© The Journal of Risk and Insurance, 1992 Vol. LIX, No 3, 426-442

An Empirical Bayes Approach to Estimating Loss Ratios

Joan Lamm-Tennant
Laura T. Starks
Lynne Stokes*

ABSTRACT

The empirical Bayes model is explored as a technique for estimating key financial variables that are meaningful to regulators, policyholders, and security analysts of property-liability insurers. The time series mean and two empirical Bayes models of the loss ratio are evaluated for four lines of business—auto liability, auto physical damage, medical malpractice, and fire. Tests were conducted of the forecasting ability of the three estimators. The empirical Bayes process performed well in the forecasting experiments for all four lines of business, especially medical malpractice. Empirical Bayes methods perform better for small firms, which have higher loss ratio variances than large firms. Forecasting tests using incurred losses suggest that the results might be improved by using a longer time series and recognizing the autoregression of the loss process.

Introduction

The accuracy of estimates of future loss costs plays a fundamental role in determining the underwriting profits of property-liability insurers. A standard measure of loss costs used by regulators, auditors, policyholders, and security analysts is the loss ratio. The well-known cyclical pattern in loss ratios over time (e.g., Witt, 1977, 1978) has led to a series of analyses of underwriting profit cycles. The underwriting cycle in the United States

* Joan Lamm-Tennant is Associate Professor in the Department of Finance at Villanova University. Laura T. Starks is Associate Professor in the Department of Finance and Lynne Stokes is Associate Professor in the Department of Management Science and Information Systems, both at the University of Texas at Austin. The authors thank Billy Charlton for excellent research assistance. The helpful comments of Pat Brockett, Larry Cox, David Cummins, Jim Garven, Thomas Kozik, Greg Niehaus, Bob Witt, and the participants of the 1990 ARIA meetings are appreciated. The article has also benefited from the comments of two anonymous referees and an associate editor. The authors are especially grateful to Bob Witt and the Gus Wortham Memorial Chair in Risk and Insurance at the University of Texas at Austin for making available the A. M. Best data tapes. Laura T. Starks acknowledges research support from the Graduate School of Business, University of Texas at Austin.

is about six years (see Cummins and Outreville, 1987; Venezian, 1985; Smith and Gahin, 1983).

Explanations for the cycles fall into two categories. The first type suggests that the insurance markets are unstable such that prices fail to converge on an equilibrium due to periods of destructive competition followed by cutbacks in supply (Berger, 1988) or ratemaking with limited information (Venezian, 1985). For example, Brockett and Witt (1982) point out that an explanation for the autoregressive behavior of loss ratios is that premiums are based in part on past losses.

A second type of model explains the cycle in terms of the insurance market response to external events, such as the liability insurance crisis (Harrington, 1988), interest rate changes (Doherty and Kang, 1988; Doherty and Garven, 1991), and institutional and regulatory rigidities (Cummins and Outreville, 1987).

This article contributes to the analysis of underwriting profits and cycles by exploring the short-term forecasting of measures of underwriting profits. Specifically, the empirical Bayes model is proposed as a methodology for estimating loss ratios—the ratio of incurred losses to earned premiums.[1]

The position taken is that of an outsider (such as a security analyst, rating agency, auditor, or insurance regulator) attempting to gauge the financial performance of an insurer. We use the loss ratio rather than the profit ratio because loss costs are the primary source of uncertainty in the determination of insurance profits. Although this analysis focuses on the loss ratio, any financial measure could be estimated using the empirical Bayes method. Hence, this analysis is intended to be an example of a technique appropriate for predicting financial measures when evaluating the financial performance of the property-liability insurer.

The empirical Bayes approach is particularly appropriate for the case of loss ratios reported by the property-liability insurance industry because of data availability and comparability. The by-line loss ratio can be derived for a large number of insurers across time using the A. M. Best tapes. The empirical Bayes methodology was developed for just such a multiparameter estimation problem (Efron and Morris, 1973, 1975, 1977; Morris, 1983) and is designed for situations in which many similar parameters are to be estimated but the information on each may be weak. The procedure "borrows strength" (Tukey, 1963) from the whole set of data for estimating each parameter. Thus, empirical Bayes estimators gain their advantage over frequentist estimators by using information about all parameters in estimating each individual parameter, which is intuitively reasonable if the parameters are similar. In the case of loss ratios, although there is information across many firms, the information on each firm's by-line loss ratios

[1] The inverse of the loss ratio indicates the cost per dollar of losses required to administer the insurance and is frequently modeled as the proportionate loading or transaction costs of insurance (see Doherty and Kang, 1988; Cummins and Outreville, 1987).

may be weak. The information within the firm is considered weak when the historic trend is short or if the business is not widely underwritten such that the historic trend has high variance. An empirical Bayes model suggests that the estimate of a firm's by-line loss ratio may borrow strength from information provided by the experience reported across all firms with business in that line.

The empirical Bayes model has been previously applied for actuarial purposes. For example, Morris and Van Slyke (1978) developed an empirical Bayes method for pricing insurance claims. Similarly, some applications of credibility theory have employed the empirical Bayes methodology (see, for example, Neuhaus, 1984; Norberg, 1980). But the concern here is one of an outsider with limited information. In this regard, our application of the empirical Bayes methodology to measures of profitability in the property-liability insurance industry is unique.

The remainder of this article is organized as follows: First, the data and methodology are described. The next section contains the presentation and discussion of the results, followed by an examination of the sensitivity of the results to changes in the assumptions and methods. The final section contains the conclusions.

Data and Methodology

The initial data set consists of premiums earned and losses incurred for four lines of business—auto liability, auto physical damage, medical malpractice, and fire—for all property-liability insurance companies listed on the A. M. Best tapes for 1980 through 1987. Since the data are derived from the annual statements, the loss ratios are based on incurred losses and earned premiums. In principle, it would be better to use accident year incurred losses, because calendar year losses reflect adjustments in reserves for previous years as well as the loss experience of any particular year. An argument for using calendar year losses in the present application is that accident year losses are not as accessible to outsiders as calendar year losses. However, the empirical Bayes approach could be easily applied to accident year data.

Each line of business sample consists of only those companies with complete data for the eight-year period. The aim is to predict the loss ratio for each line of each company in 1987, given their loss ratios in the previous seven years. Assuming that it is constant across time, the mean of each company's yearly loss ratio distribution for a line would be a reasonable predictor. These means are not known, but they can be estimated from the available data. A logical approach would be to use the time series mean of the first seven yearly loss ratios for each company as an estimator of its mean loss ratio for that line, and thus as a predictor for the eighth year. The original prediction problem can then be described as an estimation problem, in which the parameters to be estimated are population means, and the estimators are time series means.

Unfortunately, estimators of means based on samples of size seven are frequently so unstable that they do not perform well, especially for lines with large variances, such as medical malpractice. Thus, the recommended approach to the prediction problem is to use an empirical Bayes estimator of each company's loss ratio mean as the predictor for its 1987 loss ratio.

The Empirical Bayes Model

The empirical Bayes model and methodology is now briefly described. Let $\mu_1, \ldots, \mu_k$ denote the parameters of interest, which are means of k distinct populations. Let y_{ij}, $j = 1, \ldots, n_i$ denote a time series sample from the ith of these populations, and define $\bar{y}_i = \Sigma_{j=1}^{ni} y_{ij}/n_i$ as the time series mean. Assume that

$$\bar{y}_i | \mu_i \sim N(\mu_i, V) \text{ for } i = 1, \ldots, k \text{ independently,} \tag{1}$$

with V_i known, where $N(\cdot,\cdot)$ denotes a normal distribution having the specified mean and variance. Then $V_i = \sigma_i^2/n_i$, where $\sigma_i^2) = Var(y_{ij})$. Statement (1) describes the direct information available from the data about the parameter μ_i. Further assume that the μ_is are independently distributed as

$$\mu_i \sim N(\mu, A),\ i = 1, \ldots, k. \tag{2}$$

Statement (2) is the prior distribution of the parameters $\mu_1, \ldots, \mu_k$. The parameters of this distribution are not assumed known as they would be in a full Bayesian model; instead they are to be estimated from the data. Standard Bayes calculations (for example, Winkler, 1972) show that the marginal distribution of the $\bar{y}_i$s can be determined from equations (1) and (2) to be

$$\bar{y}_i \sim N(\mu, A + V_i), \text{ for } i = 1, \ldots, k, \text{ independently,} \tag{3}$$

and that the posterior distribution for each μ_i is then

$$\mu_i | \bar{y}_i \sim N((1\text{-}B)\bar{y}_i + B_i\mu, (1\text{-}B_i)V_i), \text{ for } i = 1, \ldots, k, \text{ independently,} \tag{4}$$

where

$$B_i = V_i/(A+V_i). \tag{5}$$

The posterior mean, $E(\mu_i|\bar{y}_i) = (1\text{-}B_i)\bar{y}_i + B_i\mu$, is called a Bayes estimate for μ_i, but this estimator is not available since μ and the B_is are unknown. They can be estimated from the $\bar{y}_i$s by considering their marginal distribution shown in equation (3), however, and their estimates ($\hat{\mu}$ and $\hat{B}_i$) substituted for the unknown values in $E(\mu_i|\bar{y}_i)$. The resulting statistic,

$$\hat{\mu}_{EBi} = (1\text{-}\hat{B}_i)\ y_i + \hat{B}_i\hat{\mu}, \tag{6}$$

is known as an empirical Bayes estimator. A method for estimating $\hat{\mu}$ and the $\hat{B}_i$s is described in the Appendix.

In the special case in which it may be assumed that for all i, $V_i = V$, $\hat{\mu}_{EBi}$ from equation (6) becomes

$$\hat{\mu}_{EBi} = (1-\hat{B})\, \bar{y}_i + \hat{B}\hat{\mu}, \tag{6a}$$

where $\hat{B}$ estimates $B = V/(A+V)$. This situation—the equal variance case—can occur when both the variance (σ_i^2) and the sample size (n_i) from the k populations are identical. In the equal variance case, the estimators of μ and B reduce to

$$\hat{\mu} = \sum_{i=1}^{k} \bar{y}_i/k$$

and

$$\hat{B} = \frac{(k-3)}{(k-1)}[V/(V+\hat{A})],$$

where $\hat{A} = \sum_{i=1}^{k}(\bar{y}_i - \hat{\mu})^2/(k-1) - V)$. $\hat{\mu}_{EBi}$ as defined for the equal variance case is known as the James-Stein estimator (James and Stein, 1960). It can be shown that

$$E(\hat{\mu}_{EBi} - \mu_i)^2 \leq E(\bar{y}_i - \mu_i)^2 \tag{7}$$

for i = 1,..., k, where the expectation is taken over the joint distribution of $\bar{y}_i$ and μ_i, and

$$E[\sum_{i=1}^{k}(\hat{\mu}_{EBi} - \mu_i)^2] \leq E[\sum_{i=1}^{k}(\bar{y}_i - \mu_i)^2], \tag{8}$$

where the expectation is taken over the data alone. This latter property should appeal even to frequentists, since it states that the empirical Bayes estimators dominate the $\bar{y}_i$s as estimators of the μ_is, even when taking the parameters as being held fixed, if the measure of performance is the sum of squared errors over all parameters.

Similarities exist between credibility theory and this empirical Bayes application. Credibility theory generally estimates future claims through the use of a weighted average of actual claims for the group under consideration and expected claims based on prior or similar experience (Miller and Hickman, 1975; Klugman, 1992). Jewell (1975) and others noted that credibility theory can be at least an approximation of a linearized Bayesian forecasting method, where the actuary provides the prior distribution to be used. Norberg (1980) extended Jewell's results to demonstrate that the credibility estimators can be derived as an analogue to the empirical Bayes estimator. Thus, rather than using an actuary's determination of the appropriate prior distribution for calculating expected claims, Norberg's model explicitly uses information about similar classes of claims.

The empirical Bayes framework is used for this application, where μ_i in equation (1) denotes the mean of the yearly loss ratios for the ith company in a given line, and $\bar{y}_i$ is the time series mean of its loss ratios over the seven years 1980 through 1986. The estimator $\hat{\mu}_{EBi}$ as shown in equation (6) is then a weighted average of $\bar{y}_i$ and $\hat{\mu}$, the estimated prior mean of the loss ratio process for all companies in the line. The weight depends on the relative precision of the available information about μ_i and μ. The precision of the direct information about μ_i is measured by $V_i = \sigma_i^2)/n_i$, while that of μ is measured by A. In this application, since $n_i = 7$ for all companies in the line, the size of V_i is determined by the variability of the loss ratios across years for company i in a given line of business. The prior variance A describes the variability of loss ratio means among the companies within a line of business and is a reflection of the similarity among the μ_is. The larger V_i is relative to A (i.e., the more unreliable is the information in $\bar{y}_i$ compared to the uniformity of companies within the line), the more the empirical Bayes estimator relies on information about the prior mean. Conversely, the smaller V_i is relative to A, the more heavily the estimator weights the individual time series mean estimate.

Assumptions and Limitations

Although the theory has assumed V_i to be known, in practice it is not. However, σ_i^2 can be estimated from the data in the usual way, as $\hat{\sigma}_i^2 = \sum_{j=1}^{ni}(y_{ij} - \bar{y}_i)^2/(n_i - 1)$, and the resulting estimate of V_i substituted in equation (5). This procedure performs well when the n_is are large. In this application, however, $n_i = 7$ for all i, so the estimates of V_i may be poor. Although the scarcity of observations might imply that $\hat{\mu}_{EBi}$ would not perform as well in this application as the theory would predict, as long as it can be assumed that the loss ratio variances were constant across the companies within a line (that is, that $\sigma_i^2 = \sigma^2$) , then V can be estimated by $\hat{V} = \hat{\sigma}^2/n$, where

$$\hat{\sigma}^2 = \sum_{i=1}^{k} \sum_{j=1}^{ni} (y_{ij} - \bar{y}_i)^2 / (\sum_{i=1}^{k} n_i - k). \qquad (9)$$

This pooled estimator of the variance has excellent precision in this application. The problem with this approach is the uncertainty of the assumption of equal variances for all companies within a line (for evidence on this assumption, see Lamm-Tennant and Starks, 1991). The analysis was carried out using both approaches, and similar results were obtained.

Another potential problem with models (1) and (2) for this application is the normality assumptions. The assumption of normally distributed losses or loss ratios is known to be problematic (e.g., Cummins, 1991). However, the time series mean of seven possibly non-normal observations (which is the object of statements (1) and (3)) may still be approximately normal, as justified by the Central Limit Theorem (but see Brockett, 1983, for a discussion of the problem with this assumption). A check of the marginal normality of the $\bar{y}_i$s was made by examining a normal probability plot for each line. The plots confirmed that assumption (3) is reasonable for the medical malpractice and auto physical damage lines but showed that $\bar{y}_i$ has a distribution with a longer than normal right tail in the auto liability and, to a lesser degree, the fire lines.[2]

Empirical Bayes estimation performs best when the ensemble of means (the k means that are estimated together) are as similar as possible. For that reason, one would not want to form an ensemble by pooling across lines of business, since variability across lines is larger than variability within. It also would be beneficial to exclude from the ensemble those companies whose average loss ratios are known to be extremely unusual. To control for the influence of outliers and nonrepresentative firms, the upper and lower five percentiles of the distributions of loss ratio means and loss ratio variances were excluded from the estimation process. The extremes for the 1987 data were not eliminated, but any absolute loss ratio greater than 5.0 was eliminated.[3] The resulting sample comprised 297 firms in the auto liability line, 315 in auto physical damage, 75 in medical malpractice, and 343 in fire.

Comparison of Forecast Performance

One measure used to compare the forecast performance of the empirical Bayes estimators and the time series mean is the mean squared error of prediction for the estimators:

$$\mathrm{MSE}(\hat{y}_{i,87}) = \sum_{i=1}^{k} (\hat{y}_{i,87} - y_{i,87})^2/k,$$

where $\hat{y}_{i,87}$ is the predictor of $y_{i,87}$, the observed loss ratio for the ith company in the line in 1987. Because squared error may not always be the loss function of interest, the number of companies within each line for which each predictor works best (that is, the number of companies for which each predictor lies closest to the true value) was also examined.

Results

Based on the 1980-1986 loss ratio data, three techniques were used to predict the 1987 loss ratio for the companies within each of the four lines of business. Two are empirical Bayes estimators of the mean, and one is the time series mean. The two empirical Bayes procedures differ in their assumptions about V_i, the individual firm's variance over time. For the first predictor, $\hat{\mu}_{EBi}$, the V_is are allowed to differ, and the estimator defined in equation (A.2) is used. For the second predictor, μ^*_{EBi}, it is assumed that $V_i = V$ for $i = 1,\ldots, k$. $\hat{\mu}^*_{EBi}$ is defined in equation (6a).

---

[2] Given the skewed data for the auto liability and fire lines, we also tried a power transformation with no improvement in the results.

[3] Ratios of that size can be attributed to reserve errors, data transcription errors, withdrawal of a company from a given line of business, and other problems indicating that the firm is not a full market participant. Such outliers can unduly influence the mean squared error criterion.

Table 1 describes the distributions of the observed loss ratios for 1987 for each of the four lines of business. The loss ratios in the medical malpractice line were most dispersed in 1987, suggesting that it is the most difficult line to predict. The interquartile range of the medical malpractice line was 0.52, while the interquartile range of the other three lines ranged only from 0.16 to 0.23.

Table 1
1987 Loss Ratios: Summary Statistics

	k	Min	10%	25%	50%	75%	90%	Max
Auto Liability	297	−0.46	0.54	0.64	0.74	0.82	0.96	1.91
Auto Physical Damage	315	−1.34	0.34	0.45	0.54	0.61	0.69	.98
Medical Malpractice	75	−0.26	0.24	0.49	0.65	1.01	1.68	1.98
Fire	343	−1.73	0.16	0.29	0.40	0.52	0.67	.62

Note: Numbers in column headings refer to percentiles of the distribution of loss ratios across time.

Table 2 reports the results of the mean squared error comparison for the three predictors for the four lines of business. It shows that MSE($\hat{y}_{1,87}$) for $\hat{y}_{1,87} = \bar{y}_i$, $\hat{\mu}_{EBi}$, and $\hat{\mu}^*_{EBi}$. The relative savings in mean squared error per company in the line from each of the empirical Bayes procedures is also included in Table 2. Relative savings is defined as $RS(\hat{\mu}_{EBi}, \bar{y}_i) = [MSE(\bar{y}_i) - MSE(\hat{\mu}_{EBi})] / MSE(\bar{y}_i)$, and similarly for $\hat{\mu}^*_{EBi}$. The relative savings of both of the empirical Bayes predictors were positive for all four lines of business, with $\hat{\mu}_{EBi}$ performing better than $\hat{\mu}^*_{EBi}$ in every case. The performance of the two empirical Bayes procedures was similar for the auto liability, auto physical damage, and fire lines, where the gain from the new methods was small. For the medical malpractice line, however, the gain was substantial. Further, the empirical Bayes estimator for the unequal variance assumption $\hat{\mu}_{EBi}$ performed considerably better than did the estimator for the equal variance case $\hat{\mu}^*_{EBi}$, with an improvement in relative savings from 59 to 74 percent.

Table 2
Mean Squared Forecast Errors: Loss Ratios

	Mean Squared Errors of:				
	$\bar{y}_i$	$\hat{\mu}^*_{EBi}$	$\hat{\mu}_{EBi}$	$RS(\hat{\mu}^*_{EBi}, y_i)$	$RS(\hat{\mu}_{EBi}, \bar{y}_i)$
Auto Liability	.144	.138	.135	4.17%	6.25%
Auto Physical Damage	.046	.045	.044	1.87	3.74
Medical Malpractice	1.715	.703	.452	59.00	73.60
Fire	.136	.130	.127	4.55	7.08

Note: Numbers shown in the columns headed $\bar{y}i$, $\hat{\mu}^*_{EBi}$, and $\hat{\mu}_{EBi}$ are mean square forecast errors (MSEs) for each of these loss ratio predictors. $\bar{y}_i$ is the time series mean of the annual loss ratios calculated for 1980 through 1986. $\hat{\mu}^*_{EBi}$ is the empirical Bayes estimate of the annual loss ratio calculated for 1980 through 1986, assuming that the loss ratio variance differs across firms. $\hat{\mu}_{EBi}$ is the empirical Bayes estimate of the annual loss ratio calculated for 1980 through 1986, assuming that the loss ratio variance is identical across firms. $RS(\hat{\mu}_{EBi}, \bar{y}_i)$ is the relative savings of the empirical Bayes estimate over the time series mean calculated as $[MSE(\bar{y}_i) - MSE(\hat{\mu}_{EBi})]/MSE(\bar{y}_i)$, where MSE is the mean squared error, and k is the number of companies in the line.

The assumption of equal variances across companies within a line is not advisable, since $\hat{\mu}_{EBi}$ performed better than $\hat{\mu}^*_{EBi}$ in all cases. It seems that the poor estimates of V_i from the short series of data on each company are less problematic than an unwarranted assumption of equal variance. Given this result, the remaining analysis is limited to the unequal variance case.

Table 3 reports, for each line of business, the estimated prior parameters μ and A under the unequal variance model. The table also shows the average of the estimated individual variances and shrinkage factors. The most striking observation from the table is that medical malpractice loss ratios are on average much less stable over time within a firm (large average V_i) than the other three lines. This is consistent with Table 1. In addition, the medical malpractice line has the largest prior variance. Thus, medical malpractice has the most uncertainty across time and across firms. In contrast, the most stable line, both over time and among firms, is auto physical damage. As one would expect, auto physical damage is less uncertain than medical malpractice in the pricing and underwriting process due to the nature of the underlying peril.

Table 3
Empirical Bayes Parameter Estimates
(Unequal Variance Model)

	$\hat{\mu}$ Prior Mean	$\hat{A}$ Prior Variance	Average $\hat{V}_i$ Conditional Variance	Average $\hat{B}_i$ Shrinkage Factor
Auto Liability	0.73	0.008	0.004	.28
Auto Physical Damage	0.65	0.005	0.002	.24
Medical Malpractice	0.83	0.147	11.413	.36
Fire	0.50	0.015	0.042	.21

Note: $\hat{\mu}$ is from equation (A.2). $\hat{A}$ is from Equation (A.1). $\hat{V}_i = \hat{\sigma}_i^2/n_i$, where $\hat{\sigma}_i^2$ is defined in equation (9).

Sensitivity of the Results

This section analyzes the sensitivity of the results to variations in the conditions. The applicability of the mean squared error criterion is considered, and an alternative criterion to measure the performance of the empirical Bayes estimator against the performance of the time series mean is used. Because there may be systematic differences in by-line loss ratios due to the amount of business in the line, the results are tested recognizing differences in by-line premiums earned across firms. Finally, given the autoregressive tendencies of the loss ratios, the robustness of the results is tested by applying the methodology to losses.

Alternative Performance Measure

The superior performance on the mean squared error criterion for an empirical Bayes procedure is not surprising in light of equation (8). The mean squared error criterion provides a measure of the average performance of each estimator. An alternative criterion is a measure of the performance of the estimators across the firms. This second measure of comparison is shown in Table 4, which displays the number and percentage of companies within each line of business for which the empirical Bayes predictor is closer to $y_{i,87}$ than its competitor $\bar{y}_i$. The advantage for the $\hat{\mu}^*_{EBi}$ is maintained for all four lines of business.

Table 4

Predictive Accuracy Based on Number of Companies

	Auto Liability	Auto Physical Damage	Medical Malpractice	Fire
Number of companies	297	315	75	343
Percent of Companies for which $\lvert\bar{y}_i - y_{i,87}\rvert > \lvert\hat{\mu}_{EBi} - y_{i,87}\rvert$	53.9	52.1	58.7	56.3

Note: $\bar{y}_i$ is the time series mean of the annual loss ratios calculated for 1980 through 1986. $\hat{\mu}_{EBi}$ is the empirical Bayes estimate of the annual loss ratio calculated for 1980 through 1986, assuming that the loss ratio variance differs across firms.

Smaller and Larger Firms

The heterogeneity of the by-line loss ratio across firms is also examined. The empirical Bayes methodology improves with greater homogeneity in the variables to be estimated. The concern is that there may be systematic differences in firms' variances of loss ratios due to differences in the premiums earned in the line. To check the performance of the methodology with presumably more homogeneity, the data are divided into two subsets according to the premiums volume in each line of business. Smaller premium firms were those with premiums in the line less than the median; those with premiums in the line greater than the median were considered larger premium firms.

Table 5 reports the empirical Bayes estimators for each line with the firms divided into smaller and larger premium groups. In all four lines of business the conditional variance is substantially greater for the smaller premium firms than for the larger premium firms. This is particularly true for the medical malpractice line, where the conditional variance is 16.0784 for smaller premium firms as compared to 0.1798 for larger premium firms. Hence, the concern for heterogeneity is justified. Observation of the by-line prior mean for smaller premium firms relative to the larger premium firms reveals no meaningful differences in the mean loss ratio across the two groups. The empirical Bayes methodology is applied separately within each size group.

Table 5
Empirical Bayes Parameters Estimates by Size of Firm
(Unequal Variance Case)

	$\hat{\mu}$ Prior Mean	$\hat{A}$ Prior Variance	Average $\hat{V}_i$ Conditional Variance	Average $\hat{B}_i$ Shrinkage Factor
Smaller Firms				
Auto Liability	0.70	0.010	0.055	.30
Auto Physical Damage	0.64	0.005	0.032	.32
Medical Malpractice	0.81	0.333	16.078	.36
Fire	0.46	0.013	0.054	.28
Larger Firms				
Auto Liability	0.75	0.006	0.022	.25
Auto Physical Damage	0.65	0.005	0.008	.14
Medical Malpractice	0.86	0.083	0.180	.16
Fire	0.54	0.013	0.030	.18

Note: Note: $\hat{\mu}$ is from equation (A.2). $\hat{A}$ is from Equation (A.1). $\hat{V}_i = \hat{s}_i^2/n_i$, where $\hat{s}_i^2$ is defined in equation (9). Firms with premium volume less than the median for a given line are classified as smaller premium firms, while those with premium volume greater than the median are classified as larger firms.

This method of analysis seemed intuitive since estimating loss ratios may be more difficult for companies with relatively small amounts of business than for those with larger amounts of business. That is, the forecasted future by-line loss ratio for a firm with a substantial amount of business in a line may be more stable due to less volatile historical results, better diversification, and possibly a competitive advantage in selection and underwriting. The outsider interested in estimating loss ratios for a firm with relatively few premium dollars reported in a line of business may therefore benefit the most from relying upon the empirical Bayes estimator. To test this hypothesis, the mean squared errors on the maximum likelihood estimators are compared with the empirical Bayes estimators.

Given the larger shrinkage factors for smaller premium firms, one would expect a greater benefit from the empirical Bayes estimator, and that is, in fact, what is found. Table 6 shows the improvement in the mean squared errors of the empirical Bayes estimates of 1987 loss ratios over the maximum likelihood estimates (the time series means). As expected, there is more forecasting error for smaller premium companies than for larger premium firms. This holds true across all lines for both estimation techniques. For example, for auto liability, the mean squared error for the maximum likelihood estimator is almost 16 times greater for the smaller premium firms than for the larger premium firms.

Table 6
Mean Squared Forecast Errors by Size of Firm
(Unequal Variance Case)

	Mean Squared Error of: y_i	μ_{EBi}	$RS(\mu_{EBi}, y_i)^a$ %
Smaller Firms			
Auto Liability	.271	.256	5.63
Auto Physical Damage	.071	.068	3.56
Medical Malpractice	3.251	.777	76.09
Fire	.226	.213	5.90
Larger Firms			
Auto Liability	.017	.016	10.66
Auto Physical Damage	.021	.020	5.16
Medical Malpractice	.138	.124	10.70
Fire	.046	.038	18.50

Note: y_i is the time series mean of the annual loss ratios calculated for 1980 through 1986. μ_{EBi} is the empirical Bayes estimate of the annual loss ratio calculated for 1980 through 1986, assuming that the loss ratio variance differs across firms. $RS(\mu_{EBi}, y_i)$ is the relative savings of the empirical Bayes estimate over the time series mean calculated as $[MSE(y_i) - [MSE(\mu_{EBi})]/MSE(y_i)$, where MSE is the mean squared error.
[a]Numbers in this column cannot be replicated from numbers in the table due to rounding.

Although the absolute improvement in the mean squared error is greater for the smaller premium firms across all lines, the results for medical malpractice are particularly interesting.[4] Our initial concern was primarily with medical malpractice since it exhibits meaningful differences in the variance of the loss ratio across firms based upon amount of premiums in the line. As reported in Table 6, the relative savings (percent improvement) for smaller premium firms in medical malpractice is 76.09, as compared to 10.66 for larger premium firms. When all firms are modeled, assuming firms have different loss ratio variances, the relative savings for medical malpractice is 73.6 percent. One can conclude that, for medical malpractice, the improvement provided by empirical Bayes is largely attributed to firms with smaller amounts of premiums in the line.

The Autoregressive Nature of the Loss Ratio

An additional qualification is that our results do not directly take into account the finding of underwriting cycles research that underwriting profits follow a second-order autoregressive process (Brockett and Witt, 1982; Venezian, 1985). An interesting avenue for future research would be to

---

[4] **Because the base mean squared error is so much greater for the small premium firms than for the large premium firms, the relative savings tends to be greater for the latter group.**

compare the predictions of empirical Bayes methods with those of a second-order autoregressive model.[5] It would also be interesting to compare empirical Bayes to the time series and econometric methods studied by Cummins and Griepentrog (1985) in forecasting the paid claim cost data used in insurance pricing.

Although not a precise correction for the autoregressive nature of the loss ratio, employing the empirical Bayes process on losses incurred allows a partial check of the effect of the autoregressive problem. This is because the strength of empirical Bayes can be evaluated in estimating a variable related to the loss ratio. If the empirical Bayes method offers more improvement in the estimation process of losses incurred, then the implication is that, given the data, a correction for autoregression may improve the strength of the empirical Bayes estimate of the loss ratio.

There are additional considerations in using the empirical Bayes technique on losses incurred. First, an adjustment must be made for claim cost inflation.[6] Since the lines may be affected differentially by inflation, components of the Consumer Price Index (CPI) for each line were used rather than the generalized index. The data were obtained from the Citibase data retrieval system. The proxy for claim cost inflation for auto physical damage was the CPI index for automobile body work maintenance and repair; for fire, the CPI index for housing; and for auto liability and medical malpractice, the CPI index for professional medical services.

The second consideration is the potential distortion caused by differences in size of losses across the firms in a line. One of the strengths of the empirical Bayes methodology is that known differences in the μ_is can be taken into account by replacing μ, the mean of the prior, by a regression function $\mathbf{x}_i\beta$, where $\mathbf{x}_i$ is a vector of known explanatory variables, and β is a vector of unknown regression coefficients to be estimated from the data. In this case, the empirical Bayes estimator is

$$\hat{\mu}_{EBi} = (1 - \hat{B}_i)Y_i + \hat{B}_i(\mathbf{x}_i\beta),$$

where $\beta = (\mathbf{X'D\,X})^{-1}(\mathbf{X'D\,Y})$, $\mathbf{X}$ is the design matrix whose rows are the $\mathbf{x}_i$s, $\mathbf{D}$ is a diagonal matrix having $1/(V_i+\hat{A})$ as its ith element, and $\mathbf{Y}' = (Y_1,\ldots, Y_k)$. (Other slight changes are made in the estimation procedure as well; see Morris [1983] for details.) Thus, the empirical Bayes estimators can be constructed to shrink Y_i toward its own predicted (modeled) mean, rather than toward a constant mean. To control for differences in the size of the losses, the prior mean was modeled as a function of premiums earned.

---

[5] The limited number of time series observations in our sample (seven) precluded meaningful tests of the autoregressive model.

[6] Research indicates that specialized indices, such as the Consumer Price Index for body work, perform poorly as predictors of paid claim cost inflation (see Cummins and Nye, 1980). Broad based indices such as the Consumer Price Index on wage rates typically perform better. In the absense of prior information on predictors of incurred losses, we elected to use specialized price indices.

Estimates of losses incurred were then calculated using the estimator shown above, with

$$x_i\beta = \beta o + \beta_1 * [\text{premiums earned}_i].$$

The relative savings shown by the empirical Bayes estimate of losses incurred is positive for three of the lines analyzed--auto liability, auto physical damage, and medical malpractice (see Table 7). For fire, the relative savings was slightly negative. In terms of the percentage of firms, the empirical Bayes estimate is closer to the observed 1987 losses incurred for 54 to 80 percent of the firms, depending upon the line of business. These results favor the empirical Bayes process over the maximum likelihood estimate. When compared to the results derived from estimating loss ratios, for at least two of the lines--auto liability and auto physical damage—the empirical Bayes model performs significantly better when autoregressive tendencies are not present in the variable to be estimated. This supports the assertion that, given the data, a correction for the autoregressive tendencies found in the loss ratio may improve the strength of the empirical Bayes process.

Table 7

Mean Squared Forecast Errors: Losses Incurred

	Mean Squared Error of			Percent of Companies for Which
	$\bar{y}_i$	$\hat{\mu}_{EBi}$	$RS(\hat{\mu}_{EBi}, \bar{y}_i)^a$	$\lvert\bar{y}_i - y_{i,87}\rvert > \lvert\hat{\mu}_{EBi} - y_{i,87}\rvert$
Auto Liability	3.92 E+14	2.63 E+14	32.8	74.8
Auto Physical Damage	6.78 E+13	3.97 E+13	41.5	79.6
Medical Malpractice	1.13 E+14	0.95 E+14	16.1	54.5
Fire	4.02 E+12	4.15 E+12	–3.3	64.1

Note: $\bar{y}_i$ is the time series mean of the annual loss ratios calculated for 1980 through 1986. $\hat{\mu}_{EBi}$ is the empirical Bayes estimate of the annual loss ratio calculated for 1980 through 1986, assuming that the loss variance differs across firms. $RS(\hat{\mu}_{EBi}, \bar{y}_i)$ is the relative savings of the empirical Bayes estimate over the time series mean calculated as $[MSE(\bar{y}_i) - [MSE(\hat{\mu}_{EBi})]/MSE(\bar{y}_i)$, where MSE is the mean squared error.

[a]Numbers in this column cannot be replicated from numbers in the table due to rounding.

Conclusions and Recommendations

The empirical Bayes model is used to estimate the loss ratio for four lines of business. Using loss ratio data for 1980 through 1986, three candidate predictors of the 1987 loss ratio are derived: the time series mean and two empirical Bayes procedures, which differ in their assumptions about the individual firm's variance. The performance of the three predictive procedures are compared using the mean squared prediction error. An alternative evaluation criterion is based upon the number of companies within each line for which the empirical Bayes predictor is closer than the time series mean to the reported 1987 loss ratio. The two empirical Bayes procedures result in more precise estimates than the time series mean for all lines of business, although the gain from empirical Bayes is small for auto liability, auto physical damage, and fire. For medical malpractice, the gain is substantial, and the empirical Bayes procedure that allows for unequal variances across firms performs considerably better than the estimator assuming equal variance.

To evaluate the sensitivity of the results, the number and percentage of firms within each line of business for which the empirical Bayes predictor (unequal variance assumption) is closer to the 1987 loss ratio were calculated. Again, for all lines of business the empirical Bayes estimator performs well for the majority of firms when compared with the time series mean estimate. A second sensitivity analysis considers systematic differences in the by-line loss ratio according to the amount of business in the line. These results are particularly interesting since there is evidence that firms with relatively low premium dollars in a line may benefit the most from relying upon the empirical Bayes estimator. This result is intuitive and could have practical applicability. Finally, because of the autoregressive nature of the loss ratio, the strength of the empirical Bayes approach is evaluated in estimating an alternative variable, losses incurred. These results imply that a longer time series of data might correct for the autoregressive tendencies of the loss ratio and improve the strength of the process.

Regulators, security analysts, and auditors should consider the empirical Bayes model for forecasting key financial variables that impact the profitability of the property-liability insurer. Future research might address the percent or average improvement provided by the empirical Bayes estimator of the forecasted loss ratio, which could be decomposed further and attributed to either the nature of the business or the number of participants. Given additional data, one might consider a correction for autoregression, as well as an evaluation of other key financial variables.

Appendix

The mean and variance of the prior distribution, μ and A, can be estimated by iteratively solving the equations

$$\hat{A} = \sum_{i=1}^{k} W_i\left[\frac{k}{k-1}(\bar{y}_i - \mu)^2 - V_i\right] / \left(\sum_{i=1}^{k} W_i\right), \tag{A.1}$$

where $W_i = 1/(V_i + \hat{A})$, and

$$\hat{\mu} = \left(\sum_{i=1}^{k} W_i\bar{y}_i\right)/\left(\sum_{i=1}^{k} W_i\right). \tag{A.2}$$

$\hat{A}$ should be set to zero if its estimate becomes negative. This method is attributable to Morris (1983).

References

A. M. Best Company. A. M. Best Property-Liability Data Tapes (Oldwick, N.J.: A. M. Best).

Berger, Lawrence A., 1988, A Model of the Underwriting Cycle in the Property/Liability Insurance Industry, *Journal of Risk and Insurance*, 50: 298-306.

Brockett, Patrick L., 1983, On the Misuse of the Central Limit Theorem in Some Risk Calculations, *Journal of Risk and Insurance*, 50: 727-731.

Brockett, Patrick L. and Robert C. Witt, 1982, The Underwriting Risk and Return Paradox Revisited, *Journal of Risk and Insurance*, 49: 621-627.

Cummins, J. David, 1991, Statistical and Financial Models of Insurance Pricing and the Insurance Firm, *Journal of Risk and Insurance*, 58: 261-302.

Cummins, J. David and Gary L. Griepentrog, 1985, Forecasting Automobile Insurance Paid Claims Costs Using Econometric and ARIMA Models, *International Journal of Forecasting*, 1: 203-215.

Cummins, J. David and David J. Nye, 1984, Inflation and Property-Liability Insurance: Causes, Consequences, and Solutions, in: John D. Long, ed., *Insurance Issues* (Malvern, Penn.: American Institute).

Cummins, J. David and J. François Outreville, 1987, An International Analysis of Underwriting Cycles in Property-Liability Insurance, *Journal of Risk and Insurance*, 54: 246-262.

Doherty, Neil and James R. Garven, 1991, Interest Rates, Financial Structure, and Insurance Price Cycles, in: J. D. Cummins, S. E. Harrington, and R. Klein, eds., *The Underwriting Cycle in Property-Liability Insurance* (Kansas City, Mo.: National Association of Insurance Commissioners).

Doherty, Neil and Han Bin Kang, 1988, Interest Rates and Insurance Price Cycles, *Journal of Banking and Finance*, 12: 199-214.

Efron, Bradley and Carl Morris, 1973, Stein's Estimation Rule and Its Competitors—An Empirical Bayes Approach, *Journal of the American Statistical Association*, 68: 117-130.

Efron, Bradley and Carl Morris, 1975, Data Analysis Using Stein's Estimator and Its Generalizations, *Journal of the American Statistical Association*, 70: 311-319.

Efron, Bradley and Carl Morris, 1977, Stein's Paradox in Statistics, *Scientific American*, 236: 119-127.

Harrington, Scott E., 1988, Prices and Profits in the Liability Insurance Market, in: Robert E. Litan and Clifford Winston, eds., *Liability: Perspectives and Policy* (Washington, D.C.: Brookings Institution).

James, W. and Charles Stein, 1960, Estimation with Quadratic Loss, *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, Volume 1 (Berkeley: University of California Press), 361-379.

Jewell, William, 1975, Model Variations in Credibility Theory, in: Paul M. Kahn, ed., *Credibility Theory and Applications* (New York: Academic Press), 193-244.

Klugman, Stuart A., 1992, *Bayesian Statistics in Actuarial Science with Emphasis on Credibility* (Boston: Kluwer Academic Publishers).

Lamm-Tennant, Joan and Laura T. Starks, 1991, Ownership Structure of the Property-Liability Insurance Industry: An Analysis of the Differences, Working Paper, Villanova University and University of Texas at Austin.

Miller, Robert B. and James C. Hickman, 1975, Insurance Credibility Theory and Bayesian Estimation, in: Paul M. Kahn, ed., *Credibility Theory and Applications* (New York: Academic Press), 249-270.

Morris, Carl, 1983, Parametric Empirical Bayes Inference: Theory and Applications, *Journal of the American Statistical Association*, 78: 47-55.

Morris, Carl and Lee Van Slyke, 1978, Empirical Bayes Methods for Pricing Insurance Claims, *Proceedings of the Business and Economics Statistics Section*, (Washington, D.C.: American Statistical Association), 579-582.

Neuhaus, Walther, 1984, Inference about Parameters in Empirical Linear Bayes Estimation Problems, *Scandinavian Actuarial Journal*, 67: 131-142.

Norberg, Ragnar, 1980, Empirical Bayes Credibility, *Scandinavian Actuarial Journal*, 63: 177-194.

Smith, Michael L. and Fikry S. Gahin, 1983, The Underwriting Cycles in Property and Liability Insurance (1950-1978), Paper presented to the 1983 Risk Theory Seminar, Helsinki, Finland.

Tukey, John W., 1963, Borrowing Strength in a Diversified Situation, Princeton University Working Paper, Statistical Techniques Research Group, Princeton, New Jersey.

Venezian, Emilio, 1985, Ratemaking Methods and Profit Cycles in Property and Liability Insurance, *Journal of Risk and Insurance*, 52: 477-500.

Witt, Robert C., 1977, The Automobile Insurance Rate Regulatory System in Illinois: A Comparative Study, Illinois Insurance Laws Study Commission.

Witt, Robert C., 1978, The Competitive Rate Regulatory System in Illinois: A Comparative Study by State, *CPCU Journal*, 31: 151-162.