Statistical Trend Estimation with Application to Workers Compensation Ratemaking

Frank Schmid

Abstract

Motivation. Estimating the trend of the severity, frequency, and loss ratio rates of growth is an integral part of NCCI ratemaking. The time series from which such trend estimation has to be derived are typically short and volatile, comprising only 19 observations or, equivalently, 18 rates of growth. Thus, separating signal (i.e., trend) from (white) noise is particularly challenging.

Method. NCCI has developed a Bayesian Statistical Trend model that is geared toward extracting the trend in short and high-volatility time series. This model has been optimized by minimizing the root mean squared prediction error across NCCI states using three-year hold-out periods (as the applicable forecasting horizon is typically around three years).

Results. We present trend estimates for severity, frequency, and loss ratio rates of growth for an unidentified state. The model is robust to outliers and delivers stable, yet time-varying trend estimates.

Conclusions. The statistical properties of the model are conducive to rate stability and, at the same time, allow the practicing actuary to recognize changes in trend.

Availability. The model runs in WinBUGS 1.4.3 (www.mrc-bsu.cam.ac.uk/bugs) within the R (www.r-project.org) package R2WinBUGS (http://cran.r-project.org). WinBUGS is administered by the MRC Biostatistics Unit, University of Cambridge, UK; R is administered by the Technical University of Vienna, Austria. WinBUGS and R are GNU projects of the Free Software Foundation and hence available free of charge.

Keywords. Trend and loss development; Bayesian methods; time series; Workers Compensation.

1. INTRODUCTION

Estimating the trend of the frequency, severity, and loss ratio rates of growth is an integral part of NCCI ratemaking. The time series on which such trend estimation rests are typically short, comprising only 19 observations or, equivalently, 18 rates of growth. Further, these time series may display high degrees of volatility. Thus, separating signal (i.e., trend) from (white) noise is critical for discerning the trend. To achieve this objective, NCCI has developed and implemented a Bayesian state-space model that is designed to elicit the trend in short and volatile time series. This model has been optimized by minimizing the root mean squared prediction error (RMSPE) across NCCI states using three-year hold-out periods (as the applicable forecasting horizon typically amounts to about three years).

The Statistical Trend model runs in WinBUGS 1.4.3 (www.mrc-bsu.cam.ac.uk/bugs) within the R (www.r-project.org) package R2WinBUGS (http://cran.r-project.org). WinBUGS is administered by the MRC Biostatistics Unit, University of Cambridge, UK; R is administered by the Technical University of Vienna, Austria. WinBUGS and R are GNU projects of the Free Software Foundation and hence available free of charge.

1.1 Research Context

Forecasting is a signal extraction and signal extrapolation problem. Measurement errors cause the quantities of interest (such as the rates of growth of indemnity severity, medical severity, and frequency) to be observed with (presumably Gaussian) noise, thus obscuring the signal. In forecasting, the signal is the quantity of interest, because it is the signal on which future observations center. Specifically, it is the objective of a forecasting model to educe from historical observations the process that generates the unobservable signal. Because a forecasting model replicates the data-generating process of the signal (as opposed to replicating the observations themselves), its quality cannot be judged by the (in-sample) fit to the observed data, as gauged, for instance, by the R^2 or similar measures of goodness of fit. In fact, good fit to heretofore observed data harbors the risk of overfitting. Such overfitting may imply that the forecasts do not center on the signal, thus giving rise to potentially large forecasting errors. The risk of overfitting affords parsimony a critical role in time series modeling.

As an example, consider a game of dice. In any roll of a pair of dice, the expected value of the outcome is 7. This expected value is the signal, which manifests itself as the mean outcome as the number of rolls goes to infinity. The signal offers an unbiased forecast for any future toss; the difference between the observations and the signal is noise. Thus, among all possible forecasting models, the one that simply produces the time-invariant signal of 7 as its forecast has the lowest expected root mean squared prediction error. Yet, this forecasting algorithm offers the worst in-sample fit possible, as the model has no explanatory power with regards to the variation of the outcome around the expected value. Not surprisingly, a least-squares regression of the 36 possible outcomes on the time-invariant signal reveals an R^2 equal to zero.

Two common properties in time series are nonstationarity and mean reversion. In the example above, nonstationary is equivalent to a time-varying mean; instead of invariably equaling 7, this mean would change in time. As will be argued below, in workers compensation, the frequency rate of growth (and, as a result, the loss ratio rate of growth) should be treated as nonstationary.

Mean reversion, on the other hand, implies that the outcomes of the mentioned rolls of dice are not independent draws, thus causing serial correlation. In games of chance, such mean reversion is associated with the gambler's fallacy, which rests on the (erroneous) belief that below-average outcomes of past rolls of dice are to be corrected by above-average outcomes in the future (instead of simply being diluted). Although the business cycle may introduce mean reversion in the severities and frequency rates of growth, such mean reversion is likely to be weak and, more importantly, cannot be expected to improve the quality of the forecasts in short non-stationary time series due to lack of precision in estimating such mean reversion.

Traditionally, NCCI estimates trends using the exponential trend approach, which is a linear regression of logarithmic levels on a sequence of integers that indicate time.

1.2 Objective

Recent advances in statistical modeling offer ways of dealing with the problem of estimating trend rates of growth from times series that are short, volatile, and potentially nonstationary. The state-space modeling framework, along with the Metropolis-Hastings algorithm for estimating Bayesian models by means of Markov-Chain Monte Carlo (MCMC) simulation, makes such sophisticated statistical modeling available to the practicing actuary.

1.3 Outline

What follows is the presentation of a multiequation model for forecasting the trend in the rates of growth of the indemnity and medical severities, frequency, and the respective loss ratios. This model is then applied to a "paid" data set of an unidentified U.S. state for the time period 1987–2005. The last section offers conclusions and guidance for implementation of this model in actuarial practice.

2. BACKGROUND AND METHODS

In the context of NCCI ratemaking, frequency is defined as the ratio of the developed (to the 5th report) number of claims to the developed (to the 5th report), on-leveled (to the current loss-cost or rate level, depending on the case), and wage-adjusted premium. Severity is defined as the ratio of the developed, on-leveled, and wage-adjusted loss to the developed (to the 5th report) number of claims. When defined in such way, the product of indemnity (medical) severity and frequency equals the indemnity (medical) loss ratio. In consequence, the logarithmic rate of growth of the loss ratio equals the sum of the logarithmic rates of growth of frequency and the applicable severity; in what follows, this property is referred to as the add-up constraint.

The model estimates the five rates of growth (the two severities, frequency, and the two loss ratios) simultaneously. Covariances among these variables account for common shocks. For instance, the severities and frequency are subject to common shocks because they share the wage adjustment; further, the severities and frequency share components of the on-leveling for changes in benefits levels. The joint estimation of the growth rates also allows for implementing the mentioned add-up constraint. This constraint ensures that, at any point in time, the estimated rates of growth of the indemnity (medical) severity and frequency are consistent with the estimated rate of growth of the indemnity (medical) loss ratio.

The model uses logarithmic rates of growth, because conventional rates of growth have a lower bound at minus 1 and, hence, violate the assumption of normality. These logarithmic rates of growth are then transformed into conventional rates of growth to obtain the forecast rates of growth and, after adding 1, the NCCI trend factors. Further, the (conventional) rates of growth are compounded over the multiyear forecasting horizon or, equivalently, the NCCI trend period; the number of years of this trend period is typically not an integer. Adding 1 to the compounded rates of growth delivers the NCCI adjustment factors. The purpose of the adjustment factor is to scale up the levels of the variables of interest (the severities, frequency, and the loss ratios) from the end of the experience period (i.e., the time period for which there are observations available) to the end of the trend period (i.e., the end of the forecasting horizon).

Note that transforming logarithmic rates of growth into conventional rates of growth necessitates a bias-adjustment when such transformation is done on the expected value; this is because, for a normally distributed variable x, $\mathrm{E}[e^{x}] = e^{\mathrm{E}[x]+\sigma^2/2}$. Because the model is estimated by means of Monte Carlo simulation, such transformation happens "draw by draw" (instead of on the expected value) and, thus, no bias-adjustment is necessary.

For the time period 1988-2005, Chart 1 shows for an unidentified state the (conventional) rates of growth of the indemnity and medical severities. Chart 2 displays the growth rate for frequency. Finally, Chart 3 presents the growth rates of the corresponding loss ratios.

Although Charts 1 through 3 are not necessarily representative of NCCI states, they are typical in that they are indicative of nonstationarity (i.e., time-variation in the mean) in the growth rate of frequency, but less so in the severities. Note that because the sum of two time series is nonstationary if at least one of the individual series is nonstationary, a time-varying mean in the growth rate of frequency implies time-varying means in the growth rates of both loss ratios. For instance, as Chart

2 shows, the growth rate of frequency was higher at around the year 1990 than it was ten years later; and because the variation in the mean growth rate of frequency was not offset by a change (in opposite direction) of the growth rates of the severities, such variation is mirrored in the means of the growth rates of the loss ratios (see Chart 3).

Time series can be checked for nonstationarity, but such unit root tests have little power for short time series; as a consequence, these tests favor the null hypothesis of a (pure) random walk (see, for instance, Hamilton [4]). As will be argued below, stationarity and nonstationarity are limiting cases of a smoothed random walk. Frequently, neither stationarity (a time-invariant mean) nor a (pure) random walk is an appropriate assumption for forecasting models that rely on short and volatile time series.

Chart 1: Growth Rates of Indemnity and Medical Severities, Policy Years 1988–2005

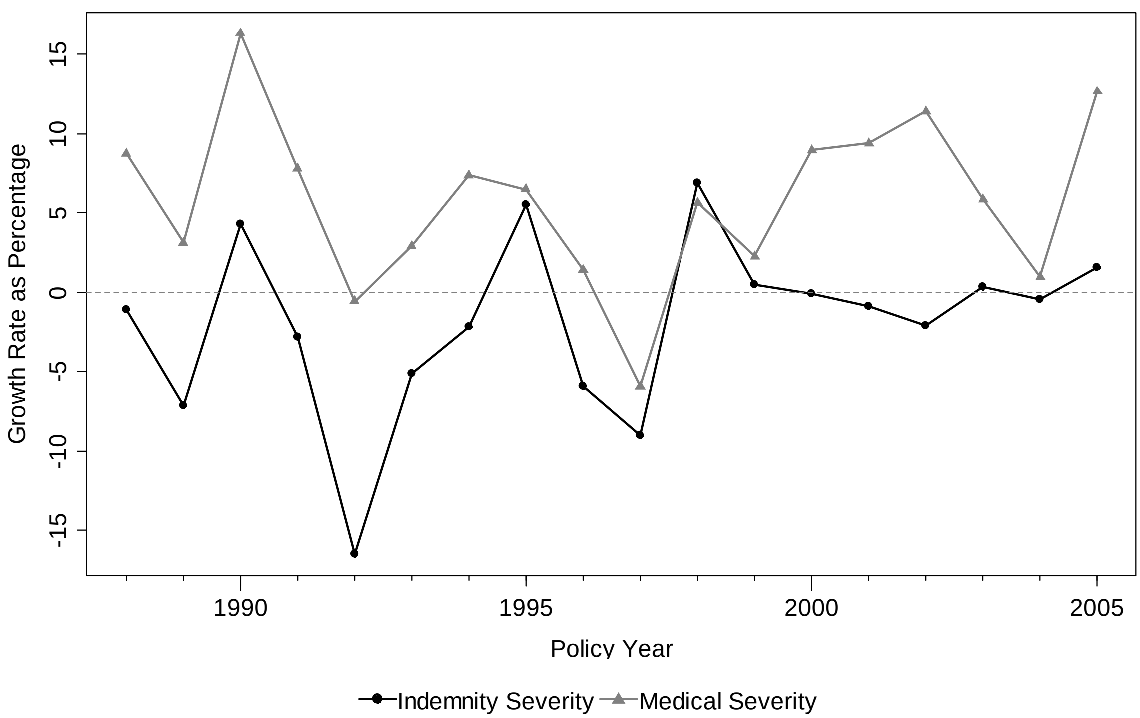

Another property frequently present in time series is serial correlation. Such serial correlation may originate in mean reversion, as caused by the business cycle. First, the rate of frequency growth may be hypothesized to vary with economic activity as the least productive workers are the last to be hired in an economic upturn and the first to be laid off in a downturn. Second, wage growth is a (lagging) indicator of economic activity; hence, the wage-adjusting of losses (severities) and premium (frequency) may introduce mean reversion into the severities and frequency series. On the other hand, the business cycle in the United States has been fairly shallow over the past 20 years; there were only two official recessions (1990/91 and 2001) according to the NBER Business Cycle Dating Committee (http://www.nber.org/cycles/cyclesmain.html) and, according to the Federal Reserve Bank of Saint Louis Fred2 database (https://research.stlouisfed.org/fred2), only one-quarter of negative GDP growth. In conclusion, discerning a shallow mean-reverting process in time series as short and volatile as those depicted in Charts 1 through 3 harbors the risk of overfitting and is likely to add little predictive power to the forecasts.

Chart 2: Growth Rate of Frequency, Policy Years 1988–2005

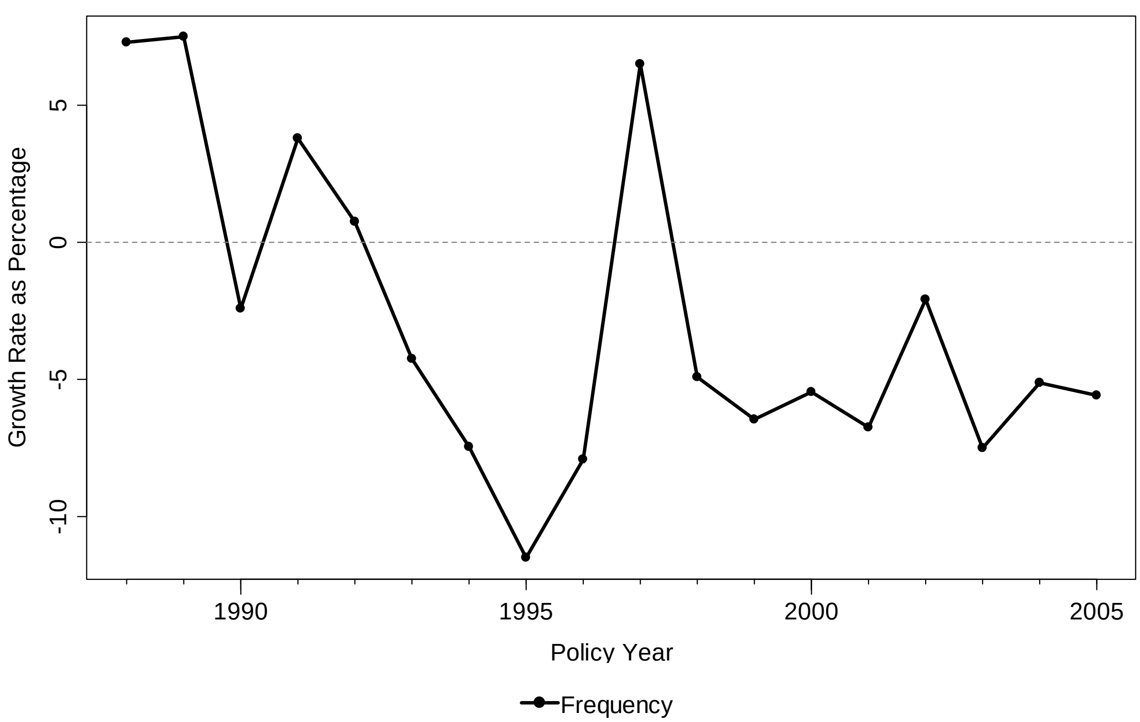

As mentioned, the forecasting model makes use of the concept of the smoothed random walk. For illustration, a simple model of a smoothed random walk may be written as follows:

$$y_t \sim \mathrm{N}(x_t, \sigma_y^2)\ ,\ \ t = 1, \ldots, T \tag{1.1}$$

$$x_t \sim \mathrm{N}(x_{t-1}, \sigma_x^2)\ ,\ \ t = 2, \ldots, T \tag{1.2}$$

where $\mathrm{N}(\mu, \sigma^2)$ indicates a normal distribution with mean μ and finite variance σ^2. Equation (1.1) states that the variable y_t is observed with measurement noise σ_y^2 around the unobservable mean x_t; in state-space format, this equation is called the measurement equation. Equation (1.2) states that the mean x_t is time-varying as described by a random walk with innovation variance σ_x^2; in state-space format, this equation is called the transition equation.

Chart 3: Growth Rates of Indemnity and Medical Loss Ratios, Policy Years 1988–2005

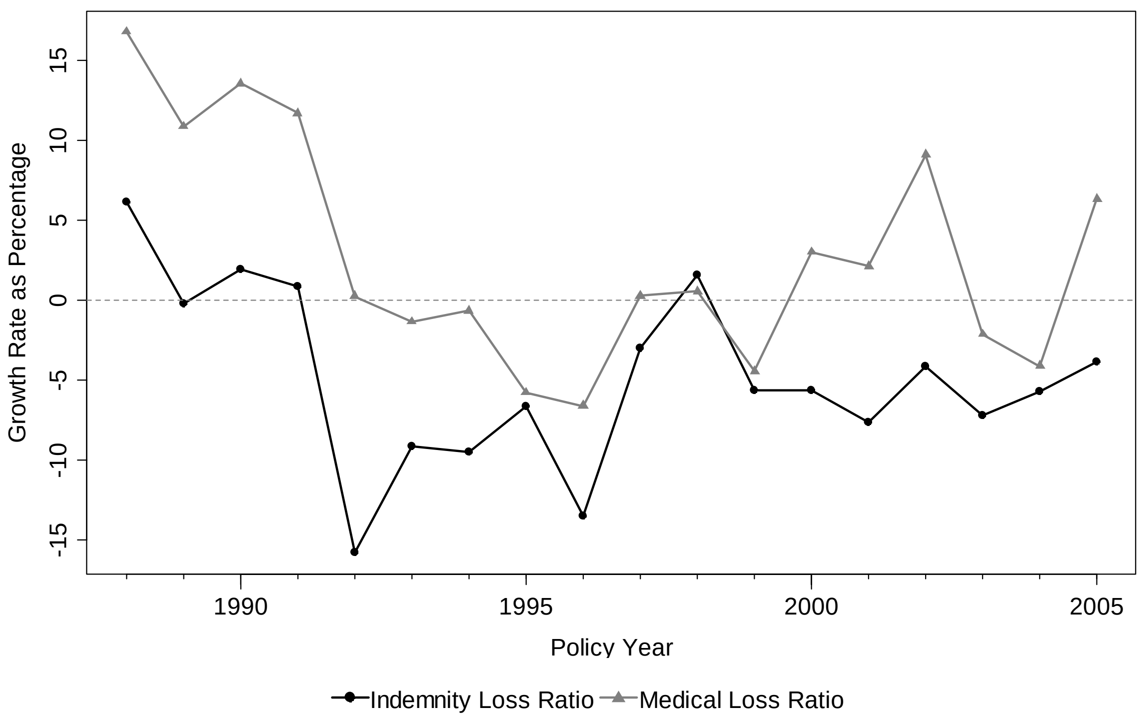

There are two limiting cases to model (1.1–1.2), one of which is the case of stationarity, and the other one is the (pure) random walk. We obtain stationarity by setting the innovation variance σ_x^2 in the transition equation to zero:

$$y_t \sim \mathrm{N}(x_t, \sigma_y^2)\ ,\ \ t = 1,...,T \tag{2.1}$$

$$x_t = x_{t-1}\ ,\ \ t = 2,...,T \tag{2.2}$$

Alternatively, we obtain the limiting case of a pure random walk by setting the measurement noise σ_y^2 to zero:

$$y_t = x_t\ ,\ \ t = 1,...,T \tag{3.1}$$

$$x_t \sim \mathrm{N}(x_{t-1}, \sigma_x^2)\ ,\ \ t = 2,...,T \tag{3.2}$$

In the general case where neither time-variation in the mean (nonstationarity) nor measurement noise can be excluded, model (1.1-1.2) applies. In such a general model, it is necessary to determine how much of the time-variation in the dependent variable y_t should be assigned to noise (σ_y^2); the remainder represents innovation (σ_x^2). This allocation decision, which determines the degree of smoothing of the dependent variable, may be assigned to an algorithm such as the Kalman filter (as discussed in Evans and Schmid [2]; for a general discussion of the Kalman filter, see, for instance, Hamilton [4]). Note that for any given set of observations y_t, $t = 1,...,T$, there is only one degree of freedom in determining the optimal degree of smoothing, as choosing σ_y^2 determines σ_x^2, and vice versa.

Unfortunately, the Kalman filter does not necessarily deliver the optimal degree of smoothing; in short and volatile time series in particular, the Kalman filter assigns more time variation to innovations in the mean than is conducive to minimizing the forecasting error.

2.1 The Model

The model to be introduced in this section is Bayesian. Such Bayesian models have a number of advantages over frequentist approaches, among which is the ease at which even very complex models can be estimated. For instance, if there were missing values in the severity, frequency, or loss ratio series, the model shown below would interpolate of its own accord, based on the estimated random walk properties.

The model is estimated using the Metropolis–Hastings algorithm, which computes the (posterior) distributions of the model parameters by means of Markov–Chain Monte Carlo simulation. For an introduction to Bayesian modeling see Gelman et al. [3].

Equation (4) below represents a system of transition equations for the rates of severity and frequency growth, which describe (smoothed) random walks; the innovations to these variables (i.e., the changes to their means) follow a multivariate normal distribution. Equation (5) states that the initial values for the three mentioned growth rates are also draws from a multivariate normal; the expected values of this normal are zero, but the covariance matrix imposes little restrictions on the means of their posterior distributions. Equation (6) describes the measurement equations, inclusive of the add-up constraint; in the measurement equations, the model fits to the observed values the process that is stated in Equation (4). Equations (7) through (10) describe the prior distributions for the covariance matrices of the initial states, the innovations, and the measurement noise; these covariance matrices are modeled on Wishart distributions.

$$\begin{pmatrix} z_{s_{ind},t} \\ z_{s_{med},t} \\ z_{fr,t} \end{pmatrix} \sim \mathrm{N}\left(\begin{pmatrix} z_{s_{ind},t-1} \\ z_{s_{med},t-1} \\ z_{fr,t-1} \end{pmatrix}, \mathbf{\Omega}_1 \right), \; t = 2,\dots,T \tag{4}$$

$$\begin{pmatrix} z_{s_{ind},1} \\ z_{s_{med},1} \\ z_{fr,1} \end{pmatrix} \sim \mathrm{N}\left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \mathbf{\Omega}_2 \right) \tag{5}$$

$$\begin{pmatrix} s_{ind,t} \\ s_{med,t} \\ fr_t \\ lr_{ind,t} \\ lr_{med,t} \end{pmatrix} \sim \mathrm{N}\left(\begin{pmatrix} z_{s_{ind},t} \\ z_{s_{med},t} \\ z_{fr,t} \\ z_{s_{ind},t} + z_{fr,t} \\ z_{s_{med},t} + z_{fr,t} \end{pmatrix}, \mathbf{\Omega}_3 \right), \; t = 1,\dots,T \tag{6}$$

$$\mathbf{\Omega}_i \sim \mathrm{W}(\mathbf{R}_i, 1000), \; i = 1,3 \tag{7}$$

$$\boldsymbol{\Omega}_i \sim \mathrm{W}(\mathbf{R}_i, 10)\ ,\ i = 2 \tag{8}$$

$$\mathbf{R}_1 = \mathbf{R}_2 = 0.01 \times I_{3\times 3} \tag{9}$$

$$\mathbf{R}_3 = 0.2 \times I_{5\times 5} \tag{10}$$

where N and W indicate normal and Wishart distributions, respectively. The variables $s_{ind,t}$ and $s_{med,t}$ are the logarithmic rates of growth of indemnity and medical severities, respectively. The variable fr_t is the logarithmic rate of growth of frequency, and the variables $lr_{ind,t}$ and $lr_{med,t}$ are the logarithmic rates of growth of the indemnity and medical loss ratios, respectively. $I_{3\times 3}$ and $I_{5\times 5}$ symbolize identity matrices. The larger the diagonal elements of $\mathbf{R}_3$ are, the greater the degree of smoothing is. The matrix $\mathbf{R}_i (i = 1,\ldots,3)$ is a scale matrix that "represents an assessment of the order of magnitude of the covariance matrix" $\boldsymbol{\Omega}_i^{-1} (i = 1,\ldots,3)$ (WinBUGS [5]). (Note that the WinBUGS notation for the normal distribution makes use of the precision matrix, which is the inverse of the covariance matrix.)

If (and only if) there is a predictable upswing in future economic activity, the model employs a covariate (explanatory variable). In this case then, the rate of frequency growth is modeled as the sum of a (smoothed) random walk and a standard regression component; this standard regression component hosts the covariate. The covariate of choice is the change in the rate of unemployment. As argued, in an economic upswing, the growth rate of frequency can be expected to rise as currently employed labor is utilized more intensively and currently unemployed labor is put back to work. Predictable upswings in economic activity typically happen in the wake of natural disasters; an example of such an event is Hurricane Katrina in 2005.

When including a covariate, equations

$$\begin{pmatrix} z_{s_{ind},t} \\ z_{s_{med},t} \\ \bar{z}_{fr,t} \end{pmatrix} \sim \mathrm{N}\left(\begin{pmatrix} z_{s_{ind},t-1} \\ z_{s_{med},t-1} \\ \bar{z}_{fr,t-1} \end{pmatrix},\ \boldsymbol{\Omega}_1 \right),\ t = 2,\ldots,T \tag{11}$$

$$\begin{pmatrix} z_{s_{ind},1} \\ z_{s_{med},1} \\ \bar{z}_{fr,1} \end{pmatrix} \sim \mathrm{N}\left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \mathbf{\Omega}_2 \right) \tag{12}$$

substitute for Equations (4) and (5), and the following two equations are added to the model:

$$\lambda \sim \mathrm{N}(0, \tau_\lambda) \tag{13}$$

$$z_{fr,t} = \bar{z}_{fr,t} - \lambda \cdot \delta_t, \; t = 1,\ldots,T \tag{14}$$

where the variable δ_t is the t to $t-1$ (accident year; policy year: $t+1$ to $t-1$) difference in the rate of unemployment. The prior for the parameter λ is a normal distribution with an expected value of zero and a variance (τ_λ^{-1}) that must be chosen sufficiently small to prevent λ from picking up noise.

The model generates forecasts by moving the (logarithmic) rates of growth of frequency and the severities forward according to the innovation variances of the random walks described in Equation (4) (or Equation (11), if applicable), based on the estimated innovation covariance matrix $\hat{\Omega}_1^{-1}$. As is the case with all estimated parameters of the models, the forecasts come as posterior distributions, the means of which serve as the point estimates. The posterior distributions of the forecast trend and adjustment factors offer credible intervals, based on the chosen probability level (e.g., 95 percent). These credible intervals differ in important ways from traditional, frequentist confidence intervals. Whereas in frequentist statistics the true value either lies within a given confidence interval or not, the (Bayesian) credible interval is indeed a probabilistic statement about its location; see Carlin and Louis [1]. Note that the credible intervals are statements about the *trend* rates of growth, rather than the realized rates of growth (which are the sum of trend and noise).

3. RESULTS AND DISCUSSION

In what follows we apply the model (without a covariate, that is, Equations 4 through 10) to an unidentified U.S. state. The observations for the severities, frequency, and the loss ratios run from policy years 1987 through 2005, which renders 18 rates of growth (1988–2005).

As mentioned, the objective of the model is to generate trend factors, which are estimates of the trend rates of growth. By means of scaling up these trend factors to the trend period (i.e., the forecasting horizon), we obtain the adjustment factors.

NCCI typically computes adjustment factors not just for the final year, but also for the penultimate and antepenultimate years of the experience period. For instance, if the experience period ends with policy year 2005, then these alternative adjustment factors attach to the policy years 2004 and 2003, respectively; the corresponding alternative trend periods end on the same point on the calendar year axis as does the trend period that attaches to policy year 2005. (Note that the model is estimated only once; in the example above, this means that the trend factors that attach to policy years 2004 and 2003 are based on the same estimation as the trend factor that attaches to policy year 2005, thus utilizing all observations of the experience period.)

For a given policy year, the trend period runs from the midpoint of the policy year to the midpoint of what is known at NCCI as the effective period. The effective period is defined as the period in which the filed rate or loss cost (depending on the case) is going to be in effect. The midpoint of a policy year or an effective period is based on the assumed monthly premium distribution; such premium distribution may vary across states. As mentioned, the trend period attaches to the final year of the experience period, with alternative trend periods attaching to the penultimate and antepenultimate policy years. For the case at hand, this final year of the experience period is policy year 2005, and the trend period equals 3.001 years, rounded to the third decimal. Correspondingly, the trend period that attaches to the penultimate (antepenultimate) policy year of the experience period is 4.001 (5.001) years of length.

When the change in the rate of unemployment is used as a covariate for frequency growth, then this variable is measured by the two-year difference of the rate of unemployment for policy years and by the first (i.e., one-year) difference of the rate of unemployment for accident years. For instance, for policy year 2005, the pertinent two-year difference is the change in level between the end-of-calendar-year 2006 and the end-of-calendar-year 2004 values. For accident year 2005, the first difference is the change in level between the end-of-calendar-year 2005 and the end-of-calendar-year 2004 values. These end-of-calendar-year numbers of the unemployment rate refer to the final quarter of the year and are averaged over the three months in the quarter. (We average the rate of unemployment because only quarterly forecasts for the rate of unemployment are available.) In determining the trend estimates for the unidentified state discussed below, no covariate was employed.

As mentioned, the model presented above must be calibrated such that it minimizes the prediction error. This calibration is done by choosing the optimal degree of smoothing, as it manifests itself in the diagonal elements of the scale matrix $\mathbf{R}_3$; the prediction error is gauged by the RMSPE. To determine the optimal degree of smoothing, we ran the model with a holdout period of 3 years for all NCCI states with an array of smoothing parameters; the diagonal elements of $\mathbf{R}_3$ (which determine the degree of smoothing) were varied equidistantly on a logarithmic scale. The 3-year holdout period corresponds (approximately) to the applicable trend periods (of typically little more, but sometimes little less than 3 years). As shown in Chart 4, the RMSPE, aggregated across all NCCI states varies systematically with the degree of smoothing (which is represented by an index, not the actual magnitude of the diagonal elements of $\mathbf{R}_3$); the prediction error is large for little smoothing (low index values), because little smoothing entails a great deal of fitting to noise; also, the prediction error is large for extensive smoothing (high index values), because a high degree of smoothing insufficiently accommodates the nonstationarity of the underlying growth series.

Based on data from an unidentified state, the model is estimated using WinBUGS 1.4.3 within the R package R2WinBUGS. We sample 50,000 times, following a burn-in of 50,000 iterations.

The results for the severities, the frequency and the loss ratios are displayed in Charts 5 through 7. The dashed vertical lines in these charts indicate the beginning of the trend period, which attaches to the final year of the experience period (policy year 2005).

Chart 5 displays the actual, fitted, and forecast trend growth rates for indemnity and medical severities. The mean rates of severity growth show little time variation, although the indemnity growth rate is slightly trending up. The chart demonstrates that, for both series, the forecasts are not sensitive to the respective final observed value, which is desirable as any observed value contains potentially a great deal of noise.

Chart 6 depicts the actual, fitted, and forecast trend growth rates for frequency. Here, there is clearly evident a downtrend in the mean rate of growth. Also, note the insensitivity of the model to the outlier in the year 1997.

Chart 7 exhibits the actual, fitted, and forecast trend growth rates for the indemnity and medical loss ratios. The medical loss ratio trend rate of growth clearly follows the trend in frequency growth, while the indemnity loss ratio trend rate of growth is also influenced by the uptrend in the trend growth rate of indemnity severity.

Charts 8 and 9 present the posterior distributions for the estimated trend factors and adjustment factors for frequency, the indemnity and medical severities, and the indemnity and medical loss ratios.

Chart 4: Root Mean Squared Prediction Error (RMSPE) as a Function of the Degree of Smoothing

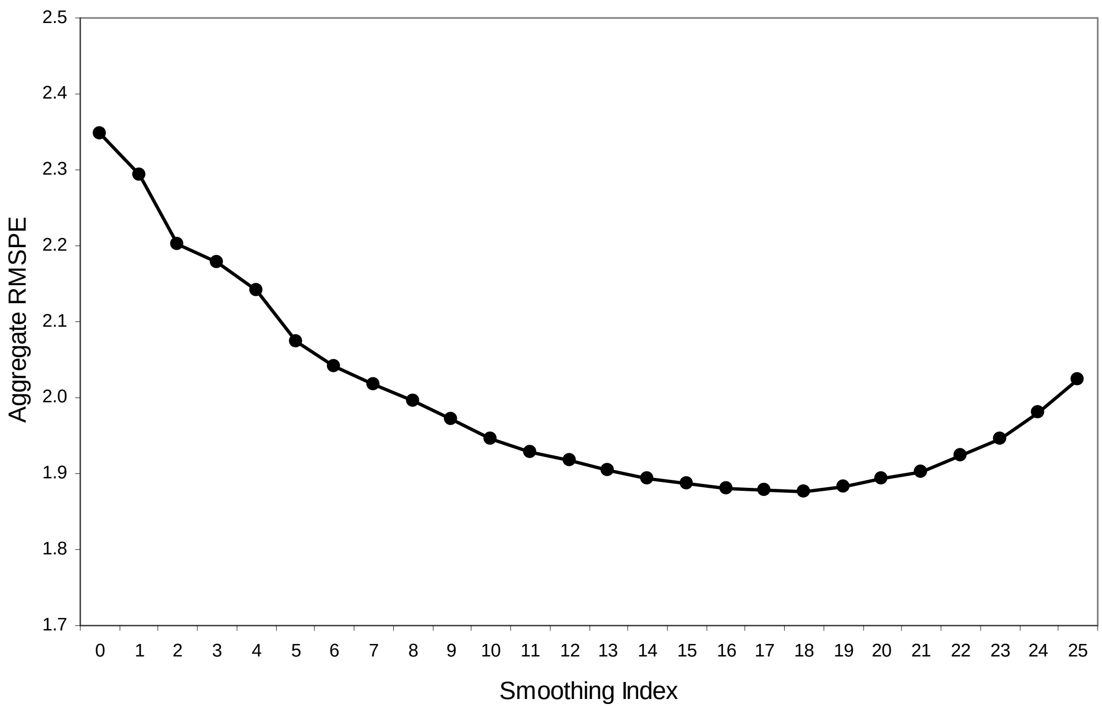

Chart 5: Growth Rates of Indemnity and Medical Severities (Actual, Fitted, and Forecast), Policy Years 1988–2005 (Forecasts: 2006–2009)

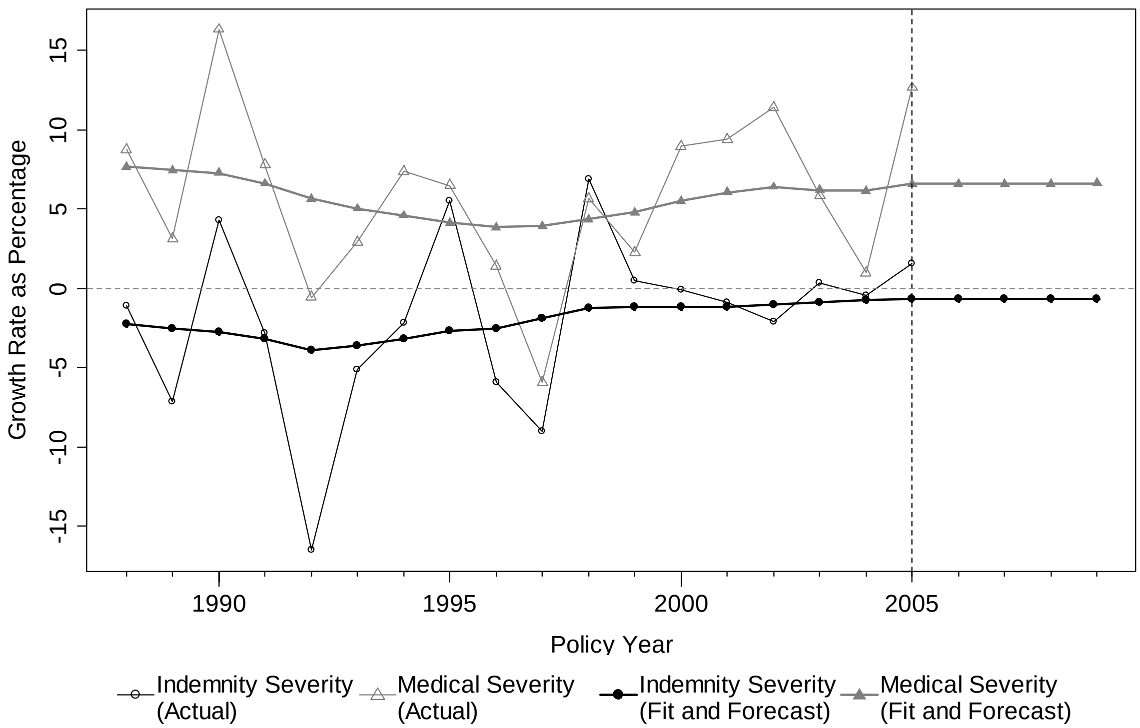

Chart 6: Growth Rate of Frequency (Actual, Fitted, and Forecast), Policy Years 1988–2005 (Forecasts: 2006–2009)

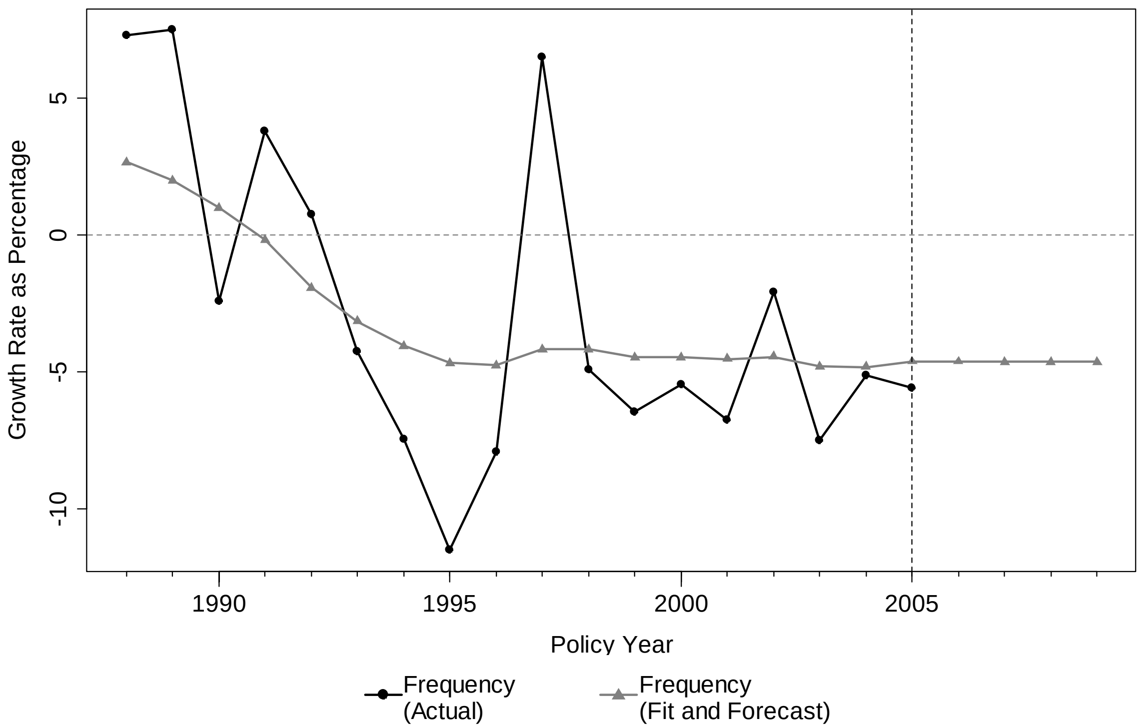

Chart 7: Growth Rates of Indemnity and Medical Loss Ratios (Actual, Fitted, and Forecast) Policy Years 1988–2005 (Forecasts: 2006–2009)

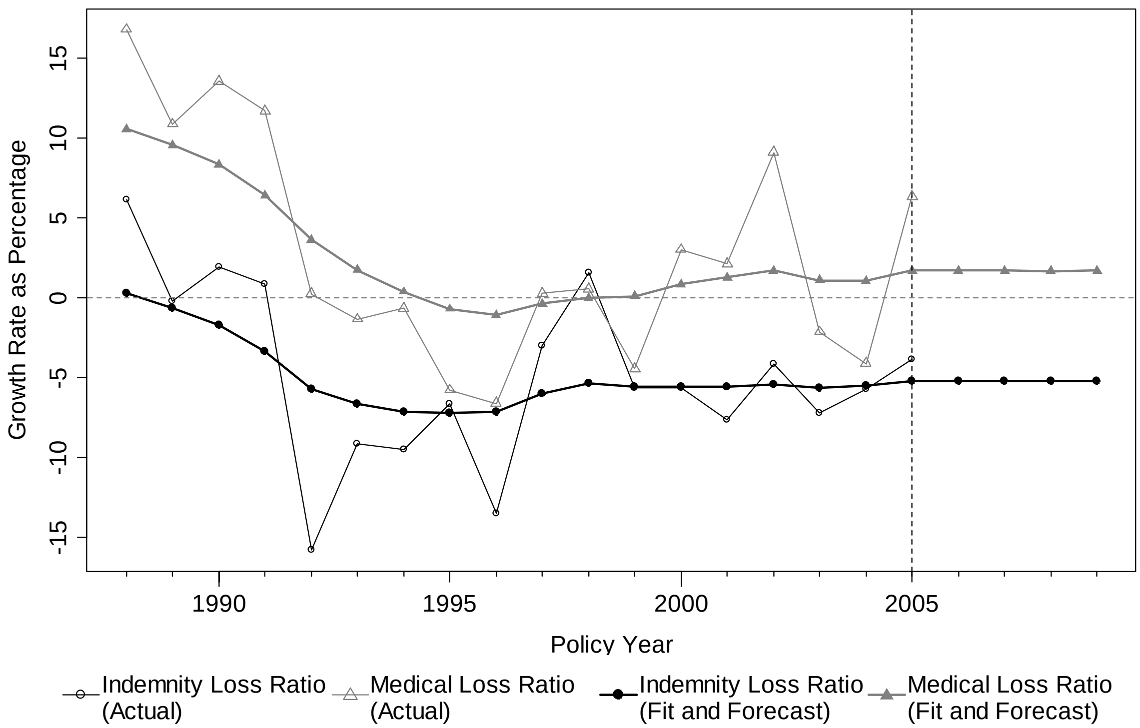

Chart 8: Posterior Densities for Trend Factors

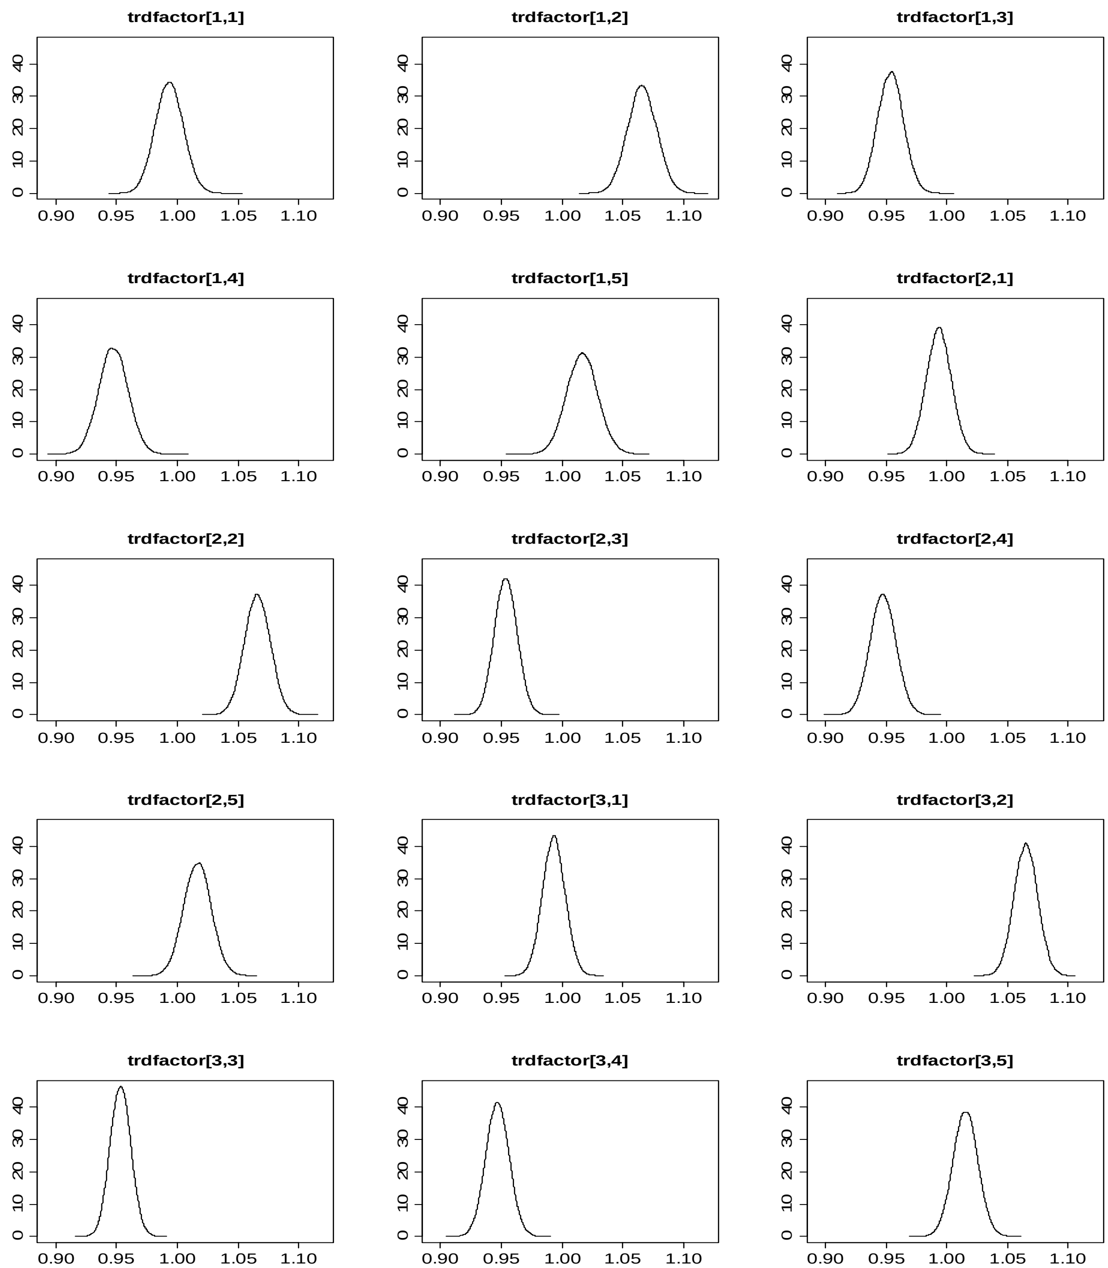

Note: The first index in brackets refers to the policy year of the experience period at which the trend factor attaches (1: final; 2: penultimate; 3: antepenultimate). The second index represents the series (1: indemnity severity; 2: medical severity; 3: frequency; 4: indemnity loss ratio; 5: medical loss ratio).

Chart 9: Posterior Densities for Adjustment Factors

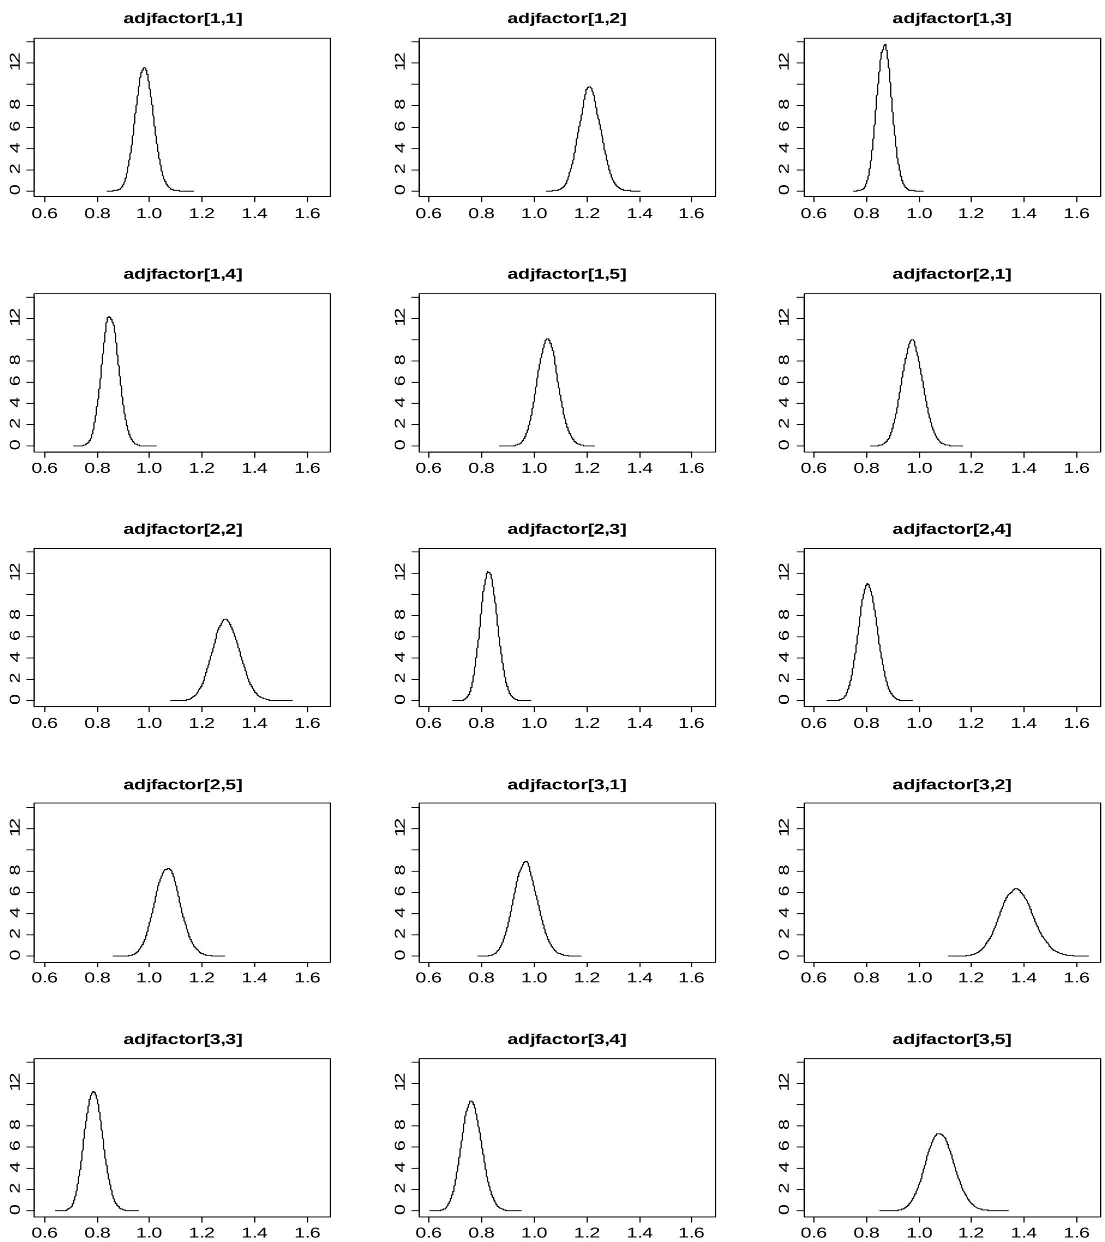

Note: The first index in brackets refers to the policy year of the experience period at which the adjustment factor attaches (1: final; 2: penultimate; 3: antepenultimate). The second index represents the series (1: indemnity severity; 2: medical severity; 3: frequency; 4: indemnity loss ratio; 5: medical loss ratio).

Table 1 exhibits the trend and adjustment factors, along with 95 percent credible intervals. Note that these intervals are not necessarily symmetric around the forecast values.

Table 1: Trend Factors and Adjustment Factors

Policy Year Paid							
Frequency							
Year	Trend Period	TF Lower Bound	**Mean Trend Factor(TF)**	TF Upper Bound	AF Lower Bound	**Mean Adjustment Factor(AF)**	AF Upper Bound
2003	5.001	0.9370	0.9533	0.9698	0.7223	0.7878	0.8579
2004	4.001	0.9357	0.9537	0.9721	0.7664	0.8276	0.8929
2005	3.001	0.9334	0.9537	0.9745	0.8131	0.8677	0.9253
Indemnity Severity							
Year	Trend Period	TF Lower Bound	**Mean Trend Factor(TF)**	TF Upper Bound	AF Lower Bound	**Mean Adjustment Factor(AF)**	AF Upper Bound
2003	5.001	0.9751	0.9932	1.0120	0.8817	0.9674	1.0600
2004	4.001	0.9735	0.9934	1.0140	0.8980	0.9746	1.0560
2005	3.001	0.9713	0.9935	1.0160	0.9164	0.9809	1.0490
Medical Severity							
Year	Trend Period	TF Lower Bound	**Mean Trend Factor(TF)**	TF Upper Bound	AF Lower Bound	**Mean Adjustment Factor(AF)**	AF Upper Bound
2003	5.001	1.0460	1.0651	1.0840	1.2530	1.3718	1.4970
2004	4.001	1.0450	1.0660	1.0870	1.1930	1.2921	1.3960
2005	3.001	1.0430	1.0660	1.0890	1.1340	1.2118	1.2930
Indemnity Loss Ratio							
Year	Trend Period	TF Lower Bound	**Mean Trend Factor(TF)**	TF Upper Bound	AF Lower Bound	**Mean Adjustment Factor(AF)**	AF Upper Bound
2003	5.001	0.9282	0.9468	0.9656	0.6888	0.7617	0.8395
2004	4.001	0.9268	0.9474	0.9683	0.7378	0.8063	0.8792
2005	3.001	0.9243	0.9475	0.9709	0.7897	0.8509	0.9151
Medical Loss Ratio							
Year	Trend Period	TF Lower Bound	**Mean Trend Factor(TF)**	TF Upper Bound	AF Lower Bound	**Mean Adjustment Factor(AF)**	AF Upper Bound
2003	5.001	0.9954	1.0153	1.0350	0.9772	1.0800	1.1900
2004	4.001	0.9945	1.0166	1.0390	0.9783	1.0689	1.1650
2005	3.001	0.9919	1.0166	1.0420	0.9759	1.0512	1.1310

Note: The trend period is measured in years. The interval between upper and lower bounds covers 95 percent of the probability mass of the distribution of the forecast.

4. CONCLUSIONS

NCCI has developed a Bayesian statistical model for estimating the trend rates of growth of the indemnity and medical severities, frequency, and the indemnity and medical loss ratios in the context of ratemaking. The model is purpose-built for short, volatile and potentially nonstationary time series and calibrated to minimize the prediction error. Further, the model accounts for common shocks, is robust to outliers, and is capable of interpolating where observations for the mentioned rates of growth are missing. Finally, by means of incorporating an add-up constraint, the model ensures consistent forecasts for the five time series in question.

Acknowledgment

Thanks to Harry Shuford for comments and to Chris Laws, Jose Ramos, and Manuel de la Guardia for research assistance.

5. REFERENCE S

[1] Carlin, Bradley P., and Thomas A. Louis, *Bayes and Empirical Bayes Methods for Data Analysis,* 2nd ed., **2000**, Boca Raton (FL): Chapman & Hall/CRC.
[2] Evans, Jonathan P., and Frank Schmid, *Forecasting Workers Compensation Severities and Frequency Using the Kalman Filter*, Casualty Actuarial Society *Forum*, **2007**: Winter, 43–660.
[3] Gelman, Andrew, John B. Carlin, Hal S. Stern, and Donald B. Rubin, *Bayesian Data Analysis,* 2nd ed., **2004**, Boca Raton (FL): Chapman & Hall/CRC.
[4] Hamilton, James D., *Time Series Analysis,* **1994**, Princeton (NJ): Princeton University Press.
[5] WinBUGS, *User Manual: WinBUGS 1.4.3*, **2007**, http://www.mrc-bsu.cam.ac.uk/bugs/.

Abbreviations and notations

MCMC, Markov-Chain Monte Carlo
NCCI, National Council on Compensation Insurance
RMSPE, Root Mean Squared Prediction Error

Biographies of the Author

Frank Schmid, Dr. habil., is a Director and Senior Economist at the National Council on Compensation Insurance.