Quantifying Model Risk

CAS Annual Meeting – Session C3

Glenn Meyers

November 17, 2009

My Nomination for Honorary Actuary

- “There are ***known knowns***. These are things we know that we know. There are ***known unknowns***. That is to say, there are things that we now know we don’t know. But there are also ***unknown unknowns***. These are things we do not know we don’t know.” – D. Rumsfeld, February 12, 2002
- This talk deals with “***known unknowns***” and how our knowledge is influenced by data.
- ***Unknown Unknowns*** – Buy “sleep insurance.”

You have two models, Blue and Green.
You are uncertain which one applies.
How do you reflect this uncertainty?

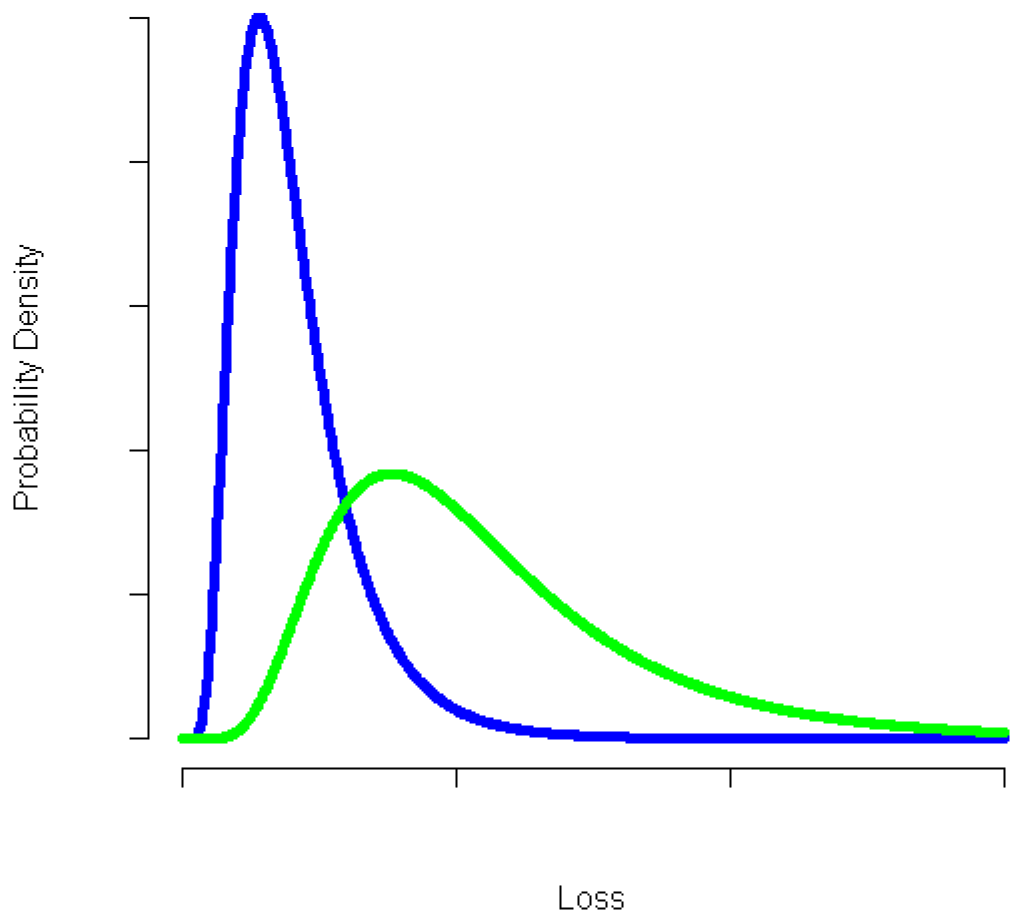

Take a mixture of the two

Red = 0.75 × Blue + 0.25 × Green

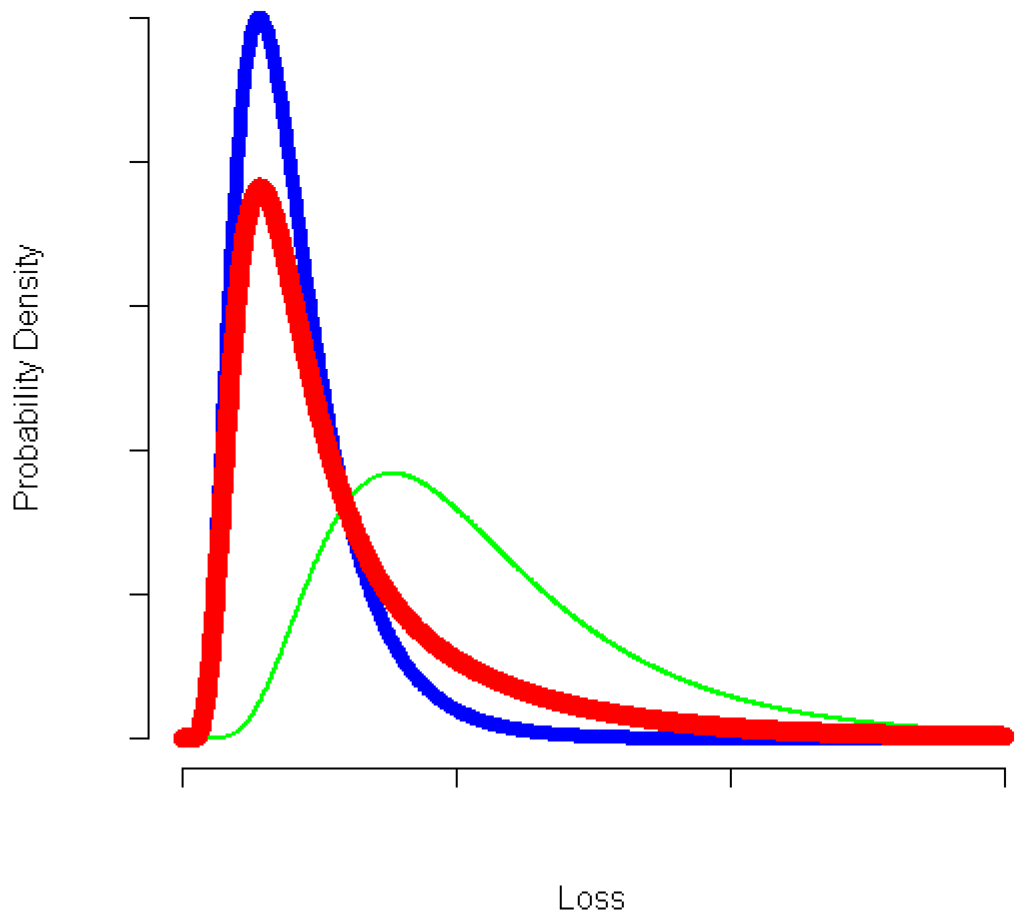

Take a mixture of the two

Red = 0.50 × Blue + 0.50 × Green

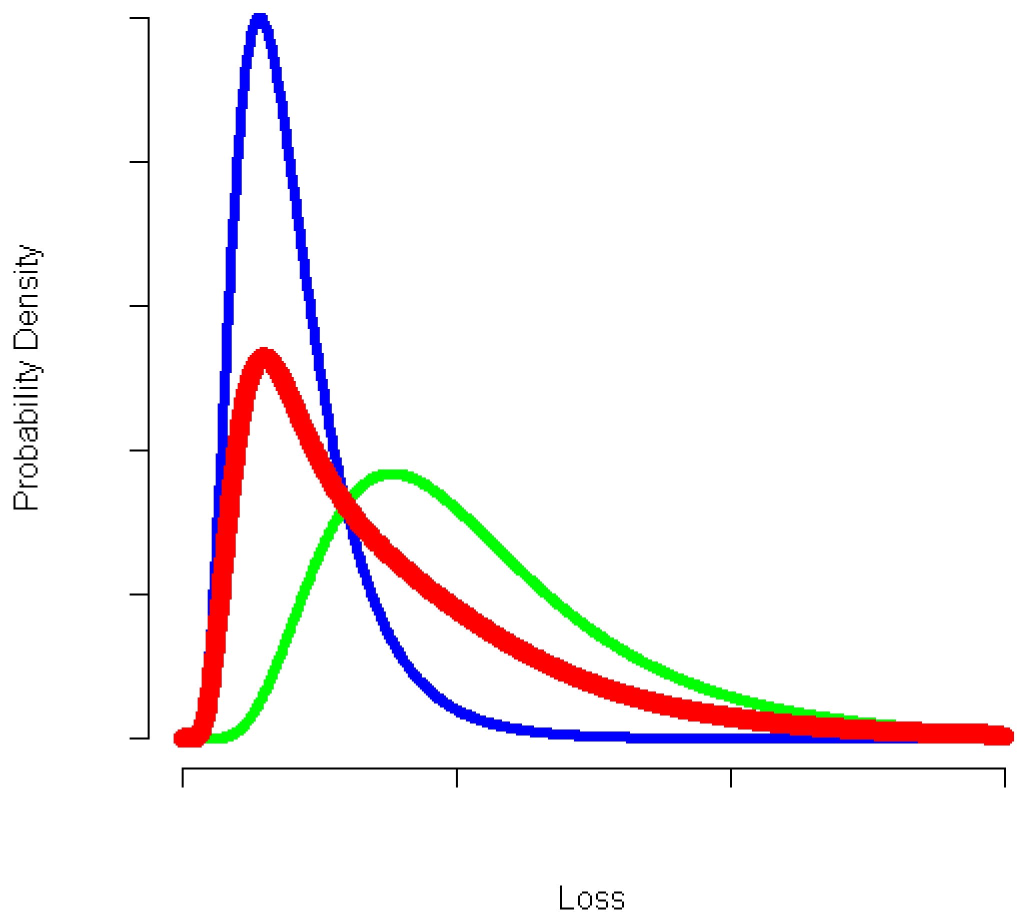

Take a mixture of the two

Red = 0.25 × Blue + 0.75 × Green

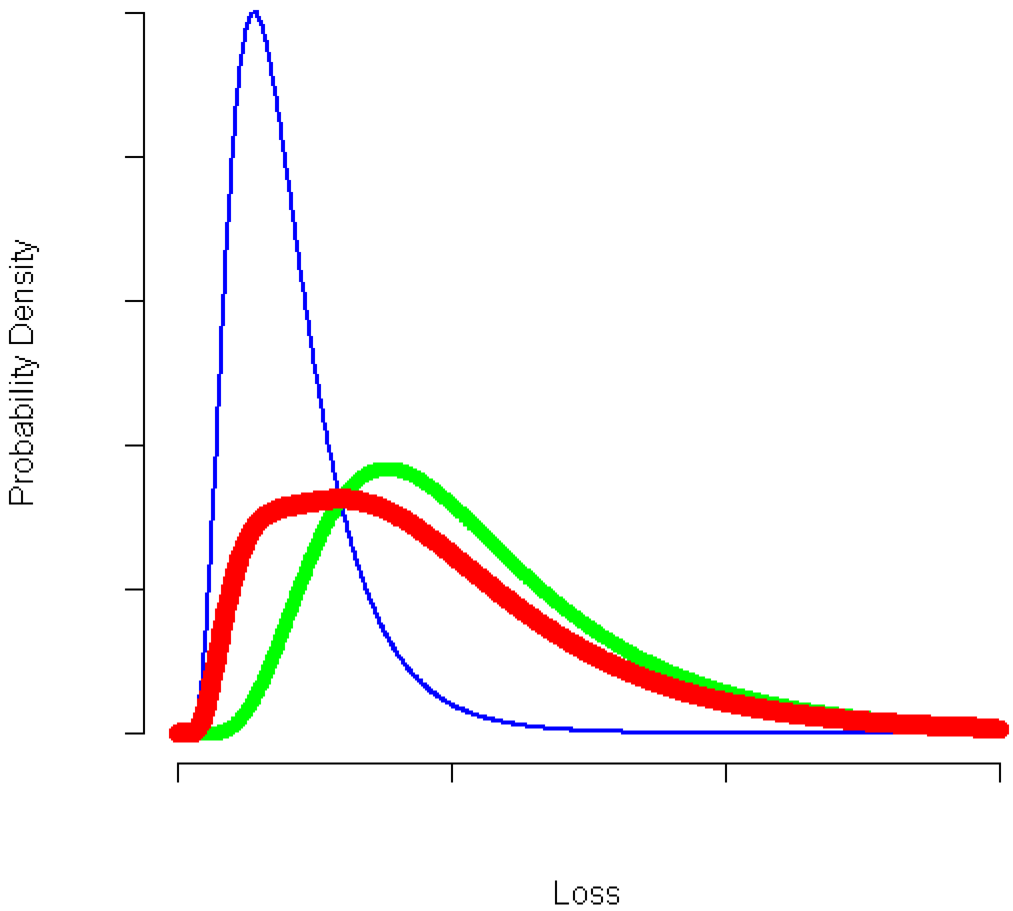

Identify Distributions

- Blue = Lognormal?
- Green = Gamma?
- Does it matter?
- Blue ~ Lognormal with $\mu = 7.5$ and $\sigma = 0.5$
- Green ~ Lognormal with $\mu = 8.5$ and $\sigma = 0.5$
- I want to discourage any distinction between “model risk” and “parameter risk”

Mixing Many Distributions

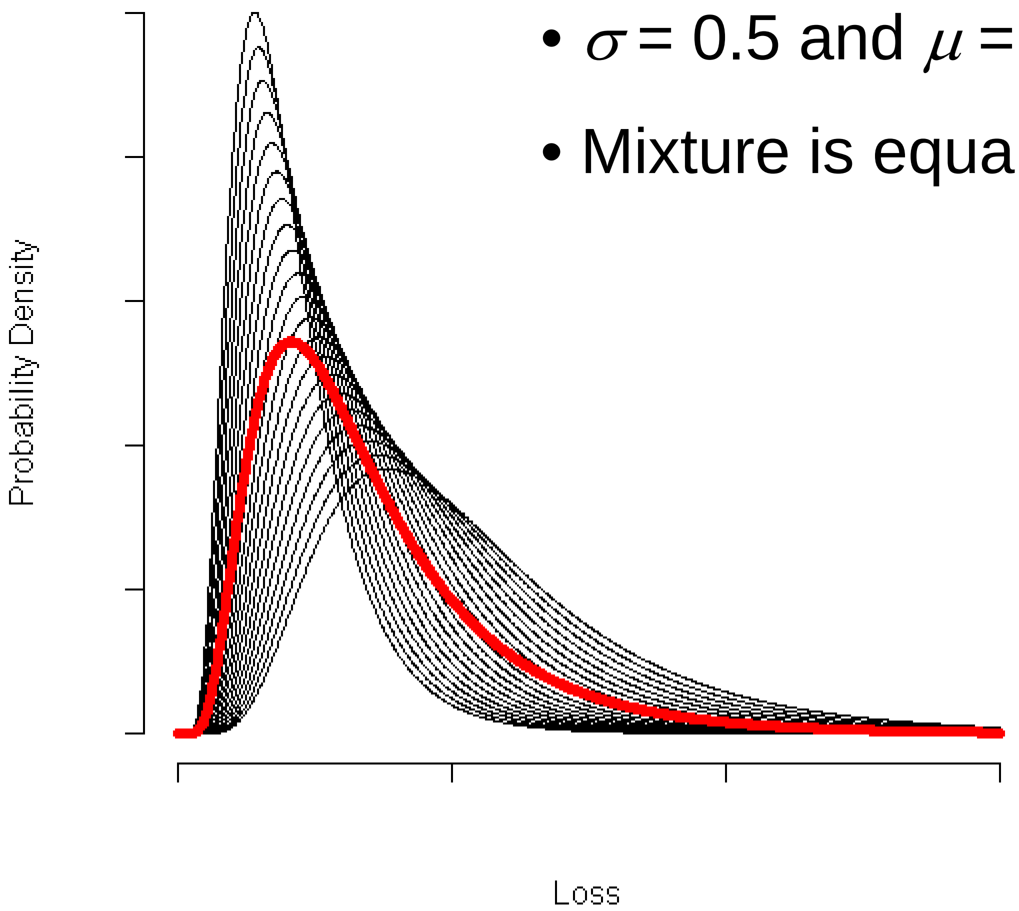

- σ = 0.5 and μ = 7.5 – 8.5
- Mixture is equally weighted

Issues in using Mixtures

- Given that we have data
  - Choosing the mixing distributions
  - Choosing the mixing weights
- I will illustrate with a simple one-dimensional example, and follow up with links to more complicated examples.

Choosing the Mixing Distributions

- Start with conventional goodness of fit testing
  - PP Plots
  - Kolomogorov-Smirnov test
  - Chi-Square goodness of fit
  - etc
- Need not restrict to single model such as lognormal
- Pass on examples for short presentation

Find a Range for Parameters

- My nomination for the second most important theorem in statistics

The likelihood ratio test

- Suppose you have a model and a maximum likelihood estimate *k*-vector $\hat{\mathbf{p}}$
- You want a range for the "true" parameter vector $\mathbf{p}$

The Likelihood Ratio Test

Test $H_0 : \mathbf{p} = \mathbf{p}^*$ against $H_1 : \mathbf{p} \neq \mathbf{p}^*$

Theorem 2.10 in Klugman, Panjer & Willmot

If H_0 is true then:

$$\ln LR \equiv 2\left[\ln L\left(\hat{p}; x\right) - \ln L\left(p^*; x\right)\right]$$

has a χ^2 distribution with k degrees of freedom.

- 99% critical value for k = 1 is 6.63

Set of μ's for which H_0 is accepted

Simulation #1 – 25 Data Points

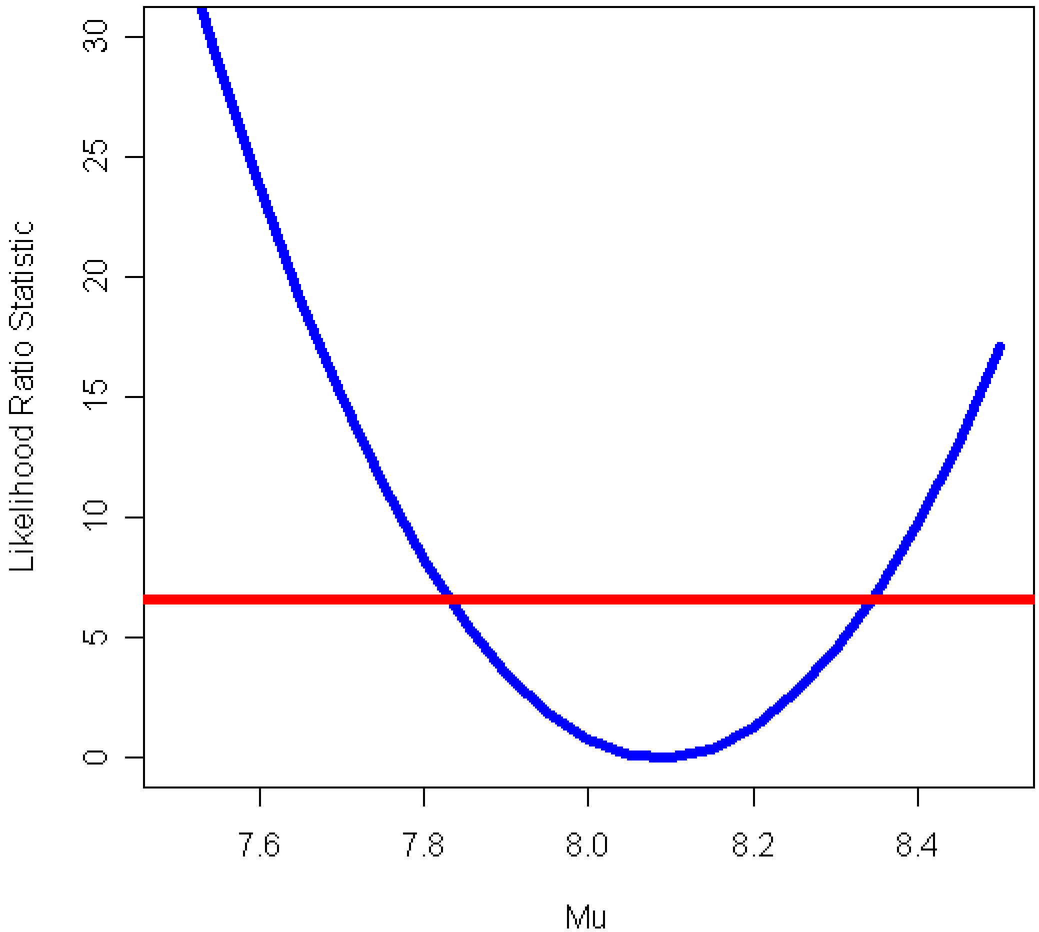

Set of μ's for which H_0 is accepted

Simulation #2 – 25 Data Points

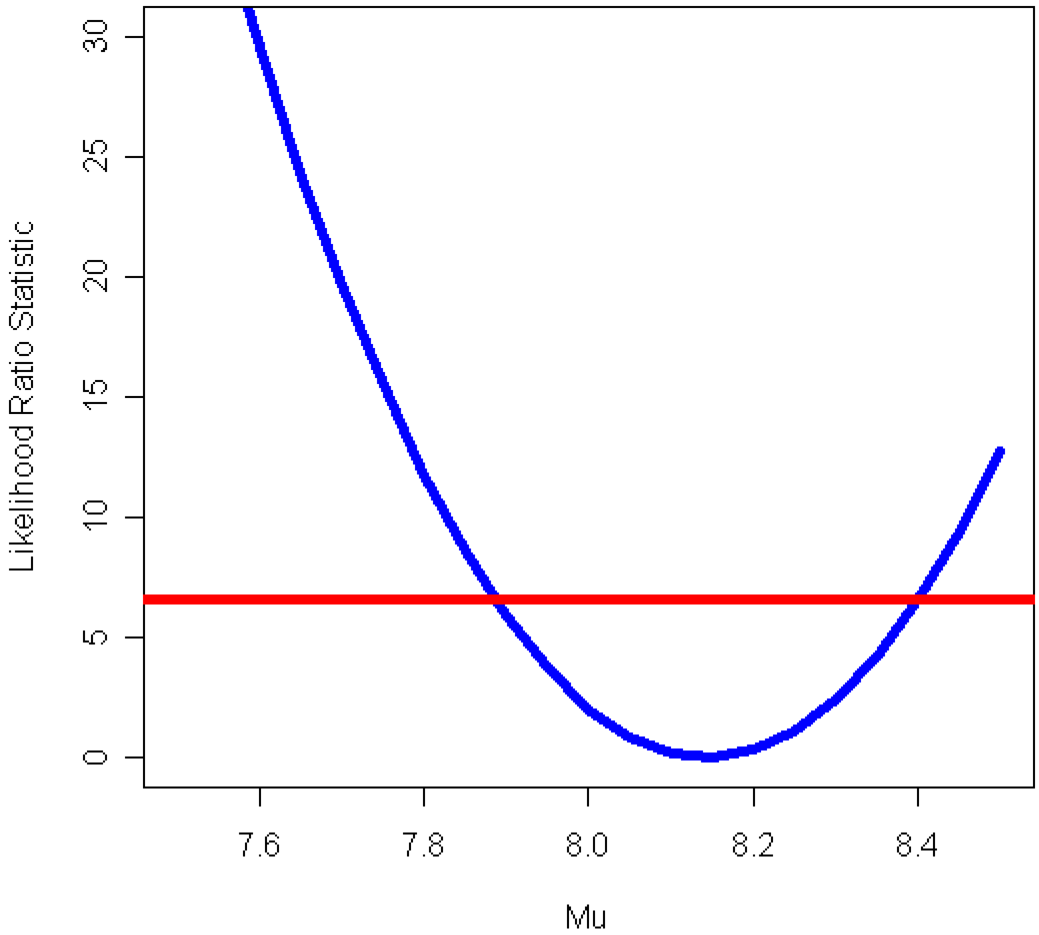

Set of μ's for which H_0 is accepted

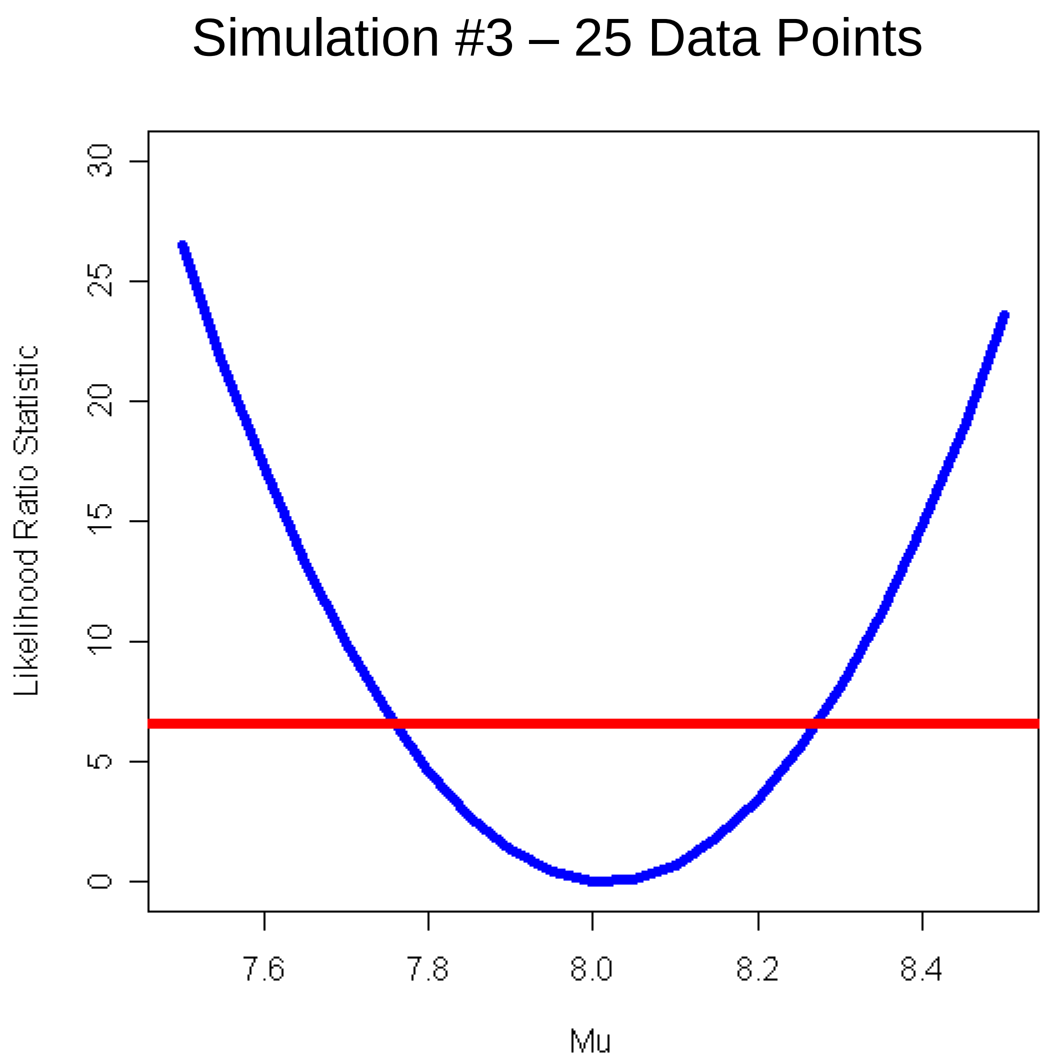

Set of μ's for which H_0 is accepted

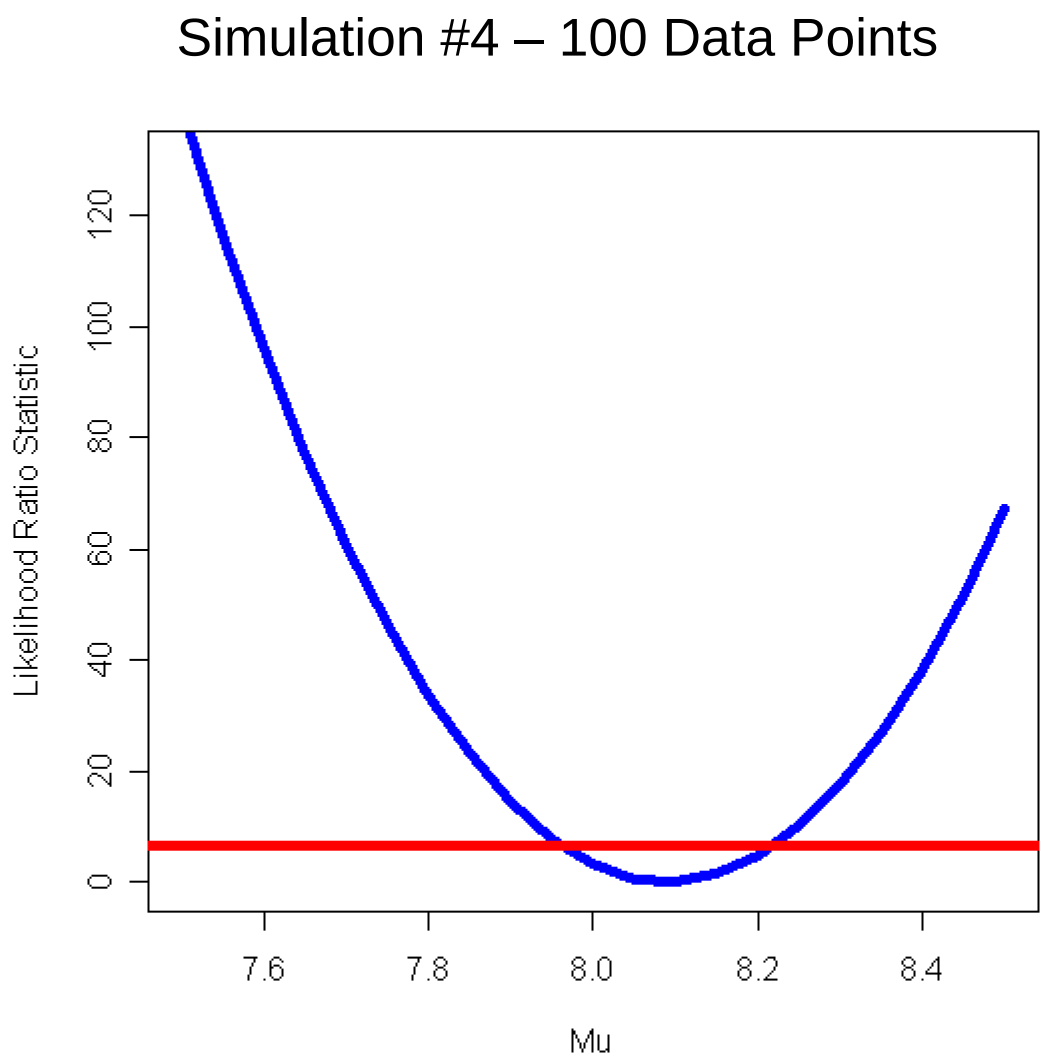

Choosing Weights for the Mixture

- My nomination for the most important theorem in statistics

Bayes Theorem

- Likelihood = Pr{Data|Model}
- Set Weight = Pr{Model|Data}

Using Bayes' Theorem

- Then using Bayes' Theorem, calculate the posterior probability of each μ given the data.
- Assume prior models are equally likely in this example.

$$\text{Posterior}\left\{\text{model} \mid \text{data}\right\}$$

$$\propto$$

$$\Pr\left\{\text{data} \mid \text{model}\right\} \times \text{Prior}\left\{\text{model}\right\}$$

Plot of Posterior Probability of μ
25 Observations

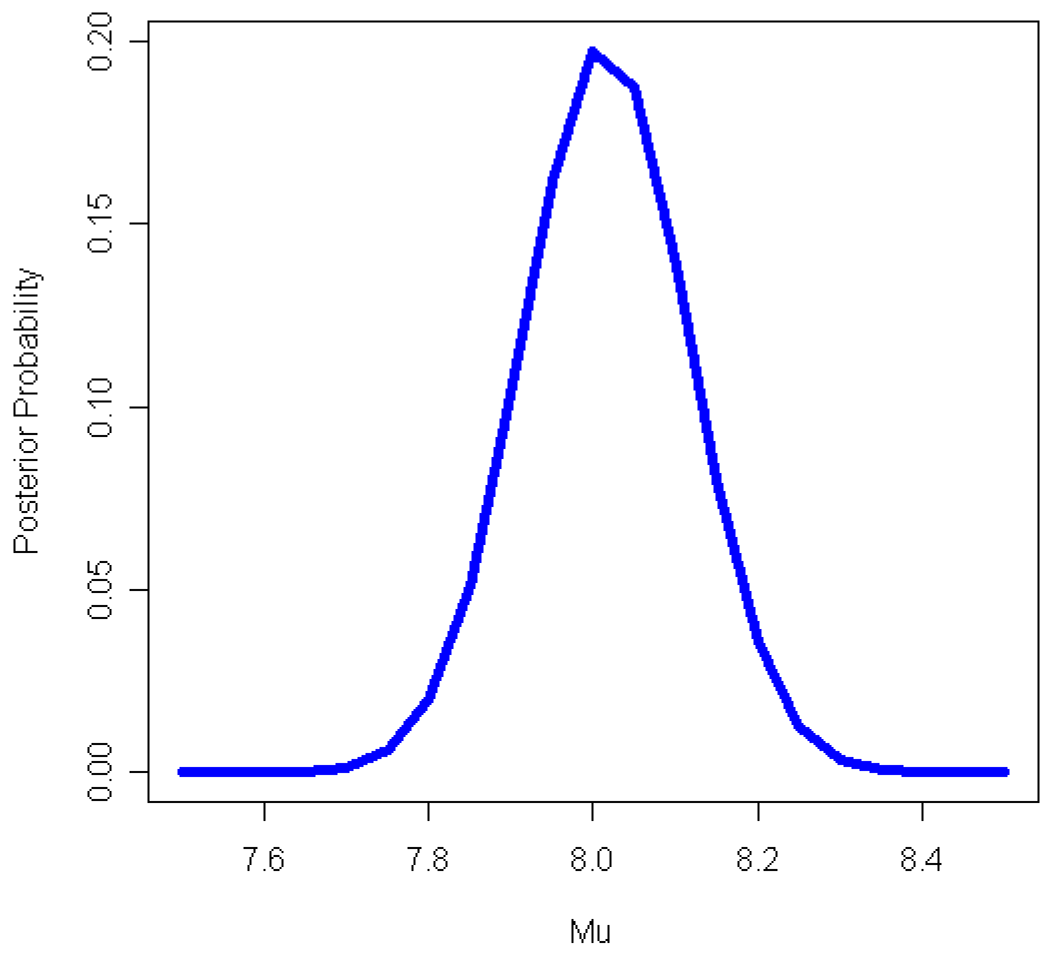

Posterior Weighted Mixture of Models
25 Observations

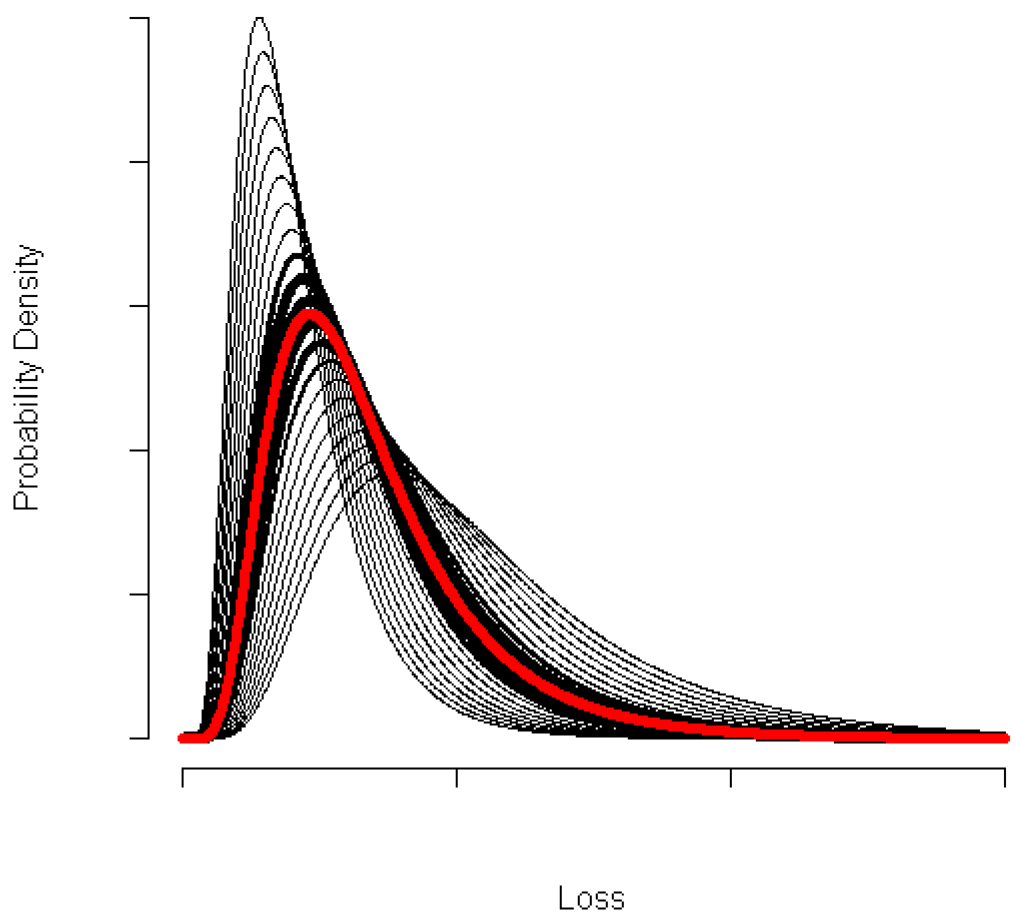

Mixing Many Distributions

- $\sigma = 0.5$ and $\mu = 7.5 – 8.5$
- Mixture is equally weighted

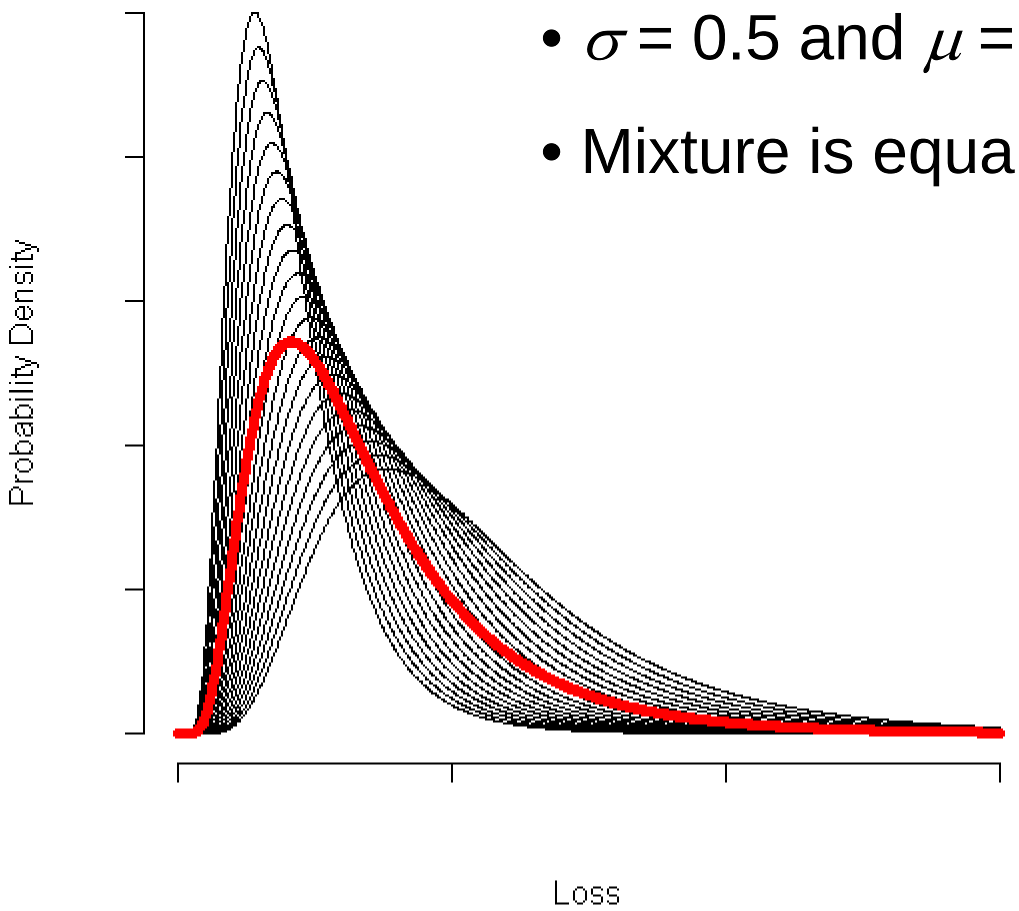

Quantities of Interest

- No real interest in the posterior probability weighted mixture of distributions.
- "Expected Reinsurer Deficit" is of more interest.

Expected Reinsurer Deficit

$$ERD = \int_{E[X]}^{\infty} \left(x - E[X]\right) \cdot f(x)dx$$

Computing *ERD* for a Mixture

- Calculate mean for each model and then mix to get the overall mean.
- Calculate integral for each model with the overall mean and then mix.

$$M = \sum_{i=0}^{20} Posterior_i \cdot e^{\mu_i + \sigma^2/2}$$

$$ERD = \sum_{i=0}^{20} Posterior_i \cdot \int_M^{\infty} \left(x - M\right) \cdot f_i(x) \cdot dx$$

Plot of *ERD* Calculations

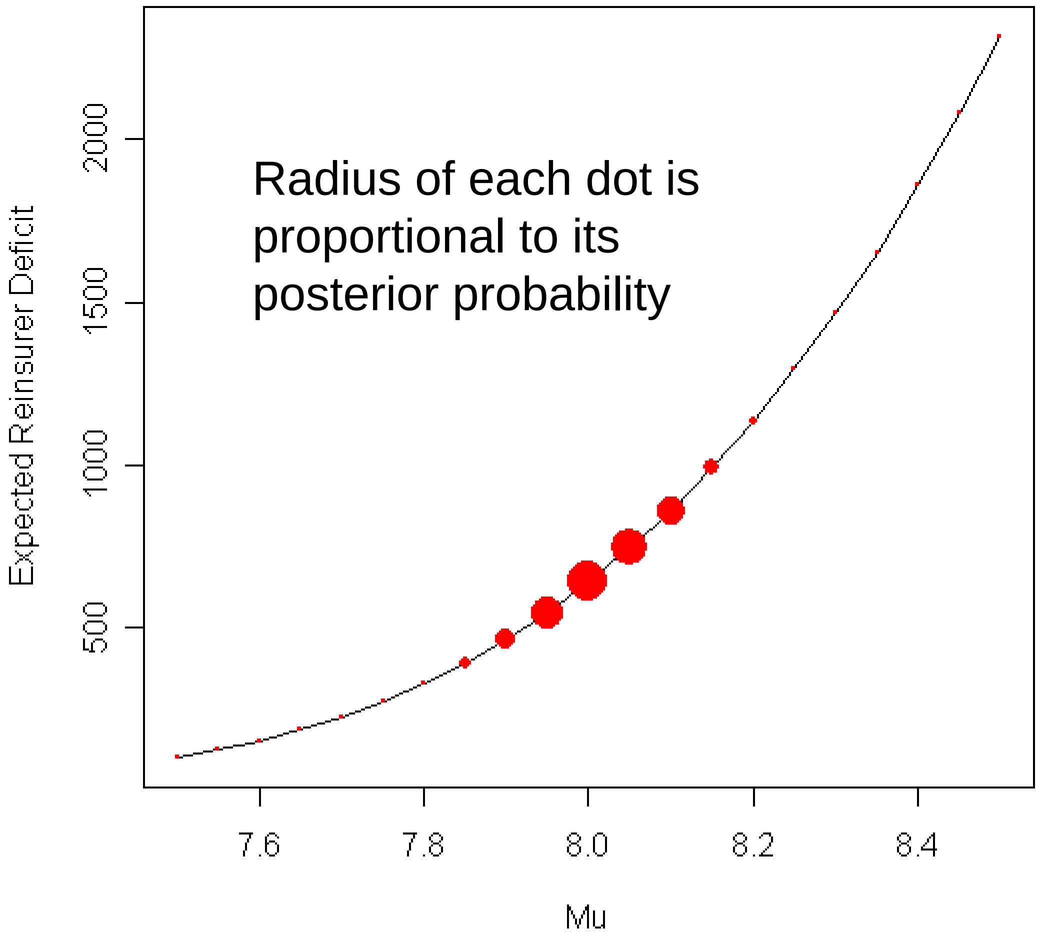

Contrast Mixture with Maximum Likelihood Estimate

Mixture

- *ERD* = 693
- *Mean* = 3,446
- *ERD* ÷ *Mean* = 20.1%

Maximum Likelihood

- *ERD* = 677
- *Mean* = 3,428
- *ERD* ÷ *Mean* = 19.7%

More Elaborate Examples

- 2005 COTOR Challenge

  http://www.casact.org/cotor/index.cfm?fa=round3
- Models had *log-t* distributions
- 4d parameter space
  - μ, σ, trend and degrees of freedom
- Uniform prior distributions on parameters
  - Academic example

More Elaborate Examples

- On Predictive Modeling for Claim Severity – *CAS Forum*, Summer 2005
  http://www.casact.org/pubs/forum/05spforum/05spf215.pdf
- Models had mixed exponential distributions derived from fits on large insurers
- Fixed parameters for each model
- Equal prior probability for each model
  - Real example

More Elaborate Examples

- “Proxies”

  http://www.actuaries.org/ASTIN/Colloquia/Helsinki/Papers/S4_21_Myers.pdf

- Uses Bayes’ Theorem and a loss reserve triangle to reweight 5,000 loss reserve models.
- Priors determined from MCMC scenarios of 50 large insurers
- See my *Variance* paper “Stochastic Loss Reserving with the Collective Risk Model” in Session P3 tomorrow.

Summary of Methods to Quantify Model Risk

- *Carefully, and with considerable thought*
  - Choose models that might describe the distribution of possible outcomes.
  - Assign prior probabilities to each model
- With Bayes' Theorem, calculate the posterior probability of each model given the data you have.
- Calculate quantities of interest (e.g. *ERD*) in terms of a mixture of models, i.e. the derived model risk.

Some object to assigning prior probabilities.

- Actuaries routinely render “Actuarial Statement of Opinions” (ASOP)
- Considerable thought and analysis can go into these actuarial opinions.
- Similar thought and analysis should go into prior distributions.
- Prior probabilities are transparent.