Stat 446
Advanced Short-Term Actuarial Modeling

Chapter 5: Credibility

Outline

1. Introduction and Review

2. Bayesian Credibility
   - Bayesian Estimation and Bayesian Premium
   - Bayesian Credibility: Conjugate Prior

3. Bühlmann Credibility
   - Bühlmann Model
   - Bühlmann-Straub Model

4. Empirical Bayes Parameter Estimation
   - Nonparametric Estimation
   - Semiparametric Estimation

Introduction

- We have learned the limited fluctuation credibility theory to combine a policyholder's past experience with related manual rates for ratemaking.
- While this approach is effective in practice, the choices of full credibility standard are arguably subjective.
- We further introduce a model-based credibility method, namely the Bühlmann credibility or greatest accuracy credibility, derived from the Bayesian credibility theory.
- We will also introduce non/semi-parametric approaches to deal with cases where the conditional loss density and the prior distribution are left largely unspecified.

Credibility Question

- The insurer has to figure out *how credible* the policyholder's own experience is.
  - The more past policyholder information the insurer has, the more credible its own experience.
  - Competitive considerations may force the insurer to give full (using the policyholder's own experience only and not the manual rate) or nearly **full credibility** to a given policyholder in order to retain the business.
- Another application: if we lack past data, how do we combine the limited past experience with other related manual rates?

Credibility Setup

- Suppose that a policyholder has experienced X_j claims/losses in past experience period $j \in \{1, 2, \ldots, n\}$.
- We may also consider X_j to be the loss from the jth policy in a group.
- If $\mathbb{E}[X_j] = \xi$ and $\mathrm{Var}(X_j) = \sigma^2$, the sample average satisfies

$$\mathbb{E}\left[\bar{X}\right] = \xi, \qquad \mathrm{Var}\left(\bar{X}\right) = \frac{\sigma^2}{n}$$

- Given a manual premium M, how do we decide the value of ξ?
  - **No credibility**: ignore past data and simply charge M.
  - **Full credibility**: ignore M and charge $\bar{X}$.
  - **Partial credibility**: some combination of M and $\bar{X}$.

Limited Fluctuation Credibility

- We "lean toward" the choice of $\bar{X}$ if the experience is stable.
- To quantify the stability: select $r > 0$ and $0 < p < 1$ such that

$$\Pr\left(-r\xi \leq \bar{X} - \xi \leq r\xi\right) \geq p.$$

Full credibility is assigned if the above condition is met.

  - r close to 0 and p close to 1 (commonly $r = 0.05$ and $p = 0.9$)

Limited Fluctuation Credibility (cont.)

- Let $\lambda_0 = \left(z_{(p+1)/2}/r\right)^2$. The condition for full credibility is

$$\frac{r\xi}{\sigma/\sqrt{n}} \geq z_{(p+1)/2}, \qquad \text{Var}\left(\bar{X}\right) = \frac{\sigma^2}{n} \leq \frac{\xi^2}{\lambda_0}$$

- Or, the number of exposure units required for full credibility is

$$n \geq \lambda_0 \left(\frac{\sigma}{\xi}\right)^2$$

Partial Credibility

- If full credibility is inappropriate, it may be desirable to use both the observed experience $\bar{X}$ and the external manual premium M.
- An intuitively appealing approach is the weighted average:

$$P_c = Z\bar{X} + (1 - Z)M, \quad Z \in [0, 1]$$

- Choices of the credibility factor Z
  - $Z = n/(n + k)$, where k is to be determined.
  - **Square root rule**: $Z = \min\left\{\sqrt{n/n_0}, 1\right\}$, where n_0 is the standard for full credibility, i.e.,

$$\frac{\xi^2}{\lambda_0} = \mathrm{Var}(P_c) = Z^2\frac{\sigma^2}{n} \quad \Rightarrow \quad Z = \min\left\{\frac{\xi}{\sigma}\sqrt{\frac{n}{\lambda_0}}, 1\right\}$$

Bayesian Methodology

- Recall that the concern is whether a given policyholder is homogeneous to others in the class.
- The **Bayesian approach** assumes that the risk level of each policyholder is characterized by a risk parameter θ, which varies by policyholder and has **prior distribution** $\pi(\theta)$.
- Given that all observable underwriting characteristics have already been used, θ may be viewed as representative of the residual, unobserved factors that affect the risk level.
- The losses X have a conditional distribution on θ, $f_{X|\Theta}(x|\theta)$.
- Naturally, we can revise our inference of $\pi(\theta)$ based on past data $\mathbf{X} = (X_1, X_2, \cdots, X_n)^\top$, and further improve our prediction of X_{n+1}.

Posterior Distribution

- For a sample $\mathbf{X} = (X_1, X_2, \cdots, X_n)^\top$ with prior $\pi(\theta)$, the joint density of $\mathbf{X}$ and Θ is

$$f_{\mathbf{X},\Theta}(\mathbf{x}, \theta) = f_{\mathbf{X}|\Theta}(\mathbf{x}|\theta)\pi(\theta)$$

  and the marginal density of $\mathbf{X}$ is given by

$$f_{\mathbf{X}}(\mathbf{x}) = \int f_{\mathbf{X}|\Theta}(\mathbf{x}|\theta)\pi(\theta)\, d\theta$$

- Therefore, given past claim data $\mathbf{X} = \mathbf{x}$, the **posterior distribution** of Θ is

$$\pi_{\Theta|\mathbf{X}}(\theta|\mathbf{x}) = \frac{f_{\mathbf{X}|\Theta}(\mathbf{x}|\theta)\pi(\theta)}{\int f_{\mathbf{X}|\Theta}(\mathbf{x}|\theta)\pi(\theta)\, d\theta}$$

Posterior Distribution (cont.)

- A **Bayes estimate** of θ is the value $\hat{\theta}$ that *minimizes the expected posterior loss* $\mathbb{E}\left[I(\hat{\theta}, \Theta) \,|\, \mathbf{X} = \mathbf{x}\right]$. It depends on the loss function:

Loss function	$I(\hat{\theta}, \theta)$	Bayes estimate
Squared error	$(\hat{\theta} - \theta)^2$	posterior mean
Absolute	$\lvert\hat{\theta} - \theta\rvert$	posterior median
Zero-one	0 if $\hat{\theta} = \theta$, 1 otherwise	posterior mode

- Bayesian estimation applies to frequency parameters (e.g., a Poisson mean), severity parameters (e.g., an exponential mean), and functions of them, such as the expected aggregate loss $\mathbb{E}[S \,|\, \Theta] = \mathbb{E}[N \,|\, \Theta]\,\mathbb{E}[X \,|\, \Theta]$.

Credibility Intervals

- A $100(1-\alpha)\%$ **credibility interval** for Θ is an interval (a, b) with

$$\Pr(a \leq \Theta \leq b \,|\, \mathbf{X} = \mathbf{x}) \geq 1 - \alpha.$$

  Unlike a confidence interval, this is a direct probability statement about the parameter.

- **Equal-tailed** interval: a and b are the $\alpha/2$ and $1-\alpha/2$ percentiles of the posterior distribution.

- **Highest posterior density (HPD)** interval: the set $\{\theta : \ \pi(\theta \,|\, \mathbf{x}) \geq k\}$, with k chosen to give probability $1-\alpha$. It is the shortest credibility interval; for a unimodal posterior, $\pi(a \,|\, \mathbf{x}) = \pi(b \,|\, \mathbf{x})$.

- **Bayesian central limit theorem**: for large samples the posterior is approximately normal, so $\mathbb{E}[\Theta \,|\, \mathbf{x}] \pm z_{1-\alpha/2}\sqrt{\mathrm{Var}(\Theta \,|\, \mathbf{x})}$ is an approximate credibility interval.

Predictive Distribution and Bayesian Premium

- The **predictive distribution** of the future loss X_{n+1} is

$$f_{X_{n+1}|\mathbf{X}}(x_{n+1}|\mathbf{x}) = \int f_{X_{n+1}|\Theta}(x_{n+1}|\theta)\pi_{\Theta|\mathbf{X}}(\theta|\mathbf{x})\,d\theta$$

- Without past data, we predict X_{n+1} using its mean

$$\mu_{n+1} = \mathbb{E}[X_{n+1}] = \mathbb{E}[\mu_{n+1}(\Theta)] = \int \mu_{n+1}(\theta)\pi_{\Theta}(\theta)\,d\theta,$$

  where $\mu_{n+1}(\theta) = \mathbb{E}[X_{n+1}|\Theta = \theta]$.

- Incorporating the data, the **Bayesian premium** is

$$\mathbb{E}[X_{n+1}|\mathbf{X} = \mathbf{x}] = \int \mu_{n+1}(\theta)\pi_{\Theta|\mathbf{X}}(\theta|\mathbf{x})\,d\theta$$

The Posterior Kernel

- The **kernel** of the posterior plays an important role in determining the distribution:

$$\pi_{\Theta|\mathbf{X}}(\theta|\mathbf{x}) \propto f_{\mathbf{X}|\Theta}(\mathbf{x}|\theta)\pi(\theta) = \left[\prod_{j=1}^{n} f_{X_j|\Theta}(x_j|\theta)\right]\pi(\theta)$$

- Keep only the factors that involve θ. If the kernel matches the form of a known density, the posterior is that distribution, and its parameters can be read off directly.
- Otherwise, standardize the kernel (divide by its integral over θ) to find the proper posterior density.

Example: STAM Sample Q226

Example

For a particular policy, the conditional probability of the annual number of claims given $\Theta = \theta$, and the probability distribution of Θ are as follows:

Number of Claims	0	1	2
Probability	2θ	θ	$1 - 3\theta$

θ	0.10	0.30
Probability	0.80	0.20

One claim was observed in Year 1.
Calculate the Bayesian estimate of the expected number of claims for Year 2. [1.0714]

(Exercises: STAM Sample Qs 11, 203 & 260)

Example: STAM Sample Q247

Example

An insurance company sells three types of policies with the following characteristics:

Type of Policy	Proportion of Total Policies	Annual Claim Frequency
I	5%	Poisson with mean 0.25
II	20%	Poisson with mean 0.50
III	75%	Poisson with mean 1.00

A randomly selected policyholder is observed to have a total of one claim for Year 1 through Year 4.
For the same policyholder, calculate the Bayesian estimate of the expected number of claims in Year 5. [0.6794]

(Exercises: STAM Sample Qs 29, 55, 60, 136, 170 & ASTAM S23 Q3(d), S24 Q5(e))

Example: STAM Sample Qs 64 & 45

Example

For a group of insureds, you are given:

(i) The amount of a claim is uniformly distributed but will not exceed a certain unknown limit θ.

(ii) The prior distribution of θ is $\pi(\theta) = \frac{500}{\theta^2}$, $\theta > 500$.

(iii) Two independent claims of $x_1 = 400$ and $x_2 = 600$ are observed.

1. Calculate the probability that the next claim X_3 will exceed 550. [0.3125]
2. Calculate the Bayesian premium, $\mathbb{E}[X_3|x_1, x_2]$. [450]

Example: STAM Sample Q76

Example

You are given:

(i) The annual number of claims for each policyholder follows a Poisson distribution with mean θ.

(ii) The distribution of θ across all policyholders has probability density function:

$$\pi(\theta) = \theta e^{-\theta}, \quad \theta > 0$$

(iii) $\int_0^\infty \theta e^{-n\theta}\, d\theta = \frac{1}{n^2}$

A randomly selected policyholder is known to have had at least one claim last year.
Calculate the posterior probability that this same policyholder will have at least one claim this year. [0.8148]

(Exercises: STAM Sample Qs 15, 43, 142, 157, 184 & 242)

Conjugate Prior Distributions

- Recall that the posterior is proportional to its kernel, $f_{\mathbf{X}|\Theta}(\mathbf{x}|\theta)\pi(\theta)$.
- For some combinations of model and prior, the kernel has the same form as the prior, so the posterior is in the same family with updated parameters.

Conjugate Prior Distributions

A prior distribution is a **conjugate prior** for a given model if the resulting posterior distribution is from the same family as the prior (closely related to the **linear exponential family**, such as normal, gamma, Poisson, binomial, negative binomial).

Bayesian Credibility: Poisson/Gamma

- Conjugate prior example: the number of claims X follows a Poisson distribution with mean θ, which has a $\text{Gamma}(\alpha, 1/\beta)$ prior.
- The posterior distribution is a new gamma with $\alpha_* = \sum x_j + \alpha$ and $\beta_* = n + \beta$.

$$\pi_{\Theta|\mathbf{X}}(\theta|\mathbf{x}) \propto \left[\prod_{j=1}^{n} \theta^{x_j} e^{-\theta}\right] \theta^{\alpha-1} e^{-\beta\theta} \propto \theta^{\sum x_j+\alpha-1} e^{-(n+\beta)\theta}$$

- The Bayesian premium is given by

$$\mathbb{E}[X_{n+1}|\mathbf{X} = \mathbf{x}] = \frac{\alpha_*}{\beta_*} = \frac{n}{n+\beta}\bar{x} + \frac{\beta}{n+\beta}\frac{\alpha}{\beta}$$

- The Bayesian premium is a weighted average of $\bar{x}$ and the unconditional mean $\mu = \alpha/\beta$, with the credibility factor $Z = \frac{n}{n+\beta}$.

Example: STAM Sample Q58

Example

You are given:

(i) The number of claims per auto insured follows a Poisson distribution with mean λ.

(ii) The prior distribution for λ has the following probability density function:

$$f(\lambda) = \frac{(500\lambda)^{50} e^{-500\lambda}}{\lambda \Gamma(50)}, \quad \lambda > 0$$

(iii) A company observes the following claims experience:

	Year 1	Year 2
Number of claims	75	210
Number of autos insured	600	900

The company expects to insure 1100 autos in Year 3. Calculate the Bayesian expected number of claims in Year 3. [184.25]

(Exercises: STAM Sample Qs 191 & 254)

Example: Bayes Estimates and Credibility Intervals

Example

The annual number of claims for a policyholder is Poisson with mean λ. The prior distribution of λ is gamma with $\alpha = 3$ and $\theta = 0.1$. The policyholder has 2 claims in 4 years. If Λ is gamma with integer α, then $2\Lambda/\theta$ is chi-square with 2α degrees of freedom; the 2.5th, 50th, and 97.5th percentiles of the chi-square distribution with 10 degrees of freedom are 3.247, 9.342, and 20.483.

(a) Determine the posterior distribution of λ. [gamma, $\alpha^* = 5$, $\theta^* = 1/14$]

(b) Calculate the Bayes estimate of λ under squared-error, absolute, and zero-one loss. [0.3571, 0.3336, 0.2857]

(c) Calculate the 95% equal-tailed credibility interval for λ. [(0.1160, 0.7315)]

(d) Calculate an approximate 95% credibility interval using the Bayesian central limit theorem. [(0.0441, 0.6702)]

Bayesian Credibility: Bernoulli/Beta

- Conjugate prior example: each trial results in a claim with probability q (Bernoulli), and Q has a $\mathrm{Beta}(a, b)$ prior.
- If we observe n Bernoulli trials with k claims, the posterior is a new beta with $a_* = a + k$ and $b_* = b + n - k$.

$$\begin{aligned}\pi_{Q|\mathbf{x}}(q|\mathbf{x}) &\propto q^k(1-q)^{n-k} \cdot q^{a-1}(1-q)^{b-1} \\ &\propto q^{a+k-1}(1-q)^{b+n-k-1}\end{aligned}$$

Bayesian Credibility: Bernoulli/Beta (cont.)

- If the next period has m trials, so that $X_{n+1} \mid Q = q \sim \text{Binomial}(m, q)$, the Bayesian premium is

$$\begin{aligned}\mathbb{E}[X_{n+1}|\mathbf{X} = \mathbf{x}] &= m\,\mathbb{E}[Q|\mathbf{X} = \mathbf{x}] = m\frac{a_*}{a_* + b_*} \\ &= m\left[\frac{n}{n+a+b}\left(\frac{k}{n}\right) + \frac{a+b}{n+a+b}\frac{a}{a+b}\right]\end{aligned}$$

- This is a credibility-weighted average of the observed claim rate k/n and the prior mean $a/(a+b)$, with credibility factor $Z = \frac{n}{n+a+b}$.
- A similar idea can be applied to the negative binomial/beta case.

Example: ASTAM Fall 2023 Q4(b)

Example

For a portfolio of m independent policies, the number of claims in a year for the ith policy, denoted I_i, is a Bernoulli random variable with parameter q. Let n denote the total observed number of claims from the portfolio in one year so that $n = \sum_{i=1}^{m} I_i$.
An actuary decides to use a Bayesian approach to model the heterogeneity of the policies. They decide to use the Beta$(2, \beta)$ distribution, with the following density function, as the prior distribution for q:

$$\pi(q) = \beta(\beta+1)q(1-q)^{\beta-1}, \quad 0 \leq q \leq 1, \ \beta > 0$$

(b) Show that the posterior distribution of q is Beta$(n+2, m+\beta-n)$.

Example: ASTAM Fall 2023 Q4(c)(d)

Example

The Bayes estimate of q under the squared error loss function can be expressed in credibility form as $Z\times$ sample mean $+(1-Z)\times$ prior mean.

(c) Show that $Z = \frac{m}{m+\beta+2}$.

(d) The actuary uses the above Bayesian approach and distributions to analyze two portfolios, A and B, described in the following table:

Portfolio	Number of policies	Number of claims	β
A	100	20	20
B	100	10	10

Example: ASTAM Fall 2023 Q4(c)(d) (cont.)

Example

(d) (i) For each portfolio, calculate the credibility factor Z. [A: 0.8197; B: 0.8929]

(ii) For each portfolio, calculate the Bayes estimate of q. [A: 0.1803; B: 0.1071]

(iii) Interpret the impact of the value of β on the credibility factor. [larger β gives a smaller Z]

(Exercises: STAM Sample Qs 5, 24, & 272)

Further Remarks

- The major challenge of the Bayesian approach: it may be difficult to evaluate the Bayesian premium $\mathbb{E}[X_{n+1}|\mathbf{X} = \mathbf{x}]$ in practice, and numerical integration is often required.
- Conjugate pairs give simple posteriors. With n observations $x_1, \ldots, x_n$:

Model for $X \mid \Theta$	Prior for Θ	Posterior for Θ
Poisson(Θ)	Gamma(α, θ)	Gamma$\left(\alpha + \sum x_i,\ \dfrac{\theta}{n\theta + 1}\right)$
Bernoulli(Θ)	Beta(a, b)	Beta$(a + \sum x_i,\ b + n - \sum x_i)$
Exponential (mean Θ)	Inv. gamma(α, β)	Inv. gamma$(\alpha + n,\ \beta + \sum x_i)$
Normal(Θ, σ^2), σ^2 known	Normal(μ, a)	Normal$\left(\dfrac{\sigma^2\mu + a\sum x_i}{\sigma^2 + na},\ \dfrac{\sigma^2 a}{\sigma^2 + na}\right)$

(Exercise: ASTAM S25 Q6(a)(b) for exponential/inverse gamma.)

Linear Estimate of Bayesian Premium

- Instead of estimating the posterior $\pi(\theta|\mathbf{x})$, Bühlmann (1967) suggested directly estimating the hypothetical mean $\mu_{n+1}(\theta)$ by a **linear function of the past data**.
- That is, we restrict the estimators to be of the form $\alpha_0 + \sum_{j=1}^{n} \alpha_j X_j$, where the α_i need to be chosen.
- Note that the credibility premium $Z\bar{X} + (1-Z)\mu$ is of this form.
- We choose α_i to minimize the squared-error loss:

$$Q = \mathbb{E}\left[\left(\mu_{n+1}(\Theta) - \alpha_0 - \sum_{j=1}^{n} \alpha_j X_j\right)^2\right]$$

- Referred to as **Bühlmann credibility**, **least squares credibility**, or **greatest accuracy credibility**.

Credibility Premium

- Taking partial derivatives w.r.t. α_i, we conclude that the credibility premium (as the *best linear approximation of the Bayesian premium in the sense of squared-error loss*) is given by

$$\tilde{\alpha}_0 + \sum_{j=1}^{n} \tilde{\alpha}_j X_j$$

where $\tilde{\alpha}_j$ $(j = 0, 1, \ldots, n)$ minimize the squared error loss Q and satisfy the following **normal equations**:

$$\mathbb{E}[X_{n+1}] = \mathbb{E}[\mu_{n+1}(\Theta)] = \tilde{\alpha}_0 + \sum_{j=1}^{n} \tilde{\alpha}_j \mathbb{E}[X_j]$$

$$\mathrm{Cov}(X_i, X_{n+1}) = \sum_{j=1}^{n} \tilde{\alpha}_j \mathrm{Cov}(X_i, X_j), \quad i = 1, 2, \ldots, n.$$

The Bühlmann Model: Assumptions

- Let past losses X_j be iid conditional on Θ, with hypothetical mean and process variance:

$$\mu(\theta) = \mathbb{E}[X_j \mid \Theta = \theta], \qquad v(\theta) = \mathrm{Var}(X_j \mid \Theta = \theta)$$

- Further define
  - Expected value of the hypothetical means $\mu = \mathbb{E}[\mu(\Theta)]$
  - Expected value of the process variance $v = \mathbb{E}[v(\Theta)]$
  - Variance of the hypothetical means $a = \mathrm{Var}(\mu(\Theta))$
- The unconditional mean, variance, and covariance of the X_j's are

$$\mathbb{E}[X_j] = \mu, \quad \mathrm{Var}(X_j) = v + a, \quad \mathrm{Cov}(X_i, X_j) = a$$

- The Bühlmann credibility premium is then

$$\tilde{\alpha}_0 + \sum_{j=1}^{n} \tilde{\alpha}_j X_j = Z\bar{X} + (1 - Z)\mu,$$

where the **Bühlmann credibility factor** Z is given by

$$Z = \frac{n}{n + k}, \qquad k = \frac{v}{a} = \frac{\mathbb{E}\left[\mathrm{Var}(X_j \mid \Theta)\right]}{\mathrm{Var}\left(\mathbb{E}[X_j \mid \Theta]\right)}$$

The Bühlmann Model (cont.)

- A desired weighted-average formula of $\bar{X}$ and the collective premium μ.
- $n \uparrow \Rightarrow \ Z \uparrow 1 \ \Rightarrow$ more credit to $\bar{X}$.
- k helps capture the homogeneity level of the risk class
  - For a homogeneous population, $\mu(\Theta)$ does not vary greatly w.r.t. Θ (small variability), and a is small relative to v $\Rightarrow k \uparrow \Rightarrow Z \downarrow 0 \ \Rightarrow$ more credit to μ.
  - For a heterogeneous population, $\mu(\Theta)$ is more variable and a is large $\Rightarrow k \downarrow \Rightarrow Z \uparrow 1 \ \Rightarrow$ more credit to $\bar{X}$.

Example: STAM Sample Q219

Example

For a portfolio of policies, you are given:

(i) The annual claim amount on a policy has probability density function:

$$f(x|\theta) = \frac{2x}{\theta^2}, \quad 0 < x < \theta$$

(ii) The prior distribution of θ has density function:

$$\pi(\theta) = 4\theta^3, \quad 0 < \theta < 1$$

(iii) A randomly selected policy had claim amount 0.1 in Year 1.

Calculate the Bühlmann credibility estimate of the claim amount for the selected policy in Year 2. [0.4283]

(Exercises: STAM Sample Qs 78, 251 & ASTAM S24 Q5, S25 Q6(c)-(e))

Example: STAM Sample Q187

Example

You are given:

(i) The annual number of claims on a given policy has a geometric distribution with parameter β.

(ii) The prior distribution of β has the Pareto density function:

$$\pi(\beta) = \frac{\alpha}{(\beta+1)^{\alpha+1}}, \quad 0 < \beta < \infty,$$

where α is a known constant greater than 2.

(iii) A randomly selected policy had x claims in Year 1.

Calculate the Bühlmann credibility estimate of the number of claims for the selected policy in Year 2. $[(x+1)/\alpha]$

(Exercises: STAM Sample Qs 32, 62, 133, 151, 177, 215, & 230; ASTAM Sample Q21)

Example: STAM Sample Q190

Example

For a particular policy, the conditional probability of the annual number of claims given $\Theta = \theta$, and the probability distribution of Θ are as follows:

Number of Claims	0	1	2
Probability	2θ	θ	$1 - 3\theta$

θ	0.10	0.30
Probability	0.80	0.20

Two claims were observed in Year 1.
Calculate the Bühlmann credibility estimate of the number of claims in Year 2. [1.4455]

(Exercises: STAM Sample Qs 41, 48, 70 & 236)

Example: STAM Sample Q8

Example

You are given:

(i) Claim counts follow a Poisson distribution with mean θ.

(ii) Claim sizes follow an exponential distribution with mean 10θ.

(iii) Claim counts and claim sizes are independent, given θ.

(iv) The prior distribution has probability density function

$$\pi(\theta) = \frac{5}{\theta^6}, \quad \theta > 1.$$

Calculate the Bühlmann's k for aggregate losses. [2.25]

(Exercises: STAM Sample Qs 67, 154, 181; 35, 200, & 267; ASTAM Sample Q4, Q23)

Notation for Credibility Data

Notation

r	number of policyholders (or groups), indexed by i
n_i	number of years of data for policyholder i, indexed by j
m_{ij}	exposure (e.g., members) for policyholder i in year j
X_{ij}	average loss per exposure unit for i in year j
$m_i,\ \bar{X}_i$	$m_i = \sum_j m_{ij}$ and $\bar{X}_i = \sum_j m_{ij} X_{ij} / m_i$
$m,\ \bar{X}$	$m = \sum_i m_i$ and $\bar{X} = \sum_i m_i \bar{X}_i / m$

Example: two classes of policies.

	Year 1		Year 2		Year 3	
Class	Exposures	Avg. loss	Exposures	Avg. loss	Exposures	Avg. loss
A	40	120	60	110		
B	20	150	30	140	25	160

Here $r = 2$; $n_A = 2$ and $n_B = 3$; $m_{A1} = 40$ and $X_{A1} = 120$; $m_A = 100$ and $\bar{X}_A = 114$; $m_B = 75$ and $\bar{X}_B = 149.33$; $m = 175$.

The Bühlmann-Straub Model: Assumptions

- The Bühlmann model assumes iid policyholder's past experience, but in practice allowing for variations in exposure or size is desired.
- The **Bühlmann-Straub** model assumes that the X_j's are independent conditional on Θ, with

$$\mathbb{E}[X_j \mid \Theta = \theta] = \mu(\theta), \qquad \mathrm{Var}(X_j \mid \Theta = \theta) = \frac{v(\theta)}{m_j}$$

  with m_j being known exposure, and X_j represents the average losses of m_j exposures of the jth policyholder.
- Here $v(\theta)$ is the process variance for *one unit* of exposure; averaging over m_j exposures divides it by m_j. So $v = \mathbb{E}[v(\Theta)]$ is a variance per exposure unit, not the variance of X_j.
- The Bühlmann model is the special case $m_j = 1$ for all j; otherwise the two models are the same, and the Bühlmann-Straub credibility uses the total exposure $m = \sum_j m_j$ in place of n.

The Bühlmann-Straub Model: Assumptions (cont.)

- We still define

$$\mu = \mathbb{E}[\mu(\Theta)], \quad v = \mathbb{E}[v(\Theta)], \quad a = \mathrm{Var}(\mu(\Theta))$$

- The unconditional mean, variance, and covariance of the X_j's are

$$\mathbb{E}[X_j] = \mu, \quad \mathrm{Var}(X_j) = \frac{v}{m_j} + a, \quad \mathrm{Cov}(X_i, X_j) = a$$

The Bühlmann-Straub Model

- The Bühlmann-Straub credibility premium per exposure is then

$$\tilde{\alpha}_0 + \sum_{j=1}^{n} \tilde{\alpha}_j X_j = Z\bar{X} + (1 - Z)\mu,$$

where the **Bühlmann credibility factor** Z is given by

$$Z = \frac{m}{m + k}, \quad k = \frac{v}{a}, \quad \bar{X} = \sum_{j=1}^{n} \frac{m_j}{m} X_j$$

and $m = \sum_{j=1}^{n} m_j$ is the total exposure.

The Bühlmann-Straub Model (cont.)

- The Bühlmann model is a special case of the above Bühlmann-Straub model.
- The credibility premium may equal the Bayesian premium, i.e., **exact credibility**. This typically occurs in Bühlmann (and Bühlmann-Straub) situations involving linear exponential family members and their conjugate priors.

Example: STAM Sample Q72

Example

You are given the following data on large business policyholders:

(i) Losses for each employee of a given policyholder are independent and have a common mean and variance.

(ii) The overall average loss per employee for all policyholders is 20.

(iii) The variance of the hypothetical means is 40.

(iv) The expected value of the process variance is 8000.

(v) The following experience is observed for a randomly selected policyholder:

Year	Average Loss per Employee	# of Employees
1	15	800
2	10	600
3	5	400

Calculate the Bühlmann-Straub credibility premium per employee for this policyholder. [12]

Example: STAM Sample Q21

Example

You are given:

(i) The number of claims incurred in a month by any insured has a Poisson distribution with mean λ.

(ii) The claim frequencies of different insureds are independent.

(iii) The prior distribution is gamma with probability density function:

$$f(\lambda) = \frac{(100\lambda)^6 e^{-100\lambda}}{120\lambda}, \ \lambda > 0$$

Example: STAM Sample Q21 (cont.)

Example

(iv)

Month	Number of Insureds	Number of Claims
1	100	6
2	150	8
3	200	11
4	300	?

Calculate the Bühlmann-Straub credibility estimate of the number of claims in Month 4. [16.91]

(Exercises: STAM Sample Qs 50, 139 & 263; ASTAM Sample Q9)

Empirical Bayes Estimation

- So far we have assumed that $f_{X_j|\Theta}(x_j|\theta)$ and $\pi(\theta)$ are known, but in practice they are left largely unspecified.
- We may need to use the data at hand to estimate the structural (prior) parameters, which is called **empirical Bayes estimation**.
- We will study two cases:
  - **Nonparametric**: $f_{X_j|\Theta}(x_j|\theta)$ and $\pi(\theta)$ are largely unspecified.
  - **Semiparametric**: $f_{X_j|\Theta}(x_j|\theta)$ is of parametric form but $\pi(\theta)$ is not.
- Whether or not to select a parametric model depends on the situation at hand and judgment.

Empirical Data

- For each of $r \geq 1$ policyholders (with risk parameter θ_i), we are given observed (average) losses per unit of exposure:

$$\mathbf{X}_i = (X_{i1}, \cdots, X_{in_i})^\top$$

  with known exposure vector $\mathbf{m}_i = (m_{i1}, \cdots, m_{in_i})^\top$ and observation years n_i.
- Experience of policyholders $(\mathbf{X}_i,\ i = 1, \ldots, r)$ is assumed to be independent.
- For fixed i, $X_{ij}|\Theta_i$ are independent with conditional pdf $f_{X_{ij}|\Theta_i}(x_{ij}|\theta_i)$ for $j = 1, \ldots, n_i$.
- Alternatively, you may also interpret r as risk classes and n_i as the number of policyholders in the ith group.

Empirical Estimation of Bühlmann Model

- In the Bühlmann model: $m_{ij} = 1$ and $n_i = n > 1$.
- The unbiased estimators for μ, v, and a are

$$\hat{\mu} = \bar{X} = \frac{1}{r}\sum_{i=1}^{r}\bar{X}_i = \frac{1}{rn}\sum_{i=1}^{r}\sum_{j=1}^{n}X_{ij}, \quad \text{with } \bar{X}_i = \frac{1}{n}\sum_{j=1}^{n}X_{ij}$$

$$\hat{v} = \frac{1}{r}\sum_{i=1}^{r}\hat{v}_i = \frac{1}{r(n-1)}\sum_{i=1}^{r}\sum_{j=1}^{n}\left(X_{ij} - \bar{X}_i\right)^2,$$

$$\text{with } \hat{v}_i = \frac{1}{n-1}\sum_{j=1}^{n}\left(X_{ij} - \bar{X}_i\right)^2$$

$$\hat{a} = \frac{1}{r-1}\sum_{i=1}^{r}\left(\bar{X}_i - \bar{X}\right)^2 - \frac{\hat{v}}{n}$$

Empirical Estimation of Bühlmann Model (cont.)

- The Bühlmann credibility premium for each policyholder/risk group is

$$\hat{Z}\bar{X}_i + \left(1 - \hat{Z}\right)\hat{\mu} \quad \text{with } \hat{Z} = \frac{n}{n + \hat{v}/\hat{a}}$$

Example: STAM Sample Q12

Example

You are given total claims for two policyholders:

Policyholder	Year 1 (Y1)	Y2	Y3	Y4
X	730	800	650	700
Y	655	650	625	750

Using the nonparametric empirical Bayes method, calculate the Bühlmann credibility premium for Policyholder Y. [687.38]

(Exercises: STAM Sample Qs 38 & 270)

Empirical Estimation of Bühlmann-Straub Model

- In general, we let

$$m_i = \sum_{j=1}^{n_i} m_{ij}, \qquad m = \sum_{i=1}^{r} m_i = \sum_{i=1}^{r} \sum_{j=1}^{n_i} m_{ij}$$

- The sample mean for each policyholder/risk class and the total are

$$\bar{X}_i = \frac{1}{m_i} \sum_{j=1}^{n_i} m_{ij} X_{ij}, \qquad \bar{X} = \frac{1}{m} \sum_{i=1}^{r} m_i \bar{X}_i = \frac{1}{m} \sum_{i=1}^{r} \sum_{j=1}^{n_i} m_{ij} X_{ij}$$

- The unbiased estimator for v is

$$\hat{v} = \frac{\sum_{i=1}^{r} \sum_{j=1}^{n_i} m_{ij} \left(X_{ij} - \bar{X}_i\right)^2}{\sum_{i=1}^{r} (n_i - 1)}$$

Empirical Estimation of Bühlmann-Straub Model (cont.)

- The unbiased estimator for a is

$$\hat{a} = \frac{\sum_{i=1}^{r} m_i \left(\bar{X}_i - \bar{X}\right)^2 - \hat{v}(r-1)}{m - \frac{1}{m}\sum_{i=1}^{r} m_i^2}$$

- To keep total credibility premiums equal to total observed losses, μ is estimated by the credibility-weighted average (the estimator on the ASTAM formula sheet; the exposure-weighted mean $\bar{X}$ is also unbiased, but the SOA expects this one)

$$\hat{\mu} = \frac{\sum_{i=1}^{r} \hat{Z}_i \bar{X}_i}{\sum_{i=1}^{r} \hat{Z}_i}, \quad \text{with } \hat{Z}_i = \frac{m_i}{m_i + \hat{v}/\hat{a}}$$

- The Bühlmann-Straub credibility premium for each risk class/policyholder is

$$\hat{Z}_i \bar{X}_i + \left(1 - \hat{Z}_i\right) \hat{\mu}$$

Example: ASTAM Spring 2023 Q3(b)(c)

Example

An insurer issues auto insurance policies through two different brokers, Broker A and Broker B. Claim frequency data from the last three years is given in the following table. There were no policies issued through Broker B in Year 1.

Risk Group		Year 1 (Y1)	Y2	Y3	**Total**
Broker A	# of claims	39	28	15	**82**
	# of Policies	200	150	180	**530**
Broker B	# of claims	n/a	16	23	**39**
	# of Policies	n/a	50	60	**110**

Assume first that the insurer models the claim frequency for the two groups using Bühlmann-Straub credibility, with nonparametric empirical Bayes parameter estimation.

Example: ASTAM Spring 2023 Q3(b)(c) (cont.)

Example

(b) (i) Show that $\hat{v}$ is 0.5 to the nearest 0.1. You should calculate the value to the nearest 0.001. [0.5014]

(ii) Show that $\hat{a}$ is 0.017 to the nearest 0.001. You should calculate the value to the nearest 0.0001. [0.0172]

(c) Calculate the Bühlmann-Straub credibility estimate of the expected claim frequency for an individual policy issued through Broker B in Year 4. [0.3317]

(Exercises: STAM Sample Qs 145, 194, 223 & ASTAM F25 Q1)

Semiparametric Estimation: Poisson Model

- In some situations, it may be reasonable to assume a parametric form for the conditional distribution $f_{X_{ij}|\Theta_i}(x_{ij}|\theta_i)$.
- We may then simplify the estimation for v and a.
- If the model is Poisson, then

$$\hat{\mu} = \hat{v} = \bar{X}$$

- $\hat{a}$ can be calculated using the previous nonparametric estimators.

Example: STAM Sample Q197

Example

You are given:

(i) During a 2-year period, 100 policies had the following claims experience:

Total Claims in Years 1 & 2	0	1	2	3	4
Number of Policies	50	30	15	4	1

(ii) The number of claims per year follows a Poisson distribution.

(iii) Each policyholder was insured for the entire 2-year period.

A randomly selected policyholder had one claim over the 2-year period. Using semiparametric empirical Bayes estimation, calculate the Bühlmann estimate for the number of claims in Year 3 for the same policyholder. [0.3928]

(Exercises: STAM Sample Qs 159, 233 & ASTAM F24 Q5)