

Biost 517: Applied Biostatistics I

Emerson, Fall 2011

Homework #8 Key

December 11, 2011

Written problems: To be handed in at the beginning of class on Friday, December 7, 2007.

*On this (as all homeworks) unedited Stata output is **TOTALLY** unacceptable. Instead, prepare a table of statistics gleaned from the Stata output. The table should be appropriate for inclusion in a scientific report, with all statistics rounded to a reasonable number of significant digits. (I am interested in how statistics are used to answer the scientific question.)*

Note: I have included the Stata output I used in order to answer the questions, even though you were not supposed to.

The written problems all refer to the data on MRI changes in the brains of elderly patients as stored in the project data file on the class web pages. In all problems, provide as complete statistical inference as possible (i.e., provide point estimates, confidence intervals, and p values where possible, along with a statement of your scientific/statistical conclusions).

1. Perform an analysis to compare the mean atrophy scores across groups defined by age, while allowing that each age might have a distinct average atrophy score.

```
. regress atrophy age
-----+-----
Source |         SS          df           MS       Number of obs =      735
-----+-----+-----+-----+-----+-----
Model   |    10626.648            1    10626.648       F( 1, 733) =     69.58
Residual|   111953.156          733    152.732819       Prob > F      =    0.0000
-----+-----+-----+-----+-----+-----
Total   |   122579.804          734    167.002458       R-squared     =    0.0867
                                           Adj R-squared =    0.0854
                                           Root MSE    =   12.359

-----+-----
atrophy |         Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-----+-----+-----+-----+-----+-----
age     |    .6979831    .0836783     8.34   0.000   .5337054   .8622609
_cons   |   -16.06213    6.256186    -2.57   0.010  -28.34431  -3.779947
```

- a. Provide an interpretation for the estimated intercept. What use would you make of this estimate in this scientific setting?

Ans: The estimated **mean** atrophy for newborns is -16.1, an impossible value. This group is way outside the range of our data, and thus I would make no use of that estimate. It is merely a parameter that identifies the best fitting line in the range of the data.

- b. Provide an interpretation for the estimated slope. What use would you make of this estimate in this scientific setting?

Ans: We estimate that when comparing two age groups, the mean atrophy score differs on average by 0.698 per year difference in age, with the older group tending toward higher atrophy scores. I would use this estimate to quantify the degree of association between atrophy and age. (Note that my wording avoided claiming that this data proves that atrophy increases as people age: This was a cross-sectional study.)

- c. Using the estimated regression model, what is the best estimate of the mean atrophy score for 70 year olds.

Ans: The estimated mean atrophy for 70 year olds would be $-16.06 + 70 \times 0.6980 = 32.8$.

- d. Using the estimated regression model, what is the best estimate of the mean atrophy score for 80 year olds.

Ans: The estimated mean atrophy for 80 year olds would be $-16.06 + 80 \times 0.6980 = 39.8$.
(I could have added 6.98 to the answer for part c, as well, as that is the estimated difference in mean atrophy per 10 year difference in age.)

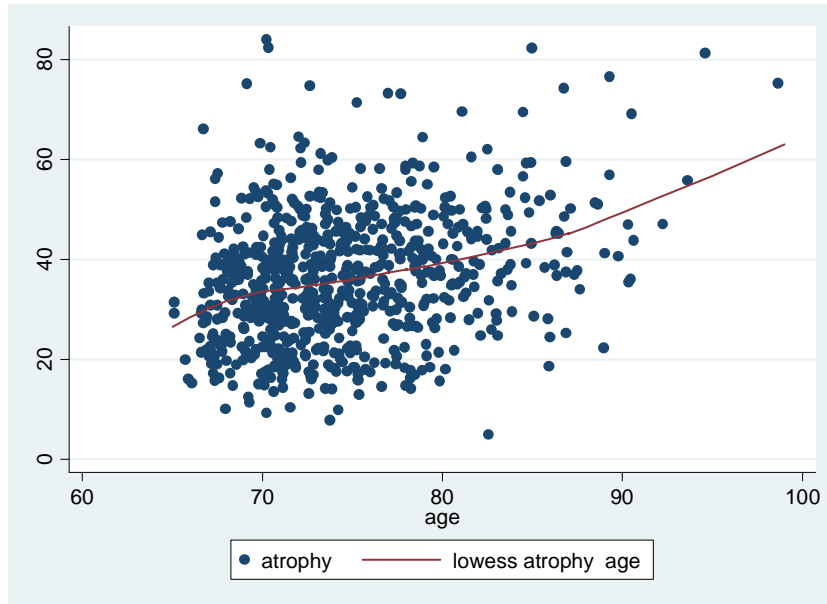
- e. Provide full inference when presuming that the variance of atrophy scores is equal across all age groups.

Ans: We estimate that when comparing two age groups, the mean atrophy score differs on average by 0.698 per year difference in age, with the older group tending toward higher atrophy scores. This result is highly unusual when there is no difference in atrophy scores by age ($P < 0.0005$). From the 95% CI, we would observe that these results were typical of situations in which the true average difference in mean atrophy were between 0.534 and 0.862 per year difference in age.

- f. Using the analysis in part e, what is your best estimate of the standard deviation of atrophy scores in each age group?

Ans: If we presume equal variances in each age group, we can use the root mean squared error, we estimate that each age group has a standard deviation of 12.4.

- g. Provide descriptive statistics that would assess (in a *post hoc* fashion) whether you believe that the estimates provided in parts c, d, and f are reliable. Explain the issues that you must consider.



Ans: In order to trust the estimated mean atrophy in individual age groups as derived from linear regression, we would need to know that the means for each age lie on a line. The lowess curve in the above graph suggests that the trend shows slight curvilinearity. Hence, some caution should be used in trusting the estimates exactly, though they would likely be in the ballpark. In order to trust a single common estimate of SD within age groups, there would need to be homoscedasticity. It is difficult to assess this exactly from the above plot due to the high number of observations in the lower age groups, but there does appear to be a little less variance in the oldest age groups. To the extent that such represents heteroscedasticity, I would be loathe to ascribe a single SD to every age group.

- h. Provide full inference when allowing that the variance of atrophy scores might be unequal across some age groups.

```
. regress atrophy age, robust
Linear regression
```

```
Number of obs = 735
F( 1, 733) = 60.12
Prob > F = 0.0000
R-squared = 0.0867
Root MSE = 12.359
```

	Coef.	Robust Std. Err.	t	P> t	[95% Conf. Interval]
atrophy					
age	.6979831	.0900192	7.75	0.000	.521257 .8747093
_cons	-16.06213	6.700595	-2.40	0.017	-29.21677 -2.907482

Ans: We estimate that when comparing two age groups, the mean atrophy score differs on average by 0.698 per year difference in age, with the older group tending toward higher atrophy scores. This result is highly unusual when there is no difference in atrophy scores by age ($P < 0.0005$). From the 95% CI, we would observe that these results were typical of situations in which the true average difference in mean atrophy were between 0.521 and 0.875 per year difference in age. (Note that the use of robust SE led to wider CI, though not markedly so.)

- i. Of the analyses considered in parts e and h, which would you prefer *a priori*.

Ans: I generally prefer to allow for the possibility that variances might be unequal across groups.

- j. Using the analysis in part h, provide an estimate and confidence interval for the difference in mean atrophy scores that might be expected between two groups that differ in age by 10 years.

Ans: We can just multiply the estimates for a one year difference in age by 10: We estimate an average difference in mean atrophy of 6.98 per 10 year difference, with a 95% CI of 5.21 to 8.75.

- k. Using the analyses performed in this problem and in problem 2, estimate the number of years difference in age that would provide the same difference in mean atrophy that is estimated for the difference in mean atrophy between men and women.

Ans: (This problem was later deleted from the assignment. At the end of the key I put some additional observations that would have led up to this question, and then I address this question..)

2. Perform an analysis to assess how the odds of being male varies across groups defined by atrophy score, while allowing that groups defined by each distinct atrophy score might have a different odds of being female. How does the inference about the association in this problem compare to the inference you provided in problem 2 (consider the P values and the comparability of the Z and t statistics)? (*Note: The Stata command `logistic male atrophy` can be used to perform logistic regression in this setting. This regression output will provide information about the ratio of the odds of being male in some atrophy group to the odds of being male in a group having an atrophy score 1 unit lower. More commonly, we would just refer to the “odds ratio associated with a 1 unit difference in atrophy scores”.*)

```
. logistic male atrophy
Logistic regression               Number of obs   =          735
                                   LR chi2(1)      =          43.53
                                   Prob > chi2     =          0.0000
Log likelihood = -487.69043        Pseudo R2      =          0.0427
```

	male	Odds Ratio	Std. Err.	z	P> z	[95% Conf. Intervals]
atrophy		1.04042	.0065502	6.29	0.000	1.027661 1.053338

- a. Provide the results of this analysis suitable for inclusion in a scientific journal.

Ans: We estimate that when comparing two atrophy groups, the odds of being male is 4.04% higher per unit difference in atrophy. This result is highly unusual when there is no difference in the distribution of sex by atrophy ($P < 0.0005$). From the 95% CI, we

would observe that these results were typical of situations in which the true odds of being male is between 2.77% higher and 5.33% higher unit difference in atrophy score. (Note how my wording here parallels the way I would discuss differences in means. Also note that this is an alternative (and less clear) way to assess an association between sex and atrophy scores—I would prefer the analyses in problem 1 in the supplemental problems below.)

- b. Using the estimated intercept and slope from your logistic regression model, provide an estimate of the odds and probability that a person with an atrophy score of 50 would be a male. (Note: The Stata command `logit male atrophy` can be used to perform logistic regression in this setting. This command will return inference on the log odds scale, rather than on the odds ratio scale.)

```
. logit male atrophy
```

```
Logistic regression                Number of obs   =          735
                                LR chi2(1)      =          43.53
                                Prob > chi2       =          0.0000
Log likelihood = -487.69043        Pseudo R2      =          0.0427
```

```
-----+-----
      male |          Coef.   Std. Err.   z    P>|z|    [95% Conf. Interval]
-----+-----
atrophy |   .0396247   .0062957   6.29   0.000   .0272854   .0519641
 _cons  |  -1.429662   .2371604  -6.03   0.000  -1.894488  -.9648362
-----+-----
```

Ans: $\log \text{odds} = -1.429662 + 50 * .0396247 = 0.551573$
 $\text{odds} = \exp(0.551573) = 1.7359816 = 1.74$
 $\text{prob} = 1.7359816 / (1 + 1.7359816) = 0.63450047 = 63.5\%$

- c. How does the inference about the association in this problem compare to the inference you would obtain from a comparison of mean atrophy scores between men and women (consider the P values and the comparability of the Z and t statistics)?

Ans: In the supplemental problem below, I use a t test to compare the mean atrophy between men and women. The t statistic is -6.6744. From the logistic regression, I obtain a Z statistic of 6.29 for the ratio for the odds of being male when comparing two groups differing by one unit in their atrophy scores. Though the resultant P values would not be exactly the same (the numerical value of the test statistics are not the same, and the t statistic would be compared to a t distribution with 733 df, while the logistic regression Z statistic would be compared to the standard normal), the close similarity of the inference will generally obtain. (If the atrophy data were normally distributed, we can show that the sex odds ratio is a function of the difference in mean atrophy.)

Additional problems that would have made problem 1k make sense:

1. Perform an analysis to compare the mean atrophy scores between men and women using the t test.
 - a. Presume that the variances would be equal for both men and women.

```
.ttest atrophy, by(male)
Two-sample t test with equal variances
```

Group	Obs	Mean	Std. Err.	Std. Dev.	[95% Conf. Interval]	
0	369	32.90515	.6363956	12.22476	31.65372	34.15658
1	366	39.08743	.6733026	12.88104	37.76339	40.41147
combined	735	35.98367	.4766699	12.92294	35.04787	36.91947
diff		-6.182283	.9262663		-8.000734	-4.363831

```
diff = mean(0) - mean(1)          t = -6.6744
Ho: diff = 0                      degrees of freedom = 733

Ha: diff < 0                      Ha: diff != 0                      Ha: diff > 0
Pr(T < t) = 0.0000                Pr(|T| > |t|) = 0.0000                Pr(T > t) = 1.0000
```

Ans: The mean atrophy for men was observed to be 39.1, while the mean atrophy for women was observed to be 32.9. This observed difference of 6.18 is highly unusual if men and women tended toward the same average degree of atrophy ($P < 0.0001$). A 95% confidence interval suggests that the observed results would not be unusual if the true average difference were such that men had atrophy scores between 4.36 and 8.00 lower than women.

- b. Allow that men and women might have different variances.

```
.ttest atrophy, by(male) unequal
Two-sample t test with unequal variances
```

Group	Obs	Mean	Std. Err.	Std. Dev.	[95% Conf. Interval]	
0	369	32.90515	.6363956	12.22476	31.65372	34.15658
1	366	39.08743	.6733026	12.88104	37.76339	40.41147
combined	735	35.98367	.4766699	12.92294	35.04787	36.91947
diff		-6.182283	.9264641		-8.001133	-4.363432

```
diff = mean(0) - mean(1)          t = -6.6730
Ho: diff = 0                      Satterthwaite's degrees of freedom = 730.335

Ha: diff < 0                      Ha: diff != 0                      Ha: diff > 0
Pr(T < t) = 0.0000                Pr(|T| > |t|) = 0.0000                Pr(T > t) = 1.0000
```

Ans: The mean atrophy for men was observed to be 39.1, while the mean atrophy for women was observed to be 32.9. This observed difference of 6.18 is highly unusual if men and women tended toward the same average degree of atrophy ($P < 0.0001$). A 95% confidence interval suggests that the observed results would not be unusual if the true average difference were such that men had atrophy scores between 4.36 and 8.00 lower than women. *(Note that in this case, there is very little difference between the inference whether I presume equal variances or I allow that variances might be unequal. This is because the estimated standard deviations are very close to each other. Furthermore, the sample sizes in each group are about the same. When either of these conditions hold, the inference for the t test presuming equal variances will be nearly the same as the inference when some allowance is made for the possibility of unequal variances. In this example, there is a very slightly higher standard error when allowing unequal variances.)*

2. Perform an analysis to compare the mean atrophy scores between men and women using linear regression.

```
. regress atrophy male
```

Source	SS	df	MS	Number of obs =	735
Model	7022.92168	1	7022.92168	F(1, 733) =	44.55
Residual	115556.882	733	157.649226	Prob > F =	0.0000
				R-squared =	0.0573
				Adj R-squared =	0.0560
Total	122579.804	734	167.002458	Root MSE =	12.556

atrophy	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]
male	6.182283	.9262663	6.67	0.000	4.363831 8.000734
_cons	32.90515	.6536311	50.34	0.000	31.62194 34.18836

- c. Provide an interpretation for the estimated intercept. How do the results of this analysis compare to your results in problem 1?

Ans: The estimated mean atrophy for women was observed to be 32.9. This agrees exactly with the sample mean for women computed in problem 1.

- d. Provide an interpretation for the estimated slope. How do the results of this analysis compare to your results in problem 1?

Ans: The estimated difference in mean atrophy scores between men and women is 6.18. This agrees exactly with the difference in sample means computed in problem 1.

- e. Provide full inference when presuming that the variances would be equal for both men and women. How do the results of this analysis compare to your results in problem 1?

Ans: The mean atrophy for women is estimated to be 32.9, with men estimated to have average atrophy scores 6.18 higher. This observed difference of 6.18 is highly unusual if men and women tended toward the same average degree of atrophy ($P < 0.0005$). A 95% confidence interval suggests that the observed results would not be unusual if the true average difference were such that men had atrophy scores between 4.36 and 8.00 lower than women. These results compare exactly with the results using the t test which presumes equal variance. *(This holds in general: Classical linear regression with a binary predictor is exactly the same as a t test which presumes equal variances.)*

- f. Provide full inference when allowing that men and women might have different variances. How do the results of this analysis compare to your results in problem 1?

```
. regress atrophy male, robust
```

Linear regression

atrophy	Coef.	Robust Std. Err.	t	P> t	[95% Conf. Interval]
male	6.182283	.9264638	6.67	0.000	4.363444 8.001122

<code>_cons</code>	32.90515	.6363992	51.71	0.000	31.65577	34.15453
--------------------	----------	----------	-------	-------	----------	----------

Ans: The mean atrophy for women is estimated to be 32.9, with men estimated to have average atrophy scores 6.18 higher. This observed difference of 6.18 is highly unusual if men and women tended toward the same average degree of atrophy ($P < 0.0005$). A 95% confidence interval suggests that the observed results would not be unusual if the true average difference were such that men had atrophy scores between 4.36 and 8.00 lower than women. These results compare approximately with the results using the t test which allows unequal variance. *(The standard error estimates agree to 5 digits, as do the limits of the CI. The difference in the standard error has to do with whether n or $n-1$ is used to calculate the sample variances. The P values and CI will differ slightly because of this difference in the SEs, as well as because different degrees of freedom are used for the t distribution: 733 for the regression with robust SE, and 730.335 when using the Satterthwaite approximation with the t test. None of these differences are material, so it is fair to regard that when regressing with a binary predictor, linear regression with robust standard errors is essentially the t test which allows unequal variances.)*

3. Perform an analysis to assess the correlation between age and atrophy scores. What is the estimated correlation? Is this estimate significantly different from 0? How does the P value from this analysis compare to the results of your analysis in problem 3?

```
. pwcorr atrophy age, sig
-----+-----+-----
      atrophy | atrophy  age
-----+-----+-----
      atrophy | 1.0000
              |
              |
      age     | 0.2944  1.0000
              | 0.0000
```

Ans: The estimated correlation between atrophy and age is 0.294, a result that is highly significant in a dataset with this sample size ($P < 0.0001$). *(This is the square root of the R^2 reported in the regression output: 0.2944 is the square root of 0.0867.)*

(Now for part k from problem 1 k.

- a. Using the analyses performed in this problem and in problem 2, estimate the number of years difference in age that would provide the same difference in mean atrophy that is estimated for the difference in mean atrophy between men and women.

Ans: The estimated difference in mean atrophy between males and females is equivalent to the estimated difference in mean atrophy per $6.182 / 0.6980 = 8.86$ year difference in age. *(Just an interesting comparison related to the frailty of males.)*