the mean (a zero mean implies no drift). The drift itself is a proxy for the long-term underwriting growth rate for an insurance company, or the business growth rate for a self-insured client. Incorporating the effects of national business cycles, as well as industry-specific cycles depending on the model's client company, into the exposure time series is an obvious enhancement since so much work has already been done by economists on modeling the cyclical nature of growth.

Note that a powerful practical use of this model for decision-making might be testing the sensitivity of ultimate costs or their present value to the growth strategy of the client company. Varying the mean growth rate, the variability of exposure growth, or the length of underwriting cycles could provide valuable insights to management.

Claim Frequency

Claim frequency may also be subject to some cycles, and may not be completely independent of exposure changes. For instance, a fast-growing manufacturing operation is probably associated with adverse selection in the labor markets, and we might expect the Workers Compensation claim frequency to be positively correlated with the exposure growth rate. Lacking strong studies of such phenomena in the general case, we assume that claim frequency from period to period is independently and normally distributed, with the mean and SD chosen appropriately.

Lags

The assumptions for individual claims bring us closer to the core competencies of our profession. Though experience in annual periods is not always uniform, particularly for property claims which tend to be seasonal, in the general model we assume that occurrences are uniformly distributed over each period. The distributions of report lags and settlement lags have received some attention in actuarial literature, and some of my intuition about these variables has probably been influenced by Weissner's paper on truncated maximum likelihood estimation [4]. In particular, a report lag represents a phenomenon similar to the waiting time for a Poisson process and should be subject to some of the same considerations. The three key assumptions underlying the Poisson process are detailed in Hogg and Klugman [5] and most college texts on mathematical statistics. Engineers model the number of defective pieces of equipment produced in a day and the expected time to failure of a piece of equipment with such process models. The gamma distribution is the theoretically correct model of the waiting time for the N th occurrence of a "failure" or Poisson event, where the failure rate is constant over time and where N is a parameter of the distribution (usually denoted α). While supporting a theoretical model of report lags as failures in the Poisson process is beyond the scope of this paper, it seems reasonable to use the gamma distribution to model the variable. The settlement lag is also essentially a waiting time variable, and it appears just as reasonable to model it with the same statistical assumptions.

Claim Severity

Modeling claim severity distributions is arguably one of the most important areas of research for the actuarial profession. Again, this paper is not intended to add anything to the body of knowledge of loss size distributions or even survey the voluminous existing literature on the subject. Nevertheless, an assumption must be chosen. For the general case presented here, we select the gamma distribution as a flexible proxy for many moderately skewed distributions.

We assume that report lag, settlement lag, and ultimate severity are mutually independent. In reality, it is almost certain that they are pairwise correlated. In fact, many actuarial techniques, particularly those embodied in structural reserving methods, depend on the assumption that larger claims are reported later and stay open longer. In order to reflect this in our model, we would have to specify a function (not necessarily in analytic or closed form) that relates them: $F(\text{report lag, settlement lag, severity}) = 0$. One way of doing this is to select the parameters of the loss severity distribution for a given claim as functions of the previously drawn report and settlement lags for the same claim, rather than determine the severity and the report and settlement lags contemporaneously. The scale and shape of the severity draw would be altered to reflect the type of claim, but the draw would still be random rather than deterministic. In my opinion, adding this aspect to the model might significantly improve the approximation of the claims process and change the test results. It is a challenge which remains for further research.

Inflation

The inflation assumptions in the general model are nonstochastic for simplicity's sake. Economic literature is replete with time series models of inflation, and it would be a relatively easy enhancement to incorporate a cyclical, random walk, or other time series model of a calendar year price index into the simulation routine.

There is also the question of how inflation affects loss payments. Butsic [6] constructed a theoretical model in order to contrast a paid loss inflation effect based on the date of accident with an effect based on the date of claim settlement. We assume that the price index at the date of claim settlement determines the nominal value of the paid losses associated with the claim (or, in Butsic's notation, alpha equals one). Therefore, nominal claim severity is positively correlated with the length of the report and settlement lags in this model.

Other Comments on Statistical Assumptions

The assumptions presented here have been justified as reasonable for an experimental case, but not rigorously supported. In a situation where historical data allows empirical examination of any of the above variables, the assumptions can be tailored to the client company by fitting a variety of distributions to the available data, selecting a model and parameters based on some "best fit" criteria, and modifying the simulation model to take account of the superior information. The simulation can also be improved if anticipated changes to the operating environment are reflected in the statistical assumptions. For example, if management has determined that a more aggressive legal department should be installed to handle an inventory of troublesome liability claims, the settlement lag assumptions and standard deviation of severity can be changed accordingly.

Most good statistical models follow the principle of parsimony - that is, the inclusion of random variables which do not add substantial information to the model and the use of distributions with more parameters than necessary to capture the shape of the probability function should be avoided. With such a complex insurance process to model, we should be sure not to include superfluous elements in the simulation algorithm, but careful to incorporate all of the key cost drivers. I have tried to avoid unnecessary complexity by advocating the use of simple two-parameter distributions for modeling the key variables and by constructing an algorithm based on the way I think the loss generation process works in practice and the major elements influencing the said process. It is certainly possible that the success of any experiment using this approach may be improved by adding some variables which I did not consider, by modifying or simplifying some of the statistical assumptions, or by removing a variable to which the results are not sensitive, but that simply adds complexity to the model. Testing for overspecification in algorithms such as standard least-squares estimation (i.e. collinearity and overidentification of instrumental variables) has been somewhat codified, but in the model described in this paper, actuarial judgment will have to prevail for now.

THE EXPERIMENT

Description

The design of the experiment is as follows:

Step 1: Create Scenario

Choose distributions and their parameters for each variable. The variables in the general case explained above are exposure growth, claim frequency, report lag, settlement lag, ultimate (present value) severity, and the claim cost inflation rate. While I have described a “general case” algorithm for the experiment presented here, only the analyst’s imagination limits what can be simulated if he/she is willing to write some code. I will do a little imagining later in the paper. Keep in mind that the number of claims in the simulated database will depend on the choices for exposure volume and claim frequency. The distributions and parameter values for the “base” scenario are shown on [Exhibit 2](#).

Step 2: Generate Data

Run the simulation routine to generate a claim database and data matrices of substantial size. In this experiment, we use 12-month accident periods, 12-month valuation intervals, ten accident periods of exposure, and ten evaluation intervals with the “tail” attaching at 120 months. A sample output claim database and actuarial data matrices are shown on [Exhibit 3, pages 1 and 2](#).

Step 3: Make Projections

Apply the aggregate and the structural reserving methods to the simulated data to project ultimate losses, at a valuation date such that the latest accident period is 12 months old. With ten annual accident periods, this would be at the end of "calendar year 10". Advance the valuation date 12 months and re-apply each reserving method. Continue to advance the valuation date 12 months at a time and re-apply the methods for a total of five iterations (in other words, at five successive calendar year-ends). The descriptions of the archetypical aggregate and structural reserving methods used in this experiment follow in the next section of the paper. In order to make the experiment consistent and repeatable, follow the general rule of using the weighted average from all of the available data to approximate values which would normally be selected by actuarial judgment (e.g. loss development factors in the aggregate method, and claim closure rates in the structural method, are selected as the loss-weighted average over all of the prior accident periods).

Step 4: Track Errors over Time

Calculate the absolute and squared projection errors of each method, which are simply the difference or squared difference between the ultimate loss projections at the valuation date and the true ultimate losses, which are known since we are simulating the data. Also calculate the error statistics on a relative basis (as a percent of the true ultimate losses). The relative error statistics can also be expressed as a ratio to the unpaid losses at the valuation date. Tabulate the error statistics from each calendar year's projections for each method and compare the reserving methods in terms of absolute (signed) error, magnitude (absolute value) of error and the speed of approach of their successive calendar year estimates to the true ultimate losses. The relative errors as proportions of unpaid losses can also be examined, but the changing denominator over time means the target is moving with each successive calendar year. The reporting format of results for a single trial of the experiment is shown in **Exhibit 6**.

Practical Considerations

How much data and how many trials are needed to draw useful conclusions? The question of credibility in general is central to actuarial science, but I will focus on practical limitations rather than theoretical support for "full credibility" of the database. In the scenario tested here, we have selected parameters which imply an average of 100 claims in each accident period, or 1000 claims in the database per trial. The actual generated number of claims is random due to the stochastic modeling of exposure and claim frequency. Credibility, under most measures, is an increasing function of the sample size. Therefore, it would be best to run an infinite number of trials in order to maximize the implied credibility of the experimental outcomes and get the best picture of the realism of the simulation algorithm and the performance of the different reserving methods, but we have limited the experiment to 25 trials for purposes of this paper. Therefore, the total number of claims simulated should be about 25,000.

I should point out that the experiment, while stochastic in nature and dynamically tested, is still static in one way - once the variables are randomly drawn and the calculations completed, the loss and claim count matrices are set. If a new claims manager was installed during calendar year 12, for example, the settlement lag parameters for open claims would not be revised. This means that were there a "perfect" reserving method which produced zero errors on a "perfect" simulated database (e.g. one that explained the real-world claims process exactly), even over an infinite number of trials, the same method would not produce zero errors when dynamically tested on a real book of claims.

One modification has been made to the experiment for the purpose of testing the performance of reserving methods. For liability lines with lengthy reserve durations, and for reinsurers who are frequently not “on the hook” until advanced ages of development, estimation of tail factors is a complex process which is usually somewhat distinct from the procedure used to develop the indicated reserves from historical data. It often involves curve-fitting and extrapolation, the results of which are then mated to the original reserving method somehow. While the subject of tail factors is extremely interesting (at least to me), I did not want any difference in tail factor estimation procedures to bias the comparative performance of the aggregate and structural reserving methods. Therefore, all of the error functions compare the “age 120 months” losses estimated at a given valuation age to the simulated losses emerged at age 120 months. In other words, we are pretending 120 months is ultimate and forgetting about any simulated losses emerging beyond this point. Depending on the characteristics of the process the analyst wishes to simulate, this simplification may cause no alarm at all or may leave the anticipated performance of the reserving methods very uncertain. In the scenario tested here, the parameters have been chosen so that there are usually no IBNR claim counts after 120 months, and only a small proportion of unsettled claims after 120 months. Since loss payments are assumed to be lump-sum at date of settlement in this model, the proportion of ultimate losses paid after 120 months is also small.

THE RESERVING METHODS

Paid Loss Development (Aggregate) Method

This is the canonical chain ladder reserving method, applied to the paid loss matrix generated by the simulation model. The concentration of actuarial judgment in this method is usually in the selection of age-to-age development factors, which may be based on historical weighted or simple averages, averages excluding extreme values, or linear or exponential trends. For a controlled test, the method must be made mechanical, and so the weighted average age-to-age factor over all observed exposure in the triangle is selected for each age interval. In my opinion, this selection is probably less likely to introduce bias into the method than simple average or “excluding high and low” selections.

In accordance with the desire to avoid the bias introduced by tail factor selections, the tail factor from 120 to ultimate has been set to unity. Therefore, the “ultimate losses” of the aggregate method are valued at 120 months rather than at a true ultimate. I stress again that aggregate methods are not applied so mechanically in everyday actuarial practice, but scientific method as well as statistical theory demand that we standardize each trial. The mechanics of the method should be well-known to actuaries; for completeness, [Exhibit 4](#) shows the setup of the calculations for multiple calendar years of application. Some of the matrices in the exhibit may be a little confusing, because it is applying the method at calendar years-end 10 through 19 all at once, but rest assured that we are using “plain vanilla” paid loss development at each valuation. The numerical values shown here and in [Exhibit 5](#) should tie together within the exhibit, but do not represent any sort of test results, only an example.

Closed Claim Cost (Structural) Method

This method is more complicated and customized for the experiment than the aggregate method. Since its mechanics involve extensive manipulation of matrices of payment and claim count data, it is tempting to resort to the notation of matrix algebra to explain it, but I think that might obfuscate more than it helps. Therefore, the method is explained in plain English and **Exhibit 5** illustrates each step. Note that it is essentially impossible to apply the method at multiple calendar years-end on one exhibit as we did with the aggregate method, so a separate sheet runs each calendar year application. The exhibit shows the setup for the first evaluation - at the end of calendar year 10.

Step 1: Project Ultimate Claim Counts (page 2)

Using the historical triangle of reported counts, select age-to-age factors for each evaluation interval and compute the age-to-ultimate factor for each accident period as the product of all the selected factors beyond the current age of the period. For each period, multiply the reported counts at the valuation date by the respective age-to-ultimate factor to get projected ultimate claim counts. In other words, apply the chain ladder development approach to claim counts. (We select the historical weighted average age-to-age factor in this experiment, just as we did for losses in the aggregate method.) An additive approach could be used instead without loss of the integrity of the structural method.

Step 2: Project the Claim Closure Pattern (page 3)

Use the triangle of incremental closed counts during each interval and the vector of projected ultimate counts. For each accident period, divide the number of closed claims in each interval by the projected ultimate for the accident period to get the historical closure density (the density percentage triangle is not shown). Examine each column of the resulting triangle of percentages to select a representative percent of projected ultimate to be closed in the interval (shown at the top of page 3). The set of selections will form a selected closure density function, except for the “tail factor” which, in this case, represents the proportion of ultimate claims closed beyond the last evaluation age in the matrix. In this experiment, to substitute for actuarial judgment as we did in selecting age-to-age factors, we select the historical weighted average closure percentage (over all applicable accident periods) for the density in each interval, and simply leave the tail factor to be 100% minus the sum of the selected percentages over all intervals.

Step 3: Project Claims Closed in Each Interval (page 3)

Subtract the current cumulative number of closed claims from the projected ultimate for each accident period to get the estimated unsettled counts for the period. Allocate the unsettled counts to future settlement intervals as follows. The selected density (incremental proportion of ultimate claims to be closed) in each future interval, divided by the remaining cumulative proportion of claims to be closed over all future intervals, should be the share of unsettled claims assigned to the given interval. This sounds complex in English but it is the “common sense” way to do the allocation. Multiply the allocation share by the number of unsettled claims for the period to get the number of unsettled claims during the interval. The only theoretical problem here is that non-integer amounts are allocated to each interval. When projecting aggregate unpaid losses over several accident periods, I believe it is more accurate to leave the fractional amounts alone, as they will be multiplied by projected severities and summed over all cells later in the method.

Step 4: Project Average Severities (pages 3 and 4)

Segregate the data triangles into columns representing age intervals or groups of intervals. On the premise that claims closing quickly have different average severities than those that stay open, calculate the historical average severity in each age group by dividing the total incremental paid losses (shown back on page 1), over all accident periods, in the age group by the number of claims closed in the age group. Once the baseline historical severity is determined for every group, project the severities of claims closed in each group in future calendar periods (the cells below the diagonal in the triangle of severities). To the extent we expect a severity trend, the future severities are projected by trending the baseline historical severities forward from the average time index of the historical column above the diagonal to the average time index in each of the cells below the diagonal (in the same column). The midpoint of accident period 1 is nominally selected as time zero. In this experiment, we match the assumed severity trend rate to the assumed rate of inflation in the simulation model to avoid specific bias due to underestimation of a nonstochastic variable.

Step 5: Project Unpaid Losses (page 4)

The goal of the previous steps was to fill in each of the cells below the diagonal of a triangle of incremental closed claim counts and a triangle of average severities. With these projections in hand, the unpaid losses are simply the sum of all the products of the corresponding cells in the closed claim triangle and the severity triangle. Note that unpaid losses are the direct output of the structural method, whereas aggregate methods usually project ultimates and subtract some known quantity to get total or IBNR reserves.

A good high-level discussion of the contrasts between aggregate and structural reserving methods is given by Wiser [7]. Basically, aggregate methods are growth models which are applicable to projecting the value of any random variable which is a function of time. Structural methods follow a process based on frequency (or failure rates) and severity (or some other measure of magnitude or intensity) to project individual occurrences which are integrated in the final step and added to known payment information to find the value of the objective function (aggregate ultimate costs). The development of a database simulation algorithm and a closed claim cost method of reserving notwithstanding, the central question considered by the experiment is how these two fundamental approaches to reserving perform over time in a controlled environment. We now look at the results of the experiment and the metrics used to answer that question.

THE RESULTS

Base Scenario

The base scenario is designed to simulate the runoff of liability claims for a stable client company in a moderately inflationary environment without the occurrence of a catastrophe in the experience period. The exposure growth (in real terms) is flat on average but ranges about 10% in either direction in a particular non-crisis year (mean of zero, SD of 5%). Accidents giving rise to liability occur evenly over the year. They are reported an average of one year after they occur, but some accidents can be latent for several years (mean of 12 months, SD of 12 months). Settlements are reached (or verdict returned) an average of two years after notification, but again some claims drag out for several years (mean of 24 months, SD of 24 months). The average claim costs \$20,000 in today's dollars, and the operations and insurance program of the client company are such that catastrophic claims will not occur during the experience period (mean of \$20,000, SD of \$20,000). The general rate of claim cost inflation applies to costs at the date of settlement and remains steady at 5% over the experience and payout periods.

Exhibit 7 presents the aggregate results of 25 simulations of ten accident periods at a time, each producing an average of 100 claims. This report is produced for each trial individually (as demonstrated in Exhibit 6), but it would be cumbersome to display all 25 pages in the paper. On page 1 of Exhibit 7, separately for the paid chain ladder and the closed claim cost methods, the true simulated losses at age 120 months (leftmost column) for each accident period, 1 through 10, are compared against the method's valuations at the end of calendar years 10, 11, 12, 13, and 14 (columns progressing to the right). The sums over all accident periods are displayed beneath in the same columns. The aggregate absolute difference between each calendar year's evaluated losses and the true losses, the relative difference as a percent of the true losses, and the squared difference from the true losses are all displayed. The same display, using true and estimated unpaid losses at each calendar year-end is shown on page 2. The only difference is that the true unpaid losses vary by evaluation date, but the true "age 120" proxies for ultimate losses do not. The absolute and squared differences in unpaid losses will be the same as on page 1 (we are only subtracting paid-to-date from both values), but the relative errors to unpaid losses will be larger as the base becomes smaller over time.

A few observations seem to jump out from the results. The structural method appears to perform much better than the aggregate method at the earliest evaluations, but the aggregate method "catches up" and eventually outperforms the structural method in terms of accuracy after calendar year 12. In Khury's framework, the aggregate error radius function has an (approximately) zero mode, but decays rapidly as a function of time. On the other hand, the structural method does not improve as rapidly in the second calendar year, but begins a downward drift afterward. In fact, the error drifts across the true value over time, progressing from an overestimation to an underestimation of the ultimate losses. This indicates either an oscillation inherent in the method or a downward bias. We might suspect a bias, though not as pronounced, in the aggregate method as well, since the estimated losses do cross the true value at the fifth evaluation.

Examining the results by accident period is also interesting. In terms of sign, the aggregate method errors do not show a strong pattern across accident periods at the first test date - some are negative, others are positive with no obvious serial correlation. The estimated losses for many individual accident periods oscillate over time, crossing the true losses and sometimes more than once. Improvements in the projections over time, on a relative basis, do not show a strong trend toward later or earlier accident periods. In terms of magnitude, however, it is clear that the worst errors occur in the most immature, leveraged accident periods - over \$12 million of the \$13 million overstatement at calendar year 10 is on accident periods 9 and 10 alone! The errors in prior accident periods jump around somewhat, but generally do not exceed \$1.5 million in either direction.

On the other hand, the structural method clearly appears to dampen the fluctuations in the immature accident periods. The errors for accident periods 6-10 are generally of smaller magnitude, and not necessarily the same direction, as those of the aggregate method. The \$2.5 million overstatement in the latest period at calendar year 10 is significant, but certainly much less than that of the aggregate method. The structural method seems to have trouble, especially in the latest accident periods, at subsequent evaluations. Downward movements of nearly \$1 million occur for accident periods 9 and 10 between calendar years 13 and 14, quite a swing at such an advanced age. The aggregate method is apparently less sensitive to slight changes in the data once the loss development pattern has been established. An interpretation might be that the structural approach might be more powerful when very little payment and claim count information has emerged for an exposure period, where claim count and trended severity estimates are (both individually and as a product) more stable than multiplicative relationships in aggregate cost data. When cost data become more mature, the aggregate relationship may be more stable than severities which have been extrapolated many periods into the future and claim closure patterns which are selected from long-past accident periods.

"Mining" both the aggregate results and the individual trial data might yield many more interesting conclusions, particularly if a much larger set of trials were run. Some possibilities for the analyst who has tailored the model to a particular client company or situation are binary (signed error) tests for bias by accident period and calendar year, plotting the error radius function by accident period over five (or more) calendar years and looking for decay patterns in more and less mature periods, and so on. I have not explored these paths in the paper for several reasons. The tools, algorithms, and reserving methods described here are intended to stimulate thinking and discussion about the process for setting up computer experiments, the key elements of a good simulation technique, and how the performance of reserving methods in general should be tested. Therefore, the focus is not on drawing conclusions. Also, I believe the imposition of conclusions regarding a new technique applied to a new reserving method would tend to stifle rather than stimulate criticism, creative thinking, and further research.

CONCLUSIONS

At the outset of the paper, I asserted some objectives. The reader, of course, must decide how well we have accomplished any or all of them, but I will analyze our progress on each front for a moment.

- We have partially filled the gap in historical literature regarding simulation approaches to generating data and structural approaches to loss reserving, just by the expositions of the simulation algorithm and the closed claim cost method. The practical utility of either is still open to question and requires further experimentation as well as attempts to actually apply the ideas to particular client companies.
- We have developed a working tool to generate the required databases according to the model we have discussed. However, many refinements should be added so that the end user has more choices about how the data will be simulated and how the output will be organized.
- We have shown that assumptions about key components or variables of the insurance function can be incorporated into a model which generates reasonable-looking actuarial data suitable for analysis with several reserving techniques. We do not yet know how well the model approximates reality for any particular segment of the insurance industry or client company.
- We have created, explained and tested a structural approach to estimating loss reserves. The assumptions underlying it are, in my opinion, reasonable, but many of them could be inaccurate in some specific situations. Also, the projections of the method could be unacceptably sensitive to some of them. Again, further research as well as practical application will help us identify the weak areas.
- We have analyzed the performance of our simulation model and reserving methods using time-dependent performance standards based on absolute, relative, and squared differences from true losses at 10 years of development. Such criteria provide unique insights which are not possible with static evaluations of a reserving technique. Dynamic testing should, however, be expanded to use more sophisticated performance criteria and parametric or nonparametric statistical tests.
- We have, starting from “scratch”, incorporated our quantitative expertise as well as knowledge of the insurance process and actuarial judgment about the variables in the process to generate estimates of ultimate losses on a volume of exposures without basing any of our functions on historical data. What we can not do without some historical data is calibrate the model to make it accurate for a particular client company. However, we have reduced our reliance on a critical volume of historical data and avoided the constraints introduced by the rigid data requirements (regular data collection intervals, consistency and accuracy of case reserves, consistency of historical payout patterns, etc.) of many canonical actuarial techniques.

Given the relatively small number of claims simulated in the experiment to date and the construction of just a baseline scenario, strong conclusions about the performance of reserving methods are difficult to draw. I have tried to make as many germane observations as possible about the “typical” single experiment run for this paper, and we do have some results which indicate that intuitions about the performance of reserving methods at successive points in time may not always be supported by empirical studies of the insurance process. In particular, structural methods may not be superior to aggregate methods mature accident periods, and both structural and aggregate methods’ estimates may not monotonically converge to the ultimate costs. In addition, both structural and aggregate reserving methods may be prone to biases if not adjusted.

The most glaring need at this point is to spend a lot more time and imagination experimenting with different scenarios, sensitivity testing the results of each scenario to key algorithmic and parameter assumptions, and simply generating a larger volume of data for each scenario (running more trials) so that we can give more credibility to the results. That being said, in the final section I will bring up a few specific ideas and issues to think about as we undertake additional research.

REFINEMENTS AND RESEARCH OPPORTUNITIES

As I mentioned before, an algorithm should be flexible enough to incorporate new ideas without requiring a complete redesign. The specificity of application of this simulation model to individual client company situations is limited only by the analyst's imagination (and programming skills or ability to marshal those colleagues with them). The following are some ideas which I have not implemented, but which cry out for experimentation.

Exposure and Inflation Time Series

We can draw from the great body of literature on modeling business, economic and monetary cycles to build a time series model of these variables containing both trend and cyclical components. Such a model could be made to "plug and play" with this simulation routine, though not without some programming work. In addition, since underwriting or company growth rates are typically a controllable variable in the eyes of management, testing the predicted value of liabilities under various assumed growth scenarios is a legitimate use of the model.

Alternative Statistical Assumptions

In the practice of using simulation models, calibration takes the place of adjusting historical data triangles. The normal and gamma distributions have been used here for simplicity and flexibility and because they are reasonable, but a variety of statistical distributions and parameter assumptions can be used to customize the model. It can be calibrated against historical measurements of the variables (by statistical or other means), coupled with managerial foresight about the expected future behavior of the key variables, developed from theoretically correct models of the variables (see the discussion about report lags above), and of course judged against the common sense of the users.

Partial Payments and Subdivision of Claim Types

For simplicity, the algorithm assumes that all liabilities associated with the claim are discharged at the date of closure, in one lump-sum check. In reality, loss adjustment expense assumptions, information about statutory benefit disbursement procedures for WC, and other refinements can be incorporated into the programming of the algorithm to simulate the partial payments which are common in some lines of insurance. The model could also be easily enhanced to randomly assign a "type" to each claim, where each type of claim exhibits different lags, severity, and timing of payments. For example, in a WC simulation, perhaps one claim in ten could be a permanent partial disability claim which generates many annuity payments rather than one lump sum, and five in ten could be medical only claims which are discharged for a constant small amount without a significant settlement lag.

Separation of Losses and Loss Adjustment Expenses

The programming of the model could be easily modified to simulate loss adjustment expense as well as loss payments. In liability lines, one way would be to open two “subclaims” for every claim drawn, with an identical report lag but different settlement lags, partial payment patterns, and ultimate severities. The first subclaim would be losses, the second legal and other adjustment expenses. Similar schemes could be programmed to track indemnity versus medical payments in WC, and other payout structures unique to certain contractual arrangements.

Relationship of Report Lag, Settlement Lag, and Severity

As actuaries, we generally assume that claims which are reported (and settled) at more advanced ages cost more on average. To my knowledge, there has been no attempt in the actuarial literature to empirically estimate the general quantitative relationship between report lags (or settlement lags) and ultimate severity. Therefore, I have no guidance on which to base a relationship which could be modeled in the simulation routine. However, this relationship is an integral part of the process we are trying to model, whether our models extrapolate from historical data or operate on a simulated database. It is an area ripe for further research.

Mechanical Refinements to the Closed Claim Cost Method

The closed claim cost method presented here has some potentially serious technical flaws. It depends heavily on the absence of partial or annuity payments and the relationship between closure age and severity. Its results are sensitive to the choice of severity trend rate and the resulting projected severity applied to unsettled claims. Its projection of the closure pattern of unsettled claims depends on good actuarial judgment as well as the consistency of past settlement patterns. Some of these weak spots may be endemic to structural reserving methods, but others could be cured with a little ingenuity. The method tested here is intended to spur thinking about structural methods and ways to integrate claim-level data into the job of loss reserving, and represents an archetypical approach, rather than a ready-made actuarial technique for a particular empirical data set.

Incorporation of Outstanding Losses

Since case reserves are a function of human judgment, one might not expect case reserve data to adapt well to a simulation algorithm. I have chosen paid loss data as the basis of the experiment, but perhaps industrywide case reserving patterns and practices could be used to create a simulated reported loss matrix. In fact, Stanard tried to incorporate case judgment in his model by including a “reserve error” random variable. I decided that making the same attempt would add unnecessary “noise” to the experiment and complicate the simulation algorithm.

Estimation of Tail Factors

As discussed above, to avoid introducing superfluous components to the experiment, we tested the methods against losses emerged at ten years of development, not the ultimate losses generated by the model. In reality, tail factors are the source of much of the parameter risk associated with a particular reserving method, and the approach used to calculate the tail factor should be subject to the same performance testing as any other component of the method. Two methods which otherwise do equally well on a given test could diverge substantially if one tail factor is estimated more efficiently than another, particularly in a line like WC or professional liability.

In aggregate methods, tail factors are usually based on extrapolation (by judgmental or statistical means) from the empirical payout or reporting patterns embodied in the age-to-age factors. The intuition is less clear for structural methods, but I suggest choosing an average severity in the tail interval, perhaps based on a multiple or trend of the average severities in prior intervals. Multiplying the selected average severity by the projected number of claims closed in the tail interval (a by-product of the projected closed claims calculation) generates an estimate of losses paid in the tail interval. Of course, like all tail factors, this estimate will be sensitive to its input variables - in particular, it will depend dollar for dollar on the selected average tail severity and be strongly influenced by the closure density which determines the number of claims closed in the tail.

This simulation model can certainly be extended to use a twenty or even fifty-year payout matrix, essentially eliminating the tail interval, but I did not want to undertake an experiment handling such a large volume of data. I encourage excess and specialty insurers and other interested parties to think about and implement the types of experiments which would prove most useful to their interests.

WORKS CITED

- [1] Stanard, James N. "A Simulation Test of Prediction Errors of Loss Reserve Estimation Techniques", *Proceedings of the Casualty Actuarial Society LXXI*, 1985, p. 124.
- [2] Khury, C.K. "Loss Reserves: Performance Standards", *Proceedings of the Casualty Actuarial Society LXVII*, 1980, p. 1.
- [3] Adler, Martin, and Kline, Charles D., Jr. "Evaluating Bodily Injury Liabilities Using a Claims Closure Model", *Evaluating Insurance Company Liabilities*, Casualty Actuarial Society Discussion Paper Program, 1988, p. 1.
- [4] Weissner, Edward W. "Estimation of the Distribution of Report Lags by the Method of Maximum Likelihood", *Proceedings of the Casualty Actuarial Society LXV*, 1978, p. 1.
- [5] Hogg, Robert V., and Klugman, Stuart A. *Loss Distributions*, pp. 21-25. New York: Wiley, 1984.
- [6] Butsic, Robert P. "The Effect of Inflation on Losses and Premiums for Property-Liability Insurers", *Inflation Implications for Property-Casualty Insurance*, Casualty Actuarial Society Discussion Paper Program, 1981, p. 58.
- [7] Wisner, Ronald F. "Loss Reserving", *Foundations of Casualty Actuarial Science*, Casualty Actuarial Society, 1990, p. 143.