Analyzing Repeated Measures Marginal Models on Sample Surveys with Resampling Methods

James D. Knoke
University of California, San Diego

Christy M. Anderson
University of California, San Diego

Gary G. Koch
University of North Carolina at Chapel Hill

Abstract

Packaged statistical software for analyzing categorical, repeated measures marginal models on sample survey data with binary covariates does not appear to be available. Consequently, this report describes a customized SAS program which accomplishes such an analysis on survey data with jackknifed replicate weights for which the primary sampling unit information has been suppressed for respondent confidentiality. First, the program employs the **Macro Language** and the Output Delivery System (**ODS**) to estimate the means and covariances of indicator variables for the response variables, taking the design into account. Then, it uses `PROC CATMOD` and **ODS**, ignoring the survey design, to obtain the design matrix and hypothesis test specifications. Finally, it enters these results into another run of `CATMOD`, which performs automated direct input of the survey design specifications and accomplishes the appropriate analysis. This customized SAS program can be employed, with minor editing, to analyze general categorical, repeated measures marginal models on sample surveys with replicate weights. Finally, the results of our analysis accounting for the survey design are compared to the results of two alternate analyses of the same data. This comparison confirms that such alternate analyses, which do not properly account for the design, do not produce useful results.

Keywords: categorical data, repeated measures marginal models, sample survey data, resampling designs, jackknifed replicate weights, SAS.

1. Introduction

A linear regression analysis of a nominal, polychotomous, repeated measures marginal model (Koch and Reinfurt 1971; Koch, Landis, Freeman, Freeman, and Lehnen 1977) with simple, random sample data can be performed with the SAS procedure `CATMOD` (SAS Institute Inc

1999). Fitting such a model with two assessments of the response variable and a single binary covariate is discussed in Stokes, Davis, and Koch (2000, Section 14.2.4). This discussion can be generalized to p binary covariates. The following SAS instructions fit such a model with four covariates to a response variable with four categories and two repeated levels:

```
proc catmod;
response marginals;
population Age Educ Sex Quit;
model Resp1 * Resp2 = Age Educ Sex|Quit Quit|_response_
    / noparm nodesign noprofile;
repeated Response 2;
```

The repeated response variable for this model is the *intensity of cigarette smoking*, coded as 1, 2, 3, or 4. Respondents, who were all current cigarette smokers, were asked the *intensity of their smoking*, *one year ago* and *now*. These variables are named `Resp1` and `Resp2` in the dataset. The generic variable in the `repeated` statement can be given any name; `Response` is simply what we chose to use. The two other uses of this word in these instructions, the `response` statement and the special variable `_response_`, are required SAS constructs. The covariates, `Age`, `Educ` (education), `Sex`, and `Quit` (whether an attempt to quit was made in the past year), were all binary and coded as 1 or 2. (Note that the SAS construct `A|B` in the model statement tells the procedure to include the *main effects* for `A` and `B` and their *interaction*.) These instructions created the following output:

The CATMOD Procedure

Data Summary

Response	Resp1*Resp2	Response Levels	16
Weight Variable	None	Populations	16
Data Set	TWO	Total Frequency	14237
Frequency Missing	0	Observations	14237

Analysis of Variance

Source	DF	Chi-Square	Pr > ChiSq
Intercept	3	51615.40	<.0001
Age	3	235.67	<.0001
Educ	3	119.22	<.0001
Sex	3	411.47	<.0001
Quit	3	722.47	<.0001
Sex*Quit	3	22.60	<.0001
Response	3	455.08	<.0001
Quit* Response	3	91.52	<.0001
Residual	72	167.64	<.0001

Certain minimum sample sizes are necessary for this analysis to produce a non-degenerate, approximately (multivariate) normal distribution with a stable, non-singular estimate of the covariance matrix. Ideally, there should be at least 10 observations in each of the *populations* defined by the 16 combinations of the four binary covariates. Several of the 16 cells of the bivariate distribution of the repeated response variable (each with four levels) should also have at least two observations, for each of the 16 *populations* (Koch *et al.* 1977). The sample size for this example is sufficiently large ($n = 14,237$) for these considerations not to be a concern. Users with smaller sample sizes, however, should be careful not break their data down into so many *populations* that these sample size recommendations are violated.

The results of an analysis of this particular model, on sample survey data with jackknife replicate weights, are discussed in Knoke, Anderson, and Burns (2006). Sets of jackknife replicate weights were provided for these data to allow for variance estimation while suppressing primary sampling unit (PSU) information from the data set for respondent confidentiality. While the residual statistic indicates a highly significant lack of fit, the significances of the terms not included in this model are substantially less than those of the included terms. Our other report (Knoke *et al.* 2006) also presents a *full model* that includes all possible terms. The *reduced model* was chosen for convenience of presentation in this report.

Ideally, sample survey data would be analyzed by statistical software specifically designed to incorporate the survey design, such as SUDAAN (Research Triangle Institute 2001) or WesVar (Westat 2002). Unfortunately, neither of these software packages, nor any others known to the authors, has a procedure that analyzes categorical, repeated measures marginal models. However, CATMOD does accept direct input of the design matrix, the desired hypothesis tests, estimated indicator variables for the response variables, and the estimated covariance matrix of the indicator variables. Mean values of the indicator variables and their variances, but not their covariances, can be correctly estimated by SUDAAN. Consequently, we prepared a customized SAS program, using the **ODS** (Output Delivery System) and the **Macro Facility**, to automate the assembly of the direct input for CATMOD necessary to take the survey design into account.

In this report we discuss the analysis of sample surveys based on resampling methods, describe the preparation of direct input for CATMOD for a simple example, and present a customized SAS program. This program analyzes a general categorical, repeated measures marginal model on sample survey data with jackknife replicate weights, via automated direct input to CATMOD. Simple modifications to this program for sample surveys with other types of resampling plans are discussed. We conclude with a comparison of the results of an analysis correctly accounting for the survey design to the results of two alternate analyses of the same data.

2. Sample surveys incorporating resampling methods

Resampling is frequently incorporated in a sample survey to facilitate the suppression of PSU information for subject confidentially. In lieu of this basic design information, sets of replicate weights, in addition to the set of full sample weights, are provided for variance estimation (Shao 1996). Our data were from California Tobacco Surveys (CTSs) conducted in 1990, 1996, and 1999 (http://ssdc.ucsd.edu/tobacco/), which provide delete-cluster, jackknifed replicate weights. In this implementation, for the rth set of replicate weights, the rth cluster has been left out (assigned a weight of zero) and the weights for the remaining clusters have

been adjusted upwards to account for the deleted cluster. For the 1990 survey there were 33 sets of replicate weights provided, and for the 1996 and 1999 surveys there were 51 sets provided. Variance estimation with these unequal numbers of replicate weights was accommodated by the technique described by Messer and Gamst (2006). This technique involved creating 84 sets of pseudo-replicate weights for all observations. For observations coming from the 1990 survey, the 33 sets of actual replicate weights were carried over unchanged into the first 33 sets of pseudo-replicate weights. The final 51 sets were all given the full sample weights. Conversely, for the 1996 and 1999 surveys, the 51 sets of actual replicate weights were carried over unchanged into the last 51 sets of pseudo-replicate weights; the remaining sets were given the full sample weights. The pseudo-replicate weights were inputted to our customized SAS program for marginal model analysis.

We also generalized the formula for the jackknife estimate of the variance (Shao 1996) to a formula for estimating the covariance matrix:

$$\hat{\Sigma} = \sum_{r=1}^{R} M_r (\hat{\theta}_r - \hat{\theta})(\hat{\theta}_r - \hat{\theta})^T,$$

where

- $\hat{\theta}$ is the estimate of the indicator variable vector, weighted by the full sample weights,

- $\hat{\theta}_r$ is the estimate of the indicator variable vector, weighted by the rth replicate weights,

- R is the number of sets of replicate weights, and

- M_r is the *multiplier*, for the rth set of replicate weights.

Had surveys not been combined, the *multiplier* would not have depended on r and would simply have been (R-1)/R. For our situation, however, M_r was taken to be 32/33 for the first 33 sets of pseudo-replicate weights (arising from the 1990 survey) and 50/51 for the last 51 sets of pseudo-replicate weights (arising from the 1996 and 1999 surveys) (Messer and Gamst 2006). Only a subsample of the respondents of the CTSs was analyzed, those who were current smokers at the time of the survey. Fortunately, replication methods of variance estimation are robust to subsample analyses (Brick, Morganstein, and Valliant 2000). However, the *multiplier* is only approximately correct for subsample analysis, because the proportion of smokers varies slightly over clusters.

Besides the delete-cluster jackknife, there are other forms of the jackknife and other survey methods based on resampling, including balanced repeated replications and its variants (Fay 1989). All resampling methods of variance estimation for sample surveys entail sets of replicate weights, and covariance estimation for other resampling methods is analogous to that for the jackknife design, with only the *multiplier* needing modification Westat (2002, Appendix A). In addition, surveys that have not had the PSU information suppressed can be analyzed by resampling methods. In this situation, the user will have to construct his own sets of replicate weights using the PSU information.

3. A simple example: One covariate, two populations

There are several steps in the preparation of direct input for CATMOD. These steps will be illustrated by a simple example derived from the example discussed in the Introduction,

involving only a single covariate, `Quit`. This single binary covariate is said to result in two populations, in **SAS** terminology (SAS Institute Inc 1999; Stokes *et al.* 2000). As before, the response variable has four categories and is assessed at two points in time. Each response variable requires three indicator variables to represent its four categories.

The first step is the creation of a special dataset that contains estimates of the indicator variable means and covariances, taking the sample survey design into account. Although the **SAS** documentation refers to these as response functions, we call them indicator variables, to avoid confusion with the categorical response variables from which they are derived and which were used in the unweighted analysis described in the Introduction. For a response variable with four categories, two replications, and one binary covariate (i.e., two populations), 12 indicator variables are needed. Using a portion of the **SAS** program described in the next section, we obtained estimates of the mean values of these indicator variables and their covariance matrix, accounting for the sample survey design, and assembled them in the **SAS** dataset `forcat`. This dataset must follow a special format to be used as direct input for `CATMOD`.

The dataset `forcat` contains 13 observations. The first observation contains the weighted estimates of the means of the indicator variables. The second through 13th observations contain the estimated covariance matrix of the indicator variables. The dataset also includes two special variables, `_name_` and `_type_`, that are required by **SAS**. `_name_` is a character variable which must be blank for the observation with the mean estimates; the user must specify values for the observations corresponding to the rows of the covariance matrix. We simply chose values `b1-b12`; these values must in turn be used in the **SAS** dataset as the names for the indicator variables. `_type_` must be a character variable with the value `parms`, for the observation with the indicator variable estimates, and with the value `cov`, for each observation corresponding to a row of the covariance matrix. Note that the covariance matrix is block diagonal, with the blocks being of dimension six (the number of indicator variables per population). The off-diagonal blocks would be exactly zero in the case of independent sampling. They were assumed to be approximately zero for this analysis, which seems a reasonable assumption given the robustness of replication methods of variance estimation to subsample survey data analysis (Brick *et al.* 2000).

We then ran an unweighted `CATMOD` analysis with the `design` and `parm` options enabled. The instructions for this run were:

```
proc catmod;
response marginals;
population Quit;
model Resp1 * Resp2 = Quit|_response_ / design noprofile;
repeated Response 2;
```

Selected output from this run is given in Table 1; it includes the 12×12 design matrix. The 12 rows of the design matrix correspond to the 12 indicator variables and occur in blocks of three. There are three rows for each response variable for each population. The 12 columns of the design matrix correspond to the 12 parameters included in the model specifications. These also occur in blocks of three, corresponding to the hypotheses to be tested, each of which has three degrees of freedom. The blocks of three rows and columns result in the design matrix consisting of blocks of 3×3 submatrices, each of which is either the identity

SAS Dataset forcat

name	b1	b2	b3	b4	b5	b6	b7	b8	b9	b10	b11	b12	_type_
b1	0.245	0.197	0.368	0.253	0.245	0.362	0.106	0.147	0.447	0.093	0.179	0.452	parms
b1	5.99-05	-2.53-05	-1.89-05	4.30-05	-1.43-05	-1.68-05	0.000	0.000	0.000	0.000	0.000	0.000	cov
b2	-2.53-05	6.12-05	-2.07-05	-1.31-05	3.52-05	-1.63-05	0.000	0.000	0.000	0.000	0.000	0.000	cov
b3	-1.89-05	-2.07-05	5.69-05	-2.04-05	-1.01-05	4.40-05	0.000	0.000	0.000	0.000	0.000	0.000	cov
b4	4.30-05	-1.31-05	-2.04-05	5.90-05	-2.54-05	-2.56-05	0.000	0.000	0.000	0.000	0.000	0.000	cov
b5	-1.43-05	3.52-05	-1.01-05	-2.54-05	5.92-05	-2.57-05	0.000	0.000	0.000	0.000	0.000	0.000	cov
b6	-1.68-05	-1.63-05	4.40-05	-2.56-05	-2.57-05	7.56-05	0.000	0.000	0.000	0.000	0.000	0.000	cov
b7	0.000	0.000	0.000	0.000	0.000	0.000	2.38-05	-8.89-06	-4.66-06	1.84-05	-4.58-06	-6.85-06	cov
b8	0.000	0.000	0.000	0.000	0.000	0.000	-8.89-06	2.45-05	-1.58-05	-7.61-06	2.20-05	-1.14-05	cov
b9	0.000	0.000	0.000	0.000	0.000	0.000	-4.66-06	-1.58-05	4.36-05	-3.11-06	-1.88-05	3.73-05	cov
b10	0.000	0.000	0.000	0.000	0.000	0.000	1.84-05	-7.61-06	-3.11-06	1.88-05	-6.38-06	-7.71-06	cov
b11	0.000	0.000	0.000	0.000	0.000	0.000	-4.58-06	2.20-05	-1.88-05	-6.38-06	3.11-05	-2.02-05	cov
b12	0.000	0.000	0.000	0.000	0.000	0.000	-6.85-06	-1.14-05	3.73-05	-7.71-06	-2.02-05	4.83-05	cov

Table 1: The forcat data set.

matrix or its negation. The parameter estimates at the bottom of Table 1 indicate which of the ordered parameters correspond to the effects included in the model. The first three parameters correspond to the intercept, which is automatically included by `CATMOD`, although not explicitly listed in the model statement. The remaining parameters, three at a time, correspond to the effects included in the model statement in the order they were listed. There are zero residual degrees of freedom for this analysis because a *full model* is specified in the model statement, unlike the *reduced model* specified in the example given in the Introduction, which included only a subset of all possible effects.

```
CATMOD Output, Illustrating The Design Matrix and Hypothesis Tests


                       The CATMOD Procedure


                         Data Summary

     Response           Resp1*Resp2  Response Levels     16
     Weight Variable    None         Populations          2
     Data Set           ONE          Total Frequency  14238
     Frequency Missing  0            Observations     14238



              Response Functions and Design Matrix


                       Function   Response
               Sample   Number    Function

                  1        1       0.22413
                           2       0.19070
                           3       0.38836
                           4       0.23583
                           5       0.24194
                           6       0.37954


                  2        1       0.09767
                           2       0.14633
                           3       0.45506
                           4       0.08629
                           5       0.17689
                           6       0.45997


              Response Functions and Design Matrix

         Function                      Design Matrix
Sample    Number    1  2  3  4  5  6  7  8  9  10  11  12


   1         1      1  0  0  1  0  0  1  0  0  1   0   0
             2      0  1  0  0  1  0  0  1  0  0   1   0
```

```
        3     0   0   1   0   0   1   0   0   1   0   0   1
        4     1   0   0   1   0   0  -1   0   0  -1   0   0
        5     0   1   0   0   1   0   0  -1   0   0  -1   0
        6     0   0   1   0   0   1   0   0  -1   0   0  -1

  2     1     1   0   0  -1   0   0   1   0   0  -1   0   0
        2     0   1   0   0  -1   0   0   1   0   0  -1   0
        3     0   0   1   0   0  -1   0   0   1   0   0  -1
        4     1   0   0  -1   0   0  -1   0   0   1   0   0
        5     0   1   0   0  -1   0   0  -1   0   0   1   0
        6     0   0   1   0   0  -1   0   0  -1   0   0   1
```

Analysis of Variance

Source	DF	Chi-Square	Pr > ChiSq
Intercept	3	55087.26	<.0001
Quit	3	835.07	<.0001
Response	3	483.81	<.0001
Quit*Response	3	88.70	<.0001
Residual	0	.	.

Analysis of Weighted Least Squares Estimates

Effect	Parameter	Estimate	Standard Error	Chi-Square	Pr > ChiSq
Intercept	1	0.1610	0.00290	3087.85	<.0001
	2	0.1890	0.00305	3847.02	<.0001
	3	0.4207	0.00387	11845.39	<.0001
Quit	4	0.0690	0.00290	567.27	<.0001
	5	0.0274	0.00305	80.61	<.0001
	6	-0.0368	0.00387	90.53	<.0001
Response	7	-0.00008	0.00127	0.00	0.9496
	8	-0.0204	0.00139	216.32	<.0001
	9	0.000977	0.00159	0.38	0.5385
Quit*Response	10	-0.00577	0.00127	20.59	<.0001
	11	-0.00517	0.00139	13.82	0.0002
	12	0.00343	0.00159	4.67	0.0307

Finally, using the output of the unweighted analysis, we assembled the following SAS instructions and ran them on the `forcat` data set:

```
proc catmod    data = forcat;
response read b1-b12;
model _f_ = (
```

```
1  0  0  1  0  0  1  0  0  1  0  0,
0  1  0  0  1  0  0  1  0  0  1  0,
0  0  1  0  0  1  0  0  1  0  0  1,
1  0  0  1  0  0 -1  0  0 -1  0  0,
0  1  0  0  1  0  0 -1  0  0 -1  0,
0  0  1  0  0  1  0  0 -1  0  0 -1,
1  0  0 -1  0  0  1  0  0 -1  0  0,
0  1  0  0 -1  0  0  1  0  0 -1  0,
0  0  1  0  0 -1  0  0  1  0  0 -1,
1  0  0 -1  0  0 -1  0  0  1  0  0,
0  1  0  0 -1  0  0 -1  0  0  1  0,
0  0  1  0  0 -1  0  0 -1  0  0  1
)(
1 2 3 = 'Intercept',
4 5 6 = 'Quit',
7 8 9 = 'Response',
10 11 12 = 'Quit*Response'
)  /   nodesign   noparm   noprofile;
```

The response statement with the `read` function tells `CATMOD` to directly read the estimated indicator variable means and covariances and to expect there to be 12 indicator variables named `b1` through `b12`. The model statement with the `_f_` specification indicates that there will be direct input. The output from these instructions is:

```
                  The CATMOD Procedure

     Response Functions Directly Input from Data Set forcat

                   Analysis of Variance

          Source          DF   Chi-Square    Pr > ChiSq

          Intercept        3     36429.12      <.0001
          Quit             3       568.13      <.0001
          Response         3       224.13      <.0001
          Quit*Response    3        40.95      <.0001

          Residual         0          .            .
```

This output contains an ANOVA table, as does the unweighted analysis output reported earlier in this section, except that instead of giving the *Data Summary* at the top of output, SAS merely says *Response Functions Directly Input from Data Set forcat*. With direct input, `CATMOD` does not have access to the information normally given in the *Data Summary*; it performs an analysis simply based on the directly-inputted design matrix and test specifications. Notice that the magnitude of all the chi-squares statistics reported in this output are less than the corresponding unweighted values and those for `Response` and `Quit*Response` are less than 50% of their corresponding unweighted values in the earlier output.

4. The general case

If a covariate is polychotomous (with q categories, say), it can be represented by $q-1$ binary covariates. With an overall total of p binary covariates there will be $k = 2^p$ populations. If the response variable has m categories, it can be represented by $m-1$ indicator variables. If the response variable is repeated r times, there will be a total of $s = rk(m-1)$ indicator variables, which can become quite a large number. For the example in the Introduction, which has $p = 4$, $k = 16$, $m = 4$, and $r = 2$, we have $s = 96$. Thus, the overall covariance matrix is of dimension s, and the design matrix has s rows. The design matrix for a *full model* would also have s columns. However, if not all effects are included (a *reduced model*) there will be fewer columns than rows. Even for reduced models, however, the design matrix can become unwieldy. For the reduced model described in the Introduction, the design matrix has dimension 96×24.

The SAS program 'WtdCatmod.sas', provided as a separate text file, analyzes a general categorical, repeated measures marginal model on survey data with delete-cluster, jackknifed replicate weights. The user merely needs to specify the dataset and the other parameters detailed in the program's opening comment. The parameter values specified in 'WtdCatmod.sas' are specific to the example described in the Introduction and should be edited by the user to reflect the specifics of his dataset. Note the required coding of the categorical variables and naming of the weights.

The parameter estimates produced by this program were verified by comparison to the available parameter estimates from the DESCRIPT procedure of SUDAAN. The estimates of the mean values of the indicator variables and their variances from the two programs were identical.

5. Other resampling methods

Minor modifications to 'WtdCatmod.sas' are required to analyze survey data employing other resampling methods. Our example, which combines surveys with differing numbers of sets of replicate weights, requires a change to the VarMult variable in the VectorProd data step. The existing instruction:

```
VarMult =  (&NRepWts - 1) / &NRepWts;
```

should be replaced by:

```
if rep <= 51 then VarMult = (51-1)/51;
    else VarMult = (33-1)/33;
```

to implement this particular jackknife scheme (Messer and Gamst 2006).

Other resampling plans similarly require modifications to the VarMult variable. For example, for the *balanced repeated replicate (BRR)* or *half-sample* plan, replace the existing instruction by:

```
VarMult = 1 / NRepWts;
```

The appropriate *multiplier* for other resampling methods is discussed in Westat (2002), Appendix A, where `VarMult` is denoted c (page A-5) and the delete-cluster jackknife is termed *JK1*.

6. Effects of the failure to account for the survey design

It is well known that if the sample survey design is ignored and an unweighted analysis is performed, neither means nor variances will be correctly estimated. If an analysis with normed weights (full sample weights divided by their mean, which results in frequencies summing to the sample size instead of to the population size) is performed, means will be correctly estimated, but variances will not (Brogan 1998). Since variance estimates follow a χ^2 distribution, it is unlikely that χ^2 statistics would be accurately evaluated unless the design is incorporated in the analysis. Table 2 reports results of an unweighted analysis, a normed weighted analysis, and the analysis accounting for the survey design, for the reduced model of the Introduction, and confirms this presumption.

Although the results of all three analyses reported in Table 2 are qualitatively similar (the empirical significance levels of all included variables were < 0.0001), the values of the χ^2 statistics themselves differ widely, by as much as a factor of 2. The empirical significance levels from both the unweighted analysis and the analysis with normed weights consistently (with one exception) underestimate the empirical significance levels which account for the survey design. In other words, the alternate analyses generally overstate the significance of an effect. It is also apparent that the results of the analysis with normed weights are not generally closer to the results correctly accounting for the design than are the results of the unweighted analysis.

In Knoke *et al.* (2006), we also fit a *simple model* for which the design matrix contained only 12 columns and the *full model* for which the design matrix contained 96 columns. The *reduced model* was simply chosen for convenience of presentation in this report. The differences in the values of the χ^2 statistics were more critical for the *full model* than for the *reduced model*, because there were a few statistics with borderline significance ($P \sim 0.05$) in the *full model*.

We conclude, as did Brogan (1998), that the normed weighted approach to the analysis of sample survey data cannot generally be recommended. It appears that an analysis that properly takes the design of a sample survey into account, whether the survey has the PSU

Effect	Unweighted χ^2	Weighted χ^2	Jackknifed χ^2
Age	35.7	218.2	140.6
Educ	119.2	107.1	58.8
Sex	411.5	389.3	280.1
Quit	722.5	789.2	527.5
Sex*Quit	22.7	38.5	26.6
Response	455.1	404.9	319.9
Quit*Response	91.5	69.3	64.3

All χ^2 statistics have 3 degrees of freedom and $P < 0.0001$

Table 2: Results of three different analyses of the same sample survey data.

information intentionally suppressed or not, is necessary for correct interpretation of the data. Until packaged statistical software for survey data analysis incorporates a wider selection of statistical analyses, the need for customized software for categorical marginal (linear) model analysis will continue. It is noteworthy that packaged statistical software for log-linear model analysis of categorical data, which takes sample survey design into account, is available in the current release (8.0) of SUDAAN, via the LOGLINK procedure.

7. Comment

It appears that the analysis of surveys based on resampling methods could be relatively easily incorporated into a broad range of packaged statistical analysis routines by the inclusion of a module that computes the covariance matrix based on sets of replicate weights. This module could be substituted for the usual covariance estimation module when a resampling method is specified by the user. We urge the developers of packaged statistical software to expand the breadth of statistical analyses that can be performed with automated control of the sample survey design. This need is especially apparent for surveys which employ resampling methods, which have become common in recent years due to the heightened need to protect respondent anonymity.

Acknowledgments

References

Brick JM, Morganstein D, Valliant R (2000). "Analysis of Complex Sample Data Using Replication." *Technical report*, Westat, Rockville, MD.

Brogan D (1998). "Software for Sample Survey Data, Misuse of Standard Packages." In P Armitage, T Colton (eds.), "Encyclopedia of Biostatistics," volume 5, pp. 4167–4174. Wiley, New York.

Fay RE (1989). "Theory and Application of Replicate Weighting for Variance Calculations." *Proceedings of the Section on Survey Research Methods of the American Statistical Association*, pp. 212–217.

Knoke JD, Anderson CM, Burns DM (2006). "Does a Failed Quit Attempt Reduce Cigarette Consumption Following Resumption of Smoking? The Effects of Time and Quit Attempts on The Longitudinal Analysis of Self-Reported Cigarette Smoking Intensity." *Nicotine and Tobacco Research*. In press.

Koch GG, Landis JR, Freeman JL, Freeman DH, Lehnen RC (1977). "A General Methodology for the Analysis of Experiments with Repeated Measurement of Categorical Data." *Biometrics*, **33**(1), 133–158.

Koch GG, Reinfurt DW (1971). "The Analysis of Categorical Data from Mixed Models." *Biometrics*, **27**(1), 157–173.

Messer K, Gamst A (2006). "A Note on Combining Delete-k Jackknife Variance Estimators with Differing Numbers of Replicates." Unpublished manuscript.

Research Triangle Institute (2001). *SUDAAN User's Manual, Release 8.0.* Research Triangle Park, NC.

SAS Institute Inc (1999). *SAS/STAT*® *User's Guide, Version 8.* Cary, NC.

Shao J (1996). "Resampling Methods in Sample Surveys (With Discussion)." *Statistics*, **27**, 203–254.

Stokes ME, Davis CS, Koch GG (2000). *Categorical Data Analysis Using The SAS*® *System.* SAS Institute Inc, Cary, NC.

Westat (2002). *WesVar 4.2 User's Guide.* Rockville, MD.

Affiliation:

James D. Knoke
Tobacco Control Policies Project
University of California, San Diego
Suite 310
1545 Hotel Circle South
San Diego, CA 92108, United States of America
Telephone: +1/619/294-3708
Fax: +1/619/220-0228
E-mail: jknoke@ucsd.edu