

BOOTSTRAP CONFIDENCE INTERVALS FOR THE MEAN UNDER IMPUTATION FOR MISSING DATA

Yongsong Qin¹, J.N.K. Rao² and Malgorzata Winiszewska³

ABSTRACT

Missing observations are commonly encountered in data from sample surveys due to nonresponse and imputation is used to compensate for nonresponse. The fractional imputation method leads to consistent imputed estimators of both the population mean and distribution function, unlike deterministic imputation. In this paper we develop asymptotically valid bootstrap confidence intervals for the mean, using the bootstrap percentile and bootstrap empirical likelihood methods, under fractional imputation for missing data. We consider imputation classes formed on the basis of auxiliary variables and assume MCAR mechanism within each class. Performance of the proposed confidence intervals is studied through simulations.

KEY WORDS: Missing data, fractional imputation, imputation classes, confidence intervals, empirical likelihood, bootstrap resampling.

RÉSUMÉ

Les observations manquantes sont couramment observées dans les échantillons d'enquêtes en raison de la non-réponse à une question. L'imputation est utilisée pour compenser la non-réponse. En particulier, l'imputation simple ou fractionnaire est souvent employée. Cet article développe des intervalles de confiance bootstrap asymptotiquement valides pour la moyenne et la fonction de répartition en utilisant le centile bootstrap et les méthodes de vraisemblance empirique bootstrap, sous imputation fractionnaire. L'imputation est faite indépendamment à l'intérieur des classes formées sur la base de variables auxiliaires connues. La performance des intervalles de confiance proposés est étudiée à l'aide de simulations.

MOTS CLÉS : Données manquantes, imputation fractionnaire, classes d'imputation, intervalles de confiance, vraisemblance empirique, bootstrap.

1. INTRODUCTION

Missing observations are commonly encountered in data from sample surveys due to nonresponse and imputation is used to compensate for nonresponse. To improve the accuracy of imputation in practice, units are often divided into homogenous groups, called imputation classes, such that the missing values can be imputed independently, using separate imputation procedures within each class. The sample is divided into classes according to auxiliary variables, which are associated with the variable to be imputed (Brick and Kalton, 1996). The concept of forming imputation classes is related to stratification in survey sampling (however, their goals are different) and stratification techniques can be used to form imputation classes. Haziza and Beaumont (2007) compared different methods to construct imputation classes.

Shao and Sitter (1996) proposed a bootstrap approach for handling imputed data by imputing the bootstrap samples in the same way as the original data set. In this article, we form bootstrap percentile and bootstrap-calibrated empirical likelihood confidence intervals on the mean $\mu = E(Y)$ based on data with imputation classes and propose an adjustment to Shao and Sitter's (1996) bootstrap confidence intervals utilizing the fractional imputation method (Kim and Fuller, 2004).

In Section 2, we establish the asymptotic normality of fractionally imputed estimator $\hat{\mu}$ and construct asymptotically valid bootstrap percentile confidence intervals. In Section 3, we obtain the empirical likelihood ratio statistic, establish its limiting distribution and form asymptotically valid bootstrap-calibrated empirical likelihood confidence interval on μ . We report the results of a small simulation study on the finite sample performance of the proposed confidence interval in Section 3. Proofs are omitted.

¹ Y. Qin, Department of Mathematics, Guangxi Normal University, Guilin, Guanxi, 541004, China, ysqin1965@yahoo.com

² J.N.K. Rao, School of Mathematics and Statistics, Carleton University, Ottawa, Canada, JRao@math.carleton.ca

³ M. Winiszewska, School of Mathematics and Statistics, Carleton University, Ottawa, Canada, malgorzata.winiszewska@servicecanada.gc.ca

1.1 Framework

In this paper, we focus on inference about the mean $\mu = E(Y)$ of a population Y with missing data. We assume no parametric structure on the distribution of Y except that $0 < \text{var}(Y) = \sigma^2 < \infty$. Suppose that the population Y can be divided into subpopulations, called imputation classes, P_s , $s = 1, \dots, S$, according to values of an auxiliary variable X with known distribution, and that $P(Y \in P_s) = W_s > 0$. Further, we suppose that all $\{Y_{si}, i = 1, \dots, n_s, s = 1, \dots, S\}$ have the same distribution as the population Y .

We consider imputation classes and assume a missing completely at random (MCAR) response mechanism within each class with $P(\delta_{si} = 1 | Y_{si}) = P(\delta_{si} = 1) = p_s$, $0 < p_s \leq 1$, where $\delta_{si} = 1$ if Y_{si} is observed and $\delta_{si} = 0$ if Y_{si} is missing in class s . Thus, we have an i.i.d. sample of incomplete data in $P_s = \{(Y_{si}, \delta_{si}), i = 1, \dots, n_s\}$ with random frequencies $\{n_s\}$ such that $\sum_{s=1}^S n_s = n$, $s = 1, \dots, S$. Let $s_{rs} = \{i : \delta_{si} = 1\}$ and $s_{ms} = \{i : \delta_{si} = 0\}$ respectively denote the sets of respondents and non-respondents, and let r_s be the number of respondents, that is $r_s = \sum_{i=1}^{n_s} \delta_{si}$, for class s , $s = 1, \dots, S$.

1.2 Fractional Imputation with Imputation Classes

Imputation is a common method to compensate for unit nonresponse in sample surveys. A variety of imputation methods have been developed and can be divided into two classes depending on whether the imputed data is fixed, given the sample, or random (Kalton and Kasprzyk, 1986). There are advantages and disadvantages to both classes of methods; for example, deterministic methods do not preserve the distribution of the imputed variables, while random methods do. Chen, et al. (2000) indicated that deterministic imputation leads to an inconsistent estimator of the distribution function, unlike random imputation which yields a consistent estimator; on the other hand, random imputation induces imputation variance due to random selection of imputed values, which is not the case for deterministic imputation.

In this article, we use the fractional imputation method (Kalton and Kish, 1984) to deal with missing data. As a compromise to deterministic and random imputation methods, fractional imputation was designed to reduce imputation variance and yet preserve the distribution as in hot-deck imputation. Under this technique, $J \geq 1$ imputed values are randomly selected for each missing observation and a weight equal to a fraction J^{-1} of the original survey weight of each donor is assigned to each imputed value. The value of J is fixed and does not depend on sample size. It can be shown that as J increases, the imputation variance decreases and that the method leads to consistent imputed estimators of both the mean and the distribution function. The disadvantage of this method is that all $J \geq 1$ imputed values have to be stored in the data file for each missing observation, therefore in practice, J is usually small; in addition, confidence intervals require identification flags on the imputed values present in data file, which in practice may be difficult to obtain due to confidentiality reasons (Qin et al., 2008). For the simulation study in Section 4, we focus on the case of single imputation ($J = 1$).

In the case of data with imputation classes, missing data are replaced by values selected randomly from the set of respondents. The selection is done separately within each class and independently across all the classes. In particular under fractional imputation, for class s , we generate $J \geq 1$ imputed values $Y_{sij} = \bar{Y}_{rs} + \varepsilon_{sij}^*$, $j = 1, \dots, J$, for each missing $Y_{si}, i \in s_{ms}$, where $\{\varepsilon_{sij}^*, j = 1, \dots, J\}$ are drawn by simple random sampling with replacement from the set of donor residuals $\{\hat{\varepsilon}_{si} = Y_{si} - \bar{Y}_{rs}, i \in s_{rs}\}$ formed within class s with $\bar{Y}_{rs} = \sum_{i=1}^{n_s} \delta_{si} Y_{si} / r_s$. The fractionally-imputed data consists of $\{(\tilde{Y}_{si}, \delta_{si}); i = 1, 2, \dots, n_s, s = 1, \dots, S\}$, where $\tilde{Y}_{si} = Y_{si}$ if $\delta_{si} = 1$ or $\tilde{Y}_{si} = (Y_{si1} \dots Y_{siJ})$ if $\delta_{si} = 0$.

2. NORMAL APPROXIMATION CONFIDENCE INTERVALS

The estimators of the mean μ under fractional imputation with imputation classes are respectively given by

$$\hat{\mu} = n^{-1} \sum_{s=1}^S \sum_{i=1}^{n_s} \left(\delta_{si} Y_{si} + (1 - \delta_{si}) J^{-1} \sum_{j=1}^J Y_{sij} \right).$$

Let $\bar{Y}_r = n^{-1} \sum_{s=1}^S n_s r_s^{-1} \sum_{i=1}^{n_s} \delta_{si} Y_{si}$, It can be easily shown that $E^*(\hat{\mu}) = \bar{Y}_r$, where E^* denotes the expectation with respect to randomness in the imputation procedure.

2.1 Ordinary Confidence Intervals

The result on the asymptotic normality of $\hat{\mu}$ is summarized in Theorem 2.1.

Theorem 2.1 Assume that $0 < p_s \leq 1$, $W_s > 0$, $0 < \sigma^2 < \infty$, $s = 1, \dots, S$, and that there exists an $\alpha_0 > 0$ such that $E|Y|^{2+\alpha_0} < \infty$. Then, as $n \rightarrow \infty$,

$$n^{1/2} \sigma_{nm}^{-1} (\hat{\mu} - \mu) \xrightarrow{d} N(0, 1) ,$$

where $\sigma_{nm}^2 = \sigma^2 \sum_{s=1}^S W_s p_s^{-1} + \sigma_{2nm}^2$ with $\sigma_{2nm}^2 = \sum_{s=1}^S W_s (1 - p_s) (J r_s)^{-1} \sum_{i \in s_{rs}} \left(Y_{si} - r_s^{-1} \sum_{i \in s_{rs}} Y_{si} \right)^2$.

Let $\hat{W}_s = n_s / n$ and $\hat{p}_s = r_s / n_s$, $s = 1, \dots, S$. We assume that the observed response rates $\hat{p}_s = r_s / n_s$ and frequencies n_s are reported in the data file for $s = 1, \dots, S$. Based on Theorem 2.1, an ordinary normal approximation confidence interval for μ with asymptotically correct coverage probability is $\mu \in \left(\hat{\mu} - z_{\alpha/2} \hat{\sigma}_{1m} n^{-1/2}, \hat{\mu} + z_{\alpha/2} \hat{\sigma}_{1m} n^{-1/2} \right)$ where $z_{\alpha/2}$ is the upper $\alpha/2$ quantile from the standard normal distribution.. Note that for $J > 1$, the individual response identification flags, δ_{si} , are needed in the construction of confidence intervals.

2.2 Bootstrap Confidence Intervals

We use the Shao and Sitter's (1996) method, in which the bootstrap data sets are imputed in the same way as the original data set, to approximate the asymptotic distribution of $n^{1/2}(\hat{\mu} - \mu)$ since the usual bootstrap method leads to invalid results with missing data (Shao and Sitter, 1996).

In the case of fractional imputation with imputation classes, the bootstrap Shao and Sitter's (1996) procedure is as follows:

1. Independently within each imputation class $s = 1, \dots, S$, draw a simple random sample $D_s^* = \{(Y_{b,si}, \delta_{b,si}), i = 1, \dots, n_s\}$ with replacement from the imputed data set $D_s = \{(\tilde{Y}_{si}, \delta_{si}), i = 1, \dots, n_s\}$. When $\delta_{b,si} = 0$, generate $J \geq 1$ imputed values $Y_{b,sij} = \bar{Y}_{b,rs} + \varepsilon_{b,sij}$, $j = 1, \dots, J$, where $\{\varepsilon_{b,sij}, j = 1, \dots, J\}$ are drawn by simple random sampling with replacement from donor residuals $\{\hat{\varepsilon}_{b,si} = Y_{b,sj} - \bar{Y}_{b,rs}, i \in s_{b,rs}\}$ with $\bar{Y}_{b,rs} = \sum_{i=1}^{n_s} \delta_{b,si} Y_{b,si} / \sum_{i=1}^{n_s} \delta_{b,si}$ and $s_{b,rs} = \{i : \delta_{b,si} = 1\}$.
2. Compute the imputed bootstrap estimators of μ from the fractionally-imputed bootstrap data:

$$\hat{\mu}_b = n^{-1} \sum_{s=1}^S \sum_{i=1}^{n_s} \left\{ \delta_{b,si} Y_{b,si} + (1 - \delta_{b,si}) J^{-1} \sum_{j=1}^J Y_{b,sij} \right\}.$$

The usual bootstrap analogue of $\hat{\mu} - \mu$ is given by $\hat{\mu}_b - \hat{\mu}$. Theorem 2.2 states that, in the presence of missing data, we need the adjusted bootstrap pivotal to approximate $n^{1/2}(\hat{\mu} - \mu)$.

Theorem 2.2 Suppose that the conditions in Theorem 2.1 are satisfied, then as $n \rightarrow \infty$,

$$\sup_{x \in R} \left| P_b \left\{ n^{1/2} (\hat{\mu}_b - \bar{Y}_r) \leq x \right\} - P \left\{ n^{1/2} (\hat{\mu} - \mu) \leq x \right\} \right| \xrightarrow{P} 0 .$$

The proposed adjustment to Shao and Sitters's (1996) statistics is given by $\mu_{nb} = \bar{Y}_r - \hat{\mu}$. Further,

$$n^{1/2} \mu_{nb} \xrightarrow{d} N\left(0, \left\{ J^{-1} \sum_{s=1}^S W_s (1 - p_s) \right\} \sigma^2 \right).$$

Note that when J is large, the adjustment becomes negligible, i.e., $n^{1/2} \mu_{nb} = o_p(1)$ as $n \rightarrow \infty$, and so the usual bootstrap statistics could be used. Also under deterministic imputation, $\mu_{nb} = 0$ which means that we could use $\hat{\mu}_b - \mu$ in place of $\hat{\mu}_b - \bar{Y}_r$. However, as we mentioned before, deterministic imputation leads to an inconsistent estimator of the distribution function of Y .

3 EMPIRICAL LIKELIHOOD CONFIDENCE INTERVALS

3.1 Ordinary Confidence Intervals

Empirical likelihood (EL) methods for constructing confidence regions under full response were explored by Owen (1988, 1990). It is known that EL confidence regions respect the range of the parameter space, they are invariant under transformations, and their shapes (symmetry) are determined by the data.

Qin et al. (2008) obtained asymptotically correct EL confidence intervals for marginal parameters under mean, random hot-deck and adjusted random hot-deck imputation methods. In this article, we extend their theory and form EL ratio for μ under fractional imputation with imputation classes.

Let $Z_{si,m}(\mu) = \delta_{si} Y_{si} + (1 - \delta_{si}) J^{-1} \sum_{j=1}^J Y_{sij} - \mu$, where $1 \leq i \leq n_s$ and $1 \leq s \leq S$. The empirical log-likelihood ratio for μ is $l_{n,m}(\mu) = -2 \sum_{s=1}^S \sum_{i=1}^{n_s} \log(np_{si,m})$, where $\{p_{si,m}, 1 \leq s \leq S, 1 \leq i \leq n_s\}$ maximize $l_m(p) = \sum_{s=1}^S \sum_{i=1}^{n_s} \log(np_{si,m})$, subject to the following constraints: $p_{si,m} > 0$, $\sum_{s=1}^S \sum_{i=1}^{n_s} p_{si,m} Z_{si,m}(\mu) = 0$ and $\sum_{s=1}^S \sum_{i=1}^{n_s} p_{si,m} = 1$.

The results on the asymptotic distribution of $l_{n,m}(\mu)$ are stated in Theorem 3.1.

Theorem 3.1 *Suppose that the conditions in Theorem 2.1 are satisfied. Then as $n \rightarrow \infty$,*

$$l_{n,m}(\mu) \xrightarrow{d} c_m \chi_1^2,$$

where $c_m = \sigma_{nm}^2 / \sigma_{2m}^2$, $\sigma_{2m}^2 = \sigma^2 \sum_s W_s \left((1 - p_s) J^{-1} + p_s \right)$.

Note that the EL ratio under imputation is asymptotically distributed as a scaled chi-square variable (Qin et al., 2008), unlike the original result under full response (Owen, 2001).

Using Theorem 3.1, a $(1 - \alpha)$ -level confidence interval on μ with asymptotically correct coverage probability, can be constructed as $\left\{ \mu \mid \left(\hat{\sigma}_2^2 / \hat{\sigma}_{1m}^2 \right) l_{n,m}(\mu) \leq \chi_\alpha^2(1) \right\}$, where $\chi_\alpha^2(1)$ is the upper α quantile of the chi-square distribution with one degree of freedom and $\hat{\sigma}_2^2 = \sum_{s=1}^S \hat{W}_s \left(J^{-1} (1 - \hat{p}_s) + \hat{p}_s \right) (r_s - 1)^{-1} \sum_{i \in S_{rs}} (Y_{si} - \bar{Y}_{rs})^2$, with $\hat{W}_s = n_s / n$, $\hat{p}_s = r_s / n_s$ and

$$\bar{Y}_{rs} = r_s^{-1} \sum_{i \in S_{rs}} Y_{si}.$$

3.2 Bootstrap Calibrated Confidence Intervals

We now approximate the asymptotic distribution of $l_{n,m}(\mu)$ using bootstrap sample data. Let

$Z_{b,si,m}(\hat{\mu}) = \delta_{b,si} Y_{b,si} + (1 - \delta_{b,si}) J^{-1} \sum_{j=1}^J Y_{b,sij} - \bar{Y}_r$ for $s = 1, \dots, S$, $i = 1, \dots, n_s$. The proposed adjusted bootstrap analog of $l_{n,m}(\mu)$ is given by $l_{b,n,m}(\hat{\mu}) = -2 \sum_{s=1}^S \sum_{i=1}^{n_s} \log(np_{si,m})$, where $\{p_{si,m}, 1 \leq s \leq S, 1 \leq i \leq n_s\}$ maximize $\sum_{s=1}^S \sum_{i=1}^{n_s} \log(np_{si,m})$ subject to: $p_{si,m} > 0$, $\sum_{s=1}^S \sum_{i=1}^{n_s} p_{si,m} Z_{b,si,m}(\hat{\mu}) = 0$ and $\sum_{s=1}^S \sum_{i=1}^{n_s} p_{si,m} = 1$.

Theorem 3.2 states that $l_{n,m}(\mu)$ can be approximated by its adjusted bootstrap analog $l_{b,n,m}(\hat{\mu})$.

Theorem 3.2 *Suppose that the conditions in Theorem 2.1 are satisfied. Then as $n \rightarrow \infty$,*

$$\sup_{x \in R} \left| P_b \{l_{b,n,m}(\hat{\mu}) \leq x\} - P \{l_{n,m}(\mu) \leq x\} \right| \xrightarrow{P} 0.$$

This result shows that the ordinary bootstrap EL statistic without adjustments cannot be used to approximate the distributions of $l_{n,m}(\mu)$ unless one lets $J \rightarrow \infty$.

4. SIMULATION STUDY

A simulation was conducted to study the performance of bootstrap confidence interval on the population mean $\mu = E(Y)$, based on singly imputed data with imputation classes. In particular, we compared the performance of the proposed adjusted bootstrap 95% confidence interval versus its ordinary counterparts based on two methods: the bootstrap percentile (BP) and the empirical likelihood (EL). The resulting confidence intervals were examined in terms of their coverage probabilities and their average lengths.

The results are based on 2000 simulations on data imputed using random imputation (i.e., fractional imputation with $J=1$) utilizing $B=3000$ bootstrap repetitions. The standard errors for simulated coverage of the 95% confidence intervals were approximately 0.01 with 2000 simulation runs. We considered three imputation classes, $S=3$, and fixed total sample size $n = n_1 + n_2 + n_3 = 300$, as well as different settings for class response probabilities (p_1, p_2, p_3). Note that under full response, the proposed adjustments cancel out. One should also be aware that as J increases, the effect of the proposed adjustment decreases.

4.1 Data Frame

The data was generated based on the simulation setup presented in Fang et al. (2009). For each simulation, we generated n values of Y from gamma distribution with shape parameter 43 and scale parameter 0.20. The sample data was divided into imputation classes according to the value of an auxiliary variable $X \in \{1, 2, 3\}$ generated by the proportional odds model $\log(P(X \leq j | Y = y) / P(X > j | Y = y)) = j + \beta y$ with $j=1, 2$ and $\beta = -0.4$. Note that the class sample sizes n_1, n_2 and n_3 were different for each simulation run. We assumed that Y_{si} is MCAR within each class, with response flags generated independently in each class from three Bernoulli distributions with corresponding success probabilities $p_s, s=1, 2, 3$.

4.2 Simulations

Ordinary versions of confidence intervals were obtained by ignoring the proposed adjustment μ_{nb} (Section 2.1). Note that the ordinary and adjusted methods produced BP confidence intervals of the same length because the proposed adjustment was present in both upper and lower bound, and so it cancelled out in the calculation of the interval length.

We used a bisection method (Wu, 2005) to obtain $\lambda_{n,m,b}$ as well as to find the lower and upper bounds of the $(1-\alpha)\%$ EL confidence intervals.

4.3 Results

Table 1 displays the coverage probabilities and average lengths of the 95% confidence intervals for the population mean, $\mu = E(Y) = 8.6$. Under full response, the coverage and lengths of the BP and bootstrap EL intervals for the population mean were very good. Generally, in the presence of missing data, the adjusted BP confidence intervals for μ led to very good coverage close to the nominal value of 95%. Compared to the ordinary methods, the adjusted methods resulted in smaller departures from the nominal level for all simulation cases. For all the missing data cases, the ordinary EL led to overcoverage while the ordinary BP produced significant undercoverage. For example, the setting with the lowest class response probabilities, generated coverage of 88% for the ordinary BP versus 97% for the ordinary EL.

In terms of the average lengths of the confidence intervals for μ , the adjusted EL intervals were slightly longer relative to the corresponding BP intervals. The adjusted EL generated shorter confidence intervals compared to their ordinary counterparts; for example, under the setting with response probabilities (.5,.6,.7), the ordinary and adjusted EL average interval lengths were respectively 0.476 and 0.441 (0.437 for BP). Generally for all methods, the resulting confidence intervals were longer under higher non-response.

Table 1: Bootstrap confidence interval coverage probability and average interval length for the mean

(p_1, p_2, p_3)	Coverage (%)				Average Length			
	OrdBP	AdjBP	OrdEL	AdjEL	OrdBP	AdjBP	OrdEL	AdjEL
(1,1,1)	95.2	95.2	95.4	95.4	0.294	0.294	0.296	0.296
(.8,.7,.6)	89.2	94.8	96.2	94.9	0.379	0.379	0.408	0.382
(.6,.7,.5)	88.1	94.7	96.6	94.9	0.410	0.410	0.442	0.413
(.5,.6,.7)	89.5	95.3	96.9	95.4	0.437	0.437	0.476	0.441
(.5,.5,.5)	87.9	95.0	97.1	95.3	0.464	0.464	0.507	0.469

REFERENCES

- Brick, J. M. and Kalton, G. (1996). "Handling missing data in survey research". *Statistical Methods in Medical Research*, **5**, 215 - 238.
- Chen, J., Rao, J.N.K. and Sitter, R.R. (2000). "Efficient random imputation for missing data in complex surveys". *Statistica Sinica*, **10**, 1153-1169.
- Fang F., Hong Q. and Shao J. (2009). "A pseudo empirical likelihood approach for stratified samples with nonresponse". *Annals of Statistics*, Vol 37, No. 1, 371-393.
- Haziza, D. and Beaumont, J-F. (2007). "On the construction of imputation classes in surveys". *International Statistical Review*, **75**, 25-43.
- Kalton, G. and Kasprzyk, D. (1986). "The treatment of missing survey data". *Survey Methodology*, **12**, 1-16.
- Kalton, G. and Kish, L. (1984). "Some efficient random imputation methods". *Communications in Statistics: Series A*, **13**, 1919-1939.
- Kim, J.K. and Fuller, W.A. (2004). "Fractional hot deck imputation". *Biometrika*, **91**, 559-578.
- Owen, A.B. (1988). "Empirical likelihood ratio confidence intervals for a single functional". *Biometrika*, **75**, 237-249.
- Owen, A.B. (1990). "Empirical likelihood ratio confidence regions". *Annals of Statistics* **18**, 90-120.
- Owen, A. B. (2001). *Empirical likelihood*. New York: Chapman and Hall.
- Qin, Y, Rao, J.N.K. and Ren, Q. (2008). "Confidence intervals for marginal parameters under imputation for item nonresponse". *Journal of Statistical Planning and Inference*, **138**, 2283-2302.

Shao, J. and Sitter, R. (1996). "Bootstrap for imputed survey data". *Journal of American Statistical Association*, **91**, 1278-1288.

Wu, C. (2005). "Algorithms and R codes for the pseudo empirical likelihood method in survey sampling". *Survey Methodology*. Vol. 31, No. 2. pp. 239-243.

Wu, C. and Rao, J.N.K. (2006). "Pseudo Empirical Likelihood Ratio Confidence Intervals for Complex Surveys". *The Canadian Journal of Statistics*, **34**, 359-375.