A Brief Review of Statistical Inference

Ryan Miller

Statistical Inference

- A model is a simplified representation intended to aide in explaining or predicting a phenomenon
  - A **statistical model** is one that involves a *probability distribution*

Statistical Inference

- A model is a simplified representation intended to aide in explaining or predicting a phenomenon
  - A **statistical model** is one that involves a *probability distribution*
- You've without a doubt seen statistical models in your Intro Stats class
  - For example, representing the sampling distribution of a differences in proportions, $\hat{p}_1 - \hat{p}_2$, using a Normal curve is an application of a statistical model

Sources of Uncertainty

- Statistical methods are used to quantify uncertainty
    - One classical example is the uncertainty introduced when taking a sample from a population

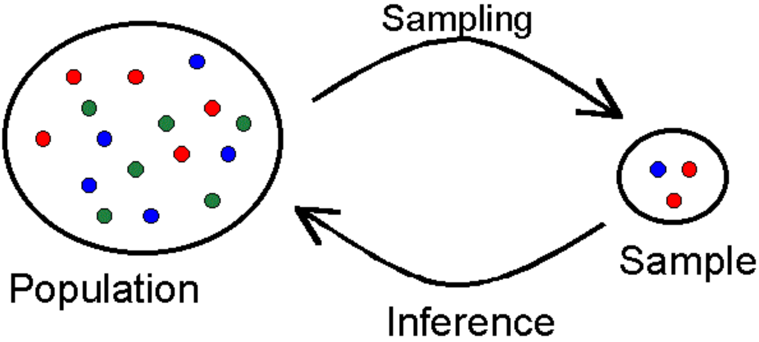

Sources of Uncertainty

- Another classical example is the uncertainty introduced via random assignment

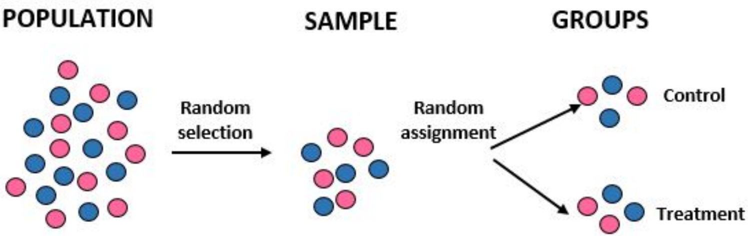

Statistical Inference

You should already be familiar with *two broad categories* of **statistical inference**

Statistical Inference

You should already be familiar with *two broad categories* of **statistical inference**

1. Confidence Interval Estimation
2. Hypothesis Testing

Confidence Intervals

- **Definition**: "Were this procedure to be repeated on numerous samples, the fraction of calculated confidence intervals (which would differ for each sample) that encompass the true population parameter would tend toward 90%."
- **Practical Interpretation**: "The values in this interval represent a statistically meaningful range of plausible values for the true population parameter we are interested in."

p-values

- **Definition**: "Under the assumption that a particular null model is correct, the *p*-value is the probability of obtaining an outcome that is at least as extreme as the results that were actually observed."
- **Practical Interpretation**: "A small *p*-value indicates the observed data are statistically incompatible with the null hypothesis"

Example (classic)

- In my Math-156 class, I discuss a study where 10-month old infants watched a series of video demonstrations involving a "climber" character attempting to reach the top of a hill
  - In some demonstrations, a "helper" character came in and assisted the "climber" in reaching the top
  - In others, a "hinderer" character came in and prevented the "climber" from reaching the top

Example (classic)

- In my Math-156 class, I discuss a study where 10-month old infants watched a series of video demonstrations involving a "climber" character attempting to reach the top of a hill
  - In some demonstrations, a "helper" character came in and assisted the "climber" in reaching the top
  - In others, a "hinderer" character came in and prevented the "climber" from reaching the top
- After watching each scenario several times, the infants got to choose between the "helper" and "hinderer" as a play toy
  - The result was 14 of 16 choosing the "helper", but is this statistically convincing evidence that the infants truly had a preference?

Example (classic)

1) Null hypothesis: infants were choosing randomly -> Null model: $\hat{p} \sim \text{Binomial}(16, 0.5)$

Example (classic)

1) Null hypothesis: infants were choosing randomly -> Null model: $\hat{p} \sim \text{Binomial}(16, 0.5)$
2) p-value $= 2*Pr(\hat{p} \geq 14/16) = 0.004$ (calculated using the null model)
3) Decision: the data provide compelling evidence that infants have a preference towards the "helper" toy

```r
binom.test(14, 16, .5)
```

```
##
##  Exact binomial test
##
## data:  14 and 16
## number of successes = 14, number of trials = 16, p-value = 0.004181
## alternative hypothesis: true probability of success is not equal to 0.5
## 95 percent confidence interval:
##  0.6165238 0.9844864
## sample estimates:
## probability of success
##                  0.875
```

Example (modeling)

- A professor collected data on homes sold in Iowa City, IA between 2005 and 2008
- We can use **simple linear regression** to model the relationship a home's assessed value and its actual sale price
  - A *statistical model* is useful here if we are interested in generalizing beyond this specific set of houses

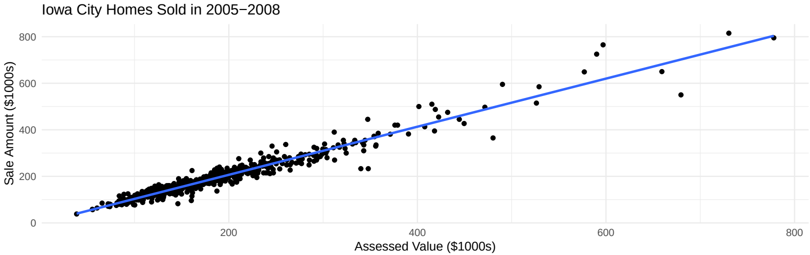

Example (cont.)

- The estimated slope in this relationship is 1.033, what does this suggest? What would a slope of 1 suggest?

Example (cont.)

- The estimated slope in this relationship is 1.033, what does this suggest? What would a slope of 1 suggest?
- The **standard error** (a measure of an estimate's uncertainty) of the slope is 0.009
  - Is it plausible that the *true slope* describing *all houses in Iowa City* is 1?

Example (cont.)

- The estimated slope in this relationship is 1.033, what does this suggest? What would a slope of 1 suggest?
- The **standard error** (a measure of an estimate's uncertainty) of the slope is 0.009
  - Is it plausible that the *true slope* describing *all houses in Iowa City* is 1?

```
## A calculating a simple 95% CI
c(1.033 - 2*0.009, 1.033 + 2*0.009)
```

```
## [1] 1.015 1.051
```

Example (cont.)

For your reference, the summary function provides numerous details pertaining to statistical inference:

```
mod <- lm(sale.amount ~ assessed, data = IC)
summary(mod)
```

```
##
## Call:
## lm(formula = sale.amount ~ assessed, data = IC)
##
## Residuals:
##     Min      1Q  Median      3Q     Max
## -152050   -7137    -347    7496  148286
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept) -1.523e+00  1.712e+03  -0.001    0.999
## assessed     1.033e+00  8.819e-03 117.142   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 20970 on 775 degrees of freedom
## Multiple R-squared:  0.9465, Adjusted R-squared:  0.9465
## F-statistic: 1.372e+04 on 1 and 775 DF,  p-value: < 2.2e-16
```

- Each row tests the hypothesis that the parameter is zero, what does the small p-value in the "assessed" row imply?

Closing Remarks

Below is brief overview of some common situations and the R functions you can use for statistical inference:

Explanatory	Response	Statistical Test	R Function
Categorical	Categorical	Chi-Squared Association	chisq.test
Categorical	Numeric	ANOVA	aov
Binary	Categorical	Difference in Proportions	fisher.test/prop.test
Binary	Numeric	Difference in Means	t.test
Numeric	Categorical	Logistic Regression	glm
Numeric	Numeric	Correlation	cor.test
Categoical		Chi-squared Good of Fit	chisq.test
Binary		Exact Binomial	binom.test
Numeric		T-Test	t.test