

Never Worry About a Normal Distribution Again!

Instead, Use Resampling Methods

Dennis R. Helsel
PracticalStats.com



© 2019 PracticalStats.com

1

Details

- If you're watching this as a recording, a pdf of these slides is available at <http://practicalstats.com/downloads/>
- This webinar may not be recorded, stored electronically, or rebroadcast to another location or time without the consent of Practical Stats
- We will answer questions at the end of the webinar. Type them into the Chat window as they occur to you.



© 2019 PracticalStats.com

2

The Problem

- Field data are usually skewed (not a normal distribution)
- Parametric tests and intervals assume data look like a normal distribution.
- When this is not true they often do not find differences that are there (low power). Intervals are often inaccurate.
- Transformations aren't the answer -- they no longer deal with the mean
- How to compute accurate confidence intervals and test differences in means with moderate sample sizes of skewed data?



© 2019 PracticalStats.com

3

The Problem: t-intervals assume the data follow a (symmetric) normal distribution


- Parametric t-intervals assume either that the data follow a normal distribution, or that there is sufficient data that the sample mean does.
- If this is not so, t-intervals will not include the true mean as often as the confidence level suggests -- a 95% confidence interval will include the true mean less than 95% of the time.
- Low ends of the confidence interval are usually too low, as the skewness and outliers at the upper end inflates the standard deviation, shooting the lower end below where the data are located (and possibly below 0).



© 2019 PracticalStats.com

4

4



5

The Solution: Resampling Methods

- Bootstrapping
 - computes confidence intervals and measures of variability
- Permutation Tests
 - test hypotheses without using an assumed theoretical distribution

Resampling Methods state that all the info about the underlying distribution is in the sample data.

They use the observed distribution of the sample data (cdf) as the distribution with which to estimate intervals or test hypotheses.


© 2019 PracticalStats.com

6

The Solution: Bootstrap intervals are more accurate

- Bootstrapped confidence intervals avoid the t-interval assumptions of normality and therefore symmetry.
- Lower end of a bootstrap interval is often shorter than a t-interval for skewed data
- Example: Mercury concentrations. Compare the 95% CIs

t interval: 4.94 to 9.54
bootstrap: 5.59 to 8.79 (shorter)



7

How Does Bootstrapping Work?

Resample from the data set, with replacement. For n=7 data:

- Randomly choose and copy 1 of the 7 concentration values and put into a second column named "Resample #1". Each value has the same chance of being chosen. Again choose any 1 of the 7 and copy into the second row of the Resample #1 column. Do this for all 7 rows, forming the first 'set' of bootstrapped data. Some observations could be chosen more than once, some not at all.
- Create thousands of Resample columns and compute the mean for each.
- Use the variability in the thousands of estimates of the resampled means to compute the ends of a confidence interval.



8

Sampling with Replacement:

Original data

Bootstrap resamples

Concentration	Bootstrap Resample #1	Bootstrap Resample #2
12	16	22
11	14	16
16	28	11
22	16	14
14	12	11
28	24	24
24	14	12

... to thousands of resamples

Mean = 18.1

Mean = 17.7

Mean = 15.7

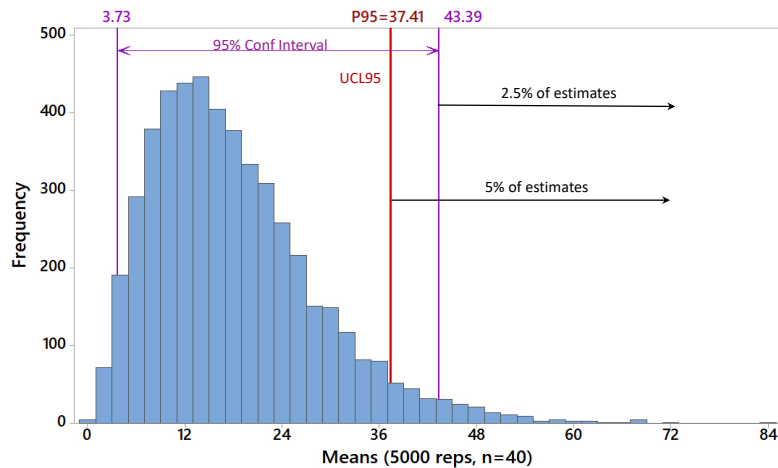


9

Example: Computing Confidence Intervals from Bootstrap Estimates of Means

The UCL95 is the 95th percentile of bootstrap estimated means.

The 2-sided 95% confidence interval endpoints are the 2.5th and 97.5th percentiles of the bootstrap estimated means.



10

Example: Computing the UCL95 by Bootstrapping

- The bootstrap estimate will reflect the shape of the observed data, not the shape of a theoretical distribution such as the normal distribution.
- Percentile Bootstrap: 1-sided 95% UCL is the 95th percentile of the several thousand means of bootstrap resamples. For a 95% two-sided interval, the 2.5th and 97.5th percentiles of the resamples become the interval endpoints (5% of estimates are outside that interval)
- BCA Bootstrap: Percentile bootstrap may not cover with 95% probability for strongly skewed data. BCA adjusts the percentile estimate for skewness. Requires more data than for the percentile bootstrap.
- With small ($n < 20$) datasets, observed data likely will not capture the breadth of the population shape. Assuming a distribution should be better than bootstrapping for $n < 20$ data sets.



11

Is It Easy to Compute Bootstrap Intervals?

`bootUCL(conc, conf=95)` ← option to change confidence level
(UCL95, conf = 95, is the default)

Bootstrap Estimate of an Upper Confidence Limit of the Mean.

Data Analyzed = conc

	XBAR	UCL	CONF	NREP
1	98.352	148.464	94.99	10000

Results using the bootUCL script from our AES course, using R software



12

Is It Easy to Compute Bootstrap Intervals?

Bootstrap results using the free PAST statistics software:

2-sided 95% CI on the mean

The screenshot shows the 'Univariate statistics' window in PAST software. It contains a table of statistics and a control panel on the right. The table has columns for 'All', 'Lower conf.', and 'Upper conf.'. The 'Mean' row is highlighted with a red border, showing a mean of 98.352 and a 95% confidence interval from 48.148 to 159.808. The control panel on the right has a checked 'Bootstrap' box, 'Bootstrap type' set to 'Percentile', and 'Bootstrap N' set to 9999. A 'Recompute' button is also visible.

	All	Lower conf.	Upper conf.
N	25	25	25
Min	1.3		
Max	580		
Sum	2458.8	1203.7	3995.2
Mean	98.352	48.148	159.808
Std. error	28.93701	15.21954	39.64802
Variance	20933.77	5790.86	39299.14
Stand. dev	144.6851	76.0977	198.2401
Median	19	4.8	110
25 prcntil	3.75	2.05	12
75 prcntil	155	41	300
Skewness	1.949363	0.6995782	3.02921
Kurtosis	4.019408	-1.115938	10.31516
Geom. mean	23.87106	11.28377	50.40279
Coeff. var	147.1094	103.8563	202.032



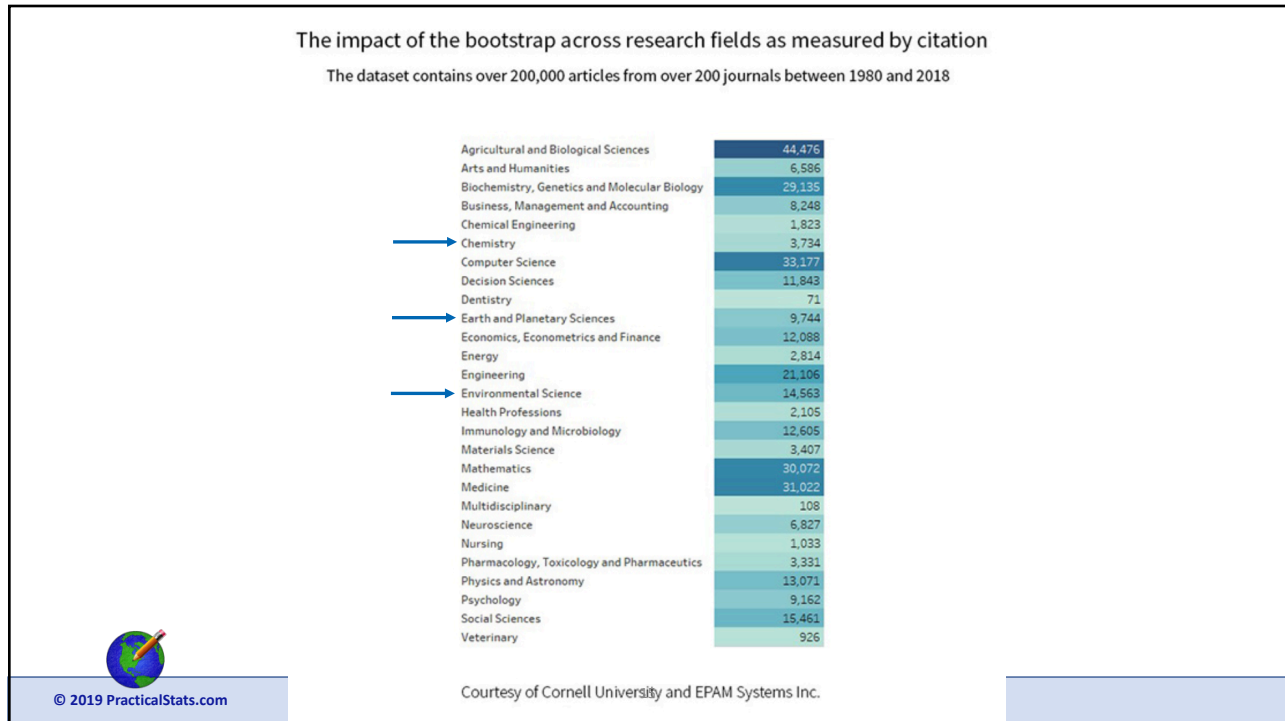
How Well-Accepted is the Bootstrap?

Announcement by the American Statistical Association and the International Statistical Institute:

The creator of the bootstrap, Bradley Efron, professor of statistics and biomedical data science at Stanford University, has been awarded the International Prize in Statistics for developing the bootstrap in 1977 to assess the uncertainty of scientific results. It has had extraordinary impact across many scientific fields.

-- Nov. 2018





15

As part of the award to Bradley Efron of the International Prize in Statistics, 2018

"While statistics offers no magic pill for quantitative scientific investigations, the bootstrap is the **best statistical pain reliever** ever produced," says XiaoLi Meng, Whipple V. N. Jones Professor of Statistics at Harvard University. "It has saved countless scientists and researchers the headache of finding a way to assess uncertainty in complex problems by providing a simple and practical way to do so in many seemingly hopeless situations."

© 2019 PracticalStats.com


16

16

Bootstrap CI on the Theil-Sen line (associated with the Mann-Kendall trend test)

Computed using the `theil-sen_UG` script for R written by Kirk Cameron, MacStat Consulting, and available in the Unified Guidance (USEPA, 2009).

No direct theoretical computation of the CI for the Theil-Sen line is available. But it can be and is bootstrapped!



17

When might a t-interval work well?

- The Central Limit Theorem (CLT) states that the sample mean follows a normal distribution regardless of the shape of the data...IF n is large. Large is a function of skewness -- more skewness, more data needed.
- USEPA (2000) stated that CLT may require up to 100 obs for skewed data such as field data. We think 70 is about right for skewness typical of environmental data. Below $n=70$, don't just hope that the CLT applies. Your t -UCL95 might exceed the population mean only 65% of time, for example, not 95%.
- However, bootstrapping is valid regardless of the data shape -- it works well for data from a normal distribution, and for $n>70$, too!



18



What about for small $n < 20$ data sets? Fit to another distribution

- Gamma, lognormal and Weibull distributions are all popular skewed distributions used in many disciplines
- The math of these shapes are included in most statistics packages, so many types of intervals and estimates may be computed using them
- Choose the one with the best GOF statistics – largest PPCC or Shapiro-Wilk test statistic or p-value for the test of fit.





Fit to another distribution

Probability plots for DOWNGRAD Moly data (using Minitab)



Shapiro-Wilk test statistic
 Normal 0.795
 Lognormal 0.956
 cube-root 0.959
 (similar to gamma dist)



PPCC
 Normal 0.879
 Lognormal 0.974
 Gamma 0.961



Summary: Bootstrapping

- Bootstrapping is widely applicable and distribution-free.
- It provides more accurate interval estimates when data are not normally distributed, the typical situation in environmental science.
- It provides estimates very similar to normal theory results when data do follow that distribution, so there is little penalty for using it.
- It is used to compute measures of variability that have no known theoretical equations for computing.
- It is widely used in a large number of scientific disciplines

Are you still worrying about whether your data follow a normal distribution or not? Bootstrapping is presented in more detail in our Applied Environmental Statistics 1 online course.




Permutation Tests

The Problem:

- Most data in environmental science have a skewed (non-normal) distribution.
- The mean is commonly used as a measure of the group center.
- Tests for differences in means (t-tests and ANOVA) are notoriously sensitive to non-normality. If given non-normal data, the tests have low power to see differences (p-values are too high)
- Results: Contamination goes unnoticed. Faulty conclusions are reached.



Is there a difference in molybdenum concentrations between the two groups?



All of the upgrad data are higher than all of the DOWNGRAD data

The DOWNGRAD data test as non-normal. The upgrad data may be, but there are only 3 observations so non-normality cannot be tested with any power to see it.



Common approach: Two sample t-test

Null Hypothesis: $\text{mean } X = \text{mean } Y$

Alternate Hyp: $\text{mean } X \neq \text{mean } Y$ (2-sided)


or

$\text{mean } X > \text{mean } Y$ (1-sided)

Assumptions: Each group's data follows a normal dist.
Each group's data have same variance



The means weren't found significantly different using the t-test



t = -2.3836
p-value = 0.14
mean DOWNGRAD: 0.25
mean upgrad: 3.93

4 vs 1/4 certainly looks different!
Did the non-normal DOWNGRAD data lessen the power of the t-test?
Did the unequal variances of the two groups lessen the power of the t-test?



The Solution: Permutation Test on Means

- H_0 : Mean[Group 1] = Mean[Group 2]
- If H_0 is true the data could be randomly reassigned to either group.
- Declare the test statistic to be the difference in the means.
- “Shuffle” the group assignment 10,000 times, or all possible ways if fewer than 10,000.
- Compute the difference in the means for each shuffle. the collection of them is a picture of the null hypothesis. Compute the percent of results equal to or more extreme than the one observed in your data (when the null hypothesis is true). That is the permutation p-value.



The Permutation Test Finds a Difference in the Means

The perm2 script will compute all possible results if there are fewer than 1000.

```
> perm2(MOLY,LOCAT)
```

```
Data analyzed = MOLY LOCAT
```

```
Group names are  DOWNGRAD upgrad
```

```
PERMUTATION TEST OF DIFFERENCE IN 2 MEANS
```

```
Number of Possible Permutations = 560 is less than 1000
```

```
ALTERNATIVE: MEAN of DOWNGRAD NOT EQUAL TO MEAN of upgrad
```

```
Diff of means = -3.686  pvalue = 0.0018  nrep = 560
```

(Remember the t-test p-value was 0.14 !)



Or Use the Version in the Coin Package of R

```
> oneway_test(MOLY~LOCAT,data=MOLY2)
```

```
Asymptotic 2-Sample Permutation Test
```

```
data: MOLY by LOCAT (DOWNGRAD, upgrad)
```

```
Z = -3.22, p-value = 0.0013
```

```
alternative hypothesis: true mu is not equal to 0
```

(Remember t-test p-value = 0.14)



In 1935, Karl Pearson said.....

- Parametric tests (t-tests, etc.) produce results that are only approximations to the true result, which would be found by computing

Permutation (or randomization) tests

- The quality of the approximation depends on whether the data fit the assumptions required.
- Permutation methods were envisioned, but just couldn't be done with the computing power (pen and paper, mostly) available then.



© 2019 PracticalStats.com

29

How are permutation tests computed?

- If the null hypothesis is true (such as: there is no difference between group means), the data could be randomly reassigned to any group. The groups are equivalent.
- Choose and compute a test statistic for the observed data, such as the difference between means. Save.
- “Shuffle” the group names many times.
- Compute the test statistic for each shuffle.
- The collection of test statistics from shuffles collectively represent the outcomes expected when the null hypothesis is true.
- The permutation p-value for the original test statistic equals the proportion of shuffled results with a test statistic \geq the one result for the observed data.



© 2019 PracticalStats.com

30

If these are the original data

Concentration	Location
6	Site 1
5	Site 1
10	Site 1
16	Site 2
8	Site 2
22	Site 2
18	Site 2



© 2019 PracticalStats.com

31


Shuffle the group names (1 of many possibilities)
Do this thousands of times or for all possible shuffles, whichever is smaller, and compute the test result for each shuffle

Concentration	Location
6	Site 2
5	Site 1
10	Site 2
16	Site 1
8	Site 2
22	Site 1
18	Site 2



© 2019 PracticalStats.com


32



33

If the mean is the objective, use a permutation test

- Perm tests make no distributional assumptions about the population sampled. (Does not require assumption of normality).
- Does not rely on the Central Limit Theorem
- Uses only the observed data and all (or many) possible rearrangements of the data
- Has as much or more power than traditional parametric tests
- Can be used with censored data (nondetects)
- Is less effected by outliers and unequal variance than are parametric tests



© 2019 PracticalStats.com

34

Number of possible permutation results increases quickly!

- With two groups of 20 observations each, there are $40!/(20!*20!) = 1.3785*10^{11}$ combinations, or 137 billion arrangements
- This is too many to compute all the possible test results in a reasonable time frame, even for today's computers
- The alternative – compute a large number, say 10,000 random shuffles producing 10,000 test results, and stop there. Assume that these 10,000 adequately represent the entire distribution of 137 billion possible test results when the null hypothesis is true




© 2019 PracticalStats.com

35

Permutation Test for Difference between k Group Means

Concentrations for 4 sites, continued.
(ANOVA had $p=0.073$)



© 2019 PracticalStats.com

36

Permutation Tests for One-Way Layout

- H_0 : All means are equal
- If H_0 is true the data could be randomly reassigned to any group.
- “Shuffle” the SITE names many times.
- Compute an F ratio for each shuffle.
- p-value equals percent of shuffled reps with an F value that exceeds the original observed F statistic



© 2019 PracticalStats.com

37

ANOVA Results Assuming Normal Distribution within Groups

ANOVA for difference in concentrations in the 4 groups:

```
> AnovaModel.1 <- aov(Concentration ~ Location, data=Conc)
```

```
> summary(AnovaModel.1)
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
Location	3	51	16.99	2.36	0.0728
Residuals	196	1411	7.20		




© 2019 PracticalStats.com

38

Normal Q-Q plot of ANOVA residuals

- Data do not follow a normal distribution




© 2019 PracticalStats.com

39

Normal Q-Q plot of ANOVA residuals – log units

- What is being tested, now that we've taken logs?



© 2019 PracticalStats.com

40

Transform by Taking Logs?

ANOVA for difference in mean log (geometric mean, or median) of concentration in the 4 groups:

```
> AnovaModel.2 <- aov(LN.of.Concentration ~ Location, data=Conc)
> summary(AnovaModel.2)
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
Location	3	53.2	17.735	4.164	0.00692 **
Residuals	196	834.8	4.259		

But this is not a test for difference in means!



© 2019 PracticalStats.com

41

Permutation Test for Whether Mean of the 4 Groups Differ Using R

```
> perm1way(Concentration, Location, nrep=10000)
```

Data Analyzed: "Concentration"

Permutation Results of the One-Way parametric ANOVA

	FOBS	PERM_PVAL	NREP
1	2.359746	0.0475	10000

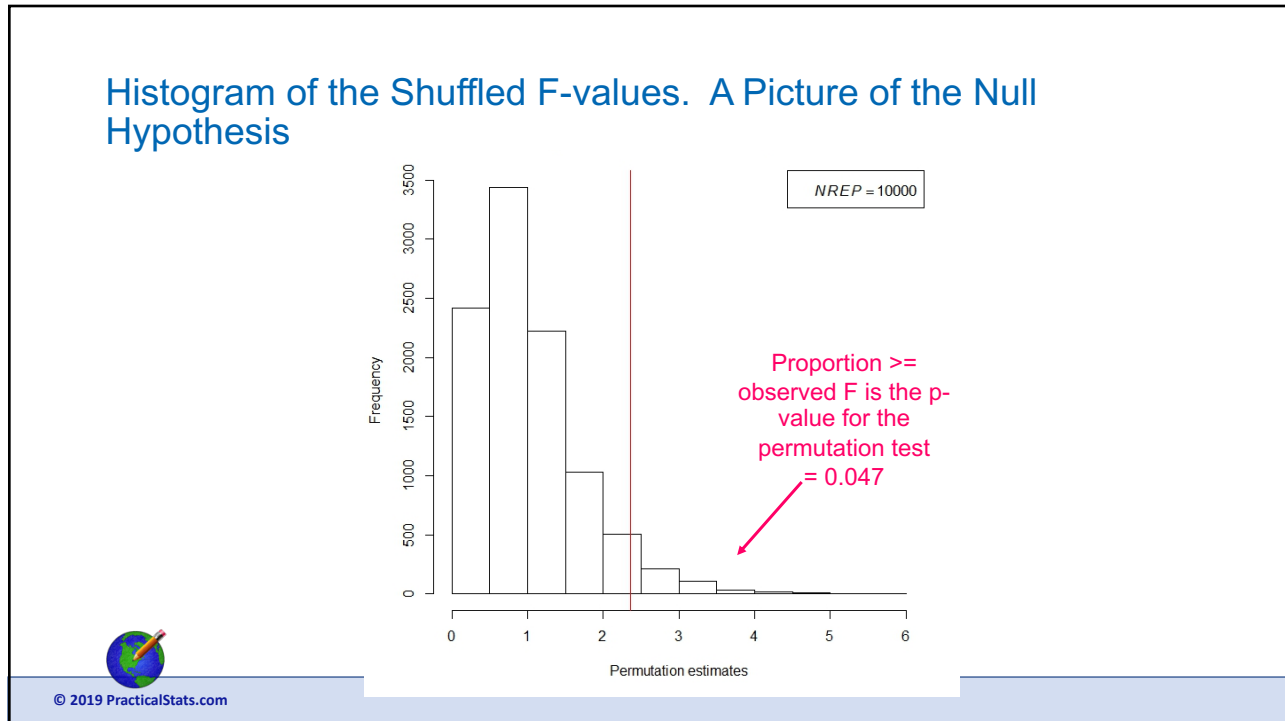
The ANOVA p-value was 0.728. The difference between 0.0728 and the permutation p-value is a measure of the power lost by applying this parametric test to non-normal data.

There were 200 observations. So the Central Limit Theorem doesn't fix everything!




© 2019 PracticalStats.com

42



43



44

Test for Matched Pairs

- Soil lead was measured at the same sites in 1996 before a major fire, and after the fire in 2001
- Any difference between the two years is attributed to the effect of the fire.
- Are mean lead concentrations before the fire different than after the fire (a two-sided test)?
- Compute the differences After–Before for pairs at the same site. Is the mean difference = 0?




© 2019 PracticalStats.com

45

Boxplot of Paired Differences

Diffs are symmetric.
Outliers, esp. #30,
don't fit a normal
distribution. Is this a
problem for the t-test?



© 2019 PracticalStats.com

46

Compute Matched Pair Perm Test

```
> permMatched(Yr1996, Yr2001)
```

The p-value for the permutation matched-pair test is = 0.0809

The normal-theory paired t test is:

Paired t-test

data: After and Before

t = 1.7406, df = 81, **p-value = 0.08555**

alternative hypothesis: true difference in means is not equal to 0

95 percent confidence interval:

-0.1186691, 1.7772057

sample estimates:

mean of the differences

0.8292683

Same result for the two tests. But now we know that the 'not significant' result is not due to non-normality. With the perm test, there's no worry.



© 2019 PracticalStats.com

47

Permutation Tests for Multivariate Data

- Perform something like a MANOVA for differences between groups, where multiple variables are used to characterize each group
- Example: atrazine concentrations are summarized as ≥ 1 or < 1 ug/L (there were multiple DLs with 1 ug/L being the highest). This forms two groups.
- Flow percentile, %corn grown in the basin, and soil conditions are some of the 7 variables used to characterize each sampling site
- If there is a significant difference in the mean of one or more of the 7 variables by atrazine group, then those variables could be used to predict the concentration grouping at another as yet unsampled site.




© 2019 PracticalStats.com

48

NMDS shows group separation, where data position reflects 7 watershed characteristics

Assignment to the <1 group increases going to the right



© 2019 PracticalStats.com

49

Permutation Tests for Multivariate Data

- Test is called Permanova (in the package sold by Primer-e) or the adonis function in the vegan package of R.
- Is like a multivariate MANOVA but computes its p-value by a permutation test.
- Test statistic is a multivariate F statistic. Data are randomly permuted several thousand times.
- No assumption of normality or equal variance/covariance matrix required.



© 2019 PracticalStats.com

50

Permanova output

```
adonis(DYPLANT~GT_1,method="euclidean", permutations = 4999)
```

Terms added sequentially (first to last)

	Df	SumsOfSqs	MeanSqs	F.Model	R2	Pr(>F)
GT_1	1	2498828	2498828	302.96	0.41848	2e-04 ***
Residuals	421	3472441	8248		0.58152	
Total	422	5971269			1.00000	


Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1



© 2019 PracticalStats.com

51

Adonis results



Permutation results by mixing group assignments. Represents the null hypothesis.

P=0.0002 (observed is higher than all 4999 permutations)



© 2019 PracticalStats.com

52

Summary: There are Permutation Tests for...

- Testing difference in variability (alternative to Bartlett's F test)
- Testing differences in matched pairs (alternative to paired t-test)
- Testing differences in means of two or more groups (alternatives to t-tests and ANOVA)
- Testing significance of regression slopes (avoiding need to transform variables)
- Testing significance of association in categories (perm tests for contingency tables)
- Testing significance of correlations between variables
- Multivariate relationships (PERMANOVA and other tests)



© 2019 PracticalStats.com

53

Summary: Permutation Tests

- Can be used instead of parametric tests, avoiding assumptions of normal distribution and equal variance
- No "pre-test" for normality required
- More power to see differences between means for skewed data than parametric tests
- Finally, a method that can see differences between means for the skewed data common to environmental sciences



© 2019 PracticalStats.com

54

Resources for Permutation Procedures

Software:

- R Stata SAS PAST

Training: (PracticalStats.com)

- Applied Environmental Statistics online class
- Untangling Multivariate Statistics (possible future) online class

Some Books:

- Statistical Methods in Water Resources, 2nd Edition. Helsel, Hirsch, Archfield, Ryberg and Gilroy. (2019)
- Introduction to the Practice of Statistics, 8th ed. Moore, McCabe and Craig, 2014
- Bootstrap Methods and their Application. Davison and Hinkley, 1997
- Mathematical Statistics with Resampling and R. Chihara and Hesterberg, 2011
- Randomization, Bootstrap and Monte Carlo Methods in Biology. B.F.J. Manly, 2007



© 2019 PracticalStats.com

55

More Resources: Newsletters in our News Archive practicalstats.com/news

Practical Stats News Archive

We send out a free newsletter about once every two months to all who sign up for them. To receive them by email when they are released, [Sign Up Here](#)

All (by date)

Click on the newsletter subject to download it.

Oct 2019	Bootstrapping
Oct 2018	Comparing Data with Nondetects to Standards, Part 2
Dec 2017	Multivariate PermTests 2 – Permanova
Oct 2017	Multivariate Permutation Tests
April 2017	What's the difference between Permutation Tests and Bootstrapping?
July 2015	Parametric vs Nonparametric vs Permutation Tests
May 2015	The power of permutation tests
Nov 2013	Bootstrapping regression equations
Jan 2008	Permutation tests – never worry about a normal distribution again!



© 2019 PracticalStats.com

56

56

Next Month's Webinar

Tuesday Nov 19 11 am Mountain time

- **Trend Analysis for Data with Nondetects**

Another topic from our online course [Nondetects And Data Analysis](#)

- Sign up for our newsletter/announcement list to get the registration link emailed to you. Our newsletter page is at <http://www.practicalstats.com/news/>
- Or check our webinars page periodically at <http://practicalstats.com/training/webinar.html> to register for it.



57

Questions?

Type them into the Chat area

I'll compile them and their answers and post it all to our Downloads page, <http://www.practicalstats.com/info2use/downloads.html>



58

PracticalStats webinars

This and other webinars are free to download from our Training Site,

<https://practicalstats.teachable.com/>

The screenshot shows the 'Practical Stats Courses' website. At the top, it says 'Statistics, Down to Earth' and has an 'Enroll now' button. Below, under 'Featured Courses', there are three course cards: 'Applied Environmental Statistics 1' (Make Sense of Your Data, \$650), 'Applied Environmental Statistics 2' (Regression Models and Trend Analysis, \$650), and 'Nondetects And Data Analysis (NADA)' (Statistical analysis of censored data, \$695). A 'View All Courses' button is at the bottom.



© 2019 PracticalStats.com

59

59

Webinars on Environmental Statistics

The screenshot shows a collection of webinar cards. The top row includes: '7 Perilous Errors in Environmental Statistics' (Stop Your Methods Into The 21st Century!, FREE), 'Stats 4 Managers' (What's new in statistics that your employees should know!, FREE), and 'Intro to R' (Getting started with the free statistics software, FREE). The bottom row includes: '40 Years' (40 Years of Water Quality Statistics, What's New, What's Not, FREE), 'How Multiple Comparison Tests Work' (Which of These Things is Different Than Others?, FREE), and a card for 'Never Worry About a Normal Distribution Again!' (coming soon).



© 2019 PracticalStats.com

60

60

Webinars on Data with Nondetects ("censored data")

<1
Introduction to Nondetects And Data Analysis
Simple things to do (and not do) with nondetects
Dennis Helsel FREE

<2
Fitting Distributions to Data with Nondetects
Fitting the most out of small datasets with NDs
Dennis Helsel FREE

<3
Testing Differences in Groups for Data with Multiple Detection Limits
Parametric and nonparametric tests of differences and multiple comparison tests for data with one or multiple DLs. Without substitution.
Dennis Helsel FREE

<4
The Mystery of Nondetects
How Censored Data Methods Work
Dennis Helsel FREE

<5
Correlation and Regression for Data with NDs
without substituting values for nondetects
Dennis Helsel FREE



© 2019 PracticalStats.com

61

61

Sign Up for our Newsletter and Receive Webinar Announcements

<http://www.practicalstats.com/news/>

Subscribe to our Newsletter

Want to be notified of new class offerings and schedules?

Want practical tips on environmental statistics delivered to your mailbox so you can stay current with other professionals?

Then sign up for our free newsletter in the "Newsletter Sign Up" area! We deliver a newsletter to your inbox about once every two months, but only to those who sign up for them. Past newsletters are ready to download from our [News Archive](#) page.

NOTE:

Our online signup is working again! If you have any problems signing up, send us an email using [our Contact Us page](#) and we'll add your address to the newsletter list, as well as try to get the online process fixed. Either way, you will then receive a confirmation email to make sure you are the one who requested to add your email to the list. Respond in the affirmative to the confirmation email and you're set to go. Don't respond and your email address is not added to our list. Your choice.

Newsletter Sign Up

Name:

E-mail:

Subscribe: Practical Stats News

Unsubscribe



© 2019 PracticalStats.com

62

62