WORKERS COMPENSATION RESERVE UNCERTAINTY

DOUGLAS M. HODES, SHOLOM FELDBLUM,
AND GARY BLUMSOHN

Abstract

The increased emphasis on solvency monitoring of insurance companies, along with the American Academy of Actuaries' vision of an expanded role for the Appointed Actuary, have stimulated reserving specialists to quantify the uncertainty in their estimates. This paper measures the uncertainty in workers compensation loss reserve indications, compares it to the "implicit interest margin" in statutory (undiscounted) reserves, and examines the implications for capital requirements.

The paper uses a stochastic simulation analysis to model the loss reserving process, with separate but interlinked components for the process risk of loss development, the parameter risk of estimating future age-to-age link ratios, and autocorrelated future interest rates. In addition, the past monetary inflation implicit in paid loss development link ratios is replaced with stochastically generated future inflation rates that are linked to both the concurrent interest rates and the previous year's differential between the inflation rate and the interest rate. Separate simulations are performed for each accident year, and loss development tail factors are generated by an inverse power curve fit to extend the development from 23 years to ultimate.

An "expected policyholder deficit ratio" procedure is used to calibrate the capital needed to guard against reserve uncertainty. Because of the statutory benefits in workers compensation, the steady payment patterns, and the long average duration of compensation reserves, the implicit interest margin in statutory reserves exceeds the *capital required to guard against the variability in the reserve estimates at a 1% expected policyholder deficit level.*

The appendices to the paper contain descriptions of the simulation procedures, as well as a comparison of the paper's conclusions with those of the NAIC's risk-based capital formula.

ACKNOWLEDGEMENT

The authors are indebted to Aaron Halpert, Stephen Lowe, Roger Bovard, Roy Morell, Richard Fein, Louise Francis, and Joseph Boor for reviews of earlier versions of this paper. Any remaining errors, of course, should be attributed to the authors alone.

1. INTRODUCTION

Actuaries have developed a host of techniques for producing point estimates of indicated reserves. Current regulatory concerns, as reflected in the NAIC's risk-based capital requirements, and developing actuarial practice, as reflected in the American Academy of Actuaries' (AAA) vision of the future role of the Appointed Actuary, now stress the uncertainty in the reserve estimates in addition to their expected values. This paper demonstrates how the uncertainty in property/casualty loss reserves may be analyzed, and it draws forth the implications for capital requirements and actuarial opinions.

Genesis of this Paper

This paper was stimulated by the NAIC's risk-based capital efforts and by the AAA's vision of the valuation actuary:

- The reserving risk charge, which measures the potential for unanticipated adverse loss development by line of business, is the centerpiece of the NAIC property/casualty risk-based capital formula, accounting for about 40% of total capital requirements before the covariance adjustment and about 50% after the covariance adjustment (see Feldblum [23]). Because good actuarial analyses of loss reserve uncertainty are still lacking, the reserving risk charges were based on simple extrapolations from past experience, with a large dose of subjective judgment to keep the results reasonable.

- The Appointed Actuary presently opines on the reasonableness of the Annual Statement's point estimates of loss and loss adjustment expense reserves. The American Academy of Actuaries envisions an expanded role, in which the actuary opines on the financial strength of a company under a variety of future conditions (see [1]). The greater the uncertainty in the reserves, the greater the range of reasonable financial conditions that the actuary must consider.

Issues Addressed

This paper focuses on the uncertainty in workers compensation loss reserves. Specifically, it addresses the following issues:

- How should the uncertainty in loss reserves be measured? In other words: How might the variability in the loss reserve estimates best be quantified?
- What insurance characteristics, such as payment patterns and contract obligations, affect reserve uncertainty?
- How does the measure of variability that underlies risk-based capital requirements differ from the measure of variability that underlies the actuarial opinion? More specifically, how does the variability of the discounted, "net" reserves (i.e., loss obligations after consideration of return premiums and additional premiums on retrospectively rated policies, valued on an economic basis) differ from the variability of the undiscounted, "gross" reserves?

The Mixing of Lines

Why concentrate on workers compensation? Why not discuss property/casualty loss reserves in general, of which workers compensation is but one instance?

This is one of the primary errors that have hampered past analyses of loss reserve variability. Many observers have contrasted short-tailed lines like homeowners and commercial property with long-tailed lines like general liability and automobile liability, and they have noted the greater reserve uncertainty associated with the latter lines of business. Consequently, they have reasoned that reserve uncertainty is associated with the length of the average payment lag (i.e., reserves with longer average payment lags have greater uncertainty).

To see the error in this reasoning, let us extend the comparison to life insurance reserves. Single premium traditional life annuities have the longest reserve duration of the major life insurance products. Yet these products have low reserving risk, since the benefits are fixed at policy inception and mortality fluctuations are low.[1]

The bulk of workers compensation loss reserves that persist more than two or three years after the accident date are lifetime pension cases. The indemnity portions of these claims are disabled life annuities, with long duration and low reserve fluctuation for large compensation carriers. For the major insurance companies, the longest workers compensation reserves often have relatively low risk.[2]

---

[1]These products have significant interest rate risk, which is indeed affected by the average payment lag of the liabilities. For the quantification of interest rate risk for property/casualty insurance companies and the implications for risk-based capital requirements, see Hodes and Feldblum [35].

[2]See Feldblum [19], which compares reserve uncertainty among four property/casualty lines of business: workers compensation, automobile liability, products liability, and property. Compare also Meyers [47], who deals with same issue: "The purpose of this paper is to continue the debate on risk loading and discounting of loss reserves."

Meyers deals with workers compensation pension reserves, which have the highest ratio of implicit interest discount to reserve uncertainty, particularly for a large portfolio

The Peculiarities of Compensation Reserves

The quantification of reserve uncertainty must begin with the characteristics of the line of business. Four aspects of workers compensation reserves that affect the level of uncertainty are dealt with in this paper:

- *Payment Lag and Discount:* The previous section noted that most compensation reserves that persist more than two or three years after the accident date are lifetime pension cases. We compared these to life annuities, which are low risk reserves for large insurance companies. But the analogy is incomplete, since the statutory accounting treatment differs for these two types of business. Life annuities are discounted at rates close to current corporate bond rates. (The statutory discount rate for single premium immediate annuities—the life insurance product most comparable to workers compensation pension cases—issued in the first half of the 1990s is about 7% per annum.)

  Most property-casualty companies discount the indemnity portion of workers compensation lifetime pension cases at 3.5% or 4% per annum, which is well below their actual investment earnings. All other claims, as well as the medical portion of life pension cases, must be shown at undiscounted values in the statutory statements. The analysis in this paper indicates that the low fluctuations in these reserves, combined

---

of reserves. We look at the distribution of age-to-age link ratios, using a lognormal assumption and a Bayesian analysis of parameter risk; Meyers looks at the distribution of ultimate pension costs, using Makeham's mortality curve, again with a Bayesian analysis of the parameter risk. We use an expected policyholder deficit analysis, using a 1% EPD ratio, to calculate capital requirements; Meyers uses a utility function analysis to calculate the needed risk load. The two methods differ, though the results are similar: the implicit interest discount overwhelms the needed risk load or capital requirement. See especially Meyers' Tables 6.1 and 6.2 on page 182. The needed risk load in Meyers' illustration is about $400,000, with some variance depending on the parameters chosen in the utility function. The implicit risk load is $34.5 million assuming no tabular discounts and $9.3 million assuming tabular discounts at a 3.5% annual interest rate.

Hayne [32] shows a method of calibrating the uncertainty in the loss reserves based on loss frequency and loss severity assumptions. Hayne demonstrates his method, but does not provide numerical illustrations based on insurance data.

with the large implicit interest margin, create enormous hidden equity in statutory balance sheets.

- *Statutory Benefits:* What about non-pension cases? Do non-pension compensation reserves have the same uncertainty as many commercial liability reserves have? After all, industry studies have found similarly strong underwriting cycles and reserve adequacy cycles in several of these lines of business.[3]

  Yes, underwriting results are driven by industry cycles, and so underwriting results vary greatly from year to year, whether in workers compensation, general liability, or automobile liability. But underwriting cycles reflect primarily the movement of premium levels, not fluctuations in loss experience. Reserve adequacy cycles are a secondary effect, driven by management desires to smooth calendar year operating results. They reflect the accounting treatment of company results, not the uncertainty inherent in the reserves themselves.[4]

  When a products liability or medical malpractice accident occurs, the claim may not be reported for some time. Even after the claim is reported, the case may not be settled until years later, and the amount of the loss liability depends on the vagaries of court decisions, societal opinion, and jury awards. This is a major source of reserve uncertainty in the liability lines of business.

  In workers compensation, most claims are reported rapidly. (It is hard for the employer to be unaware that a worker has been injured on the job and is on disability leave.) Benefits are mandated by statute, and disputes are resolved relatively quickly by administrative judges. For the major countrywide insurers with broad mixes of business, the paid-loss link ratios, or "age-to-age" factors, are stable in workers compensa-

---

[3]On the relationship between underwriting cycles and workers compensation reserve fluctuations, see Ryan and Fein [51] and Butsic [14].

[4]On the loss and premium effects of underwriting cycles, see Daykin, Pentikainen, and Pesonen [17], Cummins, Harrington, and Klein [16], and Feldblum [24].

tion, both for pension and for non-pension cases, unlike the comparable factors for the liability lines of business.[5]

- *Tail Development:* But don't workers compensation reserve estimates need large "tail factors," just as liability reserve estimates need? And aren't these tail factors highly uncertain, even as the liability tail factors are?

  Volatile commercial liability tail factors often reflect the emergence or the settlement of claims decades after the occurrence of the accident, such as toxic tort and environmental liability claims. This is true reserve uncertainty.

  Much of the volatility of workers compensation tail factors stems from two causes:

  1. First, mortality among permanently disabled workers, particularly for insurers with smaller blocks of business, is uncertain. For insurers with larger volumes of business, mortality fluctuations are less significant for annuity reserves.
  2. Second, workers compensation tail factors are affected by monetary inflation, both for cost of living adjustments to indemnity benefits and for all aspects of medical benefits. Inflation levels, especially for 30 or 40 years into the future, are extremely uncertain. This is parameter risk, not process risk, so it affects both large and small insurers.

[5]This paper emphasizes reserve estimates drawn from paid loss development methods. To avoid issues of company case reserving philosophy, we use loss payments only, not case reserves or reported losses, to quantify the uncertainty in the loss reserve estimates.

Reserving procedures based on case incurred loss development methods depend on company case reserving philosophy and stability. Some of the fluctuations in case reserves stem from different causes than the fluctuations in paid amounts. For instance, many temporary total disability claims are subsequently reclassified as permanent disability claims, causing an immediate change in the case reserve.

We do not have independent information about the reserve uncertainty inherent in case incurred reserving methods. The procedure used in this paper to quantify reserve uncertainty is not directly applicable to "incurred" methods. Analysis of the uncertainty in "paid" methods versus "incurred" methods would be a worthwhile subject for future studies.

This creates great uncertainty in the undiscounted reserve, and the actuary opining on reserve adequacy for statutory statements should consider a wide range of "reasonable" estimates. But the economic value of the reserve is less affected by long-term inflation rates for two reasons: (a) Much of the effect of high long-term inflation rate scenarios appears after 10 or 15 years, when the present value of these payments is much reduced. (b) The effect of high long-term inflation rates is often partially offset by high long-term interest rates.

- *Policy Type:* The type of insurance contract—such as "large dollar deductible" policy versus retrospectively rated policy—affects the degree of reserve uncertainty. A high percentage of the workers compensation contracts covering large employers are retrospectively rated. That is, the premium paid by the employer (the insured) is a function of the incurred losses. If loss reserves develop adversely, the insurer will collect additional retrospective premiums from the employer.

For loss-sensitive contracts, estimates of reserve uncertainty must be distinguished from their implications for capital requirements and actuarial opinions. Risk-based capital requirements reflect the equity needs of the insurer. Similarly, the envisioned future role of the appointed actuary is to opine on the financial strength of the insurer under various future conditions. To the extent that adverse loss development on a book of business is offset by favorable premium development, the financial condition of the insurer is unaffected, and there is less need for additional equity.

Other types of new insurance products have the opposite characteristics. Large dollar deductible policies and excess layers of coverage have higher reserve uncertainty per dollar of "net" reserves (i.e., reserves for the excess layer). A workers compensation reinsurer covering loss layers above high retentions may experience reserve variability unlike that experienced by a primary insurance carrier.

We may summarize the previous discussion in this section as follows. The novice actuary sees an insurer's large book of compensation reserves, notes the long payment lags and the strong underwriting cycles, and concludes: "There must be great uncertainty here. Moreover, unexpected development may severely affect the insurer's financial condition, so much capital is needed to guard against this risk." The experienced actuary replies: "No, because of the steady compensation payment patterns and the long payment lag of these claims, the reserving risk is low enough that it is outweighed by the implicit interest margin in the reserves."

2. MEASURES OF UNCERTAINTY

We have differentiated between the inherent uncertainty in reserve estimates and the accounting illusions caused by discretionary adjustments of reported reserves. Similarly, we may differentiate between actuarial measures of reserve uncertainty and regulatory measures of reserve uncertainty.

The Solvency Regulator and the Actuary

Suppose that the solvency regulator sees wide fluctuation in reported reserve levels and concludes that there is great uncertainty in the reserve estimates. The company actuary responds that the actual reserve indications have been stable. The shift in reported reserve levels from year to year stems simply from a desire to smooth calendar year earnings (see Ryan and Fein [51]).

> "*What difference does that make?*" replies the solvency regulator. *"We are concerned that the reported reserves may not be sufficient to cover the loss obligations of the company. What difference does it make whether the insufficiency stems from an inherent uncertainty in the reserve indications or from discretionary adjustment of the reported reserves?"*

We must differentiate between two types of reserve fluctuations:

- The valuation actuary tells the company's management how much capital it needs to guard against unexpected adverse events. Suppose the actuary's reserve analysis yields a point estimate of $800 million with a range of $650 million to $950 million, and the company is reporting $700 million on its statutory statements. The actuary's recommendation might be that the company needs $250 million of capital: $100 million for reserve "deficiencies" (the difference between the point estimate and the held reserves) and $150 million for reserve uncertainties.[6]
- The solvency regulator can not easily distinguish between adverse loss development stemming from unanticipated random occurrences and adverse loss development stemming from reserve inadequacies. The regulator estimates the variability of reported reserves and applies this figure to some base number. The base number might be the company's reported reserves (if the regulator believes that they are adequate) or an independent estimate of the company's reserve needs (if the regulator lacks confidence in the company's financial statements).

Regulators concerned with reserve uncertainty take the second viewpoint. Our primary interest in this paper is with the uncertainty inherent in the reserve indications themselves, the first viewpoint.

The difference is not in the *magnitude* of the uncertainty, but in the *method* of quantifying the uncertainty.

- The solvency regulator begins with the reserves reported by companies. How the companies determined these reserves, and

[6]In practice, the implicit interest margin in statutory reserves should be included in the valuation actuary's recommendation. To complete the illustration in the text, the actuary might add that there is $200 million of implicit interest margin in the statutory reserves, so only $50 million of capital is needed on an economic basis.

whether the reported reserves accurately reflect the actuary's indications, is irrelevant.

- The actuary examines the factors used to quantify reserve needs, such as age-to-age "link ratios," to determine the uncertainty in the reserve indications. How the company deviates from the reserve indications in its financial statements is not relevant to measuring the uncertainty inherent in the reserves.

Statistical and Financial Measures

We use several measures of reserve uncertainty in this paper: standard deviations, percentiles, and "expected policyholder deficits." The "expected policyholder deficit" (EPD) concept developed by Robert Butsic [13] is used here as a yardstick for the uncertainty in the reserve estimates. The EPD ratio allows us to translate "reserve uncertainty" into a "capital charge," thereby transforming an abstruse actuarial concept into concrete business terms. In Appendix A of this paper, we also discuss the "worst case year" concept used to measure reserve uncertainty for the reserving risk charge in the NAIC risk-based capital formula.[7]

Some readers will rightfully ask: "The NAIC worst case year concept is a simple but arbitrary accounting yardstick that is not supported by financial or actuarial theory. Why include it even in the appendix of an actuarial paper?"

The answer is important. This paper demonstrates that the implicit interest margin in full-value workers compensation reserves exceeds the capital needed to guard against unexpected reserve volatility. Some readers, aware of the 11% workers compensation reserving risk charge in the NAIC's risk-based capital formula, may mistakenly conclude that the "regulatory" and "actuarial" approaches to this problem yield different answers.

This is not so. The NAIC "regulatory" approach yields a similar result to that arrived at here. However, the workers com-

[7]For the NAIC worst case year concept, see Kaufman and Liebers [41] or Feldblum [23].

pensation charges were subjectively modified to produce capital requirements that seemed more reasonable to some regulators.[8] In fact, the apparent "unreasonableness" of the NAIC formula indications to these regulators stemmed from a misunderstanding of statutory accounting and of the risks of workers compensation business, not from any artifacts in the risk-based capital formula. A full discussion of the NAIC approach to reserve uncertainty embodied in the risk-based capital formula is presented in Appendix A.

3. THE QUANTIFICATION OF UNCERTAINTY

Attempts to measure reserve "uncertainty" often dissolve for failure to make clear (i) what exactly we seek to measure and (ii) how we ought to measure it.

This paper combines three elements to analyze the uncertainty of loss reserve estimates:

- A statistical *procedure* to quantify the uncertainty, relying on a stochastic simulation of the loss reserve estimation process.
- A *yardstick* to measure the uncertainty, relying on the expected policyholder deficit ratio.
- The *intuition* that explains the source of the reserve uncertainty, focusing on payment patterns, interest rates, and inflation rates.

Actuarial Procedures

Loss reserve estimates stem from empirical data, such as reported loss amounts or paid loss amounts, combined with actuar-

[8]For example, upon re-examining the workers compensation reserving risk charge in November 1996, using the NAIC formula but with more accurate figures and no subjective adjustments, the American Academy of Actuaries task force on risk-based capital found that the appropriate charge should be −12%, not the +11% in the NAIC risk-based capital *Instructions*. However, the AAA task force noted that any worker's compensation charge less than +10% would be politically infeasible to implement, so no effort was made to change the formula.

ial procedures, such as chain ladder development methods. Loss reserve uncertainty stems from both of these components.

- Random loss fluctuations may cause past experience to give misleading estimates of future loss obligations, and systemic changes (such as managed care) create uncertainty about future patterns.
- Imperfect actuarial analysis of the data may lead to invalid reserve estimates.

The two causes are intertwined. The ideal reserving actuary is ever watchful of data anomalies and will adjust the reserving procedures to avoid the most likely distortions (see, for instance, Berquist and Sherman [5]).

In this paper we do not measure the uncertainty stemming from imperfect actuarial practice. Rather, we assume a standard reserving technique that is often used for workers compensation; namely, a paid loss chain ladder development method.[9]

In practice, reserving actuaries use a variety of techniques. Even when employing a paid loss chain ladder development method, rarely does the reserving actuary follow the method by rote, with no analysis of unusual patterns. To the extent that actuarial judgment improves the reserve estimate, this paper overestimates the reserve uncertainty. To the extent that actuarial judgment masks the true reserve indications, this paper might underestimate the reserve uncertainty.

This paper measures the uncertainty inherent in the empirical data used to produce actuarial reserve estimates. It does not attempt to measure the uncertainty added or subtracted by the quality of actuarial analysis.

[9] We chose this technique, rather than a reported loss chain ladder development technique or Bornhuetter–Ferguson (expected loss) techniques, because it is dependent on claim payment patterns, and not on individual company case reserving practices. Thus, we are measuring the uncertainty caused by fluctuations in actual claim patterns, and not by changes in company case reserving practices.

Empirical Data

How should we measure the uncertainty inherent in the empirical data? The two extremes are described below, neither of which is sufficient by itself.

- We may simulate experience data, develop reserve indications, then continue the simulation to see how accurately the indications forecast the final outcomes.[10] This method is entirely theoretical. The amount of "uncertainty" depends on the simulation procedure. If the simulation procedure is firmly grounded in actual experience, the method works well. If the simulation procedure is chosen more for its mathematical tractability than for its empirical accuracy, the results may not mirror reality.
- We may look at actual experience, develop reserve indications at intermediate points in time, and then compare the indications with the final outcomes.[11] This method is "practical"—so practical, in fact, that the uncertainty measurements are often distorted by historical happenstance.[12]

A good actuarial procedure charts a middle course. We use stochastic simulation of the experience data to ensure statistically valid results. But the simulation parameters are firmly grounded in 25 years of actual paid loss histories from the country's largest workers compensation carrier.[13]

---

[10]See, for instance, Stanard and the Robertson discussion [56].

[11]This is the procedure used by the NAIC risk-based capital formula to estimate reserve uncertainty by line of business.

[12]See the report of the American Academy of Actuaries Task Force on Risk-Based Capital [44].

[13]Some reviewers of earlier drafts of this paper have questioned: Perhaps this insurer has more stable paid loss triangles than other insurers have, because of its size, because of its claim settlement practices, or because of its diversified mix of business. This is a valid comment. Small regional insurers may have different degrees of volatility in their reserve estimates. In particular, smaller insurers have greater process variance in the occurrence of lifetime pension cases, many of which have large total costs, both indemnity and medical. Expansion into new classifications or new states may similarly increase the uncertainty in the reserve estimates. See also the following footnote.

We describe the three elements of the analysis: (i) the stochastic simulation, (ii) the expected policyholder deficit ratio "yardstick," and (iii) the explanatory factors.

The Stochastic Simulation

We begin with 25 years of countrywide paid loss workers compensation experience, separately for indemnity and medical benefits, for accident years 1970 through 1994. From these data we develop 20 columns of paid loss "age-to-age" link ratios, as shown in Exhibits C-1 and C-2.[14]

We fit each column of "age-to-age" link ratios to lognormal curves, determining "mu" (μ) and "sigma" (σ) parameters for each. We perform 10,000 sets of simulations to generate the age-to-age factors that drive the simulated loss payments.

Twenty-five accident years yields 24 columns of "age-to-age" factors. The last four columns contain too few historical factors,

[14] Analysis of the uncertainty inherent in workers compensation loss reserve estimates must be grounded in actual workers compensation experience. The empirical data is the experience of the country's largest workers compensation carrier, with about 10% of the nation's experience during the historical period. To ensure confidentiality of the data, the dollar figures are normalized to a $100 million indicated undiscounted reserve.

Upon reviewing an earlier version of this paper, Stephen Lowe pointed out that "Because of its large market share, [your company's] experience probably does not respond to changes in mix of business by hazard group or state... . For smaller companies, changes in mix of business may add uncertainty beyond what is captured in your model." Similarly, the *Proceedings* referees for this paper write "For many companies, especially those with changes in mix of the type of business they write (different classes, different states) or changes in claims administration practices, the factors are not so stable."

This view is consistent with Allan Kaufman's recommendation that a "small company charge" be added to the risk-based capital formula because small companies experience greater fluctuation in underwriting results and in adverse reserve development. For political reasons, the small company charge was never added to the risk-based capital formula (see Feldblum [23]). In a review of the 1994 risk-based capital results, Barth [2], a senior research associate in the NAIC's research department, similarly concludes that "the R4 RBC i.e., (reserving risk) for companies with large reserves may be higher than necessary, relative to smaller companies."

Lowe, Kaufman, Barth and the *Proceedings* referees are correct. Small companies, or companies entering new markets or developing new products, may experience greater reserve uncertainty than implied here. This paper shows a method for quantifying reserve uncertainty, and it applies the method to the historical data of one particular insurer. To estimate the uncertainty of their own reserve estimates, readers should apply the methods described here to their own company's data. The numerical results in this paper can not necessarily be applied indiscriminately to other insurers.

so instead of fitting these columns to lognormal curves we include these development periods in the "inverse power curve" tail.[15] See Appendix C for a full description of the reserve estimation and simulation procedures.

Standard reserving methods, which forecast best-estimate future age-to-age link ratios, assume that the same factor will recur in each subsequent accident year. In actuarial parlance, when one "squares the triangle," the same age-to-age link ratios appear in each column for all subsequent accident years.

The procedure in this paper uses separate simulations for each subsequent accident year. We are simulating *actual* reserve development, where the process risk in each future accident year is independent of that in the other accident years.

Types of Risk

We categorize risk into two types: process risk and parameter risk (Freifelder [29], Miccolis [48]). We illustrate these components of risk with the fitting procedure described above.

Process Risk: Suppose that we *knew* that the observed (historical) link ratios came from a probability distribution with a mean of μ and a variance of σ^2, or "pdf (μ,σ^2)." For the stochastic analysis, we simulate new realizations of pdf (μ,σ^2).

In this case, we know the *expected* value with certainty. The uncertainty in the reserve estimate derives from the randomness of loss occurrences and loss settlements: that is, from the process risk in loss payments.

Parameter Risk: In truth, we do not know with certainty the expected value of the link ratios or the particular distribution from which they are a realization. We make two assumptions: (a) that the actual link ratios realized in the past and which will be realized in the future come from some distribution and (b)

[15] In addition, the Kreps parameter risk estimation procedure used in this paper does not work when there are only a few historical data points.

that this distribution has a particular form (such as lognormal). We estimate the parameters of the distribution from the historical values that we have observed.

This paper uses a parameter risk procedure developed by Kreps [42]. Using a Bayesian analysis, Kreps shows how to simulate from an unknown lognormal distribution based on a limited sample of data points.

The Kreps procedure is complex. To avoid repeating the mathematics of the Kreps paper [42], we simply note our choice of parameters for the Bayesian prior (readers interested in this subject should refer to that paper). Appendix C of this paper shows the equations we used to quantify the parameter risk. Appendix F of this paper provides a lay explanation of the parameter risk method, without attempting to reproduce the mathematics.

To use the Kreps procedure, one must assume a Bayesian prior distribution. Kreps uses a uniform distribution for the "translation" parameter (μ) and a distribution for the "scaling" parameter (σ) that depends on the user's prior assumptions, as reflected in a θ parameter. If the prior is uniform, then $\theta = 0$. The more conventional choice, if one is using a power-law prior, is to have $\theta = 1$. However, as Kreps pointed out to us (and as our own tests showed), "the conventional choice seems to give large values unreasonably often, given the nature of the business." He noted that $\theta = 2$ generally gives more reasonable results.[16]

Our simulations use $\theta = 2$. Even with this assumption, we found that the simulations occasionally yielded "unreasonable" results. By "unreasonable" we mean that workers compensation payments are based on statutory rules and are generally paid over the duration of a disability. Unlike some general liability claims, one rarely finds huge and unexpected lump-sum payments. Consequently, it is unreasonable to find a link ratio of say 3.0 as the factor for 15 years to 16 years of development.

[16]Kreps has also suggested that one might look for another distribution as a prior, based on our actuarial judgments about the business (private communication).

And yet, on rare occasions, that is what the simulations produce. These rare anomalies greatly affect the mean of the distribution, as well as measures of variability, like the standard deviation and the expected policyholder deficit.

Part of using actuarial judgment is to judge when the numbers being produced by mechanical formulas are not reasonable and to adjust the formulas so the results accord with insurance practice. In our case, we set a rule that if any simulated link ratio fell more than 50 standard deviations above the mean, the simulation is eliminated. In other words, we are trying to eliminate only the most extreme of the unreasonable simulations.

One might be concerned that a rule of this type would eliminate the "high" cases and thus would bias the results downwards. In fact, we found that the rule resulted in insignificant difference in the median result, or even in the 95th percentile of the distribution, and in most cases, the change in the mean was less than 1%. However, the change in the standard deviation and the expected policyholder deficit was more significant, and the results after eliminating the "outliers" are more reasonable.[17]

---

[17] An alternative procedure to quantify parameter uncertainty, which we have also tested on our data, is a procedure developed by Dickson and Zehnwirth [18]. The mean of the sample, μ, is an unbiased estimator of the mean of the distribution. If the distribution has a variance σ^2, and the sample has "n" observations, then the mean of the sample, as an estimator of the true mean of the distribution, has a variance of σ^2/n.

We want to use the sample data to simulate future realizations of the link ratios. The distribution from which these link ratios derive has a variance of σ^2. Furthermore, the whole distribution is "moving around" with a variance of σ^2/n. The total variance of the distribution from which we should simulate future realizations therefore has a variance of $\sigma^2 + \sigma^2/n$. The mean of this distribution is the sample mean, μ, which is an unbiased estimator of the true mean, as noted above. In sum, we must simulate from pdf $(\mu, \sigma^2 + \sigma^2/n)$, not from pdf (μ, σ^2).

Hayne [32] suggests a similar procedure: if the estimate of the μ of the lognormal is assumed to be unknown but to have a normal distribution with mean μ and variance σ'^2, then the final distribution is lognormal with parameters $(\mu, \sigma^2 + \sigma'^2)$.

Dickson and Zehnwirth [18] refer to these two distributions as the fitted curve and the predictive curve. The fitted curve is the best estimate of the probability distribution function; it does not include parameter variance. The predictive curve is the distribution function that one must use to simulate future realizations. It includes parameter variance, which reflects the uncertainty in the choice of parameters for the fitted curve. Our results using the Dickson–Zehnwirth procedure were similar to those using the Kreps [42] procedure. Consequently, we do not show the Dickson–Zehnwirth results in the text.

Shifting Distributions: The parameter risk discussed above assumes that there is a true distribution from which the observed link ratios are drawn, though we do not know this distribution. An additional source of variance is a shift in the true distribution, whether during the past historical period or during the future predictive period. For instance, the increasing involvement of attorneys in workers compensation claims during the 1980s may have contributed to the rising paid loss link ratios during this period, thus shifting the mean and perhaps also the variance of the distribution function. The change in the mix of claims from temporary total disability to permanent partial disability would similarly increase the mean and variance of the distribution (see Kaufman [40]). Conversely, the introduction of managed care in the 1990s may lead to a decrease in the mean of the paid loss link ratios and perhaps also their variance during this decade.

Mahler [46] refers to this as "shifting risk parameters." In his analysis of experience rating plan credibilities, Mahler divides the total expected claim variance into "within variance" and "between variance," and he includes the risk stemming from shifting risk parameters in the "within variance." We proceed similarly in our analysis. Following a suggestion by Mahler (private communication), we divide risk into process risk, specification risk, and parameter risk, where specification risk represents the risk of shifting risk parameters. The variance of the historical age-to-age link ratios stems from both process risk and from specification

---

Mathematically sophisticated readers may note some simplifications here, which are dealt with more fully in the Dickson and Zehnwirth paper. In particular, when we used the Dickson–Zehnwirth procedure, we assumed a lognormal prior distribution with known variance for the mean of the lognormal distribution (see Dickson and Zehnwirth [18, section 2.3, p. 4]. Dickson and Zehnwirth use normal distributions in their paper. As Zehnwirth has explained to the authors (private communication), "the predictive equation is lognormal, with a normal prior for the mean (μ) of the corresponding normal. The prior for $\exp(\mu)$, the median of the lognormal, is a lognormal. The prior for the mean of the lognormal, $\exp(\mu + 0.5 \times \sigma^2)$, is also a lognormal (scaled)."

Dickson and Zehnwirth also provide a parallel derivation for the predictive equation when the observed mean of the lognormal distribution comes from a Gamma prior with unknown variance. The predictive distribution is then a t-distribution, as shown in section 2.4 (pp. 4–5) and Appendix 2 (pp. 17–18) of Dickson and Zehnwirth [18]. See also Francis [27] for a similar comparison of normal and "t" distributions.

risk. Similarly, our quantification of future process includes both process risk and specification risk.

Tail Development

The paid loss development for 25 years is based on observed data. Workers compensation paid loss patterns extend well beyond 25 years. For each simulation, we complete the development pattern as follows:

- Given the 20 paid loss "age-to-age" link ratios from the set of stochastic simulations on the fitted lognormal curves, we fit an inverse power curve to provide the remaining "age-to-age" factors (see Sherman [52]). This fit is deterministic.
- The length of the development period is chosen (stochastically) from a uniform distribution of 30 to 70 years. The paid loss development is truncated at the stochastically selected age.

Because the simulated age-to-age link ratios in the first 20 development periods differ by accident year, the tail factors also differ by accident year.

4. INFLATION AND DISCOUNTING

We are primarily concerned with the economic values, or discounted values, of the reserves, not with undiscounted amounts. The exhibits here show results for undiscounted values in addition to discounted values, because statutory accounting requires the reporting of undiscounted reserves, and the Statement of Actuarial Opinion relates to the statutory figures. Butsic [13], however, emphasizes that his expected policyholder deficit (EPD) procedure, which is used here as one method of quantifying reserve uncertainty, is properly used only when balance sheet entries are stated on an economic basis, thereby avoiding "measurement bias." The EPD ratios are shown for the discounted values, not for undiscounted values.

Standard reserving procedures, when used to estimate discounted reserves, assume a fixed discount rate for unpaid losses.

Similarly, these procedures assume a fixed inflation rate for future loss payments during each development period that equals the inflation rate implicit in the historical age-to-age link ratios.

The treatment of inflation in this paper is more complex. Because of the long loss payment patterns, inflation strongly affects ultimate loss amounts. The effects on reserve variability depend on the manner in which inflation affects the loss amounts. For workers compensation, inflation affects medical benefits through the payment date. In about half of the U.S. jurisdictions, indemnity payments that extend beyond two years have cost of living adjustments (COLA's) that depend on inflation, so inflation affects the indemnity reserves as well.[18]

We use two methods for incorporating the effects of inflation into our simulation. One method leaves the effects of inflation implicit in the simulated link ratios. The other method segregates inflation from "real dollar" development and explicitly simulates future inflation rates. The two methods are described below.

- Unadjusted paid loss development patterns combine true development with the effects of inflation. That is to say, inflation is implicit in each paid loss age-to-age link ratio.[19] Were we to choose a single "best-estimate" link ratio for each development period, that would implicitly fix future inflation at the rate implicit in that "best-estimate" link ratio. Since we stochastically

---

[18]On the effects of inflation through the "payment date" versus through the "accident date," see Butsic [11], and the discussion by Balcarek.

The statutory rules for cost of living adjustments for indemnity benefits vary greatly by state. Some states have no COLA adjustments. Among the states which do have COLA's, most apply them only to disabilities extending beyond a certain time period, such as two years. In addition, many of these states cap the COLA's at specific levels, such as 5% per annum.

Properly quantifying the effect of the COLA adjustments on workers compensation indemnity reserve indications requires extensive work. For this paper we applied the stochastic inflation model to medical benefits only, where a single index can be used countrywide.

[19]For instance, the link ratio from 12 to 24 months equals the cumulative paid losses at 24 months divided by the cumulative paid losses at 12 months. A higher inflation rate during this development period raises the 24 month figure compared to the 12 month figure.

simulate the link ratios for each future accident year, we have a stochastic projection of inflation rates.

The simulated link ratios are independent of the simulated interest rates, so the implicit inflation rates are also independent of the interest rates. Although this is appropriate for link ratios, it may not be reasonable for inflation rates.

- In the second method, we deal with inflation by (a) stripping out past medical inflation from the historical loss triangles, thereby converting the figures to "real dollar terms," (b) determining "age-to-age" link ratios from the deflated loss amounts, and (c) simulating future inflation patterns and building them back into the projected (future) link ratios.

  Future inflation is simulated based on an autoregressive model that links the inflation rate both with the concurrent interest rate in the future scenarios, and with the discrepancy between the previous year's inflation rate and interest rate. The procedures used for doing this are described below.

Interest Rates

A stochastic model operates by first generating either interest rates or inflation rates—generally by some type of autoregressive function—and then generating the other index by a stochastic model with a partial dependence upon the first index.[20] Numerous methods of generating future interest paths have been developed. We used two of the simpler interest rate generators: an adaptation of the Wilkie/Daykin model, which has been used by the British Solvency Working Party, and the Cox, Ingersoll, Ross (CIR) model, which is used by many financial analysts in the United States. The generators produced comparable results. We describe the equations and results for the Cox, Ingersoll, Ross interest rate generator, which we have used for most of

[20]See Wilkie [58], Daykin, Pentikainen, and Pesonen [17], and the summary and discussion by Francis [28]. For an application to workers compensation reinsurance commutations, see Blumsohn [7].

our simulations. The procedures for the Wilkie/Daykin model are described in the previous version of this paper [34].

We begin with interest rates, simulating short rates for the CIR model, and then we simulate medical inflation rates.

The model begins by postulating a continuous process for interest rates. CIR decomposes the change in the short rate over an instantaneous period of time into a mean-reverting deterministic part and a Brownian motion stochastic part that is proportional to the square root of the current interest rate. That is

$$\partial r = a(b - r)\partial t + \sigma\sqrt{r}\partial Z,$$

where a is the mean-reverting parameter, b is the long-term average interest rate, σ is the annual volatility of the interest rates, and ∂Z is a standard Wiener process.[21] For our runs of the interest rate generator, we used parameters of

- $a = 0.2339$,
- $b = 0.050$,
- $\sigma = 0.0854$.

As a continuous time interest rate process, the CIR model has a "self-reflecting barrier" at $r = 0$. Interest rates cannot become negative, since if the interest rate process ever touches the line $r = 0$, the volatility is zero at that point and the interest rate reverts toward $a \times b$. In addition, CIR model provides for greater volatility as the interest rate becomes larger, which accords with our expectations about interest rate movements.

To run the continuous time CIR model in our simulation, we used monthly increments, with $a = 0.2339/12 = 0.0195$ and with $\sigma = 0.0854/(\sqrt{12}) = 0.0249$.

Some investment analysts concerned with short term bond options dislike equilibrium models, like the CIR model or the

[21] For an introduction to the CIR interest rate process, see Hull [38, Chapter 21].

Wilkie/Daykin model, that do not reproduce the current yield curve. Various arbitrage-free models have been proposed for securities trading operations that depend on interest rate expectations. For long-term dynamic financial analyses—like the quantification of uncertainty in loss reserves—equilibrium models seem satisfactory, and their parsimony perhaps make them preferable.

Inflation Rates

As noted above, there are two methods for dealing with inflation. Traditional reserving methods assume a continuation of the inflation rates implicit in the historical age-to-age link ratios. This procedure takes no account (i) of the autocorrelation in inflation rates or (ii) of the partial correlation with interest rates.

For the analysis in this paper, we strip inflation out of the historical age-to-age paid loss link ratios, and we stochastically simulate future inflation rates.

If we desired to simulate future inflation independently of future interest rates, we might use a procedure analogous to the autoregressive interest rate model, such as

$$\begin{aligned}\text{inflation rate} = {} & \text{average inflation rate} \\ & + \beta^{*}(\text{last year's inflation rate} - \text{average inflation rate}) \\ & + \text{an error term.}\end{aligned}$$

Similarly, one could use a formula analogous to the CIR model for inflation rates. The parameters in each model would differ, of course, such as the average rates, the β coefficient, the form of the error term, the volatility parameter, and the starting value.

The stochastic inflation rate path would be independent of the stochastic interest rate path, even over the long term. Since interest rates and inflation rates are in fact correlated, the resulting scenario set would have many unrealistic elements.

Instead, we construct the autocorrelated model to include the current interest rate. There are no "standard" models for the dual generation of interest rates and inflation rates. We have used a model developed by Kreps, namely:

$$\text{Inflation}_t = c + d^*(\text{inflation}_{t-1}) - e^*(\text{Interest rate}_{t-1}) + f^*(\text{interest rate}_t) + \text{error}(t).$$

The fitted parameters are:

$$c = 1.33\%, \qquad d = 0.546, \qquad e = 0.264, \qquad f = 0.484.$$

The error term is normal, with a mean of zero and a standard deviation of 1.83%.

Inflation and Loss Development

To separately account for the effects of inflation on reserve development, we make the following adjustments to the data:[22]

- We convert the paid medical losses to real dollar amounts, using the medical component of the CPI. We then determine paid loss age-to-age link ratios from the deflated figures, we fit lognormal curves to each column of historical link ratios, and we run the simulation 10,000 times to determine the future link ratios.
- For each simulation, we stochastically generate a future interest rate path and a future inflation rate path, using the models described above.
- For each set of simulated link ratios and future inflation rates, we determine two required reserve amounts:
    1. The undiscounted (full value) reserves, using the link ratio and the inflation rate scenarios, and

[22] For a similar adjustment to reserving point estimates, see Richards [50, p. 387]: "These steps are designed to factor out the effects of inflation from historical loss data prior to forecasting, forecast the reserve using the current methodology and then replace the effects of inflation including an assumption of future inflation."

TABLE 1

INFLATION IMPLICIT IN LINK RATIOS; UNCORRELATED ACCIDENT YEARS

	Average Reserve Amount	Standard Deviation of Reserve	95th Percentile of Reserve	5th Percentile of Reserve	Capital Needed for 1% EPD Ratio
Undiscounted	100.0	14.5	125.0	80.4	—
Discounted	57.4	6.4	68.4	47.3	6.4

2. The discounted reserves, using the link ratio, inflation rate, and interest rate scenarios.

5. RESULTS

Table 1 shows results when inflation rates are not simulated separately; rather, the effects of future inflation are implicit in the simulated link ratios. Table 2 shows the results when inflation is removed from the historical link ratios and independently generated inflation rate paths are used for future years.

Exhibits 1 and 2 show the shapes of the probability distributions for the discounted and the undiscounted reserves. Exhibit 1, like Table 1, has no separate simulation of future inflation rates. Rather, the inflation implicit in the historical link ratios is presumed to continue into the future. Exhibit 2, like Table 2, uses the separate stochastic model for future inflation rates, as discussed above.

In Table 1, the average full value reserves are normalized to \$100 million to facilitate the interpretation of the figures. The average discounted reserves are \$57.4 million, with a standard deviation of \$6.4 million. The 5th percentile of the distribution of required reserves is \$47.3 million, and the 95th percentile is \$68.4 million. To achieve a 1% EPD ratio, capital of \$6.4 million is needed, above the \$57.4 million of assets needed to support the expected (discounted) loss payments.

TABLE 2

INDEPENDENTLY GENERATED INFLATION RATES;
UNCORRELATED ACCIDENT YEARS

	Average Reserve Amount	Standard Deviation of Reserve	95th Percentile of Reserve	5th Percentile of Reserve	Capital Needed for 1% EPD Ratio
Undiscounted	84.2	14.4	109.9	64.9	—
Discounted	49.8	4.6	57.6	42.6	4.1

Table 2 shows the corresponding figures when the future inflation rates are stochastically generated. The following items are noteworthy:

- The average discounted reserve decreases to \$49.8 million, with a standard deviation of \$4.6 million. High inflation scenarios, which strongly affect medium and long duration loss payments, have a lesser effect on discounted reserves. Moreover, high long-term inflation rates are often partially offset by high long-term interest rates.
- Nominal losses decrease to \$84.2 million, since we are projecting lower future inflation than is implicit in the historical loss triangle.
- The capital needed to achieve a 1% EPD ratio declines from \$6.4 million to \$4.1 million. The rationale is similar to that mentioned in the preceding paragraph. The high inflation scenarios that increase the capital requirement when inflation is implicit in future link ratios have a dampened effect when future inflation rates are linked to future interest rates.[23]

[23] When reserves are fully discounted, interest rate risk rises. This is particularly true for lines of business that are inflation sensitive, where the ultimate value of the loss payments depends on inflation up to the payment date. When inflation accelerates, nominal loss payments increase and market values of bonds decrease (if interest rates are linked to inflation rates). For further discussion of the capital required for interest rate risk, as well as the interplay with the capital required for reserving risk, see Hodes and Feldblum [35].

TABLE 3

INDEPENDENTLY GENERATED INFLATION RATES; CORRELATED ACCIDENT YEARS

	Average Reserve Amount	Standard Deviation of Reserve	95th Percentile of Reserve	5th Percentile of Reserve	Capital Needed for 1% EPD Ratio
Undiscounted	84.5	18.4	119.2	62.9	—
Discounted	49.8	5.7	59.9	41.7	6.5

For Tables 1 and 2, the age-to-age link ratios are separately simulated for each future accident year. In other words, from each column of historical link ratios, we fitted a lognormal curve from which to simulate the future link ratios. We did not simulate a single link ratio which would be applied to all accident years that had not yet reached that stage of development. Rather, we separately simulated link ratios for each future accident year.

This assumes that the development in each accident year is independent of the development in other accident years at the same maturity. To test the results if the opposite assumption is made—namely, that the development at any given maturity is the same in all future accident years—we simulated a single age-to-age link ratio for each maturity and used it for all accident years. The results are shown below in Table 3. Since this procedure assumes perfect correlation among accident years, high or low link ratios are repeated in all accident years and the reserve uncertainty increases.

Uncertainty and Discounting

A common view is that discounted reserves are simply smaller than undiscounted reserves, but they exhibit the same degree of variability. This is not correct. As Exhibits 1 and 2 show, the probability distributions for undiscounted reserves are wide, whereas the corresponding probability distributions for discounted reserves are far more compact. The rationale for this is two-fold.

- First, much of the reserve variability comes from uncertainty in distant tail factors, which strongly wag estimates of undiscounted reserves but have less effect on discounted reserve estimates.
- Second, when using stochastic inflation rate paths with strong autocorrelation, much additional reserve variability results from high or low inflation scenarios. For discounted reserves, part of this variability is offset by corresponding high and low interest rate scenarios.

The magnitude of the difference between the two distributions depends on the parameters of the interest rate generator and the stochastic inflation process. The greater the volatility of interest rate and inflation rates, and the stronger the correlation between them, the greater the difference between the nominal and present value distributions.

Because statutory accounting mandates that insurers hold undiscounted reserves, we have shown results both for discounted reserves and for undiscounted (or "nominal") reserves in the exhibits. In particular, the means, standard deviations, and percentiles of the distributions are shown for both nominal and discounted reserves, though the capital requirements based on the expected policyholder deficit of 1% are applicable only to the discounted values. (See the discussion below in the text and in Appendix B regarding the expected policyholder deficit.) Moreover, the difference between the discounted and undiscounted reserve amounts is the "implicit interest margin" in the reserves, which is important for assessing the implications of the reserve uncertainty on the financial position of the insurance company.

Assumptions and Results

It is instructive to consider the relative reserve variability resulting from the different assumptions. Specifically, will the independent generation of future inflation rate paths increase or decrease reserve variability?

We begin with the results for our "base case," and we consider how each change in assumptions affects the estimated uncertainty. The base case assumes that:

- Link ratios are generated stochastically, incorporating both process risk and parameter risk.
- For the discounted reserves, autocorrelated interest rate paths are generated stochastically.
- Future inflation rates are not generated independently. Rather, the inflation embedded in the observed link ratios is assumed (implicitly) to continue into the future.

For *nominal* reserves, the independent generation of stochastic inflation rate paths adds an additional element of variability to the reserves. Accordingly, the standard deviation of the nominal reserve distribution is higher when inflation rates are independently generated. The coefficient of variation for the base case (Table 1) is about 14.5%, whereas it is about 17.1% when inflation rates are independently generated (Table 2).

For *discounted* reserves, the opposite is true. In the base case, the reserve discount rates are generated independently of the link ratios, in which the inflation rates are implicitly embedded, so reserve variability is high. When inflation rates are generated independently of the link ratios, they are correlated with the stochastically generated interest rates, and their effects partially offset each other, thereby dampening the reserve variability. For the discounted reserves, the capital ratio required for a 1% expected policyholder deficit ratio is 11.1% for the base case, while it is 9.2% when inflation rates are independently generated.

The implications of these results are important for the capital structure of a workers compensation insurer. In our illustration, the average undiscounted required reserves developed from a traditional reserve analysis, with no independent generation of future inflation rates, is \$100 million. Most companies use tabular discounts for lifetime pension indemnity benefits, and some companies do not fully account for inflation of medical benefits. For most companies, the held statutory reserves would be between \$80 million and \$90 million.[24]

The average discounted required reserve is \$49.8 million. The implicit interest margin in the statutory reserves is about \$25 million to \$40 million.[25]

The capital required to achieve a 1% EPD ratio because of reserve uncertainty is about \$4.1 million, which is less than a fifth of the implicit interest margin in the statutory reserves. In other words, most insurers would need no additional capital to support the uncertainty in their workers compensation reserves.[26]

A common view is that workers compensation reserve estimates are highly uncertain, because of the long payment lags and because of the unlimited nature of the insurance contract form. This uncertainty creates a great need for capital to hedge against unexpected reserve development.

In fact, the risks in workers compensation lie elsewhere. There is great *underwriting* uncertainty in workers compensation, and regulatory constraints on the pricing and marketing of this line of business have disrupted markets and contributed to the financial distress of several carriers. But once the policy term has expired and the accidents have occurred, less uncertainty remains. The difference between the economic value of the reserves and the reported (statutory) reserves, or the implicit interest margin, is generally greater than the capital needed to hedge against reserve uncertainty.[27]

---

[24]The *Proceedings* reviewers have pointed out that some companies do not carry full value reserves, even on the statutory blank. For such companies, the held statutory reserves would be lower.

[25]The size of the implicit interest margin depends on the prevailing interest rates; it is larger in the high interest rate environments of the 1980's and smaller in the low interest rate environments of the 1990's.

[26]As noted earlier, some additional capital would be needed to support the default risks, market risks, and interest rate risks on the assets supporting the reserves.

[27]The implications for capital allocation to lines of business are important; for a full discussion, see Hodes, et al., [36]. For companies that carry adequate statutory reserves,

The EPD Yardstick

Several elements of our analysis may require further explanation. The following sections provide brief qualitative discussions of certain aspects of the analysis. The appendices provide more complete quantitative descriptions, as well as full documentation of our procedures.

As a yardstick to measure reserve uncertainty, we use the "expected policyholder deficit" (EPD) ratio developed by Butsic [13] for solvency applications. The EPD ratio allows us to:

- Compare reserve uncertainty across different lines of business,
- Compare reserve uncertainty with either explicit margins in held reserves or with the "implicit interest margins" in undiscounted reserves,
- Quantify the effects of various factors (such as the presumed variability of future inflation rates or the premium sensitivity on loss sensitive contracts) on reserve uncertainty, and
- Translate actuarial concepts of reserve uncertainty into more established measures of financial solidity.[28]

The Expected Policyholder Deficit

Were there no uncertainty in the future loss payments, the insurer need hold funds just equal to the reserve amount to meet its loss obligations. Since future loss payments are not certain, funds equal to the expected loss amount sometimes will suffice to meet future obligations, and sometimes they will fall short. The "policyholder deficit" is this shortfall.

---

the capital needed to support compensation reserves is negative, though positive capital is needed to support workers compensation underwriting operations. This is in contrast to the statutory accounting procedures used in many surplus allocation procedures in insurance pricing models. See, for instance, Feldblum [22], and particularly the Cummins/NCCI dispute there on the proper funding of the underwriting loss in the internal rate of return model.

[28]For a full discussion of the use of the EPD yardstick for measuring uncertainty, see Appendix B.

When the present value of the future loss obligations is less than the funds held by the insurance company to meet these obligations, the policyholder deficit is zero. When the present value of the future loss obligations is greater than the funds held, the policyholder deficit is the difference between the two. The expected policyholder deficit (EPD) is the average deficit over all scenarios, weighted by the probability of each scenario. In the analysis here, the expected deficit is the average deficit over all simulations, each of which is weighted equally.

Let us illustrate with the workers compensation reserve simulations in this paper. Suppose first that the company holds no capital besides the funds supporting the reserves. For the discounted analysis, the average reserve amount is \$49.8 million (see Table 2). About half the simulations give reserve amounts less than \$49.8 million. In these cases, the deficit is zero. The remaining simulations give reserve amounts greater than \$49.8 million; these give positive deficits. The average deficit over all 10,000 simulations is the EPD. The "EPD ratio" is the ratio of the EPD to the expected losses, which are \$49.8 million in this case.

Clearly, if the probability distribution of the needed reserve amounts is "compact," or "tight," then the EPD ratio is relatively low. Conversely, if the probability distribution of the needed reserve amounts is "diffuse" —that is, if there is much uncertainty in the loss reserves—then the EPD ratio is relatively great.

We have two ways of proceeding:

- We could assume that the company holds no assets besides those needed to support the expected loss obligations, and compare EPD ratios for different lines of business or operating environments.
- We may "fix" the EPD ratio at a desired level of financial solidity and determine how much capital is needed to achieve this EPD ratio.

The second approach translates EPD ratios into capital amounts, so we follow this method. We use a 1% EPD ratio as our benchmark, since Butsic notes that the reserving risk charges in the NAIC property-casualty insurance company risk-based capital formula are of similar magnitude as the charges needed for a 1% EPD ratio.[29]

Suppose the desired EPD ratio is 1%. If the reserve distribution were extremely compact, then even if the insurer held no capital beyond that required to fund the expected loss payments, the EPD ratio might be 1% or less. If the reserve distribution is more diffuse, then the insurer must hold additional capital to achieve an EPD ratio of 1%. The greater the reserve uncertainty, the greater the required capital.

Trends and Correlations

Two additional issues are of importance to reserving actuaries: correlations among link ratios and trends in link ratios.

- *Correlations*: The simulation procedure assumes that a particular link ratio is independent of the other link ratios in the same row. If the link ratios are not independent, the results may be overstated or understated.

  For instance, suppose that accident year 1988 shows a high paid loss link ratio from 24 to 36 months. Should one expect a higher than average link ratio or a lower than average link ratio from 36 to 48 months?

  The answer depends on the cause of the high 24 to 36 month link ratio. If it is caused by a speeding up of the payment pattern, but the ultimate loss amount has not changed, then one should expect a lower than average link ratio from 36 to 48 months. If it is caused by higher ultimate loss amounts

[29]For private solvency monitoring analyses, Butsic suggests that a higher ratio may be appropriate, such as 0.1%; see Butsic [13].

(e.g., because of lengthening durations of disability for indemnity benefits or because of greater utilization of medical services), then one should expect a higher than average link ratio from 36 to 48 months.[30]

- *Trends*: Our procedure uses unweighted averages of the link ratios in each column. During the 1980s, industry-wide paid loss link ratios showed strong upward trends, though this trend ceased in the early 1990s.[31] How would the recognition of such trends affect the variability of the reserves estimates as discussed here?

These two issues are related. First, the observed correlations among the columns of link ratios in the historical data result from the trends in these link ratios. When the trends are removed, the correlations largely disappear. Second, the trends affect the proper reserve estimate. The reserving actuary must investigate these trends and their causes, and then project their likely effect on future loss payments. That is not our interest in this paper. Rather, we ask: "What is the inherent variability in the reserve estimation process itself?"[32]

---

[30]For further explanation, see the discussion by H. G. White to Bornhuetter and Ferguson [8], as well as Brosius [10]. Compare also Holmberg [37, p. 254]:

> *There are different reasons we might expect development at different stages to be correlated. For instance, if unusually high loss development in one period were the result of accelerated reporting, subsequent development would be lower than average as the losses that would ordinarily be reported in those later periods would have already been reported. In this instance, correlation between one stage and subsequent stages would be negative. Positive correlation would occur if there were a tendency for weaker-than-average initial reserving to be corrected over a period of several years. In that case, an unusually high degree of development in one period would be a warning of more to come.*

Holmberg looks at incurred loss development. (To circumvent the effects of company case reserving practices on the variability of reserve estimates, we use paid loss development in this analysis.) Hayne [33] also discusses the possible correlations in the reserve estimation procedure, though he deals with them in a different fashion.

[31]See Feldblum, [25, section 7, and the references cited therein].

[32]To incorporate trends in this model, one would restate ("detrend") each column of historical link ratios to the current calendar year level before fitting these observed link ratios to a lognormal curve (see Berquist and Sherman [5]).

Let us take each of these issues in turn.

- *Correlations among columns*: Suppose one has two columns of observed link ratios, each from accident years 1971 through 1993, from 12 to 24 months and from 24 to 36 months, and that they are *not* correlated. We then apply a strong upward trend to both columns. That is, we increase the accident year 1972 link ratios by 1.02, the accident year 1973 link ratios by $(1.02)^2$, the accident year 1974 link ratios by $(1.02)^3$, and so forth.

  The resulting link ratio show a strong positive correlation. Indeed, we observe such a correlation in the historical link ratios used in our simulation. But if we remove the trend, the correlation disappears.

  This trend was caused primarily by the increasing liberalization of workers compensation benefit systems between the mid-1970s and the late 1980s. This liberalization, along with its associated effects (increasing paid loss link ratios, statewide rate inadequacies, growth of involuntary markets) ceased by the early 1990s, and has even reversed in many jurisdictions. The advent of managed care, along with workers compensation reforms in several state legislatures, may lead to further reduction in paid loss link ratios.

- *Correlations among years:* The chain ladder reserving technique involves "squaring the triangle." From each column in the observed triangle of age-to-age link ratio, we estimate a future link ratio, which is applied to all cells in that column of the triangle of future link ratios. When determining point estimates of indicated reserves, it is appropriate to use the same projected "best estimate" link ratio for all future accident years (i.e., for all the remaining cells in each column).

  The analysis here is different. We are not simulating a reserve estimate, or a reserve indication. Rather, we are simulating the potential future realization of loss development. In any simulation, the actual development will differ by accident year.

This is particularly important when studying reserve uncertainty. Our concern is not simply to quantify the expected development but to measure the variability of this development. Thus, when performing a stochastic analysis to determine reserve variability, it is proper to separately simulate the projected link ratios for each future accident year.[33]

For instance, suppose we have accident years 1970 through 1994, valued through December 31, 1994, and we are simulating the link ratios for the 48 months to 60 months development period. We need projected link ratios for accident years 1991, 1992, 1993, and 1994. We perform the stochastic simulation using the predictive curve four times, to give independently simulated link ratios for these four accident years. Similarly, once we have the projected link ratios, we fit inverse power curves to each accident year, to generate separate tail factors for each year.

Practicing actuaries may wonder about the materiality of this issue: does the increase in simulations increase or decrease the resultant reserve variability, and how large is this increase or decrease?

Consider the difference between (i) simulating once and using the same projected link ratio for all four accident years and (ii) simulating four times, once for each future accident year. The more separate (independent) pieces there are in each simulation of the total reserve requirements (as in the latter procedure), the tighter will be the distribution of the total reserve requirement. The fewer separate pieces there are in each simulation of the total reserve requirement (as in the former procedure), the greater will be the effect of individual

[33]The statement in the text is true if the variability stems from process risk. For the parameter risk component of the variability, one might argue that it is more proper to simulate once and to use the same factor each future accident year.

"outlying" factors, and the distribution of the total reserve requirement will be more widely spread.

Thus, the use of separate simulations decreases the estimated reserve variability. The effect is small, though, since there are many independent development periods in each simulation. The figures are shown in Table 3.

- *Trends*: Yes, there were trends, at least in the 1980s. Moreover, there are multiple reserving methods. The mark of the skilled actuary is to take the various reserve indications and the manifold causes for discrepancies among them and to project an estimate as close as possible to the true, unfolding loss payments.

  In our analysis, we have used the full column of observed link ratios to fit the lognormal curve, and then we have compared the simulated loss payments with their averages. Had we incorporated the "trends," and had we ignored old link ratios (because they are not relevant for today's environment), we might have produced tighter reserve distributions.

  If one places faith in the skills of reserving actuaries, then the use of a solitary reserving method overstates the uncertainty of the reserving process. Suppose the simulation produces actual loss payments considerably higher than the reserve estimate. Oftentimes, the experienced actuary would have noted signs that the paid loss estimate was underestimating the actual reserve need, and that other methods were giving higher indications. By combining the indications from several methods, the actuary might come closer to the actual reserve need, thereby reducing the uncertainty in the estimates.

Perhaps uncertainty can be reduced by actuarial judgments of trends and by actuarial weighing of various indications. The concern of this paper is more fundamental: even in rote applications of basic reserving techniques, how much uncertainty is produced by the fluctuations in loss data?

Federal Income Taxes

We have ignored income taxes, since their effect is uniform for most scenarios. Federal income taxes reduce the potential profits of the insurance company, but they also reduce the potential losses.

Suppose we determined that if there were no income taxes, an insurer has a 5% chance of exhausting its surplus because of the variability in loss reserves. Then with an income tax rate of 35%, the chance of exhausting its surplus is less than 5% for this insurer.

In effect, the U.S. government acts as a pro-rata reinsurer for all the company's business. It takes 35% of the revenue, and it pays 35% of losses plus expenses.

The risk on any particular insurance contract is not affected by federal income taxes. Rather, the contract is reduced in size: all revenues and expenditures are multiplied by 65%. Similarly, the variability in the loss reserves is not affected by federal income taxes. Rather, the reserves are simply reduced in size by a factor of 65%. Yardsticks such as percentiles or the coefficient of variation are not affected by federal income taxes.

Yardsticks such as the probability of ruin and the expected policyholder deficit ratio, however, relate reserve variability to the company's capital. The capital is on a post-tax basis, so the federal income tax rate is relevant. In addition, since the expected losses are on the company's books, taxes have already been paid on the assumption that these will be the ultimate losses. This means that the company's surplus reflects taxes at the expected level of losses. If one needs a certain amount of capital to pass a given "probability of ruin" test or a given "EPD ratio" test when one does not take into account federal income taxes, then one needs only 65% as much capital to pass the same test if one *does* take into account federal income taxes.[34]

[34] Similarly, Butsic [13] recommended that the charges in the NAIC risk-based capital formula be reduced for the offsetting effects of federal income tax recoupments, though his proposal was never implemented.

Because the potential federal income tax returns are affected by a host of factors, including the amount of taxes paid in the past three years and the amount of taxable income in the insurance enterprise's other operations, we have stated all our results on a pre-tax basis. For comparative analyses, a pre-tax basis is sufficient, such as for comparing reserve uncertainty among lines of business or among different policy forms. Practicing actuaries measuring capital requirements, however, should convert the results to a post-tax basis, using the particular tax situation of their own company or client.

6. STATUTORY BENEFITS

For the insurer from which these data were drawn, workers compensation reserves have about the same average payment lags as general liability GL reserves. There is great uncertainty in this company's GL reserves, as an equivalent analysis to that shown in this paper would show.[35] The causes of the GL reserve uncertainty illuminate the reasons for the compactness of the workers compensation reserve distribution.

- *IBNR Emergence*: Many GL claims are not reported to the insurer until years after the accident. For toxic tort and environmental impairment exposures, claims are still being reported decades after the exposure period (see, for instance, ISO [39] or Simpson, Smith, and Babbitt [53]). In contrast, most workers compensation accidents are known to the em-

[35] A full actuarial study of reserve uncertainty would apply the techniques used in this paper to all lines of business and compare the reserve distributions, EPD ratios, or capital requirements among them. The analysis must take into account the factors specific to each line that affect reserve fluctuations. For instance, just as we examine loss sensitive contracts for workers compensation, we must examine latent injury claims, such as those stemming from asbestos and pollution exposures, for general liability. For lines of business like general liability, results about reserve uncertainty can not always be generalized, since company practices vary so widely: some companies write premises and operations coverage for retail establishments, while other companies insure large manufacturing concerns; some companies are inundated by asbestos claims, while other companies have few of these cases. The extent of such analysis, of course, puts it beyond the scope of this paper.

ployer within days of the accident, and insurance companies are notified soon thereafter.

- *Claim Payment Patterns*: General liability losses depend upon judicial decisions and jury awards. Ultimate costs may not be known until years after the claim has been reported to the insurer. Even cases settled out-of-court are often settled "on the courthouse steps," after pre-trial discovery and litigation efforts have provided good indications of the expected judicial outcome.

Workers compensation benefits, in contrast, are fixed by statute, both in magnitude and in timing. The benefits may be determined either by agreement between the insurer and the injured worker, or by a workers compensation hearing officer. The major uncertainty in indemnity benefits is the duration of disability on non-permanent cases and the mortality rates on permanent cases. For sufficiently large blocks of business, both of these have relatively compact distributions. The major uncertainty for medical benefits is the rate of inflation and the extent of utilization of medical services. Over a large enough block of business, these risks also have relatively compact distributions, particularly when reserves are discounted.[36]

Butsic [12, p. 179], summarizes this view as follows:

> *For example, Workers Compensation reserves should have a lower risk than Other Liability reserves, even though the average payment durations are about the same, because Workers Compensation loss reserves consist partly of fixed, more predictable, life pension benefits.*

[36] Changes in the workers compensation system may either increase or decrease the reserve uncertainty. For instance, the advent of managed care may increase the uncertainty of ultimate loss payments, since the efficacy of managed care is not well known. It is equally possible that managed care will decrease reserve uncertainty, since the medical benefits may become easier to estimate. Our analysis partially incorporates this "specification risk" (to use Mahler's [46] term) in the process risk of the lognormal distribution (see the discussion above).

This paper provides the statistical support for the workers compensation half of this citation from Butsic.

7. CONCLUSIONS

Casualty actuaries have developed numerous methods of estimating required loss reserves. But reserves are uncertain, and actuaries are now being asked to quantify the uncertainty inherent in the reserve estimates.

Many past attempts to address this subject have foundered on one of two shoals. Some attempts are silver vessels of pure theory: loss frequencies are simulated by Poisson functions, loss severity is simulated by lognormal distributions, inflation is simulated by Brownian movements, and the results are much prized by hypothetical companies. Other attempts are steel vessels of actual experience: actual reserve changes, taken from financial statements, reveal how companies have acted in the past, though they offer imperfect clues about the uncertainties inherent in the reserve estimation process itself.

This paper glides between the shoals. Loss reserve uncertainty must be tied to the line of business. The uncertainty in workers compensation reserves is different from the uncertainty in general liability reserves even as it is different from the uncertainty in life insurance or annuity reserves. We begin with extensive data—twenty five years of experience from the nation's premier workers compensation carrier.

These data allow the actuary to develop reserve indications. Our concerns in this paper are different. We fit these data to families of curves to develop probability distributions of required reserves. The power of stochastic simulation techniques enables us to develop thousands of potential outcomes that are solidly rooted in the empirical data.

The analysis shows that workers compensation reserves, when valued on a discounted basis, have a highly compact distribution. To measure uncertainty, we use the "expected policyholder deficit" (EPD) ratio. For workers compensation, the amount of capital needed to achieve a 1% EPD ratio is only a small fraction of the "implicit interest margin" in the reserves themselves.

The vicissitudes of inflation are a major cause of workers compensation reserve fluctuations, and changes in interest rates strongly influence discounted values. This paper uses stochastically generated interest rates and inflation rates to model the reserve uncertainty.

The combination of rigorous actuarial theory with an extensive empirical database enables us to examine the uncertainty in the reserves themselves. Similar analyses should be performed for other lines of business, such as automobile insurance or general liability. Comparisons among the lines, as well as comparisons of reserve uncertainty with underwriting risks and with asset risks, would allow us to exchange preconceived notions with well-supported facts.

REFERENCES

[1] American Academy of Actuaries, "Position Statement on Insurer Solvency," *Actuarial Update*, September 1992.

[2] Barth, Michael, "Risk-Based Capital Results for the Property-Casualty Industry," *NAIC Research Quarterly* II, I, January 1996, pp. 17–31.

[3] Beard, R. E., T. Pentikainen, and E. Pesonen, *Risk Theory: The Stochastic Basis of Insurance*, Third Edition, London: Chapman and Hall, 1984.

[4] Bender, Robert K., "Aggregate Retrospective Premium Ratio as a Function of the Aggregate Incurred Loss Ratio," *PCAS* LXXXI, 1994, pp. 36–74; discussion by Howard C. Mahler, pp. 75–90.

[5] Berquist, James R., and Richard E. Sherman, "Loss Reserve Adequacy Testing: A Comprehensive Systematic Approach," *PCAS* LXIV, 1977, pp. 123–184; discussion by J. O. Thorne, *PCAS* LXV, 1978, pp. 10–34.

[6] Berry, Charles H., "A Method for Setting Retro Reserves," *PCAS* LXVII, 1980, pp. 226–238.

[7] Blumsohn, Gary, "Levels of Determinism in Workers' Compensation Reinsurance Commutations," *PCAS* LXXXVI, 1999.

[8] Bornhuetter, Ronald L., and Ronald E. Ferguson, "The Actuary and IBNR," *PCAS* LX, 1973, pp. 165–168.

[9] Bowers, Newton L., Jr., Hans U. Gerber, James C. Hickman, Donald A. Jones, and Cecil J. Nesbitt, *Actuarial Mathematics*, Itasca, Illinois: Society of Actuaries, 1986.

[10] Brosius, J. Eric, "Loss Development Using Credibility," CAS Part 7 examination study note, December 1992.

[11] 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, pp. 51–102; discussion by Rafal J. Balcarek, pp. 103–109.

[12] Butsic, Robert P., "Determining the Proper Interest Rate for Loss Reserve Discounting: An Economic Approach," *Evaluating Insurance Company Liabilities*, Casualty Actuarial Society Discussion Paper Program, 1988, pp. 147–188.

[13] Butsic, Robert P., "Solvency Measurement for Property-Liability Risk-Based Capital Applications," *Journal of Risk and Insurance* 61, 4, December 1994, pp. 656–690.

[14] Butsic, Robert P., "The Underwriting Cycle: A Necessary Evil?," *The Actuarial Digest* 8, 2, April/May 1989.

[15] Cook, Charles F., "Trend and Loss Development Factors," *PCAS* LVII, 1970, pp. 1–26.

[16] Cummins, David J., Scott E. Harrington, and Robert W. Klein, *Cycles and Crises in Property/Casualty Insurance: Causes and Implications for Public Policy*, National Association of Insurance Commissioners, 1991.

[17] Daykin, Chris D., Teivo Pentikainen, and M. Pesonen, *Practical Risk Theory for Actuaries*, First Edition, Chapman and Hall, 1994.

[18] Dickson, David C. M., and Ben Zehnwirth, "Predictive Aggregate Claims Distributions," Research Paper No. 27, Centre for Actuarial Studies, Department of Economics, The University of Melbourne, Australia, February 1996.

[19] Feldblum, Sholom, "Author's Reply to Discussion by Stephen Philbrick of Risk Loads for Insurers," *PCAS* LXXX, 1993, pp. 371–373.

[20] Feldblum, Sholom, "Completing and Using Schedule P," *Regulation and the Casualty Actuary*, edited by Sholom Feldblum and Gregory Krohm, NAIC, 1997.

[21] Feldblum, Sholom, Discussion of Teng and Perkins: "Estimating the Premium Asset on Retrospectively Rated Policies," *PCAS* LXXXV, 1998.

[22] Feldblum, Sholom, "Pricing Insurance Policies: The Internal Rate of Return Model," Casualty Actuarial Society Part 10A Examination Study Note, May 1992.

[23] Feldblum, Sholom, "NAIC Property/Casualty Insurance Company Risk-Based Capital Requirements," *PCAS* LXXXIII, 1996, pp. 297–435.

[24] Feldblum, Sholom, "Underwriting Cycles and Business Strategies," Casualty Actuarial Society *Forum*, Spring 1990, pp. 63–132.

[25] Feldblum, Sholom, "Workers' Compensation Ratemaking," Casualty Actuarial Society Part 6 Study Note, September 1993.

[26] Fitzgibbon, Walter J., Jr., "Reserving for Retrospective Returns," *PCAS* LII, 1965, pp. 203–214.

[27] Francis, Louise A., "A Model for Combining Timing, Interest Rate, and Aggregate Loss Risk," *Valuation Issues*, Casualty Actuarial Society Discussion Paper Program, 1989, pp. 155–216.

[28] Francis, Louise A., "Modelling Asset Variability in Assessing Insurer Solvency," *Insurer Financial Solvency*, Casualty Actuarial Society Discussion Paper Program, 1992, II, pp. 585–656.

[29] Freifelder, R. L., *A Decision Theoretic Approach to Insurance Ratemaking*, Homewood, IL: Richard D. Irwin, 1976.

[30] Gillam, William R., and Richard H. Snader, "Fundamentals of Individual Risk Rating," 1992, available from the CAS.

[31] Greene, Howard W., "Retrospectively-Rated Workers Compensation Policies and Bankrupt Insureds," *Journal of Risk and Insurance* 7, 1, September 1988, pp. 52–58.

[32] Hayne, Roger, "Application of Collective Risk Theory to Estimate Variability in Loss Reserves," *PCAS* LXXVI, 1989, pp. 77–110.

[33] Hayne, Roger, "A Method to Estimate Probability Level for Loss Reserves," Casualty Actuarial Society *Forum* I, 1994, pp. 297–356.

[34] Hodes, Douglas M., Sholom Feldblum, and Gary Blumsohn, "Workers Compensation Reserve Uncertainty," Casualty Actuarial Society *Forum*, Summer 1996, pp. 61–149.

[35] Hodes, Douglas M., and Sholom Feldblum, "Interest Rate Risk and Capital Requirements for Property-Casualty Insurance Companies," *PCAS* LXXXIII, 1996.

[36] Hodes, Douglas M., Sholom Feldblum, and Tony Neghaiwi, "The Financial Modeling of Property-Casualty Insurance Companies," *North American Actuarial Journal*, July 1999, pp. 41–69.

[37] Holmberg, Randall D., "Correlation and the Measurement of Loss Reserve Variability," Casualty Actuarial Society *Forum* I, 1994, pp. 247–278.

[38] Hull, John C., *Options, Futures, and Other Derivatives*, Fourth Edition, Englewood Cliffs, NJ: Prentice Hall, 2000.

[39] Insurance Services Office, *Superfund and the Insurance Issues Surrounding Abandoned Hazardous Waste Sites*, December 1995.

[40] Kaufman, Allan M., "Evaluating Workers Compensation Trends Using Data by Type of Disability," *Trends*, Casualty Actuarial Society Discussion Paper Program, 1990, pp. 425–461.

[41] Kaufman, Allan M., and Elise C. Liebers, "NAIC Risk Based Capital Efforts in 1990–91," *Insurer Financial Solvency*, Casualty Actuarial Society Discussion Paper Program, I, 1992, pp. 123–178.

[42] Kreps, Rodney, "Parameter Uncertainty in (Log)Normal Distributions," *PCAS* LXXXIV, 1997, pp. 553–580.

[43] Lee, Yoong-Sin, "The Mathematics of Excess of Loss Coverages and Retrospective Rating—A Graphical Approach," *PCAS* LXXV, 1988, pp. 49–78.

[44] Lowe, Stephen P., "Report on Reserve and Underwriting Risk Factors," Casualty Actuarial Society *Forum*, Summer 1993, pp. 105–171.

[45] Mahler, Howard C., Discussion of Bender "Aggregate Retrospective Premium Ratio as a Function of the Aggregate Incurred Loss Ratio," *PCAS* LXXXI, 1994, pp. 75–90.

[46] Mahler, Howard C., "An Example of Credibility and Shifting Risk Parameters," *PCAS* LXXVII, 1990, pp. 225–282.

[47] Meyers, Glenn G., "Risk Theoretic Issues in Loss Reserving: The Case of Workers Compensation Pension Reserves," *PCAS* LXXVI, 1989, pp. 171–192.

[48] Miccolis, Robert S., "On the Theory of Increased Limits and Excess of Loss Pricing," *PCAS* LXIV, 1977, pp. 27–59.

[49] Philbrick, Stephen, "Accounting for Risk Margins," Casualty Actuarial Society *Forum* I, 1994, pp. 1–90.

[50] Richards, William F., "Evaluating the Impact of Inflation on Loss Reserves," *Inflation Implications for Property/Casualty Insurers*, Casualty Actuarial Society Discussion Paper Program, 1981, pp. 384–400.

[51] Ryan, Kevin M., and Richard I. Fein, "A Forecast for Workers Compensation," *NCCI Digest* III, Issue IV, December 1988, pp. 43–50.

[52] Sherman, Richard, "Extrapolating, Smoothing, and Interpolating Development Factors," *PCAS* LXXI, 1984, pp. 122–192; discussion by Stephen Lowe and David F. Mohrman, *PCAS* LXXII, 1985, p. 182; author's reply to the discussion, p. 190.

[53] Simpson, Eric M., W. Dolson Smith, and Cynthia S. Babbitt, "P/C Industry Begins to Face Environmental and Asbestos Liabilities," *BestWeek*, January 1996.

[54] Simon, LeRoy J., "The 1965 Table M," *PCAS* LII, 1965, pp. 1–45.

[55] Skurnick, David, "The California Table L," *PCAS* LXI, 1974, pp. 117–140.

[56] Stanard, James N., "A Simulation Test of Prediction Errors of Loss Reserve Estimation Techniques," *PCAS* LXXII, 1985, pp. 124–153.

[57] Teng, Michael T. S., and Miriam Perkins, "Estimating the Premium Asset on Retrospectively Rated Policies," *PCAS* LXXXIII, 1996, pp. 611–647.

[58] Wilkie, A. D., "A Stochastic Investment Model for Actuarial Use," *Transactions of the Faculty of Actuaries*, Vol. 39, 1986, p. 341.

EXHIBIT 1

Distributions of Reserves
Without Inflation Adjustments

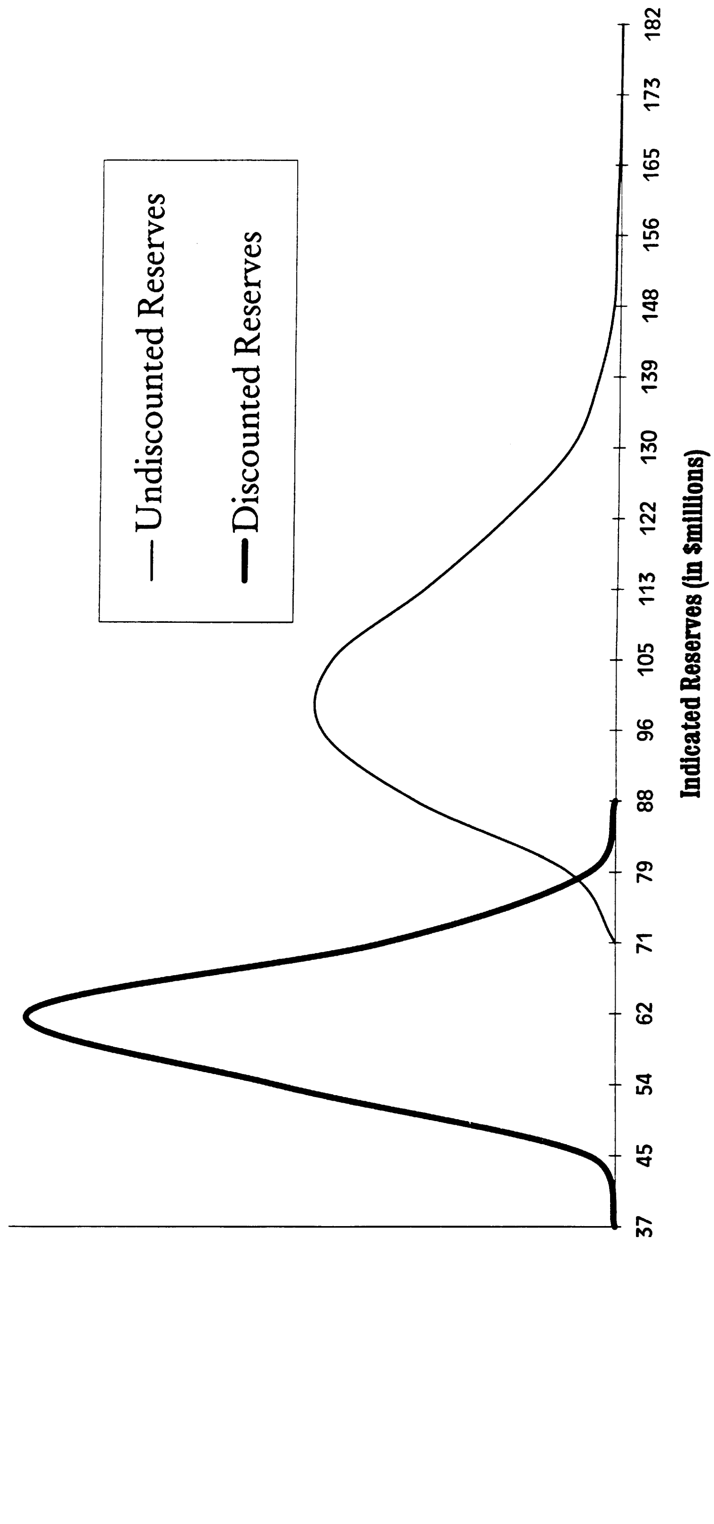

EXHIBIT 2

DISTRIBUTIONS OF RESERVES
WITH AND WITHOUT INFLATION ADJUSTMENTS

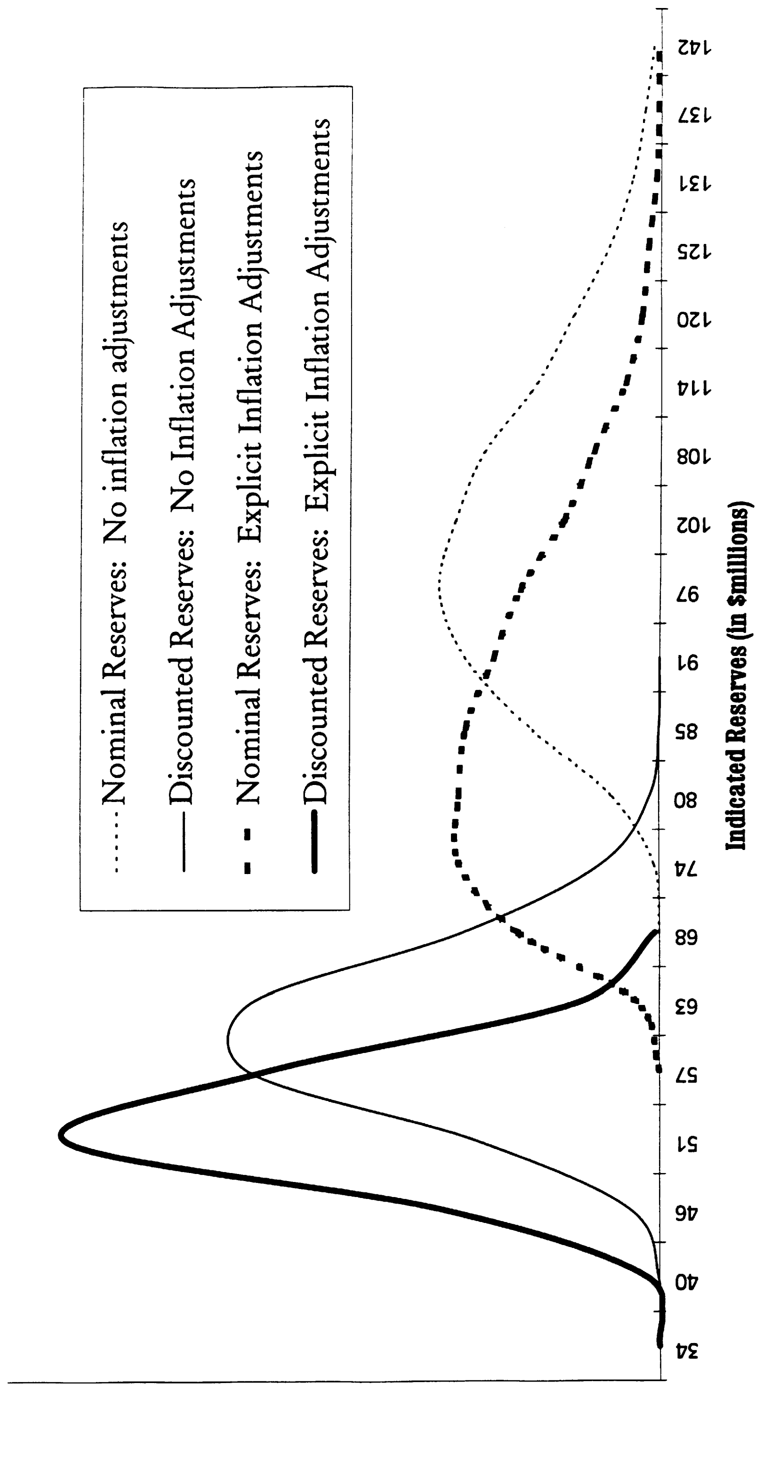

APPENDIX A

WORKERS COMPENSATION RESERVES AND RISK-BASED CAPITAL REQUIREMENTS

The text of this paper distinguishes between "regulatory measures" of reserving risk, as used in the NAIC's risk-based capital formula, and "actuarial measures" of reserving risk, as quantified here. The analysis in this paper shows that the volatility inherent in workers compensation reserve estimates is well below the implicit interest margin in statutory (undiscounted) reserves. The NAIC risk-based capital formula, however, has a reserving risk charge of 11% for workers compensation, even after incorporation of the expected investment income on the assets supporting the reserves.

An actuary unfamiliar with the development of the workers compensation reserving risk charge in the risk-based capital formula might conclude that "regulatory measures" of workers compensation reserving risk give high capital charges whereas "actuarial measures" give low charges. This is not correct. The risk-based capital formula gives a low charge for workers compensation reserving risk, even as the actuarial analysis in this paper provides. The final 11% charge in the risk-based capital formula is an *ad hoc* revision intended to provide more "reasonable" capital requirements.

The workers compensation reserving risk charge was one of the most contested aspects of the risk-based capital formula, and the derivation of the final 11% charge was never publicly revealed. This appendix explains the issues relating to the workers compensation reserving risk charge, and it shows the charge resulting from the NAIC "worst-case year" method.

Adverse Development and Loss Reserve Discounting

The reserving risk charge in the risk-based capital formula bases the capital requirements on the historical adverse loss development in each line of business. The "worst-case" industry-wide adverse loss development as a percentage of initial reserves is determined from Schedule P data, and this figure is then reduced by a conservative estimate of expected investment income.

For workers compensation, the original risk-based capital formula produced a charge of 0.4%.[37] The 1992 Best's *Aggregates and Averages* shows a gross "worst-case year" adverse development of 24.2%, as derived in Exhibit A-1.

Two considerations related to loss reserve discounting complicate the estimation of the reserving risk charge for workers compensation.

- Statutory accounting conventions for property/casualty insurers are conservative, particularly with regard to the reporting of loss reserves. The Annual Statement shows undiscounted reserves, leaving a large margin in the reserves themselves, particularly for long-tailed lines of business.

  In other words, property/casualty insurers have two potential margins to ensure adequacy of loss reserves: an implicit interest margin in the reserves themselves, and an explicit capital requirement provided by the reserving risk charge. To avoid "double counting," the risk-based capital formula offsets the implicit interest margin against the explicit reserving risk charge.

- The "double margin" occurs when reserves are reported on an undiscounted basis. But some property/casualty reserves are reported on at least a partially discounted basis. For instance, many carriers use tabular discounts for workers compensation lifetime pension claims. The special statutory treatment of workers compensation lifetime pension cases necessitates adjustments to the reserving risk charge.

---

[37]For a full description of the risk-based capital reserving risk charges, see Feldblum [23].

Both the NAIC Risk-Based Capital Working Group and the American Academy of Actuaries task force on risk-based capital spent months working on these two topics. The issues are complex, and no clear explanation is available for either regulators or for industry personnel. To clarify the issues, this appendix discusses the treatment of the implicit interest margin in statutory reserves and the adjustments needed for tabular loss reserve discounts in workers compensation.

Payment Patterns and Discount Rates

The amount of the implicit interest margin, or the difference between undiscounted (full-value) reserves and discounted (economic) reserves, depends on two items: the payout pattern of the loss reserves and the interest rate used to discount them.

For most lines of business, the NAIC risk-based capital formula uses the IRS loss reserve payment pattern along with a flat 5% discount rate. These choices were made for simplicity. Using the IRS discounting pattern avoids the need to examine loss reserve payout patterns, and using a flat 5% discount rate avoids the need to examine investment yields. For some lines of business, these choices are acceptable proxies for good solvency regulation. For workers compensation, greater complexities arise.

- *Payment Pattern*: The IRS procedure assumes that all losses are paid out within 15 years. Moreover, the pattern is based on the industry data for the first 10 years as reported in Schedule P.

For short-tailed lines of business, this is not unreasonable, since most losses are indeed paid out before the Schedule P triangles end. Workers compensation reserves, however, have a payout schedule of about 50 years, since permanent total disability cases—which are a small percentage of the claim count but a large percentage of the dollar amount—extend for the lifetime of the injured worker.

- *Discount Rate*: For its discount rate, the IRS uses a 60 month rolling average of the federal midterm rate, which is defined as the average yield on outstanding Treasury securities with maturities between 3 and 9 years. Since 1986, the IRS discount rate has ranged between 6% and 8%.

Actual portfolio yields have been about 100 to 200 basis points higher, since insurance companies invest not only in Treasury securities but also in corporate bonds, common stocks, real estate, and mortgages. However, these latter investment vehicles have additional risks, such as default risks, market risks, and liquidity risks. As a loss reserve discounting rate, many casualty actuaries would prefer the 6% to 8% "risk-free" Treasury rate to the 8% to 10% portfolio rate, particularly for statutory financial statements which emphasize solvency.

The NAIC risk-based capital formula uses a flat 5% discount rate. A variety of justifications have been given, such as:

- The 5% rate is simple, obviating any need to examine actual investment yields and cutting off any arguments about the "appropriate" rate.
- The 5% rate adds an additional margin of conservatism, since it is 1 to 3 points lower than the corresponding IRS rate.

For lines of business where the implicit interest margin in the reserves is small, the difference between the 5% NAIC rate and the 6% to 8% IRS rate is not that important in setting capital requirements. For a line of business like workers compensation, however, where the discount factor ranges from 60% to 83%, depending upon the assumptions, the choice of discount rate has a great effect.

We begin the analysis below with the current NAIC risk-based capital assumptions to see the unadjusted charge produced by the formula. We then turn to actual payment patterns and investment yields to address the fundamental questions: "What is the risk associated with workers compensation loss reserves? And how much capital ought insurance companies to hold to guard against this risk?"

The IRS Discount Factor

The IRS determines the loss reserves payout pattern by examining the ratio of paid losses to incurred losses by line of business for each accident year from Part 1 of Schedule P. The data are drawn from Best's *Aggregates and Averages*, and the payout pattern is redetermined every five years.

Schedule P shows only 10 years of data, though several lines of business, such as workers compensation, have payout schedules extending up to 50 years. The IRS allows an extension of the payout pattern beyond the 10 years shown in Schedule P for up to an additional 6 years. The extension of the payout pattern does not rely on either empirical data or financial expectations. Rather, the payout percentage in the tenth year is repeated for each succeeding year until all reserves are paid out.

Accident Years vs. Aggregate Reserves

The IRS determines a discount factor for each accident year. The risk-based capital formula uses a single discount factor for all accident years combined. Thus, one must use a weighted average of the discount factors, based on the expected reserves by accident year.[38]

Exhibit A-2 shows the workers compensation payment pattern using the IRS procedures and the Best's *Aggregates and Averages* Schedule P data.

- The left-most column shows the payment year. Because workers compensation reserves are paid out so slowly, the IRS extends the payment schedule for the full 16 years. It is still far too short, particularly for lifetime pension cases.

[38]For simplicity, the calculations in this paper assume that the volume of workers compensation business is remaining steady from year to year. A theoretical refinement would be to use the actual volume of industrywide workers compensation reserves in each of the past ten years, though there is no significant difference in the result.

FIGURE 1

WORKERS COMPENSATION PAYOUT PATTERNS

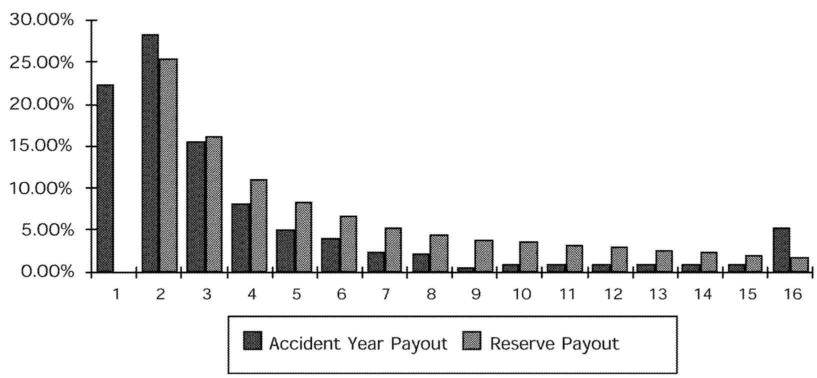

- The middle column shows the payment schedule for an individual accident year. This payment schedule says that 22.34% of an accident year's incurred losses are paid in the first calendar year, 28.36% in the next calendar year, and so forth.
- The right-most column shows the payment schedule for the aggregate reserves, assuming no change in business volume over the 16 year period. This payment schedule says that 25.42% of the reserves will be paid in the immediately following calendar year, 16.14% in the next calendar year, and so forth.

Figure 1 shows the payout patterns for an individual accident year and for the aggregate reserves. The horizontal axis represents time since the inception of the most recent accident year. The accident year payout pattern begins with the first losses paid on the policy, soon after the inception of the accident year. The valuation date of the reserves in the graph is the conclusion of the most recent accident year, so the payout pattern begins in the second year since inception.

The payout pattern is combined with an annual interest rate to give the discount factor, or the ratio of discounted reserves to undiscounted reserves. With an interest rate of 5% per annum, the discount factor for the reserves is 82.98%. The risk-based capital formula would therefore indicate a reserving risk charge of

$$[1.242 \times 82.98\%] - 1 = 3.06\%.$$

The 3% reserving risk charge depends upon the conservative 5% annual interest rate and the short IRS payment pattern. More realistic interest rates and payment patterns, even when still containing margins for conservatism, lead to a negative charge. We discuss these in conjunction with tabular loss reserve discounts below.

Discounted Reserves

What if an insurer holds discounted reserves, or partially discounted reserves? How should the reserving risk procedure described above be modified to account for the reserve discount?

This question is most relevant for workers compensation. Statutory accounting normally requires that insurers report undiscounted, or full-value, reserves. An exception is made for workers compensation lifetime pension cases, where insurers are allowed to value indemnity (lost income) reserves on a discounted basis. State statutes often mandate conservative discount rates, usually between 3.5% and 5% per annum, with the most common being 4%. These reserve discounts are termed "tabular" discounts, since they are determined from mortality tables, not from aggregate cash flow analyses.

Adverse Development and Interest Unwinding

The combination of three factors—(a) adverse development, (b) the unwinding of interest discounts, and (c) weekly claim payments—produces intricate results that are difficult even for the most technically oriented readers to follow. So let us begin with a simple example, which illustrates the concepts discussed above.

Suppose we have one claim, which will be used for determining both the "worst case" adverse loss development and the interest discount factor. The claim occurred in 1987, and it will be paid in 1997 for \$10,000.

Suppose first that the company accurately estimates the ultimate settlement amount and sets up this value at its initial reserve. Adverse loss development in this "worst case year" is 0%. Since there is a substantial implicit interest offset—the claim is paid 10 years after it occurs—the final reserving risk charge would be negative. In practice, there are no negative charges in the NAIC risk-based capital formula, since all charges are bounded below by 0%.

How large is the offset for the implicit interest discount? For a claim paid ten years after it occurs with a 5% per annum discount factor, the offset is $1 \div (1.05)^{10} = 61.39\%$. The final reserving risk charge in this simplified illustration is 38.61%.

What if the company holds the reserve on a discounted basis, using a 4% per annum discount rate? In 1987, the company sets up a reserve of $[\$10{,}000 \div (1.04)^{10}]$, or \$6,756. In 1988, the discounted reserve increases to $[\$10{,}000 \div (1.04)^{9}]$, or \$7,026. In 1989, the discounted reserve increases to $[\$10{,}000 \div (1.04)^{8}]$, or \$7,307.

The increases in the held reserve, from \$6,756 to \$7,026 in 1988, and from \$7,026 to \$7,307 in 1989, stem from the "unwinding" of the interest discount. However, they show up in Schedule P of the Annual Statement just like any other adverse development.[39]

Figure 2 shows the unwinding of the 4% interest discount over the course of the ten years that the reserve is on the company's

[39]This was true for the *pre-1995* Schedule P, when Part 2 was net of tabular discounts, though it was gross of non-tabular discounts. In 1995 and subsequent Annual Statements, Part 2 of Schedule P is gross of all discounts, so the unwinding of the interest discount no longer shows up as adverse development (see Feldblum [20]). The NAIC risk-based capital reserving risk charges were derived from the 1992 Schedule P.

FIGURE 2

UNWINDING OF INTEREST DISCOUNT

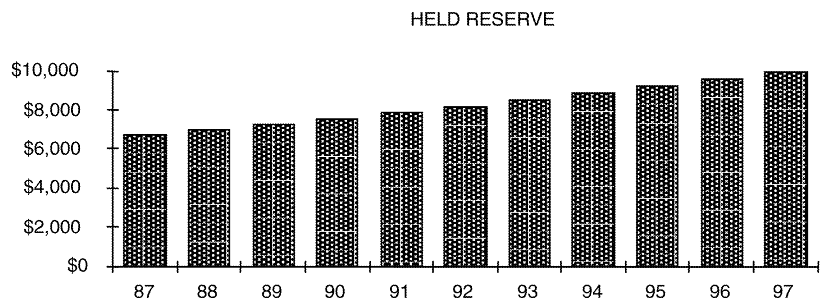

books. Between 1987 and 1992, the held reserve increases from \$6,756 to \$8,219, for observed adverse loss development during this period of 21.67% [= $(8,219 - 6,756) \div 6,756$].

The unwinding of the interest discount during 1987 through 1992 is reflected in the observed adverse development, so it is picked up by the NAIC calculation of the reserving risk charge. That is,

- A valuation basis that uses undiscounted reserves shows no adverse loss development on this claim.
- A valuation basis that uses reserves discounted at a 4% annual rate shows 21.67% of observed loss development.

The higher risk-based capital reserving risk charge generated by the discounted reserves is offset by the lower reserves held by the company.

Future Interest Unwinding

The unwinding of the interest discount continues from 1992 through 1997. Since this future unwinding is not yet reflected in the Schedule P exhibits of historical adverse loss development, a modification of the standard reserving risk charge calculation is needed.

What adjustment is needed? Consider the assumptions underlying the reserving risk charge. The reserving risk charge implicitly says:

> *Let us select the "worst case" adverse loss development that happened between 1983 and 1992, and let us assume that it might happen again.*

This procedure assumes that the 1992 reserves are adequate. That is to say, we should not *expect* either adverse or favorable development of the 1992 reserves.[40]

This is the proper assumption for the risk-based capital formula. The observed adverse loss development is meant to capture unanticipated external factors that cause higher or lower settlement values for insurance claims. A line of business may show adverse loss development even if the initial reserves were properly set on a "best estimate" basis. If a company is indeed holding inadequate reserves, it is the task of the financial examiners of the domiciliary state's insurance department to correct the situation. This is not the role of the generic risk-based capital formula.

If the reserves are valued on a discounted basis, however, they will continue to show (apparent) adverse development until all the claims are settled. In the example above,

- The unwinding of the interest discount between 1987 and 1992 is reflected in the observed adverse loss development, and no further adjustments are needed.
- The unwinding of the interest discount between 1992 and 1997 is not reflected anywhere, so an adjustment to the calculation procedure must be made.

---

[40] We do not *expect* either adverse or favorable development of the 1992 reserves. The risk-based capital requirement guards against *unexpected* adverse development of the reserves.

Alternative Adjustments

There are two ways to make this adjustment: either in the "worst case year" industry adverse loss development or in the offset for the implicit interest discount.

- *Adverse loss development*: One might add the expected future unwinding of the interest discount that will occur after the final valuation date to the "worst case year" observed adverse loss development. In the example above, the observed adverse loss development from 1987 to 1992 is $1,464, giving a factor of +21.7% as a percentage of beginning reserves. We expect further adverse loss development of $1,781 from 1992 to 1997 because of continued unwinding of the interest discount. The total adverse loss development is therefore $3,245, or +48.0% as a percentage of beginning reserves.
- *Implicit interest discount*: The further unwinding of the actual interest discount in the reserves may be used to reduce the offset for the implicit interest discount. In the example above, the observed adverse loss development is offset by ten years of implicit interest discount at a 5% annual rate. However, there are five years of unwinding of the actual 4% interest discount that are still to come (1992 through 1997), and that are not reflected in the observed adverse development.

In our illustration, ten years of implicit interest discount at a 5% annual rate gives a discount factor of 61.4%. Five future years of actual interest unwinding at a 4% annual rate gives a discount factor of 82.2%. The interest margin that should offset the "worst case year" adverse loss development is the *excess* of the implicit interest cushion over the actual interest discount, or 74.7% [= 61.2% ÷ 82.2%].

Diversity and Other Obstacles

In practice, the needed adjustments for tabular discounts are difficult to determine for a variety of reasons.

- *Industry Practice*: There is great disparity among insurance companies in the use of tabular reserve discounts. The prevalent practice is to use tabular discounts on indemnity benefits for lifetime pension cases. But there are companies that do not use tabular reserve discounts at all, and that report aggregate loss reserves on a full-value basis.[41]

- *Pension Identification*: Some companies show tabular discounts only for claims that have been identified as lifetime pension cases. Other companies show tabular discounts for the expected amount of claims that will ultimately be coded as lifetime pension cases.

  The distinction between "identified" and "unidentified" lifetime pension cases is analogous to the distinction between "reported" and "IBNR" claims. A workers compensation claim may be reported to the company soon after it occurs, but it may remain "unidentified" as a lifetime pension case for several years.

- *Indemnity vs. Medical Benefits*: Workers compensation benefits comprise two parts: indemnity benefits, which cover the loss of income, and medical benefits, which cover such expenses as hospital stays and physicians' fees.

  Lifetime pension cases may show continuing payments of both types. For instance, an injured worker who becomes a quadriplegic may receive a weekly indemnity check for loss of income as well as compensation for the medical costs of around-the-clock nursing care.

  Some insurers will discount only the indemnity benefits, since the weekly benefits are fixed by statute.[42] Other insurers will discount the medical benefits as well, since the payments

[41] More precisely, the case reserves generally show the tabular discounts. However, these discounts are "grossed up," or eliminated, by the actuarial "bulk" reserves.

[42] In some states, the indemnity benefit may depend on cost of living adjustments, so the amounts are not entirely "fixed."

are regular and do not vary significantly, even if they are not fixed by statute.

- *Interest Rates*: The interest rate used for the tabular reserve discounts varies by company and by state of domicile. Some companies use a 3.5% annual rate, since this is the interest rate used in the NCCI statistical plan. Several New York and Pennsylvania domiciled companies use a 5% annual rate, since this is the rate permitted by statute in these states. Other companies may use a 4% annual rate, since this is the most common rate in other state statutes.

Pension Discounts

The 3.06% reserving risk charge calculated above uses the conservative 5% interest rate in the risk-based capital formula and the short IRS payment pattern.

As we have discussed above, the NAIC reserving risk charge presumes that loss reserves are reported at undiscounted values. If reserves are valued on a discounted basis—as is true for certain workers compensation cases—then one expects future "adverse development," so the NAIC procedure is incomplete.

What is the expected effect of tabular discounts on the reserving risk charge for workers compensation? Analysts unfamiliar with workers compensation are tempted to say: *It should increase the charge.*

This would indeed be true if lifetime pension cases had the same payment pattern as other workers compensation claims and the only difference between pension cases and other compensation claims were that the pension cases are reported on a discounted basis whereas the other compensation claims are reported on an undiscounted basis. But this is not so. In fact, the very reason that tabular reserve discounts are permitted for lifetime pension cases is that they are paid slowly but steadily over the course of decades.

In other words, to properly incorporate tabular discounts into the workers compensation reserving risk charge, two changes are needed:

- One must increase the "worst case year" adverse development to include the future unwinding of the interest discount on the pension cases. Alternatively, one may adjust the "implicit interest discount" offset to account for the discount already included in the reported reserves.
- One must adjust the payout pattern from the IRS sixteen year pattern to the longer pattern appropriate for lifetime pension cases.

The net effect is to reduce the reserving risk charge. In fact, the indicated charge becomes negative, so it would be capped at 0% by the NAIC formula rules.

This is expected. The NAIC risk-based capital formula imposes a reserving risk charge when the "worst case" adverse development exceeds the implicit interest margin in the reserves. For lines of business like products liability and non-proportional reinsurance, the potential adverse development may far exceed the implicit interest margin, so companies must hold substantial amounts of capital to guard against reserving risk. For workers compensation "non-pension" cases, the mandated statutory benefits reduce the risk of adverse development while the slow payment pattern increases the implicit interest discount, so that the latter almost entirely offsets the former, resulting in the 3% charge calculated above with the RBC formula's exceedingly conservative assumptions. For workers compensation lifetime pension cases, true adverse development practically disappears, since mortality rates do not fluctuate randomly, and only the unwinding of the tabular discount remains. Because of the extremely long payout pattern for lifetime pension cases and the low interest rate allowed for tabular discounts, the implicit interest margin in lifetime pension reserves is well in excess of the "worst case" adverse development.

To calculate the appropriate reserving risk charge for workers compensation, after taking into consideration the tabular discounts on lifetime pension cases, we make the two adjustments discussed above.

- We replace the IRS payment pattern with a 50 year payment pattern derived from the historical experience of the nation's largest compensation carrier. At a 5% per annum interest rate, the present value of the reserves is 65.6% of the ultimate value, as shown in Exhibit A-3.[43]
- We increase the "worst case year" adverse development to incorporate the future interest unwinding on lifetime pension cases. The observed "worst case year" adverse development is 24.2% of initial reserves, from the 1985 statement date to the 1992 statement date. This includes the unwinding of tabular interest discount between 1985 and 1992. The post-1992 unwinding of interest discount on these pension cases adds between 6% and 8% to this figure. To be conservative, we use the 8% endpoint, giving a total adverse development of 34.1%.[44]
- The resulting reserving risk charge is $(1.341 \times 0.656) - 1$, or -14.1%. In other words, industry-wide workers compensation reserves have always been adequate on a discounted basis, even during the worst of years.

---

[43] Are statistics from a single carrier, no matter how large, a valid proxy for industry-wide figures? For loss ratios, expense ratios, and profit margins they are not appropriate, since each carrier has its own operating strategy. But workers compensation payment patterns are determined by statute; they do not differ significantly among companies, assuming that they have a similar mix of business by state. In November 1996, the American Academy of Actuaries task force on risk-based capital verified the pattern shown in the exhibits here, using data from eight large workers compensation carriers.

[44] For the unwinding of the tabular interest discount, it is no longer appropriate to use a single company's experience as a proxy for the industry. Insurers vary in whether they use tabular discounts at all, what types of benefits they apply the discounts to, and what interest rate they use to discount the reserves. The "6% to 8%" range in the text results from extended observation of reserving practices in workers compensation, along with detailed analysis of one company's own experience. With the reporting of tabular discounts in the 1994 Schedule P, more refined estimates of industry-wide practice may soon be available.

EXHIBIT A-1

NAIC METHOD

Consolidated Industry 1992 Schedule P, Part 2D (Workers Compensation)
Incurred Losses and ALAE

	1983	1984	1985	1986	1987	1988	1989	1990	1991	1992
All Prior	18,141,872	18,124,544	18,133,835	18,522,890	18,876,893	19,168,300	19,695,156	20,083,948	20,568,671	21,085,073
1983	10,285,007	10,518,014	10,615,001	10,800,631	10,904,709	11,053,667	11,087,456	11,163,710	11,309,445	11,364,446
1984		11,935,500	12,483,704	12,996,457	13,398,843	13,641,258	13,807,452	13,890,249	14,025,270	14,170,486
1985			13,506,212	14,148,315	14,560,000	15,036,193	15,289,931	15,451,016	15,648,178	15,824,280
1986				15,657,270	16,137,074	16,618,301	16,738,489	16,875,565	17,142,404	17,341,361
1987					18,543,543	18,630,232	18,849,648	18,945,479	19,228,271	19,492,604
1988						21,144,056	21,525,659	21,824,122	22,103,365	22,403,642
1989							23,337,805	23,983,219	24,549,997	24,863,843
1990								25,687,116	26,642,155	26,948,591
1991									27,107,842	27,477,716
1992										25,391,687
Total Incurred	28,426,879	40,578,058	54,738,752	72,125,563	92,421,062	115,292,007	140,331,596	167,904,424	198,325,598	226,363,729
Latest View of Incurreds	32,449,519	46,620,005	62,444,285	79,785,646	99,278,250	121,681,892	146,545,735	173,494,326	200,972,042	226,363,729
Adverse Development	4,022,640	6,041,947	7,705,533	7,660,083	6,857,188	6,389,885	6,214,139	5,589,902	2,646,444	0

Consolidated Industry 1992 Schedule P, Part 3D (Workers Compensation)
Paid Losses and ALAE

	1983	1984	1985	1986	1987	1988	1989	1990	1991	1992
All Prior	0	3,644,371	6,120,130	7,945,566	9,435,974	10,681,184	11,778,823	12,725,216	13,559,492	14,283,870
1983	2,595,880	5,414,887	6,989,233	8,016,341	8,714,150	9,206,673	9,565,906	9,842,305	10,035,523	10,182,271
1984		3,098,456	6,475,507	8,592,479	9,951,477	10,858,138	11,474,586	11,929,005	12,268,081	12,519,733
1985			3,307,517	7,223,536	9,609,598	11,175,251	12,188,709	12,884,942	13,409,641	13,785,641
1986				3,399,423	7,693,744	10,430,068	12,205,296	13,319,794	14,100,702	14,642,580
1987					3,823,180	8,916,751	12,037,953	13,992,209	15,236,501	16,062,480
1988						4,517,537	10,522,224	14,272,224	16,542,809	17,976,821
1989							4,923,056	11,851,679	16,021,809	18,519,232
1990								5,283,149	12,856,717	17,435,376
1991									5,481,562	12,644,529
1992										4,795,009

Consolidated Industry 1992 Schedule P, [{Part 2D}–{Part 3D}] (Workers Compensation)
Loss and ALAE Reserves

	1983	1984	1985	1986	1987	1988	1989	1990	1991	1992
All Prior	18,141,872	14,480,173	12,013,705	10,577,324	9,440,919	8,487,116	7,916,333	7,358,732	7,009,179	6,801,203
1983	7,689,127	5,103,127	3,625,768	2,784,290	2,190,559	1,846,994	1,521,550	1,321,405	1,273,922	1,182,175
1984		8,837,044	6,008,197	4,403,978	3,447,366	2,783,120	2,332,866	1,961,244	1,757,189	1,650,753
1985			10,198,695	6,924,779	4,950,402	3,860,942	3,101,222	2,566,074	2,238,537	2,038,639
1986				12,257,847	8,443,330	6,188,233	4,533,193	3,555,771	3,041,702	2,698,781
1987					14,720,363	9,713,481	6,811,695	4,953,270	3,991,770	3,430,124
1988						16,626,519	11,003,435	7,551,898	5,560,556	4,426,821
1989							18,414,749	12,131,540	8,528,188	6,344,611
1990								20,403,967	13,785,438	9,513,215
1991									21,626,280	14,833,187
1992										20,596,678
Reserves Held (1)	25,830,999	28,420,344	31,846,365	36,948,218	43,192,939	49,506,405	55,635,043	61,803,901	68,812,761	73,516,187
Adverse Development (2)	4,022,640	6,041,947	7,705,533	7,660,083	6,857,188	6,389,885	6,214,139	5,589,902	2,646,444	0
(2)/(1)	**15.6%**	**21.3%**	**24.2%**	**20.7%**	**15.9%**	**12.9%**	**11.2%**	**9.0%**	**3.8%**	**0.0%**
Worst Year Development:			**24.2%**							

EXHIBIT A-2

IRS Payment Pattern (1992–1996)

Payment Year	Payment Pattern (Single Accident Year)	Payment Pattern (Stationary Book)
	Accident Year Payout	Reserve Payout
1	22.34%	0.00%
2	28.36%	25.42%
3	15.49%	16.14%
4	8.23%	11.07%
5	5.14%	8.37%
6	4.16%	6.69%
7	2.41%	5.33%
8	2.31%	4.54%
9	0.52%	3.78%
10	0.96%	3.61%
11	0.96%	3.30%
12	0.96%	2.98%
13	0.96%	2.67%
14	0.96%	2.35%
15	0.96%	2.03%
16	5.25%	1.72%

EXHIBIT A-3

WORKERS' COMPENSATION PAYMENT PATTERN

Year	Payment Pattern (Single Accident Year)	Payment Pattern (Stationary Book)
1	0.190	
2	0.213	0.127
3	0.127	0.094
4	0.083	0.074
5	0.057	0.061
6	0.041	0.052
7	0.032	0.045
8	0.025	0.041
9	0.021	0.037
10	0.016	0.033
11	0.014	0.031
12	0.013	0.028
13	0.011	0.026
14	0.010	0.025
15	0.009	0.023
16	0.009	0.022
17	0.009	0.020
18	0.007	0.019
19	0.006	0.018
20	0.006	0.017
21	0.006	0.016
22	0.005	0.015
23	0.006	0.014
24	0.005	0.013
25	0.005	0.013
26	0.004	0.012
27	0.004	0.011
28	0.004	0.010
29	0.004	0.010
30	0.004	0.009
31	0.004	0.009
32	0.003	0.008
33	0.003	0.008
34	0.003	0.007
35	0.003	0.006
36	0.003	0.006
37	0.003	0.006
38	0.003	0.005
39	0.003	0.005
40	0.003	0.004
41	0.003	0.004
42	0.003	0.003
43	0.003	0.003
44	0.002	0.003
45	0.002	0.002
46	0.002	0.002
47	0.002	0.001
48	0.002	0.001
49	0.002	0.001
50	0.002	0.000
Total (Excluding first 12 months)	0.810	1.000
Present Value @ 5%	0.767	0.656

APPENDIX B

THE "EXPECTED POLICYHOLDER DEFICIT" YARDSTICK

Quantifying Reserve Uncertainty

Reserve uncertainty is a slippery concept, difficult to grasp and even more difficult to quantify. The actuary's skill is in forming a "best estimate" that accords with the data and that is appropriate for the particular business environment, such as the insurance marketplace for the premium rates, a statutory financial statement for the reserve requirements, or a merger transaction for the company valuation.

Quantifying reserve uncertainty is complex. A statistician might discuss reserve uncertainty as a probability distribution. One might show the mean of the distribution, its variance, and its higher moments; one might show various percentiles; one might even try to fit the empirical distribution to a mathematical curve. Accordingly, the exhibits in this paper show the mean, the standard deviation, the 95th percentile, and the 5th percentile of each of the distributions.

Capital Requirements

In recent years, state and federal regulators have been setting capital requirements for financial institutions, such as for banks and insurance companies. In theory, "risk-based capital requirements" relate the capital requirements to the uncertainty in various balance sheet items. In practice, most of the risk-based capital formulas that have been implemented in recent years use crude, generic charges that are based more on *ad hoc* considerations of what constitutes a "reasonable" charge than on rigorous actuarial or financial analyses.

Risk-based capital theory, however, is a siren for some actuaries and academicians, who have examined the relationship between uncertainty and capital requirements. In an ideal risk-based capital system, capital requirements should be calibrated among the balance sheet items in proportion to the risk that each poses to the company's solvency. Suppose a company has $100 million of bonds and $100 million of loss reserves, and the theoretically correct risk-based capital system says that the company needs $5 million of capital to guard against the uncertainty in the bond returns and $15 million of capital to guard against the uncertainty in the loss reserve payments. Then we can say that the uncertainty in the loss reserve portfolio is "three times as great" as the uncertainty in the bond portfolio.

Of course, we don't really mean that "uncertainty" is an absolute quantity that can be three times as great as some other figure. Rather, our measuring rod gives us a figure that we use as a proxy for the amount of uncertainty.

Moreover, our interest is not in absolute capital requirements but in the *relative* uncertainty among the company's various components. The regulator must indeed calibrate the absolute capital requirements, deciding between (i) $5 million of capital for bond risk and $15 million of capital for reserve risk versus (ii) $10 million of capital for bond risk and $30 million of capital for reserve risk. For the measurement of uncertainty, however, we are most interested in relative figures, such as the relative amount of capital needed to guard against reserve risk versus the amount needed to guard against bond risk, or the percentage reduction in capital for business written on loss sensitive contracts.

Calibrating Capital Requirements

There are two "actuarial" methods of calibrating capital requirements.

- The "probability of ruin" method says: How much capital is needed such that the chance of the company's insolvency during the coming time period is equal to or less than a given percentage?

- The "expected policyholder deficit" method says: How much capital is needed such that the expected loss to policyholders and claimants during the coming time period—as a percentage of the company's obligations to them—is equal to or less than a given amount?[45]

In this paper, we use the "expected policyholder deficit" (EPD) approach. The results would be no different if we used a "probability of ruin" approach.

Computing the Expected Policyholder Deficit

The "expected policyholder deficit" is a relatively new concept, having first been introduced in 1992. This appendix provides a brief explanation of the EPD analysis used in the paper.

Let us repeat the underlying question. The EPD analysis says: "Given a probability distribution for an uncertain balance sheet item, how much capital must the company hold such that the ratio of the expected loss to policyholders to the obligations to policyholders is less than or equal to a desired amount?" The format of the analysis depends on the type of probability distribution.

- For a simple discrete distribution, we can work out by hand the exact capital requirement. The type of simple discrete distribution that we illustrate below never occurs in real life. We use it only as a heuristic example, since the same procedure is used in our simulation analysis.
- If the empirical probability distribution can be modeled by a mathematically tractable curve, a closed-form analytic expression for the EPD can sometimes be found. In his previously cited paper, Butsic [14] does this for the normal and lognor-

[45]The "probability of ruin" method is explained in Daykin, Pentikainen, and Pesonen [17]. Probability of ruin analyses have long been used by European actuaries; see especially Beard, Pentikainen, and Pesonen [3] and Bowers, Gerber, Hickman, Jones, and Nesbitt [9]. The "expected policyholder deficit" method is explained in Butsic [13].

mal distributions, which can serve as reasonable proxies for many balance sheet items.

- The distributions in this paper are derived by means of stochastic simulation. Each distribution results from 10,000 Monte Carlo simulations. We determine the amount of capital needed to achieve a desired EPD ratio, as explained below.

Let us begin with the first case, the simple discrete distribution, to illustrate how the analysis proceeds. The extension to the full stochastic simulation merely requires greater computer power; there is no difference in the structure of the analysis.

Scenarios and Deficits

The distributions used in this paper are based on 10,000 simulations each. Think of this as 10,000 different scenarios. In fact, however, these simulations are *stochastic*. We do not know what these simulations are until after they have been realized. In other words, there are an infinite number of *possible* scenarios, 10,000 of which will be realized in the simulation.

To clarify the meaning of the "expected policyholder deficit," let us assume that an insurer with $250 million of assets faces two possible scenarios:

- In the *favorable* scenario, the company's interpretation of its insurance contracts will be upheld by the courts, and it must pay losses of $200 million.
- In the *adverse* scenario, the company's interpretation will *not* be upheld by the courts, and it must pay losses of $300 million.

Suppose also that there is a 60% chance of the favorable scenario being realized and a 40% chance of the adverse scenario being realized.[46]

---

[46]In the simulation analysis in this paper, only reserves are uncertain; assets are not uncertain. However, the same type of analysis applies to both assets and liabilities. Indeed, a more complete model would examine the external (economic and financial) factors that lead to variability in ultimate loss reserves, and it would analyze their effects on asset values as well.

What is the expected policyholder deficit? In the favorable scenario, the company has a positive net worth at the end. Since we are concerned only with deficits, a positive outcome of any size is considered a \$0 deficit.

In the adverse scenario, the final deficit is a \$50 million deficit, or −\$50 million. Since there is a 40% chance of an adverse outcome, the *expected* policyholder deficit is

$$\$0 \text{ million} \times 60\% + (-\$50 \text{ million} \times 40\%) = -\$20 \text{ million}.$$

The EPD Ratio

The definition of the EPD ratio is:

$$\text{EPD ratio} = (\text{expected policyholder deficit}) \div (\text{expected loss}).$$

In the example above, there is a 60% chance of a \$200 million payment to claimants and a 40% chance of a \$300 million payment to claimants. Thus, the expected loss is:

$$(\$200 \text{ million} \times 60\%) + (\$300 \text{ million} \times 40\%) = \$240 \text{ million}.$$

The EPD ratio is:

$$\$20 \text{ million} \div \$240 \text{ million} = 8.33\%.$$

Consistency

We use a 1% expected policyholder deficit ratio to determine the capital requirements. We use 1% to be consistent with the charges in the NAIC risk-based capital formula. In memoranda submitted to the American Academy of Actuaries task force on risk-based capital, Butsic estimates that the overall industrywide reserving risk charge in the NAIC risk-based capital formula amounts to approximately a 1% EPD ratio.

This allows us to compare the workers compensation loss reserve uncertainty to other sources of insurance company risk. If one believes that the overall capital requirements in the NAIC risk-based capital formula are reasonable, so a 1% EPD ratio is appropriate, then the degree of workers compensation loss reserve uncertainty measured in this paper can be viewed in light of the other NAIC capital requirements. As Butsic [12] says:

> *The amount of risk-based capital for each source of risk (e.g., underwriting, investment, or credit) must be such that the risk of insolvency (or other applicable impairment) is directly proportional to the amount of risk-based capital for each source of risk.*

Capital Requirements

We illustrate the calculation of capital requirements with the example given above. The capital required depends on the EPD ratio that the company (or the solvency regulator) seeks to maintain. We use a 1% target EPD ratio for this illustration.

If the company holds no capital, then its EPD ratio equals:

$$\text{(expected policyholder deficit)} \div \text{(expected loss)}$$

$$= \$20 \text{ million} \div \$240 \text{ million} = 8.33\%.$$

This exceeds the 1% target EPD ratio. The company must hold sufficient capital such that its revised EPD, or EPD*, satisfies the relationship:

$$\text{EPD}^* \div \text{(expected loss)} = \text{EPD}^* \div \$240 \text{ million} = 1\%,$$

$$\text{or} \qquad \text{EPD}^* = \$2.4 \text{ million}.$$

In the favorable scenario, the company already has sufficient funds to pay the losses. Adding capital will not change the policyholder deficit. In the adverse scenario, the company's assets are not sufficient to pay the losses. Adding capital will reduce the policyholder deficit. To achieve an EPD* of \$2.4 million, we solve:

$$40\% \times (\text{current assets} + \text{additional capital} - \text{liabilities})$$
$$= -\$2.4 \text{ million},$$
$$40\% \times (\$250 \text{ million} + \text{additional capital} - \$300 \text{ million})$$
$$= -\$2.4 \text{ million},$$
$$-\$50 \text{ million} + \text{additional capital}$$
$$= -\$6.0 \text{ million},$$
$$\text{additional capital} = \$44 \text{ million}.$$

Since the current assets are \$250 million, the additional capital required is \$44 million, and the expected losses are \$240 million, the total capital requirement for the company is \$250 million + \$44 million – \$240 million = \$54 million.

Full Simulation

The full analysis in this paper proceeds in the same fashion. The 10,000 simulations are run, each of which produces a "realization" for the loss amount. The average of these 10,000 realizations is the expected loss. The probability of each realization is 0.01%.

We first assume that the asset amount equals the expected loss, and we determine the loss payment and the deficit in each realization.

- If the loss amount is less than the asset amount, then the loss payment equals the loss amount, and the deficit is zero.
- If the loss amount exceeds the asset amount, then the loss payment equals the asset amount, and the deficit is the difference between the loss amount and the asset amount.

We sum the deficits in the 10,000 realizations, and we divide by 10,000. This gives the expected policyholder deficit. We then divide by the expected loss amount to give the EPD ratio.

If the probability distribution for the loss reserves is extremely compact, then the EPD ratio may be less than 1% even if no capital is held. For instance, suppose that the probability distribution is uniform over the range \$100 million ± \$4 million. Then the expected policyholder deficit is 1% if no capital is held.[47] This makes sense—if the loss payments are practically certain, there would be little need for surplus to support the reserves.

In practice, of course, the loss payments are not certain, and the EPD ratio would be greater than 1% if no capital is held. We proceed iteratively. We add capital and redetermine the loss payment and deficit in each scenario. This gives a new expected policyholder deficit and a new EPD ratio. If the EPD ratio still exceeds 1%, we must add more capital. If the EPD ratio is now less than 1%, we can subtract capital. With sufficient computer power, we quickly converge to a 1% EPD ratio.

[47] If the actual loss is less than \$100 million, then the deficit is zero. If the actual loss exceeds \$100 million, then the deficit is uniform over [\$0, \$4 million], for an average of \$2 million. The expected deficit over all cases is therefore \$1 million, for an EPD ratio of 1%.

APPENDIX C

THE SIMULATION PROCEDURE

Casualty actuaries are accustomed to providing point estimates of indicated reserves. The traditional procedures—such as a chain ladder loss development using 25 accident years of experience, supplemented by an "inverse power curve" tail factor—provide a sound basis for estimating workers compensation reserve needs. The actuary's task is to examine the historical experience for trends, evaluate the effects of internal (operational) changes on case reserving practices and settlement patterns, and forecast the likely influence of future economic and legal developments on the company's loss obligations.

Our perspective in this paper is different. We are not determining a point estimate of the reserve need; rather, we are determining a probability distribution for the reserve need. We use the same procedure and the same data as we would use for the point estimate: a chain ladder loss development based on 25 accident years of experience, along with a tail factor based on an inverse power curve fit. But now each step turns stochastic, and the probability distribution is determined by a Monte Carlo simulation.

The traditional procedures for determining point estimates are documented in various textbooks. This appendix shows the corresponding procedures for determining the probability distribution.

Data

We use a chain ladder *paid* loss development, since payment patterns for workers compensation are relatively stable whereas case reserving practices often differ from company to company and from year to year. This enables readers to replicate our results using their own companies' data.

We begin with accident year triangles with 25 years of cumulative paid losses, separately for indemnity (wage loss) and medical benefits. Indemnity and medical benefits have different loss payment patterns, and they are affected by different factors. For instance, medical benefits are strongly affected by medical inflation and by changes in medical utilization rates.

From the historical data we determine paid loss "age-to-age" factors (or "link ratios"). Exhibit C-1 shows 20 columns of paid loss age-to-age factors for countrywide indemnity plus ALAE benefits. For instance, the column labeled "12–24" shows the ratio of cumulative paid indemnity losses at 24 months to the corresponding cumulative paid indemnity losses at 12 months for each accident year. Similarly, Exhibit C-2 shows the paid loss age-to-age factors for countrywide medical benefits.

Point Estimates versus Realizations

The reserving actuary, when determining a point estimate, would examine these factors for trends. For a point estimate, the reserving actuary might use an average of the most recent five factors, instead of an average of all the factors in the column.

In this paper, our goal is to estimate the uncertainty in the reserve indications. Just as there was an upward trend in the age-to-age factors during the 1980s, there may be subsequent upward or downward trends in the 1990s. We therefore use the entire column of factors in our analysis. An "outlying" factor that is not a good estimator of the expected future value is an important element in measuring the potential variability of the future value.

We want to use the historical factors to simulate future "realizations." We do this by fitting the observed factors to a curve, thereby obtaining a probability distribution for the "12 to 24" age-to-age factors. Note carefully—this is *not* the probability distribution of the loss reserves, which will be the *output* of the simulation and which is *not* modeled by any mathematical function. This is the probability distribution of the age-to-age factors, which is the *input* to the simulation and is modeled by a curve.

Lognormal Curve Fitting

In this analysis, we used lognormal curves, which gave good fits to the data. Exhibit C-3 shows the curve fitting procedure for the first column of "indemnity plus ALAE" age-to-age factors.

For the lognormal curve, the probability distribution function is

$$f(x) = \frac{e^{-.5(\ln(x)-\mu/\sigma)^2}}{x\sigma\sqrt{2}\pi}$$

and the cumulative distribution function is

$$F(x) = \Phi\left(\frac{\ln(x)-\mu}{\sigma}\right)$$

We fit the function with the "development" part of the link ratios, or the "age-to-age factor minus one," as shown in Column 2 of Exhibit C-3. Column 3 shows the natural logarithms of the factors in Column 2. We use the method of moments to find the parameters of the fitted curve. The "mu" (μ) parameter is the mean of the figures in Column 3, and the "sigma" (σ) parameter is the standard deviation of the figures in Column 3.

We do the same for each "age-to-age" development column. The fitted parameters shown in the box in Exhibit C-3 are carried back to the final two rows in Exhibit C-1. Thus, each column has its own lognormal probability distribution function. We do this for development through 252 months. There is still paid loss development after 252 months, but there is insufficient historical experience to generate the factors, so we use an inverse power curve to estimate the loss development "tail" (discussed below).

For each run, we use a random number generator (Excel's built-in "RAND" function) to obtain simulated "age-to-age" factors in each column. Column 3 of Exhibit C-4 shows the results of one simulation for indemnity plus ALAE payments. For instance, the simulated age-to-age factor for 12 to 24 months of development is 2.312. The simulations for each of the 20 columns are independent of each other. For instance, the simulated 1.401 factor for "24 to 36" months in Column 3 of Exhibit C-4 is independent of the simulated 2.312 factor for "12 to 24" months.[48]

Parameter Variance

Two types of variance affect the simulation of future age-to-age link ratios: process variance and parameter variance.

- Process variance is the variance caused by the random nature of insurance losses. Even if the expected link ratios were known with certainty, the observed link ratios would differ from them because more losses than expected or less losses than expected might be paid in any given period.
- Parameter variance reflects the actuary's uncertainty about the expected losses. We estimate the probability distribution of the age-to-age link ratios from historical data. Our estimate may not be perfectly accurate; that is, we may have misestimated the parameters of the fitted probability distribution.

Quantifying Parameter Variance

To quantify parameter variance, we use a model developed by Kreps [42]. We assume that the observed age-to-age link ratios in a given development period come from the same lognormal probability distribution. We estimate the parameters of the fitted distribution as documented above, giving a lognormal distribution with parameters μ and σ.

[48] Our analysis assumes independence among columns. Dependence among columns may raise or lower the reserve variability, depending on whether the columns are positively or negatively correlated with each other. See the text of this paper for further discussion of trends in "age-to-age" factors on any observed correlations between columns, and see Holmberg [37] for methods of quantifying these correlations.

Think of our problem in the following fashion. We want to use our fitted distribution to simulate new observations. The actual future value may differ from the expected (mean) value because of process variance. In addition, we are uncertain whether we have chosen the proper expected (mean) value.

The Kreps procedure works as follows:

We fit a lognormal distribution to each column of the triangle. We assume that there are n age-to-age factors in the column, and that these are represented by $x_1, x_2 \ldots x_n$.

1. Calculate

$$\mu_0 = (1/n) \times \Sigma \ln(x_i)$$
$$\sigma_0 = \{(1/n) \times \Sigma[\ln(x_i) - \mu_0]^2\}^{0.5}.$$

These are the maximum likelihood estimators that would typically be used for simulation in the absence of parameter uncertainty.

2. Generate 3 random variables:

i) z, which has a standard normal distribution,

ii) w, which has a Chi-squared distribution with parameter $(n + \theta - 1)$,

iii) $v \times (n + \theta - 2)^{0.5}$, which has a t distribution, with parameter $(n + \theta - 2)$.

The value of θ depends on the Bayesian prior that is used. If the prior is a uniform distribution, $\theta = 0$. A power-law prior gives $\theta = 1$. The lower the value of θ, the more the effect of the parameter uncertainty. In private correspondence, Kreps pointed out to us that selecting $\theta = 0$ or 1 can give unreasonable results more often than one would like. In experimenting with this modeling, we found the same thing: every so often, the model would generate a gigantic age-to-age factor that would be totally unreasonable, given the nature of the business. Consequently, we used $\theta = 2$. (On rare occasions, even this gave unreasonable results—see the discussion below.)

3. Calculate

$$z_{\text{eff}} = v + z^*\{n^*(1+v^2)/w\}^{0.5}.$$

4. Calculate

$$x = \exp[\mu_0 + \sigma_0^* z_{\text{eff}}].$$

x is then a single simulation from the lognormal, taking parameter uncertainty into account.

As noted above, we occasionally found that an "unreasonably" large age-to-age factor would be generated. These factors were so large that they ended up dominating the simulation results. To eliminate these unreasonable cases, we set a rule that if any of the simulated age-to-age factors was more than 50 standard deviations from the mean, then that whole simulation would be eliminated. We are dealing with 20 years of workers compensation paid losses and are simulating a separate ATA factor for every point in the lower triangle. Also, since we have separate triangles for medical and indemnity, we are simulating 420 age-to-age factors each time. The rule applies to each one individually; in other words, if even 1 of the 420 was outside of 50 standard deviations, we threw them all out, and simulated again. Even so, we ended up throwing out the results in fewer than 3% of the cases. (In private correspondence, Kreps described this rule as "very generous" and suggested that he might have limited factors to within 10 standard deviations.)

Accident Year Correlations

In standard reserve analyses, the actuary derives a "best-estimate" age-to-age link ratio for each development period and uses that estimate for all future accident years. The actuary seeks a best-estimate reserve indication, so the best-estimate link ratios should be used for all years.

Our concern in this paper is with simulating the actual (future) development of the reserves. Each future year will have a distinct age-to-age link ratio in each period. To accurately model the future development of the reserves, we simulate separate future link ratios for each future accident year.

For example, suppose that our most recent accident year is 1994, our current valuation date is December 31, 1994, and we are simulating age-to-age link ratios from 48 to 60 months from accident years 1990 and prior. We use the simulated 48 to 60 month link ratios to develop accident years 1991 through 1994. We do four separate simulations to obtain four different link ratios for these four accident years.

Using separate simulated link ratios for each accident year assumes that the years are uncorrelated with respect to loss development. Using a single simulated link ratio for all accident years assumes that the accident years are perfectly correlated with each other. The independence assumption leads to a lower estimate of reserve uncertainty, since high development in one accident year may be offset by low development in another accident year. The dependence assumption leads to a higher estimate of reserve uncertainty, since high development in one accident year is associated with high development in all accident years.

The practical effect of using separate simulations versus using a single simulation for the link ratio for all future accident years in a given development period depends on the number of independent development periods in the simulation. The model in this paper uses 20 independent development periods plus a tail factor. Since the development periods are independent of each other, high development in one period is generally offset by low or average development in other periods. Therefore, the difference between independence among the accident years and dependence among the accident years is not great.

The tables in the text of this paper show results for both the independence assumption and the dependence assumption. The discussion in the paper uses the results for the independence assumption (i.e., for separate simulations by accident year).

Tail Development

Exhibit C-4 shows the fitting of the inverse power curve for one simulation. To clarify the procedure, let us *contrast* this with fitting an inverse power curve for a "best-estimate" reserve indication. For the "best-estimate" indication, we would use "selected" age-to-age factors in Column 3, such as averages of the factors in each column, or averages of the most recent years, or perhaps averages that exclude high and low factors. For the indemnity plus ALAE "12 to 24" months factor, the overall average is 2.685 and the average of the most recent five factors is 2.887. For a "best estimate," we would probably choose a factor such as 2.500.

In our analysis, the 20 factors in Column 3 are the results of *simulations* from the 20 fitted lognormal curves. For instance, the 2.312 factor is a simulation from the lognormal curve representing the probability distribution for the 12 to 24 month column.

From these *simulated* age-to-age factors, we fit an inverse power curve to estimate the "tail" development.[49] The inverse power curve will vary from simulation to simulation, since we have different "age-to-age" factors in each run. Moreover, the inverse power curve varies from accident year to accident year, since the simulated age-to-age link ratios vary by accident year.

The inverse power curve models the age-to-age ("ATA") factors as

$$\text{ATA} = 1 + at^{-b}$$

where "t" represents the "development year," and "a" and "b" are the parameters that we must fit. In workers compensation,

[49]For the rationale of using an inverse power curve for the tail development, see Sherman [52].

the shape of the loss payment pattern differs greatly between the first several years and subsequent years. In early years, there are many temporary total claims with rapid payment patterns. By the tenth year, most of the remaining reserves are for lifetime pension cases (fatalities and permanent total disability cases) with slow payment patterns. Therefore, we fit the inverse power curve using the simulated factors from the tenth through the 20th columns only.[50]

Columns (4) and (5) of Exhibit C-4 show the fitting procedure. Column (4) is the natural logarithm of the development year in Column (2), and Column (5) is the natural logarithm of the "simulated age-to-age [ATA] factor minus one" in Column (3). The inverse power curve can be written as

$$\ln(\text{ATA} - 1) = \ln(a) - b \times \ln(t).$$

We use a least squares procedure to determine the parameters a and b from the figures in Columns (4) and (5), giving $\ln(a) = -0.722$, or $a = 0.486$, and $b = 1.498$, as shown in the box at the bottom of Exhibit C-4.

The fitted inverse power curve provides age-to-age factors for development years 21 through 70. We don't really know how long paid loss development continues for workers compensation. Moreover, the factors are small. For development years 30 through 39 in this simulation, the age-to-age factors are about 1.002, and for development years 40 through 70, the factors are about 1.001. (The actual factors, of course, differ in the subsequent decimal places.) We therefore choose the length of the tail development stochastically; that is, the length of the total development is chosen randomly from a uniform distribution between 30 and 70 years.

---

[50]For actual reserve indications, one would probably segment the data between non-pension cases (temporary total and permanent partial cases) and lifetime pension cases (fatalities and permanent total cases).

Parameter Variance in the Tail

We have included both process variance and parameter variance in the simulated age-to-age link ratios for the first 20 development periods. The tail factors are an inverse power curve extension of each set of simulated age-to-age link ratios.

The tail factor selection procedure is a deterministic fit to the simulated age-to-age link ratios.[51] To the extent that process risk and/or parameter risk affect the variability of the age-to-age link ratios, they affect the variability of the tail factors.

One reviewer of an earlier draft of this paper wondered whether parameter variance might be incorporated independently in the tail factors. Specifically, the model currently has the following steps:

- We stochastically simulate age-to-age link ratios separately for each accident year and each development period, incorporating both process variance and parameter variance.
- We stochastically select the length of the development period, between 30 years and 70 years.
- We fit an inverse power curve to the simulated age-to-age link ratios to generate a tail factor.

The revised procedure would expand the third step in the list above as follows:

- Fit an inverse power curve to the simulated age-to-age link ratios. The inverse power curve is a two parameter family of curves. The fitting procedure gives "best estimates" for each of the two parameters.
- The current procedure considers the fitted parameters as the final values for each simulation. In place of this, assume a "structure function" for the distribution of these two param-

[51] The length of the tail development, though, is an independent stochastic choice, unrelated to the set of age-to-age link ratios.

eters. The values derived by fitting the inverse power curve would be the means of the distributions. The variance of the distribution, as well as the type of distribution, would be chosen subjectively.

- Stochastically select values for these two parameters from their assumed probability distribution. Use these simulated values of the two parameters to generate the inverse power curve tail factor.

Although this procedure is complex, it is important to consider all sources of variability, and to incorporate them, when feasible, into an actuarial model. Two factors, however, hampered the implementation of this procedure in our analysis.

- We had no *a priori* expectations about the type of structure function or the variance of the structure function.
- For the parameter risk in the link ratio estimation, we used a mathematically tractable approximation to simplify the simulation. For the parameter risk in the tail factor estimation, we are not aware of any corresponding approximation.

Thus, the procedures used in this paper do not separately incorporate parameter risk into the tail factor estimation.

Selected Factors

In the simulation shown in Exhibit C-5, the stochastic selection produced a development period of 54 years. We therefore have three sets of age-to-age factors:

- For development years 1 through 20, we use the simulated age-to-age factors generated by the lognormal curves for each column. For these development years, the "selected ATA" in Column (4) equals the "simulated ATA" in Column (2), not the "fitted ATA" in Column (3).
- For development years 21 through 53, we use the age-to-age factors from the fitted inverse power curve. For these devel-

opment years, the "selected ATA" in Column (4) equals the "fitted ATA" in Column (3).

- For development years 54 through 70, we use age-to-age factors of unity.

We now have all the age-to-age factors for this simulation. We "square the triangle" in the standard reserving fashion to determine ultimate incurred losses, and we subtract cumulative paid losses to date to obtain the required reserves. Exhibit C-6 shows the determination of the required medical reserves for one simulation. The "ultimate paids" in Exhibit C-6 are the "paid-to-date" times the "age-to-ultimate" factors, and the "indicated reserves" are the "ultimate paids" minus the "paid-to-date." The right-most two columns of Exhibit C-6 show the determination of the present value of the reserves. The "present value factors" are discussed in Appendix D, which has a full explanation of inflation effects.

We perform this simulation 10,000 times, giving a complete probability distribution of the required reserves, and we determine the mean, standard deviation, 95th percentile, and 5th percentile of this distribution. For the manner of determining the "capital required to achieve a 1% expected policyholder deficit ratio" (the right-most column of the exhibits in the text of this paper), see Appendix B.

EXHIBIT C-1

AGE-TO-AGE FACTORS FOR PAID INDEMNITY AND PAID ALAE

	12–24	24–36	36–48	48–60	60–72	72–84	84–96	96–108	108–120	120–132	132–144	144–156	156–168	168–180	180–192	192–204	204–216	216–228	228–240	240–252	252–264	264–276
1970					1.055	1.040	1.028	1.021	1.018	1.016	1.012	1.013	1.009	1.008	1.007	1.021	1.001	1.004	1.006	1.005	1.004	1.004
1971				1.094	1.055	1.041	1.026	1.024	1.016	1.011	1.010	1.011	1.007	1.010	1.012	1.009	1.005	1.006	1.005	1.005	1.004	1.006
1972			1.168	1.093	1.065	1.043	1.032	1.025	1.016	1.018	1.018	1.010	1.012	1.011	1.008	1.008	1.008	1.007	1.007	1.006	1.006	1.008
1973		1.386	1.169	1.096	1.062	1.049	1.033	1.025	1.020	1.017	1.012	1.013	1.012	1.011	1.008	1.008	1.007	1.007	1.007	1.005	1.007	
1974	2.334	1.385	1.164	1.093	1.068	1.044	1.034	1.022	1.019	1.016	1.013	1.010	1.013	1.009	1.008	1.009	1.007	1.006	1.008	1.007		
1975	2.310	1.398	1.190	1.116	1.076	1.051	1.037	1.026	1.021	1.016	1.013	1.014	1.012	1.011	1.010	1.009	1.010	1.010	1.010			
1976	2.262	1.388	1.195	1.117	1.069	1.048	1.031	1.027	1.020	1.017	1.013	1.013	1.012	1.008	1.011	1.009	1.009	1.009				
1977	2.192	1.397	1.191	1.111	1.070	1.048	1.031	1.023	1.019	1.016	1.015	1.013	1.011	1.010	1.009	1.009	1.009					
1978	2.246	1.407	1.193	1.113	1.068	1.048	1.031	1.027	1.022	1.019	1.016	1.014	1.012	1.012	1.011	1.009						
1979	2.199	1.409	1.192	1.109	1.068	1.045	1.036	1.027	1.023	1.020	1.019	1.013	1.011	1.012	1.012							
1980	2.169	1.400	1.209	1.107	1.074	1.050	1.038	1.030	1.023	1.020	1.017	1.017	1.011	1.011								
1981	2.191	1.400	1.185	1.115	1.075	1.055	1.041	1.032	1.025	1.019	1.017	1.016	1.011									
1982	2.179	1.395	1.207	1.131	1.098	1.059	1.046	1.043	1.026	1.024	1.020	1.015										
1983	2.283	1.437	1.227	1.140	1.088	1.064	1.048	1.037	1.025	1.022	1.017											
1984	2.345	1.473	1.228	1.134	1.089	1.064	1.044	1.033	1.027	1.018												
1985	2.422	1.473	1.245	1.140	1.087	1.057	1.041	1.030	1.020													
1986	2.377	1.500	1.237	1.133	1.085	1.055	1.038	1.026														
1987	2.452	1.496	1.234	1.127	1.080	1.053	1.034															
1988	2.496	1.498	1.228	1.126	1.074	1.047																
1989	2.502	1.512	1.231	1.121	1.068																	
1990	2.666	1.520	1.232	1.109																		
1991	2.529	1.507	1.217																			
1992	2.454	1.470																				
1993	2.426																					
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22
Avg of ln(ATA-1)	0.30	−0.82	−1.58	−2.16	−2.62	−3.00	−3.33	−3.59	−3.86	−4.03	−4.21	−4.34	−4.52	−4.61	−4.65	−4.64	−5.21	−5.00	−4.95	−5.19	−5.23	−5.12
Var of ln(ATA-1)	0.010	0.013	0.015	0.018	0.024	0.019	0.028	0.033	0.025	0.030	0.042	0.027	0.031	0.026	0.029	0.085	0.972	0.078	0.052	0.036	0.051	0.079
Lognormal Parameters, ignoring parameter risk:																						
mu	0.30	−0.82	−1.58	−2.16	−2.62	−3.00	−3.33	−3.59	−3.86	−4.03	−4.21	−4.34	−4.52	−4.61	−4.65	−4.64	−5.21	−5.00	−4.95	−5.19	−5.23	−5.12
sigma	0.104	0.116	0.127	0.136	0.158	0.142	0.172	0.188	0.162	0.180	0.211	0.169	0.184	0.168	0.178	0.307	1.046	0.299	0.247	0.209	0.253	0.324

EXHIBIT C-2

AGE-TO-AGE FACTORS FOR MEDICAL BENEFITS

	12–24	24–36	36–48	48–60	60–72	72–84	84–96	96–108	108–120	120–132	132–144	144–156	156–168	168–180	180–192	192–204	204–216	216–228	228–240	240–252	252–264	264–276
1970					1.013	1.011	1.010	1.008	1.005	1.007	1.006	1.008	1.007	1.007	1.005	1.006	1.001	1.007	1.005	1.007	1.008	1.009
1971				1.019	1.015	1.014	1.008	1.010	1.010	1.008	1.009	1.011	1.012	1.008	1.008	1.007	1.007	1.008	1.011	1.012	1.012	1.008
1972			1.045	1.026	1.018	1.014	1.013	1.012	1.009	1.013	1.010	1.011	1.018	1.011	1.008	1.010	1.007	1.009	1.010	1.011	1.008	1.010
1973		1.105	1.044	1.030	1.017	1.017	1.012	1.012	1.011	1.010	1.010	1.009	1.010	1.009	1.009	1.008	1.012	1.007	1.007	1.009	1.008	
1974	1.895	1.108	1.050	1.028	1.021	1.013	1.013	1.010	1.014	1.009	1.009	1.011	1.010	1.007	1.014	1.014	1.009	1.009	1.007	1.012		
1975	1.898	1.122	1.055	1.034	1.023	1.019	1.017	1.014	1.011	1.009	1.012	1.012	1.011	1.011	1.014	1.015	1.012	1.010	1.009			
1976	1.893	1.113	1.056	1.035	1.026	1.019	1.016	1.016	1.014	1.014	1.012	1.009	1.010	1.010	1.010	1.011	1.010	1.009				
1977	1.865	1.119	1.055	1.035	1.020	1.021	1.014	1.013	1.011	1.011	1.008	1.011	1.009	1.009	1.010	1.010	1.008					
1978	1.912	1.122	1.057	1.036	1.025	1.018	1.019	1.016	1.014	1.013	1.017	1.015	1.014	1.013	1.012	1.012						
1979	1.869	1.120	1.056	1.036	1.025	1.020	1.016	1.015	1.013	1.011	1.014	1.012	1.012	1.011	1.009							
1980	1.849	1.126	1.063	1.031	1.030	1.023	1.017	1.016	1.014	1.014	1.013	1.010	1.017	1.011								
1981	1.836	1.127	1.054	1.037	1.028	1.021	1.018	1.018	1.015	1.015	1.012	1.010	1.007									
1982	1.808	1.126	1.063	1.040	1.025	1.022	1.020	1.016	1.013	1.013	1.010	1.009										
1983	1.898	1.135	1.071	1.041	1.030	1.028	1.022	1.020	1.018	1.014	1.012											
1984	1.948	1.158	1.072	1.047	1.036	1.027	1.024	1.020	1.013	1.010												
1985	1.949	1.156	1.081	1.048	1.035	1.028	1.022	1.016	1.015													
1986	1.808	1.162	1.082	1.051	1.032	1.026	1.019	1.015														
1987	1.906	1.172	1.084	1.049	1.037	1.025	1.017															
1988	1.871	1.178	1.083	1.051	1.030	1.021																
1989	1.934	1.172	1.079	1.043	1.024																	
1990	1.898	1.171	1.071	1.036																		
1991	1.869	1.153	1.058																			
1992	1.773	1.126																				
1993	1.772																					
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22
Avg of ln(ATA-1)	−0.14	−1.99	−2.77	−3.31	−3.71	−3.93	−4.14	−4.25	−4.41	−4.49	−4.53	−4.56	−4.52	−4.66	−4.64	−4.63	−5.06	−4.79	−4.87	−4.61	−4.72	−4.68
Var of ln(ATA-1)	0.004	0.029	0.042	0.061	0.082	0.071	0.086	0.061	0.076	0.048	0.057	0.025	0.085	0.045	0.079	0.095	1.098	0.023	0.098	0.055	0.044	0.014
Lognormal Parameters:																						
mu	−0.14	−1.99	−2.77	−3.31	−3.71	−3.93	−4.14	−4.25	−4.41	−4.49	−4.53	−4.56	−4.52	−4.66	−4.64	−4.63	−5.06	−4.79	−4.87	−4.61	−4.72	−4.68
sigma	0.062	0.175	0.209	0.253	0.293	0.274	0.301	0.254	0.285	0.226	0.247	0.165	0.304	0.223	0.295	0.325	1.111	0.162	0.338	0.256	0.235	0.138

EXHIBIT C-3

Illustration of Fitting Lognormal Distributions to Age-to-Age Factors

	(1) 12–24 Factors for Indemnity & ALAE	(2) Age-to-Age Factor minus 1 (1) −1	(3) Natural Logs of (Age-to-Age Factors minus 1) ln (2)
1974	2.334	1.334	0.288
1975	2.310	1.310	0.270
1976	2.262	1.262	0.232
1977	2.192	1.192	0.175
1978	2.246	1.246	0.220
1979	2.199	1.199	0.181
1980	2.169	1.169	0.156
1981	2.191	1.191	0.175
1982	2.179	1.179	0.165
1983	2.283	1.283	0.249
1984	2.345	1.345	0.297
1985	2.422	1.422	0.352
1986	2.377	1.377	0.320
1987	2.452	1.452	0.373
1988	2.496	1.496	0.403
1989	2.502	1.502	0.407
1990	2.666	1.666	0.510
1991	2.529	1.529	0.425
1992	2.454	1.454	0.375
1993	2.426	1.426	0.355
Average	2.352	1.352	0.296
Variance	0.018	0.018	0.010

Fitted Lognormal	
μ_0 [= mean of the logs of (ATA-1)]	0.296
σ_0 [= standard deviation of logs of (ATA-1)]	0.099
Parameter Risk Procedure	
n (= number of ATA factors)	20
Θ	2
z [= Std normal random variable (simulated)]	−0.509
w [= Chi-square$_{(n+\Theta-1)}$ random variable (simulated)]	14.475
v [= $t_{(n+\Theta-2)}$ random variable (simulated) $\div(n+\Theta-2)^{0.5}$]	0.419
z_{eff} [= $v + z \times \{n \times (1+v^2)/w\}^{0.5}$]	−0.230
Simulated ATA [= $1 + \exp(\mu_0 + \sigma_0 \times z_{eff})$]	2.315

The simulated age-to-age factor is a single pick from a lognormal distribution with parameter risk taken into account [Note that the "1+" at the start of the expression for the simulated ATA is needed because we fit the curve to (ATA-1)] For each simulated ATA factor, we need to simulate from 3 random variables, to get z, w, and v This was done in Excel, by inverting the cumulative density functions of the respective distributions.

EXHIBIT C-4

Illustration of Fitting an Inverse Power Curve to the Simulated Age-to-Age Factors

(1) Development Period	(2) Year	(3) Simulated ATA	(4) ln(year) ln(2)	(5) ln(ATA-1) ln[(3)−1]	(6) Fitted ATA $1 + a \times (2)^{[-b]}$
12–24	1	2.312			1.486
24–36	2	1.401			1.172
36–48	3	1.208			1.094
48–60	4	1.106			1.061
60–72	5	1.072			1.044
72–84	6	1.055			1.033
84–96	7	1.040			1.026
96–108	8	1.037			1.022
108–120	9	1.029			1.018
120–132	10	1.015	2.303	−4.211	1.015
132–144	11	1.011	2.398	−4.484	1.013
144–156	12	1.013	2.485	−4.360	1.012
156–168	13	1.012	2.565	−4.439	1.010
168–180	14	1.011	2.639	−4.544	1.009
180–192	15	1.013	2.708	−4.362	1.008
192–204	16	1.008	2.773	−4.807	1.008
204–216	17	1.003	2.833	−5.770	1.007
216–228	18	1.005	2.890	−5.365	1.006
228–240	19	1.008	2.944	−4.856	1.006
240–252	20	1.007	2.996	−4.985	1.005

Fitting a least squares line to columns (4) and (5), with (5) as the dependent variable gives the following fitted paramenters:

$$\text{slope} = -1.498$$
$$\text{Intercept} = -0.722$$

Since the inverse power curve can be written in the form: $\ln(\text{ATA-1}) = \ln(a) - b\ln(t)$, we have the following parameters for the inverse power curve:

$$a = \exp(\text{intercept}) = 0.486$$
$$b = -\text{slope} = 1.498$$

EXHIBIT C-5

Illustration of Selecting Age-to-Age Factors

(1) Year	(2) Simulated ATA	(3) Fitted ATA $a = 0.486$ $b = 1.498$	(4) Selected ATA Cut-off for tail* 54	(1) Year	(3) Fitted ATA $a = 0.486$ $b = 1.498$	(4) Selected ATA Cut-off for tail* 54
1	2.312	2.626	2.312	36	1.008	1.008
2	1.401	1.576	1.401	37	1.007	1.007
3	1.208	1.314	1.208	38	1.007	1.007
4	1.106	1.204	1.106	39	1.007	1.007
5	1.072	1.146	1.072	40	1.006	1.006
6	1.055	1.111	1.055	41	1.006	1.006
7	1.040	1.088	1.040	42	1.006	1.006
8	1.037	1.072	1.037	43	1.006	1.006
9	1.029	1.060	1.029	44	1.006	1.006
10	1.015	1.052	1.015	45	1.005	1.005
11	1.011	1.045	1.011	46	1.005	1.005
12	1.013	1.039	1.013	47	1.005	1.005
13	1.012	1.035	1.012	48	1.005	1.005
14	1.011	1.031	1.011	49	1.005	1.005
15	1.013	1.028	1.013	50	1.005	1.005
16	1.008	1.026	1.008	51	1.004	1.004
17	1.003	1.023	1.003	52	1.004	1.004
18	1.005	1.021	1.005	53	1.004	1.004
19	1.008	1.020	1.008	54	1.004	1.000
20	1.007	1.018	1.007	55	1.004	1.000
21		1.017	1.017	56	1.004	1.000
22		1.016	1.016	57	1.004	1.000
23		1.015	1.015	58	1.004	1.000
24		1.014	1.014	59	1.004	1.000
25		1.013	1.013	60	1.004	1.000
26		1.012	1.012	61	1.003	1.000
27		1.012	1.012	62	1.003	1.000
28		1.011	1.011	63	1.003	1.000
29		1.010	1.010	64	1.003	1.000
30		1.010	1.010	65	1.003	1.000
31		1.009	1.009	66	1.003	1.000
32		1.009	1.009	67	1.003	1.000
33		1.009	1.009	68	1.003	1.000
34		1.008	1.008	69	1.003	1.000
35		1.008	1.008	70	1.003	1.000

*The cut off for the tail models the actuarial uncertainty in when to cut off the development from the inverse power curve The cut-off is based on a uniform distribution from 30 to 70.

EXHIBIT C-6

CALCULATION OF REQUIRED RESERVES FOR A SINGLE SIMULATION (MEDICAL PAYMENTS ONLY)

Year	Paid to Date	Age-to-Ultimate	Ultimate Paids	Indicated Reserves	Present Value Factor	Present Value of Reserves
1994	1,787,601	3.202	5,723,852	3,936,251	0.697	2,744,063
1993	3,324,538	1.778	5,910,348	2,585,810	0.579	1,496,946
1992	4,208,871	1.538	6,474,177	2,265,307	0.514	1,164,422
1991	7,017,997	1.462	10,261,961	3,243,963	0.497	1,612,068
1990	7,547,277	1.393	10,511,828	2,964,552	0.470	1,392,487
1989	7,905,743	1.348	10,655,677	2,749,934	0.452	1,243,164
1988	8,507,321	1.306	11,112,168	2,604,846	0.427	1,112,307
1987	7,629,124	1.284	9,798,726	2,169,602	0.422	915,457
1986	6,621,638	1.270	8,409,386	1,787,748	0.426	761,993
1985	5,398,367	1.250	6,746,697	1,348,331	0.418	563,797
1984	3,997,086	1.234	4,932,840	935,754	0.415	388,306
1983	3,198,587	1.222	3,908,208	709,622	0.417	295,599
1982	2,895,279	1.210	3,504,490	609,210	0.418	254,948
1981	2,929,995	1.200	3,517,101	587,106	0.422	248,033
1980	2,704,128	1.192	3,222,023	517,895	0.429	221,946
1979	2,552,368	1.181	3,013,230	460,862	0.428	197,322
1978	2,375,139	1.173	2,786,341	411,202	0.436	179,325
1977	1,986,508	1.172	2,328,957	342,449	0.463	158,711
1976	1,680,001	1.163	1,954,084	274,083	0.469	128,469
1975	1,321,413	1.159	1,531,944	210,531	0.489	103,028
1974	1,154,614	1.146	1,323,337	168,723	0.483	81,430
1973	1,004,449	1.135	1,140,181	135,733	0.478	64,937
1972	908,372	1.124	1,021,158	112,786	0.470	53,015
1971	782,100	1.118	874,591	92,491	0.478	44,228
1970	776,907	1.113	864,352	87,445	0.487	42,566
Total	90,215,423		121,527,659	31,312,236		15,468,566

APPENDIX D

INFLATION ADJUSTMENTS

For certain long-tailed lines of business, much reserve uncertainty stems from changes in the rate of inflation. For workers compensation medical benefits, as an example, the employer is responsible for physician fees, which are affected by the rate of inflation up through the date that the service is rendered.

Paid loss development analyses may overstate the uncertainty in reserve indications, particularly if one is concerned with the economic value of the reserves and not their nominal value. For instance, suppose that the cumulative paid losses *in real dollar terms* will increase by 30% over the coming year, for a "real dollar" age-to-age factor of 1.300. If inflation is high, the nominal age-to-age factor may be 1.350. If inflation is low, the nominal age-to-age factor may be 1.320.

To some extent, this is "apparent" reserve uncertainty, not real reserve uncertainty. We can get a better estimate of reserve uncertainty by

- Stripping inflation out of the historical paid losses,
- Determining "age-to-age" factors in real dollar terms,
- Using the "real dollar" factors to produce all the simulations, and
- Restoring nominal inflation, based upon a stochastically generated inflation rate path, to determine ultimate losses.[52]

Exhibit D-1 shows the procedure used to put the paid loss experience into real dollar terms (at a 1994 price level). We demon-

[52]These adjustments are equally important for standard "point estimates" of indicated reserves. Nominal dollar paid loss "age-to-age" factors have the historical inflation rate built into them (see Cook [15]). If future inflation is expected to be different from past inflation, a rote application of the paid loss chain ladder technique may give misleading reserve indications.

strate the procedure for medical benefits, which we assume to be fully inflation sensitive. Indemnity benefits, in contrast, are only partially inflation sensitive. About half the states have "cost of living" adjustments for wage loss benefits, but generally these adjustments apply only to certain cases (such as cases that extend for two years or more), and they are often capped (say, at 5% per annum).

We begin with the medical component of the Consumer Price Index, shown on the second row of Exhibit D-1. During the 1980s, the rate of increase in workers compensation medical benefits exceeded the medical CPI. This additional WC medical inflation is related to increases in utilization rates or, perhaps, to the incurral of medical services to justify claims for increased indemnity benefits.

For ratemaking, we would need a "loss cost trend factor" for workers compensation medical benefits, of which the medical CPI is but one component. For our purposes, we are concerned only with medical inflation. Changes in utilization rates remain embedded in the paid loss development factors. If the reserving actuary believes that future changes in utilization rates will differ from past changes in utilization rates, this expected difference must be separately quantified.

We must convert the *incremental* paid losses during each calendar year to their "real dollar" (calendar year 1994) values. For ease of application, the one dimensional index in the second row of Exhibit D-1 is converted to a two-dimensional triangle. For instance, the "0.76" in column (5) for accident year 1990 means that accident year losses paid between 48 and 60 months (i.e., between January 1, 1995, and December 31, 1995) must be multiplied by 0.76 to bring them to accident year 1990 levels. The 0.76 factor is derived from the inflation index: $0.76 = 1/(1.0885 \times 1.0805 \times 1.0667 \times 1.0536)$.

We now redo the entire simulation procedure as documented in Appendix C, using the paid losses that have all been adjusted to a 1994 cost level.

Inflation Rate Generator

The derivation of the stochastic medical inflation rate model is shown in Exhibit D-2. We use the medical CPI as the "monetary" inflation component of workers compensation medical benefits, since this is the index that we used to deflate the medical link ratios in Exhibit D-1.

Workers compensation medical loss cost trends are not necessarily the same as the medical CPI, whether year by year or over a long-term average, since other factors (such as utilization rates) affect medical loss cost trends. The historical link ratios are not deflated for this residual trend, so the residual trend is not added back for future periods. If the reserving actuary believes that future utilization rate trends will differ from the historical utilization rate trends, a further adjustment should be made to the simulation model.[53]

Restoring Inflation

To properly estimate reserves, we must "restore" future inflation at the rates stochastically generated for this scenario. To keep the calculations tractable, we assume (i) annual changes in interest rates and inflation rates, and (ii) mid-year loss payments.[54]

The procedure consists of the following steps:

- Remove inflation from the historical link ratios, fit them to a lognormal curve, accounting for parameter risk, and simulate future link ratios for each accident year, as in Appendix C.
- From the simulated link ratios, determine age-to-ultimate factors and payment patterns for each accident year.

[53]The advent of managed care procedures in the 1990s may warrant such an additional adjustment.

[54]Mid-year loss payments is the common proxy for loss payments spread evenly over the year. For payments after the first year, this is a reasonable approximation.

- Stochastically generate an interest rate path and an inflation rate path.
- Assume all payments are made at mid-year. Inflate the "real dollar" loss payments by the future inflation rates to determine nominal loss payments. The sum of the loss payments is the undiscounted required reserve.
- Discount the nominal loss payments by the future interest rates to determine discounted loss payments. The sum of the discounted loss payments is the discounted required reserve.

For example, suppose that in one simulation we had the following figures:

Year	Simulated Link Ratio	Development Factor	Payment Pattern	Inflation Rate	Interest Rate
1	1.776	2.446	0.409	5.7%	7.5%
2	1.105	1.378	0.317	6.3%	6.6%
3	1.057	1.247	0.076	6.2%	6.4%

The simulated link ratios are for a particular accident year in a particular simulation. The simulated development factors are the backward product of the simulated link ratios. For instance, $2.446 = 1.378 \times 1.776$.

The payment pattern is the percent of losses paid in the calendar year preceding the development factor in the adjoining cell. For instance, the development factor at the end of "year 1" is 2.446. This implies that the percent of losses paid in the first 12 months equals $1 \div 2.446$, or 40.9%. At the end of the second year, the development factor is 1.378. This implies that the percent of losses paid in the first 24 month is $1 \div 1.378$, or 72.6%. Since 40.9% of losses have been paid in the first 12 months, 31.7% of losses are paid between 12 and 24 months.

To simplify the exposition of the inflation and discounting procedures, assume that total "real dollar" losses are $1,000,000 for the most recent accident year (1994 in our example). Of this amount, \$409,000 is paid in the first twelve months, and they are not included in the loss reserves held at the end of the year.

Another \$317,000 is paid on July 1 of the following calendar year (1995 in our example). This amount is in December 31, 1994 dollars. The nominal losses paid are therefore $\$317,000 \times (1.057)^{0.5}$. The discounted dollars in this scenario equal $\$317,000 \times (1.057)^{0.5} \div (1.075)^{0.5}$.

Another \$76,000 is paid on July 1 of the next calendar year (1996 in our example). Again, this amount is in December 31, 1994 dollars. The nominal losses paid are therefore $\$76,000 \times (1.057) \times (1.063)^{0.5}$. The discounted dollars in this scenario equal $\$76,000 \times (1.057) \times (1.063)^{0.5} \div \{(1.075) \times (1.066)^{0.5}\}$.

EXHIBIT D-1

STRIPPING MEDICAL INFLATION FROM THE LOSSES

Year	1970	1971	1972	1973	1974	1975	1976	1977	1978	1979	1980	1981	1982	1983	1984	1985	1986	1987	1988	1989	1990	1991	1992	1993	1994
Medical Inflation	6.65%	6.40%	4.75%	3.65%	6.65%	10.65%	10.75%	9.55%	9.00%	8.80%	10.10%	10.85%	11.15%	10.20%	7.50%	6.25%	6.90%	7.05%	6.55%	7.10%	8.35%	8.85%	8.05%	6.67%	5.36%

Index for use in Calendar Year

(Multiplying the corresponding element in the traingle by this factor puts the loss back at the medical price level for the accident year)

Accident Year	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
1970	1.00	0.94	0.90	0.87	0.81	0.73	0.66	0.60	0.55	0.51	0.46	0.42	0.38	0.34	0.32	0.30	0.28	0.26	0.24	0.23	0.21	0.19	0.18	0.17	0.16
1971	1.00	0.95	0.92	0.86	0.78	0.70	0.64	0.59	0.54	0.49	0.44	0.40	0.36	0.34	0.32	0.30	0.28	0.26	0.24	0.22	0.21	0.19	0.18	0.17	
1972	1.00	0.96	0.90	0.82	0.74	0.67	0.62	0.57	0.52	0.47	0.42	0.38	0.35	0.33	0.31	0.29	0.27	0.25	0.24	0.22	0.20	0.19	0.18		
1973	1.00	0.94	0.85	0.77	0.70	0.64	0.59	0.53	0.48	0.43	0.39	0.37	0.34	0.32	0.30	0.28	0.26	0.24	0.22	0.21	0.19	0.18			
1974	1.00	0.90	0.82	0.74	0.68	0.63	0.57	0.51	0.46	0.42	0.39	0.37	0.34	0.32	0.30	0.28	0.26	0.24	0.22	0.21	0.20				
1975	1.00	0.90	0.82	0.76	0.70	0.63	0.57	0.51	0.46	0.43	0.41	0.38	0.36	0.33	0.31	0.29	0.26	0.24	0.23	0.22					
1976	1.00	0.91	0.84	0.77	0.70	0.63	0.57	0.51	0.48	0.45	0.42	0.39	0.37	0.35	0.32	0.29	0.27	0.25	0.24						
1977	1.00	0.92	0.84	0.77	0.69	0.62	0.56	0.52	0.49	0.46	0.43	0.41	0.38	0.35	0.32	0.30	0.28	0.26							
1978	1.00	0.92	0.83	0.75	0.68	0.61	0.57	0.54	0.50	0.47	0.44	0.41	0.38	0.35	0.32	0.30	0.29								
1979	1.00	0.91	0.82	0.74	0.67	0.62	0.59	0.55	0.51	0.48	0.45	0.41	0.38	0.35	0.33	0.31									
1980	1.00	0.90	0.81	0.74	0.69	0.64	0.60	0.56	0.53	0.49	0.46	0.42	0.39	0.36	0.34										
1981	1.00	0.90	0.82	0.76	0.71	0.67	0.62	0.59	0.55	0.51	0.46	0.43	0.40	0.38											
1982	1.00	0.91	0.84	0.79	0.74	0.69	0.65	0.61	0.56	0.52	0.48	0.45	0.42												
1983	1.00	0.93	0.88	0.82	0.77	0.72	0.67	0.62	0.57	0.53	0.49	0.47													
1984	1.00	0.94	0.88	0.82	0.77	0.72	0.67	0.61	0.57	0.53	0.50														
1985	1.00	0.94	0.87	0.82	0.77	0.71	0.65	0.60	0.56	0.53															
1986	1.00	0.93	0.88	0.82	0.76	0.69	0.64	0.60	0.57																
1987	1.00	0.94	0.88	0.81	0.74	0.69	0.64	0.61																	
1988	1.00	0.93	0.86	0.79	0.73	0.69	0.65																		
1989	1.00	0.92	0.85	0.78	0.74	0.70																			
1990	1.00	0.92	0.85	0.80	0.76																				
1991	1.00	0.93	0.87	0.82																					
1992	1.00	0.94	0.89																						
1993	1.00	0.95																							
1994	1.00																								

EXHIBIT D-2

PAGE 1

FITTING OF MODEL FOR MEDICAL INFLATION

Model: Medical inflation$_t$ =
$a\times(\text{interest rate}_t) + \beta\times[(\text{medical inflation}_{t-1}) - \alpha\times(\text{interest rate}_{t-1})]$
$+(1-\beta)\times[(\text{avg. medical inflation}) - \alpha\times(\text{avg. interest rate})] + \text{error}_t$
$\alpha = 0.484 \qquad \beta = 0.546$
α and β are chosen to minimize the sum of the squared errors in column 6

Year	(1) Medical CPI at December	(2) Annual % Increase in Medical CPI	(3) Yield on Intermediate Term Govt Bonds*	(4) Least-Squares Fit of Medical Inflation Model**	(5) Error***	(6) Squared Error****
1935	10.2					
1936	10.2	0.0%	1.3%			
1937	10.3	1.0%	1.1%	1.5%	−0.56%	0.00003
1938	10.3	0.0%	1.5%	2.3%	−2.30%	0.00053
1939	10.4	1.0%	1.0%	1.4%	−0.43%	0.00002
1940	10.4	0.0%	0.6%	1.9%	−1.88%	0.00035
1941	10.5	1.0%	0.8%	1.6%	−0.61%	0.00004
1942	10.9	3.8%	0.7%	2.0%	1.82%	0.00033
1943	11.4	4.6%	1.5%	3.9%	0.67%	0.00004
1944	11.7	2.6%	1.4%	4.1%	−1.49%	0.00022
1945	12.0	2.6%	1.0%	2.9%	−0.33%	0.00001
1946	13.0	8.3%	1.1%	3.0%	5.34%	0.00285
1947	13.9	6.9%	1.3%	6.2%	0.69%	0.00005
1948	14.7	5.8%	1.5%	5.5%	0.27%	0.00001
1949	14.9	1.4%	1.2%	4.7%	−3.31%	0.00109
1950	15.4	3.4%	1.6%	2.5%	0.83%	0.00007
1951	16.3	5.8%	2.2%	3.8%	2.06%	0.00043
1952	17.0	4.3%	2.4%	5.1%	−0.79%	0.00006
1953	17.6	3.5%	2.2%	4.1%	−0.58%	0.00003
1954	18.0	2.3%	1.7%	3.5%	−1.24%	0.00015
1955	18.6	3.3%	2.8%	3.5%	−0.14%	0.00000
1956	19.2	3.2%	3.6%	4.2%	−0.94%	0.00009
1957	20.1	4.7%	2.8%	3.5%	1.18%	0.00014
1958	21.0	4.5%	3.8%	5.0%	−0.50%	0.00003
1959	21.8	3.8%	5.0%	5.2%	−1.37%	0.00019
1960	22.5	3.2%	3.3%	3.7%	−0.48%	0.00002
1961	23.2	3.1%	3.8%	4.1%	−0.95%	0.00009
1962	23.7	2.2%	3.5%	3.7%	−1.55%	0.00024
1963	24.3	2.5%	4.0%	3.5%	−1.00%	0.00010
1964	24.8	2.1%	4.0%	3.6%	−1.54%	0.00024
1965	25.5	2.8%	4.9%	3.8%	−0.94%	0.00009
1966	27.2	6.7%	4.8%	3.9%	2.77%	0.00077
1967	28.9	6.3%	5.8%	6.5%	−0.24%	0.00001
1968	30.7	6.2%	6.0%	6.1%	0.13%	0.00000
1969	32.6	6.2%	8.3%	7.2%	−0.98%	0.00010

EXHIBIT D-2

PAGE 2

FITTING OF MODEL FOR MEDICAL INFLATION

Year	(1) Medical CPI at December	(2) Annual % Increase in Medical CPI	(3) Yield on Intermediate Term Govt Bonds*	(4) Least-Squares Fit of Medical Inflation Model**	(5) Error***	(6) Squared Error****
1970	35.0	7.4%	5.9%	5.4%	1.99%	0.00040
1971	36.6	4.6%	5.3%	6.3%	−1.76%	0.00031
1972	37.8	3.3%	5.9%	5.3%	−1.99%	0.00040
1973	39.8	5.3%	6.8%	4.9%	0.43%	0.00002
1974	44.8	126%	7.1%	5.9%	6.69%	0.00448
1975	49.2	9.8%	7.2%	9.8%	0.04%	0.00000
1976	54.1	10.0%	6.0%	7.7%	2.27%	0.00051
1977	58.9	8.9%	7.5%	8.8%	0.06%	0.00000
1978	64.1	8.8%	8.8%	8.5%	0.37%	0.00001
1979	70.6	10.1%	10.3%	8.8%	1.33%	0.00018
1980	77.6	9.9%	12.5%	10.2%	−0.24%	0.00001
1981	87.3	12.5%	14.0%	10.2%	2.29%	0.00053
1982	96.9	11.0%	9.9%	9.3%	1.74%	0.00030
1983	103.1	6.4%	11.4%	10.2%	−3.84%	0.00147
1984	109.4	6.1%	11.0%	7.1%	−1.04%	0.00011
1985	116.8	6.8%	8.6%	5.9%	0.88%	0.00008
1986	125.8	7.7%	6.9%	6.1%	1.63%	0.00027
1987	133.1	5.8%	8.3%	7.8%	−1.95%	0.00038
1988	142.3	6.9%	9.2%	6.7%	0.18%	0.00000
1989	154.4	8.5%	7.9%	6.5%	1.98%	0.00039
1990	169.2	9.6%	7.7%	7.6%	1.99%	0.00040
1991	182.6	7.9%	6.0%	7.4%	0.50%	0.00003
1992	194.7	6.6%	6.1%	7.0%	−0.40%	0.00002
1993	205.2	5.4%	5.2%	5.9%	−0.46%	0.00002
1994	215.3	4.9%	7.8%	6.7%	−1.75%	0.00030
Mean		5.4%	5.0%		0.04%	0.00033
					183% = Std Dev of errors	0.01901 = Sum of square errors

* Source: Ibbotson Associates: Stocks, Bonds, Bills, and Inflation, 1995 Edition

** Column 4 = α[Col. 3 for current year]+β[Col. 2 for previous year−α(Col. 3 for previous year)]+(1 − β) [Avg. of Col. 2 − α(Avg. of Col. 3)]

*** Column 5 = Column 2−Column 4

**** Column 6 = {Column 5}^{2}

Fitted α and β minimize the sum of column 6.

The error term for the model is a normal distribution, with mean = 0.00% and standard deviation = 1.83%

APPENDIX E

LOSS-SENSITIVE CONTRACTS

In the text of this paper, we examine the uncertainty in the loss reserves. In practice, reserve uncertainty varies with the type of insurance contract. For instance, high-level workers compensation excess-of-loss covers, as well as large dollar deductible policies offered to large employers, have greater reserve uncertainty, particularly in the early policy years when the insurer's estimated liability is subject to great variation.

For business written on loss-sensitive contracts, such as retrospectively rated plans for large workers compensation risks or reinsurance treaties with sliding scale reinsurance commissions, the opposite is true. Companies are concerned with the uncertainty in the net reserves, or the future loss payments after adjustment for retrospective premiums and variable commissions.[55]

Large dollar deductible policies are relatively new, and we do not yet have the requisite data to estimate the reserve uncertainty. In addition, the slow payment patterns of workers compensation excess covers and of large dollar deductible policies will delay the empirical quantification of their reserving risk.

In contrast, we have relatively complete data on loss-sensitive contracts. Moreover, the effects of loss sensitive contracts on reserve uncertainty has become a significant regulatory and actuarial issue in recent years. The NAIC risk-based capital formula contains an offset of 15% to 30% to the reserving risk charge for business written on loss-sensitive contracts (Feldblum [23]). In

[55]The discussion here assumes familiarity with retrospective rating plans and with their parameters, such as loss limits, premium maximums, and premium minimums, as well as with standard reserving techniques for retrospective premiums. More detailed information on the retrospective rating plan pricing parameters may be found in Simon [54], Skurnick [55], Lee [43], Gillam and Snader [30], Bender [4], and Mahler [45]. The retrospective premium reserving techniques that underlie the analysis in this paper are discussed in Fitzgibbon [26], Berry [6], Teng and Perkins [57], and Feldblum [21].

1995, a new Part 7 was added to Schedule P of the Fire and Casualty Annual Statement to quantify the risk-based capital loss-sensitive contract offset and to measure the premium sensitivity to losses on loss-sensitive contracts (Feldblum [20], [21]).

This appendix presents an analysis of reserving risk on retrospectively rated policies. Insurers writing excess layers of coverage or large dollar deductible policies should perform a similar analysis on those policy types.

When the retrospective rating plan contains loss limits or premium maximums and minimums, reserving risk remains, though it is dampened. These plans are more risky in some ways and less risky in other ways than traditional first dollar coverages are. The "pure insurance portion" of the plan is more risky, since

- The consideration paid by the insured is the "insurance charge", and
- The benefits paid by the insurer are the difference between (a) the value of the uncapped and unbounded premium and (b) the value of the capped and bounded premiums.[56]

The "pure insurance portion" is like excess-of-loss reinsurance, where the loss limit provides coverage like that of per-accident excess-of-loss and the premium bounds provide coverage like that of aggregate excess-of-loss. The variability of reserves for excess layers of coverage, per dollar of reserve, is generally greater than the corresponding variability of reserves for first dollar coverage.

If the retrospectively rated policy is considered as a whole—(both the insurance portion and the "pass-through" portion)—the retrospectively rated plan is less risky, per dollar of loss, than

[56] "Caps" refer to the loss limits; "bounds" refer to the premium maximums and minimums. "Ratable losses" are paid by the insurer but reimbursed by the employer, so there is no insurance risk. Acquisition expenses, underwriting expenses, and adjustment expenses are paid by the insurer but reimbursed in the basic premium and in the loss conversion factor, again eliminating much of the risk to the insurer.

traditional first dollar coverage. In fact, if there are no loss limits and no maximum or minimum bounds on the premium, then the insurance contract becomes simply a financing vehicle and the insurance company serves as a claims administrator, not as a risk-taker. There is no underwriting or reserving uncertainty at all, though there is still "credit risk" (see Greene [31]).

Premium Sensitivity

How potent are loss sensitive contracts in reducing "net" loss reserve uncertainty? (By "net" loss reserve uncertainty, we mean the variability in the insurer's total reserves, or loss reserves minus retrospective premium reserves. The "accrued retrospective premium reserves" are carried as an asset on statutory financial statements, whereas loss reserves are carried as a liability.) The answer depends on the "premium sensitivity" of the plan; that is, the amount of additional premium generated by each additional dollar of loss.

We quantify the net loss reserve uncertainty in the same fashion as we did earlier, by asking: "How does reserve uncertainty affect the financial condition of the insurer?" For instance, if the required reserves turn out to be 15% higher than our current estimates, how much additional funds will the company need to meet its loss obligations?

For business which is not written on loss sensitive contracts, the answer is simple. The additional funds needed equal the additional dollars of loss minus the amount of any implicit interest cushion in the reserves.

For business written on loss sensitive contracts, the answer is more complex, as the following illustration shows. Suppose that the indicated workers compensation reserves are $800 million. As a conservative range to guard against reserve uncertainty, the valuation actuary chooses an upper bound of $1,050 million as the worst case reserve estimate. The actuary estimates that there would be about $200 million of implicit interest margin in this scenario, so the capital needed to guard against reserve uncertainty is $50 million.[57]

Suppose now that half of the company's workers compensation business is written on retrospectively rated policies, of two types:

- Large accounts have plans with wide swings; loss limits and premium maximums are high, so each additional dollar of loss generates about a dollar of premium.
- Small and medium-size accounts have plans with narrower swings. Loss limits and premium maximums are lower and constrain the retro premiums. On average, each additional dollar of loss generates about 65¢ of additional premium.

For the entire book of retrospectively rated contracts, the premium sensitivity is 80%; that is, each additional dollar of loss generates about 80¢ of additional premium.

How much capital should this insurer hold to guard against reserve uncertainty? Suppose the needed reserves increase to the "worst case" scenario of $1,050 million. Half of this business is written on retrospectively rated plans, and the average premium sensitivity is 80%. In other words, of the adverse loss development of $250 million, $125 million occurs on retrospectively rated business. With a premium sensitivity of 80%, adverse loss development of $125 million generates $100 million of additional premium.

We add the $100 million of additional premium to the $200 of implicit interest margin to arrive at a solvency cushion of $300 million. Since the worst case adverse loss development is $250 million, the company already has a $50 million surplus

[57]For the illustration, we assume that the company wishes to hold a margin for reserve uncertainty even greater than the implicit interest margin. The text of this paper shows that for workers compensation, this implies a very low EPD ratio.

solvency cushion in the carried reserves, so no additional capital is needed.[58]

In sum, loss sensitive contracts have potent implications for the quantification of reserve uncertainty. We examine this subject from two perspectives:

- A theoretical perspective, showing the factors affecting the risks in loss sensitive contracts, and
- A simulation perspective, showing the effects of loss sensitive contracts on our measures of reserve uncertainty.

Underwriting Risk and Reserving Risk

Before turning to reserve uncertainty, let us broaden our inquiry and ask: "To what extent do retrospectively rated policies mitigate underwriting uncertainty in general?" We can answer this question empirically, by comparing the variability of standard loss ratios and net loss ratios on a large and mature book of retrospectively rated workers compensation policies.

- *Standard loss ratios* are incurred losses divided by standard earned premium. These loss ratios are influenced by random loss occurrences and premium rate fluctuations, and they vary considerably over time.
- *Net loss ratios* are incurred losses divided by the final earned premiums, as modified by retrospective adjustments. These adjustments counteract both the random loss occurrences and the

[58]An adjustment is needed to bring the accrued retrospective premiums to present value. The magnitude of this adjustment depends on the type of retrospective rating plan. For "paid loss" retro plans, the additional premium is collected when the losses are paid, so the present value of the retro premium is less than $100 million. For "incurred loss" retro plans, the additional premium is billed and collected when the case reserves develop adversely, so a smaller adjustment is needed. In this illustration, the implicit interest margin in the loss reserves is $200 million÷$1,050 million, or 19%. If all the retro plans in this illustration are paid loss retros, and the additional premium is collected when the losses are paid, the present value of the additional premiums is $81 million.

fluctuations in manual rate levels, so the net loss ratios should be more stable over time.

Exhibit E-1 shows these loss ratios for retrospectively rated policies issued by a large workers compensation insurer. Only mature policies are used in this comparison, to ensure that the net loss ratios are not subject to significant additional retrospective adjustments.[59]

As expected, the mean loss ratios are similar for standard and net—77.0% for standard and 78.8% for net. (The net loss ratios are slightly higher, since more retrospective premiums are returned than are collected.) The variances and standard deviations, however, differ greatly. The standard loss ratios show a variance of 46.9% and a standard deviation of 68.5%. Retrospective rating dampens the fluctuations in the loss ratios, leading to a variance of 11.2% and a standard deviation of 33.4%.

Reserve Uncertainty

Exhibit E-1 deals with (prospective) underwriting risk, or the risk that future underwriting returns will be lower than anticipated. Let us return now to reserving risk. We ask "To what extent is adverse development on existing losses mitigated by loss sensitive contracts?"

To resolve this issue, we must know the premium sensitivity of the retrospective rating plans, or the amount of additional premium received for each dollar of additional loss. Let us examine the variables that affect the premium sensitivity: the plan parameters, the current loss ratio, and the maturity of the reserves.[60]

---

[59]The exhibit in this paper, along with the variances and standard deviations, was produced by Miriam Perkins. An earlier exhibit from the same book of business, produced by Dr. J. Eric Brosius, was provided by the authors to the American Academy of Actuaries task force on risk-based capital. It was used by the Tillinghast consulting firm to support the recommendations of the task force regarding the loss-sensitive contract offset to the reserving and underwriting risk charges in the NAIC risk-based capital formula.

[60]Compare Bender [4, p. 36]: "The aggregate premium returned to a group of individual risks that are subject to retrospective rating depends upon the retrospective rating formula, the aggregate loss ratio of the risks, and the distribution of the individual risks' loss ratios around the aggregate."

Plan Parameters

If the retrospective rating plan had no loss limits and no constraints on the final premium, the premium sensitivity would equal the loss conversion factor times the tax multiplier, which is generally equal to or greater than one. In most cases—and particularly for smaller risks—the loss limits and the premium maximums constrain the swing of the plan, and the premium sensitivity is lower than one.

Generally, larger insureds choose retrospective rating plans with wide swings, while smaller insureds choose more constrained plans. To quantify premium sensitivity, therefore, the book of business should be divided into relatively homogeneous groups by size of risk, such as between medium sized risks and "national accounts."[61] (Small risks rarely use retrospective rating plans.)

The differences are dramatic. National accounts in our own book of business, with annual premium of $2 million or more per risk, almost always have wide swing plans, and the average premium sensitivity is close to one. Medium sized risks in our

---

There are several additional items which should also be examined for a complete analysis of the effects of loss-sensitive contracts on reserve uncertainty. As noted earlier, we should look at the effects of "incurred loss" retros versus "paid loss" retros on the implicit interest margin in the accrued retrospective premiums. To be conservative, we assume here that all plans are paid loss retros; since the additional loss payments and the additional premium collections occur at the same time, we simply net them out. Incurred loss retros would show even greater dampening of the loss reserve uncertainty; since the premiums have less implicit interest margin, the effective premium sensitivity is greater than a nominal dollar analysis indicates.

In addition, a complete analysis should look at the effects of the plan parameters on the credit risk of the company and on the size of the implicit interest margin. The accrued retrospective premiums are a receivable, not an investable asset. As is true for losses, they are held on statutory financial statements at ultimate value, not at present value. If loss reserves are backed by accrued retrospective premiums, then either these premium reserves should be reduced to present value or the implicit interest margin in the loss reserves should be reduced.

[61] This subdivision of the data by size of insured or by "underwriting market" is generally available in company files. Of course, if the company keeps data by type of plan (wide swing plans vs. narrow swing plans and so forth), this more accurate subdivision is preferable.

book of business, with more constrained plans, have an average premium sensitivity of about 65%.[62]

Loss Ratio

The premium maximum and the loss limits constrain the swing of the plan. Ideally, we wish to know whether adverse loss development causes the retrospectively rated premium on each policy to hit the premium maximum or the loss to hit the loss limit. However, we do not have information on each individual change in reported losses. Actuaries estimate from aggregates, not from details. We must determine which aggregate statistics are suitable predictors of the average amount of retrospective premium that will be collected.

Given the parameters of any retrospectively rated plan, the loss ratio determines whether the retrospective premium will be capped at the maximum. Given a distribution of loss ratios in a book of business, all of which are written on similar retrospectively rated plans, we can estimate the percent of plans that will hit the maximum premium. If the *shape* of this distribution does not depend significantly upon the average loss ratio of the book of business, and if we know the average loss ratio, then we can determine the percent of plans that will hit the maximum premium.

The general rule is that *premium sensitivity declines as the aggregate loss ratio increases.* During poor underwriting years,

[62]These are empirical figures, using actual ratios of retrospective premium collected to historical loss development. Bender [4], using theoretical relationships based on the NCCI's "Table M," estimates premium sensitivity for various risk sizes. Bender's analysis is a useful check on our procedure, but it is not a substitute. His analysis posits that the Table M relationships are correct and that compensation carriers actually use the NCCI Table M insurance charges to price their retrospectively rated policies. In practice, insurers use a variety of plans for their large insureds, and they often negotiate the loss limits, premium maximum, and plan parameters in each case for their national accounts.

As emphasized in Howard Mahler's [45] discussion of Bender's paper, the premium sensitivity is strongly dependent on the size of the risk. Bender analyzes primarily small risks, where the premium sensitivity is weak. The sensitivity rises rapidly with the size of the risk; see especially Bender's [4] Table 5 on page 50, which shows the "slope" of the plan as a function of the "loss group," and Mahler's [45] comments on pages 76–78.

when loss ratios are higher, adverse loss development leads to less additional premium than in good underwriting years, when loss ratios are lower.

Reserve Maturity

In workers compensation, adverse loss development at early maturities stems from delayed reporting of some cases and primarily from the reclassification of less serious cases to more serious cases. For instance, almost all lower back sprains and strains are initially classified as short-term temporary total cases. Significant case reserve development is expected in the first two or three years, as some of these claims develop into permanent partial or permanent total cases. Much of this development is within the "ratable" area of the retrospective rating plan; for instance, a $10,000 claim is reclassified as a $100,000 claim, so premium sensitivity is high.

At later maturities, adverse loss development stems primarily from re-estimation of the costs of permanent cases. For a plan with low or even moderate loss limits, most of the adverse loss reserve development after five or six years occurs in the "non-ratable" portion of the retrospective rating plan. For instance, a $300,000 claim may be re-estimated at $400,000, when it becomes evident that the worker will not soon be returning to work. For plans sold to medium-sized employers, the premium sensitivity for this change is generally low.

Furthermore, many companies "close" their retrospective rating plans after, say, six or seven years, with a final accounting between the company and the insured. Adverse development occurring after this date would not affect the retrospective premiums.[63]

[63] Retrospectively rated plans sold to large accounts are frequently kept open for longer periods. In fact, plans sold to "national accounts" are often kept open indefinitely, or at least until the insurer and the employer agree on a final reckoning.

Effects on the Simulation

For the simulation, we use premium sensitivity factors based on observed long-term patterns by market and by reserve maturity in our countrywide book of business.[64] From the empirical data we produce two curves, each showing premium sensitivity by reserve duration, one for national accounts and one for medium-sized risks. We weight these two curves by the volume of business in these two markets.

In the simulation analysis, we first repeat the steps outlined earlier. Based upon historical experience, we estimate (deterministically) the amount of case reserves associated with each cumulative paid loss amount at each duration. From the change in reported losses, we determine the change in retrospective premiums, and thereby the change in "net reserves."

The effects of loss sensitive contracts vary greatly by type of plan and by company practice. Several reviewers of drafts of this paper have pointed out to us: "Your company writes primarily large accounts and uses highly sensitive, wide swing plans. For this type of business, the net reserve uncertainty is clearly mitigated. What about other companies, which use less sensitive plans, recognize the adverse development later, and close their plans after several years? Would they also show a significant reduction in net reserve uncertainty?"

Accordingly, we made three adjustments, to model the loss sensitive contracts often used for medium-sized risks:

- We assume that the retrospective plans are relatively insensitive. For the most recent accident year, the assumed premium

[64] To avoid undue complexity, we do not consider aggregate loss ratios in the simulation analysis. To incorporate the aggregate loss ratio dimension, we would have to evaluate the effect of each simulated link ratio on the new accident year loss ratio and determine a new premium sensitivity factor for every cell in every simulation. Moreover, since we are using paid loss age-to-age factors, we would have to convert paid loss ratios to incurred loss ratios. The benefits from these refinements are far less than the additional effort.

sensitivity is 49%, with the sensitivity factor decreasing for each older accident year.

- We assume that most adverse development is recognized late, when premium sensitivity is lower.
- We assume that the plans are closed, on average, about five to ten years after policy inception. With the late recognition of the adverse development and the relative early closure of the plans, even the limited premium sensitivity is markedly reduced for older accident years.

We ran corresponding stochastic simulations for the loss-sensitive book of business. Even with the assumptions listed above, the projected reserve distribution is more compact, and there is less "reserve uncertainty." Specifically, the use of loss sensitive contracts reduces the standard deviation of the reserve realizations by about 35%, and it reduces the capital needed for a 1% EPD ratio by about 20%.

EXHIBIT E-1

Workers Compensation
Retrospectively Rated Policies

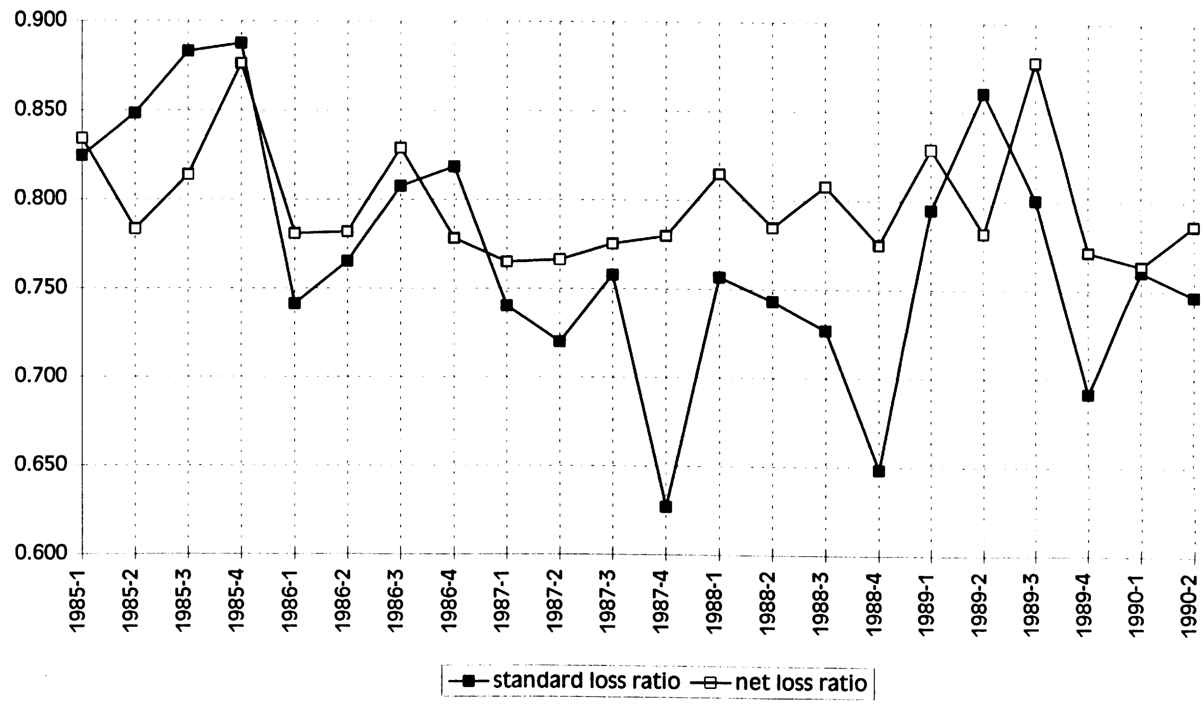

APPENDIX F

PARAMETER UNCERTAINTY IN RESERVE ESTIMATES: THE KREPS PROCEDURE

The analysis in this paper estimates the uncertainty in workers compensation loss reserves. The text and the other appendices explain the method and its rationale, and they provide the simulation equations in sufficient detail that practicing actuaries can replicate our results. Most elements of our procedure are easily visualized, so that the intuition behind each step is clear.

This is less true of the Kreps parameter risk estimation process. The procedure itself is relatively new, having first appeared in the 1997 issue of the *Proceedings of the CAS*. The simulation equations that are shown in Appendix C are taken directly from Kreps [42], which provides the justification for this process. These equations are not self-explanatory, and we have not reproduced the derivations that Kreps provides. Moreover, the magnitude of the parameter risk depends on the choice of the Bayesian prior selected by the analyst, which can be a difficult decision. To aid the reader in understanding our procedure, this appendix provides an intuitive overview of parameter risk and of the Kreps estimation process.

Actuaries generally distinguish between two sources of uncertainty: process risk and parameter risk. Process risk is the risk that actual results will differ from our expected results because of random loss occurrences. Parameter risk is the risk that our expected results are not the true expected results because we have misestimated the parameters of our distributions.

Process risk can generally be estimated directly, as long as one properly identifies all the sources of process risk. In the analysis in the text of this paper, we consider the process risk from age-to-age link ratios, from loss development tail factors, from future interest rates, and from future inflation rates. Parameter risk is more difficult to quantify. Some actuaries would argue that it is impossible to quantify completely, since any estimate of parameter risk relies on assumptions about the nature of the distributions.

In this paper, we use a procedure developed by Kreps [42] to estimate parameter risk. The mathematically adept reader is referred to Kreps's 1997 *Proceedings* paper, which is the basis for the simulations which we use. Appendix C shows the equations that we used in the simulations in incorporate parameter risk. Kreps provides a similar but independent analysis of Homeowners reserve uncertainty, using lognormal distributions of paid loss age-to-age link ratios. Kreps uses fewer data points and a more diffuse Bayesian prior, thus magnifying the parameter risk compared to the process risk. However, workers compensation has much larger paid loss development factors than Homeowners, and the development extends over a much longer period, so the total reserve uncertainty is greater in our analysis than in his.

This appendix does not purport to summarize Kreps's paper, which is already a succinct and clear exposition of a complex topic. Rather, this appendix provides a non-mathematical "intuitive" explanation of what we are doing. It explains where the parameter uncertainty resides in our analysis, what aspects of the parameter uncertainty we purport to measure, how we do so, and what choices we make in the estimation process.

Parameter Risk

Process risk and parameter risk are frequently discussed in relation to policy pricing, particularly for estimating needed profit margins and risk loads. We briefly summarize the pricing distinction between these two sources of risk, and then we extend the distinction to loss reserving.

In traditional ratemaking, the pricing actuary estimates the mean of future loss costs. This mean is based on both historical data, such as two or three years of experience, and various adjustment factors, such as development factors and trend factors.

The traditional procedure gives an expected mean for future loss costs, frequently called a "best-estimate."[65] The traditional procedure does not indicate how much uncertainty is associated with the expected future loss costs.

The uncertainty can arise from two sources: process risk and parameter risk. The pricing actuary is setting a premium rate, which considers only the expected value of the future loss costs. But losses are random occurrences, and actual losses will almost never precisely equal the expected losses. Process risk is the risk that actual losses will differ from the true expected losses.

The total pricing uncertainty, however, is the risk that actual losses will differ from our estimate of future loss costs, not from the true expected loss costs. Parameter risk is the risk that our estimate of future loss costs differs from the true expected loss costs. Parameter risk arises because the components of our pricing procedure are estimates, not known values. This is clear for such items as trend factors, since we can only estimate the effects of monetary inflation and other "social" influences on insurance losses. This is equally true, though, for our historical data. The pricing actuary begins with past experience, which he or she trends to a future policy period. In truth, the pricing actuary wishes to begin with the expected past experience, or the losses that were expected in the historical experience period. Sometimes the actual past losses are the best estimate of the expected past losses. At other times, the pricing actuary makes explicit corrections to actual past experience; the smoothing of catastrophe experience and the credibility weightings of historical loss ratios are two examples of this. Parameter risk includes the risk that

[65] In fact, this estimated mean may not be the true "best estimate"; that is, it may not be the true mean of the estimated distribution. This is because the distributions used to generate the future loss costs, such as the distribution of historical losses, the distribution of development factors, and the distribution of trend factors, are often highly skewed and correlated. For example, the trend factor used in ratemaking is the product of several annual trend rates, and these rates are autocorrelated. The mean of the product of several skewed and correlated distributions is not the same as the product of the means of these distributions.

the historical experience was not the expected experience even in the past.

Parameter Risk: Reserving

Loss reserve estimates show the same two sources of uncertainty. Chain ladder loss development methods derive age-to-age link ratios from past experience and use them to estimate future development. Process risk is the risk that actual loss development link ratios experienced in the future will differ from the true expected link ratios, since the occurrence of IBNR claims, the durations and the extent of disability on known claims, and the decisions of hearing officers and courts on contested claims are all unknown factors that influence the ultimate losses.

Traditional reserve analyses use the average historical link ratios as estimates of future ones, adjusted perhaps for outlying observations, "high" and "low" values, and systematic changes in claims operations or in the insurance environment. In this paper, we do not project "best-estimate" age-to-age link ratios. Instead, we use the historical link ratios to estimate the distribution from which future link ratios may emerge. We assume that the actual link ratios in any given development period are members of a lognormal family. We fit the parameters of the lognormal curve for each development period from the historical observations.

Parameter risk may take several forms. Some types of parameter risk are dealt with in other parts of our simulation procedure. For instance, the traditional reserve analysis is hampered by the possibility that changes in inflation rates will modify the distribution of link ratios. Our simulation procedure makes this risk explicit by stochastically generating future inflation rate paths.

Another type of parameter risk is the risk that the distribution of age-to-age link ratios is better modeled by some other curve, not by a lognormal. Curve families differ in their skewness and in the thickness of their tails, which affect the future (simulated) link ratios.

This risk definitely exists; the distributions of link ratios are presumably not perfectly lognormal. To a large extent, this risk is implicitly incorporated in our parameter risk estimation procedure, since the family of all lognormal distributions probably covers most of the variability in the actual future link ratios.[66] However, the reader should be aware that we have assumed that the distribution of link ratios is lognormal.

The parameter risk that we model here is the risk that we have incorrectly chosen the parameters of the lognormal distribution. If we had an unlimited number of observations from a distribution, we would be fully confident that the fitted distribution was indeed the true distribution. With the small sample of observations in actual reserving practice, the parameters of the fitted distribution may differ from the parameters of the true distribution.

There are other possible reasons for an incorrect choice of parameters. Perhaps we chose parameters which were correct for the historical period, but the distribution has since changed. A workers compensation insurance analogue would be a change in the types of claims over time. For instance, temporary total disability claims have low paid loss link ratios, whereas permanent partial disability claims have higher paid loss link ratios. If the mix of claims has been changing from temporary total to permanent partial, this will cause a change in the overall paid loss link ratios.

In his analysis of experience rating plan credibilities, Mahler [46] divides the total expected variance into two parts: the within variance and the between variance. He further divides the within variance into two parts: the process risk for any individual, and the change of the individual's distribution over time (the fluctuation of risk parameters over time). The standard techniques

[66] As noted by Hayne [32, p. 96], "Estimates of parameter variability may address some of the uncertainty inherent in the choice of a particular distribution for the model."

for estimating within variance usually incorporate both of these types of risk.

We have followed Mahler's approach in our analysis. We have estimated the distribution of link ratios from the full historical experience. To the extent that this distribution has been changing over time, the historical observations exhibit more variance than would otherwise be the case. The process risk estimated in our paper includes both the process risk from a stable distribution as well as the risk stemming from changing distributions over time, which Mahler terms "specification risk" (private communication).

The parameter risk incorporated in our analysis is the risk that the historical sample of observed link ratios does not accurately reflect the parameters of the true distribution. The magnitude of this parameter risk depends on three items: (i) the size of the sample, (ii) the variance of the sample observations, and (iii) our prior knowledge (or our assumed prior knowledge) of the distribution of link ratios. These factors have a strong effect on our results. We explain the intuition by illustration.

Suppose that we are estimating paid loss link ratios for 24 months to 36 months. The historical experience gives us 5 observations, of 1.400, 1.450, 1.600, 1.425, and 1.500. The average of these numbers is 1.475. We presume that these observations come from a distribution with a mean of 1.475.

With only five observations, none of which is exactly 1.475, our estimate of this mean is hardly certain. The true mean is probably close to 1.475, but it could be 1.500, 1.525, or even 2.500. The more observations we have, the more confidence we would have that the true mean is close to the sample mean. In our parameter risk quantification procedure, fewer observations we have, the greater the parameter risk, and the greater the reserve uncertainty.

Similarly, the variation in our observations also affects our confidence in the sample mean. Suppose that instead of the five observations in our illustration, we had five observations of 1.200, 1.150, 1.450, 1.900, and 1.675. The sample mean is still 1.475, but now we have less confidence that the true mean is close to 1.475. We might think now that the true mean is probably between 1.200 and 1.700. Conversely, if our observations were 1.470, 1.475, 1.480, 1.473, and 1.477, we would have greater confidence that the true mean is about 1.475.

This is a simplistic explanation; the mathematically precise version is Bayesian estimation. The chance of obtaining five observations of 1.470, 1.475, 1.480, 1.473, and 1.477 from a distribution with a mean of 1.475 and a small variance is much greater than the chance of obtaining these same five observations from a distribution with a mean of 1.600 and a larger variance. If the five observations are 1.200, 1.150, 1.450, 1.900, and 1.675, the chance of obtaining these observations from a distribution with a mean of 1.475 is still greater than the chance of obtaining them from a distribution with a mean of 1.600, but it is no longer than much greater.

In Bayesian analysis, we are concerned not just with the mean and variance of our observations. Bayesian analysis looks at every individual observation. That is, we examine the likelihood of obtaining each observation from the universe of lognormal distributions.

Our prior expectations of the true mean of the distribution also affects the parameter risk. Suppose that we knew absolutely nothing about link ratios. We have no prior expectations at all. For all that we know, the true mean might lie anywhere from $-\infty$ to $+\infty$. The sample of five observations tells us something about the true mean, but we are not about to rule out any possibilities yet.

Suppose, however, that we are experienced reserving actuaries. We have a good feel for the expected link ratio in this development period for this book of business. Even before seeing any observations, we are certain that the true mean is between 1.000 and 2.000. From our reserving experience, we are fairly confident that the mean is between 1.400 and 1.600. Given the actual observations, we are much more confident that the true mean is about 1.475.

Let us return to lognormal distributions of link ratios. The intuition behind the Kreps estimation procedure for parameter risk does not depend on the type of distributions. However, the mathematics leading to Kreps's quantification equations shown in Appendix C assume a lognormal or a normal distribution of the variable which we are estimating.[67]

With our sample observations (the historical link ratios), we fit a lognormal curve and we determine the fitted parameters μ and σ. Because we have only a limited number of observations for each development period (between 5 and 25), there is significant parameter risk; that is, our fitted μ and σ parameters may not be the parameters of the true distribution. We turn to Bayesian analysis. We take the universe of lognormal distributions, and we say: "For each member of this universe of lognormal distributions, but is the chance that it would produce a sample like the one which we observe?" This is a standard likelihood question, and Kreps uses a negative loglikelihood test. Bayesian analysis allows us to invert this relationship and to say: "For the given sample of observations, what is the chance that the true distribution is any given member of the universe of lognormal distributions?"

Fitted Distribution and Predictive Distribution

To clarify what is happening, we must distinguish between the fitted distribution and the predictive distribution. Suppose that we had an infinite number of observations, so there is no error stemming from small sample size. That is to say, if all the

[67]The equations in Appendix C are for a lognormal distribution. The equations for a normal distribution are similar.

observations come from the same distribution, then the mean of the sample is almost certainly the mean of the distribution.[68]

We use the sample to fit the lognormal distribution. There is no parameter uncertainty here (or, more accurately, the parameter uncertainty is 0%), so we use the fitted distribution to generate additional values for our stochastic simulation. In this case, the fitted distribution is also the predictive distribution.

Suppose instead that we have a finite sample. Once again, we fit a lognormal distribution. Our fitted lognormal may be the exact same distribution that we fit with the infinite sample. With the finite sample, though, there is parameter risk. That is, we are not certain that the parameters of the fitted curve are indeed the parameters of the true distribution.

In this case, we do not generate future realizations from the fitted curve. The fitted curve is the most likely true distribution, but it is not the only possible true distribution. In fact, with continuous parameters, as is true in the illustrations in this paper, there are an infinite number of potential distributions.

Think of our Bayesian analysis as telling us the chance that each possible lognormal distribution is the true distribution. That is, the Bayesian analysis gives us a distribution of lognormal curves. Think of our simulations as a two stage process. First we simulate from this distribution of lognormal curves to get the particular curve that we will use. We then simulate from this lognormal distribution to get a future observation.

The "two stage process" was simply a manner of speaking; we do not actually simulate in two stages. We are simulating

---

[68]"Almost certainly" means with 100% confidence. This is not the same as "definitely." Statistically, we can be 100% sure that the mean of the sample is the mean of the distribution, yet the two means can certainly be different, even widely different. As a heuristic example, suppose that the distribution is all integers between 1 and 10. The mean of this distribution is 5.5. The probability of an observation being greater than 5 is 50% . It is clearly possible for every observation to be greater than 5, though the probability of an infinite stream of such observations is 50% to the infinite power, or 0%. This is an example where the mean of an infinite sample differs from the mean of the distribution, though the probability of this happening is 0%.

in a single stage, but we are not simulating from a lognormal distribution. We are simulating from another distribution, from a distribution with more parameters than a lognormal has.[69] This is the predictive distribution, which is used to generate future observations.

What is this distribution from which are simulating, this predictive distribution? There is a particular distribution, though it depends not only on the historical observations and the assumption that they are members of a lognormal distribution, but also on the Bayesian prior that we use in the analysis. We could consider this question empirically, as a heuristic exercise; we can't actually do this in practice. That is, we simulate several thousand, or several million, observations, and we examine the new sample to determine what distribution it comes from.

This method is good for thought experiments only; it is not feasible. Instead, Kreps shows the analytic solution: the maximum likelihoods, the Bayesian analysis, the negative loglikelihood procedure, and the formulation of the predictive distribution. One might think: "The result must be awfully complex." Yes, it is complex in the general case. But if we assume that the distribution is a normal or lognormal distribution, and if we make certain assumptions about the Bayesian prior, then the mathematics is tractable, and Kreps obtains simple equations for the simulation. These are the equations shown in Appendix C.

One view sometimes heard on this subject runs as follows: "We know that our observations come from a lognormal distribution; this is the assumption underlying the whole procedure. We are not certain about the parameters of this lognormal distribution because of the small sample size of our historical observations. This is the source of the parameter risk. This parameter risk concerns the values of the parameters of the lognormal dis-

[69]The number of "parameters" of this distribution depends on our prior assumptions about the universe of lognormal curves, or our "Bayesian prior"; we get to this in a moment.

tribution; it is not a question of what type of distribution the observations come from. The predictive distribution may not be the same as the fitted distribution, but it still must be a lognormal distribution."

This argument is specious. The predictive distribution is not a lognormal; in fact, it is not even a two parameter distribution. What kind of distribution is it? That depends on the Bayesian prior that we use in the analysis.

Bayesian Priors

We have made several references already to Bayesian priors; it's time that we defined what we're talking about. Suppose that we knew that the link ratios come from a lognormal distributions, but that we have no prior information at all about what type of lognormal distribution it is. That is to say, we know that the link ratios come from a lognormal distribution with parameters μ and σ, but we have no assumptions about what μ and σ might be. Mathematically, we say that our prior assumption about the distribution of the μ parameter is that it is uniform over all numbers. It is just as likely that it equals 1 as that it equals 100 or one million. The σ parameter must be positive, but that is the only assumption that we make, so the prior distribution is uniform over all positive numbers. In statistical jargon, we say that we have a diffuse prior. Think of this as our having no prior assumptions about the universe of lognormal distributions; every one is just as reasonable as another.

Could we use this diffuse universe of lognormal distributions as our predictive distribution? That is, if we have no observations at all, could we use this diffuse universe of lognormals? Of course not. All we know is that the desired numbers come from a lognormal distributions, but this could be any lognormal distribution at all. The predictive distribution is so diffuse that it has infinite variance. The simulations will not converge, no matter how many simulations we use.

The preceding statement warrants further explanation, since this is a problem even for simulations which do converge. Suppose that we have no observations, and we have no prior assumptions, so we simulate from the diffuse universe of lognormals. Think of this in the two stage process: we first pick parameters μ and σ by choosing a real number for μ and a positive number for σ. We have set no bounds for these numbers; they could be anything. We then simulate a realization from this lognormal; this realization is unbounded. No matter how many realization we use, the expected mean of our realizations is unbounded.

If we have some observations, the Bayesian analysis makes our posterior universe of lognormal distributions less diffuse. If our five observations are 1.400, 1.450, 1.600, 1.425, and 1.500, then it is much more likely that the true lognormal distribution has a mean of 1.475 than that it is has a mean of 10 or of 100.

Parameters for the Bayesian Prior

In practice, a completely diffuse Bayesian prior is often unworkable; moreover, it sometimes fails to make sense even in theory. To clarify the procedure used in this paper, we must examine the method of choosing the Bayesian prior in the Kreps procedure. Kreps determines μ_0 and σ_0 from the observations, and he calculates a negative loglikelihood from these values for a lognormal with parameters μ and σ (equation 2.25 on page 558). To simplify the analysis, he rescales the problem by defining normalized variables v and y such that:

$$\mu = \mu_0 + v\sigma_0$$

and

$$\sigma = y\sigma_0.$$

The Bayesian prior for the distribution of μ and σ can be restated as a prior assumption for the distribution of v and y.

Kreps [42, pp. 559–560]:

> We take a Bayesian approach and use diffuse prior distributions for v and y. Since v runs along the full axis from minus infinity to plus infinity, the prior used is just 1. Since y runs along the semi-axis, the suggested prior is proportional to $1/y^{\theta}$ where θ is either 0 or 1, depending on one's preference. The choice $\theta = 1$ emphasizes small values of y and corresponds to the assumption that the prior distribution of $\ln(y)$ is flat; the choice $\theta = 0$ assumes that the prior distribution of y is flat. Venter has emphasized that any choice of prior has strong implications. Ideally, the nature of the data being fitted would give some clues as to proper priors.

The comment by Venter referred to above is that "on a semi-axis a flat prior corresponds to assuming that it is as likely for the variable to lie between a million and a million and one as it is for the variable to lie between zero and one, and that it is infinitely more likely to be excess of any finite amount than to be less than that amount" (Kreps [42, footnote 7]).

Even with a θ of 1, our simulations produced unreasonable results. The text of our paper explains what we mean by "unreasonable." After much discussion with Dr. Kreps, we used a θ of 2. Dr. Kreps sums up the theory as follows (private communication):

> On pages 83–74 of section 3.2.2 of *Statistical Decision Theory and Bayesian Analysis*, second edition, by James O. Berger (Springer, 1980), there is the section "Noninformative Priors for Location and Scale Problems" which outlines the arguments and the problems with the Bayesian priors. The crude result is that for a location parameter, the density is 1 and for a scale parameter it is $1/\theta$. Berger goes on to talk about the Jeffreys results in the next section, which in the case

> of normals reduce to powers of sigma. Which power depends on what you like, but the choice theta = 2 is actually the computational Jeffreys result even if Jeffreys himself prefers theta = 1. So you take your choice; personally I think we always know something about the data and a noninformative prior is something like laziness on our part.

For workers compensation paid loss link ratios, we know a great deal about the data. Simply picking a value of θ is indeed laziness. The problem, however, is two-fold. First, we have great difficulty conceptualizing what any value of θ means for the universe of lognormals as potentials distributions for paid loss link ratios. Yes, we can state the mathematics, but we have difficulty visualizing whether a $\theta = 2$ is more reasonable than a $\theta = 5$ or vice versa. Second, if we use other ways of stating our prior assumptions, we can't work these assumptions into Kreps's equations.

Our final choice is summarized in the text of the paper. We chose a θ of 2, to ensure as diffuse a Bayesian as practicable for our application, and we discarded the extreme realizations with means more than 50 standard deviations away from the overall average. This may not be the ideal procedure, but we do not even know if it is too conservative or too liberal.

The Kreps parameter risk estimation procedure had one additional effect on our method. We noted above that the variance of our predictive distribution depends on both the Bayesian prior and the number of observations ("n"). Kreps discusses this problem in terms of the variance of z_{eff}, where z_{eff} is the effective deviate of ln(x), where x is the variable which we are simulating. Kreps shows that for $n + \theta \leq 4$, the variance of z_{eff} is infinite, and he notes that "this formula also tempts one to choose $\theta = 5$ so that $\text{var}(z_{\text{eff}}) = 1$ for all n" (page 564). Similarly, in discussing the standard deviation of the underlying distribution, Kreps says:

"The standard deviation does not exist if $n + \theta \leq 4$, but goes to zero as the sample size increases" (page 561).

This is the problem of convergence discussed earlier. Kreps [42, p. 561] says:

> In simulation situations if the underlying distribution does not have a finite variance then the mean of the simulation will not converge, because the mean of the simulation itself will have an infinite standard deviation. In practice, this shows up as occasional large jumps in the mean, even with millions of simulations (in fact, no matter how many simulations are done).

We choose $\theta = 2$. We deal with the variance problem by using only 20 columns of age-to-age link ratios, so that we always have a sufficient number of observations. For development beyond the 21st year, we use the inverse power curve tail factor approximation.

Conclusion

Neither the Kreps paper nor this paper is the definitive word on parameter risk. Even with the Kreps procedure, the analyst must choose a Bayesian prior based upon his or her own reserving knowledge and prejudices. Nevertheless, the thrust or the Kreps paper is that parameter risk is a significant source of reserve uncertainty. Our analysis illustrates this uncertainty, though we do not even pretend to have authoritatively measured it. However, by choosing a relatively diffuse Bayesian prior, and by discarding only those realizations that were extremely far from the sample mean, we have presumably erred on the side of caution, by overestimating the parameter risk.