

UCLA

Department of Statistics Papers

Title

Counting the Homeless in Los Angeles County

Permalink

<https://escholarship.org/uc/item/6d28n81g>

Authors

Berk, Richard

Kriegler, Brian

Ylvisaker, Donald

Publication Date

2005-10-27

Counting the Homeless in Los Angeles County*

Richard Berk, Brian Kriegler, and Donald Ylvisaker
Department of Statistics, UCLA

October 27, 2005

Abstract

Over the past two decades, a variety of methods have been used to count the homeless in large metropolitan areas. In this paper, we report on a recent effort to count the homeless in Los Angeles County, one that employed the sampling of census tracts. A number of complications are discussed, including the need to impute homeless counts to areas of the County not sampled. We conclude that, despite their imperfections, estimated counts provided useful and credible information to the stakeholders involved.

1 Introduction

During the fall of 2004, the county of Los Angeles began a project to count the homeless. The effort was directed by the Los Angeles Homeless Services Authority (LAHSA), a City-County Joint Powers Authority independent of local government with a mandate to “address the problems of the homelessness on a regional basis” (www.lahsa.org/). A non-profit firm, Applied Survey Research (www.appliedsurveyresearch.org/), was awarded a contract to conduct the count, and brought us in as statistical consultants.

In this paper, we address the following questions. First, what was the political context in which the study was undertaken? Homeless counts are always controversial because of the range of stakeholders involved and the

*Thanks go to Peter Connery, Robin Conerly, Hope Malcom Maltz, and Peter Theodore of Applied Survey Research for doing all of the really hard work involved in this project.

fiscal resources connected to the estimates. Second, what design was used to undertake the count? A number of approaches had been used in the past with widely varying success. Third, how well was the design implemented? Earlier studies made clear that there might be a lot of slippage between the research design and its implementation. Fourth, what were the findings? It was anticipated going in that Los Angeles County might well have the largest number of homeless individuals of any metropolitan area in the country. So, the stakes were high. And fifth, what more general lessons might be learned from the Los Angeles experience? There were numerous decision points during a lengthy study period that might now be reviewed with profit. We chose as well to take advantage of the uncommon opportunity to inspect the improvement in model prediction of homeless totals as the aggregation level increases, truth being known. Here, the size of the observed data set allows us to do this in a fairly substantial setting.

2 Past Attempts to Count the Homeless

Estimates of the number of homeless depend on a clear definition of exactly what homelessness entails (Cordray and Pion, 1991) and on effective methods to apply the definition. Both steps have been controversial (Burt, 1989; Chemlinskym 1991; Kondratas, 1991; Rossi, 1989). Neither is near resolution to our knowledge. Here we focus on methods to estimate the number of homeless because definitions of homelessness ultimately depend on what kinds of homelessness the political process is prepared to accept and pay for.

It is probably fair to say that efforts to apply rigorous methods for counting the homeless began with the study conducted by Peter Rossi and his colleagues for the city of Chicago in 1985 (Rossi et al., 1987). Prior to that time, homeless estimates were based on the judgement of homeless advocates, or on interviews with local experts. In the Rossi study, Chicago was partitioned into city blocks and other geographical areas, together with sites like bus stations and airports, and characterized by the likelihood of finding homeless people in them. These areas were then sampled proportional to that probability. On a single night, enumerators were sent into the sampled areas to record all of the homeless individuals found. The street counts were supplemented with counts from homeless shelters.

A key concern with the Rossi approach is the ability of enumerators to obtain accurate information. Counting some people more than once is a

problem in principle, but the dominant worry is failing to find homeless individuals who should be counted. In a series of studies conducted in collaboration with the Bureau of the Census during the night of March 20-21, 1990, a clever effort was undertaken to document the undercount (Taeuber and Siegel, 1991; Martin et al., 1997). In each of five major cities, including Los Angeles,¹ approximately 60 “decoy” homeless individuals were put on the street in areas where other homeless were likely to be found. These confederates were instructed to report whether or not they had been counted by an enumerator. Between 22% and 67% of the decoys were counted, depending on the city (Wright and Devine, 1992), and the likelihood of a significant undercount was confirmed.

Articles on the experience in each city are contained in a special issue of the *Evaluation Review*, edited by James D. Wright (1992). Each article addressed the reasons why enumerators might miss a large number of homeless individuals and what might be done to improve matters. A common thread through all of the articles was the need to better train and supervise enumerators. With better training and supervision, there was considerable optimism that the undercount could be substantially reduced.

The use of decoys also raised the possibility of adjusting the homeless counts. The proportion of decoys counted could be used in a capture-recapture sampling framework to “correct” a count (Laska and Meisner, 1993; Martin et al., 1997; Schindler et al., 2001). Thus one would infer that when the proportions found were smaller, the undercount was larger. However, some strong and untestable assumptions are required to produce reliable estimates.

There are alternatives to the Rossi strategy, but these too have limitations. One popular approach is to build on the capture-recapture idea. For example, Cowan (1991) exploits information in the records of service-providers. If a person recorded to use a service on a given day is recorded to return the next day, that person has been “recaptured.” So, the number “captured” on day 1, can be up-weighted by the inverse of the proportion “recaptured” on day 2 to arrive at an overall estimate. However, this would only apply to services that “empty out” between day 1 and day 2 (e.g., a soup kitchen), and assumes a) that homeless individuals do not enter or leave the catchment area between the two days and b) that all homeless individuals in the catchment area have the same non-zero probability of using the service

¹The others were New York, Chicago, Phoenix, and New Orleans.

in question.

Another popular strategy is to use statistical methods to impute what a set of homeless counts should be (Hudson, 1997). For example, homeless counts at the census tract level can be regressed on other features of those census tracts. The resulting regression equation can be used to predict homeless counts when they are unavailable, as long as the same regressors can be utilized. Much the same approach can be used with data from service providers. The number of homeless individuals using the services during some specified periods and other variables (e.g., whether it is the holiday season) are taken as regressors when a model for the street counts is constructed. Future numbers of service users and other predictors can then be employed to impute street counts. Just as with capture-recapture methods, strong assumptions are required. For example, one must make the case that the regression model developed is in some sense “right.” One must also assume that the initial regression results can be generalized to the census tracts or catchment areas for which counts need to be imputed.

In summary, all of the efforts to count the homeless are flawed in various ways. The methods apparently preferred by the Census Bureau depend on getting enumerators to do their job well (Martin et al., 1997). The alternatives depend on models of one form or another, and are only as good as the models themselves. The ways to make the census results more credible seem clear and, if resources are available, viable in practice. The modeling approaches are more difficult to improve because so many of the key assumptions are effectively untestable and/or lead to no practical remedies.

3 Setting the Stage

Los Angeles County was established in 1850 as a relatively small agricultural and trading region along California’s coast between Santa Barbara and San Diego. Over the next four decades, its boundaries changed several times until, in 1889, it arrived at its current dimensions. As of January, 2005, over 10 million people were estimated to live within Los Angeles County’s 4,084 square miles, making the County the most populous in the nation with more residents than all but eight states.

Because of sheer size alone, the prospect of trying to count the number of homeless individuals was daunting. The usual sorts of political pressures complicated matters further. To begin with, the RFP from the county had

been circulated in the early Spring of 2004, and several worthy proposals were submitted shortly thereafter. There was apparently some reluctance within LAHSA to proceed; the study would be costly for a county that was struggling to provide the range of public services that it had in the past. It was only after HUD insisted that future funding would depend on credible measure of need that the contract was awarded. By then, several months were lost because the original HUD deadline remained.

Once an award was made, LAHSA mobilized quickly. Still, information necessary for an effective research design took several months to acquire. For example, a critical consideration in any counting effort would be information on where large numbers of homeless individuals were likely to be found, but moving beyond fragmented and anecdotal accounts took many weeks. Moreover, there were a number of important political considerations that were difficult to resolve. In particular, it was apparent from the start that it would be impossible with the resources available to send enumerators to all areas in the county. A sampling strategy was needed that was scientifically credible and yet managed to take the most important concerns of stakeholders into account. With a large and heterogeneous set of stakeholders, arriving at a consensus set of concerns was never fully achieved.

These and other considerations suggested a design strategy that would balance several competing risks. For example, one could not take as the whole truth a priori accounts of where large numbers of homeless individuals would be found. At the same time, it would be foolhardy to ignore such information. In addition, prudence dictated building into the design as much flexibility as possible. Beyond a homeless estimate for the county as a whole, there was a stated need for counts from certain areas within the county. Yet little guidance was provided about what those areas were likely to be. One implication was that a way had to be found to provide counts in geographical areas to which enumerators were not sent.

4 Sampling Design

The geographical area to be studied was determined by the boundaries of Los Angeles County, and the observational units were to be census tracts. Census tracts were selected because they are well defined and because they are a spatial unit for which considerable information was already available.

Los Angeles County is composed of 2054 census tracts. Efficient use of

resources dictated selecting a sample of tracts to produce an overall county estimate of the number of homeless. Homeless services for the County are in part organized into eight Service Provision Areas (SPAs). These contiguous areas, sporting such names as San Fernando Valley, Metro Los Angeles, and South Bay, suggest some of the unique features of the county, and ensure that each has strong advocates. Thus, an early attempt to sample in proportion to the size of homeless populations, as they were understood from earlier studies, foundered on politics, and the sample was stratified by SPA. Power calculations were undertaken to determine the approximate overall sample size required for desired precision.

There were discussions early in the process about providing estimated homeless counts for a variety of geographical areas within the county. The SPAs are an example, the incorporated cities in Los Angeles County are another, the five supervisorial districts yet another.² No decisions were made before the sample had to be drawn about specific areas within the county where estimated homeless counts would be needed.

Information provided by LAHSA indicated that 211 tracts could be anticipated to have a large number of homeless individuals and should be included with certainty among the tracts to be studied. From the remaining tracts, a stratified random sample of 299 was drawn. In this, each SPA formed a stratum with the number of tracts sampled in it proportional to the number of tracts it contained — SPAs with more tracts were more heavily sampled.

The design had positive features. First, by selecting with certainty the tracts anticipated to have large numbers of homeless individuals, one was assured of obtaining counts where they were most needed. Second, these tracts were then, in effect, removed from the population from which a random sample was to be drawn. This dramatically reduced the variability in counts in the population to be sampled, and led to a substantial gain in the precision of the results later obtained. Third, stratifying by SPA and sampling the number of tracts within each proportional to the actual number tracts in it did not differ markedly from a design which could be thought of as having been chosen proportional to size — politics did not interfere greatly with efficiency of estimation.

²In addition to Los Angeles city proper, Los Angeles County includes 87 incorporated cities such as Santa Monica, West Hollywood, Beverly Hills, Huntington Beach, Lynwood, Van Nuys, Burbank, and Inglewood. Los Angeles County is governed by a group of elected County Supervisors. Each represents an area of the County.

5 Implementation of the Study

Field work was undertaken from midnight to the early morning hours on January 25, 26 and 27, 2005. Enumerators were paired with homeless individuals who were hired at \$10 an hour to serve as “guides” in each census tract. Census tracts are large spatial areas, and it can be difficult to know where homeless individuals are likely to be found. The guides were to help in the search.

Shortly before the sample was drawn, we were told that the cities of Long Beach, Pasadena and Glendale would not cooperate in the study even though they were part of Los Angeles County and even though their numbers were to be included in the overall county estimate. Their counts would have to be imputed. This raised serious concerns because all three were alleged to contain a substantial number of homeless, and we had no way of knowing how well counts obtained for the rest of the county could be extended into these areas.

After the enumeration was completed, we learned that 33 tracts from SPA 4 had been selected for enumeration even though they were not a part of the random sample proposed. Tracts actually sampled contained long stretches of consecutively numbered tracts, and the decision was made to treat the unplanned sample tracts in the same manner as the ones that had earlier been chosen with certainty. Setting aside these tracts reduced the size of the estimation problem, but also ensured that one had an inadequate sample size for dealing with it. SPA 4, Metro Los Angeles, is thought of as a high homeless area, and the segment of it attended to through sampling was not well served by the process.

We were not made aware of other major problems with implementation. Indeed, the procedure for obtaining sample counts was said to have proceeded smoothly.

6 Sampling Estimates at the SPA Level

Street counts were provided for each tract sampled or selected with certainty. Shelter counts then contributed to a total count. In this total, we considered that uncertainty attached only to the sampling process, and assigned no uncertainty to the sampled tract, selected tract and shelter counts obtained. The estimated number of homeless in each SPA attributable to its sampled

tracts was computed as

$$\hat{\tau} = \frac{N}{n} \times \sum_{i=1}^n c_i, \quad (1)$$

where N is the number of sample-eligible tracts in a SPA, n is the number of sampled tracts in which counts were obtained, and c_i is the the number of homeless counted in sampled tract i . For example, if N were 100, n were 25, and the count from 25 sampled tracts was 200, the estimate for that SPA would be 800. Adding in the tract counts from the relevant excluded tracts and shelters produced an estimate for each SPA. Summing over SPAs produced an estimate for the county as a whole.

To obtain an estimate of the variance of the estimated count for a given SPA, first consider the (unknown) variance of a randomly selected tract count in the SPA using

$$\text{var}(c) = \frac{1}{N} \sum_{i=1}^N (c_i - \bar{c})^2, \quad (2)$$

where \bar{c} is the mean of the counts for the SPA tracts and N is the total number of tracts. Then the (unknown) variance of the estimated count is

$$\text{var}(\hat{\tau}) = (N/n)^2 \times n \times \text{var}(c) \times \frac{(N - n)}{(N - 1)}, \quad (3)$$

where the far right term is the finite population correction. An estimate of this variance was obtained by replacing $\text{var}(c)$ in (2) with its sample analogue. The square root of the estimate of $\text{var}(\hat{\tau})$ was then the estimated standard error (SE) for that SPA's estimated count due to the sampling of tracts. The margin of error for each SPA was taken to be twice the estimated standard error for that SPA. For example, using the same numbers as above and supposing the sample estimate of $\text{var}(c)$ to be 100, $\text{SE}(\hat{\tau}) = \sqrt{16 \times 25 \times 100 \times .76} = 174$. If one further supposed the excluded tracts and shelters in the SPA contributed a count of 400 to go with the observed 800, the margin of error for that SPA would be expressed as 1200 ± 348 .

The standard error for the county as a whole, is the square root of the sum of the estimates of $\text{var}(\hat{\tau})$ for each SPA. The margin of error for the county as a whole was taken to be twice the standard error for the county as a whole. This is usefully viewed as a percentage error compared with the estimated total.

7 Imputing Homeless Counts at the Tract Level

There was an interest in providing estimates of the number of homeless for all individual tracts in which no counts were available. Three strategies were considered for imputing the count in a given tract.

1. Using the SPA average — For each non-sampled tract, we use the mean count of the sampled tracts in the SPA in which the tract is found as the imputed value.
2. Using matched tracts — For each tract, we match it to tracts with similar values for correlates of homelessness and use the average of the counts in the matched tracts as the imputed value. Use of the SPA average as above is just a special case of this; one views tracts as matched if they are in the same SPA.
3. Using correlates of homelessness — We construct a statistical model of how various features of tracts are related to homelessness and use the model’s predictions as the imputed values. The previous strategy is a special case - categories of values there are replaced by observed values here.

When tract-level values are aggregated to large spatial units, all three methods, sensibly applied, will give similar results. The census tracts for which imputed values are needed are those that were not designated *a priori* as tracts with large number of homeless people. Insofar as this is correct, there is likely to be less variability in the counts in these tracts. After suitable aggregation, what variability there is tends to cancel out.

Such averaging out is not guaranteed to occur, however. It depends, in particular, on how many tracts are aggregated and which tracts they are. Should interest reside in a small number of tracts, the three methods can give somewhat different results (see Section 10 for some evidence on this point).

We rather reluctantly left the first strategy behind. The reasons for doing so included the following: (i) little data survived in SPA 4, a particularly unfortunate choice due to its inner city location; (ii) it was necessary to extrapolate to tracts in Long Beach, Pasadena and Glendale that were not in the sampling frame; and (iii) it was difficult to ignore possible homeless indicators that stakeholders “know” are of value in understanding homeless

counts. Accordingly, the second and third strategies seemed likely to perform better than the first since they used more (presumably useful information) than the SPA membership noted in the first method. Between the matching and the modeling approaches, modeling seemed to perform a bit better in the manner in which available tract information was put to use.

Straightforward sampling estimates of homeless totals are given by SPA below for subsequent comparison with modeled totals. Going forward, we continue to make use of ordinary sampling estimates as standards by which model estimates might be judged, as applicable.

8 Results

8.1 Estimates for the Service Provision Areas and the County as a Whole

Table 1 shows the results obtained from sampling theory, discussed in Section 4, as aggregated to the SPA and county levels. The figure of nearly 65,000 homeless individuals in the County was within the range that most stakeholders found credible. For the county as a whole, the standard errors implied a level of precision that was acceptable. The SPA standard errors varied substantially as a percentage of the estimated homeless count. Some margins of error were as large as plus or minus 30%. No one was happy with the larger figures, but they were a result of resource constraints and earlier design tradeoffs that on balance still seem appropriate.

8.2 Estimates at the Tract Level

We spent considerable time examining variables, commonly associated with homelessness, for the purpose of implementing the second and third imputation methods. These included such US census data variables as median income and the percentage of housing units vacant, and such land use variables as the percentage of land devoted to commercial activities, industry, or residences.

By and large, homelessness correlates performed as expected: homeless counts were higher, for example, in tracts with lower median income, in tracts with a higher percentage of vacant dwellings, and in tracts where a higher percentages of land was used for commercial or industrial purposes.

SPA	Count-Selected	Count-Sampled	Count-Shelter	Total	SE
1	275 (8 tracts)	370 (11 tracts)	419	2813	555
2	347 (20 tracts)	1034 (61 tracts)	1570	8816	1025
3	802 (23 tracts)	995 (51 tracts)	870	8091	1510
4	5931 (86 tracts)	99 (6 tracts)	4544	13429	935
5	1496 (21 tracts)	479 (21 tracts)	1613	5937	395
6	3635 (40 tracts)	1511 (29 tracts)	2017	13936	1420
7	419 (24 tracts)	438 (39 tracts)	912	4049	290
8	506 (22 tracts)	801 (48 tracts)	1686	7315	460
Total	13411 (244 tracts)	5727 (266 tracts)	13631	64386	4875

Table 1: Estimated Tract Totals and Standard Errors Using SPA Averages

These relationships were sometimes highly non-linear, suggesting a possible tipping effect. Thus there is a suggestion of thresholds that, if passed, tip a census tract dramatically toward having a large number of homeless people. The previous studies we reviewed assumed linear or quadratic regression relationships (e.g., Hudson, 1997) and might be usefully revisited.

The modeling technique of choice, success-driven, was that of random forests (Breiman, 2001). First of all, the relationships between some of the predictors and the counts were non-linear in ways that we could not anticipate and then easily capture in parametric form. Additive smoothers constructed through generalized additive models looked more promising, but abrupt changes in the counts over some small regions of the predictors led to over-smoothing in some places and under-smoothing in others. The flexibility of random forests seemed best suited to these data, and we adopted this methodology.

Several random forest models performed similarly and, for many purposes, it would not matter a lot which one was used for imputation. They differed a bit in the homelessness correlates used, but not in any surprising manner. For example, median income tends to be lower when the percentage of renters is higher; there is some redundancy between the two, and it does not matter greatly which correlate is used. The differences seen in various (sensible) models looked at are not likely to matter much for most policy purposes.

It should be emphasized that model search was never intended to be exhaustive. We relied more on predictors thought of as familiar in this context than on the elicitation of predictors by data-driven selectors. In all, perhaps

SPA	Count-Selected	Count-Sampled	Count-Shelter	Total
1	275 (8 tracts)	370 (11 tracts)	419	2707
2	347 (20 tracts)	1034 (61 tracts)	1570	9220
3	802 (23 tracts)	995 (51 tracts)	870	7505
4	5931 (86 tracts)	99 (6 tracts)	4544	16042
5	1496 (21 tracts)	479 (21 tracts)	1613	5634
6	3635 (40 tracts)	1511 (29 tracts)	2017	12453
7	419 (24 tracts)	438 (39 tracts)	912	5673
8	506 (22 tracts)	801 (48 tracts)	1686	9175
Total	13411 (244 tracts)	5727 (266 tracts)	13631	68409

Table 2: Estimated Tract Totals and Standard Errors Using Tract Correlates

a dozen models were considered with care. To then single out a model to provide the homeless count by tract, we invoked several criteria - among them were the faithfulness of predictors when compared to observed counts at the tract level, their face validity in the homeless context in Los Angeles, the harmony of predicted totals with outside estimates of homeless numbers in Long Beach, Pasadena and Glendale, along with the parsimonious nature of the model.

The model we ultimately put forward used just three predictors for homeless count by tract: median household income, percent vacancy and percent residential use. When these model tract predictions are aggregated to the SPA and county levels, Table 2 is the result.

When compared with the sampling totals given in Table 1, predicted totals are found to be well within sampling error in SPAs 1, 2, 3, 5 and 6. The difference seen in SPA 4 might be traced to the heavy reliance placed on a small sample for predicting street count in unsampled tracts; there is some evidence to support the (higher) model answer given here.

In SPAs 7 and 8, it may be that one sees the effect of non-participation by Long Beach in the identification of selected tracts, and in the selection of sample tracts; the sample extrapolation thinks of Long Beach as a participant, while the model picks up on characteristics of Long Beach tracts and proposes a higher count accordingly. In particular, the increase of some 4000 homeless, in going from sample-based estimation to model prediction, is largely accounted for by increases in SPAs 7 and 8.

In a later section on aggregation, we return to the differences seen between

observed counts and those produced by the above model, and we contrast the counts produced by these strategies with those predicted by two larger models.

8.2.1 Standard Errors for the Tract-Level Imputed Values

There is precious little written on statistical inference for data mining procedures in general and random forests in particular. Yet, standard errors for the imputed counts at the tract level could be helpful to stakeholders. As a useful approximation, we wrapped a bootstrap procedure around random forests and produced a large number of imputed counts for each tract. The bootstrap sampling was done in a fashion that mirrored the stratified sampling employed when the tracts were originally selected.

The set of imputed counts for each tract can be exploited in at least three ways. First, and most conservatively, the standard deviation associated with the realized imputed counts for each tract conveys approximately how large a typical disparity would be between the mean imputed value and an imputed value from a random sample of the data. A smaller standard deviation implies a more stable estimate over samples. Ideally, the standard deviation is small compared to the size of the imputed count.

Second, one can treat the set of bootstrapped imputed counts for each tract as an estimate of a sampling distribution. Then, the standard deviation becomes a standard error. If one assumes that the sampling distribution is normally distributed, conventional confidence intervals can follow. However, a skewed sampling distribution is likely for imputed counts near zero.

Third, if one is prepared to treat the bootstrapped imputed counts as an estimate of a sampling distribution, one can use the empirical sampling distribution to construct confidence intervals directly. We employed the percentile method, knowing of no formal results to justify getting any fancier. And indeed, we know of no formal results in this context to justify any method for obtaining confidence intervals.

Figure 1 is a histogram of the standard deviations for the bootstrap imputed census tract counts. The vast majority of the standard deviations are less than 10 with a plurality of less than 5. It is likely that standard deviations of 5 or less would imply sufficient stability to users of tract level

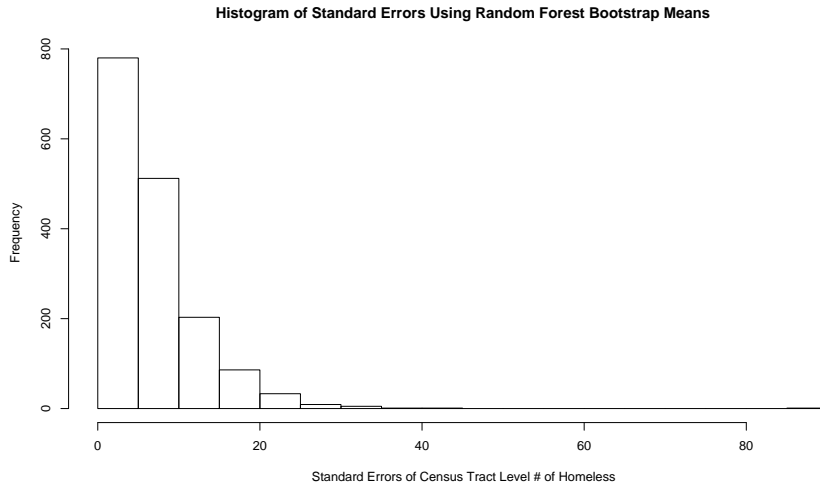


Figure 1: Histogram for Bootstrap Standard Deviations for the Imputed Counts

estimates. Standard deviations of 10 or less might also be readily acceptable. However, there are a few very large standard standard deviations that would likely be seen as problematic. A lot would dependent on the counts with which they were associated. Generally, larger standard deviations are associated with larger counts, as one would expect, but how large a standard deviation is too large would necessarily depend on what stakeholders planned to do with the information.

Figure 2 is a histogram of the percentile 95% confidence intervals associated at the census tract level with each imputed count. Given the long tailed distribution of the standard deviations, a long tailed distribution for the size of the confidence intervals is no surprise. The majority of confidence are probably acceptable, but some are probably not.

One can better sense of the potential problems from Figure 3 in which a random sample of tract-level imputed counts and percentile confidence intervals are shown. It is again apparent that the size of the confidence

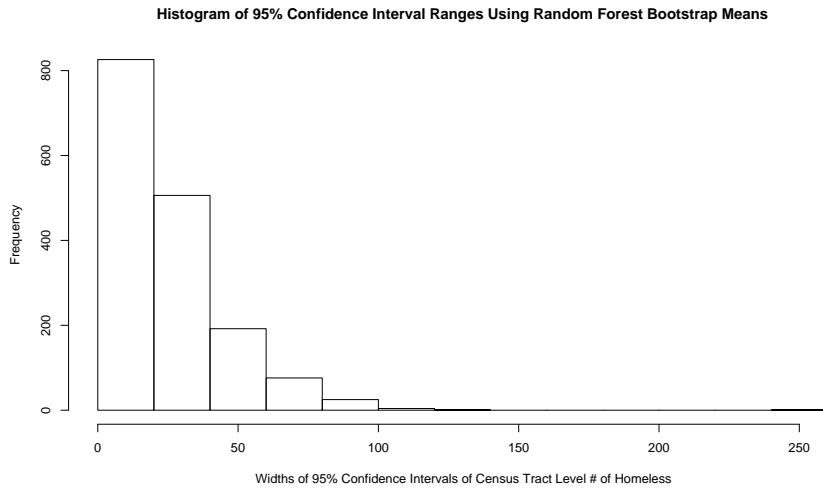


Figure 2: Histogram of Confidence Interval Size for the Imputed Counts

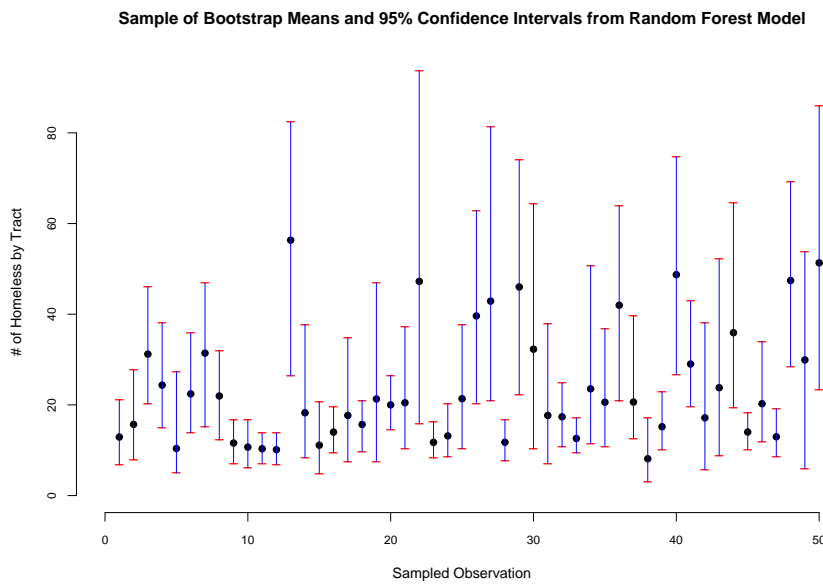


Figure 3: A Random Sample of Confidence Intervals for the Imputed Counts

SPA	Model 0	Model 1	Model 6	Sample	Standard Error
1	2557	2361	2707	2813	555
2	9400	9768	9220	8816	1025
3	7538	7603	7505	8091	1510
4	16378	16249	16042	13429	935
5	6051	5720	5634	5937	395
6	12377	11898	12453	13936	1420
7	5003	5561	5673	4049	290
8	9268	8586	9175	7315	460
Total	68568	67746	68409	64386	4875

Table 3: Model Comparisons Using Estimated SPA Totals

intervals vary substantially. The majority are small relative to their imputed counts. But some are not, and a few as so large as to perhaps make the imputed count useless. A lot would be depend on what might be done with such counts.

9 The Aggregation of Modeled Tract Counts

Tract counts gain interest as they are aggregated to recognizable levels. Here we review SPA totals, bring in predictions of homeless populations in some incorporated cities in the county, and contrast sets of figures produced by four methods — by sampling as in Table 1, by the model used for Table 2, and by invoking two alternative (larger) models.

We refer to the proposed model as Model 6. It uses the tract variables median household income, percent residential, and percent vacancy. We have considered the introduction of geographic coordinates, and Model 1 uses median household income, percent commercial, percent industrial, together with latitude and longitude. Finally, as a gauge, we let Model 0 incorporate the predictors of Model 1 together with percent owner occupied, percent vacant and percent minority. The models thus use eight, five and three predictors.

Consider Service Provider Areas again in Table 3, in the form of sample estimates and their standard errors set alongside model predictions.

Models are smoothers so it is not surprising that the sampling estimate in Table 3 is, with one exception, the largest or the smallest total in each row.

Cities	Model 0	Model 1	Model 6	Sample	Standard Error
Palmdale	558	434	517	662	134
Burbank	504	453	431	401	99
Alhambra	411	408	474	502	145
West Hollywood	166	156	147	128	26
Santa Monica	1535	1501	1497	1496	48
Lynwood	658	656	701	762	164
Whittier	428	412	415	403	29
Inglewood	1154	1160	1095	971	67

Table 4: Model Comparisons Using Some Estimated City Totals

The discrepancies between sample numbers and model predictions recall the discussion below Table 2 as it compares Model 6 and sample totals. Indeed, there are no incisive contrasts between models in Table 3 — perhaps that is the message at this level of aggregation.

Still Service Provider Areas are hardly the stuff of legend, so we give predictions for the size of homeless populations in some incorporated cities that are more recognizable entities. They are first chosen from SPAs 1 through 8, and presented in that order in Table 4. (To produce these numbers, we have shared down sample and predicted figures according to percentage of tract area in a city — it happens occasionally that different cities share a given tract in significant proportions.)

Table 5 is a continuation of Table 4, but it features the non-participating cities where predictions were sought despite the absence of the relevant tracts from the sampling universe.

Arguably, Model 6 does the best job of the three models in tracking the sample numbers based on SPA membership. Still, we do not feel the need to make a strong case for this on the grounds that the model is indeed competitive with the others, and simpler.

There is the suggestion that Models 0 and 1 are overly sensitive to location. Inglewood, for example, abuts the South Central Los Angeles area and has close to the highest minority population of any city in the country - 19% white in the 2000 census. At the same time, it is a city with a sizable percentage of homeowners (there were some 11,000 single family owner-occupied homes in a population of roughly ten times that many people in 2000), and it holds to a fair amount of stability. Here Models 0 and 1 overshoot the

Cities	Model 0	Model 1	Model 6	Sample	Standard Error
Glendale	800	1000	741	627	196
Long Beach	2902	2381	3143	2041	184
Pasadena	563	536	624	606	220

Table 5: Model Comparisons Using Totals for Non-Participating Cities

sampling estimate by a considerable margin.

An overarching comment is that there can be little doubt that local homeless policy plays a large role in the ability to effect their census. Thus, Santa Monica attracts homeless residents by paying close attention to them, by taking a rather liberal stance in their behalf, and by providing a relatively safe and clean environment.

For consistency with Table 4, we have used the same tract assignment file to produce Table 5 for the three non-participating cities. A second file detailing such assignments was the one used for "official" counts, see Point ?? in Section ?? for more details about this.

Here, Model 6 looks most like the sample estimates built on SPA averages in Glendale and Pasadena. In the case of Long Beach, use of the SPA-wide average provides a decidedly low sample estimate in that previous study of Long Beach homeless had come up with a population size of close to 6000. A shelter count of 296 was reported for Long Beach, but this figures to be considerably short of a full census. Even more to the point, the lack of pre-identified, populous tracts in Long Beach ensured a heavy undercount. Then Model 6 does the best job of recognizing, while not correcting, the situation brought on by non-participation.

10 A Model Validation Study

Setting tract-level predictions alongside the counts that were obtained on the sampled tracts, we can carry out a study of the improved accuracy of modeled totals as aggregation sample size increases within this pool of tracts. The results can be taken as suggestive of what various consumers will encounter when they rely on modeled totals across unsampled tracts, of interest for political and other reasons. Moreover, it is instructive to carry out a validation exercise not commonly available in practical settings — one in which truth is known.

Going further, we can and do compare the performances of different models according to the same aggregation criteria. In particular, Models 0, 1 and 6 mentioned above are entered as competitors, and we bring in Model S: assign the SPA-wide average to the appropriate constituents in any given collection of tracts, then sums.

While tract sample numbers have gone into the building of the models, we do not see that this taints the present study. We looked at perhaps a dozen models, as was discussed in Section 7, and the choice among them was built on criteria that had no overt connection with model totals over collections of sampled tracts.

Since we combine only sampled tracts, our constructs are, of necessity, artificial. Two methods of tract aggregation are considered for the purposes of this exercise: random selection, and selection with the aid of a proximity device that is described below. Other choices are possible of course, but we have not seen much in the way of qualitative difference over alternative sampling devices.

To be specific, we start with a file that is 259 x 5 - tracts by models plus truth (six tracts from SPA 4 have been set aside). We go through the aggregation sample sizes 4, 8, 16, 32 and 64. For each size, we first draw samples at random and without replacement from the 259 tracts, a total of 500 of them. Total homeless in the tracts drawn is regarded as small if below the median seen in the 500 draws, as large if above the median. Then, for each of the four models, we report the median absolute percent error of prediction, MAPE, as well as the percentage of cases in which predictions undershoot small totals, US, and the percentage of cases in which predictions overshoot large totals, OL. The first statistic allows comparisons to be made across models and through sample sizes, the other two quantify to some extent the imbalance inherent in the smoothing of the data through modeling.

Subsequently, for each sample size n , a tract number is selected at random, and a selection is made of n of the $n + 8$ tracts closest in distance to it. The manner of selection here is termed geographic, and performance statistics are kept as before. The results may be seen in Table 5, all entries being percentages.

For any fixed model, Table 5 shows a rough 30% decrease in median absolute percent error on a doubling of sample size. On this scale, the model ordering is the one expected based on the size of the model, Model S having notably poor performance in this regard. There is a decided difference between the ability to straddle large and small targets that emphasizes smooth-

Size	Random Selection				Geographic Selection			
	Model	MAPE	US	OL	Model	MAPE	US	OL
4	0	23	5	36	0	19	15	44
	1	24	2	34	1	19	11	40
	6	27	4	32	6	24	14	38
	S	46	4	32	S	39	10	29
8	0	17	8	18	0	14	16	43
	1	18	6	18	1	16	10	33
	6	20	7	23	6	18	11	33
	S	32	6	14	S	28	15	27
16	0	13	8	28	0	10	18	37
	1	13	7	22	1	10	12	30
	6	15	8	26	6	11	12	29
	S	23	8	21	S	20	20	27
32	0	11	8	22	0	7	18	37
	1	11	10	16	1	8	14	21
	6	12	14	17	6	9	13	26
	S	16	6	12	S	17	31	21
64	0	7	8	25	0	4	17	56
	1	7	12	16	1	4	18	17
	6	8	10	16	6	4	32	37
	S	11	10	10	S	8	41	23

Table 6: Results with Aggregated Census Tracts

ness more on the low than on the high end, and some differences between models on this score. When it comes to the presumably more realistic geographic constructs, models can perform quite differently at the higher levels of aggregation, though this can be an artifact of a selection method that is able to associate quite disparate tracts when distance is even moderately large. In any event, one sees in these models a tendency to overestimate the homeless when they are relatively few, to underestimate the size of the problem when there are relatively many. Referring to Table 4, an inference that the homeless problem in Whittier (15 unsampled tracts at about 10 homeless per tract) is overstated, while in Burbank (17 unsampled tracts at about 25 per tract) it is understated, is hugely tenuous, but the general point is perhaps of interest to the policy maker.

Comparing the left and right hand sides of Table 5, one sees a rather substantial increase in accuracy that attaches to the structured aggregates. Moreover, there is an increased ability of the models to pick up on high and low totals. This is comforting as one imagines that geography is not at great remove in most uses of the tract predictions.

Again in this setting, as was the case in the previous section, Model 6 holds up well against its larger competitors, and it dominates the simple Model S that it replaced.

11 Some Details and Caveats

1. Among the preselected tracts, one finds counts in the 600's. Thus, 1/1000 of the tracts in the county brought in 2% of the estimated total of homeless individuals. Even among sampled tracts, where counts are expected to be more homogeneous, one finds numerous counts in excess of 100 even though preselected tracts, which are meant to include heavy concentrations, average less than 60.
2. In view of the prominent use of the counts found in six sampled tracts from SPA 4, it is worth remarking on what was found in them. These six were quite homogeneous in homeless count, averaging less than 20, with a largest value of 40. Put these figures alongside the counts found in the 33 tracts that were mistakenly treated as sampled tracts: counts averaged more than 30, with a largest value of 265. There is at least the suspicion that the six sample values were from the low end

of the distribution of homeless counts over the tracts that were not preselected.

3. One enumerated tract had a missing street count, so the initial 299 sample tracts were ultimately reduced to 265 after the SPA 4 difficulty was discovered.
4. As background information, we were given homeless totals for Glendale (472), Long Beach (5845) and Pasadena (879) that had been found at various times in other studies. We took the figures provided as rough guidelines for the assessment of our projections into the three areas from which tracts were not available for sampling. In Table 4, for the sake of consistency, we used a file that assigned proportions of cities to tracts according to acreage. Earlier, based on a second file that assigned tracts to these three communities, the selection of Model 6 over Model 1 was based in part on estimates that more closely resembled the numbers at hand: in the same order, 987, 813 and 3004 under Model 1, 748, 895 and 3777 under Model 6.
5. Predictions made here at the tract level take no special account of the non-homogeneous nature of the county. One approach to this might consider breaking the county into zones somewhere between the SPA and county levels, then fitting different models to such zones. From observed counts only, one tentative split incorporates a “city” zone consisting of SPAs 4 and 6, a “suburban” zone consisting of SPAs 5, 7 and 8, and a “valley” zone consisting of SPAs 1, 2 and 3. The statistical evidence for such a split comes from the size of concentrations of the homeless: the first zone has the largest by far, the second zone has significant, but moderate by comparison, concentrations; the third zone consists of fairly homogeneous tracts. It could be, though, that access to Long Beach data would suggest aligning SPA 8 with SPAs 4 and 6 rather than with SPAs 5 and 7.
6. Without input to the undercount issue, we took the position that predictions applied only to the “countable” population, and we harbor no illusions that all eligible persons were indeed found. The homeless population of Los Angeles is often reported to be in the neighborhood of 85,000 to 90,000. The increase in estimated homeless is attributable to a separate telephone survey that was conducted, and to estimates of

the homeless in jails, hospitals and the like. We were not involved in these aspects of the study.

12 Conclusions

A variety of methods have been employed to estimate the number of homeless individuals in particular metropolitan areas. Experience gained in this project reinforces our belief that street counts undertaken in sampled geographical areas can be used to construct useful and cost-effective estimates. Effectiveness is enhanced by sampling strategies that take serious account of available information on locations where the homeless tend to congregate. Some efficiency of estimation was lost to political pressure in the Los Angeles project, but not to any substantial degree.

A major consequence of sampling can be the need to impute counts for areas not selected. For this study, estimated counts were required for each census tract because stakeholders were planning to combine census tracts in ways that could not be anticipated. We suspect that stakeholders in other jurisdictions will often want similar flexibility.

The imputation of counts to unsampled tracts can benefit from modeling in this context. We do not, however, believe that any specifics of the present choices carry over to other times or places. One can easily find, as we have in a related project, that what is “known” in advance about the location of homeless individuals and the correlates of homelessness is inaccurate. That, in turn, will affect the details of how the imputed counts are constructed. New projects are likely to profit most from up-to-date local information and a reasonably clean slate.

Aggregation of model predictions over sampled and unsampled units can improve accuracy that in this case is sufficient to help stakeholders. But a lot depends on which areas are combined, not just the number of areas. A key factor is that the fitted values tend to overestimate low counts and underestimate high counts. Bias can be a problem, but perhaps it is a mild one when an aggregate covers a decent mix of what figure to be high and low count tracts

Finally, the level of professionalism exhibited in the field and during follow-up operations in the Los Angeles project was substantial and vital to the quality of the results. Given this, we can take the view that the estimations produced are of real value. Indeed, the counts were forwarded to

HUD shortly after the data analysis was completed.

References

- Anderton, D. (1991) "Using Local Longitudinal Records to Estimate Transient and Resident homeless Populations." *Housing Policy Debates* 2(3): 883-898.
- Breiman, L. (2001a) "Random Forests." *Machine Learning* 45: 5-32.
- Burt, M. (1989) *America's Homeless: Numbers, Characteristics, and Programs that Serve Them*. Washington, D.C.: Urban Institute.
- Cowen, D.D. (1991) "Estimating Census and Survey Undercounts Through Multiple Service Contacts." *Housing Policy Debates* 2(3): 869-882.
- Chelimsky, E. (1991) "Politics, Policy Making, Data, and the Homeless." *Housing Policy Debates* 2(3): 683-697.
- Cordray, D.S., and G. M. Pion (1991) "What's Behind the Numbers? Definitional Issues in Counting the Homeless." *Housing Policy Debates* 2(3): 587-616.
- Hudson, C.G., (1997) "Estimating Homeless Populations through Structural Equation Modeling." *Journal of Sociology and Social Welfare* 25(2): 136-154.
- James, F.J. (1991) "Counting Homeless Persons with Surveys of Users of Services for the Homeless." *Housing Policy Debates* 2(3): 733-753.
- Kondratas, A. (1991) "Estimate and Public Policy: The Politics of Numbers." *Housing Policy Debates* 2(3): 631-647.
- Laska, E.M., and M. Meisner (1993) "a Plant-Capture Method for Estimating the Size of a Population from a Single Sample." *Biometrics* 49: 209-220.
- Martin, E., Laska, E, Hopper, K., Meisner, M. and J. Wanderling. (1997) "Issues in the Use of a Plant-Capture Method for Estimating the Size of the Street Dwelling Population." *Journal of Official Statistics* 13(1): 59-73.
- Rossi, P.H. Wright, J.D., Fisher, G.A., and G. Willis (1987) "The Urban Homeless: Estimating Composition and Size." *Science* 235: 1336-1341.

- Rossi, P.H. "Strategies for Homeless Research in the 1990s." *Housing Policy Debates* 2(3): 1029-1055.
- Rossi, P.H. (1989) *Down and Out in America: the Origins of Homelessness*. Chicago: University of Chicago Press.
- Schindler, E., Griffin, R., and A. Navarro (2001) "Sampling and Estimation for the Homeless Population." Washington, D.C.: U.S. Bureau of the Census.
- Tauber, C.M., and P. Siegel (1991) "Counting the Nation's Homeless Population in the 1990 Census." In C.M. Tauber (ed.) *Enumerating Homeless Persons: Methods and Data Needs*. Washington, D.C.: U.S. Bureau of the Census.
- Wright, J.D. and J.A. Devine (1992) "Counting the Homeless: The Census Bureau's 'S-Night' in Five U.S. Cities." *Evaluation Review* 16(4): 355-364.