

Prof. David Draper
Department of Statistics
University of California, Santa Cruz

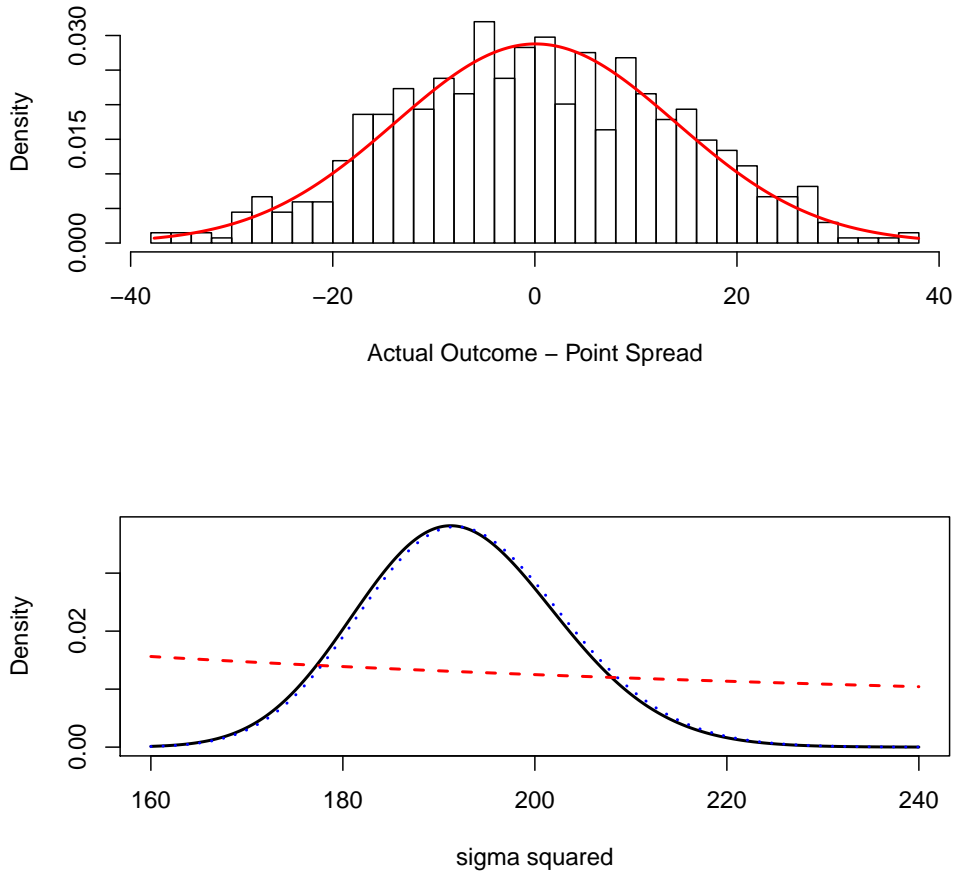
STAT 206: Quiz 3 [55 points]

Name: _____

(Inference for the variance in the Gaussian sampling model with known mean) People in Las Vegas who are experts on the National Football League provide a *point spread* for every football game before it occurs, as a measure of the difference in ability between the two teams (and taking account of where the game will be played). For example, if Denver is a 3.5–point favorite to defeat San Francisco, the implication is that betting on whether Denver’s final score minus 3.5 points exceeds or falls short of San Francisco’s final score is an even-money (50/50) proposition. The top panel of Figure 1 below (based on data from Gelman et al. (2014)) presents a histogram of the differences $d = (\text{actual outcome} - \text{point spread})$ for a representative sample of $n = 672$ professional football games in the early 1980s, with a Normal density superimposed having the same mean $\bar{d} = 0.07$ and standard deviation (SD) $s = 13.86$ as the sample (if this distribution didn’t have a mean that’s close to 0, the experts would be *uncalibrated* and you could make money by betting against them). You can see from this figure that the sampling model $(D_i | \sigma \mathcal{G} \mathcal{B}) \stackrel{\text{iid}}{\sim} N(0, \sigma^2)$ is reasonable for the observed differences d_i ($i = 1, \dots, n$), in which \mathcal{G} stands for the Gaussian sampling distribution assumption (which we’re making after looking at the data (the cheating approach to sampling distribution specification) and is therefore not part of \mathcal{B}). **Hint: You will find it much easier in this problem to work with the Gaussian variance parameter σ^2 if You define $\theta = \sigma^2$ and focus all of Your calculations on θ .**

- (1) Write down the likelihood and log likelihood functions for θ in this model [10 points]. Show that $\hat{\theta} = \hat{\theta}_{MLE} = \frac{1}{n} \sum_{i=1}^n d_i^2$, which takes the value 191.8 with the data in Figure 1, is both sufficient and the maximum likelihood estimator (MLE) for θ [10 points]. Plot the log likelihood function for θ in the range from 160 to 240 with these data, briefly explaining why it should be slightly skewed to the right. [5 points]

Figure 1: *Top panel: Differences d_i between observed and predicted American football scores, 1981–1984, with the Normal distribution with the same mean and SD superimposed (solid red curve); bottom panel: prior, likelihood and posterior for θ with the data in the top panel and the improper prior.*



- (2) The conjugate prior for θ in this model turns out to be (You don't have to show this) the *Scaled-Inverse- χ^2* distribution,

$$\theta \sim \chi^{-2}(\nu_0, \sigma_0^2), \quad \text{i.e.,} \quad p(\theta) = c \theta^{-(\frac{\nu_0}{2}+1)} \exp\left(-\frac{\nu_0 \sigma_0^2}{2\theta}\right), \quad (1)$$

in which ν_0 is the prior sample size and σ_0^2 is a prior estimate of θ (this is actually a re-parameterization of a special case of the Inverse Gamma family of densities). In an attempt to create a low-information-content prior, people sometimes work with a version of equation (1) obtained by letting $\nu_0 \rightarrow 0$, namely $p(\theta) = c_0 \theta^{-1}$. The resulting prior is *improper* (in the usual sense: that it integrates to $+\infty$), but it turns out that posterior inferences will be sensible nonetheless (even with sample sizes as small as $n = 1$). Show that with the general $\chi^{-2}(\nu_0, \sigma_0^2)$ prior in equation (1), the posterior for θ with the Gaussian likelihood is $\chi^{-2}\left(\nu_0 + n, \frac{\nu_0 \sigma_0^2 + n \hat{\theta}}{\nu_0 + n}\right)$, and conclude therefore that with the improper prior the posterior distribution is $\chi^{-2}(n, \hat{\theta})$ [5 points].

- (3) The bottom panel of Figure 1 plots the prior, likelihood, and posterior densities on the same graph using the data in the top panel of Figure 1 and taking $c_0 = 2.5$ for convenience in the plot. Get R (or some equivalent environment) to reproduce this figure (You'll need to identify which member of the χ^{-2} family the likelihood density is); include Your code as part of Your answer [5 points]. Explicitly identify the three curves, and briefly discuss what this plot implies about the updating of information from prior to posterior in this case study [10 points].

Programming notes: If You're working with R, You can use the `scaled.inverse.chisq.density` function I wrote for you, which is posted on the course web page. If You're writing code in a language other than R, note that — because the data values in this example lead to astoundingly large and small numbers on the original scale, which then need to be multiplied — it's necessary to do all possible computations on the log scale and wait to transform back to the original scale until the last possible moment (look at my R function to see what I mean). You'll need to be careful to use the correct normalizing constant c in equation (1), which can be found in Appendix A of the Gelman et al. (2014) book, and You'll need a function that computes $\log \Gamma(x)$ (such a function should be built into Your programming environment; *don't* compute $\Gamma(x)$ and take its logarithm [You'll get an overflow error with the data in this problem]).

- (4) It can be shown (You don't need to show this) that (as long as $\nu > 4$)

$$(\theta | \mathcal{B}) \sim \chi^{-2}(\nu, \sigma^2) \quad \rightarrow \quad E(\theta) = \left(\frac{\nu}{\nu - 2} \right) \sigma^2 \quad \text{and} \quad V(\theta) = \left[\frac{2\nu^2}{(\nu - 2)^2 (\nu - 4)} \right] \sigma^4. \quad (2)$$

Use this to compute the posterior mean and SD for θ with this data set (using the improper prior), and compare the posterior mean with the MLE, briefly explaining why they're so similar in this case [10 points].