Midterm Examination #1

Econ 103, Statistics for Economists

September 29th, 2014

You will have 70 minutes to complete this exam. Graphing calculators, notes, and textbooks are not permitted.

I pledge that, in taking and preparing for this exam, I have abided by the University of Pennsylvania's Code of Academic Integrity. I am aware that any violations of the code will result in a failing grade for this course.

Name: _____

Student ID #: _____

Signature: _____

Question:	1	2	3	4	5	Total
Points:	40	15	35	20	30	140
Score:						

Instructions: Answer all questions in the space provided, continuing on the back of the page if you run out of space. Show your work for full credit but be aware that writing down irrelevant information will not gain you points. Be sure to sign the academic integrity statement above and to write your name and student ID number on *each page* in the space provided. Make sure that you have all pages of the exam before starting.

Warning: If you continue writing after we call time, even if this is only to fill in your name, twenty-five points will be deducted from your final score. In addition, two points will be deducted for each page on which you do not write your name and student ID.

1. Mark each statement as True or False. If you mark a statement as False, provide a brief explanation. If you mark a statement as True, no explanation is needed.

   (a) (4 points) In a double-blind, randomized controlled trial, neither the patients participating in the study nor the statistician analyzing the results knows who was given the placebo and who was given the real drug.

   (b) (4 points) In large populations that are approximately bell-shaped, roughly 95% of observations will lie within one standard deviation of the mean.

   (c) (4 points) The average deviation of a data set from its mean is always zero.

   (d) (4 points) If the correlation between x and y is positive, then it must be smaller than the covariance between x and y.

   (e) (4 points) The complement rule is one of the axioms of probability.

(f) (4 points) The intuition behind the addition rule is simply this: don't double-count $A \cap B$ when calculating the probability of $A \cup B$.

(g) (4 points) A random variable is neither random nor a variable: it is a fixed function.

(h) (4 points) If X is a random variable, the CDF $F(x_0)$ of X gives the probability that X exceeds a specified threshold x_0.

(i) (4 points) The support set of the Bernoulli random variable is $\{0, 1\}$.

(j) (4 points) Let X be a random variable with support set $\{-1, 0, 1\}$ and probability mass function $p(-1) = 1/2, p(0) = 1/4, p(1) = 1/4$. Then $E[X] = 0$.

2. Suppose I flip a fair coin three times. Let A be the event that I get a heads on the first toss and B be the event that I get tails on the third toss. When listing outcomes of the experiment, use the notation `[T/H][T/H][T/H]`. For example, `THT` indicates tails on the first toss, heads on the second, and tails on the third.

   (a) (3 points) How many basic outcomes are there in the sample space for this example?

   (b) (3 points) Which basic outcomes make up the event $A \cap B$?

   (c) (3 points) Which basic outcomes make up the event $A \cup B$?

   (d) (3 points) Which basic outcomes make up the event $(A \cup B) \cap (A \cap B)$?

   (e) (3 points) Calculate the conditional probability of $A \cap B$ given $A \cup B$.

3. An R dataframe called `height.data` records the annual earnings in US dollars, height in inches, and sex of 1192 individuals. In the sample, the mean earnings are \$20,400 and the mean height is 67 inches. Here are the first few rows of the dataframe:

```
  earn height    sex
 50000     74   male
 60000     66 female
 30000     64 female
 50000     63 female
 51000     63 female
  9000     64 female
```

(a) (4 points) Suppose I were to use a linear regression of the form $\widehat{y} = a + bx$ to predict `earn` from `height`. What would be the units of a? What would be the units of b?

(b) (3 points) Write out the full R command you would use to calculate a and b from the previous part using the data contained in `height.data`.

(c) (4 points) The results from the preceding part are $\widehat{y} = -60000 + 1200x$. Who would you predict will earn more: someone who is 5 feet tall or someone who is 6 feet tall? What difference in earnings would you predict for these two individuals?

Name: _____ Student ID #: _____

(d) (8 points) Suppose I were to create an R vector called `height.center`, as follows
`height.center <- height.data$height - mean(height.data$height)`
and then run a linear regression predicting `earn` from `height.center`. What would be the regression intercept? Explain your answer.

(e) (8 points) Write R code to create two dataframes: `males` contains only the observations from `height.data` for which `sex` is `male`, and `females` contains only the observations from `height.data` for which `sex` is `female`. Then use these dataframes to calculate the average height and average earnings *separately* for each group.

(f) (8 points) The results of the commands from the preceding part are as follows:

	`females`	`males`
mean `earn`	\$18000	\$30000
mean `height`	65 in	70 in

Based on all the results presented above, do you think there is a causal relationship between height and income? Why or why not? Explain briefly.

Name: _____ Student ID #: _____

4. (a) (15 points) Write a function called `tip.calculator` that calculates a restaurant tip. (Don't worry about taxes or rounding your results to the nearest cent.) Your function should take two inputs: `bill` is the restaurant bill *excluding tip* in dollars and cents, e.g. `34.50`, and `percent` is the desired tip in percentage points, e.g. `18` for 18%. Your function should return a dataframe with columns named `bill`, `percent`, `tip`, and `total`. The first two elements `bill` and `percent` are the function inputs while `tip` is the tip in dollars and cents and `total` is the total bill *including tip*. For example, if I input `45` for `bill` and `20` for `percent`, your function should return:

```
bill percent tip total
  45      20   9    54
```

(b) (5 points) After creating the `tip.calculator` function, suppose I entered the following commands at the R console:

```
x <- c(1, 10, 100)
y <- c(100, 10, 1)
tip.calculator(bill = x, percent = y)
```

Write out in full the output that R will generate from the last command, namely `tip.calculator(bill = x, percent = y)`.

5. Approximately 80% of all emails sent over the internet are spam. About 10% of spam emails contain the word "viagra" compared to 1% of non-spam emails. About 5% of spam emails contain the word "herbal" compared to 3% of non-spam.

   (a) (20 points) Assume that the occurrences of words in emails are independent for both spam and non-spam. If an email contains both the words "viagra" and "herbal" what is the probability that it is spam?

   (b) (10 points) After completing your calculations, you learn an additional piece of information: approximately 14.5% of spam emails contain the word "herbal" *or* "viagra." Does this new information support or contradict the assumption that words appear independently in emails? Explain.