Comparing Men's and Women's Salaries at MSU-Bozeman: Analyses of 2008-2009 Salary Data

Introduction

In Fall 2009, in an informal report to the Women's Faculty Caucus using the 2008-2009 faculty salary data, it was reported that "(w)omen make an average of $10,000 less than men across campus." However, no other factors were considered when computing this difference in average salaries, such as faculty rank (Assistant, Associate, Full), college and department of employment, or years spent at current position. In order to provide a more accurate picture of potential gender based differences in average salary, these additional factors must be considered, and appropriate models must be used. For this project, three separate analyses of the data set are run using three different models. The first analysis involves a naive multiple linear regression model which ignores department differences; the second utilizes a hierarchical, likelihood-based, linear model allowing for department differences, and the third uses a Bayesian hierarchical model, also allowing for the nested structure of departments within colleges. Results are reported from each of these modeling strategies, and comparisons among the different methods are discussed. Ultimately, the question of interest is whether or not there is evidence of an inequity in mean salary between men and women at MSU-Bozeman. Additionally, it is of interest whether there is evidence that any gender differences in mean salary depend upon faculty rank.

Exploratory Data Analysis

From initial plots and tables, it is evident that there are relatively large differences between the number of male and female faculty at Montana State University - Bozeman as well as differences in ranges of salary between male and female faculty. The table below shows the average salaries and numbers of male and female faculty members at each rank.

Table 1: Average salaries and numbers of faculty at each rank by gender

	Rank		
	Assistant	Associate	Full
Male	\$57,734 (n=87)	\$62,615 (n=94)	\$83,972 (n=141)
Female	\$53,548 (n=67)	\$59,750 (n=67)	\$81,736 (n=21)
Difference	\$4,186	\$2,865	\$2,236

Inference to all faculty members at MSU-Bozeman is, therefore, complicated by these differences. Additionally, because gender cannot be randomly assigned to a faculty member, it cannot be inferred that a faculty member's gender caused him or her to receive a higher or lower salary than his or her gender counterpart. We are simply assessing whether there does indeed appear to be statistical evidence of a difference. Furthermore, though this is a census of faculty member's salaries, the real interest does not lie in salary differences based on gender *for this exact population.* Rather, it is of interest whether the system in place for this group of faculty members or some other future group of MSU-Bozeman faculty members lends itself to creating a salary discrepancy between male and female faculty members. We can consider the people in the survey to be from a larger group of potential faculty members to investigate the possible inequity in salaries in order to make a statement about the current pay structure. The null hypothesis is, then, that there is no inequity in the system: that there is no difference between men's and women's salaries (after accounting for other variables). The statistical analysis then addresses how likely it is that simple chance is responsible for any observed difference in salaries if there is truly no inequality in salaries.

Initial plots are examined for patterns and outliers. Interaction plots exhibit interesting trends: when averaging across departments and colleges, men's salaries at each rank are higher than women's salaries. When averaging across all ranks, men's salaries at each college are higher than women's salaries. However, this pattern is not retained when college and department are accounted for.

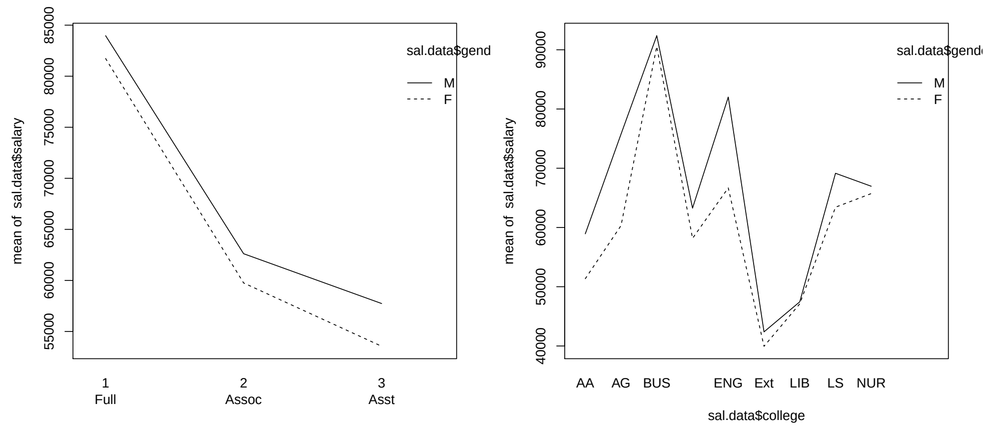

Figure 1: Average salaries by gender and rank, averaged across all departments and colleges (left). Average salaries by gender and college, averaged across all ranks (right)

From boxplots and scatterplots of the data, it is evident that there are two full professors, both male, with salaries which were much larger than other faculty members in their respective departments (near 2 standard deviations away from the means). In fact, these professors earn approximately double the average salary of others in their respective departments.

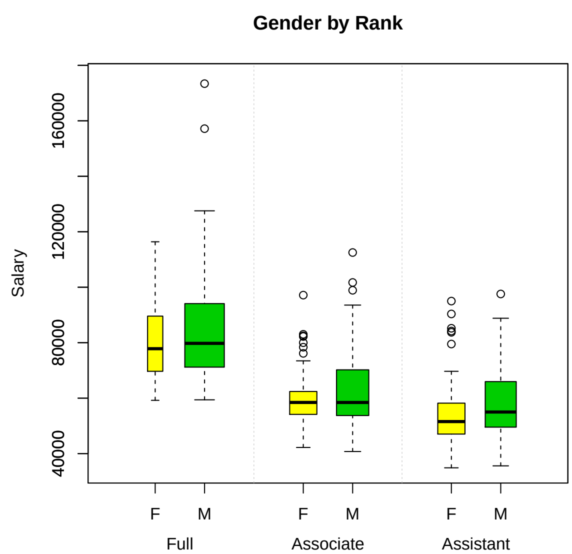

Figure 2: Salaries by gender and rank, averaged across all colleges. The width of each box represents the number of observations in each category, with wider boxes representing more observations.

These outliers are excluded from most analyses, not because they are considered erroneous reports, nor because diagnostics flag them as influential, but because each was substantially above the mean/median for their college and department for unknown reasons. Therefore, these outliers can potentially be considered as coming from a separate population, such as administration. Furthermore, we want to ensure that any results and conclusions reported are not simply due to one or two unusual observations. For the Multiple Linear Regression Analysis only, results are provided with the outliers included and excluded.

Several explanatory variables for the response are investigated for inclusion in the models. "Experience" is quantified by calculating the number of days since the faculty member was assigned his or her current rank to January 1, 2010. Gender and faculty rank as of 2009, as well as college and department of employment are also considered. From other exploratory plots, it is evident that mean salary differences between men and women depend on college and department within college. For example, the following plot illustrates the range of differences between male and female faculty members' salaries across departments within the College of Arts and Architecture. Note that while male salaries appear higher, on average, for the Architecture and Film & Photography departments, salaries between males and females in the Art and Music departments appear quite comparable.

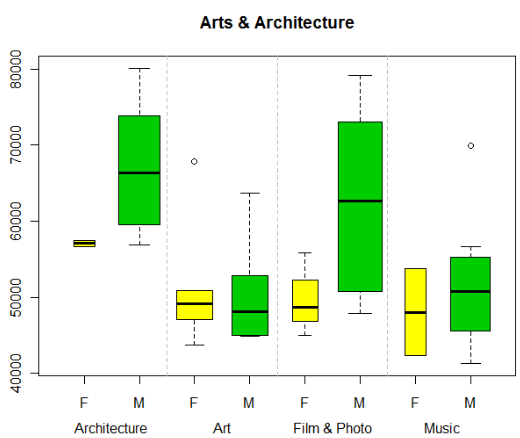

Figure 3: Average salaries by gender and department within the College of Art and Architecture, averaged across all ranks.

Naive Multiple Linear Regression

Statistical Procedures Used

After examining exploratory plots, a linear regression model is fit whereby mean salary is modeled as a function of experience, gender, faculty ranking, college of employment, and interactions which are deemed potentially meaningful, like gender-by-rank, gender-by-college, and the three-way interaction experience-by-rank-by-college. Plots of the residuals (with the two outliers included) exhibit increasing spread as fitted values increase. Additionally, from the Normal Q-Q plot, there is a clear departure from normality in the upper tail.

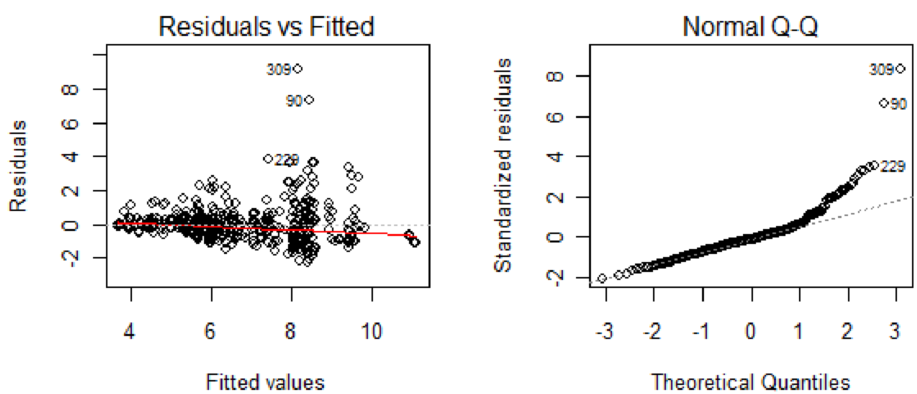

Figure 4: Plots of residuals from a model with salary modeled by experience, gender, faculty rank, college, and interactions (gender-by-rank, gender-by-college, and experience-by-rank-by-college).

The pattern of increasing variance as fitted values increase is evidence of non-constant variance in the residuals. Therefore, a transformation is considered. A Box-Cox plot indicates that a reciprocal transformation might be warranted, yet a natural log transformation is decided upon due to its ease of interpretation; it is more appropriate when discussing multiplicative differences in salaries. Residuals from the transformed model (without outliers) exhibit less of a departure from normality (though a departure from normality in the upper tail is still evident), but because multiple linear regression inference is robust to departures from normality, the upper tail departures visible in the Normal Q-Q plot do not indicate a violation of assumptions. Not surprisingly, residuals still exhibit an increasing spread as fitted values increase, indicating some remaining non-constant variance. This is, however, consistent with having fewer observations at lower salaries, making the departure from constant variance less severe than than it initially appears in the plot.

With this modeling strategy, department of employment is not taken into account. It is reasonable to assume that professors of similar rank and experience within the same department earn more similar salaries than two professors from different departments. Therefore, it is very likely that the assumption of independence between observations is violated, which puts inferences into question. For subsequent analyses, the use of more sophisticated models accounting for both department and college will help to ensure the independence assumption is adequately met.

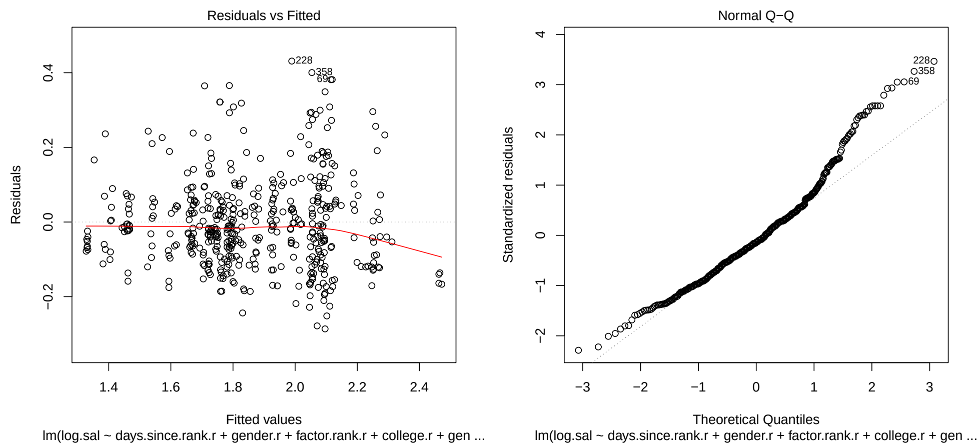

Figure 5: Diagnostic plots from a model with log-transformed salary and terms for gender, rank, experience, college, and a gender-by-rank interaction (outliers excluded).

Through the use of AIC, Likelihood Ratio Tests, and judgment in tying the questions of interest to the analysis, the model is pared down to include experience, gender, rank, and college of employment, and the interaction term of gender-by-rank. The multiple linear regression model used for inference is:

$$log(y_{ijkl}) = \mu + \beta x_i + \alpha_j + \tau_k + \gamma_l + (\alpha\tau)_{jk} + \epsilon_{ijkl}$$

with $\epsilon_{ijkl} \overset{iid}{\sim} N(0, \sigma^2)$

for $i = 1, 2, \ldots, 477$, $j = 1, 2$, $k = 1, 2, 3$, $l = 1, 2, \ldots, 9$ and where β represents the slope on experience, α represents the gender effect, τ represents the rank effect, γ represents the college effect, and $(\alpha\tau)$ represents the gender-by-rank interaction.

Summary of Statistical Findings

There is suggestive but inconclusive evidence that the ratio of male to female median salaries depends on rank (Full, Associate, or Assistant), after accounting for experience and college of employment (two-sided p-value=0.048 from an F-statistic of 3.054 on 2 and 460 d.f.). If we assume there is no interaction, then there is no evidence of a difference in median salary between men and women after considering faculty rank, college of employment, and experience (two-sided p-value=0.205 from a t-statistic of 1.270 on 462 d.f.). Under this no-interaction model, it is estimated that the median salary for men is 1.018 times that of the median salary for women when holding college, rank, and time spent at current rank fixed (95% confidence interval from 0.99 to 1.047 times). Translating into dollars, for a fixed rank, college, and amount of experience, if the median male salary is $50,000, the median female salary is estimated to be $884 less ($49,116). A 95% confidence interval is from $505 more than the median male salary of $50,000 to $2,245 less than median male salary of $50,000.

Allowing for separate gender differences by rank (including the interaction term), multiplicative factors and estimates of median salaries at different ranks can be calculated. The tables below provide these estimates and confidence intervals for female faculty members at each of the ranks, given that the median male salary at each rank is $50,000 and holding experience and college of employment fixed. Table 3 provides analogous estimates using the data set with the two outliers included. Notice that estimates of median female salaries for the Associate and Assistant ranks remain approximately the same when the outliers are included but that the estimated median salary for females at a Full rank is now below the median salary for males when the outliers are included, which is consistent with the outliers being male Full professors.

Table 2: Estimates and 95% Confidence Intervals of Median Salaries for Female Faculty Members at Each Rank, Assuming a Median Salary of $50000 for Males at Each Rank (no outliers)

	Rank		
	Full	Associate	Assistant
Multiplicative Factor	0.998	0.987	1.058
Mult. Factor 95% CI	(0.94, 1.06)	(0.947, 1.03)	(1.015, 1.102)
Male	$50,000	$50,000	$50,000
Female	$50,100	$50,658	$47,259
CI	$47,170-$53,191	$48,543-$52,798	$45,372-$49,261

Table 3: Estimates and 95% Confidence Intervals of Median Salaries for Female Faculty Members at Each Rank, Assuming a Median Salary of $50000 for Males at Each Rank (outliers retained)

	Rank		
	Full	Associate	Assistant
Multiplicative Factor	1.008	0.988	1.058
Mult. Factor 95% CI	(0.946, 1.075)	(0.945, 1.032)	(1.012, 1.105)
Male	$50,000	$50,000	$50,000
Female	$49,603	$50,607	$47,259
CI	$46,512-$52,854	$48,449-$52,910	$45,249-$49,407

It must be stressed that this analysis is naive in several ways. No interaction between college and gender is included in the model since the term was not statistically significant; therefore, there is an added assumption of equal salary differences between genders within each college. This may not be an appropriate assumption because the model does not account for department of employment within college (because of the limitations of a multiple linear regression model). Furthermore, because department is excluded as a nested factor, one major model assumption is violated: independence among residuals from the model after accounting for rank, gender, experience, and college of employment. Again, it is likely that professors of a given rank, experience, and gender are more similar to each other within departments than across departments.

Therefore, a more appropriate model will include terms for rank, experience, gender, and college, as well as a term for department nested within college.

Multilevel Regression

For this analysis, a likelihood-based multilevel regression model is employed using the lme package in the computing program R (Pinheiro, J., Bates, D., DebRoy, S., Sarkar, D. and R Development Core Team, 2010). Unlike multiple linear regression, multilevel regression allows for explicit modeling of hierarchical, or nested, structures. The structure is incorporated by placing distributions on parameters, similar to a traditional "random effects" model. Gelman and Hill (2007) note that this hierarchical modeling is, among other things, a valuable tool for the study of variation among coefficients and recommend employing multilevel modeling when analyzing data sets such as this one.

Statistical Procedures Used

Recognizing that the same covariates from the previous model are still pertinent, a model including faculty rank, experience at current rank, college, gender, and a gender-by-rank interaction term is initially fit. As discussed and justified in earlier sections, the two outliers are excluded from this analysis. Department of employment is now taken into account by allowing a different intercept for each department within college. Again, a log transformation is employed, and assumptions are checked via diagnostics as in the previous analysis.

Thus, the multilevel regression model used for inference is:

$$log(y_{ijklm}) = \mu + \beta x_i + \alpha_j + \tau_k + (\alpha\tau)_{jk} + \gamma_l + \delta_{m(l)} + \epsilon_{m(ijkl)}$$

for $i = 1, 2, \ldots, 477$, $j = 1, 2$, $k = 1, 2, 3$, $l = 1, 2, \ldots, 9$, and $m = 1, 2, \ldots, n_d$ where

$\epsilon_{m(ijkl)} \overset{iid}{\sim} N(0, \sigma^2)$,

$\boldsymbol{\gamma} \sim N(\mathbf{0}, \boldsymbol{\Psi_1})$ where Ψ_1 is the 9×9 variance-covariance matrix for the college effect, with σ^2_{coll} on the diagonal and allowing for different correlations among colleges

$\boldsymbol{\delta} \sim N(\mathbf{0}, \mathbf{\Psi_2})$ where Ψ_2 is the 46×46 variance-covariance matrix for department-within-college effect, with $\sigma^2_{dept(coll)}$ on the diagonal and allowing for different correlations among departments

Again, β represents the slope on experience, α represents the gender effect, τ represents the rank effect, γ represents the college effect, and now, δ represents the department-within-college effect.

Summary of Statistical Findings

There is suggestive but inconclusive evidence that the ratio of male to female median salaries depends on rank, after accounting for experience and variability within college and department (p-value = 0.043 from a Chi-squared statistic of 6.29 on 2 d.f.). It is estimated, for example, that the median salary for male full professors is 0.993 times that of the median salary of female full professors, after accounting for experience and department within college and department (95% confidence interval from 0.946 to 1.042 times the median female full professor's salary). In terms of dollars, for a fixed department and college as well as experience, if the median salary for a male full professor is $50,000, the median salary for a female full professor is estimated to be $50,352, $352 more (95% confidence interval from $2,854 more to $2,015 less than the median male's salary of $50,000). Estimates and confidence intervals based on a median salary of $50,000 for male faculty at ranks of Associate and Assistant are included in the table below.

Table 4: Estimates of Median Salaries for Female Faculty Members at Each Rank, Assuming a Median Salary of $50000 for Males at Each Rank

	Rank		
	Full	Associate	Assistant
Multiplicative Factor	0.993	0.977	1.035
Male	$50,000	$50,000	$50,000
Female	$50,352	$51,177	$48,309
CI	$47,985-$52,854	$49,505-$52,910	$46,729-$49,900

Note that these estimates and confidence intervals exhibit similar trends to those from the previous analysis. Again, estimates of median salaries for Full and Associate female professors have corresponding confidence intervals which cover $50,000, yet the confidence interval for the median salary estimate for a female Assistant professor does not. Something to note, however, is that these confidence intervals are a bit narrower than the intervals from the naive multiple regression model. We have essentially included a blocking variable (department), which groups similar observations together, thereby decreasing variability in gender differences.

Bayesian Hierarchical Modeling

In the previous analysis, the necessity of including department as a random effect nested within college does not permit the addition of a department-by-gender effect. Therefore, the assumption must be made that the gender differences remain constant across departments within a college. From exploratory plots, this isn't the case. For example, in the College of Arts and Architecture (Fig. 3), it is clear that in the Art Department median salaries between men and women are much more similar than they are in the Film and Photography Department. With a Bayesian hierarchical model, this dependence of differences in mean salaries is easily accounted for in conjunction with an appropriate nesting structure. Gender differences within departments are considered related to and nested within gender differences in the respective colleges. This results in a different weighting of the within-department differences than is obtained from the likelihood-based hierarchical model. The Bayesian hierarchical model is reported below; however, a diagram may aid in understanding

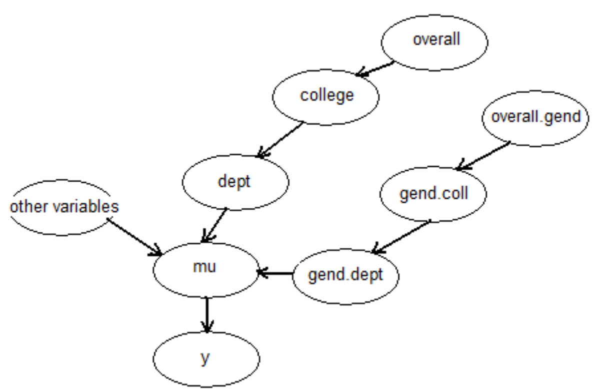

Figure 6: An illustration of one Bayesian hierarchical model

Model

In this model, $\log(y_{ijklm})$ is explained by experience and rank (τ_k) with adjustments for mean salary in a department relative to the college of employment ($\mu_{sal_{m(l)}}$) and a mean gender difference in salaries within a department nested within a mean gender difference in salaries in the respective college ($\alpha_{gend_{m(l)}}$). In this way differences variability between male and female salaries among departments as well as among colleges is allowed and modeled hierarchically. Furthermore, for those departments with few faculty members or very few faculty members of differing genders, the differences in salaries are mediated by an overall difference in salaries between genders at the college level. This way, information on a large number of observations (i.e. the college-level mean difference in salaries between males and females) is used to assist in generating estimates for gender differences in salaries at the department level. When there are several faculty members of both genders in a department, not as much "assistance" is required from the college.

Distributions are placed on the parameters, and diffuse prior distributions are placed on the hyperparameters. The model is fit using WinBUGS (Medical Research Council Biostatistics Unit, 2004) and the R package R2WinBUGS (Sturtz, S., Ligges, U., and Gelman, A., 2005). To be consistent with the parameterization of the Normal distribution in WinBUGS, the inverse of the variance parameter (the precision) is used in parameterizing the distribution, rather than the variance parameter, σ^2, itself ($\theta = \frac{1}{\sigma^2}$).

$$\log(y_{ijklm}) \sim N(\mu_{obs}, \theta_{obs})$$

$$\text{where} \qquad \mu_{obs} = \beta x_i + \tau_k + \mu_{sal_{m(l)}} + \alpha_{g_{m(l)}}$$

$$\text{and} \qquad \tau_k \sim N(0, \theta_k),$$

$$\mu_{sal_{m(l)}} \sim N(\mu_{sal_l}, \theta_{sal_l})$$

$$\alpha_{g_{m(l)}} \sim N(\alpha_{g_l}, \theta_{g_l})$$

with

$$
\begin{aligned}
\beta &\sim N(0, 1/10000) \\
\theta_k &= \frac{1}{\sigma_k^2} \text{ with } \sigma_k \sim \text{UNIF}(0, 10) \\
\theta_{obs} &= \frac{1}{\sigma_{obs}^2} \text{ with } \sigma_{obs} \sim \text{UNIF}(0, 10)
\end{aligned}
$$

$$
\begin{aligned}
\mu_{sal_l} &\sim N(\mu_{sal_{overall}}, \theta_{sal_{overall}}) \\
\theta_{sal_l} &= \frac{1}{\sigma_{sal_l}^2} \text{ with } \sigma_{sal_l} \sim \text{UNIF}(0, 10) \\
\mu_{sal_{overall}} &\sim N(0, 1/10000) \\
\theta_{sal_{overall}} &= \frac{1}{\sigma_{sal_{overall}}^2} \text{ with } \sigma_{sal_{overall}} \sim \text{UNIF}(0, 10)
\end{aligned}
$$

$$
\begin{aligned}
\alpha_{g_l} &\sim N(\alpha_{g_{overall}}, \theta_{g_{overall}}) \\
\theta_{g_l} &= \frac{1}{\sigma_{g_l}^2} \text{ with } \sigma_{g_l} \sim \text{UNIF}(0, 10) \\
\alpha_{g_{overall}} &\sim N(0, 1/10000) \\
\theta_{g_{overall}} &= \frac{1}{\sigma_{g_{overall}}^2} \text{ with } \sigma_{g_{overall}} \sim \text{UNIF}(0, 10)
\end{aligned}
$$

for $i = 1, 2, \ldots, 477$, $g = 1, 2$, $k = 1, 2, 3$, $l = 1, 2, \ldots, 9$, $m = 1, 2, \ldots, n_l$, with n_l being the number of departments within college l. In the model, β represents the slope on experience, τ represents the rank effect, $\mu_{sal_{m(l)}}$ represents the effect on mean salary of department within college, μ_{sal_l} represents the mean college-level salary, and $\mu_{sal_{overall}}$ represents the mean university salary. Additionally, $\alpha_{g_{m(l)}}$ represents the effect of gender on salary within a department-in-college, α_{g_l} represents the effect of gender on salary within a college, and $\alpha_{g_{overall}}$ represents the effect of gender on salary at the university level.

Multiple chains are run for each parameter, to compare and calculate the Gelman-Rubin statistic ($\hat{R}$) to aid in assessing convergence. $\hat{R}$ compares within chain variability to between chain variability to help assess convergence (Gelman and Hill, 2007, p. 358). The convergence diagnostics all indicate that chains adequately converged to a common distribution. It should be noted that this is not an exhaustive examination of model adequacy and convergence. Rather, the model fitting is preliminary, and, as a consequence, the results should be treated as such. For example, model fit was not assessed with posterior predictive checks to comparing data simulated from the model to the observed data nor were sensitivity to different starting values and prior distributions assessed. Before reporting formal conclusions from this model, these steps should be taken.

Summary of Statistical Results

No interaction term between rank and gender was included in the Bayesian model in part due to coding difficulties, though this could be included in the future. Therefore, a single estimate of an overall difference in salary between males and females is reported.

It is estimated that the median salary for a male faculty member is approximately 1.01 times the median salary of a female faculty member, after accounting for rank, experience, and department-within-college of employment (95% posterior interval from 0.95 to 1.06 times the median female faculty member salary). To put this in monetary terms, if a median salary for a male is $50,000, the median salary for a female is estimated to be $49,504 (95% posterior interval from $47,170 to $52,632). Note that this estimate of 1.01 is quite similar to the estimate of 1.018 generated from the naive multiple linear regression model with no interaction term between rank and gender. Though the Bayesian posterior interval is a bit wider, the results are the same from both analyses.

It is also of interest to examine estimates and 95% posterior intervals for multiplicative gender differences in salaries between men and women for each of the different departments. In Figure 7, these estimates and intervals are shown for many of the 46 departments on campus. Department names are omitted from the summary to protect the privacy of individuals.

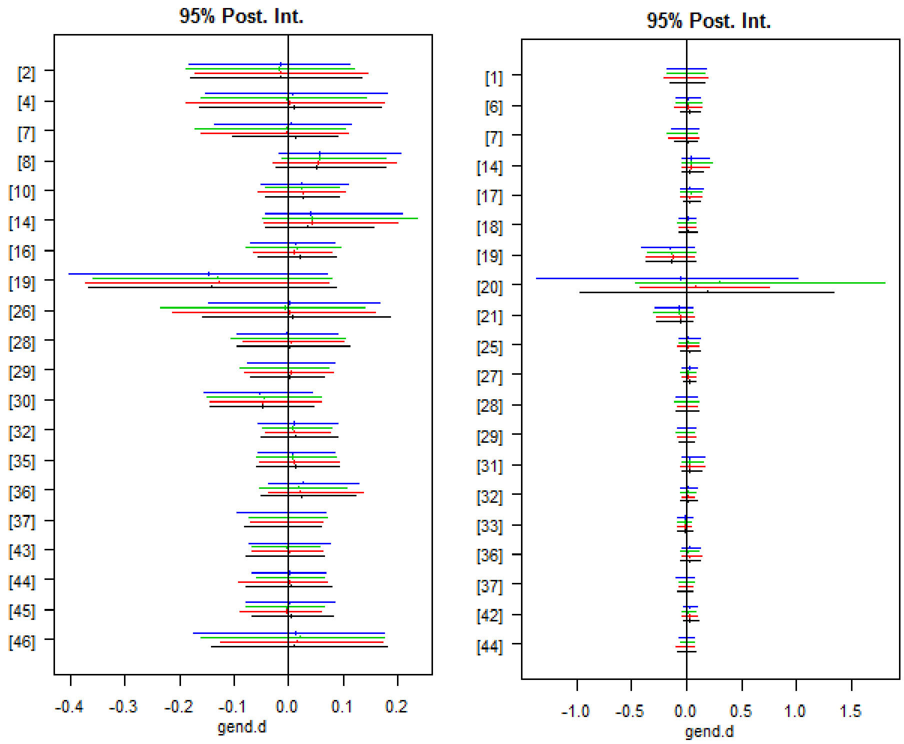

Figure 7: Plot of 95% posterior intervals of coefficients from many of the 46 departments (4 chains for each). Department names are masked with numbers due to the identifiability of some faculty members in certain departments.

Notice that all of the posterior intervals cover 0 (which, when exponentiated, is 1), meaning that in each of the departments, it is plausible that the median male salary could be the same as the median female salary. This is consistent with all the results reported thus far. There are, however, a few interesting things to note:

1. In Department 8, the posterior intervals barely cover 0. In this department, there were several male Full professors (more than ten) and only one female Full professor. Despite having the third greatest amount of experience, the female professor earned the smallest salary. A male full professor with comparable experience earned over $40,000 more. Of the other Assistant and Associate professors in Department 8, only one was female, but her salary was consistent with others of the same experience level.

2. In Department 19, which has some of the largest posterior intervals, there were very few faculty members, and only one female faculty member. However, her salary was the largest

despite having spent the fewest days at her current rank.

3. It is also of interest to investigate Department 20 because it appears that there was some trouble with convergence. When a department has few female faculty members or few faculty members overall, the model "borrows" information from the mean difference in salaries between genders from the respective college. However, with Department 20, there are many faculty members but no female faculty members. Additionally, there are several cases of salary inversion within ranks in this department. When the model borrows information from the college level to estimate the gender difference, not only does it have to contend with the relatively large number of observations in Department 20 and the pay discrepancies among faculty members, but it also must borrow information from a college where only 10% of the faculty members are female. Therefore, there is clearly a lot of uncertainty in the estimates of gender differences.

A plot with 95% posterior intervals for multiplicative gender differences in salaries between men and women in colleges can also be examined:

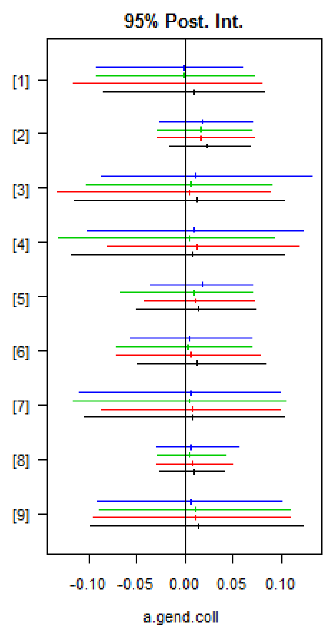

Figure 8: Plot of 95% posterior intervals of coefficients from all colleges (4 chains for each).

Again, notice that all of the posterior intervals cover 0; however, it is also of note that, with the exception of College 1 (College of Arts & Architecture), the estimates for all the other colleges are above 0. This means that though there is no statistical evidence of a difference in median college-level salaries between males and females, males are estimated to have larger median salaries than women, with the exception of one college. Additionally, it is interesting to compare the widths of the posterior intervals. Again, if there are few female faculty members in a college, the model borrows information from the overall, university-wide mean difference in salaries between men and women, with college-level estimates pulling in toward the overall mean.

Discussion

Statistical findings using each of the three models (naive linear, multilevel, Bayesian hierarchical) are approximately the same. In the first two analyses, confidence intervals on the multiplicative factors contain 1, indicating there is not enough evidence to declare that male and female faculty members truly have different median salaries assuming all other variables are fixed. Additionally, the posterior interval on the overall multiplicative estimate from the Bayesian model also contains 1, agreeing with the results from the other models. However, each of the models tells a slightly different story.

In the naive multiple regression model, the nesting of department within college was ignored. Every faculty member within a particular college was treated the same, and the gender difference in mean salaries was calculated by averaging over the inherent differences among departments. Furthermore, an interaction between gender and college was not included, thus forcing any difference in salary between men and women to be the same across all colleges.

When the hierarchical structure was accounted for in the second analysis, it was done so using random effects, allowing a different intercept for each department within college. However, the gender differences were assumed to be constant across departments within a particular college. Though this model is an improvement upon the previous one, there is still more that can be done to properly account for the nested structure while allowing for varying differences between genders across departments within a college.

Truly, it seems most appropriate to consider that each department's mean salary is related to its respective college's mean salary and that each department may have different variability in salaries between males and females than other departments in the same college. With the Bayesian hierarchical model, the mean salary of a person depends upon the mean salary of his/her department, which in turn depends on the mean salary of his/her college (in addition to also depending on rank and experience), which in turn depends on the overall mean of the university. Furthermore, it is reasonable to think that the difference in mean salaries between a male and a female faculty member depends upon the difference in mean salaries in the department of their employ, which depends on the difference in mean salaries in the college. This results in departmental estimates of gender differences shifted slightly toward the college-level estimate of a mean difference. The college-level estimate is then shifted toward the overall university estimate of a mean difference in salaries. The degree of shifting, or shrinkage, towards a higher level mean depends on the sample sizes at each department and college, as mentioned earlier.

From each of the three analyses, estimates (or a posterior distribution) were generated. However, it is paramount that a model adequately captures the structure of and uncertainty in the data. The Bayesian hierarchical model provides the most freedom for modeling the nested structure, especially for the gender differences. It allows for the inclusion of distributions on the parameters to model the data appropriately and permits us to borrow information from other departments within the same college, particularly for those departments with few, if any, female faculty members. In the end, there is no "correct" model, but a "best" model should be sought: one which represents the structure of the data, meets the assumptions, and is able to be implemented and interpreted by the user.

Acknowledgments

Many thanks to Dr. Chris Fastnow for providing this salary data set, without which, these analyses would not be possible.

I would also like to express my gratitude to Dr. Megan Higgs for her assistance in focusing this project, her patient help with coding, her thoughtful questions, her excellent editing skills, and her consistent support. I'm grateful to have had the privilege to learn from her and to have had her as an advisor.

References

- Gelman, A. and Hill, J. (2007), *Data Analysis Using Regression and Multilevel/Hierarchical Models*, Cambridge University Press.
- Pinheiro, J., Bates, D., DebRoy, S., Sarkar, D. and the R Development Core Team (2010), "nlme: Linear and Nonlinear Mixed Effects Models" (R package version 3.1-97).
- R Development Core Team (2010). *R: A language and environment for statistical computing*, Vienna: R Foundation for Statistical Computing. Available at http://www.R-project.org.
- Ramsey, F. L. and Schafer, D. W. *Statistical Sleuth: A Course in Methods of Data Analysis* (2nd ed.), Pacific Grove, CA: Duxbury, Thomson Learning.
- Robison-Cox, J. (2010), *Mixed Models and Advanced Regression* (lecture notes), Bozeman, Montana: Montana State University, Department of Mathematical Sciences.
- Spiegelhalter, D., Thomas, A., Best, N. and Lunn, D. (2003), *WinBUGS User Manual, Version 1.4* [online], Cambridge: Medical Research Council Biostatistics Unit. Available at www.mrc-bsu.cam.ac.uk/bugs.
- Sturtz, S., Ligges, U., and Gelman, A. (2005), "R2WinBUGS: A Package for Running WinBUGS," *R. Journal of Statistical Software*, 12, 3, 1-16.