BALANCED HALF-SAMPLE REPLICATION WITH AGGREGATION UNITS

Steven Kaufman, National Center for Education Statistics
Room 422e, 555 New Jersey Ave., Washington, D.C. 20208

KEY WORDS: Half-sample replication, Bootstrapping, Variances, aggregation units

Introduction

Given a list of sampling units (frame), most sample designs select a sample proportional to a known variable and collect data for the selected units (i.e., the sampling and collection unit are the same). One example of this type of design is sampling schools proportional to the number of teachers in the school, and collecting school data from the selected schools. It is well known that balanced half-sample replication provides appropriate variance estimates for such designs. A slightly different type of design is when the collection unit is an aggregation of the sampling units. In the example above, if school district data is also collected from all districts with sampled schools then this is an example of the second type of design. In this case, the school district (collection unit) is an aggregation of the schools (sampling unit) belonging to the district.

The question that this paper addresses is whether balance half-sample replication is appropriate when the collection unit is an aggregation of the sampling units. Using the usual BHR design assumptions (i.e., two units independently selected per stratum, the replicates are fully balanced, and the collection and sampling units are the same) then the following is true concerning BHR:

$$E(V_{BHR}(X_n))=V(X_n)=V((X_{n/2}^1 + X_{n/2}^2)/2)$$

$$=1/2V(X_{n/2}^1)$$

E: is expectation with respect to all possible samples

V_{BHR}: is the BHR variance estimate

X_n: is a linear estimate based on the full sample of n units

V: is the true variance estimator

$X_{n/2}^1$: is the half-sample replicate estimate of X based on the first unit selected within each stratum

$X_{n/2}^2$: is the same thing as $X_{n/2}^1$ except the estimate is based on the second unit selected in each stratum

This says that BHR assumes the true variance is inversely proportional to the sample size. It is this property of BHR that might not be true when the collection unit is an aggregation of the sampling units. If the inverse of the selection probability is used as the weight then the possibility of a biased variance estimator can be seen by looking at the form of the selection probability.

When the selection and collection units are the same, the selection probabilities are usually linear with respect to the sample size. If this is true with the aggregation unit selection probabilities, one might expect BHR to work well too.

The selection probability for an aggregate unit, A, has the following form:

$$P_A=1-\prod_{Ah\in A}(1-P_{Ah})$$

P_{Ah}: is the selection probability for the part of aggregate A that is within stratum h.

If aggregation A is composed of one sampling unit for each of two stratum then:

$$P_A=P_1 + P_2 - P_1P_2$$

P_i: is selection probability for unit i

This selection probability is not linear with respect to the sample size, assuming the P_is are linear in the sample size. Hence, BHR may be inappropriate. However, if P_1P_2 is small relative to P_is then BHR may provide a reasonable approximation. Whether BHR is appropriate, depends on the distribution of the P_As.

One possible alternative to BHR is the bootstrap variance estimator. Bootstrap samples use the full sample to approximate the distribution of the frame. From this approximate frame, bootstrap samples are selected using the initial sample size. Since the bootstrap estimates are based on the initial sample size, the half-sample assumption that the variance is inversely proportional to the sample size is not necessary.

The goal of this paper is to investigate, how well BHR estimates the variance when the collection unit is an aggregation unit. Two weighting schemes will be tested - 1) the inverse of the selection probability and 2) a weighting scheme that is linear with respect to the number of selection units selected in the aggregation unit. In addition, a bootstrap variance estimator will be tested using the inverse of the selection probability as a weight.

I'll show, using simulations with a weight based on the inverse of the selection probability, that for the National Center of Education Statistics' Teacher Demand and Shortage survey, BHR works reasonably well for most states. For eight states, BHR does not provide reasonable variance estimates. For these few states, a bootstrap estimator provides reasonable estimates. Based on simulations, I will also show when the second weighting scheme is used, BHR appears to provide unbiased variance estimates.

Simulations

The Teacher Demand and Shortage Survey (TDS) is one of four linked surveys the Center produces to study the critical aspects of teacher supply and demand, the composition of the administrator and teacher work force, and the status of teaching and schooling generally. School districts, schools, administrators, and teachers are all surveyed through a common sample design. These surveys are called the Schools and Staffing Surveys (SASS) and are designed to provide state estimates. The focus of this paper is the initial sampling unit, schools; and the survey of school districts, the aggregation unit. The simulations will be based on data from two frames - the SASS public school frame and a frame of all matching public school districts.

The school frame will be used to select multiple school samples using a design similar to the SASS school sample. Each school sample will be matched to the district frame to produce the district sample, as is done for the TDS survey. The district frame has teacher, student, graduate, and school counts. From each district sample, estimates and their BHR variances will be produced. The true variances can be estimated and compared to the BHR estimates.

Design of SASS School and TDS Surveys[1]

The school survey uses NCES's public school file as the frame. The frame is stratified by state by school level (elementary, secondary and combined). The school sample is selected using a systematic probability proportionate to size procedure. The measure of size is the square root of the number of teachers in the school. The school districts that include a sampled school comprise the TDS district sample. In order to simplify the computation of the district selection probabilities, the schools are sorted by school district within each stratum.

This design does not satisfy all the BHR assumptions. The selection is done systematically, so units selected within a stratum are not independent and a finite population correction is required. In addition, more than two schools are selected per stratum. To satisfy the BHR assumptions, the simulation sample design is modified from the above TDS design in the following manner:

1) Each sampling stratum is further stratified by substrata. The substrata are chosen to be the set of all schools that could be selected in two consecutive selections from the systematic selection procedure described above (i.e., the schools within two sampling intervals). Two schools will be selected within each substratum.

   To simplify the district weight computation, a district spanning two substrata is place in only one substratum depending on which substratum contains most of the district's measure of size.

   If the cumulative school measures of size within a district is larger than two sampling intervals than the district is subdivided into pseudo-districts that are approximately equal to two sampling intervals. Each of these pseudo-districts comprise a substratum. Such districts are certainty districts. The purpose of this modification is to maintain the original school sample distribution. Otherwise, by selecting at most two schools for very large districts, more of the smaller districts will be selected then in the original design.

2) Two schools are selected with replacement within each substratum. The first selection will be assigned to panel 1, while the second will be assigned to panel 2. This maintains the BHR assumption of independent selections and eliminates any finite population adjustment for the variance that might be part of the original design.

Bootstrap Implications of Simulation Design

When the sample and collection unit are the same, the bootstrap variance estimator for simple random sampling is biased by a factor of $n_h/(n_h-1)$. Since the simulation sample design selects two units per stratum, this bias would be significant with the school as the collection unit. With the district as the collection unit, it isn't clear what the appropriate stratum sample sizes are when determining the bias. If all districts are defined completely within a single substrata then the bias will be large (i.e., the effective district sample will be close to 2 providing a bias close to $2/(2-1)=2$). If all districts are defined across substrata, the bias might be smaller (i.e., if districts are defined within three substratum with a sample of 6 schools then the bias might be close to $6/(6-1)=1.2$). In reality, the district definitions are somewhere in the middle, so the magnitude of the bias is unclear. However, I will assume there is an effective sample size in the "stratum" which can partition states that will or won't be significantly biased.

Weighting

Two weighting schemes will be analyzed - one based on the district selection probability (probability weight), and the other based on the school selection probability (expected hits weight). The sample estimate, BHR variance estimates and the estimate of true variance will be computed for each weighting scheme.

The probability weight for district d (PWT_d) is:

$$PWT_d = 1/(1-((1-p_{de})(1-p_{ds})(1-p_{dc}))^2)$$

$p_{de}: \sum_{i \in S_{de}} p_i$, S_{de} is the set of all elementary schools in district d

p_i: is the selection probability for school i.

$p_{ds}: \sum_{i \in S_{ds}} p_i$, S_{ds} is the set of all secondary schools in district d

$p_{dc}: \sum_{i \in S_{dc}} p_i$, S_{dc} is the set of all combined schools in district d

If p_{de}, p_{ds} or p_{dc} is greater than or equal to one then the district is selected with certainty and $PWT_d = 1$.

The expected hit weight for district d (HWT_d) is:

$$HWT_d = H_d/\sum_{i \in S_d} p_i$$

H_d: is the number of schools selected in district d

S_d: is the set of all schools within district d

The unbiasness of this weight follows from the fact that the expectation of the numerator (the expected number of schools selected within a district) is equal to the denominator (the sum of all school selection probabilities within a district).

BHR should be unbiased using the expected hits weight because any linear district estimate can be written as a normalized school estimate.

Let X_d be a district variable and suppose we want to estimate the total value of X_d within the set of all districts in some set U, say all urban districts then:

$$\sum_{d\in U} HWT_d\ X_d = \sum_{d\in U} (H_d/\sum_{i\in S_d} p_i)\ X_d$$

$$= \sum_{i\in Uk_d} (p_i/ \sum_{i\in S_d} p_i)\ X_d/p_i,$$ Uk_d is the set of all sampled schools within districts in U

This is now written as a school estimate weighted by $1/p_i$, normalized by $p_i/\sum_{i\in S_d} p_i$, and where X_d is assigned to every school within district d. Since this is a school estimate, the BHR variance estimate should be unbiased.

Balanced Half-sample Replicates

The selected schools are placed into half-sample replicates using the usual textbook methodology[2]. The r^{th} district half-sample replicate is defined to be the set of districts that have schools in the r^{th} school half-sample replicate. Since the SASS replicates are based on 48 replicates, the simulations will be based on 48 replicates. The district replicate weights are:

For the probability weight, the replicate weight is:

$$RPWT_d = 1/(1-((1-p_{de}/2)(1-p_{ds}/2)(1-p_{dc}/2))^2)$$

The probabilities are divided by 2 because with half the sample, each school has half the chance of being selected.

For the expected hits weight, the replicate weight is:

$$RHWT_{ds} = H_{rd}/\sum_{i\in S_d} (p_i/2)$$

H_{rd}: is the number of schools within replicate r and district d.

District Bootstrap Samples

The idea behind the bootstrap samples is to use the sample weights from the selected units to estimate the distribution of the school and district frames. From the estimated bootstrap frame, B bootstrap samples can be selected using the simulation TDS design. For each selected school i in district d the weights say, PWT_d districts should be generated for the bootstrap frame. The PWT_d districts should have a total cumulative school measure of size equal to W_ip_i, where W_i is the school sampling weight $(1/p_i)$. The bootstrap frame and selection are described below for a specific sample.

1) Generate a file of selected schools. If a school is selected twice, it is on the file twice.
2) Divide each school into PWT_d bootstrap-districts (indexed by bd), each with a $(W_ip_i)/PWT_d$ school bootstrap measure of size. If PWT_d is not an integer then the bootstrap-district representing the noninteger part has a $C_dW_ip_i/PWT_d$ school bootstrap measure of size, where C_d is the noninteger part of PWT_d.
   If a selected district has selected schools in the elementary and secondary strata then the bd^{th} bootstrap-district generated in the elementary stratum should match to the bd^{th} bootstrap-district in the secondary stratum. This relationship should exist for all school levels that are selected for the district. Since this relationship exists for the selected districts, it is important to reflect it in each bootstrap-district.
   The sum of the school bootstrap measures of size for school i is W_ip_i, which is the appropriate representation based on the school weight. The number of districts being represented by school i is PWT_d, which is the appropriate representation based on the district weight.
   Each bootstrap-district within a stratum could be divided into W_i bootstrap-schools. Since the school is only a selection unit, not a unit of analysis, it's only required to know which district is selected and not which school. To compute the bootstrap-district weights, the bootstrap-school selection probabilities would be summed within a bootstrap-district, anyway. This would yield the same results as the procedure described above. One method is computational less intensive.
3) Using the frame generated in step 2 and assuming two units are independently selected, within the substrata, proportional to the bootstrap measures of size, compute the bootstrap-district weight, $BPWT_{bd}$. Let u denote a selection unit on the frame and p_u the selection probability for u.

   If pd is a bootstrap-district representing an integer part of PWT_d then:

$$BPWT_{bd} = 1/(1-((1-p_{bde})(1-p_{bds})(1-p_{bdc}))^2)$$

p_{bde}: $\sum_{u\in S_{bde}} p_u$, S_{bde} is the set of all elementary units in bootstrap-district pd

p_{bds}: $\sum_{u\in S_{bds}} p_u$, S_{bds} is the set of all secondary units in bootstrap-district bd

p_{bdsc}: $\sum_{u\in S_{bdc}} p_u$, S_{bdc} is the set of all combined units in bootstrap-district bd

If p_{bde}, p_{bds} or p_{bdc} is greater than or equal to one then the bootstrap-district is selected with certainty and $BPWT_{bd} = 1$.

If pd is a bootstrap-district representing a non-integer part of PWT_d then:

$$BPWT_{bd} = C_d/(1-((1-p_{bde})(1-p_{bds})(1-p_{bdc}))^2)$$

4) With the frame and bootstrap selection probabilities define in step 2, independently select two units per substratum proportional to bootstrap measures of size. The weights for the selected bootstrap-districts are defined in step 3.

5) Since the available data is defined by the districts selected in the original sample, a bootstrap-district weight indexed by d ($BPWT_d$) is required:

$BPWT_d = \sum_{bd\in S_{dB}} BPWT_{bd}$, S_{dB} is the set of all bd∈d selected in the B^{th} bootstrap

6) Repeat steps 4 and 5 until B bootstrap

samples are selected. Since there are 48 balanced half-sampled replicates, there will be 48 bootstrap samples.

Sample Estimate

For each of the simulation samples, totals, averages and ratios will be computed within each of the fifty states and the District of Columbia. The averages are average number teachers per district and average number of schools per district. The ratios are the ratio of the number pupils to the number teacher and ratio of the number of teachers to number of schools. The totals are number of student, teachers, graduates, schools and districts. The student, teacher and graduate totals are highly correlated with the measure of size, while the school and district totals have a lower correlation with the measure of size. For each of the 90 simulation samples, 459 sample estimates and respective sample variances are computed (51 states * 9 estimates). The average of these estimates across the 90 simulations is an estimate of the expectation of the respective sample estimate. It is these averages that are the building blocks of this analysis. An estimate of the true variance for the sample estimates can be obtained by computing the simple variance of the sample estimates across the 90 simulations. The expected values for the sample variances can now be compared with the estimate of the true variances.

A number of other analysis statistics are required. They are described below.

Analysis Statistics

Confidence Coefficient

To measure the accuracy of the variance estimates, a one sigma two tailed confidence coefficient is computed by determining what proportion of the time the population estimate is within the respective confidence interval. If the variance estimates are appropriate then the confidence coefficients should be close 0.68. One sigma confidence coefficients are used because there aren't enough simulations (90 of them) to accurately measure the 5% tail.

BHR Bias Indicator

The main task of this paper is to measure whether the BHR assumption, that the true variance is inversely proportional to the sample size, is violated. If it is violated, what is the impact on the district variance estimates. The following statistic can be used to partition states into those that will or won't be significantly biased.

BHR Bias Indicator = $V/(1/2(V_1 + V_2)/2)$

- V: is the simple variance of the 90 simulation sample estimates that are computed from districts selected in both panels, using the methodology of interest (PWT_d or HWT_d weighting schemes),
- V_1: is the simple variance of the 90 simulation sample estimates that are computed from the districts selected in panel 1, using the methodology of interest,
- V_2: is the simple variance of the 90 simulation sample estimates that are computed from the districts selected in panel 2, using the methodology of interest.

The numerator is an estimate of the true variance and the denominator is another estimate of the true variance assuming the true variance is inversely proportional to the sample size. Hence, the ratio should be close to one when the true variance is proportional to the sample size.

Within each state, this ratio is computed for each of the nine sample estimates. An average state ratio is then computed using the weights described in the 'Weighting the Estimate' section below. When producing the tables, this state average is assigned to each of the nine state estimates and associated statistics.

In the tables, B, the bias indicator is partitioned into three sets:

Bias Indicator (B)	Expected Result
B > 1.05	BHR underestimate the variance
1.05 > B > 0.95	BHR provides appropriate variance estimates
B < 0.95	BHR overestimates the variance

Bootstrap Bias Indicator

As stated before, the concern with the bootstrap variance estimator is its biased when two units are selected within a stratum, as is the case, with the simulation design. With districts, this bias is difficult to measure. However, the bias should be larger in states that have more districts solely defined in only one substratum. To measure this, the proportion of each state's districts that are totally within a single substratum is calculated.

In the tables, the proportion of districts in 1 stratum is divided into three groups:

Bootstrap Bias Indicator (B)	Expected Results
B > 0.2	most bias
0.2 > B > .08	some bias
B < 0.08	least bias

This proportion will be used as a potential bias indicator for the bootstrap variance estimates.

S_p/σ_p, S_e/σ_e and S_b/σ_b

Besides the confidence coefficient, the ratio of the average estimated standard error (probability or expected hits weight with BHR; or probability weight with Bootstrapping), across the 90 simulation, over the estimated true standard error (probability or expected hits weight with BHR; or probability weight with Bootstrapping) is another measure of the accuracy of the variance estimates.

σ_p/σ_e

Since the accuracy of BHR variance estimation for two weighting schemes are being compared, it is important to know which has the smaller standard error, irrespective of whether the BHR techniques work. This is done by looking at the ratio of the estimated true standard error using the probability weight (σ_p) divided by estimated true standard error using the expected hits weight (σ_e).

Weighting the Estimates

Each of the statistics described above is computed for the nine estimate within each of the 51 states. These 459 estimates with their respective sample variance estimates, estimated true variances and other statistics are summarized by type of estimate - averages, ratios and totals. Since there are differential numbers of these type of estimates (five totals for every two averages and ratios), an important consideration is how these estimates should be weighted. Within a state, the estimates are equally weighted by estimate type with high and lower correlated totals being weighted equally. All summary statistics in the tables are weighted averages using the weights describe above.

Results

Probability Weight and BHR Variances

When analyzing table 1, BHR overestimates the true standard errors. Where the bias is expected to be positive, averages, ratios and total all have a large upward bias, with confidence coefficients as high as 81%, on average. Where the bias is not expected to be positive, there is still a positive bias, but at a more acceptable level. The confidence coefficients using the true standard errors are all close to 0.68, so the difference is caused from the BHR procedures and not from the distribution of the estimates. There are eight states, where the bias is expected to be positive.

Expected Hit Weight and BHR Variances

Table 2 is the same as tables 1 except the expected hits weight is used instead of the probability weight. The results for the expected hits weight are very different than the probability weight results. There doesn't appear to be any significant bias with the expected hits weight. The BHR standard errors are all close to the true standard error. The largest difference occurs where the bias is expected to be positive, in which case, BHR overestimates the standard error by 10% for averages. The BHR confidence coefficients are all close to the coefficients based on the estimated true standard error.

Probability and Expected Hits Weight

Since the BHR variance estimate are less bias using the expected hits weight, one must ask whether estimates using the expected hits weight are as reliable as the estimates based on the probability weight. If the answer is yes, then the expected hits weight should be used instead. Table 1 and 2 provide the ratio of the true standard error using the probability weight over the true standard error using expected hits weight (σ_p/σ_e).

For averages and ratios the probability weight estimates have smaller standard errors. In table 1, the gains in precision range from an average of 2% to an average of 19% over the precision of expected hits weight estimates. In table 2, the gains range from an average of 3% to an average of 12%.

For totals, the expected hit weight estimates have smaller standard errors. In table 1, the probability weight estimate's precision ranges from an average of 12% to an average of 48% larger than the expected hits weight's precision. In table 2, the probability weight's loss of precision ranges from an average of 26% to an average of 34%.

For two out of the three types of estimates the probability weight estimates are better than the expected hits weight estimates. Overall, the probability weight is better. However, if totals are the primary interest then the expected hits weight provides the best estimates. Since none of the totals in the simulation study are uncorrelated with the selection measure of size, performance of such totals is unknown. If the expected hits weight performs poorly with uncorrelated totals, it may not be advisable to use the expected hits weight.

Probability Weight and Bootstrap Variances

Table 3 uses the proportion of districts within a state that are solely in 1 stratum as a bias indicator, to compare the bootstrap standard error to the true standard error. Where the bias is expected to be smallest, the bootstrap standard error estimator using the probability weight provides good standard error estimates. The bootstrap estimates are on average 10%, 6% and 2% smaller than the estimated true standard error respectively for averages totals and ratios. The confidence coefficients are 0.66, 0.72 and 0.71 on average. Where the bias is not expected to be smallest, the bootstrap estimator doesn't do as well, and underestimates the true standard error.

As stated before the BHR variance estimator does not work in eight states. In these eight states, the bootstrap variance estimator did work well.

Overall, the bootstrap standard error estimates preform poorly. This poor performance seems to be related to the inherent n/(n-1) bias of bootstrap variance estimator. When a state has all of their district solely in one stratum, this bias will be large because the sample design only selects two units per stratum. This implies that the bootstrap variance will be 1/2 the true variance. When a state has few district solely in one stratum, the bias is small and the result show this.

Using a sample design that selects more than two units per stratum should improve the bootstrap variance estimates.

Distribution of District's Selection Probabilities

In the introduction, I suggested the selection probability distribution would determine whether BHR would provide reasonable variance estimates. For this simulation, BHR does not work well when more than 20% of the district selection probabilities are larger than 0.95. Other surveys, using an aggregation collection methodology, should review the selection probabilities of the aggregation units. If more than 20% of the probabilities are larger than 0.95 then the BHR variances may be biased.

Conclusions

This simulation study has shown that when the collection unit is an aggregation of the selection units then BHR may not provide reasonable variance estimates. If the weight is based on the aggregation unit's selection probability then the bias can be large when more than 20% of the probabilities are larger than 0.95. BHR assumes that the true variance is inversely proportional to the sample size. This assumption is not necessary true with this design and it appears that the violation of this assumption is the cause of BHR bias in eight states.

If the expected hits weight is used then the variances do not appear to be biased. However, average and ratio estimates derived using the expected hits weight are not as precise, as estimates based on the aggregation unit's selection probability. If totals are the only estimates of importance then the expected hits weight is better.

Using the simulation design and bootstrap procedure described in this paper with the probability weight, some state's variances are best using the bootstrap methodology. In these states, the effective sampling sizes in the "stratum" are large enough to introduce only a small bias. However, bootstrapping did not work for most states because the effective sample sizes in a "stratum" are to small.

Ongoing Activities

Currently, I am trying to extend the bootstrap methodology to a systematic probability proportionate to size selection procedure where n (>2) schools are selected per stratum. With a larger stratum sample size, I'm hoping the bootstrap bias will be smaller. So far, the preliminary results are encouraging when compared with BHR variance estimates.

References

1. Kaufman, Steven (1991). 1988 Schools and Staffing Survey Sample Design and Estimation. Technical Report NCES 91-127, U.S. Department of Education, National Center for Education Statistics.

2. Wolter, K. M.(1985). Introduction to Variance Estimation. New York: Springer-Verlag, p110-145

Table 1 -- Probability weight BHR standard errors by estimate type and bias indicator

Bias Indicator	Estimate Type	σ_p / σ_e	s_p / σ_p	One s_p Conf. Coeff	One σ_p Conf. Coeff
Neg Bias	average	0.98	1.04	0.67	0.68
Unbiased	average	0.96	1.10	0.69	0.67
Pos Bias	average	0.94	1.24	0.76	0.68
Neg Bias	total	1.48	1.06	0.69	0.69
Unbiased	total	1.35	1.15	0.73	0.68
Pos Bias	total	1.12	1.43	0.81	0.68
Neg Bias	ratio	0.94	1.12	0.71	0.69
Unbiased	ratio	0.89	1.18	0.74	0.68
Pos Bias	ratio	0.81	1.42	0.80	0.68

Table 2 -- Expected hits weight BHR standard errors by estimate type and bias indicator

Bias Indicator	Estimate Type	σ_p / σ_e	s_e / σ_e	One s_e Conf. Coeff	One σ_e Conf. Coeff
Neg Bias	average	0.93	0.94	0.60	0.66
Unbiased	average	0.97	1.03	0.67	0.68
Pos Bias	average	0.96	1.10	0.70	0.68
Neg Bias	total	1.31	0.98	0.65	0.67
Unbiased	total	1.34	1.03	0.68	0.68
Pos Bias	total	1.26	1.07	0.68	0.67
Neg Bias	ratio	0.88	0.95	0.64	0.68
Unbiased	ratio	0.88	1.02	0.66	0.69
Pos Bias	ratio	0.88	1.07	0.69	0.67

Table 3 -- Bootstrap standard errors by estimate type and bias indicator

Bias Indicator	Estimate Type	Bias Indic. Size	s_b / σ_b	One s_b Conf. Coeff	One σ_b Conf. Coeff
High Bias	average	0.48	0.79	0.56	0.69
Some Bias	average	0.14	0.77	0.53	0.67
Least Bias	average	0.02	0.90	0.66	0.73
High Bias	total	0.48	0.92	0.63	0.69
Some Bias	total	0.14	0.97	0.65	0.68
Least Bias	total	0.02	1.06	0.72	0.71
High Bias	ratio	0.48	0.85	0.58	0.70
Some Bias	ratio	0.14	0.91	0.63	0.69
Least Bias	ratio	0.02	1.02	0.71	0.72