

STAT 113

Testing for Associations Between Variables

Colin Reimer Dawson

Oberlin College

October 26, 2017

Outline

Testing for a Difference Between Two Groups

Binary Response

Quantitative Response

Testing for an Association

Outline

Testing for a Difference Between Two Groups

Binary Response

Quantitative Response

Testing for an Association

Handout: Penguins and Heavy Metal

Penguins

- Hypotheses:

$$H_0 : p_{control} - p_{metal} = 0 \text{ (population proportions)}$$

$$H_1 : p_{control} - p_{metal} \neq 0$$

- Statistic:

$$\hat{p}_{control} - \hat{p}_{metal} \text{ (sample proportions)}$$

- Operative question: is difference in sample proportions “too large” to have occurred by chance (random assignment) alone?
- (“Too large”: a difference that large is too unlikely if H_0 is true)
- To assess likelihood of a difference that large (P -value), repeatedly randomize assignment to experimental groups, fixing outcomes, and see how often a difference that large appears.

Key Idea: Randomizing Association

- We can restate a hypothesis about a difference as a more general hypothesis about an *association*
- If two groups differ, then explanatory and response are *associated*; if not, they aren't

H_0 : X and Y have no association

H_1 : X and Y have some association

- Idea: Simulate H_0 by holding fixed the single variable distributions, randomizing the *pairings*

Outline

Testing for a Difference Between Two Groups

Binary Response

Quantitative Response

Testing for an Association

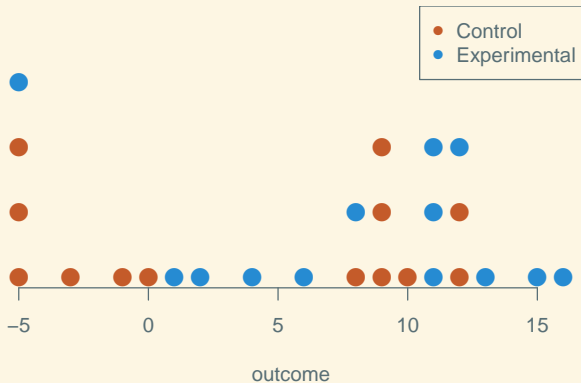
Testing for a Difference Between Two Means

We want to know whether two depression drugs (experimental vs current best practice) achieve different outcomes on a quantitative mood scale. We randomly sample 24 depressed individuals, randomly split them into experimental and control groups, and assess the change in their mood at the end of 3 months.

1. What are the explanatory and response variables?
2. What population parameter(s) is/are of interest?
3. Having collected data (12 cases in each group), how would you construct a randomization distribution using cards? (What goes on the cards? What do you do with them?)
4. How would you compute the P -value using your randomization distribution?


Example: Experimental vs. Control Drug

Here's the real data




Example: Experimental vs. Control Drug

Along with real group means




Example: Experimental vs. Control Drug

Here's one random scrambling




Example: Experimental vs. Control Drug


Here's another...



Example: Experimental vs. Control Drug




Example: Experimental vs. Control Drug




Example: Experimental vs. Control Drug




Example: Experimental vs. Control Drug



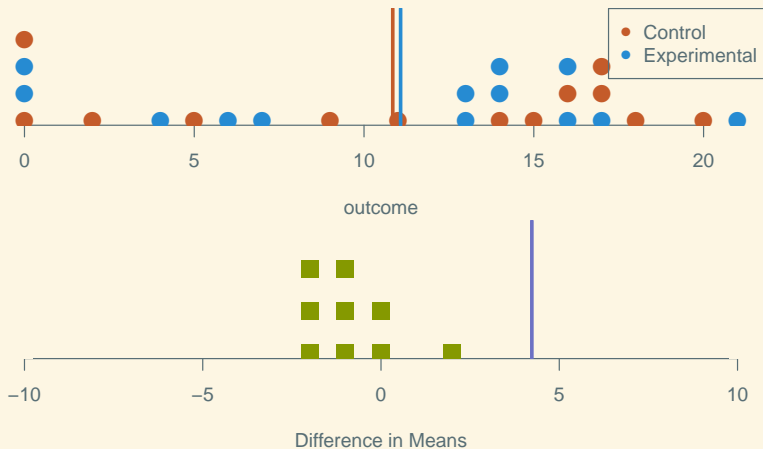
Example: Experimental vs. Control Drug



Example: Experimental vs. Control Drug



Example: Experimental vs. Control Drug




Example: Experimental vs. Control Drug



Example: Experimental vs. Control Drug

Here's the distribution after lots of replicates



Not quite enough evidence to reject H_0 at $\alpha = 0.05$

Outline

Testing for a Difference Between Two Groups

Binary Response

Quantitative Response

Testing for an Association

Randomization Procedure to Test an Association

1. Suppose H_0 is true: explanatory and response variable are not associated
2. Scramble explanatory variable across cases (fixing response variable) simulating random pairings of unassociated variables.
3. Compute the statistic of interest based on “fake” groupings.
4. Repeat steps 2-3, many times, storing the statistic each time, to generate the **randomization distribution** based on H_0 .
5. Locate the *observed* statistic in the randomization distribution, and find the proportion of *simulated* differences at or beyond it
6. This proportion is the P -value.