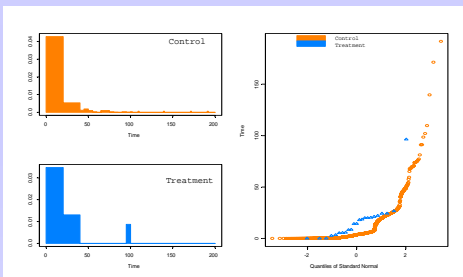


Why Resample?

- Fewer assumptions
 - Normality
 - Equal variances
- Greater accuracy
 - Typically \sqrt{n} smaller errors in size or coverage probability
- Flexible
- Nice pictures (teaching, consulting)
- What-if analysis (planning clinical trials)

Example: Compare Two means


- $n_1 = 1664, \bar{x}_{bar1} = 8.4$
- $n_2 = 23, \bar{x}_{bar2} = 16.5$
- Is the difference statistically significant?
- Highly skewed



Command line:

- `save.perm = permutationTestMeans(Trial, Arm)`
- `plot(save.perm); print, ...`

Menu interface




Permutation tests (any distribution)

- P-value 0.0183

Usual t-test (assumes Gaussian population)

- P-value 0.0045

T-test assumes symmetry, is off by a factor of 4!



Example: Logistic Regression

Predict Kyphosis

- *Kyphosis* (present or absent)
- Age of child
- Number of vertebrae in operation
- Start of range of vertebrae

Classical summary

	Value	Std. Error	t value
(Intercept)	-2.03693225	1.44918287	-1.405573
Age	0.01093048	0.00644419	1.696175
Start	-0.20651000	0.06768504	-3.051043
Number	0.41060098	0.22478659	1.826626




- Can use t-statistics for test or SE for confidence intervals, assuming Normal/large sample

Bootstrap

- Don't require Normal/large sample, estimate from data
- Here see highly non-normal sampling distribution; t tests and confidence intervals are not reasonable
- More-accurate confidence intervals; in general errors are $O(1/n)$, vs. $O(1/\sqrt{n})$ for t intervals
- Graphical bootstrap helps understand variability in predictions

Permutation Test

- Compare relationship between response and predictors as a group
- Unfortunately, cannot use to test significance of a single predictor (use bootstrap tilting instead)



Planning Trial – Larger Sample Size

- Training data size n , prospective trial size N
- With bootstrap: draw samples of size N from data
 - Use shape of data to estimate shape of population
 - Sample with different size to mimic trial design



Planning Trial – Testing Assumptions

- To test sensitivity of outcomes to different assumptions
 - Different parameters, e.g. hazard ratio, effect size $\mu_1 - \mu_2$, ratio μ_1 / μ_2
- Draw bootstrap samples from population(s) that satisfy the assumptions

How to create those population(s)?


- Bad – modify the observations
- Instead – create weighted empirical distribution that satisfy the assumptions, using maximum likelihood (empirical likelihood/bootstrap tilting)
- Example: assume $\mu_1 = 1.2 \mu_2$
 - Bad – shift observations left or right; gives distributions with the wrong mode, and shifting left could create negative observations (negative survival times). More subtle but very important – this ignores mean-variance relationships for skewed populations.
 - Instead – weight observations so weighted means satisfy the relationship

$$\Sigma W_1 X_{1j} = 1.2 \Sigma W_2 X_{2j}$$

- Maximum likelihood: maximize

$$\Pi W_1 \Pi W_2$$

- subject to assumed constraint and $\Sigma W_1 = \Sigma W_2 = 1$ (in S+Resample, use the bootstrap tilting functions to do this)
- Draw bootstrap samples from the weighted distributions.



Example: Parameters in Logistic Regression

- Bad – modify observations (how would one even start???)
- Instead – weight observations so weighted logistic regression satisfies assumptions.
- Note – weighting also provides a way to test hypotheses when permutation testing is impossible, e.g. testing a single coefficient; choose weights to satisfy null hypothesis.

S+Resample

- Most flexible and powerful software for resampling:
 - Bootstrap (many sampling options)
 - Permutation Tests
 - Jackknife
 - Cross Validation
 - Prediction Errors
 - Influence
- Features bootstrap tilting
 - 17 or 37 times fewer bootstrap replications
 - Confidence intervals, what-if analysis, & hypothesis testing
- Command line & graphical user interface

S+Resample GUI Reference

Overview of the S+Resample Library

Welcome to the S+Resample library!

There are three main ways to use the library: using modified versions of the graphical interface for certain specific applications such as comparing two means, using resampling menus (bootstrap, jackknife, influence, permutation tests) which support a wide variety of statistics, or using the Command line.

Supplements to Other Menus for Specific Applications

There are currently resampling supplements to six statistical analysis menus:

- **Summary Statistics**: Statistics ► Data Summary ► Summary Statistics...
- **Correlation**: Statistics ► Data Summary ► Correlations/Resample...
- **Proportions**: Statistics ► Compare Samples ► Counts and Proportions ► Proportions
- **One-sample t test**: Statistics ► Compare Samples ► One Sample ► t Test
- **Two-sample t test**: Statistics ► Compare Samples ► Two Samples ► t Test
- **Linear Regression**: Statistics ► Regression ► Linear/Resample...

Menus for Resampling, for use with General Statistics

There are three menus for different kinds of resampling, which may be used with very general statistics:

- **Bootstrap**: Statistics ► Resample ► Bootstrap...
- **Jackknife and Influence**: Statistics ► Resample ► Bootstrap...
- **Permutation Tests**: Statistics ► Resample ► Permutation Test...

Command Line

There are a large number of resampling functions available from the command line. The following are some of the most commonly used functions:

- `bootstrap`
- `bootstrap` (parametric bootstrap)
- `bootstrap` (smoothed bootstrap)
- `permutationTest`
- `permutationTest` (permutation test for comparing two samples)
- `permutationTestMeans` (permutation test for the difference between means of two samples)
- `crossVal` (cross-validation)
- `bootTilt` (parametric bootstrap tests)
- `bootPrad` (bootstrapping prediction error)

Generic functions such as `print`, `plot`, `summary`, `qnorm`, operate on the objects created by the resampling functions.

Other capabilities include confidence intervals:

- `limits.abc`
- `limits.bca`
- `limits.tilt`
- `limits.tilt` (B intervals using resampling standard errors)
- `bootstrapI` (bootstrap I intervals)
- percentiles: `limits.percentile`
- diagnostics: `jackAfterBootstrap`, `tiltAfterBootstrap`, `reweight`

Additional Information

For additional information, see

- the help files,
- the S+Resample Manual available under the Help ► Online Manuals menu,
- other help in this online reference manual, and
- the "ReleaseNotes.txt" and "tutorial.txt" documents in the "doc" folder of the library.