Comparing NHL Salaries: A Random Permutation Test

Dr. Laura Schultz
Rowan University

Are salaries "fair" in the National Hockey League (NHL)? If so, we would expect that salaries are comparable for the right-wing and left-wing positions. To address this question, I downloaded data from the *USA Today* salary database for the 2007-2008 NHL season and used JMP to run a random permutation test.

2007NHLSalary.jmp

PermutationTest.jsl

Permutation Tests

- Also called randomization tests, re-randomization tests, or exact tests

GENERAL PROCEDURE FOR PERMUTATION TESTS

To carry out a permutation test based on a statistic that measures the size of an effect of interest:

1. Compute the statistic for the original data.

2. Choose permutation resamples from the data without replacement in a way that is consistent with the null hypothesis of the test and with the study design. Construct the permutation distribution of the statistic from its values in a large number of resamples.

3. Find the *P*-value by locating the original statistic on the permutation distribution.

From Moore, McCabe, & Craig (2009)

Advantages of Random Permutation Tests

- ☑ "Sexier" than traditional nonparametric tests
- ☑ Can be used for any test statistic, regardless of whether or not its distribution is known
- ☑ More general hypotheses than those used for *t*-tests

  For example:
  - H_0: Right-wing and left-wing hockey players have the same salary distribution.
  - H_1: The salaries of right-wing and left-wing hockey players are systematically different.
- ☑ Very accurate *P*-values if enough permutations are used
- ☑ Useful for checking the normality assumption of the two-sample *t*-test

Primary Disadvantage

- ☑ Custom computer code is required.
  - My JMP script can only be used for hypothesis tests involving two independent samples.

Left-Wing (n = 114) and Right-Wing (n = 114) Salary Distributions

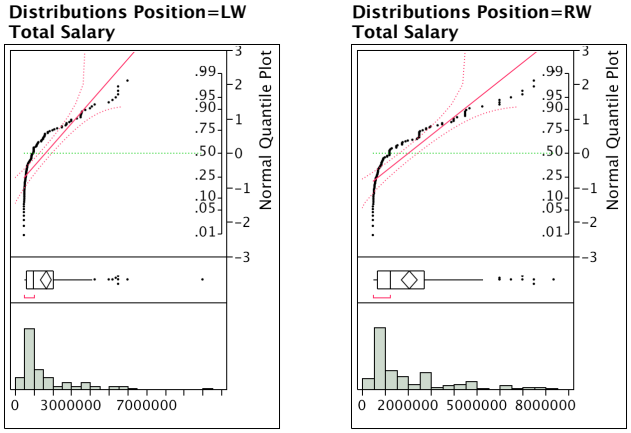

Procedure

- First, find the difference $\bar{x}_{RW} - \bar{x}_{LW}$ for the original sample data. We find that $\bar{x}_{RW} - \bar{x}_{LW} = \$400{,}102$. This is our test statistic.
- We don't compute *t* or any other standardized statistic.
- To perform a permutation test, we randomly regroup the salary data into two groups that are the same sizes as the two original samples (114 LW and 114 RW). This is consistent with the null hypothesis that left-wing and right-wing player salaries have the same distribution. Each salary appears once in each resample, but some salaries move from the LW group to the RW group and vice versa.
- Open up the script named *PermutationTest.jsl* and the data table named *2007NHLSalary.jmp*.
- Bring the script window to the front, and then click on **Edit** > **Run Script**.

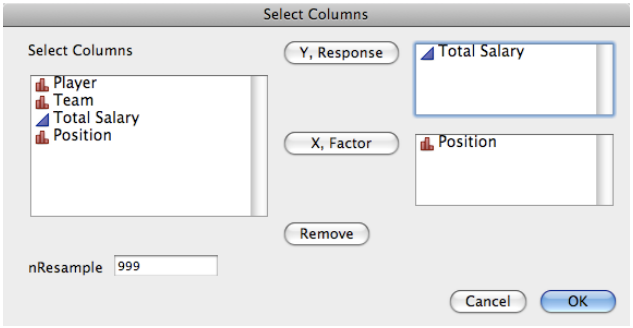

- After a short delay, the script will generate a summary table containing 999 resample means for the two groups.
- Add a new column named "Difference (RW - LW)" to the summary table and define it as *Total Salary_RW - Total Salary_LW*.
- The last row in the table includes the test statistic. Select this row, then click on Rows > Exclude/Unexclude.

For each permutation resample, the script calculates the means for the "new" RW group and the "new" LW group. Then, we find the difference between the two means, $\bar{x}_{RW} - \bar{x}_{LW}$. We repeat the process 999 times and construct the permutation distribution of $\bar{x}_{RW} - \bar{x}_{LW}$. The permutation distribution estimates the sampling distribution under the condition that H_0 is true.

Permutation Distribution of $\bar{x}_{RW} - \bar{x}_{LW}$

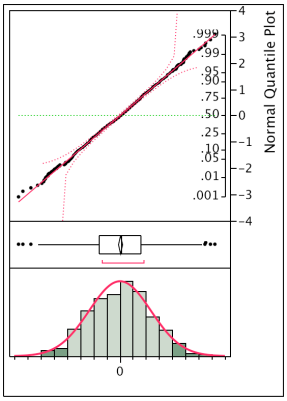

Goodness-of-Fit Test
Shapiro-Wilk W Test

W	Prob<W
0.998965	0.8550

Note: Ho = The data is from the Normal distribution. Small p-values reject Ho.

- Click on **Tables > Sort**. Sort your table in ascending order by the *Difference (RW - LW)* column.

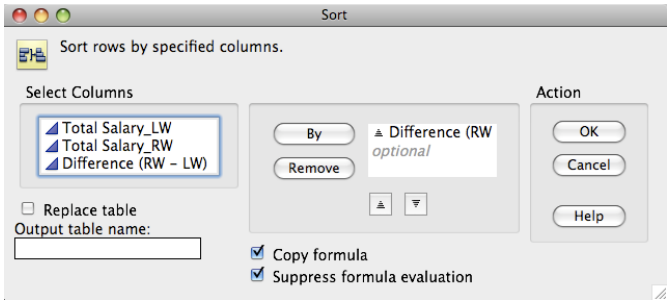

- The *P*-value is the proportion of the 999 resamples that yield $|\bar{x}_{RW} - \bar{x}_{LW}|$ values that are at least as extreme as the observed difference of 400192.018.
- For a two-sided test, we need to count how many resamples have differences that are either ≤ -400192.018 or ≥ 400192.018.
- To improve the accuracy of the *P*-value, add 1 to both the numerator and the denominator.

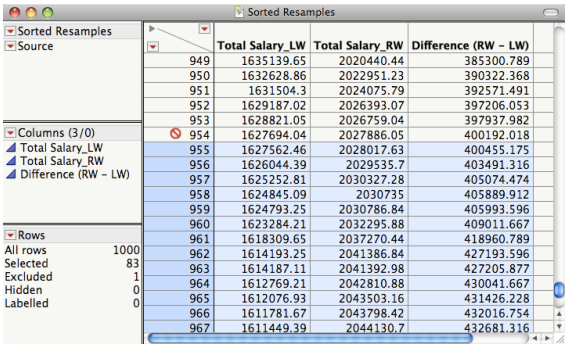

- When I ran the script, there were 83 resamples that resulted in $|\bar{x}_{RW} - \bar{x}_{LW}|$ that was at least as extreme as 400192.018. (Your results may vary.)
- The two-sided *P*-value for my random permutation test is

$$\frac{83+1}{999+1} = .084$$

Moore, McCabe, & Craig (2009) on the robustness of the two-sample *t*-test procedure:

"When the sizes of the two samples are equal and the distributions of the two populations being compared have similar shapes [not necessarily normal], probability values from the *t* table are quite accurate for a broad range of distributions when the sizes are as small as $n_1 = n_2 = 5$."

6th Edition, p. 456

t Test
RW–LW
Assuming unequal variances

Difference	400192	t Ratio	1.720009
Std Err Dif	232669	DF	220.1226
Upper CL Dif	858735	Prob > |t|	0.0868
Lower CL Dif	−58351	Prob > t	0.0434
Confidence	0.95	Prob < t	0.9566

(Compare to *P* = .084)

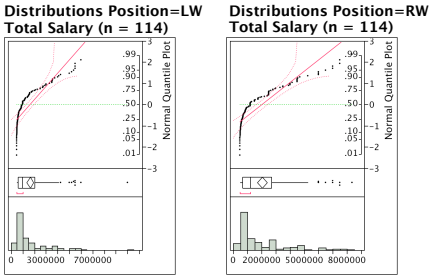

Contact Information:

Laura Schultz
Department of Mathematics
Rowan University
201 Mullica Hill Road
Glassboro, NJ 08028

schultzl@rowan.edu

http://users.rowan.edu/~schultzl