User Guide

for

psurvey.analysis, version 2.9

Probability Survey Data Analysis Functions

by

Thomas Kincaid

December 31, 2005

Introduction

The functions included in **psurvey.analysis** are intended for analysis of probability surveys. The functions were written for the U.S. Environmental Protection Agency's Environmental Monitoring and Assessment Program (EMAP; Messer et al., 1991) for analysis of probability surveys of environmental resources of interest (Dias-Ramos et al., 1995). Although the function are applicable for a wide range of environmental survey designs, the **psurvey.analysis** functions were written for analysis of data generated by a generalized random-tesselation stratified (GRTS) sampling design. For further discussion of the GRTS design see Stevens and Olsen (2004). The functions in **psurvey.analysis** can analyze finite (discrete units, zero-dimensional), linear (one-dimensional), and areal (two-dimensional) resources. Examples of these resource are lakes in the United States (a finite resource), rivers and streams in Oregon (a linear resource), and Chesapeake Bay (an areal resource). The functions can accommodate stratified and unstratified designs, both of which can utilize single-stage or two-stage sampling. Analytical capabilities accommodate both categorical and continuous data. For categorical data, estimates of proportion and size of each category (class) can be obtained. For a finite resource, size is the number of units in the resource. For an extensive (linear or areal) resource, size is the measure (extent) of the resource, i.e., length, area, or volume. For continuous data, estimates of the cumulative distribution function (CDF) and percentiles can be obtained in addition to estimation of the population mean, total, variance, and standard deviation. Optionally, for continuous data, estimation of the deconvoluted CDF and estimation of percentiles using the deconvoluted CDF are available.

Survey Design Options

As mentioned in the introduction, the **psurvey.analysis** functions can accommodate both stratified and unstratified designs. When the design is stratified, then the stratum code for each site must be provided to the functions. The stratum codes are examined by the functions to ensure that more than one unique stratum code exists, otherwise the data is analyzed as an unstratified design. In addition, when removal of missing values results in a single unique stratum code, the data is analyzed as an unstratified design.

For a two-stage design, the stage one sampling unit (primary sampling unit or cluster) codes must be input to the functions. For a stratified design, the stage one sampling unit codes must identify the stratum code in addition to the stage one sampling unit code within a stratum.

Since **psurvey.analysis** is intended for analysis of probability surveys, the design weight (i.e., inverse of the inclusion probability) for each site must be input to the functions. For a two-stage design, both stage one and stage two weights must be provided. A function (**adjwgt**) is included in **psurvey.analysis** that adjusts initial survey design weights when survey design implementation results in use of oversample sites or when it is desired to have final weights sum to a known size of the resource.

Since the default choice for variance estimation in **psurvey.analysis** requires the x-coordinate and y-coordinate of each site, a function (**marinus**) is included in **psurvey.analysis** that converts coordinates measured as latitude and longitude to a coordinate system more appropriate for calculation of the distance metric used by the variance estimator. The alternate choice for variance estimator does not require coordinates. In addition, for a two-stage design, coordinates are required for both the stage one and stage two sampling units.

Data Analysis Options

The primary functions in **psurvey.analysis** for data analysis are **category.est**, **cdf.est**, **cdf.decon**, and **total.est** (see the Function Descriptions section). Discussion regarding options for these functions follows. The other functions for data analysis are **cdf.test** and **relrisk**, which currently are called directly by the user. For further discussion of **cdf.test** and **relrisk** see the R help documentation for **psurvey.analysis**.

The user can provide the know size of a resource to the functions in **psurvey.analysis**, which is used to adjust the estimate of a total, e.g., the size of a resource in a set of categories. For a stratified design, the known size of a resource can be provided for each stratum. In addition, the size of a resource, either a known value or an estimated value when a known value is not provided, are used as stratum weights for calculating estimates for a stratified design. The known size of a resource also is used for calculation of finite population and continuous population correction factors (see the discussion that follows).

For a finite resource size-weighted analysis is accommodated in **psurvey.analysis**. An example of a size-weights is the surface area of each lake in the Northeast region of the U.S., where the set of lakes in the region is treated as a finite resource. The user must provide the size-weight for each sampling unit. For a two-stage design, size-weights are required for the stage one and stage two sampling units. The size-weights are used to scale the design weights for calculation of estimates. The user can provide the known sum of the size weights for the resource, which is used to adjust the estimate of a total. For a stratified design, the known sum of the size weights for the resource can be provided for each stratum. In addition, the sum of the size-weights for a resource, either a known value or an estimated value when a known value is not provided, are used as stratum weights for calculating estimates for a stratified design.

Use of finite population and continuous population correction factors in variance estimation is accommodated in **psurvey.analysis**. In order to calculate the factors for a single-stage design, the user must provide the known size of the resource and a support value for each sampling unit, where support is equal to one for a sampling unit from a finite resource and is equal to the size of the sampling unit for an extensive resource. If the single-stage design is stratified, then the known size of the resource must be provided for each stratum. For a two-stage design the user must provide the number of stage one sampling units in the resource, the known size of each stage one sampling unit, and a support value for each stage two sampling unit. If the two-stage design is stratified, then the number of stage one sampling units in the resource must be provided for each stratum, and the

known size of each stage one sampling unit must be identified both with a stratum code and the stage one sampling unit code.

The default choice for variance estimation in **psurvey.analysis** is the local mean variance estimator. The alternate choice for variance estimation is the simple random sampling (SRS) variance estimator, which uses the independent random sample approximation to calculate joint inclusion probabilities. For additional information regarding the local mean variance estimator see Stevens and Olsen (2003).

Function Descriptions

Functions in **psurvey.analysis** are organized in a hierarchical structure composed of four levels. Functions in the first and second levels are intended for use with a set of response variables and indicators from a probability survey. The first level function creates an object that can be passed to the second level functions. The second level functions organize input and output for analysis and can be called by the user without use of the top level function. The second level functions call the third level functions, which implement data analysis algorithms with support of the fourth level functions. The third level functions can be called by the user for analysis of an individual response variable or indicator. In addition, two of the third level functions (**adjwgt** and **marinus**) are utilized to modify specific survey design variables prior to data analysis. A short description of each function in the top three levels is provided. Functions in the fourth level are not intended for access by the user. Further details regarding the functions is provided in subsequent sections and in the R help documentation for **psurvey.analysis**.

First Level Function

psurvey.analysis

This function creates an object of class psurvey.analysis that contains all of the information necessary to use the functions in the **psurvey.analysis** library to analyze data generated by a probability survey. Output from this functions can be passed directly to the second level functions.

Second Level Functions

cat.analysis

This function organizes input and output for analysis of categorical data generated by a probability survey. Input can be either an object belonging to class psurvey.analysis or through use of the other arguments to the function. Third level function **category.est** is called by this function.

cont.analysis

This function organizes input and output for analysis of continuous data generated by a probability survey. Input can be either an object belonging to class psurvey.analysis or through use of the other arguments to this function. Third level functions **cdf.est**, **cdf.decon**, and **total.est** are called by this function.

Third Level Functions

adjwgt

This function adjusts initial survey design weights when implementation results in use of oversample sites or when it is desired to have final weights sum to a known size of the resource. Adjusted weights are equal to initial weight times the frame size divided by the sum of the initial weights. The adjustment is done separately for each weight adjustment category. This function is not called by a second level function.

category.est

This function estimates proportion (expressed as percent) and size of a resource in each of a set of categories and can also be used to estimate proportion and size for site status categories. Standard errors of the category estimates and confidence bounds are calculated. This function is called by second level function **cat.analysis**.

cdf.est

This function calculates an estimate of the CDF for the proportion (expressed as percent) and the total of a response variable, where the response variable may be defined for either a finite or an extensive resource. Optionally, for a finite resource, the size-weighted CDF can be calculated. In addition, percentiles are estimated. Standard errors of the CDF and percentile estimates and confidence bounds are calculated. This function is called by second level function **cont.analysis**.

cdf.decon

This function calculates an estimate of the deconvoluted CDF for the proportion (expressed as percent) and the total of a response variable, where the response variable may be defined for either a finite or an extensive resource. Optionally, for a finite resource, the size-weighted CDF can be calculated. In addition, percentiles are estimated. Standard errors of the CDF and percentile estimates and confidence bounds are calculated. This function is called by second level function **cont.analysis**.

cdf.test

This function calculates the Wald, Rao-Scott first order corrected (mean eigenvalue corrected), and Rao-Scott second order corrected (Satterthwaite corrected) statistics for categorical data to test for differences between two CDFs (Kincaid, 2004). The functions calculates both standard versions of those three statistics, which are distributed as chi-squared random variables, plus modified version of the statistics, which are distributed as F random variables. This function is not called by a second level function.

marinus

This function converts x-coordinates and y-coordinates measured in units of latitude and longitude, i.e., geographic coordinates measured in decimal degrees, to coordinates in the equidistant, cylindric map projection measured in units of kilometers. The projection center is defined as the midpoint in latitude-longitude space. The map projection is here named after Marinus of Tyre. This function is not called by a second level function.

relrisk

This function calculates the relative risk estimate for a 2x2 table of cell counts defined by a categorical response variable and a categorical explanatory (stressor) variable for an unequal probability design. Relative risk is the ratio of two probabilities: the numerator is the probability that the first level of the response variable is observed given occurrence of the first level of the stressor variable, and the denominator is the probability that the first level of the response variable is observed given occurrence of the second level of the stressor variable. The standard error of the log of the relative risk estimate and confidence limits for the estimate also are calculated. This function is not called by a second level function.

total.est

This function calculates estimates of the population total, mean, variance, and standard deviation of a response variable, where the response variable may be defined for either a finite or an extensive resource. In addition standard errors of the population estimates and confidence bounds are calculated. This function is called by second level function **cont.analysis**.

write.object

This function writes the contents of an object, which may be either a data frame or a matrix, to a plot. This function is not called by a second level function.

Data Input

Overview

Although the first level function provides the most flexibility, data entry is similar for the first and second level functions. Arguments to the first and second level functions provide information for the following categories: (1) sites to be included in the analysis, (2) identification of sets of populations and subpopulations, (3) survey design variables, (4) response variables, and (5) additional variables specifying analytical options. As necessary, site IDs are used to connect the various arguments. An extensive descriptions follows regarding data entry for the first level function. For the second level functions, differences in data entry from those described for the first level function are noted. Data entry for the third level functions is not described. For first, second, and third level functions, arguments are checked for errors and for compatibility of input values. In addition, for arguments indexed by site IDs, missing values are removed from the argument, and corresponding values are removed from all other arguments indexed by site IDs.

First Level Function (psurvey.analysis)

Information regarding sites to be included in the analysis is provided by argument **sites**, which is a data frame consisting of two variables: the first variable is site IDs and the second variable is a logical vector indicating which sites to use in the analysis. If this data frame is not provided, then it will be created, where (1) site IDs are obtained either from the **design** argument, the **siteID** argument, or both (when **siteID** is a formula); and (2) all sites will be used in the analysis. The default value for **sites** is NULL.

Information identifying sets of populations and subpopulations for which estimates will be calculated is provided by argument **subpop**, which is a data frame. The first variable in **subpop** is site IDs, and each subsequent variable identifies a Type of population, where the variable name is used to identify Type. A Type variable identifies each site with one of the subpopulations of that Type. If this data frame is not provided, then it will be created, where (1) site IDs are obtained either from the **design** argument, the **siteID** argument, or both (when **siteID** is a formula); and (2) a single Type variable named All.Sites that consists of all sites will be created. The default value for **subpop** is NULL.

Information regarding survey design variables is provided by argument **design**, which is a data frame, or by the individual design variable arguments to the function. Individual design variables may be provided as a vector of values or as a formula, where the formulas are interpreted using the **design** data frame. If **design** is not provided, then it will be created from the values for the individual design variables in the argument list. The default value for **design** is NULL. If values for the individual variables are not provided, then the variables in **design** should be named as follows: (1) siteID – site IDs; (2) wgt – final adjusted weights, which are either the weights for a single-stage sample or the stage two weights for a two-stage sample; (3) xcoord – the x-coordinates for location, which are either the x-coordinates for a single-stage sample or the stage two x-coordinates for a

two-stage sample; (4) ycoord – the y-coordinates for location, which are either the y-coordinates for a single-stage sample or the stage two y-coordinates for a two-stage sample; (5) stratum – the stratum codes; (6) cluster – the stage one sampling unit codes; (7) wgt1 – the final adjusted stage one weights; (8) xcoord1 – the stage one x-coordinates for location; and (9) ycoord1 – the stage one y-coordinates for location. Names of the nine individual design variable arguments are the same as the default names for the variables in the **design** data frame. Using formulas to input design variables allow the user to supply names for those variables rather than using the default names. Values always are required for design variables **siteID** and **wgt**. Values for **xcoord** and **ycoord** are required when using the local mean variance estimator (see the discussion for argument **vartype**), but are not required for the SRS variance estimator. If a stratified sampling design is used, than values must be provided for design variable **stratum**. Similarly, if a two-stage sampling design was used, than values must be provided for design variables **cluster**, **wgt1**, **xcoord1**, and **ycoord1.** The default value for the **design** data frame and for the individual design variable arguments is NULL.

Information regarding categorical response variables is provided by **data.cat,** which is a data frame, and **type.cat**, which is a vector. The first variable in **data.cat** is site IDs, and subsequent variables are response variables. Missing data (NA) is allowed in **data.cat**. Argument **type.cat** is a vector that provides the type of each categorical response variables, which is either "Status" indicating site status or "Category" indicating resource category. The names attribute for **type.cat** must be set to identify the response variable names. If **data.cat** is supplied and **type.cat** is not provided, then each response variable is assigned type "Category". The default value for **data.cat** and **type.cat** is NULL.

Information regarding continuous response variables is provided by **data.cont**, which is a data frame. The first variable in **data.cont** is site IDs, and subsequent variables are response variables. Missing data is allowed. The default value for **data.cat** is NULL.

Other arguments to the functions provide information required for optional analyses. Arguments **sigma** and **var.sigma** provide information for CDF deconvolution, where **sigma** is a vector of measurement error variance values, and **var.sigma** is a vector of variances for the measurement error variance values. When **sigma** is provided, it is not necessary to provide **var.sigma**, in which case **sigma** is treated as a known quantity and variability of the deconvolution procedure that is due to estimating **sigma** is ignored. Both **sigma** and **var.sigma** must have the names attribute set to identify the continuous response variable names. Missing data is allowed. The default value for **sigma** and **var.sigma** is NULL.

Information regarding know size of the resource is provided by **popsize**, and information regarding the know sum of the size-weights of the resource is provided by **unitsize**. Both arguments must be in the form of a list containing an entry for each Type of population in the subpop data frame, where NULL is a valid entry for a population Type. The list must be named using the variable names for population Types in subpop. If a population Type doesn't contain subpopulations, then the element of the list is either a single value for an unstratified sample or a vector containing a value for each stratum for a stratified sample, where the vector must have the names attribute set to identify

the stratum codes. If a population Type contains subpopulations, then the element of the list is a list containing an element for each subpopulation, where the list is named using the subpopulation names. The element for each subpopulation will be either a single value for an unstratified sample or a named vector of values for a stratified sample. The default value for **popsize** and **unitsize** is NULL.

Information required for calculation of finite and continuous population correction factors is provided by arguments **N.cluster**, **popsize**, **stage1size**, and **support**. **N.cluster** and **stage1size** are applicable to two-stage sampling designs. **N.cluster** provides the number of stage one sampling units in the resource. For a stratified sample **N.cluster** must be a vector containing a value for each stratum and must have the names attribute set to identify the stratum codes. Argument **stage1size** is a vector containing the known size of each stage one sampling unit and must have the names attribute set to identify the stage one sampling unit codes. For a stratified sample, the names attribute for **stage1size** must be set to identify both stratum codes and stage one sampling unit codes using a convention where the two codes are separated by the # symbol, e.g., "Stratum 1#Cluster 1". Argument **popsize** was discussed in a preceding paragraph. Argument **support** provides the support value for each site and is always required for calculation of population correction factors. For a sampling unit from a finite resource, **support** is a vector of ones; and for an extensive resource, it is a vector containing the size of the sampling unit associated with each site.

Argument **vartype** controls the choice of variance estimator, where "Local" indicates the local mean variance estimator and "SRS" indicates the SRS estimator. The default value for **vartype** is "Local".

Argument **conf** provides the confidence level that prescribes the Normal distribution multiplier used in calculating confidence bounds. The default value for **conf** is 95%.

Argument **pctval** provides the set of values at which percentiles are estimated by functions **cdf.est and cdf.decon**. The default set of values for **pctval** is: 5, 25, 50, 75, and 95.

Second Level Functions (cat.analysis and cont.analysis)

Data input for the second levels functions can be either an object belonging to class psurvey.analysis, i.e., output from the first level function **psurvey.analysis**, or through use of the other arguments to these functions. When data input is not accomplished through use of an object belonging to class psurvey.analysis, format for data entry is similar to the first level function. When an object of class psurvey.analysis is not provided, then values must be supplied for the **sites**, **subpop**, and **design** data frames plus either the **data.cat** data frame and the **type.cat** vector for function **cat.analysis** or the **data.cont** data frame for function **cont.analysis**. Unlike the first level function, individual design variables cannot be input to the second level function, which means that only the default names are allowed in the **design** data frame, and design variable cannot be input using formulas. The following arguments that were discussed previously can be input to the second level functions: **N.cluster**, **popsize**, **stage1size**, **support**, **swgt**, **swgt1**, **unitsize**, **vartype**, **conf**, and

pctval. In addition, values for arguments **sigma** and **var.sigma** can be input to function **cont.analysis**.

Third Level Functions

Regarding data entry for the third level functions, consult the entry for each function in the R help documentation for **psurvey.analysis**.

Data Analysis Algorithms

Data analysis algorithms are carried out by the third level functions with support of the fourth level functions. Descriptions follows for each type of data analysis carried out by the functions in **psurvey.analysis**. In addition, discussion is provided regarding issues that are common to each type of data analysis.

Categorical Data Analysis (category.est)

Categorical data analysis is carried out by function **category.est**. Proportion estimates are calculated using the Horvitz-Thompson ratio estimator, i.e., the ratio of two Horvitz-Thompson estimators. The numerator of the ratio estimates the size of a category. The denominator of the ratio estimates the size of the resource. When either the size of the resource or the sum of the size-weights of the resource is provided, the classic ratio estimator is used to calculate size estimates, where that estimator is the product of the known value and the Horvitz-Thompson ratio estimator. When neither the size of the resource nor the sum of the size-weights of the resource is provided, the Horvitz-Thompson estimator is used to calculate the size estimates.

CDF and Percentiles Estimation (cdf.est and cdf.decon)

Function **cdf.est** carries out CDF and percentile estimation. Function **cdf.decon** carries out estimation of the deconvoluted CDF and percentile based on the deconvoluted CDF. The simulation extrapolation deconvolution method (Stefanski and Bay, 1996) is used to remove the effect of measurement error variance from the CDF of the response variable. When function **cdf.est** or **cdf.decon** is called directly, the user can supply the set of values at which the CDF is estimated. For the CDF of a proportion, the Horvitz-Thompson ratio estimator is used to calculate the CDF estimate. For the CDF of a total when either the size of the resource or the sum of the size-weights of the resource is provided, the classic ratio estimator is used to calculate the CDF estimate. For the CDF of a total when neither the size of the resource nor the sum of the size-weights of the resource is provided, the Horvitz-Thompson estimator is used to calculate the CDF estimate. In addition, the functions use the estimated CDF to calculate percentile estimates and approximate confidence bounds for the percentile estimates.

CDF Inference (cdf.test)

Function **cdf.test** carries out inference regarding the difference between two CDFs. The user supplies the set of upper bounds for defining the classes for the CDFs. The Horvitz-Thompson ratio estimator is used to calculate estimates of the class proportions for the CDFs. Note that function **cdf.test** currently is not written to handle either stratified designs or two-stage designs.

Population Total, Mean, Variance, and Standard Deviation Estimation (total.est)

Estimation of the population total, mean, variance, and standard deviation is carried out by function **total.est**. The Horvitz-Thompson estimator is used to calculate the total, variance, and standard deviation estimates. The Horvitz-Thompson ratio estimator is used to calculate the mean estimate.

Relative Risk (relrisk)

Estimation of relative risk is carried out by function **relrisk**. The relative risk estimate is computed using the ratio of a numerator probability to a denominator probability, which are estimated using cell and marginal totals from a 2x2 table of cell counts defined by a categorical response variable and a categorical stressor variable. An estimate of the numerator probability is provided by the ratio of the cell total defined by the first level of response variable and the first level of the stressor variable to the marginal total for the first level of the stressor variable. An estimate of the denominator probability is provided by the ratio of the cell total defined by the first level of response variable and the second level of the stressor variable to the marginal total for the second level of the stressor variable. Cell and marginal totals are estimated using the Horvitz-Thompson estimator. The standard error of the log of the relative risk estimate is calculated using a first-order Taylor series linearization (Sarndal et al., 1992).

Analysis of Stratified Designs

For a stratified design, separate estimates and standard errors are calculated for each stratum, which are used to produce estimates and standard errors for all strata combined. Strata that contain a single value are removed. For a stratified design, when either the size of the resource or the sum of the size-weights for the resource is provided for each stratum, those values are used as stratum weights for calculating the estimates and standard errors for all strata combined. For a stratified design when neither the size of the resource nor the sum of the size-weights of the resource is provided for each stratum, estimated values are used as stratum weights for calculating the estimates and standard errors for all strata combined.

Analysis of Two-Stage Designs

For a two-stage design, both stages must be accommodated for calculating estimates and standard errors. For calculation of estimates, the product of the stage one and stage two weights is utilized in

the estimation process. For estimation of standard errors, the total and variance of the total is calculated for each stage one sampling unit, where the stage two weights are used in the estimation process. Next, variance of the stage one sampling unit totals is calculated using the stage one weights in the estimation process. Then the weighted sum of the estimated variance of the stage one sampling unit totals is calculated using the stage one weights. The standard error estimate is obtained by adding the estimated variance of the stage one sampling unit totals and the weighted sum of the estimated variance of the stage one sampling unit totals. Depending upon the quantity being estimated, e.g., a proportion estimate, the standard error estimate is scaled by an appropriate factor.

Function Output

First Level Function (psurvey.analysis)

This function outputs a list of class psurvey.analysis. Only those sites indicated by the logical variable in the **sites** data frame are retained in the output. The **sites**, **subpop**, and **design** data frames will always exist in the output. At least one of the **data.cat** and **data.cont** data frames will exist. Depending upon values of the input variables, other elements in the output list may be NULL. The output list is composed of the following elements: (1) the **sites** data frame; (2) the **subpop** data frame; (3) the **design** data frame; (4) the **data.cat** data frame; (5) **type.cat** – the type of categorical response variables; (6) the **data.cont** data frame; (7) **N.cluster** – the number of stage one sampling units in the resource; (8) **popsize** – the known size of the resource; (9) **stage1size** – the known size of the stage one sampling units; (10) **support** – the support for each sampling unit; (11) **swgt** – the size-weight for each site; (12) **swgt1** – the stage one size-weight for each site; (13) **unitsize** – the known sum of the size-weights of the resource; (14) **stratum.ind** – a logical value that indicates whether the sample is stratified, where TRUE indicates a stratified sample and FALSE indicates not a stratified sample; (15) **cluster.ind** – a logical value that indicates whether the sample is a two-stage sample, where TRUE indicates a two-stage sample and FALSE indicates not a two-stage sample; (16) **pcfactor.ind** – a logical value that indicates whether the population correction factor is used during variance estimation, where TRUE indicates use the population correction factor and FALSE indicates do not use the factor; (17) **swgt.ind** – a logical value that indicates whether the sample is a size-weighted sample, where TRUE indicates a size-weighted sample and FALSE indicates not a size-weighted sample; (18) **vartype** – the choice of variance estimator; (19) **conf** – the confidence level; and (20) **pctval** –the set of values at which percentiles are estimated.

Second Level Functions (cat.analysis and cont.analysis)

Function **cat.analysis** outputs a data frame of population estimates for all combinations of subpopulation Types, subpopulations within Types, response variables, and categories within each response variable. The data frame provides estimates for proportion and size of the categories. Standard error estimates and confidence interval estimates also are included.

Function **cont.analysis** outputs a list containing either three or five data frames of population estimates for all combinations of population Types, subpopulations within Types, and response variables. The data frames containing deconvoluted CDF estimates and deconvoluted percentile estimates are only included in the output list when input values of measurement error variance are provided to the function. CDF and percentile estimates are calculated for both proportion and size of the population. Standard error estimates and confidence interval estimates also are calculated. The five data frames are: (1) **CDF** – a data frame containing the CDF estimates, (2) **Pct** – a data frame containing the percentile estimates, (3) **CDF.D** – a data frame containing the deconvoluted CDF estimates, (4) **Pct.D** – a data frame containing the deconvoluted percentile estimates, and (5) **Tot** – a data frame containing the total, mean, standard deviation, and variance estimates.

Third Level Functions

Regarding output for the third level functions, consult the entry for each function in the R help documentation for psurvey.analysis.

References

Diaz-Ramos, S., D.L. Stevens, Jr., and A.R. Olsen. 1995. EMAP Statistics Methods Manual. EPA/620/R-96/002, U.S. Environmental Protection Agency, National Health and Environmental Effects Research Laboratory, Corvallis, Oregon.

Kincaid, T.M. 2004. Testing for differences between cumulative distribution functions from complex environmental surveys. Survey Methodology (in revision).

Messer, J.J ., R. A. Linthurst, and W. S. Overton. 1991. An EPA program for monitoring ecological status and trends. Environmental Monitoring and Assessment 17:67-78

Särndal, C.-E., B. Swensson, and J. Wretman. 1992. *Model Assisted Survey Sampling*. Springer-Verlag, New York.

Stefanski, L.A. and J.M. Bay. 1996. Simulation extrapolation deconvolution of finite population cumulative distribution function estimators. Biometrika 83: 496-517.

Stevens, D.L., Jr., and Olsen, A.R. 2003. Variance estimation for spatially balanced samples of environmental resources. Environmetrics 14: 593-610.

Stevens, D.L., Jr. and Olsen, A.R. 2004. Spatially-balanced sampling of natural resources. Journal of American Statistical Association 99: 262-278.