Homework 1

MLRM 2009

Instructions. These questions are based mainly on material in G&H , Chapter 2. You may wish to read the Guided Tour handout for this chapter in the Statistics Handout section on the website. Answer the following questions. Show your R code, your input, and your output. Feel free to ask me for hints if you get stumped.

1. Review the definition of the normal curve functions `pnorm`, `qnorm`, and `dnorm` in R. Review the syntax for defining a function in R. Then define a function `inorm(a,b,mu,sigma)`. This function should compute the probability of obtaining a result between `a` and `b` in a normal distribution with mean `mu` and standard deviation `sigma`. This function should work regardless of the ordering of `a` and `b`. (Hint, use the `abs` function.) It should also use the standard normal curve (i.e., `mu=0, sigma=1`) by default if `mu` and `sigma` are not specified.

   To verify the correctness of your function, use it to demonstrate the following results:

   (a)
   ```
   > inorm(0,1)
   [1] 0.3413447
   ```
   (b)
   ```
   > inorm(1,0)
   [1] 0.3413447
   ```
   (c)
   ```
   > inorm(2,1)
   [1] 0.1359051
   ```
   (d)
   ```
   > inorm(130,145,100,15)
   [1] 0.02140023
   ```
   (e)
   ```
   > inorm(350,500,500,100)
   [1] 0.4331928
   ```

2. Create a normal curve function that computes the *upper tail probability* for a particular value. This is the probability of obtaining a value larger than or equal to x in a normal distribution with mean μ and standard deviation σ. This function will have a call of the form `unorm(x,mu,sigma)`. Demonstrate that your function works by producing the following results

   (a)
   ```
   > unorm(3)
   [1] 0.001349898
   ```
   (b)
   ```
   > unorm(2)
   [1] 0.02275013
   ```
   (c)
   ```
   > unorm(3,2,2)
   [1] 0.3085375
   ```
   (d)
   ```
   > unorm(145,100,15)
   [1] 0.001349898
   ```

3. *Binomial Distribution.* John believes the long run probability that he will beat Phil in chess is .8. If this is true, the probability remains constant, and games are independent, what is the probability that, if John and Phil play 10 games, *Phil* will win *more than 3*?

4. (a) Suppose two normal distributions have possibly different means and standard deviations. Write a function with a call of the form

   `tail.odds(x,mu.a,sigma.a,mu.b,sigma.b)`

   to compute the *tail odds.* This is the probability of exceeding a value x in distribution A divided by the probability of exceeding x in distribution B. (Hint: Use your `unorm` function). Test your function on the following example.

   Example. Suppose a cutoff value of 780 is needed to qualify for a position. Group A has a mean of 500 and a standard deviation of 118. Group B has a mean of 500 and a standard deviation of 100. What are the odds that a person who qualifies will be in Group A (as opposed to group B)?

Answer.

```
> tail.odds(780,500,118,500,100)
[1] 3.453822
```

(b) Use the `curve` function in R to plot the tail odds in the above situation for the range $700 \leq x \leq 900$. What do you see? Keeping in mind that the two groups have identical means, discuss some implications of this result.

5. *Poisson Distribution.* Your factory produces large batches of electronics components. Component Z-42 has an excellent track record. Over the last 5 years, failures have occurred at the rate of 4.55 items per batch. What is the probability that the next batch you produce will have 8 or more failures?