

STAT 6200 — Design of Experiments for Research Workers
Lab 5 – Due Thursday, November 18

Example — Weight Loss

In a dietary study, 14 of 20 subjects lost weight. If weight is assumed to fluctuate up or down by chance, then the probability of losing weight would be $p = 1/2$.

1. Test whether the diet was effective in the sense that it resulted in more people losing weight than would have occurred by chance alone.

Answer: We are interested in testing

$$H_0 : p = p_0 \quad \text{versus} \quad H_A : p > p_0, \quad \text{where } p_0 = .5.$$

The normal approximation to the binomial can be used to test this hypothesis using a z test. However, before considering the z test, let's think about how this test can be done exactly, based on the binomial distribution.

Let X = the number of people who lose weight, out of $n = 20$. We observed $X = 14$ for a sample proportion of $\hat{p} = x/n = 14/20 = .7$.

As usual, we decide whether the population parameter p is equal to the null value $p_0 = .5$ by looking at how far our estimate of the population parameter $\hat{p} = .7$ is from the null value of $.5$.

Equivalently, we can look at how far

$$n\hat{p} = 20(.7) = 14 = X$$

is from

$$np_0 = 20(.5) = 10 = E(X|H_0 : p = p_0 \text{ is true}).$$

- So, to decide whether to reject H_0 , we look at the strength of the evidence against H_0 (the p -value) provided by the fact that $\hat{p} = .7$ is greater than $p_0 = .5$, or, equivalently, provided by the fact that we observed $X = 14$ successes (people losing weight) when we only expected 10 under H_0 .

That is, the p -value for our test is the probability of getting $X = 14$ successes under the null hypothesis that $X \sim \text{Bin}(n, p_0) = \text{Bin}(20, .5)$.

We know how to calculate such a probability:

$$\begin{aligned} p = P(X \geq 14) &= 1 - P(X < 14) = 1 - \left\{ \binom{20}{0}.5^0(1-.5)^{20-0} + \binom{20}{1}.5^1(1-.5)^{20-1} \right. \\ &\quad \left. + \cdots + \binom{20}{13}.5^{13}(1-.5)^{20-13} \right\} \\ &= 1 - .9423 = .0577 \end{aligned}$$

- Since $p = .0577 > \alpha = .05$, we would not reject H_0 . There is marginal evidence that the diet is effective, but the result does not quite reach significance.

- The value .9423 in the last calculation, the cumulative probability of getting 13 or fewer successes from a $Bin(20, .5)$ distribution, can be obtained from Minitab. Select

Calc → Probability Distributions → Binomial...

and then click “Cumulative probability”, set “Number of trials:” to 20, “Probability of success:” to .5, and “Input constant:” to 13”. Then click OK.

Now consider how we would use the normal approximation to estimate this p -value. Recall that the normal distribution with the same mean and variance as the binomial we want to approximate is used.

That is, for large enough sample size, the $Bin(n, p)$ distribution is well approximated by the $N(np, np(1 - p))$ distribution.

- That is, $X \sim Bin(20, .5)$ has about the same distribution as Y where $Y \sim N(20(.5), 20(.5)(1 - .5)) = N(10, 5)$.

- Therefore, our p -value is still $p = P(X \geq 14)$ where $X \sim Bin(n, p_0) = Bin(20, .5)$, but we approximate this p -value as

$$\begin{aligned} p &= P(X \geq 14) \approx P(Y \geq 14) \quad \text{where } Y \sim N(np_0, np_0(1 - p_0)) = N(10, 5) \\ &= P\left(\frac{Y - np_0}{\sqrt{np_0(1 - p_0)}} \geq \frac{14 - np_0}{\sqrt{np_0(1 - p_0)}}\right) \\ &= P\left(Z \geq \frac{14 - 10}{\sqrt{5}}\right) \quad \text{where } Z \sim N(0, 1) \\ &= P(Z \geq 1.7889) = P(Z \leq -1.7889) = .0368 \end{aligned}$$

- This last probability can be obtained in Minitab by selecting

Calc → Probability Distributions → Normal...

and then click “Cumulative probability”, set “Mean:” to 0, “Standard deviation:” to 1, and “Input constant:” to -1.7889”. Then click OK.

- Note that this normal-approximation-based p -value is slightly in error (the exactly correct value was $p = .0577$, computed above). In fact, we go from a non-significant result to a significant result as a result of the error in approximation.

- Also note that the normal-approximation-based p -value can also be obtained more directly in Minitab. Here are the steps: select

Stat → Basic Statistics → 1 Proportion...

and then click “Summarized data”, set “Number of trials:” to 20, and “Number of events:” to 14. Then click “Options...”, set “Confidence level:” to 95.0, “Test proportion” to .5 (this is the value of p_0), “Alternative:” to “greater than”, and place a check next to “Use test and interval based on normal distribution”. Then click OK twice, and you’ll get the p -value we just obtained: $p = .037$.

The normal approximation to the binomial can be improved by using what is known as a **continuity correction**. This correction adjusts for the fact that we are approximating the binomial, a discrete distribution, with the normal, a continuous distribution.

To understand the continuity correction, recall that the normal p.d.f. doesn’t give the probability of observing any single value (that probability is 0, for a continuous distribution like the normal). Instead it gives the probability associated with a range of values.

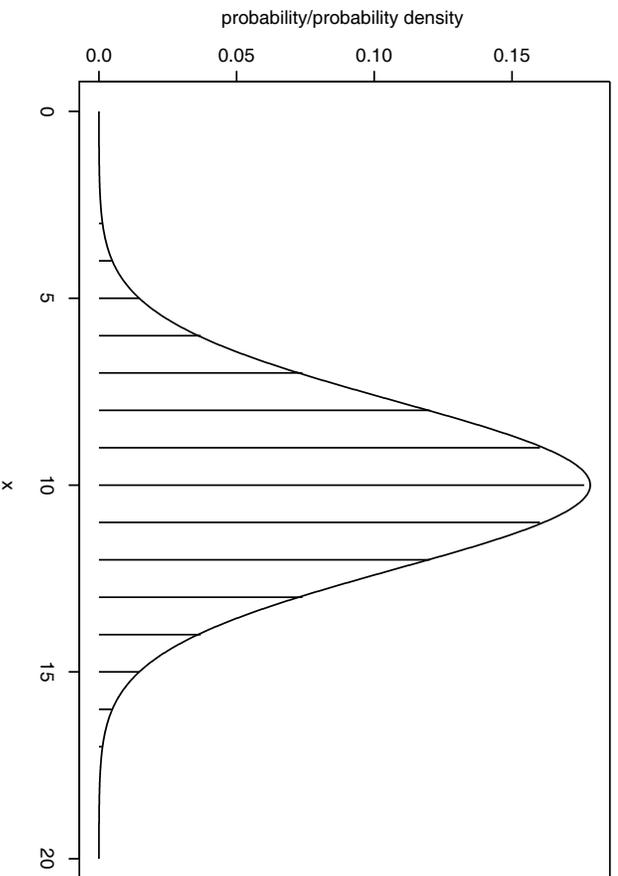
So, for instance, to estimate the probability of getting exactly $X = 14$ successes from a $Bin(n = 20, p_0 = .5)$ distribution, we would not use

$$P(X = 14) \approx P(Y = 14), \quad \text{where } Y \sim N(np_0, np_0(1 - p_0)) = N(10, 5)$$

because $P(Y = 14) = 0$.

Below is the $Bin(n = 20, p_0 = .5)$ probability function with the $N(np_0 = 10, np_0(1 - p_0) = 5)$ p.d.f. superimposed on top.

Bin(20,.5) and N(10,5) distributions



- In this plot, the vertical lines are at $0, 1, \dots, 20$, the only possible values that X can take on. Each line has height equal to the probability of that value according to the $Bin(20, .5)$ distribution.
- The smooth bell-curve is the $N(10, 5)$ distribution. Obviously, this distribution follows the binomial probabilities closely.

The best normal approximation to $P(X = 14)$, say, is not $P(Y = 14)$ but instead

$$P(14 - \frac{1}{2} \leq Y \leq 14 + \frac{1}{2}) = P(13.5 \leq Y \leq 14.5)$$

Similarly, if we want to approximate our p value, which was given by $P(X \geq 14)$, with the normal distribution, it is best to use

$$\begin{aligned} P(X \geq 14) &= P(X = 14) + P(X = 15) + \dots \\ &\approx P(13.5 \leq Y \leq 14.5) + P(14.5 \leq Y \leq 15.5) + \dots \\ &= P(Y \geq 13.5) \\ &= P(Z \geq \frac{13.5 - np_0}{\sqrt{np_0(1 - p_0)}}) = P(Z \geq \frac{13.5 - 10}{\sqrt{5}}) \\ &= P(Z \geq 1.5652) = P(Z \leq -1.5652) = .0588 \end{aligned}$$

- The last calculation above was done in Minitab.
- Notice, that with this continuity correction, our approximate normal-based p -value of .0588 is much closer to the true value of $p = .0577$ we calculated directly from the binomial distribution.
- Here the continuity correction for $P(X \geq x)$ involved subtracting $\frac{1}{2}$ from x . That is, we used $P(X \geq x) \approx P(Y \geq x - \frac{1}{2})$. Note that if we had wanted $P(X \leq x)$ we would have added $\frac{1}{2}$. That is, we would have used $P(X \leq x) \approx P(X \leq x + \frac{1}{2})$.
- Continuity corrections generally improve the normal approximation to the binomial. However, their effect becomes negligible as the sample size n grows.
- Minitab implements the normal approximation **without** the continuity correction.
- Finally, note that for two-sided alternatives, the p -value is twice the one-sided p -value (unless this value turns out to be ≥ 1 , in which case the p -value is rounded down to 1), using either the normal approximation or the exact binomial approach.

2. Now form a 95% confidence interval for p , the probability of losing weight on the diet.

Answer: In the case of a confidence interval or one-sided confidence limits, it is also possible to get an exact answer using the binomial distribution. However, exactly how this is done is somewhat complicated, so we will show how to get the exact answer with Minitab and not discuss the computational details at all. We will discuss the normal approximation approach.

An approximate (normal-based) $100(1 - \alpha)\%$ CI for p is given by

$$\hat{p} \pm z_{1-\alpha/2} \sqrt{\hat{p}(1 - \hat{p})/n}$$

In our case, we want a 95% interval, so $\alpha = .05$ and $z_{1-\alpha/2} = z_{.975} = 1.96$. In addition, $\hat{p} = .7$ so our 95% CI is

$$.7 \pm 1.96 \sqrt{.7(1 - .7)/20} = (.499, .901)$$

- We can obtain this result using Minitab through the following steps: select
Stat → Basic Statistics → 1 Proportion...

and then click “Summarized data”, set “Number of trials:” to 20, and “Number of events:” to 14. Then click “Options..”, set “Confidence level:” to 95.0, “Test proportion” to .5 (this is not necessary for a confidence interval), “Alternative:” to “not equal” (for a two-sided confidence interval, rather than a one-sided confidence bound), and place a check next to “Use test and interval based on normal distribution”. Then click OK twice, and you’ll get the CI we just obtained: (.499, .901).

- To get the exact answer, just repeat the previous steps, but do not place a check next to “Use test and interval based on normal distribution”. The resulting exact 95% CI is (.457, .881).

Note that it is possible to improve on the approximate normal-based CI we just computed by using a continuity correction. With a continuity correction, a (slightly better) $100(1 - \alpha)\%$ CI for p is given by

$$\hat{p} \pm \left[z_{1-\alpha/2} \sqrt{\hat{p}(1 - \hat{p})/n} + \frac{1}{2n} \right]$$

- In our problem the continuity-corrected normal-based 95% interval is

$$.7 \pm \left[1.96 \sqrt{.7(.3)/20} + \frac{1}{2(20)} \right] = (.474, .926)$$

This turns out to not be much better than the non-continuity corrected interval in this problem.

- Again, Minitab does not implement the continuity correction.

Exercise:

3. Ounsted (1953) presents data about cases with convulsive disorders. Among the cases there were 82 females and 118 males.
- Compute the p value for the test of the hypothesis that a case is equally likely to be of either sex using exact methods.
 - Compute the p value for the test of the hypothesis that a case is equally likely to be of either sex using the normal approximation to the binomial without a continuity correction.
 - Compute the p value for the test of the hypothesis that a case is equally likely to be of either sex using the normal approximation to the binomial with a continuity correction.
 - Compare your answers in parts a, b, c.
 - Obtain a 99% CI for p , the population proportion of convulsive cases that are male using exact methods.
 - Obtain a 99% CI for p , the population proportion of convulsive cases that are male using the normal approximation without a continuity correction.
 - Obtain a 99% CI for p , the population proportion of convulsive cases that are male using the normal approximation with a continuity correction.
 - Compare your answers in parts e, f, g.
- For exams in this course, I will not expect you to know how to implement the continuity correction, but at a minimum, I want you to have seen it, and to know that it exists.