

Better Estimates from Binned Income Data: Interpolated CDFs and Mean-Matching

Paul T. von Hippel,^a David J. Hunter,^b McKalie Drown^b

a) University of Texas at Austin; b) Westmont College

Abstract: Researchers often estimate income statistics from summaries that report the number of incomes in bins such as \$0 to 10,000, \$10,001 to 20,000, ..., \$200,000+. Some analysts assign incomes to bin midpoints, but this treats income as discrete. Other analysts fit a continuous parametric distribution, but the distribution may not fit well. We fit nonparametric continuous distributions that reproduce the bin counts perfectly by interpolating the cumulative distribution function (CDF). We also show how both midpoints and interpolated CDFs can be constrained to reproduce the mean of income when it is known. We evaluate the methods in estimating the Gini coefficients of all 3,221 U.S. counties. Fitting parametric distributions is very slow. Fitting interpolated CDFs is much faster and slightly more accurate. Both interpolated CDFs and midpoints give dramatically better estimates if constrained to match a known mean. We have implemented interpolated CDFs in the "binsmooth" package for R. We have implemented the midpoint method in the "rpme" command for Stata. Both implementations can be constrained to match a known mean.

Keywords: Gini; inequality; income brackets; grouped data

SURVEYS often ask respondents to report income in brackets, or *bins*, such as \$0 to 10,000, \$10,001 to 20,000, ..., \$200,000+. Even if respondents report exact incomes, incomes may be binned before publication, either to protect privacy or to summarize the income distribution compactly with the number of incomes in each bin. Table 1 gives a binned summary of household incomes in Nantucket, the richest county in the United States.

Binning presents challenges to investigators who want to estimate simple summary statistics such as the mean, median, or standard deviation or inequality statistics such as the Gini coefficient, the Theil index, the coefficient of variation, or the mean log deviation. Researchers have implemented several methods for calculating estimates from binned incomes.

The simplest and most popular approach is to assign each case to the midpoint of its bin by using a robust pseudo-midpoint for the top bin, whose upper bound is typically undefined (e.g., Table 1). The weakness of the midpoint approach is that it treats income as a discrete variable, but the method also has several strengths. The midpoint method is easy to implement and runs quickly. Midpoint estimates are also "bin-consistent" (von Hippel, Scarpino, and Holas 2015) in the sense that midpoint estimates will get arbitrarily close to their estimands if the bins are sufficiently numerous and narrow.

Another approach is to fit the bin counts to a continuous parametric distribution. Popular distributions include two-, three-, and four-parameter distributions from the generalized beta family, which includes the Pareto, log normal, Weibull, and

Citation: von Hippel, Paul T., David J. Hunter, and McKalie Drown. 2017. "Better Estimates from Binned Income Data: Interpolated CDFs and Mean-Matching." *Sociological Science* 4: 641-655.

Received: September 23, 2017

Accepted: October 8, 2017

Published: November 15, 2017

Editor(s): Jesper Sørensen, Stephen Morgan

DOI: 10.15195/v4.a26

Copyright: © 2017 The Author(s). This open-access article has been published under a Creative Commons Attribution License, which allows unrestricted use, distribution and reproduction, in any form, as long as the original author and source have been credited. 

Table 1: Household incomes in Nantucket.

Minimum	Maximum	Households	Cumulative distribution
\$0	\$10,000	165	5%
\$10,000	\$15,000	109	8%
\$15,000	\$20,000	67	9%
\$20,000	\$25,000	147	13%
\$25,000	\$30,000	114	17%
\$30,000	\$35,000	91	19%
\$35,000	\$40,000	148	23%
\$40,000	\$45,000	44	24%
\$45,000	\$50,000	121	28%
\$50,000	\$60,000	159	32%
\$60,000	\$75,000	358	42%
\$75,000	\$100,000	625	59%
\$100,000	\$125,000	338	69%
\$125,000	\$150,000	416	80%
\$150,000	\$200,000	200	86%
\$200,000		521	100%

Each bin's population is estimated from a one-in-eight sample of households who took the American Community Survey in 2006 through 2010. Incomes are in 2010 dollars.

Dagum distributions, among others (McDonald and Ransom 2008). One implementation fits up to 10 distributions and selects the one that fits best. An alternative is to calculate a weighted average of income statistics across several candidate distributions (von Hippel et al. 2015).

A strength of the parametric approach is that it treats income as continuous. A weakness is that even the best-fitting parametric distribution may not fit the bin counts particularly well. If the fit is poor, the parametric approach is not bin consistent; that is, even with an infinite number of bins, each infinitesimally narrow, a parametric distribution may produce poor estimates if it is not a good fit to the underlying distribution of income.

A practical weakness of the parametric approach is that it is typically implemented using iterative methods, which can be slow. The speed of a parametric fit may be acceptable if you fit a single distribution to a single binned data set, but runtimes of hours are possible if you fit several distributions to thousands of binned data sets—such as every county or school district in the United States. Other computational issues include nonconvergence and undefined estimates. These issues are rare but inevitable when you run thousands of binned data sets (von Hippel et al. 2015).

Neither approach—midpoint or parametric—is uniformly more accurate. With many bins, midpoint estimates are better because of bin consistency, whereas with fewer than eight bins, parametric estimates can be better because of their smoothness (von Hippel et al. 2015). Empirically, the parametric and midpoint approaches produce similarly accurate estimates from typical U.S. income data with

15 to 25 bins. Both methods typically estimate the Gini within a few percentage points of its true value. This is accurate enough for many purposes but can lead to errors when estimating small differences or changes such as the 5 percent increase in the Gini of U.S. family income that occurred between 1970 and 1980 (von Hippel et al. 2015).

A potential improvement is to fit binned incomes to a flexible nonparametric continuous density. Like parametric densities, a nonparametric density treats income as continuous. Like the midpoint method, a nonparametric density can be bin consistent and fit the bin counts as closely as we like.

Unfortunately, past nonparametric approaches have been disappointing. A nonparametric approach using kernel density estimation had substantial bias under some circumstances (Minoiu and Reddy 2012). A nonparametric approach using a spline to model the log of the density (Kooperberg and Stone 1992) had even greater bias (von Hippel, Holas, and Scarpino 2012), though it was not clear whether the bias came from the method or its software implementation.

In this article, we implement and test a nonparametric continuous method that simply connects points on the empirical cumulative distribution function (CDF). The method, which we call *CDF interpolation*, outperforms its predecessors in both speed and accuracy. The method can connect the points using line segments or cubic splines. When cubic splines are used, the method is similar to “histospline” or “histopolation” methods, which fit a spline to a histogram (Wahba 1976; Morandi and Costantini 1989). But histosplines are limited to histograms that have bins of equal width (Wang 2015). CDF interpolation is a more general approach that can handle income data for which the bins have unequal width and the top bin has no upper bound (e.g., Table 1).

We have implemented CDF interpolation in our R package *binsmooth* (Hunter and Drown 2016), which is available for download from the Comprehensive R Archive Network (CRAN).

Our results will show that statistics estimated with CDF interpolation are slightly more accurate than estimates obtained using midpoints or parametric distributions. In addition, CDF interpolation is much faster than parametric estimation, though not as fast as the midpoint method.

We also show that the differences between methods are dwarfed by the improvement we get if we constrain a method to match the grand mean of income, which the U.S. Census Bureau often reports alongside the bin counts. If we constrain either the interpolated CDF or the midpoint method to match a known mean, we get dramatically better estimates of the Gini. Our *binsmooth* package can constrain an interpolated CDF to match a known mean, and our new version (2.0) of the *rpme* command for Stata can constrain the midpoint method to match a known mean as well.

In the rest of this article, we define the midpoint, parametric, and interpolated CDF methods more precisely, then compare the accuracy of estimates in binned data summarizing household incomes within U.S. counties. We also show how much estimates improve if we have the mean as well as the bin counts.

Methods

Binned Data

A binned data set, such as Table 1, consists of counts n_1, n_2, \dots, n_B specifying the number of cases in each of B bins. The total number of cases is $T = n_1 + n_2 + \dots + n_B$. Each bin b is defined by an interval $[l_b, u_b)$, $b = 1, \dots, B$, where l_b and u_b are the lower and upper bound of income for that bin. The bottom bin often starts at zero ($l_1 = 0$), and the top bin may have no upper bound ($u_B = \infty$).

A Midpoint Method

The oldest and simplest way to analyze binned data is the midpoint method, which within each bin b assigns incomes to the bin midpoint $m_b = (l_b + u_b)/2$ (Heitjan 1989 for a review). Then statistics such as the Gini can be calculated by applying sample formulas to the midpoints m_b , weighted by the counts n_b .

When the top bin has no upper bound, we must define a pseudo-midpoint for it. The traditional choice is $\mu_B = l_B\alpha/(\alpha - 1)$, which would define the arithmetic mean of the top bin if top-bin incomes followed a Pareto distribution with shape parameter $\alpha > 1$ (Henson 1967). The problem with this choice is that the arithmetic mean of a Pareto distribution is undefined if $\alpha \leq 1$ and grows arbitrarily large as α approaches 1 from above.

A more robust choice is the harmonic mean $h_B = l_B(1 + 1/\alpha)$, which is defined for all $\alpha > 0$ (von Hippel et al. 2015). We use the harmonic mean in this article. We estimate α by fitting a Pareto distribution to the top two bins and calculating the maximum likelihood estimate from the following formula (Quandt 1966):

$$\hat{\alpha} = \frac{\ln((n_{B-1} + n_B)/n_B)}{\ln(l_B/l_{B-1})}. \tag{1}$$

If the survey provides the grand mean μ , then we need not assume that top-bin incomes follow a Pareto distribution. Instead, we can calculate

$$\hat{\mu}_B = \frac{1}{n} \left(T\mu - \sum_{b=1}^{B-1} n_b m_b \right), \tag{2}$$

which would be the mean of the top bin if the means of the lower bins were the midpoints m_b . Then $\hat{\mu}_B$ can serve as a pseudo-midpoint for the top bin.

When estimated in this way, it occasionally happens (e.g., in 4 percent of U.S. counties) that the top bin's mean $\hat{\mu}_B$ is slightly less than its lower bound l_B . This is infelicitous, but $\hat{\mu}_B$ can still be used, and the resulting Gini estimates are not necessarily bad.¹

The midpoint methods described in this section are implemented by the *rpme* command for Stata (Duan and von Hippel 2015), in which *rpme* stands for "robust Pareto midpoint estimator" (von Hippel and Powers 2017). Except for mean-matching, the approach is also implemented in the *binequality* package for R (Scarpino, von Hippel, and Holas 2017).

Fitting Parametric Distributions

The weakness of the midpoint method is that it treats income as discrete. An alternative is to model income as a continuous variable X that fits some parametric CDF $F(X|\theta)$. Here θ is a vector of parameters, which can be estimated by iteratively maximizing the log likelihood:

$$\ell(\theta|X) = \ln \prod_{b=1}^B (P(l_b < X < u_b))^{n_b}, \tag{3}$$

where $P(l_b < X < u_b) = F(u_b) - F(l_b)$ is the probability, according to the fitted distribution, that an income is in the bin $[l_b, u_b)$.²

Although any parametric distribution can be considered, in practice it is hard to fit a distribution unless the number of parameters is small compared to the number of well-populated bins. Most investigators favor two-, three-, and four-parameter distributions from the generalized beta family (McDonald and Ransom 2008), which includes the following 10 distributions: the log normal, the log logistic, the Pareto (type 2), the gamma and generalized gamma, the beta 2 and generalized beta (type 2), the Dagum, the Singh-Maddala, and the Weibull.

A priori, it is hard to know which distribution, if any, will fit well. The fit of a distribution can be tested by the following goodness-of-fit likelihood ratio statistic (von Hippel et al. 2015):

$$G^2 = -2(\hat{\ell} - \sum_{b=1}^B n_b \ln(n_b/T)), \tag{4}$$

where $\hat{\ell}$ is the maximized log likelihood. Under the null hypothesis that the fitted distribution is the true distribution of income, G^2 would follow a chi-square distribution with $B^* - k$ degrees of freedom, where k is the number of parameters, $B^* = \min(B_{>0}, B - 1)$, and $B_{>0}$ is the number of bins with nonzero counts. We reject the fit of a distribution if G^2 has $p < 0.05$ in the null distribution.

In empirical data it is common to reject every distribution in the generalized beta family (Bandourian, McDonald, and Turley 2002; von Hippel et al. 2015). In addition, some distributions may fail to converge or may converge on parameter values that imply that the mean or variance of the distribution is undefined (von Hippel et al. 2015).

A solution is to fit all 10 distributions, screen out any with undefined moments, and among the distributions remaining, select the one that fits best according to the Akaike information criterion (AIC) or Bayesian information criterion (BIC):

$$\begin{aligned} AIC &= 2k - 2\hat{\ell} \\ BIC &= \ln(T)k - 2\hat{\ell}. \end{aligned}$$

Or, instead of selecting a single best-fitting distribution, one can average estimates across several candidate distributions weighted proportionately to a function of the AIC or BIC (specifically $\exp(-AIC/2)$ or $\exp(-BIC/2)$), an approach known as *model averaging* (Burnham and Anderson 2004). In general, model averaging

yields better estimates than model selection, but when modeling binned incomes, the advantage of model averaging is negligible (von Hippel et al. 2015).

Although it sounds broad-minded to fit 10 different distributions, there is limited diversity in the generalized beta family. All the distributions in the generalized beta family are unimodal and skewed to the right. Some distributions are quite similar (e.g., Dagum and generalized beta), and others rarely fit well (e.g., log normal, log logistic, and Pareto). So in practice the range of viable and contrasting distributions in the generalized beta family is small; you can fit just 3 well-chosen distributions (e.g., Dagum, gamma, and generalized gamma) and get estimates almost as good as those obtained from fitting all 10 (von Hippel et al. 2015).

We use the fitted distributions to estimate income statistics such as the mean, variance, Gini, or Theil. Sometimes the income statistic is a simple function of the distributional parameters θ , but other times the function is unknown or hard to calculate. As a general solution, it is easier to calculate income statistics by applying numeric integration to appropriate functions of the fitted distribution (McDonald and Ransom 2008).

When the grand mean is available, the distributional parameters could in theory be constrained to match the grand mean as well as approximate the bin counts. This would be difficult, though, because for distributions in the generalized beta family, the mean is a complicated nonlinear function of the parameters. We have not attempted to constrain our parametric distributions to match a known mean.

The parametric approaches described in this section are implemented in the *binequality* package for R (Scarpino et al. 2017) and the *mgbe* command for Stata (Duan and von Hippel 2015). Here *mgbe* stands for “multimodel generalized beta estimator” (von Hippel et al. 2015).

Interpolated CDFs

Because parametric distributions may fit poorly, an alternative is to define a flexible nonparametric density that fits the bin counts exactly. Our nonparametric approach is called CDF interpolation.

To understand CDF interpolation, consider that binned data define B discrete points on the empirical CDF. The empirical CDF for Nantucket is given in the last column of Table 1, which shows that 5 percent of Nantucket households make less than \$10,000, 8 percent make less than \$15,000, and so on. Formally, at an income of 0 the empirical CDF is $\hat{F}(0) = 0$, and at each bin's upper bound u_b the empirical CDF is the observed fraction of incomes that are less than u_b —that is, $\hat{F}(u_b) = (n_1 + n_2 + \dots + n_b)/T$.³

Now, to estimate a continuous CDF $\hat{F}(x)$, $x > 0$, we just connect the dots. That is, we define a continuous nondecreasing function that interpolates between the B discrete points of the empirical CDF.

Then the estimated probability density function (PDF) is just the derivative of the interpolated CDF $\hat{F}(x)$. Note that the estimated PDF “preserves areas” (Morandi and Costantini 1989). That is, the fraction of incomes that should be in a bin according to the fitted distribution ($\hat{P}(l_b < x < u_b) = \hat{F}(u_b) - \hat{F}(l_b)$) is equal to the observed fraction of incomes that are actually in that bin (n_b/T).

The shape of the PDF depends on the function that interpolates the CDF:

- If the CDF is interpolated by line segments, then the CDF is polygonal, and the PDF is a step function that is discontinuous at the bin boundaries.
- If the CDF is interpolated more smoothly by a continuously differentiable monotone cubic spline, then the PDF is piecewise quadratic—continuous at the bin boundaries and quadratic between them.

There remains a question of how to shape the CDF in the top bin, which typically has no upper bound. In our implementation, the CDF of the top bin can be rectangular, exponential, or Pareto.⁴ Each of these distributions has one parameter, which is estimated as follows:

- If the grand mean of income is known, then the parameter shaping the top bin is constrained so that the mean of the fitted distribution matches the grand mean. It occasionally happens (e.g., in 4 percent of U.S. counties) that we cannot reproduce the known mean this way because the known mean is already less than the mean of the lower $B - 1$ bins without the tail. In that case, we make an ad hoc adjustment by shrinking the bin boundaries toward the origin—that is, by replacing (l_b, u_b) with (sl_b, su_b) , where $s < 1$ is a shrinkage factor chosen so that a small tail can be added to reproduce the grand mean. The shrinkage factor is rarely less than 0.995.
- If the mean income is not known, we substitute an ad hoc estimate. We obtain that estimate by temporarily setting the upper bound of the top bin to $u_B = 2l_B$ and calculating the mean of a step PDF fit to all B bins. Then we unbound the top bin and proceed as though the mean were known.

Income statistics are estimated by applying numerical integration to functions of the fitted PDF or CDF.

The methods in this section are implemented by our *binsmooth* package for R (Hunter and Drown 2016). Within the *binsmooth* package, the *stepbins* function implements a step-function PDF (and polygonal CDF), whereas the *splinebins* function implements a cubic spline CDF (and piecewise quadratic PDF).

Recursive Subdivision

Another way to obtain a smooth PDF estimate that preserves bin areas is to subdivide the bins into smaller bins and then adjust the heights of the subdivided bins to shorten the jumps at the bin boundaries. This method, *recursive subdivision*, is implemented by the *rebin* command in our *binsmooth* package for R (Hunter and Drown 2016). Recursive subdivision is slower and more computationally intensive than CDF interpolation, and the resulting estimates are practically identical. We present the details of recursive subdivision in the Appendix.

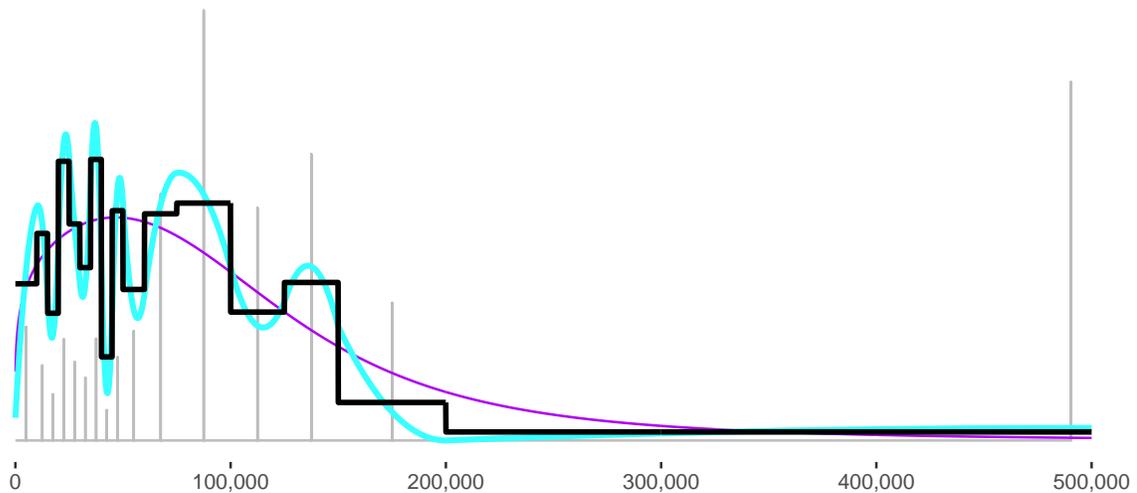


Figure 1: Different PDFs for the Nantucket data. The Dagum distribution is drawn in purple, the quadratic spline in blue, and the step function in black. Gray spikes illustrate the midpoint method.

Data and Results

Between 2006 and 2010, the American Community Survey (ACS) took a one-in-eight sample of U.S. households (U.S. Census Bureau 2014). Household incomes were inflated to 2010 dollars and summarized in binned income tables for each of the 3,221 US counties. The published bin counts are estimates of the population counts. We can approximate the sample counts by dividing the population counts by eight. Dividing counts by a constant makes no difference to any of our statistics except for the BIC and G^2 statistics that are sometimes used when fitting parametric distributions.

The U.S. Census Bureau also published means and Gini coefficients for each county (Bee 2012). These statistics were estimated from exact incomes before binning, so they are more accurate than any estimate that could be calculated from the binned data. They are sample estimates which may differ from population values, but they remain a useful standard of comparison for our binned-data estimates.

Results for Nantucket

Table 1 summarized the binned incomes for Nantucket county. Figure 1 fits several distributions to the Nantucket data.

The midpoint method is illustrated by gray spikes at the bin midpoints; the spike heights are proportional to the bin counts. The black step function is the PDF implied by a linear interpolation of the CDF, and the blue curve is the piecewise quadratic PDF implied by a cubic spline of the CDF. Both CDF interpolations fit the bin counts perfectly; in fact, their jagged appearance suggests they may *overfit* the data—a concern that we will revisit in the Conclusion. The step PDF looks slightly

Table 2: Estimates for Nantucket.

	Estimate	Mean	Gini
Without the true mean	True	\$137,811	0.547
	Midpoint (<i>rpme</i>)	\$121,506	0.464
	Parametric (<i>mgbe</i>)	\$112,960	0.453
	CDF interpolation (linear)	\$110,419	0.438
Matching the true mean	CDF interpolation (cubic spline)	\$110,419	0.433
	Midpoint		0.510
	CDF interpolation (linear)		0.537
	CDF interpolation (cubic spline)		0.525

rpme = robust Pareto midpoint estimator; *mgbe* = multimodel generalized beta estimator.

less volatile than the piecewise quadratic PDF, suggesting that the step PDF may be less overfit.

The purple curve is the Dagum distribution. The Dagum fits Nantucket better than other distributions from the generalized beta family, but it does not fit well. It fails the G^2 goodness-of-fit test, and visually it fits the bin counts poorly. For example, between \$70,000 and \$150,000, the Dagum curve suggests there should be substantially fewer households than there are, and above \$150,000, it suggests that there should be more.

Except for the Dagum distribution, all the methods in Figure 1 are calibrated to reproduce the grand mean.

Table 2 summarizes the Nantucket estimates. The true mean is \$137,000 and the true Gini is 0.547. When fit without knowledge of the true mean, every method underestimates the mean by 12 to 20 percent and the Gini by 15 to 21 percent, with the simple midpoint method coming closer than its more sophisticated competitors.

When given the true mean, the midpoint and CDF interpolation methods do much better. They still underestimate the Gini but only by 2 to 7 percent. The closest estimate is obtained by linear interpolation of the CDF. A smoother cubic spline interpolation does a little worse but still better than the midpoint method.

Although the estimates for Nantucket are less accurate overall than the estimates for most other counties, the relative performance of different methods in Nantucket is similar to what we will see elsewhere.

Results for All U.S. Counties

Figure 2 evaluates all county Gini estimates graphically by plotting the estimated Gini $\hat{\theta}_j$ of each county against the published Gini θ_j . In the bottom row, where the methods are constrained to match the published mean, the Gini estimates are close to a diagonal reference line ($\hat{\theta}_j = \theta_j$), indicating nearly perfect estimation. In the top row, where the methods do not match the mean, the estimates are more scattered, indicating lower accuracy.

We can summarize the accuracy of Gini estimates in several ways. For a single county j , the percent estimation error is $e_j = 100 \times (\hat{\theta}_j - \theta_j) / \theta_j$. Then, across all

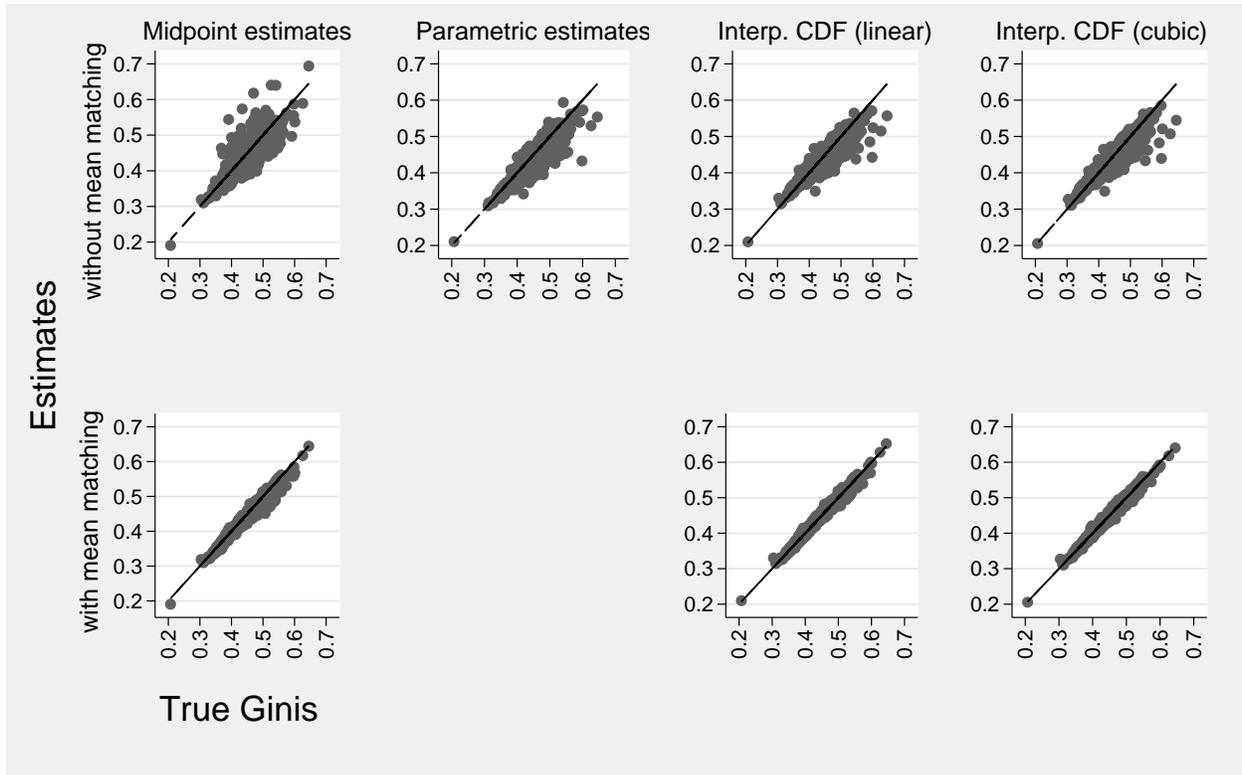


Figure 2: Accuracy of Gini coefficients from 3,221 U.S. counties estimated by different methods, with and without mean-matching.

counties, the *percent relative bias* is the mean of e_j , the *percent root mean squared error* (RMSE) is the square root of the mean of e_j^2 , and the *reliability* is the squared correlation between θ_j and $\hat{\theta}_j$.

Table 3 summarizes our findings. If the estimators ignore the published county means, then estimated Ginis have biases between 0 percent and -3 percent, RMSEs between 3 percent and 4 percent, and reliabilities between 82 percent and 88 percent. The interpolated CDF estimates have the best bias, the best RMSE, and the second-best reliability, and they are just as good with linear interpolation as with cubic spline interpolation. The parametric estimates have the best reliability but the worst bias, the worst RMSE, and by far the worst runtime, at 4.5 hours.

When the methods are constrained to match the published county means, the estimates improve dramatically. The bias shrinks to 0 to 1 percent, the RMSE shrinks to 1 to 2 percent, and the reliability grows to 98 to 99 percent. The midpoint estimates are excellent; the interpolated CDF estimates are even better, and just as good with linear interpolation as with cubic spline interpolation.

The differences among the methods are much smaller than the improvement that comes from constraining any method to match the mean. Of course, this observation is only helpful when the mean is known.

Table 3: Speed and accuracy of Gini estimates for 3,221 U.S. counties.

	Estimator	Bias	Root mean squared error	Reliability	Runtime
Without true mean	Midpoint (<i>rpme</i>)	-2%	4%	82%	4 sec
	Parametric	-3%	4%	89%	4.5 hr
	CDF interpolation (linear)	0%	3%	88%	40 sec
	CDF interpolation (cubic spline)	-1%	3%	88%	2 min
Matching true mean	Midpoint	-1%	2%	98%	7 sec
	CDF interpolation (linear)	0%	1%	99%	36 sec
	CDF interpolation (cubic spline)	-1%	1%	99%	2 min

Runtimes on an Intel i7 Core processor with a speed of 3.6 GHz. *rpme* = robust Pareto midpoint estimator.

Conclusion

CDF interpolation produces estimates that are at least a little better than midpoint or parametric estimates, whether the true mean is known or not. And CDF interpolation runs much faster than parametric estimation, though not as fast as midpoint estimation.

We initially suspected that cubic spline interpolation would improve on simple linear interpolation, but empirically this turns out to be false. In estimating county Ginis, linear CDF interpolation was at least as accurate as cubic spline interpolation.

The accuracy of linear CDF interpolation is remarkable because it implies a step function for the PDF. Step PDFs look unrealistic, especially in the top and bottom bins where the step function is flat, whereas the true distribution likely has an upward or downward slope (Cloutier 1995). Our step PDF permits a downward Pareto or exponential slope in the top bin, and our cubic spline CDF can fit an upward slope to the bottom bin. But neither of these refinements does much to improve the accuracy of Gini estimates.

The differences in accuracy among the methods are small, and they are dwarfed by the improvement in accuracy that comes from knowing the grand mean. By constraining binned-data methods to match a known mean, we can typically get county Gini estimates that are within 1 to 2 percent of the estimates we would get if the data were not binned. Our *binsmooth* package for R can constrain interpolated CDFs to match a known mean, and our *rpme* command for Stata can constrain the top-bin midpoint to match a known mean as well. We have not constrained our parametric distributions to match a known mean; we believe it would be difficult to do so.

Although the mean-constrained estimates are very accurate, there may be room for improvement when the mean is unknown. Perhaps the most promising idea is additional smoothing. As we noticed in Figure 1, interpolated CDFs can be a bit jagged and may overfit the sample in the sense that they find nooks and crannies that might not appear in another sample or in the population. Likewise, interpolated CDFs may be overfit to a specific set of bin boundaries. If the fitted CDF were a little smoother and did not quite preserve the counts of the least populous bins, it

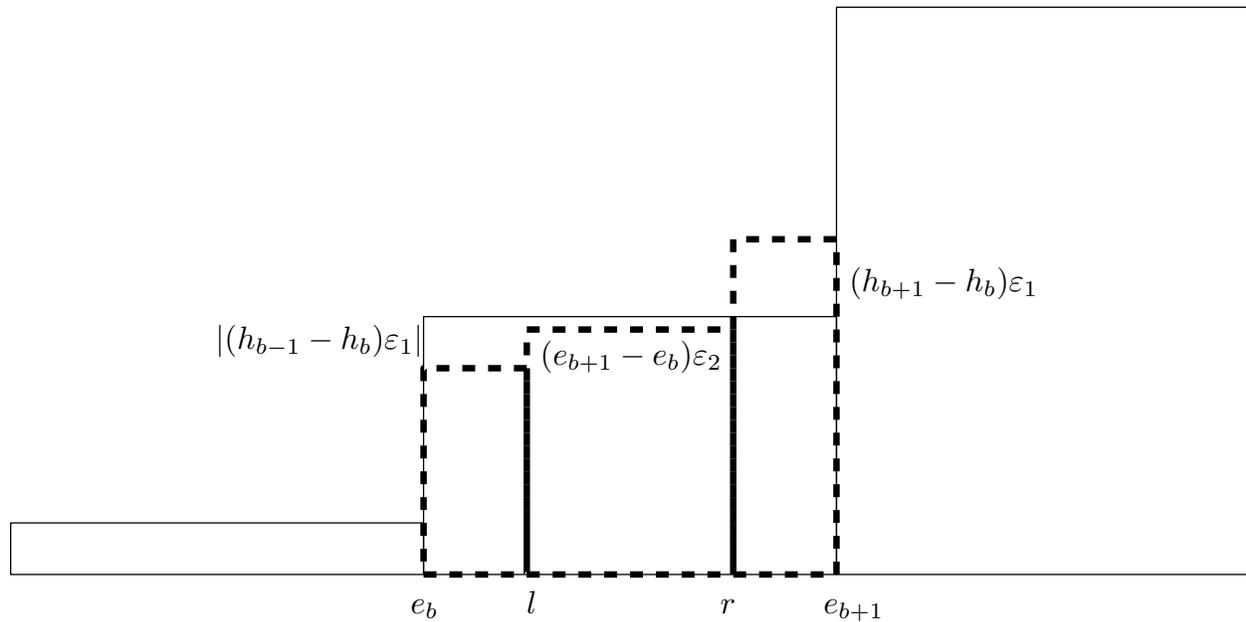


Figure 3: Bin subdivision. Each original bin is replaced by three new bins (bold, dashed) such that the bin area is preserved.

might fit the population and other samples (perhaps with different bin boundaries) a little better.

Appendix: PDF Smoothing by Recursive Subdivision

Recursive subdivision is another way to smooth the fitted PDF. Like CDF interpolation, recursive subdivision preserves bin areas, but recursive subdivision is more computationally intensive and produces very similar results. Recursive subdivision is implemented by the *recbins* function in our *binsmooth* package for R.

A slight change of notation will be helpful. Since the upper bound u_b of each bin is equal to the lower bound l_{b+1} of the next, we can think the bins as having a set of “edges” e_0, e_1, \dots, e_B , where $e_0 = 0$, and the other $e_b = u_b$ are the upper bounds of bins $1, \dots, B$.

Start by fitting a step PDF. Let h_b be the height of the step PDF in the bin $[e_b, e_{b+1})$. Given parameters $\epsilon_1 \in (0, 0.5)$ and $\epsilon_2 \in (0, 1)$, the subdivision process begins by introducing new bin edges l and r between e_b and e_{b+1} such that $(l + r)/2 = (e_b + e_{b+1})/2$ and $r - l = (e_{b+1} - e_b)\epsilon_2$. The height of the new bin on the left with edges e_b and l is then shifted horizontally by $(h_{b-1} - h_b)\epsilon_1$, and the height of the new bin on the right with edges r and e_{b+1} is shifted horizontally by $(h_{b+1} - h_b)\epsilon_1$. Finally, the new middle bin with edges l and r is shifted horizontally so that the area of the three new bins equals the area of the original bin.⁵ See Figure 3.

In order for the above formulas to apply to the top and bottom bins, we create pseudo-bins above and below them with heights of zero. This ensures that the

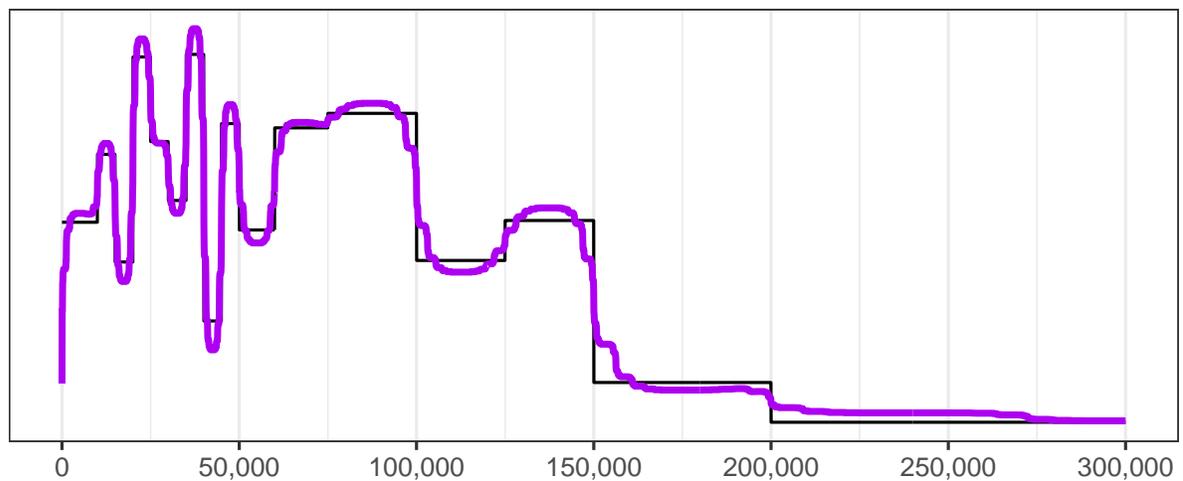


Figure 4: Recursively subdivided density for the Nantucket data (purple). A step function (black) is shown in the background.

subdivided PDF will tend toward a height of zero at the lower edge of the bottom bin and the upper edge of the top bin.

The smoothed PDF is obtained from the step PDF by applying the subdivision process to each bin, then applying the process again to each subdivided bin, and so on until the desired level of smoothness is reached. In practice, three rounds of subdivision are sufficient to produce a reasonably smooth PDF, and we found that choosing $\varepsilon_1 = 0.25$ and $\varepsilon_2 = 0.75$ produced nicely smoothed PDFs from most empirical data sets. Figure 4 shows the result of recursive subdivision in Nantucket.

Unfortunately, if the original step PDF was constrained to match a known mean, the subdivision process may cause the mean to deviate slightly. But the estimated Gini typically remains quite accurate.

Notes

- 1 It might be a little better to set $\hat{\mu}_B = l_B$ and move the other midpoints slightly to the left, but this would affect only 4 percent of counties, and those only slightly because the affected counties typically have few cases in the top bin.
- 2 This formula includes the top bin B , where $u_B = \infty$ and $F(u_B) = 1$. Some older articles (e.g., McDonald and Ransom 1979; McDonald and Xu 1995; Bandourian et al. 2002) add the following constant to the log likelihood: $\ln(T!) - \sum_{b=1}^B \ln(n_b!)$. This is not wrong, but it is unnecessary because adding a constant does not change the parameter values at which the log likelihood is maximized (Edwards 1972).
- 3 If the binned incomes come from a sample (as they do in Table 1), then the estimate $\hat{F}(u_b)$ may differ from the true CDF $F(u_b)$ because of sampling error.
- 4 In our implementation, the exponential and Pareto tails are approximated by a sequence of rectangles of decreasing heights.

⁵ We do not subdivide the bin in those rare cases in which subdivision would yield a middle bin with negative height.

References

- Bandourian, Ripsy, James B. McDonald, and Robert S. Turley. 2002. "A Comparison of Parametric Models of Income Distribution Across Countries and Over Time." *SSRN Electronic Journal*. Luxembourg Income Study Working Paper No. 305. <https://doi.org/10.2139/ssrn.324900>.
- Bee, Adam. 2012. "Household Income Inequality Within U.S. Counties: 2006–2010." American Community Survey Briefs No. ACSBR/10-18. Washington, D.C.: U.S. Census Bureau. <https://www.census.gov/prod/2012pubs/acsbr10-18.pdf>.
- Burnham, Kenneth P., and David R. Anderson. 2004. "Multimodel Inference." *Sociological Methods & Research* 33:261–304. <https://doi.org/10.1177/0049124104268644>.
- Cloutier, Norman R. 1995. "Lognormal Extrapolation and Income Estimation for Poor Black Families." *Journal of Regional Science* 35:165–171. <https://doi.org/10.1111/j.1467-9787.1995.tb01405.x>.
- Duan, Yutong, and Paul T. von Hippel. 2015. *mgbe – Multimodel Generalized Beta Estimator*. Stata Command Version 1.0. <https://ideas.repec.org/c/boc/bocode/s458189.html>.
- Edwards, A. W. F. 1972. *Likelihood*. Baltimore, MD: Johns Hopkins University Press.
- Heitjan, Daniel F. 1989. "[Inference from Grouped Continuous Data: A Review]: Rejoinder." *Statistical Science* 4:182–183. <https://doi.org/10.1214/ss/1177012603>.
- Henson, M. F. 1967. *Trends in the Income of Families and Persons in the United States, 1947-1964*. Washington, D.C.: U.S. Department of Commerce, Bureau of the Census.
- Hunter, David J., and McKalie Drown. 2016. *binsmooth: Generate PDFs and CDFs from Binned Data*. R Package Version 0.1.0. Comprehensive R Archive Network. <https://CRAN.R-project.org/package=binsmooth>.
- Kooperberg, Charles, and Charles J. Stone. 1992. "Logspline Density Estimation for Censored Data." *Journal of Computational and Graphical Statistics* 1:301. <https://doi.org/10.2307/1390786>.
- McDonald, James B., and Michael Ransom. 2008. "The Generalized Beta Distribution as a Model for the Distribution of Income: Estimation of Related Measures of Inequality." In *Modeling Income Distributions and Lorenz Curves*, pp. 147–166. New York: Springer.
- McDonald, James B., and Michael R. Ransom. 1979. "Alternative Parameter Estimators Based upon Grouped Data." *Communications in Statistics - Theory and Methods* 8:899–917. <https://doi.org/10.1080/03610927908827806>.
- McDonald, James B., and Yexiao J. Xu. 1995. "A Generalization of the Beta Distribution with Applications." *Journal of Econometrics* 66:133–152. [https://doi.org/10.1016/0304-4076\(94\)01612-4](https://doi.org/10.1016/0304-4076(94)01612-4).
- Minoiu, Camelia, and Sanjay G. Reddy. 2012. "Kernel Density Estimation on Grouped Data: The Case of Poverty Assessment." *The Journal of Economic Inequality* 12:163–189. <https://doi.org/10.1007/s10888-012-9220-9>.

- Morandi, Rossana, and Paolo Costantini. 1989. "Piecewise Monotone Quadratic Histosplines." *SIAM Journal on Scientific and Statistical Computing* 10:397–406. <https://doi.org/10.1137/0910026>.
- Quandt, Richard E. 1966. "Old and New Methods of Estimation and the Pareto Distribution." *Metrika* 10:55–82. <https://doi.org/10.1007/bf02613419>.
- Scarpino, Samuel V., Paul T. von Hippel, and Igor Holas. 2017. *binequality: Methods for Analyzing Binned Income Data*. R Package Version 1.02. Comprehensive R Archive Network. <https://cran.r-project.org/web/packages/binequality/index.html>.
- U.S. Census Bureau. 2014. "American FactFinder." <https://factfinder.census.gov/faces/nav/jsf/pages/index.xhtml>.
- von Hippel, Paul T., Igor Holas, and Samuel V. Scarpino. 2012. "Estimation with Binned Data." <https://arxiv.org/abs/1210.0200>.
- von Hippel, Paul T., and Daniel A. Powers. 2017. *rpme – Robust Pareto Midpoint Estimator*. Statistical Software Components S457962, Boston College Department of Economics. <https://ideas.repec.org/c/boc/bocode/s457962.html>.
- von Hippel, Paul T., Samuel V. Scarpino, and Igor Holas. 2015. "Robust Estimation of Inequality from Binned Incomes." *Sociological Methodology* 46:212–51. <http://journals.sagepub.com/doi/abs/10.1177/0081175015599807>.
- Wahba, Grace. 1976. "Histosplines with Knots Which Are Order Statistics." *Journal of the Royal Statistical Society. Series B (Methodological)* 38:140–151. <http://www.jstor.org/stable/2985018>.
- Wang, Bin. 2015. *bda: Density Estimation for Grouped Data*. R Package Version 5.1.6. Comprehensive R Archive Network. <https://CRAN.R-project.org/package=bda>.

Acknowledgements: Drown is grateful for support from a Tensor Grant of the Mathematical Association of America.

Paul T. von Hippel: Lyndon B. Johnson School of Public Affairs, University of Texas at Austin. E-mail: paulvonhippel.utaustin@gmail.com

David J. Hunter: Department of Mathematics and Computer Science, Westmont College. E-mail: dhunter@westmont.edu.

McKalie Drown: Department of Mathematics and Computer Science, Westmont College. E-mail: mdrown@westmont.edu.