

MODELING HIDDEN EXPOSURES IN CLAIM SEVERITY VIA THE EM ALGORITHM

Grzegorz A. Rempala* and Richard A. Derrig†

ABSTRACT

We consider the issue of modeling the latent or hidden exposure occurring through either incomplete data or an unobserved underlying risk factor. We use the celebrated expectation-maximization (EM) algorithm as a convenient tool in detecting latent (unobserved) risks in finite mixture models of claim severity and in problems where data imputation is needed. We provide examples of applicability of the methodology based on real-life auto injury claim data and compare, when possible, the accuracy of our methods with that of standard techniques. Sample data and an EM algorithm program are included to allow readers to experiment with the EM methodology themselves.

1. INTRODUCTION

Actuarial analysis can be viewed as the process of studying profitability and solvency of an insurance firm under a realistic and integrated model of key input random variables such as loss frequency and severity, expenses, reinsurance, interest and inflation rates, and asset defaults. In a modern analysis of financial models of property-casualty companies, these input variables typically can be classified into financial market variables and underwriting variables (cf., e.g., D'Arcy et al. 1997). The financial variables generally refer to asset-side-generated cash flows of the business, and the underwriting variables relate to the cash flows of the liabilities side. The process of developing any actuarial model begins with the creation of probability distributions of these input variables, including the establishment of the proper range of values of input parameters. The use of parameters is generally determined by the use of the parametric families of distributions, although the nonparametric techniques have a role to play as well (see, e.g., Derrig, Ostaszewski, and Rempala 2001). In this paper we consider an issue of hidden risk factors in modeling the input variables and point out the possible use of the celebrated expectation-maximization (EM) algorithm to uncover them. The methodology we present here is related to the general statistical methods for the analysis of the so-called *latent trait models*, but we shall not make that connection in our exposition explicitly. Readers interested in that aspect of the methodology are referred to the monographs on the topic of latent traits and classes (e.g., McCutcheon 1987) or a recent collection of papers in Hagenaars and McCutcheon (2001).

We present a way of addressing the issue of modeling hidden factors in the context of studying the loss distributions, as their proper modeling is essential in the actuarial analysis. Cummins and Wiltbank (1983) advance the notion that multivariate and dependent frequency and severity distributions are appropriate for aggregate losses and suggest that some complexity in loss modeling may be appropriate and necessary. In Cummins and Wiltbank (1985) they add the notion of heterogeneous claim types (lines of insurance) in modeling aggregate loss data. To keep things in focus we will be concerned here only with two applications to modeling the severity of loss, but the methods discussed may be applied

* Grzegorz A. Rempala, PhD, is a visitor of the Institute for Mathematics and Its Applications, University of Minnesota, and Department of Mathematics, University of Louisville, garemp01@rdos.math.louisville.edu, rempala@louisville.edu.

† Richard A. Derrig, PhD, is a Consultant with Opal Consulting LLC, 91 Fosdyke Street, Providence, RI 02906, richard@derrig.com.

easily to other actuarial problems like loss frequencies (injury types), asset returns (equity types), and asset defaults (rating subtypes). Indeed, in complex areas of actuarial decision making, many different factors are involved, some of them being stochastic (random) in nature. Taking account of each of them individually, that is, making independence assumptions, can be hurtful (see, e.g., the famous examples of serious bias introduced by the assumption of frequency/severity independence discussed by Bailey and Simon [1960] in the context of auto-insurance ratemaking). On the other hand, taking account of all interactions is intractable and is sometimes known as the *curse of dimensionality* (see, e.g., Dugas et al. 2003). In practice, knowledge of the underlying data and actuarial judgment are used to discover the most relevant of these interactions. We feel that often the *latent trait* or *hidden factor* models could be beneficial in making such judgments when attempting to improve data fit and/or making predictions into the future. In this paper we shall illustrate the use of the EM algorithm in the actuarial context via two simple examples. In the first example, we discuss the use of the EM method to impute missing values in an asset portfolio and compare its performance with that of the standard multiple imputation procedures implemented in modern statistical software. In the second example we shall screen medical bills for possible fraud or abusive practices and argue the usefulness of the EM algorithm in medical provider classification problems. The former example serves as a sample application to problems actuaries face often (records with sporadic missing values), while the latter one exemplifies screening for heterogeneity. Throughout the paper, depending upon the particular context, we shall sometimes refer to the hidden risk exposures in loss severity as “hidden factors” or in the context of a specific likelihood-based model as “hidden parameters.”

As suggested by the general statistical methodology of the latent trait models (see, e.g., McCutcheon 1987) one often can model hidden risk exposures as additional dimension(s) of the loss severity distribution. This in turn, in many cases, leads to considering mixtures of probability distributions as the model of choice for losses affected by hidden exposures with some parameters of the mixture considered missing (i.e., unobservable in practice). During the last 20 years or so there has been a considerable advancement in statistical methodologies dealing with partially hidden or incomplete data models. Empirical data imputation has become more sophisticated, and the availability of ever faster computing power has made it increasingly possible to solve these problems via both iterative and Monte Carlo-based algorithms (see, e.g., the review paper by Rubin [1996] or the monograph by Schafer [1997] and the references therein).

In what follows we proceed by describing the EM algorithm in Section 2, and by applying the algorithm to missing data and fraud detection problems in Sections 3 and 4. Section 5 concludes with some final remarks.

2. THE EM ALGORITHM

In their seminal paper Dempster, Laird, and Rubin (1977) have proposed the methodology that they have called the expectation-maximization (EM) algorithm as an iterative way of finding maximum likelihood estimates.¹ They demonstrated that the method was especially appropriate for finding the parameters of an underlying distribution from a given data set where the data were *incomplete* or had missing values. At present there are two basic applications of the EM methodology considered in the statistical literature. The first occurs when the data indeed have missing values, due to problems with, or limitations of, the data collection process. The second occurs when the original likelihood estimation problem is altered by assuming the existence of the *hidden* parameters. It turns out that both these circumstances can be described, at least initially, in the following statistical setting. Let us consider a density function (possibly multivariate) $p(\cdot|\Theta)$ that is indexed by the set of parameters Θ . As a simple example we may take p to be a univariate Gaussian density and $\Theta = \{(\mu, \sigma) \mid -\infty < \mu < \infty, \sigma > 0\}$.

¹ A full explanation of the role of the EM algorithm in missing data problems can be found in Little and Rubin (1987) or in a somewhat more mathematically advanced monograph by McLachlan and Krishnan (1997).

Additionally, we have an observed data set \mathcal{X} of size n , drawn from the distribution p . More precisely, we assume that the points of $\mathcal{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ are the realizations of some independent random variables distributed according to $p(\cdot|\Theta)$. We shall call \mathcal{X} the *incomplete* data. In addition to \mathcal{X} , we also consider a *complete* data set $\mathcal{Z} = (\mathcal{X}, \mathcal{Y})$ and specify the joint density

$$p(\mathbf{z}|\Theta) = p(\mathbf{x}, \mathbf{y}|\Theta) = p(\mathbf{y}|\mathbf{x}, \Theta) p(\mathbf{x}|\Theta). \quad (2.1)$$

As we can see from the last equality, this joint density $p(\mathbf{z}|\Theta)$ arises from considering the marginal density $p(\mathbf{x}|\Theta)$ and the specific assumptions on the relation between hidden (or missing) variables $\mathcal{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)$ and the observed incomplete data \mathcal{X} . Associated with the joint density is the joint likelihood function

$$\mathcal{L}(\Theta|\mathcal{Z}) = \mathcal{L}(\Theta|\mathcal{X}, \mathcal{Y}) = \prod_{i=1}^n p(\mathbf{x}_i, \mathbf{y}_i|\Theta),$$

which is often called the *complete* likelihood. For the sake of computational simplicity it is often more convenient to consider the logarithm of the complete likelihood

$$l(\Theta|\mathcal{Z}) = \log \mathcal{L}(\Theta|\mathcal{X}, \mathcal{Y}) = \sum_{i=1}^n \log p(\mathbf{x}_i, \mathbf{y}_i|\Theta). \quad (2.2)$$

Note that the function above may be thought of as a random variable since it depends on the unknown or missing information \mathcal{Y} , which, by assumption, is governed by an underlying probability distribution. Note also that in accordance with the likelihood principle, we now regard \mathcal{X} as constant.

The EM algorithm as described in Dempster, Laird, and Rubin (1977) consists of two steps repeated iteratively. In its *expectation* step (E-step), the algorithm first finds the expected value of the complete log-likelihood function $\log p(\mathcal{X}, \mathcal{Y}|\Theta)$ with respect to the unknown data \mathcal{Y} given the observed data \mathcal{X} and the current parameter estimates. That is, instead of the complete log-likelihood (2.2) we consider the following:

$$Q(\Theta, \Theta^{(i-1)}) = E[\log p(\mathcal{X}, \mathcal{Y}|\Theta)|\mathcal{X}, \Theta^{(i-1)}]. \quad (2.3)$$

Note the presence of the second argument in the function $Q(\Theta, \Theta^{(i-1)})$. Here $\Theta^{(i-1)}$ stands for the current value of the parameter Θ at the iteration $(i - 1)$: that is, the value that is used to evaluate the conditional expectation.

After the completion of the E-step, the second step of the algorithm is to maximize (with respect to Θ) the expectation computed in the first step. This is called the *maximization* step (M-step), at which time the value of Θ is updated by taking

$$\Theta^{(i)} = \underset{\Theta}{\operatorname{argmax}} Q(\Theta, \Theta^{(i-1)}). \quad (2.4)$$

The steps are repeated until convergence. It can be shown (via the relation (2.1) and Jensen's inequality) that if Θ^* maximizes $Q(\Theta, \Theta^{(i-1)})$ with respect to Θ for fixed $\Theta^{(i-1)}$, then

$$l(\Theta^*|\mathcal{Z}) - l(\Theta^{(i-1)}|\mathcal{Z}) \geq Q(\Theta^*, \Theta^{(i-1)}) - Q(\Theta^{(i-1)}, \Theta^{(i-1)}) \geq 0,$$

and each iteration of the procedure indeed increases the value of complete log-likelihood (2.2). Let us note that from the above argument it follows that a full maximization in the M-step is not necessary: it suffices to find any value of $\Theta^{(i)}$ such that $Q(\Theta^{(i)}, \Theta^{(i-1)}) > Q(\Theta^{(i-1)}, \Theta^{(i-1)})$. Such procedures are called GEM (*generalized EM*) algorithms. For a complete set of references see, for instance, the monograph by McLachlan and Krishnan (1997), where also the issues of convergence rates for the EM and GEM algorithms are thoroughly discussed. For some additional references and examples see also Wu (1983) or the monographs by Little and Rubin (1987) and Hastie, Tibshirani, and Friedman (2001).

3. MODELING HIDDEN EXPOSURES IN LOSS SEVERITY VIA THE EM ALGORITHM

In our two examples presented in the paper we shall consider two types of circumstances involving hidden (latent) factors. The first one will arise when, due to some problems with the data collection, parts of the observations are randomly missing from the observed data set. We consider this problem via the EM method in the particular context of multivariate (loss) models. The standard methodology for multiple data imputation is based on the three-step approach or IAC (imputation, analysis, combination; see, e.g., Rubin 1996). Here we discuss only the first step, imputation.

Although insurance has traditionally built on the assumption of independent risks, and the law of large numbers has governed the determination of premiums, the increasing complexity of insurance and reinsurance products has led over the past decade to increased actuarial interest in the modeling of dependent risks (see, e.g., Wang 1998; Embrechts, McNeil, and Straumann 2000). Recall here the use of multivariate distributions in the modeling of the multivariate aggregate claims and of the total claims process as put forth by Cummins and Wiltbank (see, e.g., Cummins and Wiltbank 1983, 1985).

In some circumstances, however, parts of the observed multivariate data may be missing. Claim-reporting systems depend heavily on the front-line adjusters to provide data elements beyond the simple payment amounts. In the absence of, or even in the presence of, system edits, daily work load pressures and the lack of interest in the coded data provide a deadly combination of disincentives for accurate and complete coding. Actuaries are quite familiar with missing data fields, which when essential to the analysis most often results in throwing the record out, thereby creating unknown “hidden” biases. Likewise, financial time series data may be interrupted, unavailable, or simply lost for securities or portfolios that are not widely tracked.

As an illustration of an application of the EM algorithm in this setting, let us consider a hypothetical example of 10 losses/gains from a two-dimensional vector of risk portfolios, which we have generated using a bivariate normal distribution. The data are presented in the following list of fictitious observed gains and losses from two risk portfolios (in thousands of dollars):

Gains	0.914	2.088	2.644	0.477	-1.940	-0.245	0.362	1.147	?	?
Losses	3.855	4.025	2.092	3.400	1.520	2.626	?	?	5.473	6.235

Note that parts of the last four observations are missing from the list. In fact, for the purpose of our example, they have been removed from the generated data. We illustrate the usefulness of the EM algorithm in estimating these missing values. Note that the observed rowwise averages are (0.681, 3.653). Actuaries often are tempted to ignore missing data and use average summary statistics for the data observed, hoping the missing data bias will be small.

If we denote by \mathcal{X} the observed (incomplete) data listed in the list above, then following our notation from the previous section, we have the complete data vector \mathcal{Z} given by

$$\mathcal{Z} = (\mathbf{z}_1 \dots \mathbf{z}_n) = (\mathbf{x}_1 \dots, \mathbf{x}_6, (x_{1,7}, y_{2,7})^T, (x_{1,8}, y_{2,8})^T, (y_{1,9}, x_{2,9})^T, (y_{1,10}, x_{2,10})^T),$$

where $\mathbf{x}_j = (x_{1,j}, x_{2,j})^T$ for $j = 1 \dots, 6$ is the set of pairwise complete observations. The missing data (corresponding to the question marks in the list) are, therefore,

$$\mathcal{Y} = (y_{2,7}, y_{2,8}, y_{1,9}, y_{1,10}). \tag{3.1}$$

Note that under our assumption of normality, equation (2.2) now becomes

$$l(\Theta|\mathcal{Z}) = -n \log(2\pi) - \frac{1}{2} \log|\Sigma| - \frac{1}{2} \sum_{j=1}^n (\mathbf{z}_j - \boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{z}_j - \boldsymbol{\mu})$$

where $n = 10$, $\boldsymbol{\mu} = (\mu_1, \mu_2)^T$ is a vector of means, and

$$\Sigma = \begin{pmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{12} & \sigma_{22} \end{pmatrix}$$

is the covariance matrix. The vector of unknown parameters, therefore, can be represented as

$$\Theta = (\mu_1, \mu_2, \sigma_{11}, \sigma_{12}, \sigma_{22}). \quad (3.2)$$

To describe the EM algorithm in this setting we need to find the particular form of $Q(\Theta, \Theta^{(i-1)})$ defined by equation (2.3). Due to the independence of the z_i 's this is equivalent, in effect, to evaluating

$$E_{\Theta^{(i-1)}} [\mathcal{Y}|\mathcal{X}] \quad \text{and} \quad E_{\Theta^{(i-1)}} [Y^2|\mathcal{X}],$$

where Y is the underlying random variable for \mathcal{Y} , assumed to be normal. From the general formulae for conditional moments of a bivariate normal variable $W = (W_1, W_2)$ with the set of parameters \mathcal{X} as above, we have that

$$\begin{aligned} E[W_2|W_1 = w_1] &= \mu_2 + \sigma_{12}/\sigma_{11}(w_1 - \mu_1), \\ \text{Var}(W_2|W_1 = w_1) &= \sigma_{22.1} = \sigma_{22}(1 - \rho^2), \end{aligned} \quad (3.3)$$

where ρ stands for the correlation coefficient. Interchanging the subscripts 1 and 2 in equation (3.3) gives the formulae for the conditional mean and variance of the distribution $W_1|W_2 = w_2$. Using the relations (3.3) and the usual formulae for ML estimators of the normal mean vector μ and the covariance matrix Σ , we may now state the EM algorithm for imputing missing data in the earlier list as follows.

Algorithm 1 (EM Version of Buck's Algorithm)

1. Define the initial value $\Theta^{(0)}$ of the set of parameters (3.2). Typically, it can be obtained on the basis of the set of complete pairs of observations (i.e., $x_1 \dots, x_6$ in the list above).
2. The E-step: given the value of $\Theta^{(i)}$, calculate via equation (3.3) the vector $\mathcal{Y}^{(i)}$ of the imputations of the missing data \mathcal{Y} given by equation (3.1):

$$\begin{aligned} y_{2k}^{(i)} &= \mu_2^{(i)} + \frac{\sigma_{12}^{(i)}}{\sigma_{11}^{(i)}} (x_{1k} - \mu_1^{(i)}) \quad \text{and} \quad y_{2k}^2{}^{(i)} = (y_{2k}^{(i)})^2 + \sigma_{22.1}^{(i)} \quad \text{for } k = 7, 8, \\ y_{1k}^{(i)} &= \mu_1^{(i)} + \frac{\sigma_{12}^{(i)}}{\sigma_{22}^{(i)}} (x_{2k} - \mu_2^{(i)}) \quad \text{and} \quad y_{1k}^2{}^{(i)} = (y_{1k}^{(i)})^2 + \sigma_{11.2}^{(i)} \quad \text{for } k = 9, 10. \end{aligned}$$

3. The M-step: given the current value of the imputed complete data vector $\mathcal{X}^{(i)} = (\mathcal{X}, \mathcal{Y}^{(i)})$, set $M_k = \sum_{j=1}^n z_{kj}^{(i)}/n$ and $M_{kl} = \sum_{j=1}^n z_{kj}^{(i)} z_{lj}^{(i)}/n$ for $k, l = 1, 2$, and calculate $\Theta^{(i+1)}$ as

$$\begin{aligned} (\mu_1^{(i+1)}, \mu_2^{(i+1)}) &= (M_1, M_2), \\ \sigma_{kl}^{(i+1)} &= M_{kl} - M_k M_l \quad \text{for } k, l = 1, 2. \end{aligned}$$

4. Repeat steps 2 and 3 until the relative difference of the subsequent values of $l(\Theta^{(i+1)}|\mathcal{X}^{(i)})$ is sufficiently small.

The above algorithm in its noniterative version was first introduced by Buck (1960), who used the method of imputation via linear regression with subsequent covariance correction to estimate means and covariance matrices of p -dimensional random vectors in the case where some parts of the vector components were missing. Note that in our iterative version of the procedure this imputation is performed in step 2 (i.e., the E-step) by replacing the missing values with the appropriate conditional expectations. For more details about Buck's imputation procedure, we refer to his original paper (Buck 1960) or to Chapter 3 of Little and Rubin (1987) or Chapter 2 of McLachlan and Krishnan (1997).

The numerical illustration of the algorithm is presented in Table 1. As we can see from the table, with accuracy of up to three significant digits, the algorithm seems to converge after about 30 steps or so, and the estimated or imputed values of equation (3.1) are given by

Table 1
Selected Iterations of the EM Algorithm for Data in List of Gains and Losses

Iteration	μ_1	μ_2	σ_{11}	σ_{12}	σ_{22}	$\gamma_{2,7}$	$\gamma_{2,8}$	$\gamma_{1,9}$	$\gamma_{1,10}$	-2Q
1	0.6764	3.5068	1.8170	0.3868	2.0671	3.4399	3.6069	1.0443	1.1867	65.7704
5	0.8779	3.6433	1.8618	0.8671	2.2030	3.4030	3.7685	1.5982	1.8978	64.7568
10	0.9279	3.6327	1.9463	0.9837	2.1724	3.3466	3.7433	1.7614	2.1061	64.5587
20	0.9426	3.6293	1.9757	1.0181	2.1639	3.3301	3.7345	1.8102	2.1683	64.5079
30	0.9435	3.6291	1.9775	1.0202	2.1634	3.3291	3.7339	1.8132	2.1722	64.5048
35	0.9436	3.6291	1.9776	1.0203	2.1633	3.3290	3.7339	1.8134	2.1724	64.5047
40	0.9436	3.6291	1.9777	1.0204	2.1633	3.3290	3.7339	1.8134	2.1724	64.5046
45	0.9436	3.6291	1.9777	1.0204	2.1633	3.3290	3.7339	1.8134	2.1725	64.5046

$$\hat{y}^{(em)} = (3.329, 3.734, 1.813, 2.173).$$

To compare this result with a standard multiple imputation procedure, we have analyzed the same data set using a much more sophisticated imputation model of *Bayesian linear regression* as given in Rubin (1987, p. 167) and implemented in the SAS PROC MI procedure or in the Multiple Imputations for Chained Equations (MICE) S-Plus library.² We have selected this particular imputation method as it is often the method of choice for imputation of multivariate continuous data similar to that presented in the list above. The results of the analysis via SAS and S-Plus functions were quite similar. Here we present the result of MICE software run for imputing missing values based on the remaining set of available observations in our list averaged over 10 separate imputations:

$$\hat{y}^{(Breg)} = (2.855, 3.112, 2.229, 3.520).$$

Both $\hat{y}^{(em)}$ and $\hat{y}^{(Breg)}$ now can be compared with the actual values removed from the list, which were

$$y = (3.362, 3.657, 1.484, 3.410).$$

It is perhaps also of interest to point out that the “naive” imputation performed via the sample average method, which ignores the dependence structure in the data, would yield in this particular case

$$\hat{y}^{(avg)} = (3.653, 3.653, 0.681, 0.681).$$

As we can see, in our example the EM method did reasonably well in recovering the missing values, and its performance was comparable to that of a much more complicated method that used multiple imputation runs and was comparably computationally intensive. Both methods were noticeably better than imputation via the sample average, which resulted in a large error for two of the four missing values.

4. MASSACHUSETTS AUTO BODILY INJURY LIABILITY DATA: FRAUD AND BUILD-UP SCREENING VIA MIXTURE MODELS

By now it is fairly well known that fraud and build-up, exaggerated injuries and/or excessive treatment, are key components of the auto injury loss distributions (Derrig, Weisberg, and Chen 1994; Cummins and Tennyson 1996; Abrahamse and Carroll 1999). Indeed, auto injury loss distributions are prime candidates for mixture modeling, for at least the different payment patterns by injury type. Even within an auto injury type as predominant as strain and sprain,³ there can be substantial differences in sub-populations arising from fraud and build-up. One common method of identifying these claims has been

² For some insight into most recent imputation methodologies and available software, see, for instance, an excellent review article by Horton and Lipsitz (2001). A good general reference in the area is Schafer (1997).

³ Currently Massachusetts insured auto bodily injury claims are upwards of 80% strain and sprain claims as the most costly part of the medical treatment. Of course, that may have a dependency on the \$2,000 threshold to file a tort claim.

to gather additional features of the claim, the so-called *fraud indicators*, and to build models to identify those bogus claims (Brockett, Xia, and Derrig 1998). The acquisition of reliable indicators, some of which may be highly subjective, is costly and may not be efficient in uncovering abusive patterns in injury claims (Crocker and Tennyson 1999). Finally, even when the indicator characteristics can be gathered efficiently during claim processing, the objective trait of fraud and build-up is clearly latent, requiring an indirect or unsupervised methods of detection (Brockett, Xia, and Derrig 1998; Brockett et al. 2002).

Another approach to uncovering fraud and build-up, perhaps grounded more in practical considerations, is to construct a filter, or screening algorithm, for medical provider bills (Derrig 2002). Massachusetts's no-fault system with a monetary threshold (\$2,000) for tort liability provides the incentive for so-called *medical mill billing practices*, that is, a tight distribution of total bills with a mean slightly in excess of \$2,000. Routinely, excessive medical bills can be reduced to "reasonable and customary" levels by computer algorithms that compare incoming bills to right censored billing distributions, with "excessive" operationally defined to be above the censoring point. Less routine is the implementation of systematic analysis of the patterns of *a provider's billing practices* (Major and Riedinger 1992). Our second application of the EM algorithm is to build a first-level screening device to uncover potential abusive billing practices and, most importantly, the appropriate set of claims to review. We perform the pattern analysis by uncovering abusive-like distributions within mixture models parametrized by the estimates obtained via the EM algorithm. An illustration of the method follows.


In the table provided in Appendix B, we present a set of outpatient medical providers' total billings on a sample of 348 auto bodily injury liability claims closed in Massachusetts during 2001. For illustration purposes, 76 claims with one provider ("A") were chosen based on a pattern typical of abusive practice, namely, an empirical kurtosis more than five times the overall average. The claims with doctor "A" providing medical treatment were then combined with medical bills in claims from a random sample of providers. The losses are recorded in thousands and are presented in column two. Column four identifies each medical billing amount as provider "A" or "other." Attempting to uncover the claims generated by provider A, we will use the EM algorithm applied to a normal mixture model for the log-claim values.

The relatively large volume of provider A's claims is clearly visible in the left panel of Figure 1, where it is presented as a portion of the overall claims. Whereas the volume of claims by itself never constitutes a basis for the suspicion of fraud or build-up, it certainly might warrant a closer look at the data in hand, especially via some type of homogeneity analysis, since the right panel in Figure 1 clearly indicates the difference between the overall claims distribution and that of the provider A. Hence, in this problem we look for a hidden exposure that could manifest itself as a nonhomogenous component of the data, but not assuming that this component is necessarily due to provider A. In fact, as the initial inspection of the *overall* data distribution does not immediately indicate nonhomogeneity, we shall not make any prior assumptions about the nature or source of the possible nonhomogeneity.

Since the standard analysis of the data by fitting a kernel density estimator⁴ (see the solid curve in the right panel of Figure 1) appears to give no definite indication of multimodality, it seems that some more sophisticated methods are needed to identify any foreign components of the claims. Whereas many different approaches to this difficult problem are possible (see, e.g., Hastie, Tibshirani, and Friedman 2001 for a review), we have chosen one that shall illustrate the applicability of the EM methodology in our setting. Namely, we attempt to fit a *mixture of normal distributions* to the log-scaled data; that is, we model the logarithm of the claim outpatient medical billing distribution as a mixture of several normal variables. The use of normal distributions here is due mostly to convenience of the EM implementation, which in more complicated real-life problems can be inappropriate. However, the principle

⁴ Here we used a standard nonparametric kernel density estimator with a Gaussian kernel. See, e.g., Bolancé, Guillen, and Nielsen (2003) for the discussion of this methodology in the context of actuarial loss models and auto claims. Thanks are due to a referee for pointing out this reference.

Figure 1
Overall Distribution of the 348 BI Medical Bill Amounts from Appendix B Compared with That Submitted by Provider A



Notes: Left panel: Frequency histograms (provider A's histogram in solid bars). Right panel: Density estimators (provider A's density in the dashed line).

that we shall attempt to describe here is, in general, applicable to any mixture of distributions, even including nonparametric ones.⁵ For an extensive discussion of the available implementations of the EM method for a variety of different parametric mixture models, see McLachlan and Peel (2000).

Let us note that the EM algorithm has been used previously in the context of auto insurance claims by Kalb, Kofman, and Vorst (1996), who used it to fit a mixture of two distributions, a discrete and a continuous one, to the claim severity data to improve parametric model fit and adjust for the claims clustering effect. Unlike claim settlement data of Kalb et al., our data set does not suggest discrete clusters of medical bills, but rather differing billing practices represented by mixtures of continuous distributions. Thus, in contrast with Kalb et al., in the current setting we are interested primarily in fitting the mixture model to *classify* observations. The difference is subtle but important, as we are less concerned with the overall goodness of fit of our model and more with its ability to distinguish between the observations arriving from the different mixture components.

To describe our method in the context of the EM algorithm, we relate the problem to our EM methodology introduced in Section 1. In our current setting we consider the set of logarithms of the BI claim medical bills as the incomplete data \mathcal{X} . According to our model assumption, we identify the underlying random variable X , of which \mathcal{X} is a realization, as a mixture of several (say, $m \geq 2$) normal distributions. For a general reference on the mixture models and their applications in statistical modeling the reader is referred to the recent monograph by McLachlan and Peel (2000). Hence, if $X_j \sim N(\mu_j, \sigma_j)$ for $j = 1, \dots, m$, then we may write⁶

⁵ The notion of fitting nonparametric distributions via likelihood methods, which at first may seem a contradiction in terms, has become very popular in statistics over the last decade. This is due to the intensive research into the so-called empirical likelihood methods (see, e.g., a recent monograph by Owen 2001 and references therein). In principle, with some modifications, the mixture approach discussed in this section and the associated EM algorithm can be applied to the empirical likelihood as well.

⁶ Note that in our notation σ denotes the variance, not standard deviation.

$$X = \sum_{j=1}^m Y_j \cdot X_j, \quad (4.1)$$

where $Y_j \in \{0, 1\}$ with $P(Y_j = 1) = \pi_j$ such that $\sum \pi_j = 1$ and the joint distribution of the binary vector (Y_1, \dots, Y_m) is multinomial with one trial ($\sum Y_j = 1$, i.e., there is only a single one among Y 's). The right-hand side of equation (4.1) is sometimes known as the *generative* representation of a mixture (see, e.g., Hastie, Tibshirani, and Friedman 2001, p. 237). Indeed, if we generate a multinomial variable (Y_1, \dots, Y_m) with probabilities of $Y_j = 1$ equal to π_j , and depending on the index j for which outcome is unity, deliver X_j , then it can be shown that the density of X is

$$\sum_{j=1}^m \pi_j p(x|\Theta_j), \quad (4.2)$$

where $p(\cdot|\Theta_j)$ is a normal density with the parameter

$$\Theta_j = (\mu_j, \sigma_j) \text{ for } j = 1, 2, \dots, m.$$

Hence X is indeed a random variable whose distribution is given in terms of a mixture of the distributions of the X_j 's. The density given by equation (4.2) is less helpful in our approach, as it does not explicitly involve the variables Y_j 's. Moreover, fitting the set of parameters,⁷

$$\Theta = (\Theta_1, \dots, \Theta_m, \pi_1, \dots, \pi_{m-1}), \quad (4.3)$$

by considering the log-likelihood of equation (4.2) is numerically difficult as it involves evaluation of the sums under the logarithm. In contrast, the representation (4.1) provides a simpler approach, which also better suits our purpose of illustrating the use of the EM methodology. In the spirit of the search for hidden exposures, we consider the (unobserved) realizations of the random vector (Y_1, \dots, Y_m) in equation (4.1) as the missing data \mathcal{Y} . Note that, unlike in the example discussed in Section 2, here we have in some sense artificially created the set \mathcal{Y} . In this setting the complete set of data is now $\mathcal{X} = (\mathcal{X}, \mathcal{Y})$ or $\mathbf{z}_j = (x_j, y_{jk})$ for $j = 1, \dots, n$, and $k = 1, \dots, m$. Here $n = 348$ is the number of observations, m is the number of components in the mixture, unspecified for now, x_j is (logarithm of) the observed medical bill value, and $y_{jk} \in \{0, 1\}$ is the auxiliary indicator variable indicating whether or not x_j arrives from the distribution of X_k . In this setting the complete log-likelihood function (2.2) takes the form

$$l(\Theta|\mathcal{X}) = \sum_{j=1}^n \sum_{k=1}^m y_{jk} \log p(x_j|\Theta_k), \quad (4.4)$$

and the conditional expectation (2.3) is given by

$$Q(\Theta, \Theta^{(i-1)}) = \sum_{j=1}^n \sum_{k=1}^m \delta_{jk} \log p(x_j|\Theta_k), \quad (4.5)$$

where

$$\delta_{jk} = E[Y_{jk}|\Theta^{(i-1)}, \mathcal{X}] = P(Y_{jk} = 1|\Theta^{(i-1)}, \mathcal{X}) \text{ for } j = 1, \dots, n; \quad k = 1, \dots, m. \quad (4.6)$$

As we can see from the above formulae, in this particular case $Q(\Theta, \Theta^{(i-1)})$ is obtained from the complete data likelihood by substituting for the unknown y_{jk} 's their conditional expectations δ_{jk} 's calculated under the current value of the estimates of Θ .⁸ The quantity δ_{jk} is often referred to as the *responsibility* of the component X_k for the observation j (see, e.g., Hastie, Tibshirani, and Friedman 2001, p. 238). This terminology reflects the fact that we may think about final δ_{jk} as the conditional (posterior) probability of the j th observation arriving from the distribution of X_k .

⁷ Note that we need to estimate only $m - 1$ proportions since $\sum \pi_j = 1$.

⁸ It may happen that some of the values y_{jk} are, in fact, available. In such cases we would take $\delta_{jk} = y_{jk}$.

Once we have replaced the y_{jk} 's in equation (4.4) by the δ_{jk} 's, the maximization step of the EM algorithm is straightforward and, applied to equation (4.5), gives the usual weighted ML estimates of the normal means, variances, and mixing proportions (see below for the formulae). However, to proceed with the EM procedure we still need to construct the initial guesses for the set of parameters (4.3). A good way to do so (for a discussion, see, e.g., Chapter 8 of Hastie, Tibshirani, and Friedman 2001; Xu and Jordan 1996) is simply to choose at random m of the observed claim values as the initial estimates of the means, and set all the estimates of the variances to the overall sample variance. The mixing proportion can be set uniformly over all components. This way of initiating the parameters ensures the relative robustness of the final estimates obtained via EM against any particular initial conditions. In fact, in our BI data example we have randomly selected several initial sets of values for the means, and in all cases have obtained convergence to the same set of estimates. Below we present the detailed EM algorithm we have used to analyze the Massachusetts auto BI data. To identify the number m of the mixture components in the model, we have used the EM method to obtain the estimates of the complete log-likelihood function (estimates based on the final values maximizing equation (4.5) plugged into the log-likelihood based on equation (4.2)) for $m = 2, 3, 4$ (we determined earlier that for $m > 4$ the BI mixture model becomes too cumbersome). The results are presented in Table 2. As can be seen from the last row of the table, $m = 3$ is the number of components minimizing the negative of the estimated log-likelihood based on equation (4.2). Henceforth, we take $m = 3$ for the BI mixture model. In our case the simple log-likelihood criterion worked well for selecting m ; however, under many circumstances one might want to weight the decrease in negative log-likelihood against the increase in the overall model complexity. For some insight into that problem see Chapter 7 in Hastie, Tibshirani, and Friedman (2001).

Algorithm 2 (The EM Algorithm for Fitting an m -Component Normal Mixture)

1. Define the initial estimate $\Theta^{(0)}$ of the set of parameters (4.3) (see discussion above).
2. The E-step: given the current value of $\Theta^{(i)}$ compute the responsibilities δ_{jk} as

$$\delta_{jk} = \frac{\pi_k^{(i)} p(x_j|\Theta_k^{(i)})}{\sum_{l=1}^m \pi_l^{(i)} p(x_j|\Theta_l^{(i)})} \quad j = 1, \dots, n \quad \text{and} \quad k = 1, \dots, m.$$

3. The M-step: compute the ML estimators of equation (4.5) as

$$\pi_k^{(i+1)} = \frac{\sum_{j=1}^n \delta_{jk}}{n} \quad \text{for} \quad k = 1, \dots, m - 1,$$

Table 2
Comparison of Mixture Fit for Different Values of m for BI Data

Parameter	m		
	2	3	4
μ_1	0.071	0.107	-0.01
μ_2	1.110	0.874	0.218
μ_3	—	1.248	0.911
μ_4	—	—	1.258
$\sigma_1^{1/2}$	1.265	1.271	1.201
$\sigma_2^{1/2}$	0.252	0.178	1.349
$\sigma_3^{1/2}$	—	0.146	0.214
$\sigma_4^{1/2}$	—	—	0.144
π_1	0.470	0.481	0.250
π_2	0.530	0.205	0.224
π_3	—	0.314	0.247
π_4	—	—	0.279
$-2\log L$	819.909	811.381	811.655

and

$$\mu_k^{(i+1)} = \frac{\sum_{j=1}^n \delta_{jk} x_j}{\sum_{j=1}^n \delta_{jk}},$$


$$\sigma_k^{(i+1)} = \frac{\sum_{j=1}^n \delta_{jk} (x_j - \mu_k^{(i+1)})^2}{\sum_{j=1}^n \delta_{jk}} \quad \text{for } k = 1, \dots, m.$$

4. Repeat steps 2 and 3 until the relative difference of the subsequent values of equation (4.5) is sufficiently small.

Let us remark that in our implementation of the above procedure for the sake of improving numerical stability as well as convergence speed we have used a relative change in the complete log-likelihood based on equation (4.2) instead of equation (4.5) as the convergence criterion. The difference seemed to have very little effect on the final estimated values of the 11 mixture parameters but did considerably improve the speed of the algorithm.

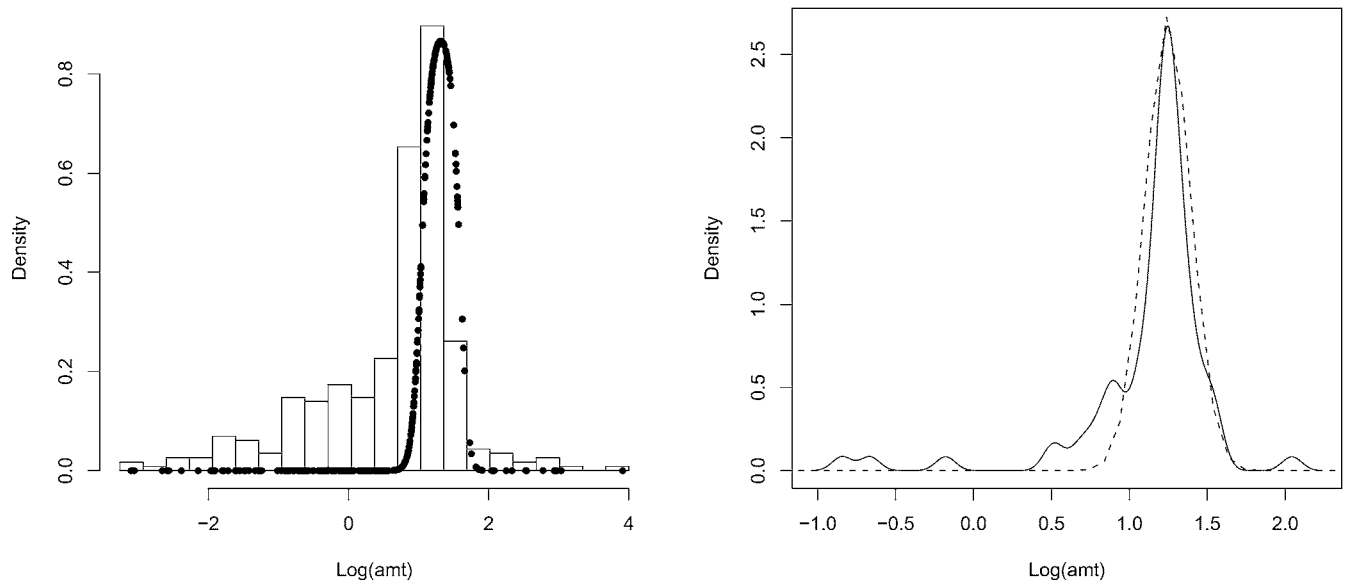
In Figures 2 and 3 we present graphically the results of the analysis of the BI data via the mixture model with $m = 3$ using the EM algorithm as described above. Some selected iterations of the EM algorithm for the three-component normal mixture are presented in Table 3. In the left panel of Figure 2 we show the fit of the normal mixture fitted to the data using Algorithm 2 (with parameter values given by the last row of Table 3). As we can see, the fit looks reasonable, and the fitted mixture distribution looks similar to the standard density estimator (solid curve in the right panel of Figure 1). The mixture components identified by the EM method are presented in the right panel of Figure 2 and clearly indicate nonhomogeneity of the data, which seems to consist of two (in fact, three) different types of claims. This is related to a high volume of claims in the interval around 1.8–4.5 thousands (corresponding to the values 0.6–1.5 on the log scale). This feature of the data is modeled by the two tall and thin (i.e., with small dispersion) components of the mixture (corresponding in our notation to X_2 and X_3 , marked as solid and dashed curves, respectively). We also note the very pronounced difference (over sevenfold) in the spread between the first and the two last components.

Figure 2
EM Fit



Notes: Left panel: Mixture of normal distributions fitted via the EM algorithm to BI data. Right panel: Three normal components of the mixture. The values of all the parameters are given in the last row of Table 3.

Figure 3
Latent Risk in BI Data Modeled by the EM Algorithm with $m = 3$



Notes: Left panel: Set of the responsibilities δ_{j3} . Right panel: The third component of the normal mixture compared with the distribution of provider A's claims ("A" claims density estimator is a solid curve).

In the left panel of Figure 3 we present the set of responsibilities (δ_{j3}) of the model (or component) X_3 as calculated by the EM algorithm superimposed on the histogram of the BI data. The numerical values of the responsibilities for each data point are also listed in the last column of the table in Appendix B. The relationship between the set of responsibilities obtained via the EM procedure and the apparent lack of homogeneity of the data, demonstrated by Figure 2, is easy to see. The high responsibilities are clustered around the claim values within two standard deviations of the estimated mean (1.25) of the tallest distribution X_3 . Hence the plot of responsibilities superimposed on the data distribution again uncovers the nonhomogeneity or the risk factor, which was initially hidden (doctor "A"). As we can see from the right panel in Figure 3, the observed nonhomogeneity may be attributed largely, as initially expected (and as the illustration intended), to the high kurtosis of "A" claims. Indeed, the superimposing of the distribution of "A" claims (solid curve) on the component X_3 (dashed curve) in the right panel of Figure 3 reveals a reasonably close match in the interval (0.8, 1.7) or so. Outside this interval the normal approximation to provider A's claims fails, mostly due to the fact that

Table 3
Selected Iterations of EM Algorithm for BI Data with $m = 3$

Iteration	μ_1	μ_2	μ_3	$\sigma_1^{1/2}$	$\sigma_2^{1/2}$	$\sigma_3^{1/2}$	π_1	π_2	π_3	$-2\log L$
1	0.229	0.785	0.885	1.172	0.89	0.843	0.35	0.329	0.321	973.115
5	-0.129	0.946	1.054	1.374	0.525	0.356	0.337	0.301	0.361	854.456
6	-0.131	0.953	1.083	1.357	0.499	0.300	0.349	0.281	0.370	839.384
10	-0.041	0.917	1.137	1.324	0.456	0.223	0.396	0.217	0.387	820.903
20	0.042	0.875	1.166	1.302	0.364	0.207	0.438	0.177	0.385	817.363
30	0.064	0.876	1.184	1.29	0.301	0.200	0.453	0.176	0.372	816.143
40	0.074	0.871	1.204	1.285	0.259	0.188	0.460	0.186	0.354	814.957
50	0.084	0.868	1.226	1.281	0.222	0.17	0.467	0.197	0.336	813.367
60	0.099	0.871	1.243	1.275	0.190	0.153	0.476	0.204	0.320	811.838
64	0.105	0.873	1.247	1.272	0.180	0.147	0.48	0.205	0.315	811.454
65	0.107	0.874	1.248	1.271	0.178	0.146	0.481	0.205	0.314	811.381

the normal mixture model employed is not sufficiently “fine tuned” in its tails to look for this particular type of distribution. The deficiency perhaps could be rectified in this particular case by incorporating some different (non-normal) components into the mixture model. However, our main task in this analysis is to uncover hidden factors (if any) and not necessarily to model them precisely, which should be done afterwards using some different, more sophisticated modeling approach depending on the type of problem in hand. See, for instance, Bilmes (1998), who presents the extension of our Algorithm 2 to the so-called general hidden Markov model (HMM). For a full review of some possible approaches to fitting the finite mixtures models and the use of the EM methodology in this context, readers are referred to the recent monograph by McLachlan and Peel (2000), which also contains some descriptions of the currently available software for fitting a variety of non-normal mixtures via the EM method.

4.1 Classification Via the Posterior Probabilities

As shown in the previous section, each run of the EM algorithm estimating an m -mixture model will produce responsibilities for each claim and for each one of the m mixture distributions. They can be interpreted as the (posterior) probability that the claim “arises” from each of the components of the mixing distributions. If for any claim the responsibility (membership) of a particular model component equals one, we say that the claim arises from that model component. When the responsibility is less than one, the claim arises partially from that component, and if the responsibility equals zero, we can say the claim does not arise from that component. In that context every claim “belongs to” each of the mixing components with measurement value equal to the responsibility. For a predetermined responsibility value (sometimes referred to as α -cut) this allows for claims classification. The principle is known as *kernel classification* and is widely used in statistical decision making (see, e.g., Hastie, Tibshirani, and Friedman 2001, Chapter 6, Section 6). The α -cut classification table is presented in Table 4 and shows the portions of “A” claims contained in each α -cut from 0.1 to 0.9 for each mixture component distribution. In particular, the α -cut analysis confirms our previous findings that “A” claims belong predominantly to the third mixing distribution (i.e., distribution of X_3). Indeed, the α -cut at about 0.5 provides us with a good indication that “A” arises from the third mixing distribution (corresponding to the value 75% in the table) but not from the first one (corresponding to 8% value only). These findings are consistent with those illustrated by Figure 3. In contrast, the second mixing distribution (distribution of X_2) does not allow us to classify “A” and “other” correctly in our three-mixture model. The low proportion of “A” claims assigned to the model X_2 , indicates that they are generally unlikely to arrive from X_2 , which may be an indication of some further nonhomogeneity among claims, even after adjusting for the type “A.” The X_2 component could therefore be the manifestation of some additional hidden factors, which again confirms the findings summarized in the previous section.⁹

4.2 Accuracy Assessment for the EM Output Via the Parametric Bootstrap

In our analysis of the BI data conducted in the previous sections we have used the numeric values of the estimated parameters (4.3) and the responsibilities (4.6). Since these values were estimated from the data via the EM algorithm, it is important to learn about their accuracy. In general, for the set of parameters (4.3) the usual approach to assessing accuracy based on the asymptotic normality of the maximum likelihood estimators can be applied here, as soon as we calculate the information matrix for Θ . This is slightly more complicated for the set of responsibilities (4.6), as they are the functions of Θ and hence require the appropriate transformation of the information matrix. However, a simpler method of obtaining, for instance, confidence intervals for the set of responsibilities and the model parameters, can be also used, based on the so-called *parametric bootstrap method* outlined in Algorithm 3 below. The method can be shown to be asymptotically equivalent to the normal approximation ap-

⁹ An analysis of the mixture model applied only to 272 “other” claims shows that X_2 has a more pronounced representation (high α -cut proportions) of (1) chiropractic and physical therapy treatment, (2) special investigations and independent medical examinations, and (3) extended treatment delays.

Table 4
Classification via Responsibilities for Different α -Cuts

α	Resp. X_1		Resp. X_2		Resp. X_3	
	A	Other	A	Other	A	Other
0.9	0.05	0.42	0.00	0.00	0.00	0.00
0.8	0.05	0.42	0.00	0.00	0.54	0.13
0.7	0.07	0.45	0.09	0.11	0.62	0.19
0.6	0.08	0.46	0.13	0.20	0.70	0.22
0.5	0.08	0.48	0.13	0.24	0.75	0.24
0.4	0.11	0.49	0.16	0.28	0.78	0.26
0.3	0.16	0.54	0.21	0.30	0.79	0.30
0.2	0.22	0.64	0.24	0.35	0.79	0.34
0.1	0.95	1.00	0.33	0.41	0.82	0.38

proach and is known to be often more reliable for smaller sample sizes or for the heavily biased estimators (which will often be the case for the responsibilities (4.6)). The algorithm describes how to obtain confidence intervals for the parameters given by equations (4.3) and (4.6) using the bootstrap. For some more examples and further discussion see, for instance, McLachlan and Peel (2000) or the recent paper by Rempala and Szatzschneider (2004), where the issue of the hypothesis testing for the number of mixture components via the parametric bootstrap method is also discussed.

Algorithm 3 (Bootstrap Confidence Intervals)

- 1 Using the values of the model parameters (4.3) obtained from the EM algorithm, generate the set of pseudo-data \mathcal{X}^* (typically of the same length as the original data \mathcal{X}).
- 2 With \mathcal{X}^* in hand, use Algorithm 2 to obtain a set of pseudo-values Θ^* .
- 3 Using the set of the original data values \mathcal{X} and Θ^* from step 2, calculate the pseudo-responsibilities δ_{jk}^* as in Algorithm 2, step 2.
- 4 Repeat the steps 1–3 a large number of times, say, B .¹⁰
- 5 Use the empirical quantiles of the distributions of pseudo-values Θ^* and δ_{jk}^* to obtain confidence bounds for Θ and δ_{jk} .

For illustration purposes, we present the set of confidence intervals for the three-mixture-normal model parameters and the responsibilities (of X_3) obtained via the above algorithm for the BI data in Tables 5 and 6. The term “bootstrap estimate” in the tables refers to the average value of the B bootstrap pseudo-values obtained in steps 2 or 3.

Table 5
Accuracy of Parameter Estimates for BI Data with $B = 1000$

Parameter	Value	Bootstrap Estimate	95% CI
μ_1	0.107	0.104	(-0.115, 0.298)
μ_2	0.874	0.871	(0.809, 0.924)
μ_3	1.248	1.249	(1.216, 1.284)
$\sigma_1^{1/2}$	1.271	1.269	(1.132, 1.389)
$\sigma_2^{1/2}$	0.178	0.175	(0.125, 0.222)
$\sigma_3^{1/2}$	0.146	0.144	(0.117, 0.174)
π_2	0.205	0.207	(0.157, 0.253)
π_3	0.314	0.317	(0.268, 0.375)

¹⁰ In our setting B needs to be fairly large, typically at least a thousand. For a discussion see, e.g., McLachlan and Peel (2000).

Table 6
Accuracy of Selected Responsibilities $\delta_{j\beta}$

No (j)	Log Claim Value	$\delta_{j\beta}$ Value	Bootstrap Estimate	95% CI
100	0.380	0.000	0.000	(3.90e-12, 2.04e-06)
200	1.031	0.410	0.396	(0.243, 0.531)
300	1.353	0.854	0.863	(0.802 0.912)

5. SUMMARY AND CONCLUSION

This paper has introduced the statistical methodology for inference in the presence of missing data, known as the EM algorithm, into actuarial settings. We have shown that this methodology is particularly appropriate for those practical situations that require consideration of the missing or incomplete data, or hidden factors. We believe that, due to its conceptual simplicity, the EM method could become a standard tool of actuarial analysis in the future. Herein we have given only some examples of its usefulness in modeling loss severity. Specifically, in modeling claim severities in the presence of missing data values, the EM algorithm was used to impute missing values in a straightforward fashion and with accuracy comparable to that of a much more sophisticated method based on Bayesian regression. The EM method also was used to uncover hidden patterns in a different set of claim severity data by means of mixture model fitting and kernel density classification. Actual auto bodily injury liability claims closed in Massachusetts in 2001 were used to illustrate a first-stage screen for abusive medical providers, and their abusive claims. The Monte Carlo-based method of assessing the accuracy of the model parameters fitted via the EM algorithm, known as the parametric bootstrap, was also presented, and the appropriate algorithm for its implementation was developed. The set of functions written in the statistical language R implementing the EM algorithms discussed in the paper have been included in Appendix A to allow readers to try different actuarial situations where missing data and hidden components might be found. A large variety of actuarial and financial applications of the presented methodology are possible, including its incorporation into models of Risk-Based Capital, Value at Risk, and general Dynamic Financial Analysis, or, in general, any area of actuarial decision making where latent traits considerations might help in modeling underlying uncertainty (see, e.g., work by Craiu and Duchesne [2004] for applications in competing risks models). We hope that our paper will promote enough interest in the EM methodology for further exploration of those opportunities.

ACKNOWLEDGMENT

The authors wish to acknowledge Dr. Krzysztof M. Ostaszewski, FSA, for his encouragement and helpful remarks at the inception of this project. They also would like to thank anonymous referees for their insightful comments and several useful references that helped us improve presentation of the material. Finally, the authors are very grateful to the Institute for Mathematics and Its Applications and especially its Director, Dr. Douglas N. Arnold, for their hospitality and support during the final stages of the project.

APPENDIX A

R FUNCTIONS

We present here the implementation of Algorithms 1 and 2 in the statistical software R, which is a freeware version of the award-winning statistical software S+ and is available from <http://www.r-project.org>. The functions below were used in the numerical examples discussed in the text.

#Algorithm 1: EM version of Buck's imputation procedure#####

```

#auxiliary function
inv<-function(m,...) solve(m,diag(rep(1,length(m[,1]))),...);
#defining matrix inverse (for compatibility with older versions of R)
#input parameters
# d-dataframe of two columns containing complete observations
# d1-list of observations with missing second coordinate
# d2-list of observations with missing first coordinate
# B-maximal number of iterations (default value 500)
# eps- convergence criterion (default value .0001)
#####

em.buck <-function(d,d1,d2,B=500,eps=.0001) {
n<-length(d[,1]);
n1<-length(d1);
n2<-length(d2);
m<-apply(d,2,mean);
R<-cov(d);
rho<-cor(d)[1,2]; nLL.old<-eps; nLL.new<-100;
w<-rbind(d,cbind(d1,rep(m[2],n1)),cbind(rep(m[1],n2),d2));
i<-1; #mainloop#
while (abs(nLL.new-nLL.old)/nLL.old>eps && i<=B) {
T1<-sum(w[,1]); T2<-sum(w[,2]); T12<-sum(w[,1]*w[,2]);
T11<-sum(w[(n+n1+1):(n+n1+n2),1]^2
+R[1,1]*(1-rho^2))+sum(w[-((n+n1+1):(n+n1+n2)),1]^2);
T22<-sum(w[(n+1):(n+n1),2]^2+R[2,2]*(1-rho^2))+sum(w[-((n+1):(n+n1)),2]^2);
R<-array(c(T11-T1^2/(n+n1+n2),T12-T1*T2/(n+n1+n2),T12-T1*T2/(n+n1+n2),
T22-T2^2/(n+n1+n2))/(n+n1+n2),c(2,2));
m<-c(T1/(n+n1+n2),T2/(n+n1+n2));
rho<-R[1,2]/sqrt(R[1,1]*R[2,2]);
w[(n+1):(n+n1),2]<-m[2]+R[1,2]*(w[(n+1):(n+n1),1]-m[1])/R[1,1];
w[(n+n1+1):(n+n1+n2),1]<-m[1]+R[1,2]*(w[(n+n1+1):(n+n1+n2),2]-m[2])/
R[2,2];
nLL.old<-nLL.new;
s<-0; for (k in 1:(n+n1+n2)) s<-(w[k,]-m)%*%inv(R)%*%(w[k,]-m)+s;
nLL.new<-2*(n+n1+n2)*log(2*pi)+s+(n+n1+n2)*log(abs(det(R)));
i<-i+1;}; #end mainloop#
print(paste(`n`='', n, n1, n2, `Theta estimates`='', m[1],m[2], R[1,1],
R[1,2], R[2,2], `iter`='',i-1,`-2LL`='',nLL.new, `rho`='',rho))
return(list(m=m, R=R,iter=i-1,LL=nLL.new,w=w));}
#output parameters: list of objects (m,R,iter,LL,w)
# m-vector of estimated means
# R-estimated covariance matrix
# iter -number of iterations until convergence
# w-concatenated dataframe of d,d1,d2 along with imputed missing values

```

Algorithm 2: EM for normal mixtures

```

#auxiliary function
#lsum<-function(a,p.new,m,s){k<-length(m); ss<-0;
for (i in 1:k) ss<-ss+p.new[[i]]*dnorm(a,m[[i]],s[[i]]);

```

```

return(ss)}
# facilitates calculation of LL in the main procedure below

# input parameters:
# a -any list of numeric data
# pi -initial estimate of mixing proportions (default value: uniform over three
# components)
# eps -desired convergence accuracy (default value .0001)
# B -maximal number of iterations allowed (default value 100)
# m -initial values of means estimates (default value: random selection from a)
# number of fitted mixing components is determined from the dimension of the
# vectors
# pi and m
#####
em.multnorm<-function(a, pi=c(1/3,1/3,1/3),eps=.0001,B=100,
  m=sort(sample(a,3)))
{n<-length(a); k<-length(m); s<-rep(sd(a),k);
i<-1; p.new<-pi;
m0<-m;
logl.old<-1;
logl.new<-sum(log(lsum(a,p.new,m,s)));
#mainloop#
while (abs((logl.new-logl.old)/logl.old)>eps && i<=B)
{g<-NULL;
for (t in 1:k) g<-rbind(g, p.new[[t]]*dnorm(a,m[[t]],s[[t]])/
  lsum(a,p.new,m,s));
m<-g%*%a/g%*%rep(1,n);
s<-sqrt(g%*%a^2/g%*%rep(1,n)-m^2);
p.old<-p.new; p.new<-g%*%rep(1,n)/n; i<-i+1;
logl.old<-logl.new;
logl.new<-sum(log(lsum(a,p.new,m,s)));
};
#end mainloop#
print(paste('`Theta estimates`',m,s,``pi=``,p.new,``iter=``,i-1,``-2LL=``,
  -2*logl.new));
return(list(m=m,s=s,pi=p.new,iter=i-1,start=m0,logl=
  -2*logl.new,resp=t(g),data=a))
}

# output parameters: list of objects (m,s,pi,iter,logl,resp)
# m - vector of estimated values of means
# s - vector of estimated values of standard deviations
# pi -vector of estimated value of mixing proportions
# iter- number of iterations until convergence
# logl- final value of -2Q
# resp- matrix of responsibilities (columns correspond to mixture components)

```

APPENDIX B

MASSACHUSETTS AUTO INSURANCE BODILY INJURY LIABILITY DATA

Below we present the set of auto insurance data discussed in the paper. Medical bill claim amounts are given in thousands. Responsibilities δ_{j3} are calculated according to Algorithm 2.

No.	Claimed Amt	Log(Amt)	Provider	Resp. δ_{β}	No.	Claimed Amt	Log(Amt)	Provider	Resp. δ_{β}
1	0.045	-3.101	Other	0.00	2	0.047	-3.058	Other	0.00
3	0.07	-2.659	Other	0.00	4	0.075	-2.590	Other	0.00
5	0.077	-2.564	Other	0.00	6	0.092	-2.386	Other	0.00
7	0.117	-2.146	Other	0.00	8	0.117	-2.146	Other	0.00
9	0.14	-1.966	Other	0.00	10	0.145	-1.931	Other	0.00
11	0.149	-1.904	Other	0.00	12	0.165	-1.802	Other	0.00
13	0.167	-1.790	Other	0.00	14	0.169	-1.778	Other	0.00
15	0.18	-1.715	Other	0.00	16	0.18	-1.715	Other	0.00
17	0.199	-1.614	Other	0.00	18	0.202	-1.599	Other	0.00
19	0.212	-1.551	Other	0.00	20	0.225	-1.492	Other	0.00
21	0.23	-1.470	Other	0.00	22	0.242	-1.419	Other	0.00
23	0.264	-1.332	Other	0.00	24	0.275	-1.291	Other	0.00
25	0.285	-1.255	Other	0.00	26	0.29	-1.238	Other	0.00
27	0.363	-1.013	Other	0.00	28	0.384	-0.957	Other	0.00
29	0.4	-0.916	Other	0.00	30	0.4	-0.916	Other	0.00
31	0.413	-0.884	Other	0.00	32	0.414	-0.882	Other	0.00
33	0.416	-0.877	Other	0.00	34	0.425	-0.856	Other	0.00
35	0.425	-0.856	Other	0.00	36	0.43	-0.844	Other	0.00
37	0.43	-0.844	A	0.00	38	0.431	-0.842	Other	0.00
39	0.45	-0.799	Other	0.00	40	0.46	-0.777	Other	0.00
41	0.486	-0.722	Other	0.00	42	0.5	-0.693	Other	0.00
43	0.5	-0.693	Other	0.00	44	0.514	-0.666	A	0.00
45	0.531	-0.633	Other	0.00	46	0.54	-0.616	Other	0.00
47	0.556	-0.587	Other	0.00	48	0.564	-0.573	Other	0.00
49	0.6	-0.511	Other	0.00	50	0.605	-0.503	Other	0.00
51	0.605	-0.503	Other	0.00	52	0.65	-0.431	Other	0.00
53	0.66	-0.416	Other	0.00	54	0.66	-0.416	Other	0.00
55	0.685	-0.378	Other	0.00	56	0.69	-0.371	Other	0.00
57	0.698	-0.360	Other	0.00	58	0.7	-0.357	Other	0.00
59	0.705	-0.350	Other	0.00	60	0.725	-0.322	Other	0.00
61	0.74	-0.301	Other	0.00	62	0.75	-0.288	Other	0.00
63	0.78	-0.248	Other	0.00	64	0.785	-0.242	Other	0.00
65	0.785	-0.242	Other	0.00	66	0.806	-0.216	Other	0.00
67	0.825	-0.192	Other	0.00	68	0.825	-0.192	Other	0.00
69	0.83	-0.186	Other	0.00	70	0.836	-0.179	A	0.00
71	0.87	-0.139	Other	0.00	72	0.9	-0.105	Other	0.00
73	0.934	-0.068	Other	0.00	74	0.95	-0.051	Other	0.00
75	0.954	-0.047	Other	0.00	76	0.956	-0.045	Other	0.00
77	0.962	-0.039	Other	0.00	78	0.97	-0.030	Other	0.00
79	0.975	-0.025	Other	0.00	80	0.988	-0.012	Other	0.00
81	1.015	0.015	Other	0.00	82	1.053	0.052	Other	0.00
83	1.058	0.056	Other	0.00	84	1.08	0.077	Other	0.00
85	1.161	0.149	Other	0.00	86	1.167	0.154	Other	0.00
87	1.195	0.178	Other	0.00	88	1.215	0.195	Other	0.00
89	1.242	0.217	Other	0.00	90	1.26	0.231	Other	0.00
91	1.295	0.259	Other	0.00	92	1.31	0.270	Other	0.00
93	1.319	0.277	Other	0.00	94	1.33	0.285	Other	0.00
95	1.34	0.293	Other	0.00	96	1.355	0.304	Other	0.00
97	1.39	0.329	Other	0.00	98	1.444	0.367	Other	0.00
99	1.455	0.375	Other	0.00	100	1.463	0.380	Other	0.00
101	1.49	0.399	Other	0.00	102	1.5	0.405	Other	0.00
103	1.542	0.433	Other	0.00	104	1.598	0.469	Other	0.00
105	1.616	0.480	Other	0.00	106	1.623	0.484	Other	0.00
107	1.64	0.495	Other	0.00	108	1.645	0.498	A	0.00
109	1.65	0.501	Other	0.00	110	1.66	0.507	Other	0.00
111	1.68	0.519	Other	0.00	112	1.695	0.528	Other	0.00
113	1.7	0.531	A	0.00	114	1.758	0.564	Other	0.00
115	1.759	0.565	Other	0.00	116	1.76	0.565	Other	0.00
117	1.896	0.640	Other	0.00	118	1.92	0.652	Other	0.00
119	1.923	0.654	Other	0.00	120	1.941	0.663	A	0.00
121	1.96	0.673	Other	0.00	122	1.972	0.679	Other	0.00
123	1.99	0.688	Other	0.00	124	2.005	0.696	Other	0.00
125	2.018	0.702	Other	0.00	126	2.02	0.703	Other	0.00
127	2.02	0.703	Other	0.00	128	2.03	0.708	Other	0.00
129	2.042	0.714	A	0.00	130	2.062	0.724	Other	0.00
131	2.063	0.724	Other	0.00	132	2.08	0.732	Other	0.00
133	2.087	0.736	Other	0.00	134	2.089	0.737	Other	0.00
135	2.1	0.742	Other	0.00	136	2.115	0.749	Other	0.01

No.	Claimed Amt	Log(Amt)	Provider	Resp. δ_B	No.	Claimed Amt	Log(Amt)	Provider	Resp. δ_B
137	2.12	0.751	Other	0.01	138	2.155	0.768	Other	0.01
139	2.159	0.770	Other	0.01	140	2.161	0.771	Other	0.01
141	2.184	0.781	A	0.01	142	2.188	0.783	Other	0.01
143	2.191	0.784	Other	0.01	144	2.196	0.787	Other	0.01
145	2.224	0.799	Other	0.02	146	2.237	0.805	Other	0.02
147	2.251	0.811	A	0.02	148	2.253	0.812	Other	0.02
149	2.288	0.828	Other	0.03	150	2.295	0.831	Other	0.03
151	2.318	0.841	Other	0.03	152	2.325	0.844	Other	0.03
153	2.325	0.844	A	0.03	154	2.335	0.848	Other	0.04
155	2.341	0.851	Other	0.04	156	2.35	0.854	Other	0.04
157	2.374	0.865	Other	0.05	158	2.39	0.871	Other	0.05
159	2.406	0.878	Other	0.06	160	2.434	0.890	Other	0.07
161	2.45	0.896	Other	0.08	162	2.453	0.897	A	0.08
163	2.468	0.903	Other	0.09	164	2.468	0.903	A	0.09
165	2.48	0.908	Other	0.10	166	2.48	0.908	A	0.10
167	2.49	0.912	A	0.10	168	2.498	0.915	Other	0.11
169	2.5	0.916	Other	0.11	170	2.5	0.916	Other	0.11
171	2.5	0.916	A	0.11	172	2.51	0.920	Other	0.11
173	2.532	0.929	Other	0.13	174	2.54	0.932	Other	0.14
175	2.543	0.933	Other	0.14	176	2.559	0.940	Other	0.15
177	2.572	0.945	Other	0.16	178	2.593	0.953	Other	0.18
179	2.601	0.956	Other	0.19	180	2.616	0.962	Other	0.20
181	2.619	0.963	Other	0.20	182	2.63	0.967	Other	0.21
183	2.635	0.969	Other	0.22	184	2.635	0.969	Other	0.22
185	2.653	0.976	Other	0.24	186	2.655	0.976	Other	0.24
187	2.675	0.984	Other	0.26	188	2.679	0.985	Other	0.26
189	2.697	0.992	Other	0.28	190	2.718	1.000	Other	0.31
191	2.73	1.004	Other	0.32	192	2.734	1.006	Other	0.32
193	2.755	1.013	Other	0.35	194	2.758	1.015	Other	0.35
195	2.773	1.020	Other	0.37	196	2.775	1.021	Other	0.37
197	2.78	1.022	Other	0.38	198	2.785	1.024	A	0.38
199	2.795	1.028	Other	0.40	200	2.805	1.031	Other	0.41
201	2.805	1.031	Other	0.41	202	2.808	1.032	A	0.41
203	2.88	1.058	Other	0.49	204	2.881	1.058	Other	0.50
205	2.881	1.058	A	0.50	206	2.924	1.073	A	0.54
207	2.93	1.075	Other	0.55	208	2.934	1.076	A	0.55
209	2.94	1.078	Other	0.56	210	2.972	1.089	Other	0.59
211	2.975	1.090	Other	0.59	212	3	1.099	Other	0.62
213	3	1.099	A	0.62	214	3.025	1.107	Other	0.64
215	3.058	1.118	Other	0.67	216	3.082	1.126	A	0.68
217	3.085	1.127	Other	0.69	218	3.095	1.130	Other	0.69
219	3.1	1.131	Other	0.70	220	3.102	1.132	A	0.70
221	3.106	1.133	Other	0.70	222	3.135	1.143	Other	0.72
223	3.17	1.154	Other	0.74	224	3.187	1.159	Other	0.75
225	3.192	1.161	A	0.75	226	3.193	1.161	Other	0.75
227	3.2	1.163	Other	0.76	228	3.21	1.166	Other	0.76
229	3.23	1.172	Other	0.77	230	3.23	1.172	Other	0.77
231	3.23	1.172	A	0.77	232	3.232	1.173	Other	0.77
233	3.235	1.174	Other	0.78	234	3.243	1.176	A	0.78
235	3.248	1.178	A	0.78	236	3.249	1.178	Other	0.78
237	3.26	1.182	Other	0.79	238	3.261	1.182	Other	0.79
239	3.272	1.185	A	0.79	240	3.29	1.191	Other	0.80
241	3.295	1.192	Other	0.80	242	3.304	1.195	Other	0.80
243	3.332	1.204	A	0.81	244	3.333	1.204	Other	0.81
245	3.338	1.205	Other	0.81	246	3.34	1.206	Other	0.82
247	3.341	1.206	A	0.82	248	3.349	1.209	A	0.82
249	3.349	1.209	A	0.82	250	3.349	1.209	A	0.82
251	3.353	1.210	A	0.82	252	3.36	1.212	Other	0.82
253	3.378	1.217	A	0.83	254	3.385	1.219	A	0.83
255	3.387	1.220	A	0.83	256	3.416	1.228	Other	0.84
257	3.429	1.232	A	0.84	258	3.438	1.235	A	0.84
259	3.444	1.237	A	0.84	260	3.469	1.244	A	0.85
261	3.473	1.245	A	0.85	262	3.473	1.245	A	0.85
263	3.475	1.246	A	0.85	264	3.477	1.246	A	0.85
265	3.505	1.254	Other	0.85	266	3.517	1.258	A	0.85
267	3.518	1.258	Other	0.85	268	3.527	1.260	A	0.85
269	3.535	1.263	A	0.86	270	3.547	1.266	A	0.86
271	3.55	1.267	Other	0.86	272	3.552	1.268	Other	0.86

No.	Claimed Amt	Log(Amt)	Provider	Resp. δ_{β}	No.	Claimed Amt	Log(Amt)	Provider	Resp. δ_{β}
273	3.567	1.272	A	0.86	274	3.57	1.273	Other	0.86
275	3.575	1.274	Other	0.86	276	3.58	1.275	Other	0.86
277	3.583	1.276	A	0.86	278	3.59	1.278	A	0.86
279	3.603	1.282	A	0.86	280	3.615	1.285	A	0.86
281	3.623	1.287	A	0.86	282	3.647	1.294	A	0.86
283	3.655	1.296	Other	0.86	284	3.655	1.296	A	0.86
285	3.658	1.297	Other	0.87	286	3.675	1.302	Other	0.87
287	3.675	1.302	Other	0.87	288	3.687	1.305	A	0.87
289	3.72	1.314	Other	0.87	290	3.72	1.314	Other	0.87
291	3.742	1.320	Other	0.87	292	3.757	1.324	A	0.87
293	3.765	1.326	Other	0.87	294	3.8	1.335	A	0.87
295	3.809	1.337	Other	0.86	296	3.848	1.348	A	0.86
297	3.857	1.350	A	0.86	298	3.867	1.352	Other	0.86
299	3.867	1.352	A	0.86	300	3.87	1.353	Other	0.86
301	3.883	1.357	Other	0.86	302	3.89	1.358	Other	0.86
303	3.905	1.362	A	0.86	304	3.907	1.363	A	0.86
305	4	1.386	Other	0.85	306	4.011	1.389	Other	0.85
307	4.039	1.396	A	0.84	308	4.065	1.402	A	0.84
309	4.095	1.410	Other	0.83	310	4.134	1.419	Other	0.82
311	4.147	1.422	Other	0.82	312	4.155	1.424	A	0.82
313	4.17	1.428	Other	0.81	314	4.179	1.430	A	0.81
315	4.2	1.435	Other	0.81	316	4.215	1.439	Other	0.80
317	4.257	1.449	A	0.79	318	4.3	1.459	Other	0.78
319	4.489	1.502	A	0.70	320	4.593	1.525	A	0.64
321	4.595	1.525	Other	0.64	322	4.63	1.533	A	0.62
323	4.653	1.538	Other	0.60	324	4.7	1.548	A	0.57
325	4.731	1.554	Other	0.55	326	4.741	1.556	A	0.55
327	4.75	1.558	Other	0.54	328	4.761	1.560	Other	0.53
329	4.81	1.571	Other	0.50	330	5.072	1.624	Other	0.31
331	5.161	1.641	Other	0.25	332	5.24	1.656	Other	0.20
333	5.64	1.730	Other	0.06	334	5.779	1.754	Other	0.03
335	6.166	1.819	Other	0.01	336	6.406	1.857	Other	0.00
337	6.725	1.906	Other	0.00	338	7.717	2.043	A	0.00
339	8	2.079	Other	0.00	340	9.5	2.251	Other	0.00
341	10.295	2.332	Other	0.00	342	12.533	2.528	Other	0.00
343	12.688	2.541	Other	0.00	344	16.043	2.775	Other	0.00
345	18.847	2.936	Other	0.00	346	19.5	2.970	Other	0.00
347	20.827	3.036	Other	0.00	348	50	3.912	Other	0.00

REFERENCES

ABRAHAMSE, ALAN F., AND STEPHAN J. CARROLL. 1999. The Frequency of Excess Claims for Automobile Personal Injuries. In *Automobile Insurance: Road Safety, New Drivers, Risks, Insurance Fraud and Regulation*, edited by Georges Dionne and Claire Laberge-Nadeau, pp. 131–50. Dordrecht: Kluwer.

BAILEY, ROBERT A., AND LEROY J. SIMON. 1960. Two Studies in Automobile Insurance Ratemaking. *ASTIN Bulletin* 1(4): 192–217.

BILMES, JEFF A. 1998. *Gentle Tutorial of the EM Algorithm and Its Application to Parameter Estimation for Gaussian Mixture and Hidden Markov Models*. International Computer Science Institute, University of California, Berkeley.

BOLANCÉ, CATALINA, MONTSERRAT GUILLEN, AND JENS P. NIELSEN. 2003. Kernel Density Estimation of Actuarial Loss Functions. *Insurance: Mathematics and Economics* 32(1): 19–36.

BROCKETT, PATRICK L., RICHARD A. DERRIG, LINDA L. GOLDEN, ARNOLD LEVINE, AND MARK ALPERT. 2002. Fraud Classification Using Principal Component Analysis of RIDITs. *Journal of Risk and Insurance* 69(3): 341–71.

BROCKETT, PATRICK L., XIAOHUA XIA, AND RICHARD A. DERRIG. 1998. Using Kohonen’s Self-organizing Feature Map to Uncover Automobile Bodily Injury Claims Fraud. *Journal of Risk and Insurance* 65(2): 245–74.

CRAIU, RADU V., AND THIERRY DUCHESNE. 2004. Inference Based on the EM Algorithm for the Competing Risk Model with Masked Causes of Failure. *Biometrika*, forthcoming.

CROCKER, KEITH J., AND SHARON TENNYSON. 1999. Costly State Falsification or Verification? Theory and Evidence from Bodily Injury Liability Claims. In *Automobile Insurance: Road Safety, New Drivers, Risks, Insurance Fraud and Regulation*, edited by Georges Dionne and Claire Laberge-Nadeau, pp. 119–31. Dordrecht: Kluwer.

CUMMINS, J. DAVID, AND LAUREL J. WILTBANK. 1983. Estimating the Total Claims Distribution Using Multivariate Frequency and Severity Distribution. *Journal of Risk and Insurance* 50: 377–403.

———. 1985. A Multivariate Model of the Total Claims Process. *ASTIN Bulletin* 14(1): 45–52.

- CUMMINS, J. DAVID, AND SHARON TENNYSON. 1996. Controlling Automobile Insurance Costs. *Journal of Economic Perspectives* 6(2): 95–115.
- D'ARCY, STEPHEN P., RICHARD W. GORVETT, J. A. HERBERS, AND THOMAS E. HETTINGER. 1997. Building a Dynamic Financial Analysis Model that Flies. *Contingencies* (November/December): 40–45.
- DEMPSTER, ALLAN P., N. M. LAIRD, AND DONALD B. RUBIN. 1977. Maximum Likelihood from Incomplete Data Using EM Algorithm. *Journal of Royal Statistical Society Series B* 39(1): 1–38.
- DERRIG, RICHARD A. 2002. Insurance Fraud. *Journal of Risk and Insurance* 69(3): 271–87.
- DERRIG, RICHARD A., KRZYSZTOF M. OSTASZEWSKI, AND GRZEGORZ A. REMPALA. 2001. Applications of Resampling Methods in Actuarial Practice. *Proceedings of the Casualty Actuarial Society* 87: 322–64.
- DERRIG, RICHARD A., HERBERT I. WEISBERG, AND XIU CHEN. 1994. Behavioral Factors and Lotteries under No-Fault with a Monetary Threshold: A Study of Massachusetts Automobile Claims. *Journal of Risk and Insurance* 61(2): 245–75.
- DUGAS, CHARLES D., YOSHUA BENGIO, NICOLAS CHAPADOS, PASCAL VINCENT, GERMAIN DENONCOURT, AND CHRISTIAN FOURNIER. 2003. Statistical Learning Algorithms Applied to Automobile Insurance Ratemaking. In *Intelligent Techniques for the Insurance Industry: Theory and Applications*, edited by L. Jain and A. F. Shapiro, pp. 137–97. Singapore: World Scientific.
- EMBRECHTS, PAUL, ALEXANDER MCNEIL, AND DANIEL STRAUMANN. 2000. Correlation and Dependence in Risk Management: Properties and Pitfalls. In *Extremes and Integrated Risk Management*, edited by Paul Embrechts, pp. 71–76. London: Risk Books.
- HAGENAARS, JACQUES, AND ALLAN C. MCCUTCHEON, eds. 2001. *Applied Latent Class Analysis*. Cambridge: Cambridge University Press.
- HASTIE, TREVOR, ROBERT TIBSHIRANI, AND JEROME FRIEDMAN. 2001. *The Elements of Statistical Learning*. New York: Springer.
- HORTON, NICHOLAS J., AND STUART R. LIPSITZ. 2001. Multiple Imputation in Practice: Comparison of Software Packages for Regression Models with Missing Variables. *American Statistician* 55(3): 244–54.
- KALB, GUYONNE R. J., PAUL KOFMAN, AND TON C. F. VORST. 1996. Mixtures of Tails in Clustered Automobile Collision Claims. *Insurance: Mathematics & Economics* 18: 89–107.
- LITTLE, RODERICK J. A., AND DONALD B. RUBIN. 1987. *Statistical Analysis with Missing Data*. New York: John Wiley & Sons.
- MAJOR, JOHN A., AND DAN R. RIEDINGER. 1992. EFD: A Hybrid Knowledge/Statistical-Based System for the Detection of Fraud. *International Journal of Intelligent Systems* 7: 687–703 (reprinted in *Journal of Risk and Insurance* [2002] 69(3): 309–24).
- MCCUTCHEON, ALLAN C. 1987. *Latent Class Analysis*. Beverly Hills: Sage Publications.
- MCLACHLAN, GEOFFREY, AND THIRIYAMBAKAM KRISHNAN. 1997. *The EM Algorithm and Extensions*. New York: Wiley.
- MCLACHLAN, GEOFFREY, AND DAVID PEEL. 2000. *Finite Mixture Models*. New York: Wiley-Interscience.
- OWEN, ART B. 2001. *Empirical Likelihood*. New York: Chapman and Hall.
- REMPALA, GRZEGORZ A., AND KONRAD SZATZSCHNEIDER. 2004. Bootstrapping Parametric Models of Mortality. *Scandinavian Actuarial Journal* (1): 53–78.
- RUBIN, DONALD B. 1987. *Multiple Imputation for Nonresponse in Surveys*. New York: John Wiley.
- . 1996. Multiple Imputation after 18+ Years. *Journal of the American Statistical Association* 91: 473–89.
- SCHAFFER, JOE L. 1997. *Analysis of Incomplete Multivariate Data*. London: Chapman & Hall.
- WANG, SHAUN. 1998. Aggregation of Correlated Risk Portfolios: Models and Algorithms. *Proceedings of the Casualty Actuarial Society* 83: 848–939.
- WU, JEFF A. 1983. On the Convergence Properties of the EM Algorithm. *Annals of Statistics* 11(1): 95–103.
- XU, LI, AND MIKE I. JORDAN. 1996. On Convergence Properties of the EM Algorithm for Gaussian Mixtures. *Neural Computation* 8: 129–51.

Discussions on this paper can be submitted until October 1, 2005. The authors reserve the right to reply to any discussion. Please see the Submission Guidelines for Authors on the inside back cover for instructions on the submission of discussions.