

Using Binary Variables to Estimate Proportions

When you compute the *mean* of a *binary* variable, it is mathematically equivalent to a proportion. Since all the x_i are equal to 0 or 1, summing them all the observations together is the same as counting the number x_i that equal to 1. The following formula for the sample mean reveals it is equal to the sample proportion

$$\bar{x} = \frac{\sum_{i=1}^n x_i}{n} = \frac{n(x_i = 1)}{n},$$

where $n(x_i = 1)$ denotes the count for the observations in the sample that are equal to 1, and the ratio $n(x_i = 1)/n$ is exactly the proportion of a sample that is equal to 1.

Similarly, the *average value from the population for a binary variable* is the proportion of times the binary variable is equal to 1 or the *probability* that $x_i = 1$.

To estimate a probability of an event occurring, one can use standard statistical methods using a *binary outcome variable*.

Example 1: suppose a national poll is used to measure whether or not someone is voting for the republican candidate for president. A binary variable is set equal to 1 if the respondent said ‘yes’, they do intend to vote for the republican candidate, and set to 0 if they intent to vote for someone else. If one wishes to estimate the probability that the Republican candidate will win with more than 50% of the vote, we can use a simple single sample t-test for a mean with the following hypotheses:

$$H_0 : \mu = 0.5$$

$$H_A : \mu > 0.5$$

Example 2: The data set, `loanapp.RData`, includes actual data from 1,989 mortgage loan applications, including whether or not a loan was approved, and a number of possible explanatory variables including variables related to the applicant’s ability to pay the loan such as the applicant’s income and employment information, value of the mortgaged property, and credit history. Also included in the data set are variables measuring the applicant’s race and ethnicity.

The code below loads the R data set, which creates a data set called `data`, and a list of descriptions for the variables called `desc`.

```
download.file("http://murraylax.org/datasets/loanapp.RData", "loanapp.RData")
load("loanapp.RData")
```

Let us estimate the proportion of loans that were approved. The variable `approve` is a binary variable equal to 1 if the loan was approved and 0 otherwise. We simply compute the mean of approved:

```
mean(data$approve)
```

```
## [1] 0.8773253
```

We can see that 87% of loan applications are approved in our sample.

Suppose someone claims that more than 85% of all mortgage loan applications are approved. We can test this claim with a single-sample hypothesis test. The null and alternative hypotheses are given by:

$$H_0 : \mu = 0.85$$

$$H_A : \mu > 0.85$$

We conduct the test with the *t.test* function:

```
t.test(data$approve, mu=0.85, alternative="greater")

##
## One Sample t-test
##
## data: data$approve
## t = 3.7138, df = 1988, p-value = 0.0001049
## alternative hypothesis: true mean is greater than 0.85
## 95 percent confidence interval:
##  0.8652171      Inf
## sample estimates:
## mean of x
## 0.8773253
```

The p-value is equal to 0.0001049 which is less than 0.05. We reject the null hypothesis and conclude we have sufficient statistical evidence that more than 85% of loan applications are approved.

We can also compute a 95% confidence interval for the proportion of loans approved:

```
loantest <- t.test(data$approve, conf.level=0.05)
loantest$conf.int
```

```
## [1] 0.8768638 0.8777867
## attr(,"conf.level")
## [1] 0.05
```

With 95% confidence we can report an interval estimate for the average approval rate between 87.69% to 87.78%.

Example 3: Let us continue using the loan approval data set and make comparisons across different groups.

Suppose we wish to estimate whether the average approval rate is lower for self-employed people versus people not self employed. The variable, `self` is a binary variable equal to 1 if the person in the observation is self employed and zero otherwise.

Our null and alternative hypotheses are given by,

$$H_0 : \mu_0 - \mu_1 = 0$$

$$H_A : \mu_0 - \mu_1 > 0$$

where μ_1 is the probability a self-employed person (`self=1`) is approved a loan and μ_0 is the probability someone who is not self-employed (`self=0`) is approved for a loan.

The alternative hypothesis has a *greater-than* sign, because our research question asks whether group 1 has a smaller proportion than group 0. Since we are subtracting the smaller proportion, the difference will be *greater than* 0.

The following call to `t.test` runs an independent-samples t-test for a difference in means depending on the value for `self`:

```
t.test(approve ~ self, data=data, alternative="greater")
```

```
##
## Welch Two Sample t-test
##
```

```
## data: approve by self
## t = 1.5811, df = 318.73, p-value = 0.05743
## alternative hypothesis: true difference in means is greater than 0
## 95 percent confidence interval:
## -0.001642264      Inf
## sample estimates:
## mean in group 0 mean in group 1
##      0.8822171      0.8443580
```

Our sample evidence shows that 88.2% of people who are not self-employed are approved for a mortgage loan and 84.4% of self-employed people are approved.

The p-value for the hypothesis test is equal to 0.0573. This is greater than 0.05, but less than 0.1. At the 10% significance level we can reject the null hypothesis. With 90% confidence, we can conclude that self-employed people on average are approved less often for mortgages than others.

We can compute a 95% confidence interval for the difference in approval rates:

```
loantest <- t.test(approve ~ self, data=data, conf.level=0.95)
loantest$conf.int
```

```
## [1] -0.009251764  0.084969991
## attr(,"conf.level")
## [1] 0.95
```

With 95% confidence, the difference in approval rate for non-self-employed people versus self-employed people is between -0.9% and 8.5%