

Binscatter Regressions

Matias D. Cattaneo
Princeton University
Princeton, NJ
cattaneo@princeton.edu

Richard K. Crump
Federal Reserve Bank of New York
New York, NY
richard.crump@ny.frb.org

Max H. Farrell
UC Santa Barbara
Santa Barbara, CA
mhfarrell@gmail.com

Yingjie Feng
Tsinghua University
Beijing, China
fengyj@sem.tsinghua.edu.cn

Abstract. We introduce the package `Binsreg`, which implements the binscatter methods developed by [Cattaneo, Crump, Farrell, and Feng \(2024b,a\)](#). The package includes seven commands: `binsreg`, `binslogit`, `binsprobit`, `binsqreg`, `binstest`, `binspwc`, and `binsregselect`. The first four commands implement binscatter plotting, point estimation, and uncertainty quantification (confidence intervals and confidence bands) for least squares linear binscatter regression (`binsreg`) and for nonlinear binscatter regression (`binslogit` for Logit regression, `binsprobit` for Probit regression, and `binsqreg` for quantile regression). The next two commands focus on pointwise and uniform inference: `binstest` implements hypothesis testing procedures for parametric specifications and for nonparametric shape restrictions of the unknown regression function, while `binspwc` implements multi-group pairwise statistical comparisons. Finally, the command `binsregselect` implements data-driven number of bins selectors. The commands offer binned scatter plots, and allow for covariate adjustment, weighting, clustering, and multi-sample analysis, which is useful when studying treatment effect heterogeneity in randomized and observational studies, among many other features.

Keywords: st0001, binscatter, binned scatter plot, nonparametrics, semiparametrics, partitioning estimators, B-splines, tuning parameter selection, confidence bands, shape and specification testing.

July 22, 2024

1 Introduction

Data visualization is a crucial step in any statistical analysis. The classical scatter plot is a fundamental tool for visualization, used for studying how an outcome y relates to a continuous covariate of interest x . However, in “big data” settings, such as the administrative data sets now common in social science and medical research, the classical scatter plot yields a dense cloud of points that is not informative. Further, there is no way to create the plot rigorously controlling for other covariates, which would be standard in any subsequent statistical analysis. A binned scatter plot, or binscatter, is a visualization method that addresses these limitations. Binscatter techniques offer flexible, yet parsimonious ways of visualizing and summarizing regression (and other) functions. Binscatters have become popular in applied microeconomics for visualization, specification testing, treatment effect heterogeneity, and other uses. See [Starr and Goldfarb \(2020\)](#), and references therein, for current usage.

However, little was known about the statistical properties of binscatter until recently. [Cattaneo, Crump, Farrell, and Feng \(2024b,a\)](#), collectively CCFE hereafter) provided the first foundational and comprehensive analysis of binscatter methods, including an array of theoretical and practical results that aid both in understanding current practices (i.e., validity or lack thereof) and in offering theory-based guidance for future use. Maintaining rigor and statistical uncertainty are crucial for trustworthy data visualization in scientific settings ([Healy 2018](#); [Schwabish 2021](#)).

This paper introduces the package `Binsreg`, which includes seven commands implementing the main methodological results in CCFE. These commands are organized as follows.

- *Estimation, uncertainty quantification, and plotting.* The command `binsreg` implements canonical and extended least squares binscatter methods, while the commands `binslogit`, `binsprobit`, and `binsqreg` implement generalized nonlinear binscatter methods (i.e., Logistic regression, Probit regression and quantile regression, respectively). All commands allow for higher-order polynomial fits within bins, smoothness restrictions across bins, and covariate adjustment for estimation, uncertainty quantification and plotting, also covering higher-order derivatives and related partial effects of interest in linear and nonlinear settings, including multi-sample comparisons.
- *Hypothesis testing and statistical inference.* The command `binstest` implements hypothesis testing procedures for parametric specifications and for nonparametric shape restrictions of the unknown regression function, while `binspwc` implements multi-group pairwise comparisons. These two commands offer the same flexibility and features as the estimation commands and therefore allow for linear and nonlinear binscatter methods with within-bin higher-order polynomial fits, across-bins smoothness restrictions, and semi-linear covariate adjustments, among several other features and options.
- *Optimal number of bins selection.* The six commands above take as input the

binning scheme to construct the binscatter approximation, which requires selecting the position of the bins as well as the total number of bins on the support of the independent variable of interest. Whenever this information is not provided, the command `binsregselect` implements data-driven selectors for the number of bins for implementation using either quantile-spaced or evenly-spaced binning/partitioning (quantile-spaced binning is chosen by default, following popular empirical practice).

The seven commands in the package `Binsreg` offer several other important functionalities for empirical work. First, the commands incorporate, by default, mass point and degrees of freedom checks and adjustments, which improve the stability of the implementation. Second, the commands allow for multi-way fixed effects and clustering estimation and inference whenever available in the underlying statistical software platform. Third, in `Stata`, the commands offer the option of estimation and inference with multi-way fixed effects and multi-way clustering via the community-distributed package `reghdfe` (Correia and Constantine 2024), and also allow for using the community-distributed package `gtools` (Caceres 2024), instead of our internal implementations, to potentially increase the speed of internal computations with ultra-large datasets. Depending on the data size and structure, the commands in the package `Binsreg` may improve implementation execution speed when (i) mass point and degrees of freedom checks are turned off, (ii) the user-written packages `reghdfe` and `gtools` instead of our internal (open-source) implementations are used (in `Stata`). See Section 2.9 for more discussion.

There exist two other user-written `Stata` commands implementing binscatter methods: `binscatter` (Stepner 2017) and `binscatter2` (Droste 2019). Both of those packages incorporate other covariates or fixed effects incorrectly, yielding invalid results. Even without additional controls, those two packages only give the binned scatter plot, only for piecewise constant estimation, and only in least squares regression, lacking all the theoretically-founded features of `Binsreg` such as nonlinear models, valid uncertainty visualization, formal specification and shape testing, group comparisons, and optimal binning selection. See CCFE for further discussion.

The rest of the article is organized as follows. Section 2 gives an overview of the main methods available in the package `Binsreg` and discusses some implementation details. Section 3 gives a numerical illustration. Section 4 concludes. The software help files contain a detailed description of all available options. The latest version of the package `Binsreg` and other related materials can be found at:

<https://nppackages.github.io/binsreg/>.

2 Methods Overview and Implementation Details

This section summarizes the main methods implemented in the package `Binsreg`. For further methodological and theoretical details see CCFE.

Given a random sample $(y_i, x_i, \mathbf{w}'_i)$, $i = 1, 2, \dots, n$, where y_i is a scalar response

variable, x_i is the scalar independent variable of interest, and \mathbf{w}_i is a d -dimensional vector of additional covariates, binscatter seeks to flexibly approximate the function

$$\vartheta_{\mathbf{w}}^{(v)}(x) = \frac{\partial^v}{\partial x^v} \eta(\mu_0(x) + \mathbf{w}'\boldsymbol{\gamma}_0), \quad (1)$$

where \mathbf{w} is some user-chosen evaluation point, and the underlying parameters $\mu_0(\cdot)$ and $\boldsymbol{\gamma}_0$ are defined by

$$(\mu_0(\cdot), \boldsymbol{\gamma}_0) = \arg \min_{\mu \in \mathcal{M}, \boldsymbol{\gamma} \in \mathbb{R}^d} \mathbb{E}[\rho(y_i; \eta(\mu(x_i) + \mathbf{w}'_i \boldsymbol{\gamma}))], \quad (2)$$

with $\rho(\cdot; \cdot)$ and $\eta(\cdot)$ user-chosen loss and (inverse) link functions, respectively, and \mathcal{M} an appropriate space of functions satisfying certain conditions. Several settings of applied interest are covered by this formulation. (For any function $f(x)$, we define $f^{(v)}(x) = d^v f(x)/dx^v$, with the usual notation $f(x) = f^{(0)}(x)$.)

- *Semi-linear regression*: $\rho(y; \eta) = (y - \eta)^2$ and $\eta(u) = u$. The parameter of interest becomes

$$\vartheta_{\mathbf{w}}^{(v)}(x) = \frac{\partial^v}{\partial x^v} \mathbb{E}[y_i | x_i = x, \mathbf{w}_i = \mathbf{w}] = \begin{cases} \mu_0(x) + \mathbf{w}'\boldsymbol{\gamma}_0 & \text{if } v = 0 \\ \mu_0^{(v)}(x) & \text{if } v \geq 1 \end{cases}.$$

For example, $\vartheta_{\mathbf{w}}(x)$ (resp. $\vartheta_{\mathbf{w}}^{(1)}(x)$) corresponds to the average (partial) effect of x on y for level $\mathbf{w}_i = \mathbf{w}$. In this setting, $\vartheta_{\mathbf{w}}^{(v)}(x) = \mu_0^{(v)}(x)$, which may be of interest in some applications.

- *Logistic/Probit regression*: $\rho(y; \eta) = -y \log \eta - (1 - y) \log(1 - \eta)$ and $\eta(u)$ denotes the (inverse) link function of Logistic or Probit regression. The parameter of interest becomes

$$\vartheta_{\mathbf{w}}^{(v)}(x) = \frac{\partial^v}{\partial x^v} \mathbb{E}[y_i | x_i = x, \mathbf{w}_i = \mathbf{w}] = \frac{\partial^v}{\partial x^v} \eta(\mu_0(x) + \mathbf{w}'\boldsymbol{\gamma}_0),$$

which coincides with the usual average (partial) effect in binary response models.

- *Quantile regression*: $\rho(y; \eta) = \ell_{\tau}(y - \eta)$ and $\eta(u) = u$, where $\ell_{\tau}(u)$ denotes the check function associated with the τ -th quantile. The parameter of interest becomes

$$\vartheta_{\mathbf{w}}^{(v)}(x) = \frac{\partial^v}{\partial x^v} Q_{\tau}(y_i | x_i = x, \mathbf{w}_i = \mathbf{w}),$$

where $Q_{\tau}(y_i | x_i = x, \mathbf{w}_i = \mathbf{w}) = \mu_0(x) + \mathbf{w}'\boldsymbol{\gamma}_0$ denotes the conditional τ -th quantile regression function of y_i given $x_i = x, \mathbf{w}_i = \mathbf{w}$.

The different parameters above, as well as many others, are determined by the choice of loss function $\rho(\cdot; \cdot)$ and (inverse) link function $\eta(\cdot)$. In the above formulation, we assume that the models are correctly specified relative to the true data generating process (i.e., relative to the assumptions on the probability distribution of the data

$(y_i, x_i, \mathbf{w}'_i)$, $i = 1, 2, \dots, n$), which is also assumed to be an i.i.d. sample. However, in many settings, the choices of $\rho(\cdot; \cdot)$ and $\eta(\cdot)$ are only working models, which may not lead to the underlying target parameter but rather only to an approximation thereof in a principled way. That is, under incorrect specification, the parameter $\vartheta_{\mathbf{w}}^{(v)}(x)$ can only be interpreted as the solution to the minimization in (2). Under dependent data, binscatter can still be applied but statistical guarantees for parameter estimation and for uncertainty quantification are not available.

2.1 Binscatter Construction

To approximate $\mu_0(x)$ and its derivatives in model (2), binscatter first partitions the support of x_i into J quantile-spaced bins, leading to the partitioning scheme:

$$\widehat{\Delta} = \{\widehat{\mathcal{B}}_1, \dots, \widehat{\mathcal{B}}_J\}, \quad \widehat{\mathcal{B}}_j = \begin{cases} [x_{(1)}, x_{(\lfloor n/J \rfloor)}] & \text{if } j = 1, \\ [x_{(\lfloor n(j-1)/J \rfloor)}, x_{(\lfloor nj/J \rfloor)}] & \text{if } j = 2, \dots, J-1, \\ [x_{(\lfloor n(J-1)/J \rfloor)}, x_{(n)}] & \text{if } j = J, \end{cases}$$

where $x_{(i)}$ denotes the i -th order statistic of the sample $\{x_1, x_2, \dots, x_n\}$, $\lfloor \cdot \rfloor$ is the floor operator, and $J < n$. Each estimated bin $\widehat{\mathcal{B}}_j$ contains roughly the same number of observations $N_j = \sum_{i=1}^n \mathbb{1}_{\widehat{\mathcal{B}}_j}(x_i)$, where $\mathbb{1}_{\mathcal{A}}(x) = \mathbb{1}(x \in \mathcal{A})$ with $\mathbb{1}(\cdot)$ denoting the indicator function. This binning approach is the most popular in empirical work but, for completeness, all commands in the package `Binsreg` also allow for evenly-spaced binning and user-specified binning. See below for more implementation details.

Given the quantile-spaced partitioning/binning scheme, for a choice of number of bins J , and a choice of loss function $\rho(\cdot; \cdot)$ and (inverse) link function $\eta(\cdot)$, the generalized nonlinear binscatter estimator of the v -th derivative $\vartheta_{\mathbf{w}}^{(v)}(x)$ of $\eta(\mu_0(x) + \mathbf{w}'\gamma_0)$ in (1), employing a p -th order polynomial approximation within each bin, imposing $(s-1)$ -times differentiability across bins, and adjusting for additional covariates \mathbf{w}_i , is

$$\widehat{\vartheta}_{\mathbf{w}}^{(v)}(x) = \frac{\partial^v}{\partial x^v} \eta(\widehat{\mu}(x) + \mathbf{w}'\widehat{\gamma}) \quad (3)$$

where

$$\widehat{\mu}^{(v)}(x) = \widehat{\mathbf{b}}_{p,s}^{(v)}(x)' \widehat{\beta}, \quad \begin{bmatrix} \widehat{\beta} \\ \widehat{\gamma} \end{bmatrix} = \arg \min_{\beta, \gamma} \sum_{i=1}^n \rho(y_i; \eta(\widehat{\mathbf{b}}_{p,s}(x_i)' \beta + \mathbf{w}'_i \gamma)), \quad (4)$$

with $s \leq p$, $v \leq p$ and $\widehat{\mathbf{b}}_{p,s}(x) = \widehat{\mathbf{T}}_s \widehat{\mathbf{b}}_{p,0}(x)$ with

$$\widehat{\mathbf{b}}_{p,0}(x) = [\mathbb{1}_{\widehat{\mathcal{B}}_1}(x) \quad \mathbb{1}_{\widehat{\mathcal{B}}_2}(x) \quad \dots \quad \mathbb{1}_{\widehat{\mathcal{B}}_J}(x)]' \otimes [1 \quad x \quad \dots \quad x^p]'$$

being the p -th order polynomial basis of approximation within each bin, hence of dimension $(p+1)J$, and $\widehat{\mathbf{T}}_s$ being a $[(p+1)J - (J-1)s] \times (p+1)J$ matrix of linear restrictions ensuring that the $(s-1)$ -th derivative of $\widehat{\mu}(x)$ is continuous.

When $s = 0$, $\widehat{\mathbf{T}}_0 = \mathbf{I}_{(p+1)J}$, the identity matrix of dimension $(p+1)J$, and therefore no restrictions are imposed: $\widehat{\mathbf{b}}_{p,0}(x)$ is the basis used for (disjoint) piecewise p -th order polynomial fits. Consequently, the binscatter $\widehat{\mu}(x)$ is discontinuous at the bins' edges whenever $s = 0$. On the other hand, $p \geq s$ implies that a p -th order polynomial fit is constructed within each bin $\widehat{\mathcal{B}}_j$, in which case setting $s = 1$ forces these fits to be connected at the boundaries of adjacent bins (leading to a continuous but nondifferentiable function), $s = 2$ forces these fits to be connected and continuously differentiable at the boundaries of adjacent bins, and so on. Enforcing smoothness on binscatter boils down to incorporating restrictions on the basis of approximation. The resulting constrained basis, $\widehat{\mathbf{b}}_{p,s}(x)$, corresponds to a choice of spline basis for approximation of $\mu_0(\cdot)$ in (2), with estimated quantile-spaced knots according to the partition $\widehat{\Delta}$. The package `Binsreg` employs $\widehat{\mathbf{T}}_s$ leading to B-splines, which tend to have very good finite sample properties.

The binscatter estimator $\widehat{\vartheta}_{\mathbf{w}}^{(v)}(x)$ in (3) is a plug-in estimator for (2). Returning to the settings of applied interest mentioned previously, we have:

- *Semi-linear regression*: $\rho(y; \eta) = (y - \eta)^2$ and $\eta(u) = u$. This case corresponds to linear least squares binscatter, where the estimator becomes

$$\widehat{\vartheta}_{\mathbf{w}}^{(v)}(x) = \frac{\partial^v}{\partial x^v} \widehat{\mathbb{E}}[y_i | x_i = x, \mathbf{w}_i = \mathbf{w}] = \begin{cases} \widehat{\mu}(x) + \mathbf{w}'\widehat{\boldsymbol{\gamma}} & \text{if } v = 0 \\ \widehat{\mu}^{(v)}(x) & \text{if } v \geq 1 \end{cases}.$$

This estimator is obtained by running the linear least squares regression of y_i on $(\widehat{\mathbf{b}}_{p,s}(x_i)', \mathbf{w}_i')$, and then constructing predicted values at (x, \mathbf{w}') for $v = 0$, or predicted values $\widehat{\mu}^{(v)}(x) = \widehat{\mathbf{b}}_{p,s}^{(v)}(x)' \widehat{\boldsymbol{\beta}}$ for $v \geq 1$. The command `binsreg` provides implementation for this case.

- *Logistic/Probit regression*: $\rho(y; \eta) = -y \log \eta - (1 - y) \log(1 - \eta)$ and $\eta(u)$ denotes the (inverse) link function of Logistic or Probit regression. This case corresponds to nonlinear Logistic or Probit binscatter, where the estimator becomes

$$\widehat{\vartheta}_{\mathbf{w}}^{(v)}(x) = \frac{\partial^v}{\partial x^v} \widehat{\mathbb{E}}[y_i | x_i = x, \mathbf{w}_i = \mathbf{w}] = \frac{\partial^v}{\partial x^v} \eta(\widehat{\mu}(x) + \mathbf{w}'\widehat{\boldsymbol{\gamma}}).$$

This estimator is obtained by running the Logit or Probit nonlinear regression of y_i on $(\widehat{\mathbf{b}}_{p,s}(x_i)', \mathbf{w}_i')$, and constructing predicted values at (x, \mathbf{w}') for $v = 0$, or derivatives thereof for $v \geq 1$. The commands `binslogit` and `binsprobit` provide implementation for these cases. These two commands only allow $v = 0$ or 1.

- *Quantile regression*: $\rho(y; \eta) = \ell_{\tau}(y - \eta)$ and $\eta(u) = u$, where $\ell_{\tau}(u)$ denotes the check function associated with the τ -th quantile. This case corresponds to nonlinear, non-differentiable quantile regression binscatter, where the estimator becomes

$$\widehat{\vartheta}_{\mathbf{w}}^{(v)}(x) = \frac{\partial^v}{\partial x^v} \widehat{Q}_{\tau}(y_i | x_i = x, \mathbf{w}_i = \mathbf{w}),$$

where $\widehat{Q}_{\tau}(y_i | x_i = x, \mathbf{w}_i = \mathbf{w}) = \widehat{\mu}(x) + \mathbf{w}'\widehat{\boldsymbol{\gamma}}$ denotes the estimate of the conditional τ -th quantile function of y_i given (x_i, \mathbf{w}_i') . This estimator is obtained by running

the quantile regression of y_i on $(\widehat{\mathbf{b}}_{p,s}(x_i)', \mathbf{w}'_i)$, and constructing predicted values at (x, \mathbf{w}') for $v = 0$, or derivatives thereof for $v \geq 1$. The command `binsqreg` provides implementation for this case.

In practice, the binscatter estimator $\widehat{\vartheta}_{\mathbf{w}}^{(v)}(x)$ needs to be evaluated at some point \mathbf{w} . Typical choices are $\mathbf{w} = \mathbf{0}$, $\mathbf{w} = \bar{\mathbf{w}} = \frac{1}{n} \sum_{i=1}^n \mathbf{w}_i$, or $\mathbf{w} = \text{median}(\mathbf{w}_i)$, with $\mathbf{0}$ denoting a vector of zeros and $\text{median}(\mathbf{w}_i)$ denoting the empirical median of each component in \mathbf{w}_i . For discrete variables in \mathbf{w} it is natural to set those components to some base category (e.g., zero for binary variables), while for continuous variables it may be better to set those components at some other value (e.g., mean or some quantile). The point of evaluation affects the visual and statistical properties of the binscatter, as discussed below.

Canonical Binscatter

Canonical binscatter, as implemented in the packages `binscatter` and `binscatter2`, corresponds to linear least squares regression ($\rho(y; \eta) = (y - \eta)^2$ and $\eta(u) = u$) with $p = s = 0$ and without covariate adjustment (i.e., not including \mathbf{w}_i in (4)). Specifically, in canonical binscatter the basis $\widehat{\mathbf{b}}_{0,0}(x)$ is a J -dimensional vector of orthogonal dummy variables, that is, the j -th component of $\widehat{\mathbf{b}}_{0,0}(x)$ records whether the evaluation point x belongs to the j -th bin in the partition $\widehat{\Delta}$. Therefore, canonical binscatter can be expressed as the collection of J sample averages of the response variable y_i , one for each bin: $\bar{y}_j = \frac{1}{N_j} \sum_{i=1}^n \mathbb{1}_{\widehat{\mathcal{B}}_j}(x_i) y_i$ for $j = 1, 2, \dots, J$. Empirical work employing canonical binscatter typically plots these binned sample averages along with some other estimate(s) of the regression function $\mu_0(x)$.

Covariate-Adjusted Binscatter

Prior work employing binscatter methods, including the packages `binscatter` and `binscatter2`, not only considered exclusively least squares regressions with $p = s = 0$, but also performed covariate adjustment by residualization. To be precise, first the residuals from the linear regressions of y_i on $(1, \mathbf{w}'_i)$ and of x_i on $(1, \mathbf{w}'_i)$ were computed, and then a canonical binscatter was estimated using those residuals. We call this approach residualized canonical binscatter.

CCFF showed that residualized canonical binscatter is very hard to rationalize or justify, and will lead to an inconsistent estimator of $\vartheta_{\mathbf{w}}^{(v)}(x)$ unless very special assumptions hold, even when the statistical model is correctly specified. In contrast, our proposed approach for covariate adjustment (4) is justified via the model (2), and is therefore principled and interpretable. Even when model (2) is misspecified, the approach to covariate adjustment employed by the package `Binsreg` enjoys a natural probability limit interpretation, while the residualization approach does not. See CCFF for more discussion, numerical examples, and technical details.

Main implementation details

The four estimation commands `binsreg`, `binslogit`, `binsprobit`, and `binsqreg` implement, respectively, least squares, Logit, Probit, and quantile regression binscatter estimators for a given choice of partitioning/binning $\hat{\Delta}$. The option `deriv()` is used to set the value of v and the option `at()` is used to set the value of \mathbf{w} in the estimator $\hat{v}_{\mathbf{w}}^{(v)}(x)$. The options `dots(p s)` and `line(p s)` generate “dots” and a “line” tracing out two distinct implementations of $\hat{v}_{\mathbf{w}}^{(v)}(x)$ with the corresponding choices of p and s selected in each case, but using the same values of v and \mathbf{w} . If `dots(T)` (or `line(T)`) is specified, `dots(0 0)` (or `line(0 0)`) is used unless the degree p or smoothness s selection is requested via the option `pselect()` or `sselect()` (see details in the next subsection).

The defaults are to estimate the level of the function ($v = 0$, `deriv(0)`) and evaluate the covariates at the mean ($\mathbf{w}_i = \bar{\mathbf{w}}$, `at(mean)`). Evaluating the covariates at different points can effect a level shift of the plotted point estimates and change the statistical uncertainty, such as confidence bands (see the supplemental appendices of CCFF). The option `at()` allows for the mean, median, and a vector of zeros, the last being useful for dummy variables or fixed effects. For high-dimensional fixed effects, the user-contributed `reghdfe` package can be called with the `absorb` option. In this case the fixed effects are normalized within `reghdfe` and the point of evaluation cannot be set.

For example, when using `binsreg`, the default implementation yields the estimate $\hat{v}_{\bar{\mathbf{w}}}(x) = \hat{\mathbb{E}}[y_i | x_i = x, \mathbf{w}_i = \bar{\mathbf{w}}] = \hat{\mu}(x) + \bar{\mathbf{w}}' \hat{\gamma}$. Thus, `dots(0 0)` leads to “dots” representing sample averages within each bin for the “long” regression with $\mathbf{w}_i = \bar{\mathbf{w}}$. In particular, if \mathbf{w}_i are not included, then the default coincides with Canonical Binscatter (i.e., the same results would be obtained using the packages `binscatter` and `binscatter2` for the same J). The line option is muted by default, and needs to be set explicitly to appear in the resulting plot: for example, the option `line(3 3)` adds a line tracing out $\hat{\mu}^{(v)}(x)$, implemented with $p = 3$ and $s = 3$, a cubic B-spline approximation of $\mu_0^{(v)}(x)$.

The common partitioning/binning used by the four estimation commands across all implementations is set to be quantile-spaced for some choice of J . The option `nbins()` sets J manually (e.g., `nbins(20)` corresponds to $J = 20$ quantile-spaced bins), but if this option is not supplied then the companion command `binsregselect` is used to choose J in a fully data-driven way, as described below. As an alternative, an evenly-spaced or user-specified partitioning/binning can be implemented via the option `binspos()`.

Several other options are available for the four estimation commands, including multi-way fixed effects and multi-way clustering adjustments. Each command has an accompanying help file with complete details.

2.2 Choosing the Number of Bins

From a statistical point of view, J is the main tuning parameter of a binscatter and, as usual for nonparametrics, one must assume $J \rightarrow \infty$ for consistent estimation. To provide an optimal, data-driven choice of J , CCFF developed valid integrated mean squared error (IMSE) approximations for generalized nonlinear binscatter in the context of model (2). These expansions give IMSE-optimal selection of the number bins J , depending on polynomial order p within bins and smoothness level s across bins, the target estimand set by the derivative order v , and the evaluation point of interest \mathbf{w} for covariate adjustment. Specifically, the IMSE-optimal choice of J is

$$J_{\text{IMSE}} = \left\lceil \left(\frac{2(p-v+1)\mathcal{B}_n(p, s, v)}{(1+2v)\mathcal{V}_n(p, s, v)} \right)^{\frac{1}{2p+3}} n^{\frac{1}{2p+3}} \right\rceil,$$

where $\lceil \cdot \rceil$ denotes the ceiling operator, $\mathcal{B}_n(p, s, v)$ and $\mathcal{V}_n(p, s, v)$ represent an approximation to the integrated (squared) bias and variance of $\hat{\vartheta}_{\mathbf{w}}^{(v)}(x)$ respectively, and the three integer choices must respect $p \geq s \geq 0$ and $p \geq v \geq 0$. The constants $\mathcal{B}_n(p, s, v)$ and $\mathcal{V}_n(p, s, v)$ depend on the partitioning scheme and binscatter estimator used. Note that the package `Binsreg` allows for the standard option `vce()` to specify variance-covariance estimation methods, which can also affect the variance constant $\mathcal{V}_n(p, s, v)$.

For simplicity, the command `binsregselect` in the package `Binsreg` implements number of bins selection based on the IMSE expansion for the linear least squares binscatter. For generalized nonlinear binscatter (Logistic, Probit, or quantile regression), the number of bins J given by the command `binsregselect` still has the “correct” rate (the same order as that of the IMSE-optimal one). Thus, confidence bands and testing procedures based on such choices of J and the robust bias correction strategy described below are still valid even in the general nonlinear case.

Both IMSE constants, $\mathcal{B}_n(p, s, v)$ and $\mathcal{V}_n(p, s, v)$, can be estimated consistently using a preliminary choice of J . Thus, our implementation offers two J selectors.

- \hat{J}_{ROT} : implements a rule-of-thumb (ROT) approximation for the constants $\mathcal{B}_n(p, s, v)$ and $\mathcal{V}_n(p, s, v)$, employing a trimmed-from-below Gaussian reference model for the density of x_i , and global polynomial approximations for the other two unknown features needed, $\mu_0^{(v)}(x)$ and $\mathbb{V}[y_i|x_i = x, \mathbf{w}_i = \mathbf{w}]$. This J selector employs the correct rate but an inconsistent constant approximation.
- \hat{J}_{DPI} : implements a direct-plug-in (DPI) approximation for the constants $\mathcal{B}_n(p, s, v)$ and $\mathcal{V}_n(p, s, v)$, based on the desired binscatter, set by the choices p and s , and employing a preliminary J . If a preliminary J is not provided by the user, then $J = \max\{\hat{J}_{\text{ROT}}, \lceil (\frac{2(p-v+1)}{1+2v}n)^{\frac{1}{2p+3}} \rceil\}$ is used for DPI implementation. This J selector employs the correct rate as well as consistent estimators of the appropriate constants. Default implementation uses \hat{J}_{DPI} whenever J is not specified.

Implementing a binscatter with $J = J_{\text{IMSE}}$ is optimal from a statistical point of view (for valid estimation, testing, and uncertainty quantification), but sometimes a fixed,

user-chosen number of bins, denoted by $J = \mathbf{J}$, may yield a more visually appealing binscatter (albeit with the caveat that the implied estimator may be inaccurate). Many applications in the past used round numbers such as 10, 20, 50, or 100 bins. Further, a fixed $J = \mathbf{J}$ can also be directly interpretable as a discretized version of x . For instance, setting $J = \mathbf{J} = 10$ yields a binscatter that allows for comparison across deciles of x , for example comparing those in the top decile of earnings to those at the bottom. Setting $J = 100$ can be used to compare across percentiles, and is commonly used when studying ranks.

To balance the potential for an appealing visualization based on a fixed $J = \mathbf{J}$ with the desire for statistical validity, CCFF developed a novel method for selecting the polynomial order p (and along with it, the smoothness s) as a function of the fixed \mathbf{J} , as opposed to $J_{\text{IMSE}} = J_{\text{IMSE}}(p, s, v)$ which does the reverse. Specifically, CCFF proposed to look for the values p^* and s^* (in a pre-specified range) such that $J_{\text{IMSE}}(p^*, s^*, v)$ is approximately equal to the researcher's chosen \mathbf{J} . That is, finding the p and s for which the chosen \mathbf{J} would be IMSE-optimal. Although $p \rightarrow \infty$ is not allowed in the theory of CCFF, and so this choice is somewhat ad-hoc, it may effectively reduce the bias of the binscatter.

Main implementation details

Unless J is manually specified, all commands in the `Binsreg` package employ the command `binsregselect` to implement ROT and DPI data-driven, IMSE-optimal selection of J for all possible choices of $p \geq v, s \geq 0$, and for both quantile-spaced or evenly-spaced partitioning/binning. For DPI implementation, the user can provide the initialization value of J via the option `nbinsrot()` or, if not provided, then \hat{J}_{ROT} is used.

As discussed above, instead of selecting the number of bins J , an alternative strategy is setting a fixed value for J and implementing (ROT or DPI) data-driven, IMSE-optimal selection of the degree of polynomial p and/or the number of smoothness constraints s . The command `binsregselect` implements this selection procedure if (i) the number of bins J is supplied via the option `nbins()` and (ii) a range for searching for the optimal p or s is supplied via the option `pselect()` or `sselect()`.

Several other options are available for the command `binsregselect`, including the possibility of generating an output file with the IMSE-optimal partitioning/binning structure selected and the corresponding grid of evaluation points, which can be used by the other six companion commands for plotting, simulation, testing, and other calculations.

2.3 Confidence Intervals

Both confidence intervals and confidence bands for the unknown function $\vartheta_{\mathbf{w}}^{(v)}(x)$ are constructed employing the same type of Studentized t -statistic:

$$T_p(x) = \frac{\widehat{\vartheta}_{\mathbf{w}}^{(v)}(x) - \vartheta_{\mathbf{w}}^{(v)}(x)}{\sqrt{\widehat{\Omega}(x)/n}}, \quad 0 \leq v, s \leq p,$$

where the binscatter variance estimator is of the usual “sandwich” form

$$\begin{aligned} \widehat{\Omega}(x) &= \widehat{\mathbf{b}}_{p,s}^{(v)}(x)' \widehat{\mathbf{Q}}^{-1} \widehat{\Sigma} \widehat{\mathbf{Q}}^{-1} \widehat{\mathbf{b}}_{p,s}^{(v)}(x) (\eta^{(1)}(\widehat{\mu}(x) + \mathbf{w}'\widehat{\gamma}))^2, \\ \widehat{\mathbf{Q}} &= \frac{1}{n} \sum_{i=1}^n \widehat{\mathbf{b}}_{p,s}(x_i) \widehat{\mathbf{b}}_{p,s}(x_i)' \widehat{\Psi}_{i,1} \widehat{\eta}_{i,1}^2, \\ \widehat{\Sigma} &= \frac{1}{n} \sum_{i=1}^n \widehat{\mathbf{b}}_{p,s}(x_i) \widehat{\mathbf{b}}_{p,s}(x_i)' \widehat{\eta}_{i,1}^2 \psi(y_i, \widehat{\eta}_{i,0})^2 \end{aligned}$$

with $\widehat{\Psi}_{i,1}$ a consistent estimator of $\Psi_{i,1} = \frac{\partial}{\partial \eta} \mathbb{E}[\psi(y_i, \eta) | x_i, \mathbf{w}_i] \Big|_{\eta=\eta_{i,0}}$, $\widehat{\eta}_{i,v}$ a consistent estimator of $\eta_{i,v} = \eta^{(v)}(\mu_0(x_i) + \mathbf{w}_i' \gamma_0)$ for $v = 0, 1$, and $\psi(y, u)$ the (weak) derivative of $\rho(\cdot; u)$ with respect to u . See CCFF for details and omitted formulas. In practice, these estimators are implemented using the base commands from the statistical software.

CCFF showed that $T_p(x) \rightarrow_d \mathbf{N}(0, 1)$ pointwise in x , that is, for each evaluation point x on the support of x_i , provided the misspecification error introduced by binscatter is removed from the distributional approximation. Such a result justifies asymptotically valid confidence intervals for $\vartheta_{\mathbf{w}}^{(v)}(x)$, pointwise in x , after bias correction. Specifically, for each x , the $(1 - \alpha)\%$ confidence interval takes the form:

$$\widehat{I}_p(x) = \left[\widehat{\vartheta}_{\mathbf{w}}^{(v)}(x) \pm \Phi^{-1}(1 - \alpha/2) \cdot \sqrt{\widehat{\Omega}(x)/n} \right], \quad 0 \leq v, s \leq p,$$

where $\Phi(u)$ denotes the distribution function of a standard normal random variable (e.g., $\Phi^{-1}(1 - 0.05/2) \approx 1.96$ for a 95% Gaussian confidence intervals), and provided the choice of J is such that the misspecification error can be ignored.

However, employing an IMSE-optimal binscatter (i.e., setting $J = J_{\text{IMSE}}$ for the selected polynomial order p) introduces a first-order misspecification error leading to invalidity of these confidence intervals, and hence cannot be directly used to form the confidence intervals $\widehat{I}_p(x)$ in general. To address this problem, we rely on a simple application of robust bias correction (Calonico et al. 2014, 2018; Cattaneo et al. 2020b; Calonico et al. 2022) to form valid confidence intervals based on IMSE-optimal binscatter, that is, without altering the partitioning scheme $\widehat{\Delta}$ used.

Our implementation employs robust bias-corrected binscatter confidence intervals as follows. First, for a given choice of p , select the number of bins in $\widehat{\Delta}$ according to $J = J_{\text{IMSE}}$, which gives an IMSE-optimal binscatter (point estimator). Then, employ

the confidence interval $\widehat{I}_{p+q}(x)$ with $q \geq 1$, which gives a valid confidence interval: $\mathbb{P}[\vartheta_{\mathbf{w}}^{(v)}(x) \in \widehat{I}_{p+q}(x)] \rightarrow 1 - \alpha$, for all x .

The same strategy is applied when the degree/smoothness selection described previously is implemented. First, for a fixed choice of J , select the “optimal” degree of polynomial p (i.e., the resulting J_{IMSE} is the closest to the chosen J). Then, increase the degree of polynomial to construct the confidence intervals $\widehat{I}_{p+q}(x)$. This idea is also employed in the construction of confidence bands and hypothesis testing procedures and can be fully implemented using the options `pselect()` and `sselect()`. Since the degree/smoothness selection is *not* the default in the package `Binsreg`, our discussion focuses on robust bias correction based on IMSE-optimal J selection.

Main implementation details

The four estimation commands `binsreg`, `binslogit`, `binsprobit`, and `binsqreg` implement confidence intervals, and report them as part of the final binned scatter plot. Specifically, the option `ci(p s)` estimates confidence intervals with the corresponding choices of p and s selected, and plots them as vertical segments along the support of x_i . If `ci(T)` is specified, `ci(1 1)` is used unless the degree p or smoothness s selection is requested via the option `pselect()` or `sselect()` as described before. The confidence interval option is muted by default, and needs to be set explicitly to appear in the resulting plot. The implementation is done over a grid of evaluation points, which can be modified via the option `cigrd()`, and the desired level is set by the option `level()`. Notice that `dots(p s)`, `lines(p s)`, and `ci(p s)` may all take different choices of p and s , which allows for robust bias correction implementation of the confidence intervals and permits incorporating different levels of smoothness restrictions.

2.4 Confidence Bands

In many empirical applications of `binscatter`, the goal is to conduct inference about the entire function $\vartheta_{\mathbf{w}}^{(v)}(x)$, simultaneously, that is, uniformly over all x values of the support of x_i . This goal is fundamentally different from pointwise inference. A leading example of uniform inference is reporting confidence bands for $\vartheta_{\mathbf{w}}(x)$ and its derivatives, which are different from (pointwise) confidence intervals. The package `Binsreg` offers asymptotically valid constructions of both confidence intervals, as discussed above, and confidence bands, which can be implemented with the same choices of (p, s) used to construct $\widehat{\vartheta}_{\mathbf{w}}^{(v)}(x)$ or different ones.

Following the theoretical work in CCFE, for a choice of p and partition/binning of size J , the $(1 - \alpha)\%$ confidence band for $\vartheta_{\mathbf{w}}^{(v)}(x)$ is:

$$\widehat{I}_p(\cdot) = \left[\widehat{\vartheta}_{\mathbf{w}}^{(v)}(\cdot) \pm \mathbf{c} \cdot \sqrt{\widehat{\Omega}(\cdot)/n} \right], \quad 0 \leq v, s \leq p,$$

where the quantile value c is now approximated via simulations using

$$c = \inf \left\{ c \in \mathbb{R}_+ : \mathbb{P} \left[\sup_x |\widehat{Z}_p(x)| \leq c \mid \mathbf{D} \right] \geq 1 - \alpha \right\},$$

with $\mathbf{D} = ((y_i, x_i, \mathbf{w}'_i) : 1 \leq i \leq n)$ denoting the original data,

$$\widehat{Z}_p(x) = \frac{\widehat{\mathbf{b}}_{p,s}^{(v)}(x)' \widehat{\mathbf{Q}}^{-1} \widehat{\boldsymbol{\Sigma}}^{-1/2}}{\sqrt{\widehat{\Omega}(x)/n}} \mathbf{N}_K^*, \quad K = (p+1)J - (J-1)s, \quad 0 \leq v, s \leq p,$$

and $\mathbf{N}_K^* \sim \mathbf{N}(\mathbf{0}, \mathbf{I})$ being a K -dimensional standard normal random vector independent of the data \mathbf{D} . The distribution of $\sup_x |T_p(x)|$, which is unknown, is approximated by that of $\sup_x |\widehat{Z}_p(x)|$ conditional on the data \mathbf{D} , which can be simulated by taking repeated samples from \mathbf{N}_K^* and recomputing the supremum each time. In other words, the quantiles used to construct confidence bands can be approximated by resampling from the standard normal random vector \mathbf{N}_K^* , keeping fixed the data \mathbf{D} (and hence all quantities depending on it). See CCFF for more details.

A confidence band covers the entire function, $\vartheta_{\mathbf{w}}^{(v)}(x)$, $(1-\alpha)\%$ of the time in repeated sampling, whenever the misspecification error can be ignored. As before, we recommend employing robust bias correction to remove misspecification error introduced by `binscatter`, that is, following the same logic discussed above for the case of confidence interval construction. To be more precise, first p is chosen, along with s and v , and the optimal partitioning/binning is selected according to $J = J_{\text{IMSE}}$. Then, valid confidence bands are constructed using $\widehat{I}_{p+q}(x)$ with $q \geq 1$: $\mathbb{P}[\vartheta_{\mathbf{w}}^{(v)}(x) \in \widehat{I}_{p+q}(x), \text{ for all } x] \rightarrow 1 - \alpha$.

Moreover, the visual appearance of the confidence band will be impacted by the chosen evaluation point \mathbf{w} , and the researcher needs to be careful when using the band as visual aids in parametric specification testing. See CCFF, and in particular Section SA-1.2 of Cattaneo et al. (2024b).

Main implementation details

The four estimation commands `binsreg`, `binslogit`, `binsprobit`, and `binsqreg` implement confidence bands, and report them as part of the final binned scatter plot. The option `cb(p s)` estimates an asymptotically valid confidence band with the corresponding choices of p and s selected, and plots it as a shaded region along the support of x_i . If `cb(T)` is specified, `cb(1 1)` is used unless the degree p or smoothness s selection is requested via the option `pselect()` or `sselect()` as described before. The confidence band option is muted by default, and needs to be set explicitly to appear in the resulting plot. The implementation is done over a grid of evaluation points, which can be modified via the option `cbgrid()`, and the desired level is set by the option `level()`. The options `dots(p s)`, `lines(p s)`, `ci(p s)`, and `cb(p s)` can all take different choices of p and s , which allows for robust bias correction implementations, as well as many other practically relevant possibilities.

2.5 Parametric Specification Testing

In addition to implementing binscatter and producing binned scatter plots, with both point estimation and uncertainty visualization, the package Binsreg also allows for formal testing of substantive hypotheses. The command `binstest` implements two types of substantive hypothesis tests about $\vartheta_{\mathbf{w}}^{(v)}(x)$: (i) parametric specification testing and (ii) nonparametric shape restriction testing. This subsection discusses the former, while the next subsection discusses the latter.

For a choice of (p, s, v) , and partitioning/binning scheme of size J , the implemented parametric specification testing approach contrasts a (nonparametric) binscatter approximation $\widehat{\vartheta}_{\mathbf{w}}^{(v)}(x)$ of $\vartheta_{\mathbf{w}}^{(v)}(x)$ with a hypothesized parametric specification of the form $\vartheta_{\mathbf{w}}(x) = \eta(M_{\mathbf{w}}(x; \boldsymbol{\theta}, \boldsymbol{\gamma}_0))$ where $M_{\mathbf{w}}(x; \boldsymbol{\theta}, \boldsymbol{\gamma}_0) = m(x; \boldsymbol{\theta}) + \mathbf{w}'\boldsymbol{\gamma}_0$ for some $m(\cdot)$ known up to a finite parameter $\boldsymbol{\theta}$, which can be estimated using the available data. Formally, the null and alternative hypotheses are, respectively,

$$\begin{aligned} \dot{H}_0 &: \sup_x \left| \vartheta_{\mathbf{w}}^{(v)}(x) - \frac{\partial^v}{\partial x^v} \eta(M_{\mathbf{w}}(x; \boldsymbol{\theta}, \boldsymbol{\gamma}_0)) \right| = 0, \quad \text{for some } \boldsymbol{\theta}, \quad \text{vs.} \\ \dot{H}_A &: \sup_x \left| \vartheta_{\mathbf{w}}^{(v)}(x) - \frac{\partial^v}{\partial x^v} \eta(M_{\mathbf{w}}(x; \boldsymbol{\theta}, \boldsymbol{\gamma}_0)) \right| > 0, \quad \text{for all } \boldsymbol{\theta}, \end{aligned}$$

for a choice of v .

For example, excluding additional covariates \mathbf{w}_i , $\widehat{\mu}(x)$ is compared to $\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i$ in order to assess whether there is a relationship between y_i and x_i or, more formally, whether $\mu_0(x)$ is a constant function. Similarly, it is possible to formally test for a linear, quadratic, or even nonlinear parametric relationship $\mu_0(x) = m(x, \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ would be estimated from the data under the null hypothesis, that is, assuming that the postulated relationship is indeed correct.

However, when additional covariates \mathbf{w}_i are included, CCFF showed that the special case of the test \dot{H}_0 vs. \dot{H}_A with $v = 0$ could give conclusions that are sensitive to the user-selected point of evaluation \mathbf{w} , which implies that the common practice of visually examining a binned scatter plot compared to a parametric specific could be misleading. See Section SA-1.2 of Cattaneo et al. (2024b) for further details.

To avoid this issue, and motivated by the fact that the central point of binscatter is to study how y_i relates to x_i , controlling for \mathbf{w}_i , we advocate reformulating the hypothesis as pertaining to the *derivative* of $\mu_0(x)$, instead of the level. For example, to test if $\mu_0(x)$ is linear, one can test if it has a constant first derivative, i.e., $H_0 : \sup_x |\mu_0^{(1)}(x) - a| = 0$ for some a , vs. $H_A : \sup_x |\mu_0^{(1)}(x) - a| > 0$. This can be viewed as a special case of the test \dot{H}_0 vs. \dot{H}_A above with $v = 1$ and the (inverse) link $\eta(\cdot)$ suppressed (i.e., apply the test to the linear index $\mu_0(x) + \mathbf{w}'\boldsymbol{\gamma}_0$ directly).

Formally, the command `binstest` employs the test statistic

$$\dot{T}_p(x) = \frac{\widehat{\vartheta}_{\mathbf{w}}^{(v)}(x) - \frac{\partial^v}{\partial x^v} \eta(M_{\mathbf{w}}(x; \widetilde{\boldsymbol{\theta}}, \widetilde{\boldsymbol{\gamma}}))}{\sqrt{\widehat{\Omega}(x)/n}}, \quad 0 \leq v, s \leq p,$$

where $(\tilde{\boldsymbol{\theta}}', \tilde{\boldsymbol{\gamma}}')$ are consistent estimates of $(\boldsymbol{\theta}', \boldsymbol{\gamma}_0')$ under the null hypothesis (correct parametric specification), and are “well behaved” under the alternative hypothesis (parametric misspecification). The researcher needs to carefully choose a proper v and decide if the test should be applied to $\eta(\mu_0(x) + \mathbf{w}'\boldsymbol{\gamma}_0)$ or $\mu_0(x) + \mathbf{w}'\boldsymbol{\gamma}_0$, following the discussion above. Then, a parametric specification hypothesis testing procedure is

$$\text{Reject } \dot{H}_0 \quad \text{if and only if} \quad \sup_x |\dot{T}_p(x)| \geq \mathbf{c}, \quad (5)$$

where $\mathbf{c} = \inf\{c \in \mathbb{R}_+ : \mathbb{P}[\sup_x |\hat{Z}_p(x)| \leq c \mid \mathbf{D}] \geq 1 - \alpha\}$ is again computed by simulation from a standard normal random vector, conditional on the data \mathbf{D} , as in the case of confidence bands already discussed. This testing procedure is an asymptotically valid $\alpha\%$ -level test if the misspecification error is removed from the test statistic $\dot{T}_p(x)$.

The command `binstest` employs robust bias correction by default: first p and s are chosen, and the partitioning/binning scheme is selected by setting $J = J_{\text{IMSE}}$ for these choices. Then, using this partitioning scheme, the testing procedure (5) is implemented with the choice $p + q$ instead of p , with $q \geq 1$. CCFF showed that, under regularity conditions, the resulting parametric specification testing approach controls Type I error with non-trivial power: for given p , $0 \leq v, s \leq p$, and $J = J_{\text{IMSE}}$,

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[\sup_x |\dot{T}_{p+q}(x)| > \mathbf{c} \right] = \alpha, \quad \text{under } \dot{H}_0,$$

and

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[\sup_x |\dot{T}_{p+q}(x)| > \mathbf{c} \right] = 1, \quad \text{under } \dot{H}_A,$$

where $q \geq 1$. This testing approach formalizes the intuitive idea that if the confidence band for $\vartheta_{\mathbf{w}}^{(v)}(x)$ does not contain the hypothesized parametric fit entirely, then the parametric fit is incompatible with the data, i.e., the null should be rejected.

Main implementation details

The command `binstest` implements parametric specification testing in two ways. First, polynomial regression (parametric) specification testing is implemented directly via the option `testmodelpoly(P)`, where the null hypothesis is $m(x, \boldsymbol{\theta}) = \theta_0 + x\theta_1 + \dots + x^P\theta_P$ and $\boldsymbol{\theta} = (\theta_0, \theta_1, \dots, \theta_P)'$ is estimated by least squares regression. For other parametrizations of $m(x, \boldsymbol{\theta})$, the command takes as input an auxiliary array/database (`dta` in Stata, or data frame in Python and R) via the option `testmodelparfit(filename)` containing the following columns/variables: grid of evaluation points in one column, and fitted values $\eta^{(v)}(m(x, \tilde{\boldsymbol{\theta}}) + \mathbf{w}'\tilde{\boldsymbol{\gamma}})$ (over the evaluation grid) for each parametric model considered in other columns/variables. The ordering of these variables is arbitrary, but they have to follow a naming rule: the evaluation grid has the same name as the independent variable x_i , and the names of other variables storing fitted values take the form `binsreg_fit*`.

The `binscatter` (nonparametric) estimate used to construct the testing procedure is set by the options `testmodel(p s)` and `deriv(v)`, and the partitioning/binning scheme

selected. If `testmodel(T)` or `testmodel()` is supplied, the default `testmodel(1 1)` is used unless the degree p and smoothness s selection is requested via the options `pselect()` and `sselect()` as described before. The option `nolink` can be used to specify if the test should be applied to the linear index directly.

2.6 Nonparametric Shape Testing

The second type of hypothesis tests implemented by the command `binstest` concern nonparametric testing of shape restrictions. For a choice of v , the null and alternative hypotheses of these testing problems are:

$$\ddot{H}_0 : \sup_x \vartheta_{\mathbf{w}}^{(v)}(x) \leq 0, \quad vs. \quad \ddot{H}_A : \sup_x \vartheta_{\mathbf{w}}^{(v)}(x) > 0,$$

that is, one-sided testing problem to the left. For example, negativity, monotonicity, and concavity of $\vartheta_{\mathbf{w}}(x)$ correspond to $\vartheta_{\mathbf{w}}(x) \leq 0$, $\vartheta_{\mathbf{w}}^{(1)}(x) \leq 0$, and $\vartheta_{\mathbf{w}}^{(2)}(x) \leq 0$, respectively. Of course, the analogous testing problem to the right is also implemented, but not discussed here to avoid unnecessary repetition.

The relevant Studentized test statistic for this class of testing problems is:

$$\ddot{T}_p(x) = \frac{\widehat{\vartheta}_{\mathbf{w}}^{(v)}(x)}{\sqrt{\widehat{\Omega}(x)/n}}, \quad 0 \leq v, s \leq p.$$

Then, the testing procedure is:

$$\text{Reject } \ddot{H}_0 \quad \text{if and only if} \quad \sup_x \ddot{T}_p(x) \geq \mathfrak{c}, \quad (6)$$

with $\mathfrak{c} = \inf\{c \in \mathbb{R}_+ : \mathbb{P}[\sup_x \widehat{Z}_p(x) \leq c \mid \mathbf{D}] \geq 1 - \alpha\}$. Misspecification errors of `binscatter` need to be taken into account in order to control Type I error. CCFP showed that for given p , $0 \leq v, s \leq p$, and $J = J_{\text{IMSE}}$ accordingly, then

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[\sup_x \ddot{T}_{p+q}(x) > \mathfrak{c} \right] \leq \alpha, \quad \text{under } \ddot{H}_0,$$

and

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[\sup_x \ddot{T}_{p+q}(x) > \mathfrak{c} \right] = 1, \quad \text{under } \ddot{H}_A,$$

for any $q \geq 1$, that is, using a robust bias correction approach. These results imply that the testing procedure (6) is an asymptotically valid hypothesis test provided that it is implemented with the choice $q \geq 1$ after the IMSE-optimal partitioning/binning scheme for `binscatter` of order p is selected.

Main implementation details

The command `binstest` implements one-sided and two-sided nonparametric shape restriction testing as follows. Option `testshapel(a)` implements one-sided testing

to the left: $\check{H}_0 : \sup_x \vartheta_{\mathbf{w}}^{(v)}(x) \leq \mathbf{a}$. Option `testshaper(a)` for one-sided to the right: $\check{H}_0 : \inf_x \vartheta_{\mathbf{w}}^{(v)}(x) \geq \mathbf{a}$. Option `testshape2(a)` for two-sided testing: $\check{H}_0 : \sup_x |\vartheta_{\mathbf{w}}^{(v)}(x) - \mathbf{a}| = 0$. The constant \mathbf{a} needs to be specified by the user.

The `binscatter` (nonparametric) estimate used to construct the testing procedure is set by the options `testshape(p s)` and `deriv(v)`, and the chosen partitioning/binning scheme. If `testshape(T)` or `testshape()` is supplied, `testshape(1 1)` is used unless the degree p and smoothness s selection is requested via the options `pselect()` and `sselect()` as described before.

2.7 Multi-Sample Estimation and Testing

The package `Binsreg` also allows for comparisons of mean, quantile, and other regression functions across different groups (or treatment arms), which can be useful for estimation and inference of treatment effects that are heterogeneous in x_i , possibly after controlling for \mathbf{w}_i . For each subsample defined by a group indicator variable, the parameter of interest can be defined as $\vartheta_{\mathbf{w},\ell}^{(v)}(x)$, which corresponds to the parameter in (1) for specific subsample $\ell = 0, 1, 2, \dots, L$.

For example, assuming that two sub-samples of the same size n are available ($L = 1$), one being a control group and the other a treatment group, all the methods discussed above can be applied to each subsample. Furthermore, the null hypothesis of no heterogeneous treatment effect is: $H_0^\Delta : \vartheta_{\mathbf{w},0}^{(v)}(x) = \vartheta_{\mathbf{w},1}^{(v)}(x)$ for all $x \in \mathcal{X}$, which captures the idea of no (heterogeneous in x_i) treatment effect across the two groups. A natural test statistic is:

$$T_p^\Delta(x) = \frac{\widehat{\vartheta}_{\mathbf{w},1}^{(v)}(x) - \widehat{\vartheta}_{\mathbf{w},0}^{(v)}(x)}{\sqrt{\widehat{\Omega}_1(x)/n + \widehat{\Omega}_0(x)/n}}, \quad 0 \leq v, s \leq p,$$

which compares the pairwise difference between the two groups, where $\widehat{\Omega}_\ell(x)$ is the variance estimator (of $\widehat{\vartheta}_{\mathbf{w},\ell}^{(v)}(x)$) for the subsample $\ell = 0, 1$. The testing procedure is:

$$\text{Reject } H_0^\Delta \quad \text{if and only if} \quad \sup_x |T_p^\Delta(x)| \geq \mathbf{c}, \quad (7)$$

with the critical value obtained as before via Gaussian approximations (resampling from a normal random vector conditional on the data). As discussed before, in practice the robust bias-corrected test statistics $T_{p+q}^\Delta(x)$ is used to eliminate misspecification bias and obtain a valid hypothesis testing procedure.

All the ideas and results above also apply to pairwise comparisons across multi-samples. In particular, estimation, uncertainty quantification and hypothesis testing can be conducted for each subsample at the time, and then hypothesis testing for pairwise comparisons can also be implemented following the results above. `CCFF` provided all the necessary theoretical background. Importantly, concerns regarding the choice of evaluation point \mathbf{w} also apply to the multi-sample testing problems: researchers need to be careful when implementing the tests and interpreting the results.

Main implementation details

Estimation and uncertainty quantification across subsamples is done using the estimation commands (`binsreg`, `binslogit`, `binsprobit`, and `binsqreg`) via the option `by()`. In addition, the command `binspwc` implements formal hypothesis testing for pairwise comparisons for the null hypothesis H_0^Δ (and analogous one-sided problems).

2.8 Extensions and Other Implementation Details

The package `Binsreg` is implemented using the base commands in the statistical software. For example, in `Stata`, `binsreg` relies on `regress` (or `reghdfe` if that option is selected), `binslogit` relies on `logit`, `binsprobit` relies on `probit`, and `binsqreg` relies on `qreg` (or `bsqreg` if bootstrapping-based standard error is selected). Furthermore, the testing commands (`binstest` and `binspwc`) also employ base commands whenever possible. This approach may sacrifice some speed of implementation, but improves substantially in terms of stability and replicability. Importantly, essentially most options available in the base commands are available in the package `Binsreg`.

This section reviews some specific extensions and other numerical issues of the package `Binsreg` and discusses related choices made for implementation, all of which can affect speed and/or robustness of the package.

Other metrics

All the results presented above employ the uniform norm, that is, focus on the the largest deviation on the support of a function. See, for example, \dot{H}_0 , \dot{H}_0 , and H_0^Δ . Our results also apply to other metrics, such as the L_p metric. In such a case, the null hypotheses, the corresponding statistics and simulated critical values will focus on an integral computation of the function of interest. For example, \dot{H}_0 is replaced by

$$\int \left| \vartheta_{\mathbf{w}}^{(v)}(x) - \frac{\partial^v}{\partial x^v} \eta(M_{\mathbf{w}}(x; \boldsymbol{\theta}, \boldsymbol{\gamma}_0)) \right|^p dx = 0, \quad \text{for some } \boldsymbol{\theta},$$

where \mathbf{p} is some positive integer no less than 1 (typically $\mathbf{p} = 2$ for squared deviations), and the corresponding critical value simulation takes the form $\mathbf{c} = \inf\{c \in \mathbb{R}_+ : \mathbb{P}[\int |\hat{Z}_{\mathbf{p}}(x)|^p dx \leq c \mid \mathbf{D}] \geq 1 - \alpha\}$. Analogous modifications are done for other hypothesis tests. Note that by construction, the L_p metric measures the integrated absolute deviation, and thus should be applied to two-sided tests only. The choice of metric is implemented in each testing command via the option `lp()`.

Mass points and minimum effective sample size

The package `Binsreg` incorporates specific implementation decisions to deal with mass points in the distribution of the independent variable x_i . The number of distinct values of x_i , denoted by N , is taken as the effective sample size as opposed to the total number of observations n . If x_i is continuously distributed, then $N = n$. However, in

many applications, N can be substantially smaller than n , and this affects some of the implementations in the package.

First, assume that J is set by the user (via the option `nbins(J)`). Then, given the choice J , the commands `binsreg`, `binslogit`, `binsprobit`, `binsqreg`, `binstest`, and `binspwc` perform a degrees of freedom check to decide whether the x_i data exhibit enough variation. Specifically, given p and s set by the option `dots(p s)` or `bins(p s)`, these commands check whether $N > N_2 + (p + 1)J - (J - 1)s$ with $N_2 = 30$ by default. If this check is not passed, then the package `Binsreg` regards the data as having “too little” variation in x_i , and turns off all nonparametric estimation and inference results based on large sample approximations. Thus, in this extreme case, the command `binsreg` (or `binslogit`, `binsprobit`, `binsqreg`) only allows for `dots(0 0)`, `ci(0 0)`, and `polyreg(P)` for any $P + 1 < N$, while the command `binstest` (or `binspwc`) does not return any results and issues a warning message instead.

If, on the other hand, for given J , the numerical check $N > N_2 + (p + 1)J - (J - 1)s$ is passed, then all nonparametric methods implemented by the commands `binsreg`, `binslogit`, `binsprobit`, `binsqreg`, `binstest`, and `binspwc` become available. However, before implementing each method (`dots(p s)`, `lines(p s)`, `ci(p s)`, `cb(p s)`, `polyreg(P)`, and the hypothesis testing procedures), a degrees of freedom check is performed in each individual case. Specifically, each nonparametric procedure is implemented only if $N > N_2 + (p + 1)J - (J - 1)s$, where recall that p and s may change from one procedure to the next.

Second, as discussed above, whenever J is not set by the user via the option `nbins()`, the command `binsregselect` is employed to select J in a data-driven way, provided there is enough variation in x_i . To determine the latter, an initial degrees of freedom check is performed to assess whether J selection is possible or, alternatively, if the unique values of x_i should be used as bins directly. Specifically, if $N > N_1 + p + 1$, with p set by the option `dots(p s)` (or `bins(p s)`) and $N_1 = 20$ by default, then the data are deemed appropriate for ROT selection of J via the command `binsregselect`, and hence \hat{J}_{ROT} is implemented. If, in addition, $N > N_2 + (p + 1)\hat{J}_{\text{ROT}} - (\hat{J}_{\text{ROT}} - 1)s$, then \hat{J}_{DPI} is also implemented whenever requested. Furthermore, the command `binsregselect` employs the following alternative formula for J selection:

$$J_{\text{IMSE}} = \left\lceil \left(\frac{2(p - v + 1)\mathcal{B}_n(p, s, v)}{(1 + 2v)\mathcal{V}_n(p, s, v)} \right)^{\frac{1}{2p+3}} N^{\frac{1}{2p+3}} \right\rceil,$$

with a slightly different constant $\mathcal{V}_n(p, s, v)$, taking into account the frequency of data at each mass point. All other estimators in the package `Binsreg`, including bias and standard error estimators, automatically adapt to the presence of mass points. Once the final J is estimated, the degrees of freedom checks discussed in the previous paragraphs are performed based on this choice.

If J is not set by the user and $N \leq N_1 + p + 1$, so that not even ROT estimation of J is possible, then N is taken as “too small.” In this extreme case, the package `Binsreg` sets $J = N$ and constructs a partitioning/binning structure with each bin containing one unique value of x_i . In other words, the support of the raw data is taken as the

binning structure itself. In this extreme case, the follow-up degrees of freedom checks based on the formula $N > N_2 + (p + 1)J - (J - 1)s$ fail by construction, and hence the nonparametric methods are turned off as explained above.

Finally, the specific numerical checks and corresponding adjustments mentioned in this subsection can be modified or omitted. This is controlled by two main options: `dfcheck()` and `masspoints()`, respectively. First, the default cutoff points N_1 and N_2 , corresponding to the degrees of freedom checks for parametric global polynomial regression and nonparametric `binscatter`, respectively, can be modified using the option `dfcheck(N1 N2)`. Second, the option `masspoints()` controls how the package `Binsreg` handles the presence of mass points (i.e., repeated values) in x_i . Specifically, setting `masspoints(noadjust)` omits mass point checks and the corresponding effective sample size adjustments, that is, it sets $N = n$ and ignores the presence of mass points in x_i (if any). Setting `masspoints(nolocalcheck)` omits within-bin mass point checks, but still performs global mass point checks and adjustments. The option `masspoints(off)` corresponds to setting both `masspoints(noadjust)` and `masspoints(nolocalcheck)` simultaneously. Finally, setting `masspoints(veryfew)` forces the package to proceed as if N is so small that all checks are failed, thereby treating x_i as if it has very few distinct values.

Clustered data and minimum effective sample size

As discussed in CCFF, the main methodological results for `binscatter` can be extended to accommodate clustered data. All three commands in the package `Binsreg` allow for clustered data via the option `vce()`. In this case, the number of clusters G is taken as the effective sample size, assuming $N = n$ (see below for the other case). The only substantive change occurs in the command `binsregselect`, which now employs the following alternative formula for J selection:

$$J_{\text{IMSE}} = \left\lceil \left(\frac{2(p - v + 1)\mathcal{B}_n(p, s, v)}{(1 + 2v)\mathcal{V}_n(p, s, v)} \right)^{\frac{1}{2p+3}} G^{\frac{1}{2p+3}} \right\rceil,$$

with a variance constant $\mathcal{V}_n(p, s, v)$ accounting for the clustered structure of the data. Accordingly, cluster-robust variance estimators are used in this case.

Minimum effective sample size

The package `Binsreg` requires some minimal variation in x_i in order to successfully implement nonparametric methods based on large sample approximations. The minimal variation is captured by the number of distinct values on the support of x_i , denoted by N , and the number of clusters, denoted by G . Thus, all three commands in the package perform degrees of freedom numerical checks using $\min\{n, N, G\}$ as the general definition of effective sample size, and proceeding as explained above for the case of mass points in the distribution of x_i .

2.9 Increasing Speed of Execution

The package `Binsreg` offers a large array of options and methods, many of which involve nonlinear estimation and/or simulations, thereby slowing down its speed of execution. Furthermore, in order to improve the stability and replicability of the package, it implements several robustness checks that may further decrease execution speed, particularly in settings with ultra-large datasets. There are, however, several options and approaches that could be used to improve the speed of execution of the package `Binsreg`.

1. *Sorted data.* The core implementations of the package `Binsreg` employ several algorithms and procedures that require sorted data along the x dimension. If the provided data is not sorted, then the package begins by sorting the data, which slows down the execution (particularly in large datasets).
 - Speed improvement: provide sorted data in x , which may substantially increase execution speed (particularly in ultra-large datasets).
2. *Data Distribution.* The methods implemented in the package `Binsreg` were developed for continuously distributed data with “enough” variation (e.g., enough degrees of freedom within and across bins, appropriate rank conditions for Gram matrices, etc.). Because empirical work may involve data with mass points and/or other irregularities that can make the default methods fail, the package `Binsreg` implements a series of robustness checks before execution (see above for details).
 - Speed improvement: use option `masspoints(off)` whenever x is known to be (close to) continuously distributed and the data exhibits “enough” regularity.
3. *Number of Bins Selection.* The package `Binsreg` selects the number of bins J in a multi-step, data-driven and optimal way, whenever the user does not provide a selection manually (via the options `nbins()` or `bynbins()`). In large datasets, estimating J may be time consuming.
 - Speed improvement: provide J manually or use option `randcut(numeric)` to speed up the process.
4. *Gtools.* The package `Binsreg` is open source and, by default, relies exclusively on base commands and functions in `Stata` (as well as in `Python` and `R`). However, some parts of this algorithm (e.g., `pctile`) may be slow in large datasets.
 - Speed improvement: employ the community-distributed package `gtools` (Caceres 2024) via the option `usegtools(on)`. This community-distributed package needs to be installed separately by the user.
5. *Other Possibilities.* The package `Binsreg` offers several other options for increasing speed of execution. First, the community-distributed package `reghdfe` (Correia and Constantine 2024) could be used when employing the `binsreg` command. Second, for uncertainty quantification and inference, the number of simulations (option `nsims()`) and the number of grid points for simulation (option `simsgrid()`) can

be decreased to improve speed, which could offer a good alternative for preliminary exploration. Finally, for ultra-large datasets, it may be advisable to begin exploratory analysis with a random sample of the data, if the goal is to increase speed of execution of the package `Binsreg`.

3 Illustration of Methods

We illustrate the package `Binsreg` using a simulated dataset, which is available in the file `binscatter_simdata.dta`. In this dataset, `y` is the outcome variable, `x` is the independent variable for binning, `w` is a continuously distributed covariate, and `t` is a binary covariate, and `id` is a group identifier. Summary statistics of the simulated data are as follows.

```
. use binsreg_simdata, clear
. sum
```

Variable	Obs	Mean	Std. dev.	Min	Max
x	1,000	.4907072	.2932553	.0002281	.9985808
w	1,000	.0120224	.5799381	-.9993055	.9973198
t	1,000	.515	.500025	0	1
id	1,000	250.5	144.4095	1	500
y	1,000	.5283884	1.727878	-5.159858	5.751276
d	1,000	.45	.4977427	0	1

3.1 Estimation, Uncertainty Quantification and Plotting

The basic syntax for `binsreg` is the following:

```
. binsreg y x w
Sorting dataset on x...
Note: This step is omitted if dataset already sorted by x.
Binscatter plot
Bin selection method: IMSE-optimal plug-in choice (select # of bins)
Placement: Quantile-spaced
Derivative: 0
```

# of observations	1000
# of distinct values	1000
# of clusters	.

Bin/Degree selection:	
Degree of polynomial	0
# of smoothness constraints	0
# of bins	21
imse, bias ²	5.420
imse, var.	1.192

	p	s	df
dots	0	0	21

The main output is a binned scatter plot as shown in Figure 1. By default, the

(nonparametric) mean relationship between y and x is approximated by piecewise constants (`dots(0 0)`). Each dot in the figure represents the point estimate corresponding to each bin, which is the canonical binscatter plot. The number of bins, whenever not specified, is automatically selected via the companion command `binsregselect`. In this case, 21 bins are used. Other useful information is also reported, including total sample size, the number of distinct values of x , bin selection results, and the degrees of freedom of the statistical model(s) employed.

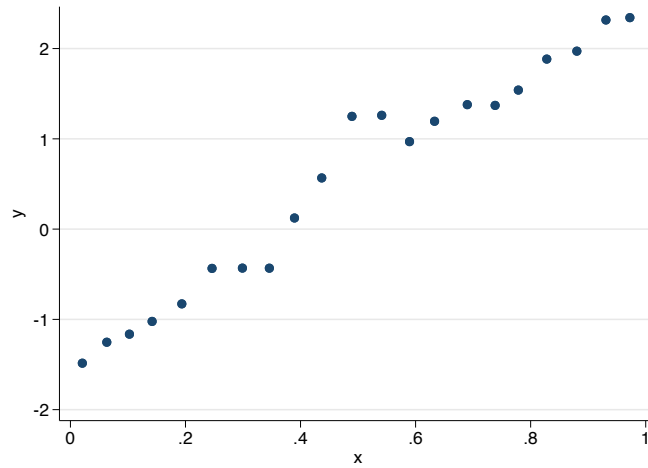


Figure 1: Canonical Binned Scatter Plot.

By default, the command `binsreg` evaluates and plots the regression function of interest $v_{\mathbf{w}}^{(v)}(x)$ at the mean of the additional covariates \mathbf{w}_i , i.e., $\mathbf{w} = \bar{\mathbf{w}}$. Users may specify a different value of \mathbf{w} , for example, the empirical median of each component in \mathbf{w}_i , via the option `at()`:

```
. binsreg y x w, at(median)
```

Users may also save the values of the additional covariates at which the binscatter estimate is evaluated in another file, and then specify the file name in the option `at()`. For example,

```
. tempfile evalcovar
. preserve
. clear
. set obs 1
Number of observations (_N) was 0, now 1.
. gen w=0.2
. gen t=1
. save `evalcovar`, replace
(file /var/folders/0b/h0w19g7d3s3dr9d_vm0qts9r0000gn/T//S_37285.000001 not found)
file /var/folders/0b/h0w19g7d3s3dr9d_vm0qts9r0000gn/T//S_37285.000001 saved as .dta
format
. restore
```

```
. binsreg y x w i.t, at(`evalcovar`)
Sorting dataset on x...
Note: This step is omitted if dataset already sorted by x.
Binscatter plot
Bin selection method: IMSE-optimal plug-in choice (select # of bins)
Placement: Quantile-spaced
Derivative: 0
```

# of observations	1000		
# of distinct values	1000		
# of clusters	.		
Bin/Degree selection:			
Degree of polynomial	0		
# of smoothness constraints	0		
# of bins	22		
imse, bias ²	4.736		
imse, var.	0.974		
<hr/>			
	p	s	df
dots	0	0	22

In this case, we control for a continuous variable w and a dummy variable generated based on the binary covariate t . We evaluate the binscatter estimate at $w=0.2$ and $t=1$, and these values are saved in the temporary file ‘evalcovar’ in advance.

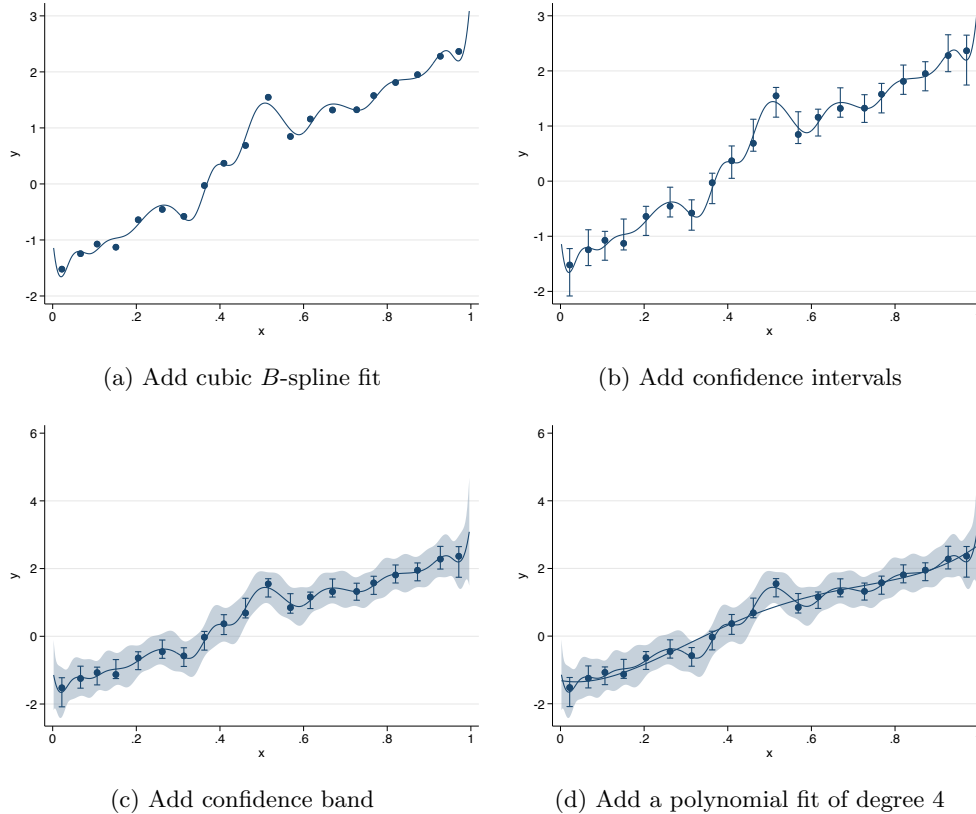
Users may specify the number of bins manually rather than relying on the automatic data-driven procedures. For example, a popular ad-hoc choice in practice is setting $J = 20$ quantile-spaced bins:

```
. binsreg y x w, nbins(20) polyreg(1)
Sorting dataset on x...
Note: This step is omitted if dataset already sorted by x.
Note: When additional covariates w are included, the polynomial fit may not always be close to the binscatter fit.
Binscatter plot
Bin selection method: User-specified
Placement: Quantile-spaced
Derivative: 0
```

# of observations	1000		
# of distinct values	1000		
# of clusters	.		
Bin/Degree selection:			
Degree of polynomial	0		
# of smoothness constraints	0		
# of bins	20		
<hr/>			
	p	s	df
dots	0	0	20
polyreg	1	NA	2

The option `polyreg(1)` adds a linear prediction line to the canonical binscatter plot, but the resulting binned scatter plot is not reported here to conserve space.

Figure 2: Binned Scatter Plot with Lines, Confidence Intervals and Bands.



The command `binsreg` allows users to add a binscatter-based line approximating the unknown regression function, pointwise confidence intervals, a uniform confidence band, and a global polynomial regression approximation. For example, the following syntax cumulatively adds in four distinct plots a fitted line, confidence intervals and a confidence band, all three based on cubic B -splines, and also a fitted line based on a global polynomial of degree 4. The results are shown in Figure 2.

```
. qui binsreg y x w, nbins(20) dots(0,0) line(3,3)
. qui binsreg y x w, nbins(20) dots(0,0) line(3,3) ci(3,3)
. qui binsreg y x w, nbins(20) dots(0,0) line(3,3) ci(3,3) cb(3,3)
. qui binsreg y x w, nbins(20) dots(0,0) line(3,3) ci(3,3) cb(3,3) polyreg(4)
```

By construction, a cubic B -spline fit is a piecewise cubic polynomial function which is continuous, and has continuous first- and second-order derivatives. Thus, the prediction line and confidence band generated are quite smooth. In this case, it is arguably under-smoothed because of the “large” choice of $J = 20$. The degree and smoothness of

polynomials can be changed by adjusting the values of `p` and `s` in the options `dots()`, `line()`, `ci()`, and `cb()`.

The command `binsreg` also allows for the standard `vce` options, factor variables, and `twoway` graph options, among other features. This is illustrated in the following code:

```
. binsreg y x w i.t, dots(0,0) line(3,3) ci(3,3) cb(3,3) polyreg(4) ///
>                                     vce(cluster id) savedata(output/graphdat) replace ///
>                                     title("Binned Scatter Plot")
Sorting dataset on x...
Note: This step is omitted if dataset already sorted by x.
Note: When additional covariates w are included, the polynomial fit may not always be close to the binscatter fit.
Note: Setting at least nsims(2000) and simsgrid(50) is recommended to obtain the final results.
> ults.
Binscatter plot
Bin selection method: IMSE-optimal plug-in choice (select # of bins)
Placement: Quantile-spaced
Derivative: 0
Output file: output/graphdat.dta
```

# of observations	1000
# of distinct values	1000
# of clusters	500
Bin/Degree selection:	
Degree of polynomial	0
# of smoothness constraints	0
# of bins	20
imse, bias ²	3.588
imse, var.	0.494

	p	s	df
<code>dots</code>	0	0	20
<code>line</code>	3	3	23
<code>CI</code>	3	3	23
<code>CB</code>	3	3	23
<code>polyreg</code>	4	NA	5

Specifically, a dummy variable based on the binary covariate `t` is added to the estimation, standard errors are clustered at the group level indicator `id`, and a graph title is added to the resulting binned scatter plot. Note that any unrecognized options for the command `binsreg` will be understood as `twoway` options and therefore appended to the final plot command. Thus, users may easily modify, for example, axis properties, legends, etc. The option `savedata(graphdat)` saves the underlying data used in the binned scatter plot in the file `graphdat.dta`.

In addition, the command `binsreg` can be used for subgroup analysis. The following command implements binscatter estimation and inference across two subgroups separately, defined by the variable `t`, and then produces a common binned scatter plot (Figure 3):

```
. binsreg y x w, by(t) dots(0,0) line(3,3) cb(3,3) ///
>                                     bycolors(blue red) bysymbols(0 T)
Sorting dataset on x...
```

Note: This step is omitted if dataset already sorted by x.
 Note: Setting at least nsims(2000) and simsgrid(50) is recommended to obtain the final res > ults.

Binscatter plot
 Bin selection method: IMSE-optimal plug-in choice (select # of bins)
 Placement: Quantile-spaced
 Derivative: 0
 Group: t = 0

# of observations	485
# of distinct values	485
# of clusters	.
Bin/Degree selection:	
Degree of polynomial	0
# of smoothness constraints	0
# of bins	20
imse, bias ²	7.092
imse, var.	0.941

	p	s	df
dots	0	0	20
line	3	3	23
CB	3	3	23

Note: Setting at least nsims(2000) and simsgrid(50) is recommended to obtain the final res > ults.

Group: t = 1

# of observations	515
# of distinct values	515
# of clusters	.
Bin/Degree selection:	
Degree of polynomial	0
# of smoothness constraints	0
# of bins	15
imse, bias ²	2.861
imse, var.	0.955

	p	s	df
dots	0	0	15
line	3	3	18
CB	3	3	18

Figure 3 highlights a difference across the two subgroups defined by the variable t , which corresponds to the fact that our simulated data add a 1 to the outcome variable for those units with $t = 1$. The colors, symbols, and line patterns in Figure 3 can be modified via the options `bycolors()`, `bysymbols()`, and `bylpatterns()`. When the number of bins is unspecified, the command `binsreg` selects the number of bins for each subsample separately, via the companion command `binsregselect`. This means that, by default, the choice of binning/partitioning structure will be different across subgroups in general. However, if the option `samebinsby` is specified, then a common binning scheme for all subgroups is constructed based on the full sample.

As described before, sometimes one would like to keep the number of bins J fixed

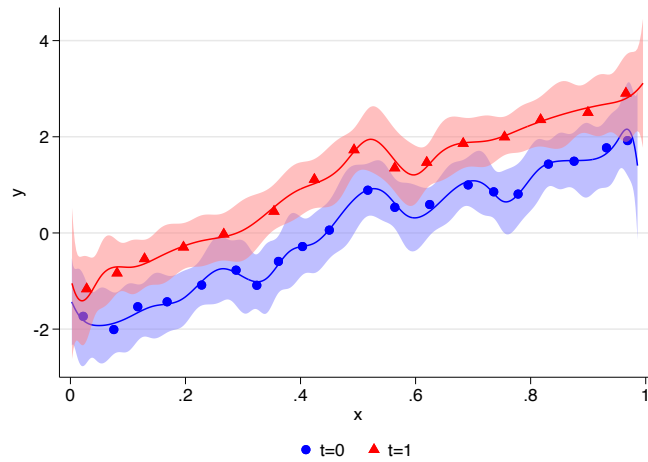


Figure 3: Binned Scatter Plot: Group Comparison

and select the degree/smoothness of the polynomial instead. The following snippet illustrates how to implement this procedure:

```
. binsreg y x w, nbins(20) line(T) ci(T) cb(T) pselect(0/3)
Sorting dataset on x...
Note: This step is omitted if dataset already sorted by x.
Note: Setting at least nsims(2000) and simsgrid(50) is recommended to obtain the final res
> ults.
```

```
Binscatter plot
Bin selection method: IMSE-optimal plug-in choice (select degree and smoothness)
Placement: Quantile-spaced
Derivative: 0
```

# of observations	1000
# of distinct values	1000
# of clusters	.
Bin/Degree selection:	
Degree of polynomial	0
# of smoothness constraints	0
# of bins	20
imse, bias ²	5.422
imse, var.	1.181

	p	s	df
dots	0	0	20
line	0	0	20
CI	1	1	21
CB	1	1	21

Here we let $J = 20$ and select the degree of polynomial p within the specified range $\{0, 1, 2, 3\}$. The resulting optimal p is 0. Accordingly, we construct “dots” and “line” based on piecewise constant estimates, and confidence intervals and a confidence band

based on linear spline estimates.

The accompanying replication files include other illustrations. For example:

- Inference based on asymptotic variance formula:

```
. binsreg y x w i.t, dots(0 0) line(3 3) ci(3 3) cb(3 3) polyreg(4)
vce(cluster id) asyvar(on)
```
- Using the community-contributed module `reghdfe`:

```
. binsreg y x w, absorb(t) dots(0 0) line(3 3) ci(3 3) cb(3 3) polyreg(4)
```
- Turning off data distribution robustness checks and using the community-contributed module `gtools`:

```
. binsreg y x w, masspoints(off) usegtools(on)
```

Next, we illustrate the command `binsqreg` for estimation and uncertainty quantification using quantile regression `binscatter` methods. The following code looks at the conditional 25-th quantile of the outcome variable.

```
. binsqreg y x w, quantile(0.25)
Sorting dataset on x...
Note: This step is omitted if dataset already sorted by x.
Binscatter plot, quantile
Bin selection method: IMSE-optimal plug-in choice (select # of bins)
Placement: Quantile-spaced
Derivative: 0
```

# of observations	1000
# of distinct values	1000
# of clusters	.
Bin/Degree selection:	
Degree of polynomial	0
# of smoothness constraints	0
# of bins	21
imse, bias ²	5.420
imse, var.	1.192

	p	s	df
dots	0	0	21

By default, quantile regression methods employ an analytic variance estimator formula, which may not perform well in applications. A more robust alternative is employing bootstrap methods:

```
. binsqreg y x w, quantile(0.25) ci(3 3) vce(bootstrap, reps(100))
Sorting dataset on x...
Note: This step is omitted if dataset already sorted by x.
Binscatter plot, quantile
Bin selection method: IMSE-optimal plug-in choice (select # of bins)
Placement: Quantile-spaced
Derivative: 0
```

# of observations	1000
-------------------	------

# of distinct values	1000		
# of clusters	.		
Bin/Degree selection:			
Degree of polynomial	0		
# of smoothness constraints	0		
# of bins	21		
imse, bias ²	5.420		
imse, var.	1.192		
<hr/>			
	p	s	df
dots	0	0	21
CI	3	3	24

The replication files also illustrate how to plot together least squares and quantile regression binscatter approximations. The final output is illustrated in Figure 4, which plots the conditional mean and its confidence band, together with the conditional 25-th and 75-th quantile regressions (i.e., conditional inter-quartile range).

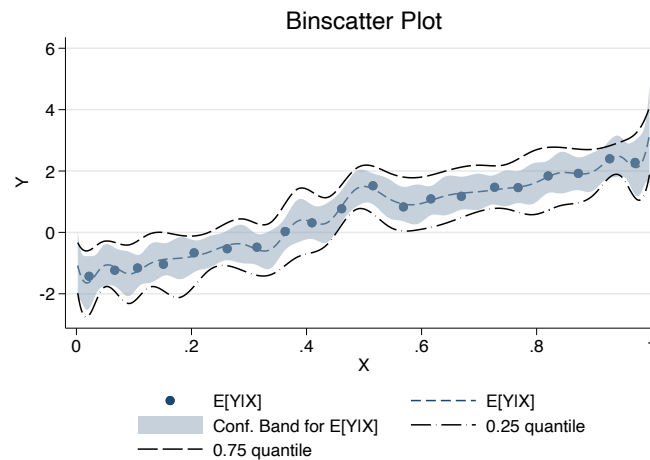


Figure 4: Binned Scatter Plot for Means and Quantiles.

Moreover, we provide a convenient option `qregopt(qreg_option)` to modify the underlying quantile regression. For example, the user can control the optimization process as follows:

```
. qui binsqreg y x w, quantile(0.25) qregopt(iterate(1000) wls(1))
```

Finally, we illustrate the command `binslogit` for binary response regression binscatter methods using logistic regression.

```
. binslogit d x w
Sorting dataset on x...
```

Note: This step is omitted if dataset already sorted by x.
 Binscatter plot, logit model
 Bin selection method: IMSE-optimal plug-in choice (select # of bins)
 Placement: Quantile-spaced
 Derivative: 0

# of observations	1000
# of distinct values	1000
# of clusters	.
Bin/Degree selection:	
Degree of polynomial	0
# of smoothness constraints	0
# of bins	12
imse, bias ²	0.172
imse, var.	0.208

	p	s	df
dots	0	0	12

3.2 Hypothesis Testing and Statistical Inference

We illustrate first the syntax and outputs of the command `binstest`. The basic syntax is the following:

```
. binstest y x w, testmodelpoly(1)
Note: Setting at least nsims(2000) and simsgrid(50) is recommended to obtain the final res
> ults.
```

Hypothesis tests based on binscatter estimates
 Estimation method: reg
 Bin selection method: IMSE-optimal plug-in choice (select # of bins)
 Placement: Quantile-spaced
 Derivative: 0

# of observations	1000
# of distinct values	1000
# of clusters	.
Bin/Degree selection:	
Degree of polynomial	0
# of smoothness constraints	0
# of bins	21

Model specification Tests:
 Degree: 1 # of smoothness constraints: 1

H0: mu =	sup T	p value
poly. degree 1	6.566	0.000

A test for linearity of the regression function $\mu_0(x)$ is implemented using the binscatter estimator. By default, a linear B -spline is employed in the inference procedure, which can be adjusted by the option `testmodel()`. In addition, when unspecified, the number of bins is selected using a data-driven procedure via the companion command `binsregselect`. The selected number of bins is IMSE-optimal for piecewise constant

point estimates by default. A summary of the sample and binning scheme is displayed, and then the test statistic and p-value are reported. In this case, the test statistic is the supremum of the absolute value of the t -statistic evaluated over a sequence of grid points, and the p-value is calculated based on simulation. Clearly, the p-value is quite small, and thus the null hypothesis of linearity of the regression function is rejected.

As emphasized before, the parametric specification test for a null hypothesis about the level may be sensitive to the choice of evaluation point \mathbf{w} . Thus, a recommended strategy to test for linearity of $\mu_0(x)$ is to check if its first derivative is a constant, which is implemented in the following:

```
. binsctest y x w, testmodelpoly(1) deriv(1)
Note: Setting at least nsims(2000) and simsgrid(50) is recommended to obtain the final res
> ults.
Hypothesis tests based on binscatter estimates
Estimation method: reg
Bin selection method: IMSE-optimal plug-in choice (select # of bins)
Placement: Quantile-spaced
Derivative: 1
```

# of observations	1000
# of distinct values	1000
# of clusters	.

Bin/Degree selection:	
Degree of polynomial	1
# of smoothness constraints	1
# of bins	6

```
Model specification Tests:
Degree: 2      # of smoothness constraints: 2
```

H0: mu =	sup T	p value
poly. degree 1	4.114	0.000

Note that for $v = 1$, the selected number of bins is IMSE-optimal for the linear B -spline estimate by default, and the test statistic based on the proposed robust bias correction strategy is constructed using a quadratic B -spline fit.

The command `binsctest` can implement testing for any parametric model specification by comparing the fitted values based on the `binscatter` estimator (computed by the command) and the parametric model of interest (provided by the user). For example, the following code creates an auxiliary database with a grid of evaluation points, implements a linear regression first, makes an out-of-sample prediction using the auxiliary dataset, and then tests for linearity based on the `binscatter` estimator by specifying the auxiliary file containing the fitted values.

```
. qui binsregselect y x w, simsgrid(30) savegrid(output/parfitval) replace
. qui reg y x w
. use output/parfitval, clear
. predict binsreg_fit_lm
(option xb assumed; fitted values)
. save output/parfitval, replace
file output/parfitval.dta saved
```



```

. use binsreg_simdata, clear
. binstest y x w, testmodelparfit(output/parfitval) lp(2) deriv(1)
Note: Setting at least nsims(2000) and simsgrid(50) is recommended to obtain the final res
> ults.
Hypothesis tests based on binscatter estimates
Estimation method: reg
Bin selection method: IMSE-optimal plug-in choice (select # of bins)
Placement: Quantile-spaced
Derivative: 1

```

# of observations	1000
# of distinct values	1000
# of clusters	.

```

Bin/Degree selection:
  Degree of polynomial      1
# of smoothness constraints 1
# of bins                   6

```

```

Model specification Tests:
Degree: 2 # of smoothness constraints: 2
Input file: output/parfitval.dta

```

H0: mu =	L2 of T	p value
binsreg_fit_lm	4.377	0.000

The first line, `binsregselect y x w, simsgrid(30) savegrid(output/parfitval)` replace, generates the auxiliary file containing the grid of evaluation points. Since the parameter of interest is only the mean relation between y and x , i.e., $\mu_0(x)$, at the out-of-sample prediction step, the testing dataset `parfitval.dta` must contain a variable `x` containing a sequence of evaluation points at which the binscatter and parametric models are compared, and the covariate `w` whose values are set to be zeros. In addition, the variable containing fitted values has to follow a specific naming rule, i.e., takes the form of `binsreg_fit*`. The companion command `binsregselect` can be used to construct the required auxiliary dataset, as illustrated above. We discuss this other command further below.

In addition to model specification tests, the command `binstest` can test for non-parametric shape restrictions on the regression function. For example, the following syntax tests whether the regression function is increasing:

```

. binstest y x w, deriv(1) nbins(20) testshaper(0)
Warning: Testing procedures are valid when nbins() is much larger than the IMSE-optimal ch
> oice. Compare your choice with the IMSE-optimal one obtained by binsregselect.
Note: Setting at least nsims(2000) and simsgrid(50) is recommended to obtain the final res
> ults.
Hypothesis tests based on binscatter estimates
Estimation method: reg
Bin selection method: User-specified
Placement: User-specified
Derivative: 1

```

# of observations	1000
# of distinct values	1000
# of clusters	.

```

Bin/Degree selection:

```

Degree of polynomial	.
# of smoothness constraints	.
# of bins	20

Shape Restriction Tests:
Degree: 2 # of smoothness constraints: 2

H0: inf mu >=	inf T	p value
0	-3.709	0.004

The null hypothesis here is that the infimum of the first-order derivative of the regression function is no less than 0. The output reports the test statistic, which is the infimum of the t -statistic over a sequence of evaluation points, and the corresponding simulation-based p-value.

The command `binstest` may implement many tests simultaneously (given the derivative of interest). For example,

```
. binstest y x w, nbins(20) testshaper(-2 0) testshapel(4) testmodelpoly(1) ///
> nsims(1000) simsgrid(30)
Warning: Testing procedures are valid when nbins() is much larger than the IMSE-optimal ch
> oice. Compare your choice with the IMSE-optimal one obtained by binsregselect.
Note: Setting at least nsims(2000) and simsgrid(50) is recommended to obtain the final res
> ults.
```

```
Hypothesis tests based on binscatter estimates
Estimation method: reg
Bin selection method: User-specified
Placement: User-specified
Derivative: 0
```

# of observations	1000
# of distinct values	1000
# of clusters	.

Bin/Degree selection:	
Degree of polynomial	.
# of smoothness constraints	.
# of bins	20

Shape Restriction Tests:
Degree: 1 # of smoothness constraints: 1

H0: sup mu <=	sup T	p value
4	-4.131	1.000

H0: inf mu >=	inf T	p value
-2	1.774	1.000
0	-10.772	0.000

Model specification Tests:
Degree: 1 # of smoothness constraints: 1

H0: mu =	sup T	p value
poly. degree 1	5.972	0.000

The above syntax tests three shape restrictions and one model specification (linearity), employing 1000 random draws from \mathbf{N}_K^* and 30 evaluation points to evaluate the supremum/infimum in the simulation.

The accompanying replication files include other illustrations. For example:

- Testing whether the median regression function is linear:
`. binstest y x w, estmethod(qreg 0.5) testmodelpoly(1)`
- Testing whether the nonlinear logistic regression function is increasing:
`. binstest d x w, estmethod(logit) deriv(1) nbins(20) testshaper(0)`

Next, consider pairwise comparison as implemented via the testing command `binspwc`. Using least square `binscatter` for the two samples identified via the binary variable `t`, we have:

```
. binspwc y x w, by(t)
Note: Setting at least nsims(2000) and simsgrid(50) is recommended to obtain the final res
> ults.
Pairwise group comparison based on binscatter estimates
Estimation method: reg
Derivative: 0
Group variable: t
Bin/Degree selection method: IMSE-optimal plug-in choice (select # of bins)
Placement: Quantile-spaced
Group 1 vs. Group 0
```

Group t=	1	0
# of observations	515	485
# of distinct values	515	485
# of clusters	.	.
Degree of polynomial	1	1
# of smoothness constraints	1	1
# of bins	15	20

```
diff = group 1 - group 0
```

H0:	sup T	p value
diff=0	5.790	0.000

Similarly, employing quantile regression `binscatter` methods, for the 40-th conditional quantile of the outcome variable, we have

```
. binspwc y x w, by(t) estmethod(qreg 0.4)
Note: Setting at least nsims(2000) and simsgrid(50) is recommended to obtain the final res
> ults.
Pairwise group comparison based on binscatter estimates
Estimation method: qreg
Derivative: 0
Group variable: t
Bin/Degree selection method: IMSE-optimal plug-in choice (select # of bins)
Placement: Quantile-spaced
Group 1 vs. Group 0
```

Group t=	1	0

# of observations	515	485
# of distinct values	515	485
# of clusters	.	.
Degree of polynomial	1	1
# of smoothness constraints	1	1
# of bins	15	20

diff = group 1 - group 0

H0:	sup T	p value
diff=0	5.023	0.000

3.3 Binning Selection

As already mentioned, all our estimation and inference commands rely on data-driven bin selection procedures via the command `binsregselect` whenever the option `nbins()` is not employed by the user. Its basic syntax is as follows:

```
. binsregselect y x w
Bin selection for binscatter estimates
Method: IMSE-optimal plug-in choice (select # of bins)
Position: Quantile-spaced
```

# of observations	1000
# of distance values	1000
# of clusters	.
eff. sample size	1000
Degree of polynomial	0
# of smoothness constraint	0

method	# of bins	df	imse, bias ²	imse, var.
ROT-POLY	18	18	3.295	1.212
ROT-REGUL	18	18	.	.
ROT-UKNOT	18	18	.	.
DPI	21	21	5.420	1.192
DPI-UKNOT	21	21	.	.

df: degrees of freedom.

The following choices of number of bins are reported: `ROT-POLY`, the rule-of-thumb (ROT) choice based on global polynomial estimation; `ROT-REGUL`, the ROT choice regularized as discussed in Section 2, or the user's choice specified in the option `nbinsrot()`; `ROT-UKNOT`, the ROT choice with unique knots; `DPI`, the direct plug-in (DPI) choice; and `DPI-UKNOT`, the DPI choice with unique knots.

The direct plug-in choice is implemented based on the rule-of-thumb choice, which can be set by users directly:

```
. binsregselect y x w, nbinsrot(20) binspos(es)
Bin selection for binscatter estimates
Method: IMSE-optimal plug-in choice (select # of bins)
Position: Evenly-spaced
```

# of observations	1000
-------------------	------

# of distance values		1000		
# of clusters		.		
eff. sample size		1000		
<hr/>				
Degree of polynomial		0		
# of smoothness constraint		0		
<hr/>				
method	# of bins	df	imse, bias ²	imse, var.
ROT-POLY
ROT-REGUL	20	20	.	.
ROT-UKNOT	20	20	.	.
DPI	22	22	5.793	1.194
DPI-UKNOT	22	22	.	.

df: degrees of freedom.

Notice that in the example above an evenly-spaced, rather than quantile-spaced, binning scheme is selected via the option `binspos(es)`. The binning used in other commands may be adjusted similarly.

In addition, as illustrated above, the command `binsregselect` also provides a convenient option `savegrid()`, which can be used to generate the auxiliary dataset needed for parametric specification testing of user-chosen models via the command `binstest`. Specifically, the following command was (quietly) used above:

```
. binsregselect y x w, simsgrid(30) savegrid(output/parfitval) replace
Bin selection for binscatter estimates
Method: IMSE-optimal plug-in choice (select # of bins)
Position: Quantile-spaced
Output file: output/parfitval.dta
```

# of observations		1000		
# of distance values		1000		
# of clusters		.		
eff. sample size		1000		
<hr/>				
Degree of polynomial		0		
# of smoothness constraint		0		
<hr/>				
method	# of bins	df	imse, bias ²	imse, var.
ROT-POLY	18	18	3.295	1.212
ROT-REGUL	18	18	.	.
ROT-UKNOT	18	18	.	.
DPI	21	21	5.420	1.192
DPI-UKNOT	21	21	.	.

df: degrees of freedom.

The resulting file, `parfitval.dta`, includes `x` and `w` as well as some other variables related to the binning scheme. The variable `x` contains a sequence of evaluation points, in this case set to 30 within each bin via the option `simsgrid()`, and the values of `w` are set to zero on purpose (this is used to generate the fitting model correctly).

When an extremely large dataset is available, the data-driven procedures for selecting the binning scheme could be very time-consuming. In such a scenario, one could use a small sub-sample to estimate the leading constants in the integrated mean squared

error (IMSE) expansions, and then extrapolate the optimal number of bins to the full sample. The following code illustrates how this method is implemented:

```
. binsregselect y x w if t==0, useeffn(1000)
Bin selection for binscatter estimates
Method: IMSE-optimal plug-in choice (select # of bins)
Position: Quantile-spaced
```

# of observations	485			
# of distance values	485			
# of clusters	.			
eff. sample size	1000			
<hr/>				
Degree of polynomial	0			
# of smoothness constraint	0			
<hr/>				
method	# of bins	df	imse, bias ²	imse, var.
ROT-POLY	20	20	3.185	0.937
ROT-REGUL	20	20	.	.
ROT-UKNOT	20	20	.	.
DPI	26	26	7.092	0.941
DPI-UKNOT	26	26	.	.

df: degrees of freedom.

In this example 485 observations with $t = 0$ are used to compute the leading constants $\mathcal{B}_n(p, s, v)$ and $\mathcal{V}_n(p, s, v)$ in the IMSE expansion, but then the reported optimal numbers of bins are calculated based on the effective sample size specified in the option `useeffn()`. This method also applies to extrapolating the optimal number of bins to a smaller sample based on a larger one.

Alternatively, the number of bins selection can be implemented using only a random sub-sample of the observations. For example, the following command selects J using, in expectation, 30% of the observations based on uniformly distributed random numbers. Repeated application of this command will produce modestly different bins selection choices depending on the sequence of realized uniform random variables.

```
. binsregselect y x w, randcut(0.3)
Bin selection for binscatter estimates
Method: IMSE-optimal plug-in choice (select # of bins)
Position: Quantile-spaced
```

# of observations	1000			
# of distance values	1000			
# of clusters	.			
eff. sample size	1000			
<hr/>				
Degree of polynomial	0			
# of smoothness constraint	0			
<hr/>				
method	# of bins	df	imse, bias ²	imse, var.
ROT-POLY	18	18	3.714	1.303
ROT-REGUL	18	18	.	.
ROT-UKNOT	18	18	.	.
DPI	23	23	7.104	1.289
DPI-UKNOT	23	23	.	.

df: degrees of freedom.

4 Conclusion

We have introduced the package `Binsreg`, which provides general-purpose software implementations of binned scatter plots in different statistical models. The package is based on Cattaneo, Crump, Farrell, and Feng (2024b) and Cattaneo, Crump, Farrell, and Feng (2024a), which yields a thorough treatment and broad applicability, but does entail several limitations. Chiefly, `Binsreg` focuses on cross-sectional settings and, by focusing on standard binscatters, univariate x for plotting purposes. However, the general idea of binning data for visualization and for inference has appeal outside these areas, and would make for useful extensions of our work. Binning is a commonly used method, for example, in empirical finance, where it is used to create portfolios for factor creation and anomaly detection (Cattaneo et al. 2020a). Extending `Binsreg` to panel data to cover such cases would be valuable. The same is true for time series data, where the binning may need to be adapted to accommodate, or elucidate, features of the dependence in the data. Finally, univariate binscatters are closely related to bivariate heatmaps, a popular visualization tool in social and physical sciences, and more broadly even in data-based journalism. It is an open but important question of how to add the statistical rigor we have brought to binscatters to the world of heatmaps, or beyond into multidimensional binning.

5 Acknowledgments

We thank Michael Droste, John Friedman, Andreas Fuster, Filippo Palomba, Paul Goldsmith-Pinkham, David Lucca, Xinwei Ma, Ricardo Masini, Jonah Rockoff, Jesse Rothstein, Ryan Santos, Jesse Shapiro, and Rocio Titiunik for helpful comments and discussions. We also thank the Editor, Stephen Jenkins, and an anonymous reviewer for their detailed comments that improved our paper. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Cattaneo gratefully acknowledges financial support from the National Science Foundation through grants SES-1947805, SES-2019432, and SES-2241575. Feng gratefully acknowledges financial support from the National Natural Science Foundation of China (NSFC) through grants 72203122, 72133002, and 72250064.

6 References

- Caceres, M. 2024. `Stata` Module `Mauricio` <https://github.com/mcaceresb/stata-gtools>.
- Calonico, S., M. D. Cattaneo, and M. H. Farrell. 2018. On the Effect of Bias Estimation on Coverage Accuracy in Nonparametric Inference. *Journal of the American Statistical Association* 113(522): 767–779.

- . 2022. Coverage Error Optimal Confidence Intervals for Local Polynomial Regression. *Bernoulli* 28(4): 2998–3022.
- Calonico, S., M. D. Cattaneo, and R. Titiunik. 2014. Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica* 82(6): 2295–2326.
- Cattaneo, M. D., R. K. Crump, M. H. Farrell, and Y. Feng. 2024a. Nonlinear Binscatter Methods. working paper .
- . 2024b. On Binscatter. *American Economic Review* 114(5): 1488–1514.
- Cattaneo, M. D., R. K. Crump, M. H. Farrell, and E. Schaumburg. 2020a. Characteristic-Sorted Portfolios: Estimation and Inference. *arXiv:1809.03584, Review of Economics and Statistics* 101(3): 531–551.
- Cattaneo, M. D., M. H. Farrell, and Y. Feng. 2020b. Large sample properties of partitioning-based series estimators. *Annals of Statistics* 48(3): 1718–1741.
- Correia, S., and N. Constantine. 2024. **Stata** Module **reghdfe** <https://github.com/sergiocorreia/reghdfe>.
- Droste, M. 2019. **Stata** Module **Binscatter2** <https://github.com/mdroste/stata-binscatter2/>.
- Healy, K. 2018. *Data Visualization: A Practical Introduction*. Princeton University Press.
- Schwabish, J. 2021. *Better Data Visualizations: A Guide for Scholars, Researchers, and Wonks*. Columbia University Press.
- Starr, E., and B. Goldfarb. 2020. Binned Scatterplots: A Simple Tool to Make Research Easier and Better. *Strategic Management Journal* 41(12): 2261–2274.
- Stepner, M. 2017. **Stata** Module **Binscatter** <https://github.com/michaelstepner/binscatter/>.

7 About the Authors

Matias D. Cattaneo is a Professor of Operations Research and Financial Engineering at Princeton University.

Richard K. Crump is a Financial Research Advisor in Macroeconomics Studies at the Federal Reserve Bank of New York.

Max H. Farrell is an Associate Professor of Economics at the University of California at Santa Barbara.

Yingjie Feng is an Associate Professor of Economics at Tsinghua University.