

Honest inference in Sharp Regression Discontinuity

Michal Kolesár

2017-06-01

Contents

Sharp RD model	1
Plots	2
Inference based on local polynomial estimates	3
Parameter space	4
Inference when running variable is discrete	6
Optimal inference	8
Specification testing	9
References	10

The package `RDHonest` implements confidence intervals for the regression discontinuity parameter considered in Armstrong and Kolesár (2017), Armstrong and Kolesár (2016), and Kolesár and Rothe (2017). In this vignette, we demonstrate the implementation of these confidence intervals using datasets from Lee (2008) and Oreopoulos (2006), which are included in the package as a data frame `lee08` and `cghs`. The datasets from Lalive (2008) and Ludwig and Miller (2007) that are used in Armstrong and Kolesár (2016), and Kolesár and Rothe (2017) are also included in the package as data frames `rebp` and `headst`.

Sharp RD model

In the sharp regression discontinuity model, we observe units $i = 1, \dots, n$, with the outcome y_i for the i th unit given by

$$y_i = f(x_i) + u_i,$$

where $f(x_i)$ is the expectation of y_i conditional on the running variable x_i and u_i is the regression error. A unit is treated if and only if the running variable x_i lies above a known cutoff c_0 . The parameter of interest is given by the jump of f at the cutoff,

$$\beta = \lim_{x \downarrow c_0} f(x) - \lim_{x \uparrow c_0} f(x).$$

Let $\sigma^2(x_i)$ denote the conditional variance of u_i .

In the Lee dataset, the running variable corresponds to the margin of victory of a Democratic candidate in a US House election, and the treatment corresponds to winning the election. Therefore, the cutoff is zero. The outcome of interest is the Democratic vote share in the following election.

The Oreopoulos dataset consists of a subsample of British workers, and it exploits a change in minimum school leaving age in the UK from 14 to 15, which occurred in 1947. The running variable is the year in which the individual turned 14, with the cutoff equal to 1947 so that the “treatment” is being subject to a higher minimum school-leaving age. The outcome is log earnings in 1998.

Some of the functions in the package require the data to be transformed into a custom `RDData` format. This can be accomplished with the `RDData` function:

```

library("RDHonest")
## Assumes first column in the data frame corresponds to
## outcome, and second to running variable
dl <- RDDData(lee08, cutoff = 0)
## Transform earnings to log earnings
do <- RDDData(data.frame(logearn = log(cghs$earnings), year14 = cghs$yearat14),
              cutoff = 1947)

```

Plots

The package provides a function `plot_RDscatter` to plot the raw data. To remove some noise, the function plots averages over avg number of observations. The function takes an `RDDData` object as an argument

```

## plot 25-bin averages in for observations 50 at most
## points away from the cutoff. See Figure 1
plot_RDscatter(dl, avg = 25, window = 50, xlab = "Margin of victory",
              ylab = "Vote share in next election")

```




Figure 1: Lee (2008) data

The running variable in the Oreopoulos dataset is discrete. It is therefore natural to plot the average outcome by each value of the running variable, which is achieved using by setting `avg=Inf`. The option `dotsize="count"` makes the size of the points proportional to the number of observations that the point averages over.

```

## see Figure 2
f2 <- plot_RDscatter(do, avg = Inf, xlab = "Year aged 14",
                  ylab = "Log earnings", propdotsize = TRUE)
## Adjust size of dots if they are too big
f2 + ggplot2::scale_size_area(max_size = 4)

```

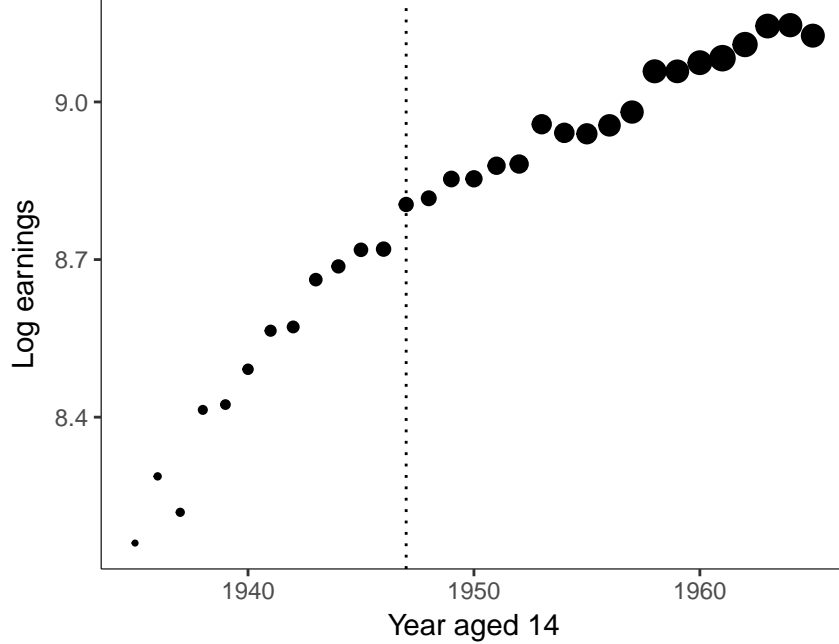


Figure 2: Oreopoulos (2006) data

Inference based on local polynomial estimates

The function `RDHonest` constructs one- and two-sided confidence intervals (CIs) around local linear and local quadratic estimators using either a user-supplied bandwidth (which is allowed to differ on either side of the cutoff), or bandwidth that is optimized for a given performance criterion. The sense of honesty is that, if the regression errors are normally distributed with known variance, the CIs are guaranteed to achieve correct coverage *in finite samples*, and achieve correct coverage asymptotically uniformly over the parameter space otherwise. Furthermore, because the CIs explicitly take into account the possible bias of the estimators, the asymptotic approximation doesn't rely on the bandwidth to shrink to zero at a particular rate.

To describe the form of the CIs, let $\hat{\beta}_{h_+,h_-}$ denote a local polynomial estimator with bandwidth equal to h_+ above the cutoff and equal to h_- below the cutoff. Let $\beta_{h_+,h_-}(f)$ denote its expectation conditional on the covariates when the regression function equals f . Then the bias of the estimator is given by $\beta_{h_+,h_-}(f) - \beta$. Let

$$B(\hat{\beta}_{h_+,h_-}) = \sup_{f \in \mathcal{F}} |\beta_{h_+,h_-}(f) - \beta|$$

denote the worst-case bias over the parameter space \mathcal{F} . Then the lower limit of a one-sided CI is given by

$$\hat{\beta}_{h_+,h_-} - B(\hat{\beta}_{h_+,h_-}) - z_{1-\alpha} \widehat{se}(\hat{\beta}_{h_+,h_-}),$$

where $z_{1-\alpha}$ is the $1 - \alpha$ quantile of a standard normal distribution, and $\widehat{se}(\hat{\beta}_{h_+,h_-})$ is the standard error (an estimate of the standard deviation of the estimator). Subtracting the worst-case bias in addition to the usual critical value times standard error ensures correct coverage at all points in the parameter space.

A two-sided CI is given by

$$\hat{\beta}_{h_+,h_-} \pm cv_{1-\alpha}(B(\hat{\beta}_{h_+,h_-})/\widehat{se}(\hat{\beta}_{h_+,h_-})) \times \widehat{se}(\hat{\beta}_{h_+,h_-}),$$

where the critical value function $cv_{1-\alpha}(b)$ corresponds to the $1 - \alpha$ quantile of the $|N(b, 1)|$ distribution. To see why using this critical value ensures honesty, decompose the t -statistic as

$$\frac{\hat{\beta}_{h_+,h_-} - \beta}{\widehat{se}(\hat{\beta}_{h_+,h_-})} = \frac{\hat{\beta}_{h_+,h_-} - \beta_{h_+,h_-}(f)}{\widehat{se}(\hat{\beta}_{h_+,h_-})} + \frac{\beta_{h_+,h_-}(f) - \beta}{\widehat{se}(\hat{\beta}_{h_+,h_-})}$$

By a central limit theorem, the first term on the right-hand side will be distributed standard normal, irrespective of the bias. The second term is bounded in absolute value by $B(\hat{\beta}_{h_+,h_-})/\widehat{se}(\hat{\beta}_{h_+,h_-})$, so that, in large samples, the $1 - \alpha$ quantile of the absolute value of the t -statistic will be bounded by $cv_{1-\alpha}(B(\hat{\beta}_{h_+,h_-})/\widehat{se}(\hat{\beta}_{h_+,h_-}))$. This approach gives tighter CIs than simply adding and subtracting $B(\hat{\beta}_{h_+,h_-})$ from the point estimate, in addition to adding and subtracting $z_{1-\alpha}\widehat{se}(\hat{\beta}_{h_+,h_-})$

The function `CVb` gives these critical values:

```
## Usual critical value
CVb(0, alpha = 0.05) # returns a list
#> bias alpha cv TeXDescription
#> 1 0 0.05 1.95996 $\alpha=0.05$
CVb(1/2, alpha = 0.05)$cv # extract critical value
#> [1] 2.18148
## Tabulate critical values for different significance
## levels when bias-sd ratio equals 1/4
knitr::kable(CVb(1/4, alpha = c(0.01, 0.05, 0.1)), caption = "Critical values")
```

Table 1: Critical values

bias	alpha	cv	TeXDescription
0.25	0.01	2.65224	$\alpha = 0.01$
0.25	0.05	2.01971	$\alpha = 0.05$
0.25	0.10	1.69558	$\alpha = 0.1$

The field `TeXDescription` is useful for plotting, or for exporting to \LaTeX , as in the table above.

Parameter space

To implement the honest CIs, one needs to specify the parameter space \mathcal{F} . The function `RDHonest` computes honest CIs when the parameter space \mathcal{F} corresponds to a second-order Taylor or second-order Hölder smoothness class, which capture two different types of smoothness restrictions. The second-order Taylor class assumes that f lies in the the class of functions

$$\mathcal{F}_{\text{Taylor}}(M) = \{f_+ - f_- : f_+ \in \mathcal{F}_T(M; [c_0, \infty)), f_- \in \mathcal{F}_T(M; (-\infty, c_0))\},$$

where $\mathcal{F}_T(M; \mathcal{X})$ consists of functions f such that the approximation error from second-order Taylor expansion of $f(x)$ about c_0 is bounded by $M|x|^2/2$, uniformly over \mathcal{X} :

$$\mathcal{F}_T(M; \mathcal{X}) = \{f : |f(x) - f(c_0) - f'(c_0)x| \leq M|x|^2/2 \text{ all } x \in \mathcal{X}\}.$$

The class $\mathcal{F}_T(M; \mathcal{X})$ formalizes the idea that the second derivative of f at zero should be bounded by M . See Section 2 in Armstrong and Kolesár (2017) (note the constant C in that paper equals $C = M/2$ here). This class is doesn't impose smoothness away from boundary, which may be undesirable in some empirical applications. The Hölder class addresses this problem by bounding the second derivative globally. In particular, it assumes that f lies in the class of functions

$$\mathcal{F}_{\text{Hölder}}(M) = \{f_+ - f_- : f_+ \in \mathcal{F}_H(M; [c_0, \infty)), f_- \in \mathcal{F}_H(M; (-\infty, c_0))\},$$

where

$$\mathcal{F}_H(M; \mathcal{X}) = \{f : |f'(x) - f'(y)| \leq M|x - y| \text{ } x, y \in \mathcal{X}\}.$$

The smoothness class is specified using the option `sclass`. CIs around a local linear estimator with bandwidth that equals to 10 on either side of the cutoff when the parameter space is given by a Taylor and Hölder smoothness class, respectively, with $M = 0.1$:

```
RDHonest(voteshare ~ margin, data = lee08, kern = "uniform",
  M = 0.1, hp = 10, sclass = "T")
#> Call:
#> RDHonest(formula = voteshare ~ margin, data = lee08, M = 0.1,
#>   kern = "uniform", hp = 10, sclass = "T")
#>
#>
#> Inference by se.method:
#>   Estimate Maximum Bias Std. Error
#> nn  6.05677      3.78224    1.19053
#>
#> Confidence intervals:
#> nn  (0.316293, 11.7973), (0.316293, Inf), (-Inf, 11.7973)
#>
#> Bandwidth below cutoff: 10
#> Bandwidth above cutoff: 10 (Bandwidths are the same)
#> Number of effective observations: 292.325
RDHonest(voteshare ~ margin, data = lee08, kern = "uniform",
  M = 0.1, hp = 10, sclass = "H")
#> Call:
#> RDHonest(formula = voteshare ~ margin, data = lee08, M = 0.1,
#>   kern = "uniform", hp = 10, sclass = "H")
#>
#>
#> Inference by se.method:
#>   Estimate Maximum Bias Std. Error
#> nn  6.05677      1.72377    1.19053
#>
#> Confidence intervals:
#> nn  (2.37473, 9.73882), (2.37476, Inf), (-Inf, 9.73878)
#>
#> Bandwidth below cutoff: 10
#> Bandwidth above cutoff: 10 (Bandwidths are the same)
#> Number of effective observations: 292.325
```

The confidence intervals use the nearest-neighbor method to estimate the standard error by default (this can be changed using the option `se.method`, see help file). The package reports two-sided as well one-sided CIs (with lower as well as upper limit) by default.

Instead of specifying a bandwidth, one can just specify the smoothness class and smoothness constant M , and the bandwidth will be chosen optimally for a given optimality criterion:

```
RDHonest(voteshare ~ margin, data = lee08, kern = "triangular",
  M = 0.1, opt.criterion = "MSE", sclass = "H")
#> Call:
#> RDHonest(formula = voteshare ~ margin, data = lee08, M = 0.1,
#>   kern = "triangular", opt.criterion = "MSE", sclass = "H")
#>
#>
#> Inference by se.method:
#>   Estimate Maximum Bias Std. Error
#> nn  5.93665      0.832259    1.29442
```

```

#>
#> Confidence intervals:
#> nn (2.95483, 8.91847), (2.97526, Inf), (-Inf, 8.89804)
#>
#> Bandwidth below cutoff: 8.84851
#> Bandwidth above cutoff: 8.84851 (Bandwidths are the same)
#> Number of effective observations: 213.463
## Choose bws optimal for length of CI, allowing for
## different bws on either side of cutoff
RDHonest(voteshare ~ margin, data = lee08, kern = "triangular",
  M = 0.1, opt.criterion = "FLCI", sclass = "H", bw.equal = FALSE)
#> Call:
#> RDHonest(formula = voteshare ~ margin, data = lee08, M = 0.1,
#> kern = "triangular", opt.criterion = "FLCI", bw.equal = FALSE,
#> sclass = "H")
#>
#>
#> Inference by se.method:
#> Estimate Maximum Bias Std. Error
#> nn 5.96024 0.880966 1.27649
#>
#> Confidence intervals:
#> nn (2.9647, 8.95578), (2.97964, Inf), (-Inf, 8.94084)
#>
#> Bandwidth below cutoff: 8.80412
#> Bandwidth above cutoff: 9.38041 (Bandwidths are different)
#> Number of effective observations: 220.335

```

It is also possible to compute the optimal bandwidths directly using the function `RDOptBW`

```

RDOptBW(voteshare ~ margin, data = lee08, kern = "triangular",
  M = 0.1, opt.criterion = "MSE", sclass = "H")
#> Call:
#> RDOptBW(formula = voteshare ~ margin, data = lee08, M = 0.1,
#> kern = "triangular", opt.criterion = "MSE", sclass = "H")
#>
#>
#> Bandwidth below cutoff: 8.84851
#> Bandwidth above cutoff: 8.84851 (Bandwidths are the same)

```

Inference when running variable is discrete

The confidence intervals described above can also be used when the running variable is discrete, with G support points: their construction makes no assumptions on the nature of the running variable (see Section 5.1 in Kolesár and Rothe (2017) for more detailed discussion).

As an example, consider the Oreopoulos (2006) data, in which the running variable is age in years:

```

## Replicate Table 2, column (10)
RDHonest(log(earnings) ~ yearat14, cutoff = 1947, data = cghs,
  kern = "uniform", M = 0.04, opt.criterion = "FLCI",
  sclass = "H")
#> Call:
#>

```

```

#> RDHonest(formula = log(earnings) ~ yearat14, data = cghs, cutoff = 1947,
#>     M = 0.04, kern = "uniform", opt.criterion = "FLCI", sclass = "H")
#>
#>
#> Inference by se.method:
#>   Estimate Maximum Bias Std. Error
#> nn 0.0790946    0.0473659  0.0678409
#>
#> Confidence intervals:
#> nn  (-0.0806132, 0.238802), (-0.0798596, Inf), (-Inf, 0.238049)
#>
#> Bandwidth below cutoff: 2
#> Bandwidth above cutoff: 2 (Bandwidths are the same)
#> Number of effective observations: 2017.08
## Triangular kernel generally gives tighter CIs
RDHonest(log(earnings) ~ yearat14, cutoff = 1947, data = cghs,
  kern = "triangular", M = 0.04, opt.criterion = "FLCI",
  sclass = "H")
#> Call:
#>
#> RDHonest(formula = log(earnings) ~ yearat14, data = cghs, cutoff = 1947,
#>     M = 0.04, kern = "triangular", opt.criterion = "FLCI", sclass = "H")
#>
#>
#> Inference by se.method:
#>   Estimate Maximum Bias Std. Error
#> nn 0.0732707    0.0594767  0.0563895
#>
#> Confidence intervals:
#> nn  (-0.0790059, 0.225547), (-0.0789585, Inf), (-Inf, 0.2255)
#>
#> Bandwidth below cutoff: 3.20207
#> Bandwidth above cutoff: 3.20207 (Bandwidths are the same)
#> Number of effective observations: 2265.83

```

In addition, the package provides function `RDHonestBME` that calculates honest confidence intervals under the assumption that the specification bias at zero is no worse at the cutoff than away from the cutoff as in Section 5.2 in Kolesár and Rothe (2017).

```

## Replicate Table 2, column (6), run local linear
## regression (order=1) with a uniform kernel (other
## kernels are not yet implemented)
RDHonestBME(log(earnings) ~ yearat14, cutoff = 1947, data = cghs,
  hp = 3, order = 1)
#> Call:
#> RDHonestBME(formula = log(earnings) ~ yearat14, data = cghs,
#>     cutoff = 1947, hp = 3, order = 1)
#>
#>
#> Confidence intervals:
#> (-0.0696559, 0.201989)

```

Let us describe the implementation of the variance estimator $\hat{V}(W)$ used to construct the CI as described in Section 5.2 in Kolesár and Rothe (2017). Suppose the point estimate is given by the first element of

the regression of the outcome y_i on $m(x_i)$. For instance, local linear regression with uniform kernel and bandwidth h corresponds to $m(x) = I(|x| \leq h) \cdot (I(x > c_0), 1, x, x \cdot I(x > c_0))'$. Let $\theta = Q^{-1}E[m(x_i)y_i]$, where $Q = E[m(x_i)m(x_i)']$, denote the estimand for this regression (treating the bandwidth as fixed), and let $\delta(x) = f(x) - m(x)'\theta$ denote the specification error at x . The RD estimate is given by first element of the least squares estimator $\hat{\theta} = \hat{Q}^{-1} \sum_i m(x_i)y_i$, where $\hat{Q} = \sum_i m(x_i)m(x_i)'$.

Let $w(x_i)$ denote a vector of indicator (dummy) variables for all support points of x_i within distance h of the cutoff, so that $\mu(x_g)$, where x_g is the g th support point of x_i , is given by the g th element of the regression estimand $S^{-1}E[w(x_i)y_i]$, where $S = E[w(x_i)w(x_i)']$. Let $\hat{\mu} = \hat{S}^{-1} \sum_i w(x_i)y_i$, where $\hat{S} = \sum_i w(x_i)w(x_i)'$ denote the least squares estimator. Then an estimate of $(\delta(x_1), \dots, \delta(x_G))'$ is given by $\hat{\delta}$, the vector with elements $\hat{\mu}_g - x_g\hat{\theta}$.

By standard regression results, the asymptotic distribution of $\hat{\theta}$ and $\hat{\mu}$ is given by

$$\sqrt{n} \begin{pmatrix} \hat{\theta} - \theta \\ \hat{\mu} - \mu \end{pmatrix} \xrightarrow{d} \mathcal{N}(0, \Omega),$$

where

$$\Omega = \begin{pmatrix} Q^{-1}E[(\epsilon_i^2 + \delta(x_i)^2)m(x_i)m(x_i)']Q^{-1} & Q^{-1}E[\epsilon_i^2 m(x_i)w(x_i)']S^{-1} \\ S^{-1}E[\epsilon_i^2 w(x_i)m(x_i)']Q^{-1} & S^{-1}E[\epsilon_i^2 w(x_i)w(x_i)']S^{-1} \end{pmatrix}.$$

Let \hat{u}_i denote the regression residual from the regression of y_i on $m(x_i)$, and let $\hat{\epsilon}_i$ denote the regression residuals from the regression of y_i on $w(x_i)$. Then a consistent estimator of the asymptotic variance Ω is given by

$$\hat{\Omega} = n \sum_i T_i T_i', \quad T_i' = (\hat{u}_i m(x_i)' \hat{Q}^{-1} \quad \hat{\epsilon}_i w(x_i)' \hat{S}^{-1}).$$

Note that the upper left block and lower right block correspond simply to the Eicker-Huber-White estimators of the asymptotic variance of $\hat{\theta}$ and $\hat{\mu}$. By the delta method, a consistent estimator of the asymptotic variance of $(\hat{\delta}, \hat{\theta}_1)$ is given by

$$\hat{\Sigma} = \begin{pmatrix} -X & I \\ e_1' & 0 \end{pmatrix} \hat{\Omega} \begin{pmatrix} -X & I \\ e_1' & 0 \end{pmatrix}',$$

where X is a matrix with g th row equal to x_g' , and e_1 is the first unit vector.

Recall that in the notation of Kolesár and Rothe (2017), $W = (g^-, g^+, s^-, s^+)$, and g^+ and g^- are such that $x_{g^-} < c_0 \leq x_{g^+}$, and $s^+, s^- \in \{-1, 1\}$. An upper limit for a right-sided CI for $\theta_1 + b(W)$ is then given by

$$\hat{\theta}_1 + s^+ \hat{\delta}(x_{g^+}) + s^- \hat{\delta}(x_{g^-}) + z_{1-\alpha} \hat{V}(W),$$

where $\hat{V}(W) = a(W)' \hat{\Sigma} a(W)$, and $a(W) \in \mathbb{R}^{G_h+1}$ denotes a vector with the g_- th element equal to s^- , $(G_h^- + g_+)$ th element equal to s^+ , the last element equal to one, and the remaining elements equal to zero. The rest of the construction then follows the description in Section 5.2 in Kolesár and Rothe (2017).

Optimal inference

For the second-order Taylor smoothness class, the function `RDHonest`, with `kernel="optimal"`, computes finite-sample optimal estimators and confidence intervals, as described in Section 2.2 in Armstrong and Kolesár (2017). This typically yields tighter CIs. Comparing the lengths of two-sided CIs with optimally

chosen bandwidths, using Silverman's rule of thumb to estimate the preliminary variance estimate used to compute optimal bandwidths:

```
2 * RDHonest(voteshare ~ margin, data = lee08, kern = "optimal",
  M = 0.1, opt.criterion = "FLCI", se.initial = "Silverman",
  se.method = "nn")$hl
#>      nn
#> 6.29408
2 * RDHonest(voteshare ~ margin, data = lee08, kern = "triangular",
  M = 0.1, opt.criterion = "FLCI", se.initial = "Silverman",
  se.method = "nn", sclass = "T")$hl
#>      nn
#> 6.64827
```

Specification testing

The package also implements lower-bound estimates for the smoothness constant M for the Taylor and Hölder smoothness class, as described in the supplements to Kolesár and Rothe (2017) and Armstrong and Kolesár (2017)

```
## Add variance estimate to the lee data so that the
## RDSmoothnessBound function doesn't have to compute
## them each time
dl <- RDHonest::RDPrelimVar(dl, se.initial = "NN")
### Only use three point-average for averages of a 100
### points closest to cutoff, and report results
### separately for points above and below cutoff
RDSmoothnessBound(dl, s = 100, separate = TRUE, multiple = FALSE,
  sclass = "T")
#>
#> Smoothness bound estimate using observations above cutoff:
#> Estimate: 0.172723, Lower CI: [0, Inf)
#>
#> Delta: 0.183302, sd=0.179281
#> E_n[f(x_1)]: 53.1291, I1=[0.0112861, 1.74936]
#> E_n[f(x_2)]: 53.1468, I2=[1.75187, 3.43698]
#> E_n[f(x_3)]: 56.1762, I3=[3.44594, 4.89676]
#>
#> Smoothness bound estimate using observations below cutoff:
#> Estimate: 0.333754, Lower CI: [0.138838, Inf)
#>
#> Delta: -0.333754, sd=0.118477
#> E_n[f(x_1)]: 44.9589, I1=[0.0308186, 1.95308]
#> E_n[f(x_2)]: 46.6076, I2=[2.00265, 3.75439]
#> E_n[f(x_3)]: 41.7689, I3=[3.78509, 5.49484]
### Pool estimates based on observations below and above
### cutoff, and use three-point averages over the entire
### support of the running variable
RDSmoothnessBound(dl, s = 100, separate = FALSE, multiple = TRUE,
  sclass = "H")
#>
#> Smoothness bound estimate:
#> Estimate: 0.229375, Lower CI: [0.0250072, Inf)
```

```
#>
#> Delta: -2.15055, sd=0.763411
#> E_n[f(x_1)]: 44.9589, I1=[0.0308186, 1.95308]
#> E_n[f(x_2)]: 46.6076, I2=[2.00265, 3.75439]
#> E_n[f(x_3)]: 41.7689, I3=[3.78509, 5.49484]
```

References

- Armstrong, Timothy B., and Michal Kolesár. 2016. “Simple and Honest Confidence Intervals in Nonparametric Regression.”
- . 2017. “Optimal Inference in a Class of Regression Models.”
- Kolesár, Michal, and Christoph Rothe. 2017. “Inference in Regression Discontinuity Designs with a Discrete Running Variable.”
- Lalive, Rafael. 2008. “How Do Extended Benefits Affect Unemployment Duration? A Regression Discontinuity Approach.” *Journal of Econometrics* 142 (2): 785–806.
- Lee, David S. 2008. “Randomized Experiments from Non-Random Selection in U.S. House Elections.” *Journal of Econometrics* 142 (2): 675–97.
- Ludwig, Jens, and Douglas L. Miller. 2007. “Does Head Start Improve Children’s Life Chances? Evidence from a Regression Discontinuity Design.” *Quarterly Journal of Economics* 122 (1): 159–208.
- Oreopoulos, Philip. 2006. “Estimating Average and Local Average Treatment Effects When Compulsory Education Schooling Laws Really Matter.” *American Economic Review* 96 (1): 152–75.