A noninformative Bayesian approach to small area estimation

Glen Meeden*
School of Statistics
University of Minnesota
Minneapolis, MN 55455
glen@stat.umn.edu

September 2001
Revised May 2002

*Research supported in part by NSF Grant DMS 9971331

SUMMARY

In small area estimation one uses data from similar domains to estimate the mean in a particular small area. This borrowing of strength is justified by assuming a model which relates the small area means. Here we suggest a noninformative or objective Bayesian approach to small area estimation. Using this approach one can estimate population parameters other than means and find sensible estimates of their precision.

AMS 1991 subject classifications Primary 62D05; secondary 62C10.

Key Words and phrases: Sample survey, small area estimation, Polya posterior and noninformative Bayes.

1 Introduction

In the standard approach to small area estimation the parameters of interest, the small area means, are assumed to be related through some type of linear model. Drawing on linear model theory one can derive estimators which "borrow strength" by using data from related areas to estimate the mean of interest. Finding a good estimate of the precision of the estimator is often difficult however. Good recent summaries of the literature can be found in Rao (1999) and Ghosh and Rao (1994).

The Bayesian approach to statistical inference summarizes information concerning a parameter through its posterior distribution, which depends on a model and prior distribution and is conditional on the observed data. In finite population sampling the unknown parameter is just the entire population and the likelihood function for the model comes from the sampling design. A Bayesian must specify a prior distribution over all possible values of the population. Once the sample is observed the posterior is just the conditional distribution of the unobserved units given the the values of the observed units computed under the prior distribution for the population. For most designs this posterior does not depend on the design probability used to select the actual sample. The Bayesian approach to finite population sampling was very elegantly described in the writings of D. Basu. For further discussion see his collection of essays in Ghosh (1988).

Assume that given the sample one can simulate values for all the unobserved units from the posterior to generate a "complete copy" of the population. Then given the simulated and observed values one can compute the value of the population mean, $N^{-1}\sum_{i=1}^{N} y_i$, for this simulated copy of the entire population. By generating many independent simulated copies of the population and in each case finding the mean of the simulated population and then taking the average of these simulated means one has an estimate of the unknown population mean. This process computes approximately the Bayes estimate of the population mean under squared error loss for the given prior. More generally by simulating many such full copies of the population one can compute, approximately, the corresponding Bayes point or interval estimates for many population parameters. The problem then is to find a sensible Bayesian model which utilizes the type of prior information available for the small area problem at hand.

The Polya posterior is a noninformative Bayesian approach to finite population sampling which uses little or no prior information about the population. It is appropriate when a classical survey sampler would be willing to use simple random sampling as their sampling design. In Nelson and Meeden (1998) the authors considered several scenarios where it was assumed that information about the population quantiles of the auxiliary variable was known a priori. They demonstrated that an appropriately constrained Polya posterior, i. e. one that used the prior knowledge about the quantiles of x, yielded sensible frequentist results. Here we will see that this approach can be useful for a variant of small area estimation problems.

We will consider a population that is partitioned into a domain D, of interest, and its complement D'. Also we suppose that it is partitioned into K areas, say $A_1, \ldots, A_K$. Let y be the characteristic of interest and x be an auxiliary variable. Suppose, using a random sample from the entire population, for some k we wish to estimate $\mu_{D,k}(y)$, the mean of y for the all units that belong to the small area $D \cap A_k$. Often the number of sampled units that belong to $D \cap A_k$ is quite small and using just these observations can lead to an imprecise estimator. As an example where this could arise imagine D is a region of a state which is broken up into counties. Each county in D is then paired with a similar county that is outside of D. Hence the kth county and its twin form the kth area and the collection of "twin" counties forms D'. Then a random sample is taken from $D \cup D'$ and one wishes to to estimate the means of the counties, or small areas, making up D.

In order to improve on this naive estimator one needs to make some additional assumptions. Here we will assume that for each unit in the sample we learn both its y and x values. For units belonging to A_k we make two assumptions which formalize the idea that the small areas, $A_k \cap D$ and $A_k \cap D'$, are similar. First we assume that the small area means of the auxiliary variable, $\mu_{D,k}(x)$ and $\mu_{D',k}(x)$, although unknown are not too different. Secondly we assume that for units belonging to A_k the distribution of y_i depends only on its x_i value and not on its membership in D or D'. Finally we assume that $\mu_D(x)$, the mean of x for all the units that belong to D, is known. Note that we do not assume that $\mu_{D,k}(x)$ and $\mu_{D',k}(x)$ are known which is often the case in small area estimation.

Here we will demonstrate that when our assumptions are true a modification of the Polya posterior yields good point and interval estimators of $\mu_{D,k}(y)$ and of the the median of y in the small area $D \cap A_k$. In section two we will briefly review facts about the Polya posterior and in section three discuss simulating from a constrained version of it. In section four we present some simulation results that indicate how it could work in practice. Section five contains some concluding remarks.

2 The Polya posterior

Consider a finite population consisting of N units labeled $1, 2, \ldots, N$. The labels are assumed to be known and to contain no information. For each unit i let y_i, a real number, be the unknown value of some characteristic of interest. The unknown state of nature, $y = (y_1, \ldots, y_N)$, is assumed to belong to some subset of N-dimensional Euclidean space, $\Re^N$. A sample s is a subset of $\{1, 2, \ldots, N\}$. We will let $n(s)$ denote the number of elements in s. A sample point consists of the set of observed labels s along with the corresponding values for the characteristic of interest. If $s = \{i_1, \ldots, i_{n(s)}\}$ then such a sample point can be denoted by (s, y_s).

Given the data the Polya posterior is a predictive joint distribution for the unobserved units in the population conditioned on the values in the sample. Given a data point (s, y_s) we now show how to generate a set of possible values for the unobserved units from this distribution. We consider an urn that contains $n(s)$ balls, where ball one is given the value $y_{s_{i_1}}$, ball two the value $y_{s_{i_2}}$ and so on. We begin by choosing a ball at random from the urn and assigning its value to the unobserved unit in the population with the smallest label. This ball and an additional ball with the same value are then returned to the urn. Another ball is chosen at random from the urn and we assign its value to the unobserved unit in the population with the second smallest label. This second ball and another with the same value are then returned to the urn. This process is continued until all $N - n(s)$ unobserved units are assigned a value. Once this is done we have generated one realization of the complete population from the Polya posterior distribution. This simulated, completed copy contains the $n(s)$ observed values along with the $N - n(s)$ simulated values for the unobserved members of the population. Hence by simple Polya sampling we have a predictive distribution for the unobserved given the observed.

One can verify that under this predicted distribution the expected value of the population mean is just the sample mean and it's posterior variance is approximately the frequentist variance of the sample mean under simple random sampling when $n(s) \geq 25$. Hence inference for the population mean under the Polya posterior agrees with standard methods. Although the design probabilities play no formal role in the inference based on the Polya posterior for it to be appropriate in the judgment of the survey sampler the values for the characteristic of interest for the observed and unobserved units need to be roughly exchangeable. This is usually the case when simple random sampling is used to select the sample.

It has been shown for a variety of decision problems that procedures based on the Polya posterior are admissible because they are stepwise Bayes. (See Ghosh and Meeden (1997).) In these stepwise Bayes arguments a finite sequence of disjoint subsets of the parameter space is selected, where the order is important. A different prior distribution is defined on each of the subsets. First the Bayes procedure is found for each sample point that receives positive probability under the first prior. Next the Bayes procedure is found for each sample point which receives positive probability under the second prior and which was not considered under the first prior. Then the third prior is considered and so on. For a particular sample point the value of the stepwise Bayes estimate is the value for the Bayes procedure for that sample point for the Bayes procedure identified in the step at which the sample point was considered. It is the stepwise Bayes nature of the Polya posterior that explains its somewhat paradoxical properties. Given a sample it behaves just like a proper Bayesian posterior but the collection of possible posteriors that arise from all possible samples comes from a family of priors not from a single prior. From the Bayesian point of view it is appropriate when one's prior beliefs about the population is that the units are roughly exchange but nothing more about them is known. The stepwise Bayesian nature of the Polya posterior also helps to explain why it yields 0.95 Bayesian credible intervals that in most cases behave approximately like 95% confidence intervals.

For more details and discussion on the theoretical properties of the Polya posterior see Ghosh and Meeden (1997). The Polya posterior is related to the Bayesian bootstrap of Rubin (1981). See also Lo (1988).

3 Simulation from the Polya posterior

The interval estimate of the population mean and point and interval estimates for other population quantities under the Polya posterior usually cannot be found explicitly. One must use simulation to find these values approximately. This is done by simulating many independent completed copies of the entire population and calculating the value of the parameter of interest for each copy. One may do this in a straightforward manner but often a well known approximation also works well. For simplicity assume the sample values y_s are all distinct and that the sampling fraction $n(s)/N$ is small. For $j = 1, \ldots, n(s)$ let λ_j be the proportion of units in a complete simulated copy of the entire population which take on the value y_{i_j}. Then under the Polya posterior $\lambda = (\lambda_1, \ldots, \lambda_{n(s)})$ has approximately a Dirichlet distribution with a parameter vector of all ones, i.e. it is uniform on the $n(s) - 1$ dimensional simplex where $\sum_{j=1}^{n(s)} \lambda_j = 1$.

We now assume that there is an auxiliary characteristic associated with each element in the population. For unit i let x_i be the value of this auxiliary characteristic. The vector of these values for the auxiliary characteristic is denoted by x. The values of x are unknown but we assume their population mean is known. This is a common situation and either the regression estimator or the ratio estimator is often used in such cases. Let x_s denote the x values of the observed units in the sample. Now the Polya posterior can be adapted to use this additional information in the following way. When creating a simulated copy of the entire population using the values $\{(y_i, x_i) : i \in s\}$ one only uses completed copies whose simulated population mean of x is equal to the known mean of x.

Simulating from a constrained Polya posterior is more difficult than simulating from the unconstrained Polya. Let μ_x^* denote the known population mean of x. Suppose s is a sample such that x_s contains values smaller and larger than μ_x^*. When this is the case an approximate solution to the problem of generating simulated copies from the Polya posterior distribution which satisfies the mean constraint is available. For $j = 1, \ldots, n(s)$ let λ_j be the proportion of units in the simulated copy of the population which have the value (y_{i_j}, x_{i_j}). (Note the x_s need not be distinct.) If we ignore the constraint for a moment then, as we observed earlier, simulation from the Polya posterior is approximately equivalent to assuming a uniform distribution for $\lambda = (\lambda_1, \ldots, \lambda_{n(s)})$ on the $n(s) - 1$ dimensional simplex where $\sum_{j=1}^{n(s)} \lambda_j = 1$. In order to satisfy the mean constraint we must select λ's at random from the set which is the intersection of the hyperplane $\sum_{j=1}^{n(s)} \lambda_j x_{i_j} = \mu_x^*$ with the simplex for λ. In general one cannot generate independent random samples from this distribution. One may, however, use the Metropolis-Hasting algorithm to generate dependent simulated copies of the population from a convergent Markov chain. For more details on this algorithm see Metropolis, et al (1953) and Hastings (1970).

Using the approximate solution based on the Dirichlet distribution allows one to finesse a bothersome technical problem which has no practical significance. That is given the sample it is often impossible to get simulated copies of the population which satisfy the mean constraint exactly. For example suppose $N = 5$ and our sample of size three yielded x values of 0, 0 and 10. Now if we know $\mu_x = 4.5$ then under the Polya posterior it is impossible to generate simulated copies of the population since the only possible values for an x value of an unobserved unit is 0 or 10. This implies that given this sample under the Polya posterior the only possible values of μ_x are 2, 4 and 6. In general even if we have generated a λ which satisfies the constraint the $\lambda_i N$'s need not be integers and hence their need not be an actual copy of the population corresponding to λ. But in real problems this should not matter very much. For one thing the mean constraint will usually only be known approximately. Furthermore for larger sample sizes the approximate nature of the simulated copies is just not important.

Recently Nelson and Meeden (1998) and Meeden and Nelson (2001) have considered a variety of problems where a constrained Polya posterior is applicable. When the population mean of x is known Meeden and Nelson (2001) presented simulations that demonstrated that the point and interval estimators of the constrained Polya posterior were nearly identical those of the regression estimator. Hence just as the regression estimator does, when estimating the population mean of y the constrained Polya posterior utilizes the information contained in knowing the population mean of x.

4 A small area problem

Consider again the small area estimation problem described in the introduction. A population is partitioned in two different ways. The first partitions the population into a domain of interest, D, and its complement D'. The second partitions it into K areas $A_1, \ldots, A_K$ where for each k we assume that the small areas $A_k \cap D$ and $A_k \cap D'$ are nonempty. Figure 1 gives a graphical representation of the population. A random sample is taken from the whole population and we wish to estimate $\mu_{D,k}(y)$, the mean of y for all the units belong to the small area $A_k \cap D$. For such problems one often assumes that for the auxiliary variable x all the means $\mu_{D,k}(x)$ and $\mu_{D',k}(x)$ are known. Here we make the weaker assumptions that $\mu_{D,k}(x)$ and $\mu_{D',k}(x)$ are unknown but not too different and that $\mu_D(x)$, the mean of x for all the units belonging to D, is known. We also assume that for units belonging to $A_k \cap D$ and $A_k \cap D'$ the distribution of y_i depends only on x_i and does not depend on whether it belongs to D or D'. In terms of Figure 1 we are assuming that the mean of x for all the units in the population which belong to the first column is known and that within each row the distribution of the units across the the two columns is roughly the same. As we will soon see this is enough to produce estimators of $\mu_{D,k}(y)$ which improve on the naive estimator.

Put Figure 1 about here

Before explaining how this is done we need a bit more notation. Let $N_{D,k}$ be the number of units in the population that belong to $D \cap A_k$. We assume that the $N_{D,k}$'s are known. For unit i let $t_i = (1, k)$ if $i \in D \cap A_k$ and $t_i = (0, k)$ if $i \in D' \cap A_k$. Then given a sample s we must use $\{(y_i, x_i, t_i) : i \in s\}$ to estimate $\mu_{D,k}(y)$. The constrained Polya posterior is now constructed in two stages. In the first stage, using the members of the sample that fall into D and their (x_i, t_i) values, we create a completed copy of D which satisfies the known mean constraint $\mu_D(x)$. In the second stage we first find for the simulated copy of D the mean of the x values for all the units belonging to $D \cap A_k$. (Remember that this set contains both observed and simulated values.) Let $\tilde{\mu}_{D,k}(x)$ denote this mean. Next using the observed sample values from $D \cap A_k$ and $D' \cap A_k$ we create a completed copy of $D \cap A_k$ which satisfies the mean constraint $\tilde{\mu}_{D,k}(x)$. By repeating this two staged process many times one can construct simulated copies of $D \cap A_k$ which use the similarity of units within the small areas $A_k \cap D$ and $A_k \cap D'$ and the information from knowing $\mu_D(x)$.

To see how this approach could work in practice we present simulation results for some constructed populations. In all the cases $K = 2$ so there are just two areas and in Figure 1 there are just four cells or four small areas. The populations will be constructed so that there are 250 units in each of the four cells. For each cell we first generate 250 values for the auxiliary variable x by taking a random sample from a gamma distribution with some shape parameter and scale parameter one. Next within each area conditioned on the x values the y values are independent observations from normal distributions where the mean of $y_i|x_i$ depends on x_i and where the the variance of $y_i|x_i$ may be constant or in some cases depends on x_i.

In the first population, pop1, the shape parameter of the gamma distribution was four in both $A_1 \cap D$ and $A_1 \cap D'$ and was six in $A_2 \cap D$ and

$A_2 \cap D'$. For units in A_1 $y_i|x_i$ was normal with mean $25 + 2x_i$ and variance 100. For units in A_2 $y_i|x_i$ is normal with mean $25 + 3x_i$ and variance 25.

Note that pop1 was generated under a model which is consistent with the assumptions underlying the constrained Polya posterior described above. In fact our method should work very well for pop1. This is because for each k the average values of the auxiliary variable in $A_k \cap D$ and $A_k \cap D'$ will be approximately equal. This is not necessary for our approach to work but if it does not work in this example then it is hard to imagine that it could work in practice. In two of the remaining populations for each k we will take the shape parameters generating the values of x in $A_k \cap D$ and $A_k \cap D'$ to be different. This is a more realistic assumption. We will also let the mean of $y_i|x_i$ be a nonlinear function of x_i and let the variance of $y_i|x_i$ depend on x_i. In all cases the form of the distribution of $y_i|x_i$ will be the same across $A_k \cap D$ and $A_k \cap D'$ for each k. This is the most crucial assumption. If this is not satisfied approximately then our method cannot work.

In the second population, pop2, the shape parameters of the gamma distributions were eight in $A_1 \cap D$, ten in $A_1 \cap D'$, six in $A_2 \cap D$ and four in $A_2 \cap D'$. For units in A_1 $y_i|x_i$ was normal with mean $25 + 2x_i$ and variance $9x_i$. For units in A_2 $y_i|x_i$ was normal with mean $25 + 3x_i$ and variance $4x_i$.

In the third population, pop3, the shape parameters of the gamma distributions were eight in $A_1 \cap D$ and $A_1 \cap D'$ and six in $A_2 \cap D$ and $A_2 \cap D'$. For units in A_1 $y_i|x_i$ was normal with mean $25 + 0.5(x_i - 8)^2$ and variance $9x_i$. For units in A_2 $y_i|x_i$ was normal with mean $25 + |x_i - 6|$ and variance $4x_i$.

In the fourth population, pop4, the shape parameters of the gamma distributions were four in $A_1 \cap D$, six in $A_1 \cap D'$, six in $A_2 \cap D$ and eight in $A_2 \cap D'$. For units in A_1 $y_i|x_i$ was normal with mean $25 + 0.5(x_i - 4)^2$ and variance $9x_i$. For units in A_2 $y_i|x_i$ was normal with mean $25 + |x_i - 6|$ and variance $4x_i$.

In the fifth population, pop5, the shape parameters for the gamma distributions were the same as those in pop2. For units in A_1 $y_i|x_i$ was normal with mean $25+0.5(x_i-9)^2$ and variance $9x_i$. For units in A_2 $y_i|x_i$ was normal with mean $25 + |x_i - 5|^{1.5}$ and variance $4x_i$.

For each of these five populations we took 500 random samples of size 80. For each sample we calculated the usual point estimates and 95% confidence intervals for $\mu_{D,1}(y)$ and $\mu_{D,2}(y)$ using just the observations that fell into the small areas. We also found approximately the point estimate and 0.95 credible interval for the constrained Polya posterior. The results are given in Table 1. In each case the constrained Polya posterior estimates were computed using 500 simulated copies of the small area. Then our point estimate is just the average of these 500 computed values and our 0.95 credible interval ranges from the 0.025 quantile to the 0.975 quantile of this set.

We see that the constrained Polya posterior yields significantly better point estimators in every case but one, $\mu_{D,2}(y)$ of pop5. Its intervals are also considerable shorter than the usual. There is some evidence that their frequency of coverage is a bit less than the usual approximate 95% normal theory intervals. In particular this is true for the small area $A_2 \cap D$ in the fifth population.

The results in Table 1 are for the small area means. In Table 2 we give similar results for the small area medians. We compared our estimates to the sample median of the set of the sampled observations that fell into the small area and the usual confidence interval for the median due to Woodruff (1952). Compared to the usual estimators the performance of the constrained Polya posterior estimators for the small area medians is even better than it was for the small area means. In every case its point estimators are better than the sample median. Its interval estimators are always shorter than Woodruff's and for most cases their frequency of coverage seems to be quite close to the nominal 0.95.

Put Tables 1 and 2 about here

5 Concluding remarks

Here we have presented a new method of "borrowing strength" when estimating parameters of a small area of a population. It makes weaker assumptions than those made by the usual approaches to such problems. It is an objective or noninformative Bayesian approach which uses no more prior information than is typically assumed by a frequentist. Simulations indicate that it should be applicable in a variety of situations and should work well especially for some of the problems which roughly satisfy the usual linear model type assumptions, often assumed in small area estimation. It has the advantage of not being restricted to estimating small area means but can estimate other parameters as well. Here we assumed that a certain mean of an auxiliary variable was known. This approach can be extended to when other parameters of an auxiliary variable are known, like the median. Also it should be possible to extend this method to situations where prior information is available for more than one auxiliary variable. In summary we believe that this is flexible approach which can yield point and interval estimators with good frequentist properties for a variety of problems.

References

[1] J. K. Ghosh, editor. *Statistical information and likelihood: A collection of critical essays by D. Basu.* Springer-Verlag, New York, 1988.

[2] M. Ghosh and G. Meeden. *Bayesian Methods for Finite Population Sampling.* Chapman and Hall, London, 1997.

[3] M Ghosh and J. N. K. Rao. Small area estimation: an appraisal. *Statistical Science*, 9:55–93, 1994.

[4] W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57:97–109, 1970.

[5] A. Lo. A Bayesian bootstrap for a finite population. *Annals of Statistics*, 16:1684–1695, 1988.

[6] G. Meeden and D. Nelson. A noninformative Bayesian approach to problems in finite population sampling when information about an auxiliary variable is present. Technical Report, June 2001.

[7] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087–1092, 1953.

[8] D. Nelson and G. Meeden. Using prior information about population quantiles in finite population sampling. *Sankhyā A*, 60:426–445, 1998.

[9] J. N. K. Rao. Some recent advances in model-based small area estimation. *Survey Methodology*, 25:175–186, 1999.

[10] D. Rubin. The Bayesian bootstrap. *Annals of Statistics*, 9:130–134, 1981.

[11] R. S. Woodruff. Confidence intervals for medians and other position measures. *Journal of the American Statistical Association*, 47:635–646, 1952.

Table 1: The average value and the average absolute error for the usual naive small area estimator and the constrained Polya posterior estimator for the small area means. Also given are the length and relative frequency of coverage for their nominal 0.95 intervals for 500 random samples of size 80 from five different populations.

Pop	Small Area	Method	Ave value	Ave aberr	Ave len	Freq. of coverage
pop1	$A_1 \cap D$	usual	33.11	1.84	9.10	0.936
		cstpp	33.20	1.30	6.37	0.934
	$A_2 \cap D$	usual	43.03	1.47	7.78	0.946
		cstpp	43.13	1.03	5.15	0.940
pop2	$A_1 \cap D$	usual	40.39	1.79	8.69	0.932
		cstpp	40.29	1.20	5.62	0.944
	$A_2 \cap D$	usual	42.13	1.48	7.50	0.944
		cstpp	41.97	1.16	5.16	0.912
pop3	$A_1 \cap D$	usual	28.57	1.97	9.85	0.936
		cstpp	28.90	1.47	6.66	0.898
	$A_2 \cap D$	usual	26.71	1.01	5.08	0.940
		cstpp	26.83	0.70	3.24	0.930
pop4	$A_1 \cap D$	usual	27.73	1.27	6.57	0.960
		cstpp	27.64	0.81	4.09	0.940
	$A_2 \cap D$	usual	27.03	0.97	5.33	0.952
		cstpp	27.03	0.65	3.32	0.934
pop5	$A_1 \cap D$	usual	29.25	1.74	9.31	0.942
		cstpp	29.30	1.26	6.16	0.930
	$A_2 \cap D$	usual	27.73	1.08	5.85	0.954
		cstpp	28.82	1.28	4.40	0.850

Table 2: The average value and the average absolute error for the usual naive small area estimator and the constrained Polya posterior estimator for the small area medians. Also given are the length and relative frequency of coverage for their nominal 0.95 intervals for 500 random samples of size 80 from five different populations.

Pop	Small Area	Method	Ave value	Ave aberr	Ave len	Freq. of coverage
pop1	$A_1 \cap D$	usual	33.88	2.01	11.48	0.944
		cstpp	33.25	1.44	7.81	0.930
	$A_2 \cap D$	usual	42.84	1.72	9.94	0.950
		cstpp	42.42	1.35	6.92	0.944
pop2	$A_1 \cap D$	usual	38.94	1.82	9.81	0.940
		cstpp	38.53	1.41	7.47	0.936
	$A_2 \cap D$	usual	40.99	1.77	8.75	0.970
		cstpp	40.33	1.38	6.36	0.914
pop3	$A_1 \cap D$	usual	27.64	1.73	9.52	0.952
		cstpp	27.73	1.24	6.46	0.958
	$A_2 \cap D$	usual	27.03	1.15	6.26	0.954
		cstpp	26.59	0.70	3.76	0.938
pop4	$A_1 \cap D$	usual	27.14	1.27	7.00	0.962
		cstpp	27.05	0.95	5.37	0.966
	$A_2 \cap D$	usual	26.84	1.07	5.99	0.960
		cstpp	26.81	0.78	4.32	0.954
pop5	$A_1 \cap D$	usual	29.10	2.06	11.01	0.956
		cstpp	28.89	1.51	8.28	0.944
	$A_2 \cap D$	usual	27.03	1.14	5.98	0.952
		cstpp	27.87	0.97	4.46	0.900

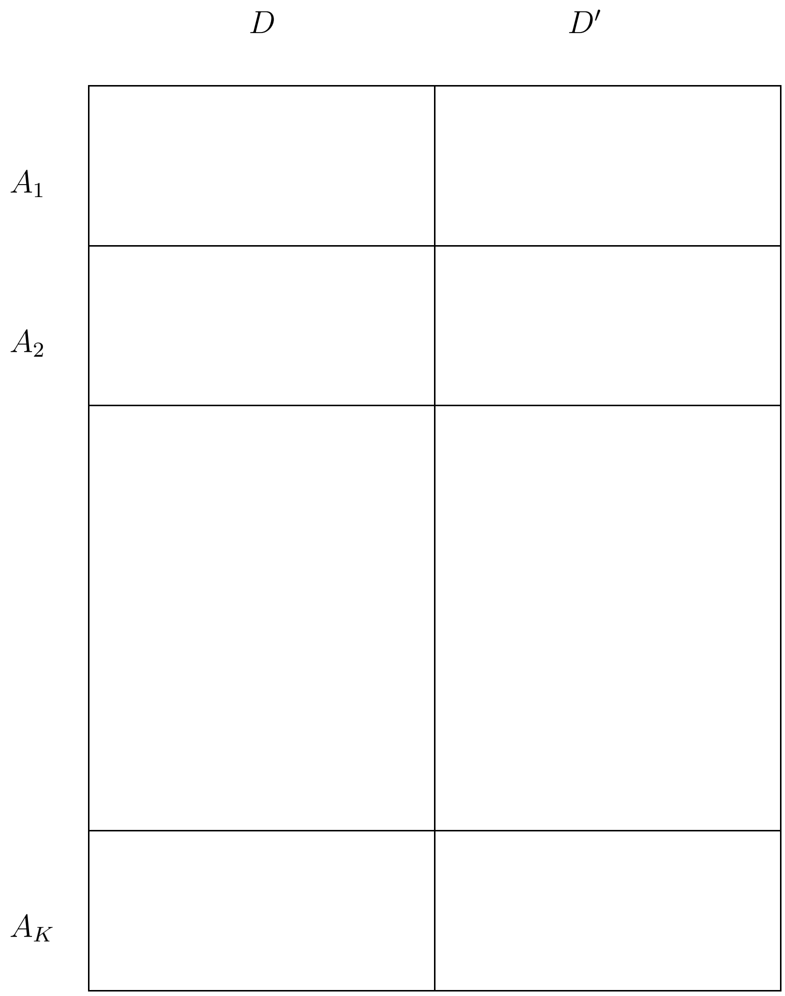

Figure 1: A population partitioned into a domain D and its complement along with a second partition of K small areas.