Using Sample Data to Draw Conclusions about Population Means

Your instructor has a can containing several hundred pennies. The collection of pennies in the can is the population that we will pretend to be interested in today. Our goal is to learn what we can about the mean age of all the pennies in the can by observing only a simple random sample of 10 pennies. In a real problem, our population of interest might be all geese in North America, all trees in Alaska, all farmers in Iowa, etc. We will use the population of pennies in the can to illustrate many important statistical concepts because this population is much easier to work with in the classroom than many other more interesting populations.

1. Draw a simple random sample of 10 pennies from the can and list the mint year for each of the 10 pennies below.

2. Convert the years above into ages by taking the current year minus the mint year.

3. Compute $\bar{Y}$ the average age of the sampled pennies. We use the statistic $\bar{Y}$ to estimate $\mu =$ the mean age of all the pennies in the can.

4. One of the interesting and useful fact is that we can use information in the sample to obtain information about how far our statistic ($\bar{Y}$) is likely to be from the population parameter (μ). We begin by estimating the standard deviation of the ages of all the pennies in the can (σ) by the sample standard deviation (s). Compute s, the standard deviation of the 10 pennies' ages.

5. Statistical theory dictates that the standard deviation of a sample average is $\sigma/\sqrt{n}$, where σ denotes the population standard deviation and n denotes the number of observations drawn randomly from the population of interest. The estimated standard deviation of a statistic is called its **standard error**. Give the formula for the standard error of a sample average.

6. Compute the standard error for $\bar{Y}$ in this example. This value is an estimate of the "typical" size of the error made when estimating μ by $\bar{Y}$.

7. The quantity
$$t = \frac{\bar{Y} - \mu}{s/\sqrt{n}}$$
is called a t-ratio. Note that a t-ratio will vary from sample to sample because it depends on statistics ($\bar{Y}$ and s) that change from sample to sample. Statistical theory dictates that under fairly general conditions discussed in Chapter 3 the distribution of t-ratio values will be approximated by a Student's t-distribution with $n-1$ degrees of freedom. Table A.2 in the back of your text gives information about the proportion of observations from a Student's t-distribution that will be less than various values. (Usually we will just say t-distribution rather than Student's t-distribution.)

   (a) What proportion of observations from a t-distribution with 1 degree of freedom will be less than 6.314?

   (b) What proportion of observations from a t-distribution with 20 degrees of freedom will be less than 1.064?

(c) What proportion of observations from a t-distribution with 20 degrees of freedom will be less than -1.725?

(d) What proportion of observations from a t-distribution with 20 degrees of freedom will be larger in magnitude than 1.725?

8. Based on our observed sample and the information in the previous problem, do you think the average age of all the pennies in the can is 5 years? In other words, does $\mu = 5$ seem plausible?

9. Formally we can test the null hypothesis $H_0 : \mu = 5$ against the alternative $H_A : \mu > 5$.

(a) Our **test statistic** is the t-ratio that we get by plugging the value of μ according to the null hypothesis in for μ in the t-ratio equation. Compute the test statistic.

(b) Our p-**value** is the probability – computed assuming the null hypothesis to be true – of obtaining a t-ratio that provides at least as much evidence in favor of the alternative hypothesis as the test statistic that we computed. Compute a p-value for this problem.

(c) Is there evidence that the average age of all the pennies is greater than 5 years? (See information about interpreting p-values on page 47.)

(d) We have just carried out what is known as a **one-sided hypothesis test**. A **two-sided hypothesis test** is concerned about alternatives on both sides of the null hypothesis ($H_A : \mu \neq 5$). The p-value for a two-sided test is just twice the p-value of a one-sided test. Write the p-value for the two-sided test below.

10. Rather than checking many individual values for μ to see if they seem plausible, we can compute a confidence interval for μ. A confidence interval directly provides a range of plausible values for μ. A $100(1 - \alpha)\%$ confidence interval for a population mean μ is given by

$$\bar{Y} \pm t_{n-1}^{(1-\alpha/2)} \frac{s}{\sqrt{n}}, \text{ where } t_{n-1}^{(1-\alpha/2)}$$

is the t-ratio larger than $100(1 - \alpha/2)\%$ of t-ratios from a t-distribution with $n - 1$ degrees of freedom. Compute a 95% confidence interval for μ.

11. The percentage attached to a confidence interval indicates the percentage of samples that will give a confidence interval that contains the true value of μ. In other words, when your class drew a random sample of 10 pennies from the can, there was a 95% chance that the sample selected would lead to an interval in question 10 that contains the population mean μ=the average age of all the pennies in the can. Your instructor knows μ because he computed it by examining all the pennies in the can. Was your sample one of the 95% that produces a good confidence interval, or did you get unlucky?