

1. Guided R exercise. Normal probabilities can be difficult to compute because the normal probability density function cannot be integrated by hand. R uses numerical integration techniques to compute these probabilities. Use the R function `pnorm` to compute  $P(64 \leq X \leq 120)$  if  $X \sim N(\mu = 88, \sigma = 18)$ . To see some information on `pnorm` open R and run the line `?pnorm` to bring up some helpful documentation. Once we have tools to calculate normal probabilities, we would like to use these tools for everything, mainly because we are lazy and we don't want to have to come up with other tools. This is one reason why we may want to approximate, for instance, binomial probabilities using the normal family of distributions. Below, I have codes to plot the histogram of samples from various binomial distributions. Overlaid on top of each plot is the corresponding normal density (i.e. the normal density with the same mean and variance as the binomial). Run my codes and comment on each of the four graphs. Which approximations work and which don't? What kind of errors do you see in the approximations? Use your observations to develop a rule for when the approximation will be good and when it will be bad. Here is a hint: Keep in mind that the normal distribution is on all of  $\mathbb{R}$  whereas a binomial random variable is always positive. Codes (you may have to edit apostrophes if you copy/paste this code):

```
# Binomial n = 10 q = 0.5, Normal overlaid with mu = 5, sigma^2 = 2.5
hist(rbinom(1000, 10, .5), freq = FALSE)
points(seq(from = -2, to = 12, by = 0.01),
```

```
dnorm(seq(from = -2, to = 12, by = 0.01), 5,  
sd = sqrt(2.5)), add = TRUE, pch = '.')
```

```
# Binomial n = 10 q = 0.1, Normal overlaid with mu = 1, sigma^2 = .9  
hist(rbinom(1000, 10, .1), freq = FALSE)  
points(seq(from = -2, to = 3, by = 0.01),  
dnorm(seq(from = -2, to = 3, by = 0.01), 1,  
sd = sqrt(.9)), add = TRUE, pch = '.')
```

```
# Binomial n = 100 q = 0.5, Normal overlaid with mu = 50, sigma^2 = 25  
hist(rbinom(1000, 100, .5), freq = FALSE, breaks = 12)  
points(seq(from = -20, to = 120, by = 0.01),  
dnorm(seq(from = -20, to = 120, by = 0.01), 50,  
sd = sqrt(25)), add = TRUE, pch = '.')
```

```
# Binomial n = 100 q = 0.1, Normal overlaid with mu = 10, sigma^2 = 9  
hist(rbinom(1000, 100, .1), freq = FALSE, breaks = 12)  
points(seq(from = -20, to = 30, by = 0.01),  
dnorm(seq(from = -20, to = 30, by = 0.01), 10,  
sd = sqrt(9)), add = TRUE, pch = '.')
```

2. Exercise 3.6 in the text.

3. Exercise 3.28 in the text.
4. Exercise 4.12 in the text.
5. Exercise 6.18 in the text.
6. Exercise 6.32 in the text. Use both the binomial distribution and the normal approximation to compute the probability and compare your two results.