

NBER WORKING PAPER SERIES

PLRD: PARTIALLY LINEAR REGRESSION DISCONTINUITY INFERENCE

Aditya Ghosh
Guido Imbens
Stefan Wager

Working Paper 35699
<http://www.nber.org/papers/w35699>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2026

We are grateful to Evan Munro for sharing results on the behavior of existing methods for RDDs under generative adversarial simulation following (Athey et al., 2024). We are also grateful to Tim Armstrong, P. Aronow, Max Farrell, Paul Goldsmith-Pinkham, Calumn Hamilton, Michal Kolesár, Christoph Rothe and Rocío Titiunik for insightful conversations and for comments on a previous version of this paper. This work was supported by the National Science Foundation under grant SES-2242876, the Office of Naval Research under grants N00014-17-1-2131 and N00014-19-1-2468, and a gift from Amazon. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w35699>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Aditya Ghosh, Guido Imbens, and Stefan Wager. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

PLRD: Partially Linear Regression Discontinuity Inference
Aditya Ghosh, Guido Imbens, and Stefan Wager
NBER Working Paper No. 35699
September 2026
JEL No. C01, C14, C20

ABSTRACT

Regression discontinuity designs have become one of the most popular research designs in empirical economics. We argue, however, that the widely used approaches to building confidence intervals in regression discontinuity designs often exhibit suboptimal behavior in practice. We propose a new estimator, the partially linear regression discontinuity (PLRD) estimator that, in set of a simulation studies carefully calibrated to twelve high-profile applications of regression discontinuity designs, has substantially lower estimation error than available comparison methods. Throughout our experiments, the confidence intervals built using PLRD are both valid and short. We also provide large-sample guarantees for PLRD. Our simulation study serves as a general template for how new econometric methods can be credibly evaluated relative to the existing alternatives by constructing simulation designs that generate synthetic data indistinguishable from the original data using the Wasserstein generative adversarial network methodology.

Aditya Ghosh
Stanford University
Department of Statistics
ghoshadi@stanford.edu

Stefan Wager
Stanford University
Graduate School of Business
swager@stanford.edu

Guido Imbens
Stanford University Graduate
School of Business
Department of Economics
and NBER
Imbens@stanford.edu

PLRD: Partially Linear Regression Discontinuity Inference

Aditya Ghosh
ghoshadi@stanford.edu

Guido Imbens
imbens@stanford.edu

Stefan Wager
swager@stanford.edu

Stanford University

August 25, 2026

Abstract

Regression discontinuity designs have become one of the most popular research designs in empirical economics. We argue, however, that the widely used approaches to building confidence intervals in regression discontinuity designs often exhibit suboptimal behavior in practice. We propose a new estimator, the partially linear regression discontinuity (PLRD) estimator that, in set of a simulation studies carefully calibrated to twelve high-profile applications of regression discontinuity designs, has substantially lower estimation error than available comparison methods. Throughout our experiments, the confidence intervals built using PLRD are both valid and short. We also provide large-sample guarantees for PLRD. Our simulation study serves as a general template for how new econometric methods can be credibly evaluated relative to the existing alternatives by constructing simulation designs that generate synthetic data indistinguishable from the original data using the Wasserstein generative adversarial network methodology.

1 Introduction

The fraction of empirical NBER working papers using regression discontinuity designs (RDDs) increased from nearly zero in 1995 to between 15% and 20% by 2015 (Currie et al., 2020, Figure 4B).¹ Despite the ubiquity of RDDs in applied work, however, there are still open questions as to econometric best practices for inference in RDDs. In particular, we find that default implementations of widely used local-polynomial methods can fail to achieve nominal coverage in realistic settings.

We assess the performance of three commonly used methods by evaluating them on simulations calibrated to twelve high-profile applications of RDDs. First, we consider conventional local linear regression (Hahn et al., 2001) with Imbens–Kalyanamaraman bandwidth (Imbens and Kalyanamaraman, 2012). We also consider two more recent proposals (Calonico et al., 2014; Armstrong and Kolesár, 2018), as implemented in the R functions `rdrobust` (Calonico et al., 2014) and `RDHonest` (Armstrong and Kolesár, 2018). These methods are representative of the current state of the art for inference in RDDs (Stommes et al., 2023), and both offer automated end-to-end confidence interval constructions that can directly be applied to raw data. The method `rdrobust` (Calonico et al., 2014) is a dominant approach to regression-discontinuity inference in current econometric practice.² The method starts by running a local linear regression (Hahn et al., 2001), and then applies a 2nd-order bias correction to mitigate curvature bias. Meanwhile, the method `RDHonest` (Armstrong

¹Up-to-date numbers are maintained at <https://paulgp.com/econlit-pipeline> (Goldsmith-Pinkham, 2024).

²This is the only general purpose regression discontinuity package listed in the “Econometrics Task View” for highly used R packages on CRAN (<https://cran.r-project.org/web/views/Econometrics.html>); and both the R (Calonico et al., 2015) and Stata (Calonico et al., 2017) versions of the function are widely used and the original paper (Calonico et al., 2014) cited is over 5000 times as of now.

Dataset	conventional			rdrobust			RDHonest			plrd		
	cvrg	width	rmse	cvrg	width	rmse	cvrg	width	rmse	cvrg	width	rmse
Lee (2008)	94.6%	0.036	0.010	94.9%	0.043	0.011	97.3%	0.055	0.012	98.4%	0.044	0.009
Cattaneo et al. (2015)	94.3%	4.955	1.316	94.9%	5.828	1.516	98.0%	6.767	1.464	96.2%	3.865	0.937
Ludwig and Miller (2007)	94.2%	5.275	1.599	94.3%	6.247	1.883	97.5%	7.017	1.948	96.5%	3.746	1.144
Meyersson (2014)	93.9%	6.309	1.732	94.3%	7.436	1.998	96.8%	7.941	1.858	96.7%	4.682	1.159
Matsudaira (2008) (reading)	82.3%	0.050	0.018	90.1%	0.059	0.018	97.9%	0.104	0.022	96.7%	0.059	0.014
Matsudaira (2008) (math)	91.7%	0.044	0.013	93.1%	0.053	0.015	98.5%	0.093	0.019	97.6%	0.051	0.011
Jacob and Lefgren (2004) (reading)	93.7%	0.069	0.020	94.0%	0.083	0.023	97.7%	0.116	0.026	97.2%	0.078	0.019
Jacob and Lefgren (2004) (math)	94.2%	0.058	0.016	95.2%	0.070	0.018	97.9%	0.122	0.026	97.6%	0.060	0.013
Oreopoulos (2006)	93.9%	0.191	0.055	95.0%	0.258	0.072	99.0%	0.304	0.060	95.9%	0.124	0.031
Akhtari et al. (2022) (4th grade)	94.3%	0.104	0.028	95.0%	0.119	0.031	97.3%	0.135	0.031	96.8%	0.096	0.022
Akhtari et al. (2022) (8th grade)	94.0%	0.109	0.030	94.6%	0.127	0.034	97.2%	0.132	0.031	94.9%	0.071	0.018
Akhtari et al. (2022) (headmaster)	94.8%	0.059	0.016	95.1%	0.069	0.018	97.2%	0.072	0.017	95.9%	0.040	0.010

Table 1: Empirical coverage (cvrg), average width, and root mean squared error (RMSE) of 95% confidence intervals from different methods, on simulated RDDs calibrated to twelve applications. We generate our simulation distributions by training a Wasserstein Generative Adversarial Neural Network (WGAN) on the original data (Athey et al., 2024). We report results for conventional local linear regression with Imbens–Kalyanamaraman bandwidth (Imbens and Kalyanamaraman, 2012); robust bias-corrected (rdrobust) confidence intervals (Calonico et al., 2014); bias-aware (RDHonest) confidence intervals (Armstrong and Kolesár, 2018); and our proposed plrd method. We highlight (in green) the entries whose coverage is not statistically significantly below nominal 95% coverage (at the 5% level using a one-sided binomial test) and (in blue) the entries with the lowest RMSE. All results are aggregated across 10,000 simulation replications.

and Kolesár, 2018) explicitly inflates the width of local linear regression confidence intervals to accommodate worst-case plausible curvature bias (rather than trying to correct for this bias).

We use generative adversarial neural networks (Goodfellow et al., 2014) to generate simulation instances that are reflective of real-world datasets (Athey et al., 2024). Specifically, we calibrate our evaluation to eight seminal regression discontinuity studies (Lee, 2008; Cattaneo et al., 2015; Ludwig and Miller, 2007; Meyersson, 2014; Matsudaira, 2008; Jacob and Lefgren, 2004; Oreopoulos, 2006; Akhtari et al., 2022).³ The generative adversarial learning approach yields simulators whose data-generating distribution closely mimics the moments and idiosyncrasies of the original datasets; see Figure 1 for an illustration. The generated simulators do not use explicit parametric models and do not enforce any built-in smoothness assumptions, thus bringing us towards a neutral comparison ground for evaluating methods that were motivated under different assumptions (Athey et al., 2024).

We summarize results for 95% confidence intervals in Table 1; see Section 4 for details on the simulation design. A first observation is that the conventional local linear regression intervals systematically undercover the true parameter. The method rdrobust⁴ achieves better—although not perfect—coverage, with nominal rates in 5 out of the 10 settings. However, it does so at the expense of wider confidence intervals than conventional local linear regression (18% to 36% wider in our applications). RDHonest⁵ in turn achieves reliable coverage throughout all designs, but at

³Some of the studies have more than one plausible outcome variables, leading to a total of twelve applications.

⁴The method rdrobust offers two options for building confidence intervals: A basic bias-corrected method which inherits the standard error estimator from the original local linear regression (and is justified by asymptotic arguments), and a robust method that incorporates a finite-sample variance inflation to account for the bias correction. The latter is the recommended one in the paper and this is the one we consider here.

⁵In principle, the method RDHonest asks the user to specify a worst-case bound on the curvature of the underlying conditional response functions up front. However, the method also has a default fall-back option where curvature




Figure 1: The original data from the Lee(2008) dataset (Lee, 2008), as well as a simulated dataset produced using a WGAN trained to the original data (Athey et al., 2024).




Figure 2: The relative RMSE and average width of 95% confidence intervals from different methods as in Table 1, relative to the minimum across the methods for each empirical application.

the cost of still-wider intervals: The median ratio between the width of these intervals and the **rdrobust** ones varies between 1.06 to 1.76 in our simulations, and between 1.25 to 2.11 relative to the conventional ones. Furthermore, these confidence intervals are often excessively conservative, thus suggesting that some of this additional width is not strictly needed.⁶ These qualitative findings do not appear to be specific to our simulation designs; for example, in Table 3 we observe similar coverage properties when applying the method to pure noise with no relationship between the outcome and the running variable.

These semi-synthetic simulation findings leave open a pragmatic question: Is it possible to design automated confidence intervals that reliably achieve nominal coverage in realistic settings, while remaining competitive with the conventional intervals in terms of interval width? We emphasize here the “realistic” qualifier. Even if it is not feasible to achieve uniformly better properties, it may be that most real-world applications are sufficiently smooth to allow for nominal coverage and narrower confidence intervals than the existing methods.

The main contribution of this paper is a method, partially linear regression discontinuity inference (PLRD), that provides an affirmative answer to this question—at least in the context of the

bounds are derived using a provided rule-of-thumb (Armstrong and Kolesár, 2020); in our experiments we use this default.

⁶In several settings, **RDHonest** achieves 98% coverage or above—and this may come at a substantial cost in width. For example, in the classical setting with no bias adjustment, Gaussian 98% confidence intervals are 18.7% wider than the conventional 95% intervals.

numerical settings we consider. As described further below, algorithmically, the PLRD estimator adapts `RDHonest` (Armstrong and Kolesár, 2018) to accommodate a partially linear treatment effect model; conceptually, however, our full approach fits within an asymptotic econometric framework as used for `rdrobust` (Calonico et al., 2014). We implement our method in the package `plrd` for R (R Core Team, 2023), which provides end-to-end functionality for regression discontinuity inference without relying on the user to specify hard-to-choose tuning parameters. As illustrated in Tables 1 and 3, our proposed `plrd` method achieves nominal and near-exact coverage throughout. It also has the most accurate point estimates across all settings considered here (Figure 2).

We emphasize that our semi-synthetic simulation setup, central to argument, comes with two qualifications. First, it relies on the availability of datasets where regression discontinuity methods are applicable. We apply it to twelve such data sets, but we encourage researchers to apply it to additional ones. Second, it relies on our generative adversarial nets yielding realistic data generating processes given the provided datasets. Well-known results in non-parametric inference imply that, fixing a sample size and given any dataset where standard regression-discontinuity estimators work well, one can construct another adversarial dataset that is statistically indistinguishable from the first but where the same estimators fail to get coverage (Low, 1997). Our approach to evaluation implicitly relies on an assumption that the real datasets we are calibrated to aren’t adversarial in this sense. Again, we encourage researchers to explore the sensitivity of our results to other generative models that are tied to existing data sets. Finally, our estimation process involves some choices of tuning parameters for which different researchers may make different choices; we seek to mitigate concerns of researcher bias by using a single set of tuning parameters across all datasets.

2 The PLRD Estimator

Consider the sharp RDD formalized using potential outcomes (Imbens and Lemieux, 2008). For each unit $i = 1, \dots, n$, we observe IID-sampled pairs (X_i, Y_i) , where $X_i \in \mathbb{R}$ is the running variable and $Y_i \in \mathbb{R}$ is the outcome of interest. Following the sharp discontinuity design, we assume there exists a cutoff $c \in \mathbb{R}$ such that treatment is assigned as $W_i = \mathbf{1}\{X_i \geq c\}$. We posit potential outcomes $\{Y_i(0), Y_i(1)\}$ such that $Y_i = Y_i(W_i)$, and our goal is to estimate the conditional average treatment effect at the cutoff, $\tau_c := \mu_1(c) - \mu_0(c)$ where $\mu_w(x) := \mathbb{E}[Y_i(w) | X_i = x]$.

It is well known that, provided the $\mu_w(x)$ are continuous in x and that the running variable has continuous support around c , our target τ_c is identified as a discontinuity in the conditional response surface at $X_i = c$ (Hahn et al., 2001):

$$\tau_c = \lim_{x \downarrow c} \mathbb{E}[Y_i | X_i = x] - \lim_{x \uparrow c} \mathbb{E}[Y_i | X_i = x]. \quad (1)$$

In order to obtain practical estimation or inference results for τ_c , though, one needs to make further regularity assumptions on the $\mu_w(x)$; our PLRD method relies on the following two assumptions.

Assumption 1. $\mu_w(x)$ are 3 times continuously differentiable in x .

Assumption 2. The conditional average treatment effect function $\tau(x) = \mu_1(x) - \mu_0(x)$ has a linear dependence in x , i.e., $\tau(x) = \tau_c + \beta(x - c)$ for some $\beta \in \mathbb{R}$.

Assumption 1 is a standard smoothness assumption; this is exactly the same as the main smoothness assumption used by (Calonico et al., 2014) to justify `rdrobust`. Assumption 2 implies that the treatment effect has a simpler dependence on the running variable than the baseline effect $\mu_0(x)$. This assumption is well in line with assumptions often made in the literature on heterogeneous treatment effect estimation (e.g., Wager, 2024, Chapter 4); however, it has not typically been used

in the RDD literature to date, which is why we emphasize this assumption in naming our method. Given these two assumptions, PLRD proceeds as follows:

1. Under Assumption 1 there exists a finite constant B such that the second derivatives $\mu_w''(x)$ are B -Lipschitz in a neighborhood of c ; and it is possible to identify such a B from data in the large-sample limit. PLRD starts by estimating B via a cubic regression (Imbens and Kalyanaraman, 2012; Calonico et al., 2014).⁷
2. We form the set of all possible choices of $\mu_0(x)$ and $\mu_1(x)$ consistent with the B -Lipschitz curvature bound estimated above and Assumption 2. Then, we use numerical convex optimization to derive a minimax linear estimator⁸ for τ_c within this class (Armstrong and Kolesár, 2018; Imbens and Wager, 2019).
3. We identify the worst-case bias of our estimator within this class, and construct “bias aware” confidence intervals for τ_c using a construction that simultaneously accounts for the standard error and potential bias of an estimator (Imbens and Manski, 2004).

Our full procedure—including some extra algorithmic details such as cross fitting—is detailed as Algorithm 1. As a pragmatic safeguard against failures of Assumption 2, we start our procedure with a hypothesis test, and if we reject this assumption with high confidence then we proceed with a variant of our method that only relies on Assumption 1; details of this flexible fallback are given in the Supplementary Material.

Our main formal result about the PLRD estimator is that it yields valid confidence intervals under our assumptions in the large-sample limit.

Theorem 1. *Consider a sharp regression discontinuity design under Assumptions 1 and 2, where the running variable X_i has a strictly positive density at the threshold c , $\sigma_i^2 = \text{Var}(Y_i | X_i)$ is bounded away from 0, i.e., $\sigma_i \geq \sigma_{\min}$ for some $\sigma_{\min} > 0$, and $\mathbb{E}[|Y_i - \mu_{W_i}(X_i)|^q | X_i = x] \leq C$ uniformly over all x , for some $q > 2$ and $C \in (0, \infty)$. Suppose furthermore that we execute the initial cubic regression used to pick B over a pilot bandwidth of width ℓ with*

$$\lim_{n \rightarrow \infty} n^{1/12} \ell = 0, \quad \lim_{n \rightarrow \infty} n^{1/7} \ell = +\infty, \quad (6)$$

and choose any $\varepsilon > 0$ in Algorithm 1. Then, level- $(1 - \alpha)$ PLRD intervals $\mathcal{I}_\alpha$ are asymptotically valid: For any $0 < \alpha < 1$,

$$\liminf_{n \rightarrow \infty} \mathbb{P}[\tau_c \in \mathcal{I}_\alpha] \geq 1 - \alpha. \quad (7)$$

We note that our algorithm does in principle require two tuning parameters, namely the pilot bandwidth ℓ used to run the initial cubic regression to choose B and the lower bound ε for $\hat{B}$. We note, however, that our algorithmic performance is largely insensitive to these choices (and simple defaults appear to provide good performance in practice); and similar objects are also required for other methods (Calonico et al., 2014; Imbens and Kalyanaraman, 2012). First, regarding ℓ , this bandwidth can in practice be chosen simply to be “large”: We only need ℓ to decay faster than $n^{-1/12}$, whereas the mean-squared error optimal bandwidth for local linear regression in our setting

⁷We use sample splitting Algorithm 1 to justify technical arguments underlying formal results given in Section 5. It is possible, however, that an alternative argument could be used to justify similar results without sample splitting; see, e.g., (Armstrong and Kolesár, 2020, Appendix E) for results of this type.

⁸Any estimator of the form $\hat{\tau} = \sum_{i=1}^n \gamma_i Y_i$ where the weights γ_i only depend on the realization of the running variables X_i is called a “linear” estimator; the minimax linear estimator is the estimator of this form with the minimax mean-squared error conditionally on the X_i . The minimax linear estimation approach was introduced by (Donoho, 1994), and first applied to RDDs by (Armstrong and Kolesár, 2018). For a textbook overview of the minimax linear estimation approach to inference for RDDs, see (Wager, 2024, Chapter 8.2).

Algorithm 1 Partially Linear Regression Discontinuity Inference (PLRD)

This algorithm provides asymptotically valid $100(1 - \alpha)\%$ confidence intervals for the conditional average treatment effect τ_c in a sharp RDD. This procedure is implemented in the R package `plrd`.

Input: Data $\{(X_i, Y_i)\}_{i=1}^n$, the cutoff c , and the confidence level $1 - \alpha$ (default is set as 95%).

Tuning Parameters: Window size ℓ for estimating the smoothness parameter B (default is set as the whole range), a lower bound ε for $\hat{B}$ (default is $\varepsilon = \text{SD}[Y] / 100$), and a confidence level $1 - \alpha'$ for curvature detection (default is set as 99.9%).

Output: The PLRD estimator $\hat{\tau}_{\text{PLRD}}$ with associated $1 - \alpha$ confidence interval $\mathcal{I}_\alpha$.

- 1: Discard all observations with $|X_i - c| > \ell$.
- 2: Perform an ANOVA test between cubic polynomials fitted with or without different curvatures at significance level $1 - \alpha'$ to detect whether different curvatures should be used.
- 3: Randomly split the data into two near-equally sized folds, with index sets I_1 and I_2 .
- 4: Estimate σ^2 by running the ordinary least-squares regression of Y_i on the interaction of X_i and W_i for $i \in I_2$. Denote this estimate as $\hat{\sigma}_{(2)}^2$ and record the residuals from this regression as:

$$r_i := Y_i - \hat{\mu}_{W_i}(X_i) \text{ for each } i \in I_2. \quad (2)$$

- 5: **if** the test for different curvatures in Step 2 is *not* rejected **then**
- 6: Use the observations in fold I_2 to estimate the Lipschitz constant B by fitting cubic polynomials with the same curvature as in (9), and denote this estimate as $\hat{B}_{(2)}$.
- 7: Enforce a lower bound $\hat{B}_{(2)} \leftarrow \max\{\hat{B}_{(2)}, \varepsilon\}$.
- 8: Let $\mathcal{R} := \{f : f(0) = f'(0) = f''(0) = 0, f'' \text{ is 1-Lipschitz}\}$, and solve the following optimization problem
via quadratic programming (see the Supplementary Material for implementation details):

$$(\hat{\gamma}_{(1)}, \hat{t}_{(1)}) := \underset{\gamma, t}{\text{minimize}} \left(\hat{B}_{(2)}^2 t^2 + \hat{\sigma}_{(2)}^2 \sum_{i \in I_1} \gamma_i^2 \right) \quad \text{subject to} \quad \sup_{\rho \in \mathcal{R}} \sum_{i \in I_1} \gamma_i \rho(X_i - c) \leq t. \quad (3)$$

- 9: Repeat Steps 6 to 8 with the roles of I_1 and I_2 swapped.
- 10: **else**
- 11: Use the observations in fold I_2 to estimate the Lipschitz constant B by fitting cubic polynomials with different curvatures (see the Supplementary Material), and denote this estimate as $\hat{B}_{(2)}$.
- 12: Enforce a lower bound $\hat{B}_{(2)} \leftarrow \max\{\hat{B}_{(2)}, \varepsilon\}$.
- 13: Let $\mathcal{R} := \{f : f(0) = f'(0) = f''(0) = 0, f'' \text{ is 1-Lipschitz}\}$, and solve the following optimization problem
via quadratic programming (see the Supplementary Material for implementation details):

$$(\hat{\gamma}_{(1)}, \hat{t}_{(1)}) := \underset{\gamma, t}{\text{minimize}} \left(\hat{B}_{(2)}^2 t^2 + \hat{\sigma}_{(2)}^2 \sum_{i \in I_1} \gamma_i^2 \right) \quad \text{subject to} \quad \sup_{\rho_0, \rho_1 \in \mathcal{R}} \sum_{i \in I_1} \gamma_i \rho_{W_i}(X_i - c) \leq t. \quad (4)$$

- 14: Repeat Steps 11 to 13 with the roles of I_1 and I_2 swapped.
- 15: **end if**
- 16: Set $\hat{\gamma}_i = \gamma_{(k), i} / 2$ if $i \in I_k$ ($k = 1, 2$), and $\hat{\gamma}_i = 0$ if $|X_i - c| > \ell$. Compute

$$\hat{\tau}_{\text{PLRD}} = \sum_{i=1}^n \hat{\gamma}_i Y_i \quad \text{and} \quad \hat{s}^2(\hat{\gamma}) = \sum_{i=1}^n \hat{\gamma}_i^2 r_i^2,$$

where the r_i are the residuals from the least squares regressions in (2).

- 17: Report the following $100(1 - \alpha)\%$ confidence interval for τ_c :

$$\mathcal{I}_\alpha := [\hat{\tau}_{\text{PLRD}} - \hat{h}_\alpha, \hat{\tau}_{\text{PLRD}} + \hat{h}_\alpha], \quad \hat{h}_\alpha := \min \{h : \mathbb{P}(|b + \hat{s}(\hat{\gamma})Z| \leq h) \geq 1 - \alpha \text{ for all } |b| \leq \hat{b}\}, \quad (5)$$

where $Z \sim \mathcal{N}(0, 1)$, and $\hat{b} := (\hat{B}_{(2)} \hat{t}_{(1)} + \hat{B}_{(1)} \hat{t}_{(2)}) / 2$.

(i.e., under Assumption 1) scales as $n^{-1/7}$ (Cheng et al., 1997). In our implementation, we by default simply set ℓ to contain the full range of the data, and this choice is what we use in all our experiments⁹ (Table 1). As such, practitioners should be able to get reliable results out of our method unless the nature of the relationship between X_i and Y_i changes completely as we get far from the cutoff—and in settings like this it would be possible to pick ℓ by plotting the data. Meanwhile, the lower bound ε for $\hat{B}$ will have no effect in large samples as long as we choose it to be small enough that $\varepsilon < |\mu'''(c)|$.¹⁰

2.1 Discrete running variables

Discrete running variables arise frequently in regression discontinuity applications. Our formal results are stated under the assumption that the running variable has a continuous density bounded away from zero near the cutoff, which does not accommodate the discrete case. In particular, when the running variable is discrete, *Step 1* is no longer guaranteed to consistently recover B under Assumption 1 alone, and thus Theorem 1 does not provide a formal guarantee for the discrete case. However, *Steps 2* and *3* of our procedure remain valid, i.e., given an appropriate B they yield partial identification intervals for τ_c with a guarantee of the type (7) conditional on that bound (Kolesár and Rothe, 2018; Imbens and Wager, 2019). Thus, PLRD is robust to discrete running variables in the following sense: The interval logic (7) is sound as long as the practitioner is willing to commit to a curvature bound—whether from the pilot regression or from background knowledge.

Table 1 illustrates the empirical performance of PLRD when the running variable is discrete: Five of our twelve benchmark cases—Matsudaira (math and reading), Jacob & Lefgren (math and reading), and Oreopoulos—use assignment rules based on test-scores, where the running variable takes values on a finite grid. This suggests that the pilot-based curvature estimate continues to behave reasonably even in the case of discrete running variables, where one would need further assumptions to formally justify the data-driven detection of B in Algorithm 1.

3 Relationship to Existing Methods

Regression discontinuity designs were originally introduced by (Thistlethwaite and Campbell, 1960). The modern econometric literature goes back to (Hahn et al., 2001), who emphasized the identification result (1) and recommended estimation τ_c via local linear regression, $\hat{\tau}_{LLR}$. In the following decade, local linear (or polynomial) regression developed into the prevalent approach to inference in RDDs, with confidence intervals produced by applying off-the-shelf tools for heteroskedasticity-robust inference directly to the regression (Imbens and Lemieux, 2008; Lee and Lemieux, 2010).

A limitation of this approach, however, is that it leads to contradictory guidance on the bandwidth choice under a non-parametric specification as given in, *e.g.*, Assumption 1. The accuracy of $\hat{\tau}_{LLR}$ depends delicately on the bandwidth h_n , with a first systematic proposal given in (Imbens and Kalyanaraman, 2012). For choices of h_n that optimize this accuracy, the bias and standard error of $\hat{\tau}_{LLR}$ are of the same order, so off-the-shelf local linear regression inference will not have valid

⁹This choice mirrors (Imbens and Kalyanaraman, 2012, p. 942, *step 2*) who seed their bandwidth-selection procedures by fitting global cubic polynomials to the full dataset; see also (Calonico et al., 2014, S.2.6, *step 0*).

¹⁰One interesting question is what happens at the critical point where $\mu'''(c) = 0$; as, in this case, our constraint does necessarily bind and we fall back to using ε as our curvature bound. The challenge is that, when $\mu'''(c) = 0$, our cubic polynomial estimator $\hat{B}$ (in step 6 of Algorithm 1) converges to 0, but the exact way it does so is left underspecified under Assumption 1. Our use of the lower bound ε guarantees validity of our confidence intervals at the expense of some excess conservativeness when $\mu'''(c) = 0$. We note that existing methods, including (Calonico et al., 2014) and (Imbens and Kalyanaraman, 2012), also do not offer guidance on how to potentially benefit from “superefficiency” in similar critical cases; i.e., our approach is again in line with current best practices.

coverage for this choice of h_n . The practitioner is thus forced to invoke “undersmoothing”, whereby they voluntarily use a non-accuracy-optimizing choice of h_n to justify validity of their inferential procedure.

In a major advance, (Calonico et al., 2014) and (Armstrong and Kolesár, 2018) introduced paradigms for inference in RDDs that formally account for the bias of local linear regression. These methods offer material improvements in the credibility of regression-discontinuity analyses over the earlier practice (Stommes et al., 2023). The main idea in (Calonico et al., 2014) is to estimate and correct for the bias of local linear regression by leveraging higher-order smoothness; we refer to this approach as the “bias-corrected” approach;¹¹ this approach has further been extended in a number of directions (Calonico et al., 2019; Cattaneo and Titiunik, 2022). In contrast, the “bias-aware” approach developed by (Armstrong and Kolesár, 2018) involves conservatively widening intervals to account for worst-case bias (without needing to actually estimate the bias); this approach has also been recently developed further (Kolesár and Rothe, 2018; Imbens and Wager, 2019; Noack and Rothe, 2024).

Interestingly, not only do the above papers under the bias-corrected vs. bias-aware paradigms explore different methodological strategies for inference in RDDs; they also focus on proving formal results under different formal paradigms. (Calonico et al., 2014) and follow-ups focus entirely on asymptotic validity results—as we do here. They fix a data-generating process, and guarantee that they eventually get coverage for this data-generating process as the sample size gets large (as is also done in (7)). In contrast, most papers written under the bias-aware paradigm specify a quantitative smoothness class for the $\mu_w(x)$ (e.g., all functions with a second derivative bounded by a pre-specified constant B), and then seek guarantees that hold in finite samples and uniformly over this whole class. Getting uniform guarantees in finite samples of course requires stronger assumptions than getting pointwise asymptotic guarantees; and for this reason papers associated with the bias-aware paradigm often start by making stronger assumptions than those associated with the bias-corrected paradigm. We emphasize, however, that this is due to the sought formal guarantees; and it’s not that bias-aware methods inherently require stronger assumptions to work than bias-corrected methods.

Our proposed method, PLRD, straddles these two literatures. Our confidence intervals are rooted in the bias-aware paradigm (Armstrong and Kolesár, 2018), and we use numerical convex optimization to construct our estimator (Imbens and Wager, 2019). In contrast, our formal results are focused on pointwise asymptotics, and the assumptions and guarantee type given in our Theorem 1 mirror those in Theorem 1 of (Calonico et al., 2014). One notable paper that also straddles these literatures is (Armstrong and Kolesár, 2020), who propose bias-aware confidence intervals for local linear regression with pointwise asymptotic guarantees. In particular, they show that if one runs local linear regression with a mean-square-error optimal bandwidth choice, then under generic conditions we can build valid, bias-aware 95% confidence intervals for τ_c as $\hat{\tau}_{\text{LLR}} \pm 2.18 \times \text{standard errors}$ rather than the conventional $\hat{\tau}_{\text{LLR}} \pm 1.96 \times \text{standard errors}$. We provide a comparison of PLRD to this approach in the Supplementary Material.

Wrapping up, PLRD can best be understood within the context of earlier developments in the bias-aware paradigm; and in fact *Steps (2) and (3)* of our algorithmic outline fall squarely within the roadmap for minimax linear inference via numerical convex optimization as in (Imbens and Wager, 2019). However, in a deviation from existing work:

- We rely on higher-order smoothness and a preliminary cubic fit to initialize tuning parameters of our algorithm; in (Armstrong and Kolesár, 2018), (Kolesár and Rothe, 2018) and (Imbens

¹¹When reporting results using the bias-correct approach, we use robust standard error estimates as recommended by (Calonico et al., 2014) and as implemented in the package `rdrobust` (Calonico et al., 2015, 2017).

and Wager, 2019), a parameter analogous to the B derived in our *Step (1)* would need to be pre-specified using subject matter knowledge or chosen using a heuristic.

- Because of this data-driven choice of B in *Step (1)*, we are not able to provide uniform, finite-sample inferential guarantees;¹² thus, like in (Calonico et al., 2014), we focus on pointwise asymptotic inferential statements.
- The specific worst-case class we use in our *Step (2)*, i.e., with B -Lipschitz curvature bounds and a linear treatment effect assumption, has not been previously considered in the literature. The most standard choice in applications of (Armstrong and Kolesár, 2018), (Kolesár and Rothe, 2018) and/or (Imbens and Wager, 2019) has been to only assume that the $\mu_w(x)$ have second derivatives bounded by M .

Regarding this last point, the earlier focus on the M -bounded curvature class appears to have been largely guided by intuitive simplicity of this class rather than any conceptual or empirical arguments. Given the empirical performance of PLRD on calibrated simulations, we suggest that our B -Lipschitz curvature class with linear treatment effects may be a good default choice in many applications.

To further examine the effect of different choices on the behavior of our estimator, we conduct an ablation analysis, reported in the Supplementary Material, comparing four alternatives to PLRD: Bias-aware intervals for local linear regression with the inflated critical value of (Armstrong and Kolesár, 2020); the analogous construction for local quadratic regression; PLRD with Assumption 2 removed while retaining the B -Lipschitz curvature structure (so that separate curvature bounds are allowed on each side of the threshold); and PLRD with Assumption 1 relaxed to second-order smoothness, which essentially recovers the *optrdd* construction of (Imbens and Wager, 2019) augmented with Assumption 2. The ablation analysis empirically demonstrates that the efficiency of PLRD is a joint consequence of both Assumptions 1 and 2.

4 Empirical Study

We now describe in more detail the experiment presented in the introduction (Table 1). As context for this experiment, Table 2 shows a replication of a number of high-profile RDD applications (Lee, 2008; Cattaneo et al., 2015; Ludwig and Miller, 2007; Meyerson, 2014; Matsudaira, 2008; Jacob and Lefgren, 2004; Oreopoulos, 2006; Akhtari et al., 2022) using a variety of methods; further description of each benchmark dataset is provided in the Supplementary Material.

A fundamental challenge in assessing methods for causal inference on real-world data is that we typically do not have access to a ground truth we can use for evaluation. In light of this challenge, (Athey et al., 2024) proposed using Wasserstein Generative Adversarial Networks (WGANs) to generate synthetic data that mimic benchmark datasets, such as to obtain simulation instances that enable comparing methods in realistic settings. Our WGAN evaluation pipeline involves the following three steps:

1. **Training:** We train WGANs on each benchmark dataset to mimic the underlying data generating process for both treated and control units.

¹²As shown in (Armstrong and Kolesár, 2018), doing so would in fact be impossible under smoothness assumptions like Assumption 1 alone. We do note that proving uniform bounds for methods involving data-driven smoothness guarantees can be possible under more restricted function classes, e.g., (Armstrong and Kolesár, 2020, Appendix E) consider such results over function classes where global polynomial approximations are guaranteed to capture worst-case local curvature of the underlying functions. We leave a study of uniformity guarantees for PLRD over restricted function classes to future work.

Dataset	conventional	rdrobust	RDHonest	plrd
Lee (2008)	0.064 ± 0.023	0.059 ± 0.026	0.058 ± 0.032	0.073 ± 0.024
Cattaneo et al. (2015)	7.416 ± 2.862	7.504 ± 3.419	8.030 ± 4.210	5.830 ± 2.127
Ludwig and Miller (2007)	-2.440 ± 2.238	-2.809 ± 2.535	-3.153 ± 2.828	-1.249 ± 1.507
Meyersson (2014)	3.019 ± 2.782	2.979 ± 3.284	3.024 ± 3.109	3.010 ± 2.093
Matsudaira (2008) (reading)	0.036 ± 0.029	0.030 ± 0.033	0.031 ± 0.046	0.017 ± 0.033
Matsudaira (2008) (math)	0.075 ± 0.025	0.072 ± 0.030	0.057 ± 0.041	0.073 ± 0.028
Jacob and Lefgren (2004) (reading)	0.130 ± 0.042	0.143 ± 0.048	0.127 ± 0.046	0.100 ± 0.040
Jacob and Lefgren (2004) (math)	0.195 ± 0.035	0.203 ± 0.039	0.187 ± 0.058	0.168 ± 0.039
Oreopoulos (2006)	0.070 ± 0.104	0.099 ± 0.143	0.070 ± 0.125	0.021 ± 0.064
Akhtari et al. (2022) (4th grade)	0.022 ± 0.047	0.018 ± 0.056	-0.020 ± 0.049	0.042 ± 0.046
Akhtari et al. (2022) (8th grade)	0.040 ± 0.058	0.049 ± 0.068	0.020 ± 0.068	0.036 ± 0.054
Akhtari et al. (2022) (headmaster)	0.274 ± 0.035	0.267 ± 0.040	0.269 ± 0.043	0.270 ± 0.033

Table 2: 95% confidence intervals produced by different methods on cited benchmark datasets.

2. **Generation:** Once trained, the WGANs generate multiple synthetic datasets that reflect the characteristics of the original data.
3. **Evaluation:** We apply the various methods, including our proposed PLRD method, to the synthetic datasets. The performance of these methods is then evaluated based on the empirical coverage and the average width of the confidence intervals.

As illustrated in Figure 1, a trained WGAN is able to simulate data whose distribution is very similar to that of the original data; thus suggesting that if a method gets reliable coverage on the simulator we can also trust it on the real data.

Our WGAN training procedure consists of the following steps. First, apply marginal CDF transforms to the data to convert (X_i, Y_i) to (U_i, V_i) , where $U_i = F_x(X_i)$ and $V_i = F_y(Y_i)$. For discrete data we break ties by randomization, and for additional stability we use a Gaussian transformation. We train two generators $\hat{g}_w(u, z)$ ($w = 0, 1$), one for treated units and one for controls, that mimic the conditional distribution of the CDF-transformed outcome V given the CDF-transformed running variable U . Finally, we invert the CDF-transforms to obtain $(\tilde{X}_i, \tilde{Y}_i)$ from the pair $(\tilde{U}_i, \tilde{V}_i)$ generated by the trained WGAN.

We calculate the ground truth $\tilde{\tau}$ for the GAN estimated distribution as follows. We first calculate the cutoff c' under the applied CDF transforms. We then sample random noise $Z_{0,q}, Z_{1,q} \in \mathbb{R}$ for $q = 1, \dots, Q$, where Q is a large number (10^7 in our simulations), generate CDF-transformed outcomes at the threshold c' , and invert them:

$$\tilde{\tau} := \frac{1}{Q} \sum_{q=1}^Q (F_y^{-1}(\hat{g}_1(c', Z_{1,q})) - F_y^{-1}(\hat{g}_0(c', Z_{0,q}))).$$

We generate $N = 1,000$ simulated datasets indexed by $i = 1, \dots, N$, and compute confidence intervals $[a_{i,k}, b_{i,k}]$, for each method k . The empirical coverage for the k -th method is given by $\hat{p}_k = N^{-1} \sum_{i=1}^N \mathbf{1}\{a_{i,k} \leq \tilde{\tau} \leq b_{i,k}\}$. Table 1 shows the obtained results.

4.1 Revisiting existing simulation designs

For completeness, we also benchmark the PLRD estimator against other existing methods on some standard simulation designs. First, we consider a pure noise setting where we generate the running

Dataset	conventional			rdrobust			RDHonest			plrd		
	cvrg	width	rmse	cvrg	width	rmse	cvrg	width	rmse	cvrg	width	rmse
Pure noise	92.3%	1.392	0.398	92.9%	1.652	0.460	96.0%	1.854	0.445	97.2%	1.011	0.228
Setting 1	88.4%	0.203	0.064	90.9%	0.240	0.071	95.4%	0.288	0.069	94.8%	0.237	0.061
Setting 2	79.7%	0.266	0.099	90.5%	0.296	0.085	96.2%	0.440	0.095	97.3%	0.389	0.090
Setting 3	90.1%	0.213	0.064	92.7%	0.246	0.069	96.2%	0.301	0.070	99.6%	0.290	0.050
Setting 4	91.8%	0.201	0.059	92.3%	0.238	0.068	96.6%	0.276	0.064	96.1%	0.178	0.045

Table 3: Empirical coverage, average width, and RMSE of 95% confidence intervals produced by different methods on pure noise and on simulations from (Calonico et al., 2014). All simulations have sample size 500, and results are aggregated across 10,000 replications. Methods are run as in Table 1.

variables as $X_i \sim \text{Unif}([-1, 1])$, the outcomes as $Y_i \sim \mathcal{N}(0, 1)$, and set treatment to $W_i = 1(\{X_i \geq 0\})$. Next, we consider simulation settings used in (Calonico et al., 2014, Section 6). In each of these settings, we draw the running variable $X_i \sim 2\text{Beta}(2, 4) - 1$, and the outcome is given by $Y_i = m(X_i) + \varepsilon_i$, where ε_i are mean-zero Gaussian noise with variance 0.1295². The form of $m(\cdot)$ varies for the different simulation settings; see (Calonico et al., 2014) for details on Settings 1 through 3. Setting 4 uses the function $m = m_{\text{quad}}$ as in (Imbens and Kalyanaraman, 2012, Section 4.6). We note that the linear treatment effect condition (Assumption 2) is not satisfied in several of these designs, so we expect our pragmatic safeguard (given in Algorithm 1) to play an important role in delivering robust performance.

Table 3 presents the empirical coverage and average width of 95% CIs produced by different methods across these simulation settings. As before, **rdrobust** falls short of nominal coverage; note that these numbers are in line with those originally reported by (Calonico et al., 2014, Table 1). **RDHonest** delivers valid coverage at the expense of wider intervals, and **plrd** delivers valid coverage with intervals whose width is competitive with **rdrobust**.

5 Formal Results

Our point estimator, i.e., the $\hat{\gamma}_i$ -weighted average of outcomes Y_i with weights as given in (3) is, ignoring cross-fitting, equivalent to a minimax linear estimator for τ_c over the class $\mathcal{M}_{\hat{B}}$ of all functions satisfying Assumption 2 and for which $\mu_0''(x)$ is globally $\hat{B}$ -Lipschitz, i.e.,¹³ $\hat{\gamma}_{\text{PLRD}} := \sum_{i=1}^n \hat{\gamma}_i Y_i$, where

$$\hat{\gamma} := \underset{\gamma}{\operatorname{argmin}} \left\{ \hat{\sigma}^2 \|\gamma\|_2^2 + \sup_{\mu_0, \mu_1 \in \mathcal{M}_{\hat{B}}} \Delta(\gamma; \mu_0, \mu_1)^2 \right\}, \quad (8)$$

$$\Delta(\gamma; \mu_0, \mu_1) := \sum_{i=1}^n \gamma_i \mu_{W_i}(X_i) - \{\mu_1(c) - \mu_0(c)\}.$$

If $\hat{B}$ were a constant chosen a-priori and $\mu_0''(x)$ were in fact globally $\hat{B}$ -Lipschitz, we could use existing results for minimax linear regression-discontinuity inference (e.g., Armstrong and Kolesár,

¹³The conditional variance of a linear estimator as in (8) is $\sum_{i=1}^n \hat{\gamma}_i^2 \sigma_i^2$ with $\sigma_i^2 = \text{Var}[Y_i | X_i]$. If we had $\sigma_i^2 = \hat{\sigma}^2$ for all $i = 1, \dots, n$, then (8) would exactly be the minimax linear estimation problem for τ_c . In the case of heteroskedasticity, the term $\hat{\sigma}^2 \|\gamma\|_2^2$ should be interpreted as a variance proxy that yields a practical estimator—and one that still maintains provable validity guarantees under heteroskedasticity; see (Imbens and Wager, 2019, p. 272) for further discussion.

2018; Imbens and Wager, 2019) to directly provide guarantees for $\hat{\tau}_{\text{PLRD}}$. But here, of course, this is not the case: Our choice of $\hat{B}$ is learned in *Step 1* of our algorithm, and Assumption 1 only implies local—not global—Lipschitz continuity in $\mu_0''(x)$.

This section records the main technical statements underlying Theorem 1. We start by verifying validity of the cubic regression in *Step 1* of our procedure. The following result is a direct consequence of standard results for local polynomial regression (e.g., Fan and Gijbels, 1996, Theorem 3.1), and implies that $\hat{B}_{(1)}$ and $\hat{B}_{(2)}$ in Algorithm 1 satisfy $\hat{B}_{(k)} \rightarrow_p B$ as $n \rightarrow \infty$.

Proposition 2. *Suppose that Assumptions 1 and 2 hold and that X_i has a density bounded away from 0 in a neighborhood of c , and $B := \max\{|\mu_0'''(c)|, |\mu_1'''(c)|, \varepsilon\}$. Run the local cubic regression:*

$$\hat{\beta} := \underset{\beta \in \mathbb{R}^6}{\operatorname{argmin}} \sum_{|X_i - c| \leq \ell} \left(Y_i - \beta_{W_i,0} - \beta_{W_i,1}(X_i - c) - \frac{\beta_2}{2}(X_i - c)^2 - \frac{\beta_3}{6}(X_i - c)^3 \right)^2, \quad (9)$$

with bandwidth ℓ scaling as in (6), and set¹⁴ $\hat{B} = \max\{|\beta_3|, \varepsilon\}$. Then, $\hat{B} \rightarrow_p B$ as $n \rightarrow \infty$.

Our core analytic result is the following central limit theorem for our PLRD estimator. Asymptotic validity of the PLRD confidence intervals as claimed in Theorem 1 then follows immediately. Proofs of Theorems 1 and 3, together with the auxiliary lemmas, are given in the Supplementary Material.

Theorem 3. *Under the conditions of Theorem 1, $\hat{\tau}_{\text{PLRD}}$ as constructed in Algorithm 1 satisfies*

$$(\hat{\tau}_{\text{PLRD}} - b(\hat{\gamma}) - \tau_c) / s(\hat{\gamma}) \xrightarrow{d} \mathcal{N}(0, 1), \quad (10)$$

where $b(\hat{\gamma}) = \sum_{i=1}^n \hat{\gamma}_i \mu_{W_i}(X_i) - \tau_c$ is the conditional bias of $\hat{\tau}_{\text{PLRD}}$, and $s^2(\hat{\gamma}) = \sum_{i=1}^n \hat{\gamma}_i^2 \sigma_i^2$ where $\sigma_i^2 = \operatorname{Var}(Y_i | X_i)$. Furthermore,

$$\limsup_{n \rightarrow \infty} (|b(\hat{\gamma})| - \hat{b}) / s(\hat{\gamma}) \leq_p 0, \quad (11)$$

i.e., the learned bias bound $\hat{b}$ in (5) is conservative for the true bias; and $s(\hat{\gamma}) = \mathcal{O}_p(n^{-3/7})$.

6 Extensions

It is often of interest to incorporate auxiliary covariates that are predictive of the outcome but unrelated to the treatment assignment. One approach to doing so is by augmenting (8) with covariate-balance conditions as linear constraints (Imbens and Wager, 2019), provided that exact balance is feasible. Another approach is to fit a flexible predictive model and replace the outcomes with the residuals from this fit (Noack et al., 2026a). Investigation of these approaches is left for future work.

Extending PLRD to fuzzy RDDs (see (Imbens and Lemieux, 2008) for a review) requires an analogue of Assumption 2 for the local average treatment effect. A linear shape restriction on this imposes a joint constraint on the numerator and the denominator, and is less transparent than in the sharp-design. An alternative formulation would instead impose partial linearity on the intent-to-treat function and treat the first-stage discontinuity as a separate object subject to its own smoothness constraint, similar in spirit to (Noack and Rothe, 2024).

Adapting PLRD to clustered data requires replacing the heteroskedasticity-robust sandwich variance in the last step of Algorithm 1 with a cluster-robust analogue, and optimizing the minimax weights to account for within-cluster correlation in the objective (Noack et al., 2026b).

¹⁴As an additional safe-guard against small sample sizes, we calculate a standard error estimate $\text{s.e.}(\hat{\beta}_3)$ and use the conservative yet asymptotically consistent estimate $\hat{B} = \max\{|\hat{\beta}_3| \pm 1.96 \cdot \text{s.e.}(\hat{\beta}_3)|, \varepsilon\}$ in our software.

Replication Files

The `plrd` R package implementing the proposed method, along with replication files, is publicly available at <https://github.com/ghoshadi/plrd>.

Acknowledgments

We are grateful to Evan Munro for sharing results on the behavior of existing methods for RDDs under generative adversarial simulation following (Athey et al., 2024). We are also grateful to Tim Armstrong, P. Aronow, Max Farrell, Paul Goldsmith-Pinkham, Calumn Hamilton, Michal Kolesár, Christoph Rothe and Rocío Titiunik for insightful conversations and for comments on a previous version of this paper. This work was supported by the National Science Foundation under grant SES-2242876, the Office of Naval Research under grants N00014-17-1-2131 and N00014-19-1-2468, and a gift from Amazon.

References

- Akhtari, M., D. Moreira, and L. Trucco (2022, February). Political turnover, bureaucratic turnover, and the quality of public services. *American Economic Review* 112(2), 442–93.
- Armstrong, T. B. and M. Kolesár (2018). Optimal inference in a class of regression models. *Econometrica* 86(2), 655–683.
- Armstrong, T. B. and M. Kolesár (2020). Simple and honest confidence intervals in nonparametric regression. *Quant. Econ.* 11(1), 1–39.
- Athey, S., G. W. Imbens, J. Metzger, and E. Munro (2024). Using wasserstein generative adversarial networks for the design of monte carlo simulations. *Journal of Econometrics* 240(2), 105076.
- Boyd, S. and L. Vandenberghe (2004). *Convex optimization*. Cambridge University Press, Cambridge.
- Calonico, S., M. D. Cattaneo, M. H. Farrell, F. Palomba, and R. Titiunik (2025). Treatment effect heterogeneity in regression discontinuity designs. arXiv preprint arXiv:2503.13696.
- Calonico, S., M. D. Cattaneo, M. H. Farrell, and R. Titiunik (2017). `rdrobust`: Software for regression-discontinuity designs. *The Stata Journal* 17(2), 372–404.
- Calonico, S., M. D. Cattaneo, M. H. Farrell, and R. Titiunik (2019). Regression discontinuity designs using covariates. *Rev. Econ. Stat.* 101(3), 442–451.
- Calonico, S., M. D. Cattaneo, and R. Titiunik (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica* 82(6), 2295–2326.
- Calonico, S., M. D. Cattaneo, and R. Titiunik (2015). `Rdrobust`: An R package for robust nonparametric inference in regression-discontinuity designs. *R J.* 7(1), 38.
- Cattaneo, M. D., B. R. Frandsen, and R. Titiunik (2015). Randomization inference in the regression discontinuity design: An application to party advantages in the us senate. *Journal of Causal Inference* 3(1), 1–24.

- Cattaneo, M. D., N. Idrobo, and R. Titiunik (2024). *A practical introduction to regression discontinuity designs: Extensions*. Cambridge University Press.
- Cattaneo, M. D. and R. Titiunik (2022). Regression discontinuity designs. *Annual Review of Economics* 14(1), 821–851.
- Cheng, M.-Y., J. Fan, and J. S. Marron (1997). On automatic boundary corrections. *The Annals of Statistics* 25(4), 1691 – 1708.
- Currie, J., H. Kleven, and E. Zwiers (2020). Technology and big data are changing economics: Mining text to track methods. In *AEA Papers and Proceedings*, Volume 110, pp. 42–48. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.
- Donoho, D. L. (1994). Statistical estimation and optimal recovery. *Ann. Statist.* 22(1), 238–270.
- Fan, J. and I. Gijbels (1996). *Local Polynomial Modelling and Its Applications*, Volume 66 of *Monographs on Statistics and Applied Probability*. London: Chapman & Hall.
- Goldsmith-Pinkham, P. (2024). Tracking the credibility revolution across fields. *arXiv preprint*, arXiv:2405.20604.
- Goodfellow, I., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio (2014). Generative adversarial nets. *Advances in neural information processing systems* 27, 2672–2680.
- Hahn, J., P. Todd, and W. V. der Klaauw (2001). Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica* 69(1), 201–209.
- Imbens, G. and K. Kalyanaraman (2012). Optimal bandwidth choice for the regression discontinuity estimator. *Rev. Econ. Stud.* 79(3), 933–959.
- Imbens, G. and T. Lemieux (2008). Regression discontinuity designs: A guide to practice. *J. Econometrics* 142(2), 615–635.
- Imbens, G. and S. Wager (2019). Optimized regression discontinuity designs. *Rev. Econ. Stat.* 101(2), 264–278.
- Imbens, G. and T. Zajonc (2009). Regression discontinuity design with vector-argument assignment rules. Unpublished manuscript.
- Imbens, G. W. and C. F. Manski (2004). Confidence intervals for partially identified parameters. *Econometrica* 72(6), 1845–1857.
- Jacob, B. A. and L. Lefgren (2004). Remedial education and student achievement: A regression-discontinuity analysis. *Rev. Econ. Stat.* 86(1), 226–244.
- Kolesár, M. and C. Rothe (2018). Inference in regression discontinuity designs with a discrete running variable. *Amer. Econ. Rev.* 108(8), 2277–2304.
- Lee, D. S. (2008). Randomized experiments from non-random selection in us house elections. *J. Econometrics*. 142(2), 675–697.
- Lee, D. S. and T. Lemieux (2010). Regression discontinuity designs in economics. *Journal of Economic Literature* 48(2), 281–355.

- Low, M. G. (1997). On nonparametric confidence intervals. *The Annals of Statistics* 25(6), 2547–2554.
- Ludwig, J. and D. L. Miller (2007). Does head start improve children’s life chances? evidence from a regression discontinuity design. *The Quarterly Journal of Economics* 122(1), 159–208.
- Matsudaira, J. D. (2008). Mandatory summer school and student achievement. *Journal of Econometrics* 142(2), 829–850. The regression discontinuity design: Theory and applications.
- Meyersson, E. (2014). Islamic rule and the empowerment of the poor and pious. *Econometrica* 82(1), 229–269.
- Noack, C., T. Olma, and C. Rothe (2026a). Flexible covariate adjustments in regression discontinuity designs. *Journal of Econometrics* forthcoming.
- Noack, C., T. Olma, and C. Rothe (2026b). Inference in regression discontinuity designs with clustered data. arXiv preprint arXiv:2603.18870.
- Noack, C. and C. Rothe (2024). Bias-aware inference in fuzzy regression discontinuity designs. *Econometrica* 92(3), 687–711.
- Oreopoulos, P. (2006). Estimating average and local average treatment effects of education when compulsory schooling laws really matter. *American Economic Review* 96(1), 152–175.
- R Core Team (2023). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Stommes, D., P. Aronow, and F. Sävje (2023). On the reliability of published findings using the regression discontinuity design in political science. *Research & Politics* 10(2), 20531680231166457.
- Thistlethwaite, D. L. and D. T. Campbell (1960). Regression-discontinuity analysis: An alternative to the ex post facto experiment. *Journal of Educational psychology* 51(6), 309.
- van der Vaart, A. W. (1998). *Asymptotic statistics*, Volume 3 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge.
- van der Vaart, A. W. and J. A. Wellner (1996). *Weak convergence and empirical processes*. Springer Series in Statistics. Springer-Verlag, New York. With applications to statistics.
- Wager, S. (2024). *Causal inference: A statistical learning approach*. Technical report, Stanford University.
- White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica* 48(4), 817–838.

A Proofs of the Main Results

A.1 Proof of Theorem 1

Define $\hat{s}^2(\hat{\gamma}) = \sum_{i=1}^n \hat{\gamma}_i^2 \hat{\sigma}_i^2$ as in Algorithm 1. It follows from (White, 1980) that $\hat{s}(\hat{\gamma})/s(\hat{\gamma}) \rightarrow_p 1$; see Lemma 15 for details. Next, fix any $\delta \in (0, 1)$. In view of (11) of the main paper and Slutsky's lemma, $\mathbb{P}(A_n) \rightarrow 1$ where

$$A_n := \{|b(\hat{\gamma})| \leq \hat{b} + \delta \hat{s}(\hat{\gamma})\}.$$

We show in Lemma 14 that for any s, h, t and $\delta > 0$,

$$\inf_{|b| \leq t + \delta s} \mathbb{P}(|b + sZ| \leq h) \geq \inf_{|b| \leq t} \mathbb{P}(|b + sZ| \leq h) - \delta,$$

where $Z \sim \mathcal{N}(0, 1)$. Using this with $s = \hat{s}(\hat{\gamma})$, $h = \hat{h}_\alpha$ and $t = \hat{b}$, we get, for every n ,

$$\begin{aligned} \mathbf{1}_{A_n} \mathbb{P}\left(|b(\hat{\gamma}) + \hat{s}(\hat{\gamma})Z| \leq \hat{h}_\alpha \mid (X_i, Y_i)_{i=1}^n\right) &\geq \mathbf{1}_{A_n} \inf_{|b| \leq \hat{b} + \delta \hat{s}(\hat{\gamma})} \mathbb{P}\left(|b + \hat{s}(\hat{\gamma})Z| \leq \hat{h}_\alpha\right) \\ &\geq \mathbf{1}_{A_n} \left(\inf_{|b| \leq \hat{b}} \mathbb{P}(|b + \hat{s}(\hat{\gamma})Z| \leq \hat{h}_\alpha) - \delta \right) \\ &\geq \mathbf{1}_{A_n} (1 - \alpha - \delta), \end{aligned} \quad (12)$$

where the last inequality follows from the definition of $\hat{h}_\alpha$. On the other hand, Theorem 3 and Slutsky's lemma imply that

$$Z_n := \frac{\hat{\tau}_{\text{PLRD}} - \tau_c - b(\hat{\gamma})}{\hat{s}(\hat{\gamma})} \xrightarrow{d} Z \sim \mathcal{N}(0, 1).$$

Consequently,

$$\begin{aligned} \liminf_{n \rightarrow \infty} \mathbb{P}(\tau_c \in \mathcal{I}_\alpha) &= \liminf_{n \rightarrow \infty} \mathbb{P}(|\hat{\tau}_{\text{PLRD}} - \tau_c| \leq \hat{h}_\alpha) \\ &= \liminf_{n \rightarrow \infty} \mathbb{P}\left(|b(\hat{\gamma}) + \hat{s}(\hat{\gamma})Z| \leq \hat{h}_\alpha\right) \\ &\geq \liminf_{n \rightarrow \infty} \mathbb{P}(\{|b(\hat{\gamma}) + \hat{s}(\hat{\gamma})Z| \leq \hat{h}_\alpha\} \cap A_n) \\ &\geq 1 - \alpha - \delta. \end{aligned} \quad (\text{using (12) and } \mathbb{P}(A_n) \rightarrow 1)$$

Since $\delta \in (0, 1)$ is arbitrary, we are through. $\square$

A.2 Proof of Theorem 3

The optimization problem in Algorithm 1 is inherently related to its convex dual (see (Boyd and Vandenberghe, 2004) for a definition). We show in Appendix E that the dual problem is given by

$$\begin{aligned} \hat{\theta}_n(\sigma^2, B) &= \underset{\theta=(\rho, \lambda) \in \Theta}{\operatorname{argmin}} \hat{L}_n(\theta; \sigma^2, B), \\ \hat{L}_n(\theta; \sigma^2, B) &:= \frac{1}{4n\sigma^2} \sum_{i=1}^n G^2(\theta; X_i - c, W_i) + \frac{\lambda_1^2}{4nB^2} + \frac{\lambda_2 - \lambda_3}{n}, \end{aligned} \quad (13)$$

where $\theta = (\rho, \lambda)$ varies in the set

$$\Theta := \{(f, \lambda) \mid f : [-\ell, \ell] \rightarrow \mathbb{R}, f(0) = f'(0) = f''(0) = 0, f'' \text{ is } \lambda_1\text{-Lipschitz}, \lambda \in [0, \infty) \times \mathbb{R}^5\}, \quad (14)$$

and

$$G(\rho, \lambda; x, w) := \rho(x) + \lambda_2 w + \lambda_3(1 - w) + \lambda_4 w x + \lambda_5(1 - w)x + \lambda_6 x^2.$$

In Algorithm 1, we use sample splitting to estimate B (cf. Proposition 2 of the main paper) and the conditional variance $\sigma^2 = \mathbb{E}[\text{Var}(Y_i \mid X_i)]$. However, for ease of exposition, we assume throughout this section that we have $\widehat{B} \rightarrow_p B$ and $\widehat{\sigma}^2 \rightarrow_p \sigma^2$ independently of the data, and

$$\widehat{\theta}_n := \widehat{\theta}_n(\widehat{\sigma}^2, \widehat{B}) = \underset{\theta \in \Theta}{\operatorname{argmin}} \widehat{L}_n(\theta; \widehat{\sigma}^2, \widehat{B}).$$

We can recover the primal solution $(\widehat{\gamma}_n, \widehat{t}_n)$ from the dual solution $\widehat{\theta}_n$ as:

$$\widehat{\gamma}_{n,i} = -\frac{1}{2\widehat{\sigma}_i^2} G(\widehat{\theta}_n; X_i - c, W_i), \quad \widehat{t}_n = \frac{1}{2\widehat{B}^2} \widehat{\lambda}_{n,1}.$$

Next, define the population-level dual problem as:

$$\begin{aligned} \theta_n^*(\sigma^2, B) &= \underset{\theta \in \Theta}{\operatorname{argmin}} L_n(\theta; \sigma^2, B), \\ L_n(\theta; \sigma^2, B) &:= \frac{1}{4\sigma^2} \mathbb{E}[G^2(\theta; X - c, W)] + \frac{\lambda_1^2}{4nB^2} + \frac{\lambda_2 - \lambda_3}{n}. \end{aligned} \quad (15)$$

Our proof strategy hinges on establishing fast-enough convergence of $\widehat{\theta}_n$ to $\theta_n^*(\sigma^2, B)$ to enable us to read-off large-sample properties of $\widehat{\tau}_{\text{PLRD}}$ from the solution to the population dual (15).

Although the empirical minimization problem (13) looks similar to a moment-estimation problem, it has some technical complications, *e.g.*,

- (i) the dual parameter $\theta = (\rho, \lambda)$ involves a function ρ and a vector λ ,
- (ii) the set Θ in (14) is *not* the Cartesian product of a function space and a vector space,
- (iii) the population dual problem in (15) involves the sample size n .

To tackle these issues, we borrow tools from empirical process theory (see, *e.g.*, (van der Vaart and Wellner, 1996) for a textbook treatment). The following result, proven in Appendix A.3, characterizes the rate of convergence of the empirical dual optimizer $\widehat{\theta}_n$ to the population dual optimizer θ_n^* ; and implies (Corollary 5) that the weights underlying $\widehat{\tau}_{\text{PLRD}}$ satisfy a Lindeberg-type condition.

Lemma 4. *Let $(\widehat{\rho}_n, \widehat{\lambda}_n) = \widehat{\theta}_n(\widehat{\sigma}^2, \widehat{B})$ be the optimizer of the empirical dual problem (13), and $(\rho_n^*, \lambda_n^*) = \theta_n^*(\sigma^2, B)$ be the optimizer of the population dual problem (15). Assume that the pilot bandwidth ℓ in Algorithm 1 satisfies (6), and $\widehat{B} \rightarrow_p B$. Then, under the conditions of Theorem 1,*

$$\|\widehat{\rho}_n - \rho_n^*\|_{L^2(P)}^2 + \|\widehat{\lambda}_{n,-1} - \lambda_{n,-1}^*\|_2^2 = \mathcal{O}_p\left(\frac{\ell^6}{n^{1/4+6/7}} + \frac{\ell^3}{n^{1+3/7}}\right).$$

Corollary 5. *Under the conditions of Theorem 1, the weights $\hat{\gamma}_i$ underlying $\hat{\tau}_{\text{PLRD}}$ satisfy*

$$\frac{\|\hat{\gamma}\|_\infty}{\|\hat{\gamma}\|_2} = \frac{\max_{1 \leq i \leq n} |\hat{\gamma}_i|}{(\sum_{i=1}^n \hat{\gamma}_i^2)^{1/2}} \rightarrow_p 0. \quad (16)$$

The proof of Corollary 5 is delegated to Appendix A.4. Corollary 5 immediately implies that the central limit theorem claimed in Theorem 3 holds; see the proof of (Imbens and Wager, 2019, Theorem 1) for details.

We delegate showing that $s^2(\hat{\gamma}) = \mathcal{O}_p(n^{-6/7})$ (Lemma 12) and $\hat{b}/s(\hat{\gamma}) = \mathcal{O}_p(1)$ (Lemma 13) to Appendix F. Finally, to verify (11) of the main paper, we note that Assumptions 1 and 2 yield $b(\hat{\gamma}) = \sum_{i=1}^n \hat{\gamma}_i \mu_{W_i}(X_i) = B \sum_{i=1}^n \hat{\gamma}_i \rho(X_i - c)$, where

$$\rho(x - c) := (\mu_0(x) - \mu_0(c) - \mu'_0(c)(x - c) - \mu''_0(c)(x - c)^2/2)/B.$$

It follows from Algorithm 1 (with cross-fitted $\hat{B}$) that $|b(\hat{\gamma})| \leq B\hat{t}$ where $\hat{t} = \hat{b}/\hat{B}$. Consequently,

$$\limsup_{n \rightarrow \infty} \frac{|b(\hat{\gamma})| - \hat{b}}{s(\hat{\gamma})} - \frac{(B/\hat{B} - 1)\hat{b}}{s(\hat{\gamma})} \leq_p 0. \quad (17)$$

Furthermore, since $\hat{B} \rightarrow_p B$ and $\hat{b}/s(\hat{\gamma}) = \mathcal{O}_p(1)$, we have

$$\lim_{n \rightarrow \infty} \frac{(B/\hat{B} - 1)\hat{b}}{s(\hat{\gamma})} =_p 0,$$

and so (17) implies (11) of the main paper. $\square$

A.3 Proof of Lemma 4

Proof. For simplicity in notation, we assume without loss of generality that the regression discontinuity threshold is $c = 0$. We split the proof into some auxiliary results, namely, Lemmas 8 to 12 in Appendix F. First, using tools from empirical process theory (see (van der Vaart and Wellner, 1996) for a book-level treatment), we prove a uniform concentration bound for the loss functions of the empirical and population dual problems. Recall that the dual parameter space is defined as

$$\Theta(\ell) := \{(f, \lambda) \mid f : [-\ell, \ell] \rightarrow \mathbb{R}, f(0) = f'(0) = f''(0) = 0, f'' \text{ is } \lambda_1\text{-Lipschitz}, \lambda \in [0, \infty) \times \mathbb{R}^5\}.$$

Define a neighborhood $\Theta_n(\ell, \delta, \varepsilon)$ of the dual optimizer θ_n^* , as follows:

$$\Theta_n(\ell, \delta, \varepsilon) := \{(\rho, \lambda) \mid (\rho, \lambda) \in \Theta(\ell), |\lambda_1 - \lambda_{n,1}^*| \leq \delta \lambda_{n,1}^* \text{ and } \|\lambda_{-1} - \lambda_{n,-1}^*\|_2 \leq \varepsilon \|\lambda_{n,-1}^*\|_2\},$$

where $\ell, \delta, \varepsilon \in (0, 1)$ are fixed. We show in Lemma 8 the following maximal inequality.

$$\mathbb{E} \sup_{\theta \in \Theta_n(\ell, \delta, \varepsilon)} |\hat{L}_n(\theta; \hat{\sigma}^2, \hat{B}) - L_n(\theta; \hat{\sigma}^2, \hat{B})| \lesssim (\lambda_{n,1}^*)^2 \frac{\ell^6}{\sqrt{n}} + \|\lambda_{n,-1}^*\|_2^2 \frac{\sqrt{\varepsilon}}{\sqrt{n}} + \lambda_{n,1}^* \|\lambda_{n,-1}^*\|_2 \frac{\ell^3}{\sqrt{n}}.$$

Next, Lemma 10 simplifies the above bound to the following.

$$\mathbb{E} \sup_{\theta \in \Theta_n(\ell, \delta, \varepsilon)} |\widehat{L}_n(\theta; \widehat{\sigma}^2, \widehat{B}) - L_n(\theta; \widehat{\sigma}^2, \widehat{B})| \lesssim (\lambda_{n,1}^*)^2 \frac{\ell^6}{\sqrt{n}} + \lambda_{n,1}^* \frac{\ell^3}{n^{3/2}} + \frac{\sqrt{\varepsilon}}{n^{5/2}}. \quad (18)$$

On the other hand, Lemma 10 tells us that

$$\sup_{\theta \in \Theta_n(\ell, \delta, \varepsilon)} |L_n(\theta; \widehat{\sigma}^2, \widehat{B}) - L_n(\theta; \widehat{\sigma}^2, B)| \lesssim \frac{(\lambda_{n,1}^*)^2}{n} \left| \frac{1}{\widehat{B}} - \frac{1}{B} \right|.$$

This, in conjunction with (18) tells us that for every $\eta \in (0, 1)$, there exists $C = C(\eta) > 0$ such that, for all sufficiently large n ,

$$\mathbb{P} \left(\sup_{\theta \in \Theta_n(\ell, \delta, \varepsilon)} |\widehat{L}_n(\theta; \widehat{\sigma}^2, \widehat{B}) - L_n(\theta; \sigma^2, B)| \leq C \left((\lambda_{n,1}^*)^2 \frac{\ell^6}{\sqrt{n}} + \lambda_{n,1}^* \frac{\ell^3}{n^{3/2}} + \frac{\sqrt{\varepsilon}}{n^{5/2}} + \frac{(\lambda_{n,1}^*)^2}{n} \right) \right) \geq 1 - \eta. \quad (19)$$

We also show in Lemma 9 that

$$L_n(\theta; \sigma^2, B) - L_n(\theta_n^*; \sigma^2, B) \geq \frac{\alpha}{4\sigma^2} \left(\|\rho - \rho_n^*\|_{L^2(P)}^2 + \|\lambda_{-1} - \lambda_{n,-1}^*\|_2^2 \right) + \frac{(\lambda_1 - \lambda_{n,1}^*)^2}{4nB^2} - R_n, \quad (20)$$

for some $\alpha \in (0, 1)$, where $0 \leq R_n \leq \ell^3(\lambda_1 + \lambda_{n,1}^*)\|\lambda_{-1} - \lambda_{n,-1}^*\|_2/2\sigma^2$. Note, for $(\rho, \lambda) \in \Theta_n(\ell, \delta, \varepsilon)$,

$$R_n \leq \frac{1}{2\sigma^2} \ell^3(2 + \delta)\lambda_{n,1}^* \varepsilon \|\lambda_{n,-1}^*\|_2 \leq \frac{3}{2\sigma^2} \lambda_{n,1}^* \|\lambda_{n,-1}^*\|_2 \ell^3 \varepsilon. \quad (21)$$

Next, we show in Lemma 11 that for every $\delta, \eta \in (0, 1)$, and $\varepsilon = O(n^{-1/4})$, it holds that

$$\mathbb{P} \left(\widehat{\theta}_n \in \Theta_n(\ell, \delta, \varepsilon) \right) \geq 1 - \eta,$$

for all sufficiently large n . The following inequalities hold whenever the event in (19) holds as well as $\widehat{\theta}_n \in \Theta_n(\ell, \delta, \varepsilon)$ with $\varepsilon = O(n^{-1/4})$, i.e., the following inequalities hold with probability at least $1 - 2\eta$ for all sufficiently large n .

$$\begin{aligned} & \frac{\alpha}{4\sigma^2} \left(\|\widehat{\rho}_n - \rho_n^*\|_{L^2(P)}^2 + \|\widehat{\lambda}_{n,-1} - \lambda_{n,-1}^*\|_2^2 \right) + \frac{(\widehat{\lambda}_{n,1} - \lambda_{n,1}^*)^2}{4nB^2} \\ & \leq \frac{3}{2\sigma^2} \lambda_{n,1}^* \|\lambda_{n,-1}^*\|_2 \ell^3 \varepsilon + L_n(\widehat{\theta}; \sigma^2, B) - L_n(\theta_n^*; \sigma^2, B) \quad (\text{using (20) and (21)}) \\ & \leq \frac{3}{2\sigma^2} \lambda_{n,1}^* \|\lambda_{n,-1}^*\|_2 \ell^3 \varepsilon + \underbrace{L_n(\widehat{\theta}; \widehat{\sigma}^2, \widehat{B}) - L_n(\theta_n^*; \widehat{\sigma}^2, \widehat{B})}_{\leq 0} + 2 \sup_{\theta \in \Theta_n(\ell, \delta, \varepsilon)} \left| \widehat{L}_n(\theta; \widehat{\sigma}^2, \widehat{B}) - L_n(\theta; \sigma^2, B) \right| \\ & \quad (\text{using triangle inequality and the definition of } \widehat{\theta}_n) \\ & \lesssim \lambda_{n,1}^* \|\lambda_{n,-1}^*\|_2 \frac{\ell^3}{n^{1/4}} + (\lambda_{n,1}^*)^2 \frac{\ell^6}{\sqrt{n}} + \lambda_{n,1}^* \frac{\ell^3}{n^{3/2}} + \frac{1}{n^{5/2+1/8}} + \frac{(\lambda_{n,1}^*)^2}{n}. \quad (\text{using (19)}) \end{aligned}$$

Finally, Lemma 10 tells us that $\|\lambda_{n,-1}^*\|_2 = O(\lambda_{n,1}^* \ell^3 + n^{-1})$, while Lemma 12 tells us that $\lambda_{n,1}^* = \mathcal{O}_p(n^{-3/7})$. This simplifies the above upper bound to the following.

$$\begin{aligned} & \frac{\alpha}{4\sigma^2} \left(\|\hat{\rho}_n - \rho_n^*\|_{L^2(P)}^2 + \|\hat{\lambda}_{n,-1} - \lambda_{n,-1}^*\|_2^2 \right) + \frac{(\hat{\lambda}_{n,1} - \lambda_{n,1}^*)^2}{4nB^2} \\ & \lesssim_\eta (\lambda_{n,1}^*)^2 \frac{\ell^6}{n^{1/4}} + \lambda_{n,1}^* \frac{\ell^3}{n^{3/2}} + \frac{(\lambda_{n,1}^*)^2}{n} \\ & \lesssim_\eta \frac{\ell^6}{n^{1/4+6/7}} + \frac{\ell^3}{n^{3/2+3/7}} + \frac{1}{n^{1+6/7}}, \end{aligned}$$

with probability at least $1 - 2\eta$, for all sufficiently large n . Finally, we can ignore the last term in the above display, since ℓ satisfies (6). This completes the proof. $\square$

A.4 Proof of Corollary 5

Proof. For the PLRD weights $\hat{\gamma}_i$ derived in Algorithm 1, we now show that

$$\frac{\|\hat{\gamma}\|_\infty}{\|\hat{\gamma}\|_2} = \frac{\max_{1 \leq i \leq n} |\hat{\gamma}_i|}{(\sum_{i=1}^n \hat{\gamma}_i^2)^{1/2}} \rightarrow_p 0.$$

Assume without loss of generality that the threshold for treatment is at $c = 0$. Recall that we use sample splitting in Algorithm 1 and write $\hat{\gamma}_i = \hat{\gamma}_{(k),i}/2$ if $i \in I_k$ ($k = 1, 2$). Thus,

$$\frac{\|\hat{\gamma}\|_\infty}{\|\hat{\gamma}\|_2} = \frac{\max_{1 \leq i \leq n} |\hat{\gamma}_i|}{(\sum_{i=1}^n \hat{\gamma}_i^2)^{1/2}} \lesssim \frac{\|\hat{\gamma}_{(1)}\|_\infty}{\|\hat{\gamma}_{(1)}\|_2} + \frac{\|\hat{\gamma}_{(2)}\|_\infty}{\|\hat{\gamma}_{(2)}\|_2}.$$

Since $\hat{\gamma}_{(1)}$ and $\hat{\gamma}_{(2)}$ are identically distributed, it suffices to show the following.

$$\frac{\|\hat{\gamma}_{(1)}\|_\infty}{\|\hat{\gamma}_{(1)}\|_2} = \frac{\max_{i \in I_1} |\hat{\gamma}_{(1),i}|}{(\sum_{i \in I_1} \hat{\gamma}_{(1),i}^2)^{1/2}} \rightarrow_p 0. \quad (22)$$

Towards showing (22), note that by Cauchy-Schwarz inequality,

$$|I_1| \sum_{i \in I_1} \hat{\gamma}_{(1),i}^2 = \sum_{i \in I_1} (2W_i - 1)^2 \sum_{i \in I_1} \hat{\gamma}_{(1),i}^2 \geq \left(\sum_{i \in I_1} (2W_i - 1) \hat{\gamma}_{(1),i} \right)^2 = 4.$$

This yields the bound $\|\hat{\gamma}_{(1)}\|_2 \gtrsim n^{-1/2}$. For the numerator in (22) use triangle inequality to write

$$\sqrt{n} \max_{i \in I_1} |\hat{\gamma}_{(1),i}| \leq \sqrt{n} \max_{i \in I_1} |\gamma_{(1),i}^*| + \sqrt{n} \max_{i \in I_1} |\hat{\gamma}_{(1),i} - \gamma_{(1),i}^*|,$$

where $\gamma_{(1),i}^* := -(\rho_n^*(X_i) + \lambda_{n,-1}^* v(X_i, W_i))/2\sigma_{(2)}^2$, with $v(x, w) = (w, 1-w, xw, x(1-w), x^2)^\top$, and

$$(\rho_n^*, \lambda_n^*) := \operatorname{argmin}_{(\rho, \lambda) \in \Theta} \frac{1}{4\sigma_{(2)}^2} \mathbb{E} [G^2(\rho, \lambda; X, W)] + \frac{\lambda_1^2}{4nB_{(2)}^2} + \frac{\lambda_2 - \lambda_3}{n},$$

with Θ as defined in (33). For any $(\rho, \lambda) \in \Theta$, $\rho^{(j)}(0) = 0$ for $j = 0, 1, 2$, and ρ'' is $\lambda_{n,1}^*$ -Lipschitz, so a Taylor expansion of $\rho(x)$ around $x = 0$ gives us

$$|\rho(x)| = |\rho(x) - \rho(0) - \rho'(0)x| = \frac{|\rho''(\xi) - \rho''(0)|}{|\xi - x|} |\xi - x| \frac{x^2}{2} \leq \lambda_{n,1}^* \frac{\ell^3}{2}. \quad (23)$$

Consequently,

$$\max_{i \in I_1} |\gamma_{(1),i}^*| \lesssim \sup_x \left| \rho_n^*(x) + v(x, 1(x > 0))^\top \lambda_{n,-1}^* \right| \lesssim \lambda_{n,1}^* \ell^3 + \|\lambda_{n,-1}^*\|_2,$$

where we used (23) to bound $\sup_x |\rho_n^*(x)|$ and Cauchy-Schwarz inequality to bound the other term. Since $B_{(2)}$ and $\sigma_{(2)}^2$ are independent of the data in the fold I_1 , Lemma 10 applies here and tells us that $\|\lambda_{n,-1}^*\|_2 = O(\lambda_{n,1}^* \ell^3 + n^{-1})$. Consequently, we have

$$\sqrt{n} \max_{i \in I_1} |\gamma_{(1),i}^*| \lesssim \sqrt{n} (\lambda_{n,1}^* \ell^3 + n^{-1}).$$

On the other hand, $\max_{i \in I_1} |\hat{\gamma}_{(1),i} - \gamma_{(1),i}^*| \lesssim \|\hat{\rho}_n - \rho_n^*\|_\infty + \|\hat{\lambda}_{n,-1} - \lambda_{n,-1}^*\|_2$. It follows from (34) that $\|\hat{\rho}_n - \rho_n^*\|_\infty \leq (\hat{\lambda}_{n,1} + \lambda_{n,1}^*) \ell^3$. Applying Lemma 4, we deduce that

$$\|\hat{\lambda}_{n,-1} - \lambda_{n,-1}^*\|_2 = \mathcal{O}_p \left(\lambda_{n,1}^* \frac{\ell^3}{n^{1/8}} + \sqrt{\lambda_{n,1}^*} \frac{\ell^{3/2}}{\sqrt{n}} + \frac{\lambda_{n,1}^*}{n^{3/4}} \right).$$

Putting this all together,

$$\sqrt{n} \max_{i \in I_1} |\hat{\gamma}_{(1),i} - \gamma_{(1),i}^*| \lesssim (\hat{\lambda}_{n,1} + \lambda_{n,1}^*) \ell^3 \sqrt{n} + \sqrt{\lambda_{n,1}^*} \ell^{3/2} + \lambda_{n,1}^* n^{-1/4}.$$

Further, we show in Lemma 12 that $\hat{\lambda}_{n,1}$ and $\lambda_{n,1}^*$ are both $O(n^{-3/7})$. Consequently,

$$\sqrt{n} \max_{i \in I_1} |\hat{\gamma}_{(1),i}| \leq \sqrt{n} \max_{i \in I_1} |\gamma_{(1),i}^*| + \sqrt{n} \max_{i \in I_1} |\hat{\gamma}_{(1),i} - \gamma_{(1),i}^*| = \mathcal{O}_p \left(n^{-3/7} \ell^3 \sqrt{n} \right) = \mathcal{O}_p \left(n^{1/14} \ell^3 \right),$$

which is $o_p(1)$ since $\ell = o(n^{-1/12})$. □

B Safeguard against Curvature Changes

While Assumption 2 provides a parsimonious and empirically realistic model for treatment effect heterogeneity in many applications, it can in principle be violated. To safeguard against such violations, our software incorporates an automatic diagnostic that tests for difference in curvatures. The diagnostic works as follows. Before running the main PLRD procedure, we conduct an ANOVA test comparing a restricted cubic polynomial model (where treatment interacts only linearly with $X_i - c$, consistent with Assumption 2) against an unrestricted model (where treatment interacts with all polynomial terms in $X_i - c$). If the test rejects at significance level α' (default 0.1%), the method proceeds with a more flexible specification that only assumes Assumption 1 but not Assumption 2. Otherwise, it uses the default specification under both assumptions. This adaptive approach helps explain the strong empirical performance of `plrd` across all simulation settings in Table 3, including those where Assumption 2 fails.

When the flexible specification is selected, estimation of the Lipschitz constant B changes to accommodate potentially different curvature on either side of the threshold. Rather than fitting a single cubic polynomial as in (9), we fit *separate* cubic regressions for treated and control units:

Proposition 6 (Consistent estimation of the Lipschitz constant B under different curvatures). *Suppose that Assumption 1 holds and that X_i has a density bounded away from 0 in a neighborhood of c , and $B := \max\{|\mu_0'''(c)|, |\mu_1'''(c)|, \varepsilon\}$. Estimate $\hat{B}$ via local cubic regression: For $w \in \{0, 1\}$ run*

$$\hat{\beta}_w := \operatorname{argmin}_{\beta_w \in \mathbb{R}^4} \sum_{\substack{W_i=w \\ |X_i-c| \leq \ell}} \left(Y_i - \beta_{w,0} - \beta_{w,1}(X_i - c) - \frac{\beta_{w,2}}{2}(X_i - c)^2 - \frac{\beta_{w,3}}{6}(X_i - c)^3 \right)^2, \quad (24)$$

with bandwidth ℓ scaling as in (6), and set $\hat{B} = \max\{|\beta_{0,3}|, |\beta_{1,3}|, \varepsilon\}$. Then, $\hat{B} \rightarrow_p B$ as $n \rightarrow \infty$.

With this estimated $\hat{B}$ in hand, the subsequent minimax optimization in Algorithm 1 adjusts the constraint set to allow for different curvature functions ρ_0 and ρ_1 on either side of the cutoff (see (4) versus (3)). The optimization problem remains computationally tractable via quadratic programming, and the remainder of the PLRD procedure—including cross-fitting, variance estimation, and confidence interval construction—proceeds similarly to the default case. We note that a formal proof of asymptotic validity under this flexible specification would follow similar arguments to those in Theorem 1, with nearly identical technical steps that are omitted here.

C Details on Baseline Methods

In this section, we give further details on the baseline methods `rdrobust` and `RDHonest` considered in our experiments, and provide further discussion of how our modeling assumptions differ from those used to-date in software for bias-aware inference in RDDs.

C.1 Bias-Corrected Confidence Intervals

Define the local polynomial estimator of order p as $\hat{\tau}_p(h_n) := \hat{\mu}_{1,p}(h_n) - \hat{\mu}_{0,p}(h_n)$, where $\hat{\mu}_{1,p}(h_n)$ and $\hat{\mu}_{0,p}(h_n)$ are the intercepts from local polynomial regressions of order p fitted separately on each side of the threshold. More precisely, for $w \in \{0, 1\}$, compute

$$\begin{aligned} \hat{\mu}_{w,p}(h_n) &:= \mathbf{e}_0' \hat{\beta}_{w,p}(h_n), \\ \hat{\beta}_{w,p}(h_n) &:= \operatorname{argmin}_{\beta \in \mathbb{R}^{p+1}} \sum_{i: W_i=w} K(|X_i - c|/h_n) \left(Y_i - \mathbf{r}_p(X_i - c)^\top \beta \right)^2, \end{aligned}$$

where $\mathbf{e}_0 = (1, 0, \dots, 0) \in \mathbb{R}^{p+1}$ is the first unit vector, $\mathbf{r}_p(x) = (1, x, \dots, x^p)^\top$, $K(\cdot)$ is a kernel, and h_n is a bandwidth. The first-order bias of $\hat{\tau}_p(h_n)$ is (Calonico et al., 2014; Imbens and Kalyanaraman, 2012)

$$B_p(h_n) = \frac{h_n^{p+1}}{(p+1)!} \left[\mu_1^{(p+1)}(c+) - \mu_0^{(p+1)}(c-) \right],$$

where $\mu_w^{(p+1)}$ denotes the derivative of $\mu_w(\cdot)$ of order $p+1$. Calonico et al. (2014) estimate the above bias using a higher-order local polynomial of order $q > p$ with a separate pilot bandwidth b_n , thus

compute the bias-corrected estimator as

$$\hat{\tau}_{p,q}^{\text{bc}}(h_n, b_n) := \hat{\tau}_p(h_n) - h_n^{p+1} \hat{B}_{n,p,q}(h_n, b_n),$$

where $\hat{B}_{n,p,q}(h_n, b_n)$ is the estimated bias term $B_p(h_n)$. The confidence interval they construct based on the bias-corrected estimator takes the following form:

$$\widehat{\text{CI}}_{1-\alpha}^{\text{bc}}(h_n, b_n) = \left(\hat{\tau}_{p,q}^{\text{bc}}(h_n, b_n) \pm z_{1-\alpha/2} \hat{\sigma}(h_n) \right), \quad (25)$$

where $z_{1-\alpha/2}$ is the critical value from the standard normal distribution, and $\hat{\sigma}(h_n)$ is the standard error (for which (Calonico et al., 2014) provide two options: plug-in and fixed matches). The above CI construction is implemented in the `rdrobust` package, with $p = 1$ and $q = 2$ as default.

The CI constructed in (25) might not perform well in finite samples because the estimator of the bias introduces additional variability. To address this, (Calonico et al., 2014) also propose an alternative, more robust CI construction using the bias-corrected estimator $\hat{\tau}_{p,q}^{\text{bc}}(h_n, b_n)$ as a starting point. They construct this robust CI by inflating the CI in (25) with a more conservative estimator of the standard error:

$$\widehat{\text{CI}}_{1-\alpha}^{\text{rbc}}(h_n, b_n) = \left(\hat{\tau}_{p,q}^{\text{bc}}(h_n, b_n) \pm z_{1-\alpha/2} \sqrt{\hat{V}_{n,p,q}^{\text{bc}}} \right), \quad (26)$$

where $\hat{V}_{n,p,q}^{\text{bc}}$ is an estimator of the asymptotic variance that accounts for both the variability in estimation and the bias-correction. In Tables 1 to 3, the columns marked as `rdrobust` correspond to the confidence intervals computed using (26).

C.2 Bias-Aware Confidence Intervals

The bias-aware inference approach (e.g., Armstrong and Kolesár, 2018, 2020; Imbens and Wager, 2019) provides a different approach for obtaining confidence intervals for τ_c . Note that the LLR estimator takes the form $\hat{\tau}_{\text{LLR}} = \sum_{i=1}^n \gamma_{i,\text{LLR}} Y_i$ for some weights $\gamma_{i,\text{LLR}}$ that are functions of the running variable. In fact, the same also holds more generally for the estimator $\hat{\tau}_p(h_n)$ defined in Appendix C.1. Next, one makes the observation that although $\hat{\tau}_{\text{LLR}}$ is motivated by a regression problem, the rest of our inference procedure hinges upon on properties that hold for any linear estimators of the form $\hat{\tau}(\gamma) = \sum_{i=1}^n \gamma_i Y_i$. Given by this observation, we can systematically do better by directly finding weights γ_i that minimize the mean squared error (MSE) of the estimator $\hat{\tau}(\gamma)$ uniformly over a class $\mathcal{M}$ of conditional response functions μ_w , i.e.,

$$\hat{\gamma} := \underset{\gamma}{\operatorname{argmin}} \sup_{\mu_w \in \mathcal{M}} \text{MSE}(\hat{\tau}(\gamma) \mid \{X_i\}_{i=1}^n). \quad (27)$$

The two most widely considered classes to date (Armstrong and Kolesár, 2018; Kolesár and Rothe, 2018; Imbens and Wager, 2019) are

$$\mathcal{M}_{\text{T2},M} = \left\{ \mu_w(\cdot) : |\mu_w(x) - \mu_w(c) - \mu'_w(c)x| \leq M(x - c)^2/2 \text{ for every } x, \text{ and } w = 0, 1 \right\}, \quad (28)$$

$$\mathcal{M}_{\text{H2},M} = \left\{ \mu_w(\cdot) : \mu'_w(x) \text{ is } M\text{-Lipschitz for all } x \text{ and } w = 0, 1 \right\}, \quad (29)$$

where $M > 0$ is a constant. The first option, considered in early work by (Armstrong and Kolesár, 2018), justifies accuracy of a 2nd-order Taylor expansion at c and allows for explicit, closed-form solutions to the minimax linear estimation problem. The second option, which simply takes all

functions to belong to a 2nd-order Hölder class, was used by (Kolesár and Rothe, 2018) and (Imbens and Wager, 2019) and also adopted for the RDHonest package by (Armstrong and Kolesár, 2018).

Fixing a value of M , we numerically solve the optimization problem (27) and obtain an upper bound on the worst-case bias of the resulting estimator as a by-product. We then follow (Imbens and Manski, 2004) to construct a bias-aware confidence interval by inflating the Gaussian confidence interval accounting for the worst-case bias of $\tau(\hat{\gamma})$ under the function class $\mathcal{M}$. These bias-aware confidence intervals provide uniformly valid coverage in finite samples, provided the function class $\mathcal{M}$ and the distribution of Y_i given X_i is Gaussian (Armstrong and Kolesár, 2018).¹⁵ In practice, however, smoothness constant M is typically unknown; both (Armstrong and Kolesár, 2020) and (Imbens and Wager, 2019) suggest using some rule-of-thumb estimates.

C.3 The Choice of Function Class

One of the key insights our paper provides is that the choice of the function class $\mathcal{M}$ in the min-max problem (8) plays a crucial role in determining the conservativeness (and thus the practical usefulness) of the resulting bias-aware method. The function classes used in the existing bias-aware literature (*e.g.*, the classes $\mathcal{M}_{T2,M}$ and $\mathcal{M}_{H2,M}$ defined in (28) and (29) respectively) are much modest in terms of assumed smoothness as compared to the the bias-corrected methods used in practice (see, *e.g.*, (Calónico et al., 2014, Theorem 1)). Here, on the other hand, use the class of all functions $\mu_w(\cdot)$ that satisfy Assumption 2 and for which $\mu''_0(\cdot)$ is $\hat{B}$ -Lipschitz, where we obtain $\hat{B}$ via a global cubic regression and cross-fitting. While this might appear as a minor modification from a theoretical point of view, it materially affects the width of our intervals in practice, as we demonstrate in Tables 1 to 3.

To illustrate the difference between the conservativeness of our function class versus the function classes used by existing bias-aware approaches, we plot in Figure 3 the regression functions corresponding to the worst-case bias obtained by our method and $\mathcal{M}_{H2,M}$ as used in RDHonest. This plot demonstrates that the classes $\mathcal{M}_{H2,M}$ can often be unreasonably broad, and include regression functions that are *ex ante* implausible or rarely occur in practice. This explains why the resulting bias-aware confidence intervals can be substantially more conservative than the bias-corrected ones (see, *e.g.*, Tables 1 to 3). The empirical success of our PLRD method that we demonstrate in Section 4 is, in fact, rooted in this choice of the function class over which one computes the worst-case MSE.

C.4 Ablation Analysis

Here we present an empirical comparison of our PLRD method with other alternative methods. We summarize our empirical results in Table 4, which complements the results in Table 1 of the main paper. Table 4 compares our PLRD method with four other alternative approaches:

- (a) Bias-aware intervals for local linear regression with inflated critical values as proposed by (Armstrong and Kolesár, 2020):¹⁶
 - (i) Fit a local linear regression using the `rdrobust` package, with default bandwidth selection rule.

¹⁵If the conditional distribution of Y_i is non-Gaussian, one can recover exact finite-sample results using the bias-aware approach by using finite-sample concentration results instead of Gaussian approximation.

¹⁶Our end-to-end implementation differs from that in (Armstrong and Kolesár, 2020), as we choose bandwidths for local linear regression using the bandwidth selector of (Imbens and Kalyanaraman, 2012) as implemented in `rdrobust`; whereas (Armstrong and Kolesár, 2020) discuss adaptive bandwidth choice using an alternative rule of thumb. Here, we chose to pick bandwidths as in `rdrobust` to enable a cleaner comparison to the conventional estimator.




Figure 3: Functions $\mu_w(\cdot)$ with worst-case bias in the RDHonest and PLRD function classes. The top two plots use the m_{quad} example from (Imbens and Kalyanaraman, 2012), while the bottom two plots use the data from (Lee, 2008). In each plot, black lines show $\mu_w(\cdot)$ functions with maximum (solid) or minimum (dotted) bias in the respective function class. This figure illustrates how the RDHonest function class may accommodate in-hindsight-unreasonable functions for several practical settings.

- (ii) Obtain the LLR estimator $\hat{\tau}_{\text{LLR}}$ with its conventional standard error estimate $\widehat{\text{SE}}_{\text{LLR}}$
- (iii) Build the confidence interval $\hat{\tau}_{\text{LLR}} \pm 2.181 \cdot \widehat{\text{SE}}_{\text{LLR}}$.
- (b) Bias-aware intervals for local quadratic regression with inflated critical values as proposed by (Armstrong and Kolesár, 2020):
 - (i) Fit a local quadratic regression using the `rdrobust` package (i.e., with $p = 2$ in Appendix C.1), with default bandwidth selection rule.
 - (ii) Obtain the LQR estimator $\hat{\tau}_{\text{LQR}}$ (defined as $\hat{\tau}_2(h_n)$ in Appendix C.1) with its conventional standard error estimate $\widehat{\text{SE}}_{\text{LQR}}$
 - (iii) Build the confidence interval $\hat{\tau}_{\text{LQR}} \pm 2.113 \cdot \widehat{\text{SE}}_{\text{LQR}}$.
- (c) PLRD without Assumption 2: Bias-aware inference operating under the class

$$\mathcal{M}_{\text{H3}, B} = \left\{ \mu_w(\cdot) : \mu_w''(x) \text{ is } M\text{-Lipschitz for all } x \text{ and } w = 0, 1 \right\}.$$

Dataset	Alt. (a)			Alt. (b)			Alt. (c)			Alt. (d)			plrd		
	cvg	width	rmse	cvg	width	rmse	cvg	width	rmse	cvg	width	rmse	cvg	width	rmse
Lee08	96.7%	0.040	0.010	95.8%	0.049	0.012	98.4%	0.044	0.009	99.9%	0.041	0.007	98.4%	0.044	0.009
CFT15	96.6%	5.513	1.316	96.0%	6.791	1.661	98.3%	5.643	1.189	99.9%	9.731	1.535	96.2%	3.865	0.937
LM07	96.5%	5.870	1.599	96.0%	7.175	2.003	97.6%	6.311	1.670	87.9%	3.526	1.182	96.5%	3.746	1.144
Mey14	96.1%	7.021	1.732	95.5%	8.696	2.173	97.0%	6.503	1.498	97.7%	7.590	1.644	96.7%	4.682	1.159
Mat08 (read.)	87.2%	0.056	0.018	95.0%	0.072	0.018	96.7%	0.059	0.014	94.6%	0.085	0.022	96.7%	0.059	0.014
Mat08 (math)	94.7%	0.049	0.013	94.9%	0.059	0.015	97.6%	0.051	0.011	92.1%	0.074	0.022	97.6%	0.051	0.011
JL04 (read.)	96.1%	0.077	0.020	94.9%	0.101	0.027	94.9%	0.083	0.021	97.3%	0.072	0.016	97.2%	0.078	0.019
JL04 (math)	96.5%	0.065	0.016	95.1%	0.084	0.022	95.0%	0.069	0.018	96.9%	0.070	0.016	97.6%	0.060	0.013
Ore06	96.3%	0.212	0.055	96.5%	0.400	0.123	98.5%	0.245	0.052	98.3%	0.180	0.038	95.9%	0.124	0.031
AMT22 (4th grade)	96.5%	0.115	0.028	94.4%	0.117	0.032	96.8%	0.096	0.022	62.8%	0.086	0.041	96.8%	0.096	0.022
AMT22 (8th grade)	96.6%	0.121	0.030	96.1%	0.147	0.037	96.8%	0.108	0.026	96.5%	0.079	0.019	94.9%	0.071	0.018
AMT22 (HM)	97.0%	0.066	0.016	96.0%	0.081	0.020	97.5%	0.060	0.014	97.1%	0.045	0.010	95.9%	0.040	0.010
Pure noise	96.2%	1.611	0.394	94.9%	2.053	0.525	97.1%	1.678	0.374	98.3%	1.049	0.223	97.2%	1.011	0.228
Setting 1	93.5%	0.238	0.063	94.8%	0.301	0.075	95.2%	0.250	0.062	88.9%	0.169	0.054	94.8%	0.237	0.061
Setting 2	87.9%	0.330	0.100	94.0%	0.313	0.080	97.3%	0.389	0.090	83.9%	0.468	0.178	97.3%	0.389	0.090
Setting 3	94.9%	0.252	0.063	95.2%	0.300	0.074	99.6%	0.290	0.050	99.9%	0.421	0.061	99.6%	0.290	0.050
Setting 4	96.2%	0.235	0.058	95.2%	0.300	0.074	99.3%	0.276	0.052	100.0%	0.368	0.048	96.1%	0.178	0.045

Table 4: Comparison of PLRD with other alternatives: (a) and (b) correspond to inflating conventional LLR and LQR confidence intervals by multiplying the standard error by 2.181 and 2.113, respectively (Armstrong and Kolesár, 2020); (c) and (d) correspond to PLRD minus Assumption 2 or Assumption 1, respectively. Other details of this table are same as for Table 1 in the main text.

This method uses the additional smoothness just as PLRD, while allowing different curvatures around the treatment threshold.

- (i) Estimate B using a global cubic regression on either side of the threshold, in the same manner as we estimate B via global cubic regression in Proposition 2.
- (ii) Follow (Armstrong and Kolesár, 2018) and (Imbens and Wager, 2019) to solve the minimax problem (8) with the function class $\mathcal{M}_{H_3, \hat{B}}$ via numerical convex optimization and compute the minimax linear estimator for τ_c within this new class.
- (iii) As a by-product of the optimization problem, obtain a bound on the worst-case bias of our estimator within this class, and construct “bias aware” confidence intervals for τ_c using the (Imbens and Manski, 2004) construction.

Algorithmically, this is similar to our Algorithm 1; we also implemented this as a part of our `plrd` package.

- (d) PLRD without Assumption 1: Bias-aware intervals operating under the intersection of the class $\mathcal{M}_{H_2, M}$ and the class of regression functions satisfying Assumption 2. In essence, this method combines `optrdd` (Imbens and Wager, 2019) with Assumption 2.
 - (i) Here we estimate M using a global quadratic regression on either side of the threshold, in the same manner as we estimate B via global cubic regression in Proposition 2.
 - (ii) Follow (Armstrong and Kolesár, 2018) and (Imbens and Wager, 2019) to solve the minimax problem (8) with the above function class via numerical convex optimization and compute the minimax linear estimator for τ_c within this new class.
 - (iii) Obtain an upper bound on the worst-case bias of our estimator within this class, and construct “bias aware” confidence intervals for τ_c using the (Imbens and Manski, 2004) construction.

Algorithmically, this is similar to the `optrdd` package by (Imbens and Wager, 2019).¹⁷ We also provide an implementation of this method as a part of our replication files.

The empirical results demonstrate that the rule-of-thumb methods (AK20 (a) and AK20 (b)) achieve valid coverage but at the cost of wider intervals, mirroring our earlier findings in Tables 1 and 3. The relaxed variants of PLRD offer insights into the role of each assumption: Dropping either Assumption 1 or Assumption 2 generally leads to more conservative intervals, suggesting that both assumptions contribute to the efficiency of PLRD.

D Additional Details on Real Data Experiments

Here we provide detailed descriptions of the real datasets analyzed in the main paper. We reference each dataset by the paper that introduced it, same as in Tables 1, 2 and 4.

1. (Lee, 2008): This dataset, frequently used for illustrating regression discontinuity methods, examines the incumbency advantage in U.S. House elections. The running variable is the margin of victory, defined as the difference in vote share between the Democratic candidate and their strongest opponent, with the threshold set at $c = 0$. The outcome of interest is the probability of winning subsequent elections. The dataset consists of 6,558 observations from U.S. House elections between 1946 and 1998. (Lee, 2008) reported a significant positive effect of incumbency on the probability of winning subsequent elections, indicating an incumbency advantage; see Table 1 and Figure 2 of (Lee, 2008) for more details.
2. (Cattaneo et al., 2015): This is similar to (Lee, 2008) except that they study the incumbency advantage in U.S. Senate elections instead of House elections. The running variable is the margin of victory, defined as the difference in vote percentages (ranging from -100 to 100) between the Democratic candidate and their strongest opponent, with the threshold set at $c = 0$. The outcome of interest is the percentage of winning subsequent elections. The dataset consists of 1,390 observations (1,297 excluding missing outcomes) from U.S. House elections between 1914 and 2010. Our findings align with the original paper (Cattaneo et al., 2015) as well as the replication in (Cattaneo et al., 2024).
3. (Matsudaira, 2008): This dataset records fifth-grade students' math and reading scores in an unspecified large school district in the northeast from 2001 to 2002. (Matsudaira, 2008) examined the impact of mandatory summer school on students' performance. The running variable is the minimum of students' test scores in math and reading, with a score below the passing threshold for either math or reading determining mandatory attendance in the summer school. The primary outcomes are post-summer school reading and math test scores. We use a version of this dataset shared by the author with K. Kalyanamaran and G. Imbens in 2008. This dataset has a slightly larger sample size $n = 68,798$ than the one used in the original paper ($n = 66,839$). Our findings mirror the results in (Matsudaira, 2008, Table 4) for fifth-grade students.
4. (Jacob and Lefgren, 2004): This dataset contains the record of third-grade student achievement in Chicago Public Schools from 1997 to 1999. (Jacob and Lefgren, 2004) investigated the effectiveness of remedial education programs on student achievement, using a regression

¹⁷In fact, over the weaker smoothness class $\mathcal{M}_{H2, M}$, Assumption 2 on partial linearity becomes much less useful, because the second derivative of the baseline effect can have (bounded) sudden jumps anyways. As such, we also expect numerical performance of this ablation baseline to be close to `optrdd`.

discontinuity design based on test score. We use the same dataset as (Imbens and Zajonc, 2009), which contains 70,831 observations after dropping entries with missing information. The running variable is the minimum of the math and reading score for an exam in the spring and the outcome is the post-remedial program reading (resp. math) test scores, reported as Rasch test scores rather than grade equivalents. Rasch test scores are normalized scores (zero mean, unit s.d.) for the population of third-grade students from 1997 to 1999. Our results are aligned with the findings in (Jacob and Lefgren, 2004) for third-grade students.

5. (Oreopoulos, 2006): This study concerns the impact of compulsory schooling laws on educational attainment. The running variable is the individual’s year of birth, with thresholds determined by changes in compulsory schooling laws. The outcome here is years of completed education. Our dataset comprises 73,954 observations, differing slightly from the original study due to variations in preprocessing; again, we use the same dataset also used in (Kolesár and Rothe, 2018) and focus on the reduced form analysis rather than a fuzzy RD approach (see Table 1, Panel C of the [corrigendum](#) to (Oreopoulos, 2006)). Our findings are in line with Table 2 of (Kolesár and Rothe, 2018).
6. (Ludwig and Miller, 2007): This study examines the effectiveness of the Head Start program on children’s health and schooling. The running variable is the 1960 poverty rate, and the threshold corresponds to the qualifying cutoff for program participation. Following (Calonico et al., 2014), we consider as outcome variable the mortality rates per 100,000 for children between 5 and 9 years old, with Head Start-related causes, for 1973–1983. After cleaning missing values, this dataset includes 2,783 observations. Table III and Figure IV (Panel A) in (Ludwig and Miller, 2007) illustrate that the Head Start program had a negative impact on the outcomes.
7. (Meyersson, 2014): This study investigates the effect of Islamic party control in local governments on women’s rights, in particular on the educational attainment of young women, focusing on Turkey’s 1994 mayoral elections. The study employs an RD design where the running variable is the Islamic margin of victory—the vote share gap between the leading Islamic and secular parties, thus with a natural cutoff of zero. The analysis is conducted at the municipality level and has 2,629 observations. We replicate part of the study where the primary outcome is the school attainment for women who were (potentially) in high school during the period 1994–2000, measured with variables extracted from the 2000 census. This data have been used as a running example in (Cattaneo et al., 2024). Our results align with the illustrations in (Cattaneo et al., 2024) and the original paper (Meyersson, 2014) — in municipalities where the Islamic party barely won, the educational attainment of women is roughly 3% higher than in municipalities where the party barely lost.
8. (Akhtari et al., 2022): This paper studies the effect of political turnover in Brazilian mayoral elections on the quality of public services—in particular, on municipal education. The running variable is the incumbent party’s vote margin (the vote-share gap between the incumbent party and its strongest opponent), with a natural cutoff at zero. We run the WGAN experiments with three outcomes: The 4th-grade and 8th-grade student test scores following Noack et al. (2026a), and headmaster replacement in municipal schools following Calonico et al. (2025).

E Implementation via Convex Optimization

In this section, we discuss how one can solve the optimization problem (8) using readily available software described in, *e.g.*, (Boyd and Vandenberghe, 2004). Essentially, our task is to solve quadratic minimization problems of the form:¹⁸

$$\underset{\gamma}{\text{minimize}} \left[\sup_{\mu_w \in \mathcal{M}_B} \left(\sum_{i=1}^n \gamma_i \mu_{W_i}(X_i) - (\mu_1(c) - \mu_0(c)) \right)^2 + \sum_{i=1}^n \sigma_i^2 \gamma_i^2 \right]. \quad (30)$$

Here, we take $\mathcal{M}_B$ of all conditional response functions $\mu_w(\cdot)$ satisfying Assumption 2 and such that $\mu_0''(x)$ is globally B -Lipschitz. Since γ_i are measurable functions of $\{X_i\}_{i=1}^n$, the bias of the estimator $\hat{\tau}(\gamma) = \sum_{i=1}^n \gamma_i Y_i$ conditional on the running variable is given by

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^n \gamma_i Y_i \mid \{X_i\}_{i=1}^n \right] - \tau_c &= \sum_{i=1}^n \gamma_i \mu_{W_i}(X_i) - \tau_c \\ &= \sum_{i=1}^n \gamma_i (W_i(\tau_c + \beta(X_i - c)) + \mu_0(X_i)) - \tau_c \\ &= \sum_{i=1}^n \gamma_i \left(\mu_0(X_i) - \mu_0(c) - \mu_0'(c)(X_i - c) - \frac{1}{2} \mu_0''(c)(X_i - c)^2 \right), \end{aligned}$$

provided

$$\sum_{i=1}^n \gamma_i = 0, \quad \sum_{i=1}^n \gamma_i W_i = 1, \quad \sum_{i=1}^n \gamma_i W_i (X_i - c) = 0, \quad \sum_{i=1}^n \gamma_i (1 - W_i)(X_i - c) = 0, \quad \sum_{i=1}^n \gamma_i (X_i - c)^2 = 0,$$

which are implicit constraints in our optimization procedure. Indeed, at values of γ for which these constraints are not all satisfied, we can construct $\mu_w(\cdot) \in \mathcal{M}_B$ that forces the squared conditional bias to blow up to ∞ . Thus, the solution $\hat{\gamma}$ to (30) must satisfy the constraints. Next, define

$$\rho(x) := \frac{1}{B} \left(\mu_0(x + c) - \mu_0(c) - \mu_0'(c)x - \frac{1}{2} \mu_0''(c)x^2 \right).$$

Note that for any $\mu_w(\cdot) \in \mathcal{M}_B$, the function $\rho(\cdot)$ belongs to the function class

$$\mathcal{R} := \{f : f(0) = f'(0) = f''(0) = 0, f'' \text{ is 1-Lipschitz}\}.$$

We can therefore rewrite the optimization problem (30) as:

$$\hat{\gamma} := \underset{\gamma}{\text{minimize}} \left[B^2 \left(\sup_{\rho \in \mathcal{R}} \sum_{i=1}^n \gamma_i \rho(X_i - c) \right)^2 + \sum_{i=1}^n \sigma_i^2 \gamma_i^2 \right].$$

¹⁸Here, the σ_i^2 should simply be taken as observation-level weights in an optimization problem. If we set $\sigma_i^2 = \text{Var}(Y_i \mid X_i)$ to be the conditional variance of the response Y_i given the running variable X_i , then the problem below is exactly the minimax linear estimation problem. When we run this the optimization as specified in Algorithm 1, we set $\sigma_i^2 = \hat{\sigma}^2$ via cross-fitting, as described in Section 2.

Finally, we make the implicit constraints explicit and use an auxiliary variable t and rewrite the above optimization problem as:

$$\begin{aligned}
(\hat{\gamma}, \hat{t}) &:= \underset{\gamma, t}{\text{minimize}} \left\{ B^2 t^2 + \sum_{i=1}^n \gamma_i^2 \sigma_i^2 \right\} \\
&\text{subject to } \sup_{\rho \in \mathcal{R}} \sum_{i=1}^n \gamma_i \rho(X_i - c) \leq t, \quad \sum_{i=1}^n \gamma_i W_i = 1, \quad \sum_{i=1}^n \gamma_i (1 - W_i) = -1, \\
&\quad \sum_{i=1}^n \gamma_i (X_i - c) = 0, \quad \sum_{i=1}^n \gamma_i (1 - W_i) (X_i - c) = 0, \quad \sum_{i=1}^n \gamma_i (X_i - c)^2 = 0.
\end{aligned} \tag{31}$$

An immediate consequence of this modified optimization problem is that we get an upper bound on the conditional bias of our PLRD estimator. Next, we discuss how one can solve the above problem using convex duality. The dual problem is given by

$$\begin{aligned}
&\underset{\rho(\cdot), \lambda}{\text{maximize}} \inf_{\gamma, t} \left\{ \sum_{i=1}^n \gamma_i^2 \sigma_i^2 + B^2 t^2 \right. \\
&\quad + \lambda_1 \left(\sum_{i=1}^n \gamma_i \rho(X_i - c) - t \right) + \lambda_2 \left(\sum_{i=1}^n \gamma_i W_i - 1 \right) + \lambda_3 \left(\sum_{i=1}^n \gamma_i (1 - W_i) + 1 \right) \\
&\quad \left. + \lambda_4 \sum_{i=1}^n \gamma_i W_i (X_i - c) + \lambda_5 \sum_{i=1}^n \gamma_i (1 - W_i) (X_i - c) + \lambda_6 \sum_{i=1}^n \gamma_i (X_i - c)^2 \right\} \\
&\text{subject to } \lambda_1 \geq 0, \lambda_2, \lambda_3, \lambda_4, \lambda_5, \lambda_6 \in \mathbb{R}, \text{ and } \rho \in \mathcal{R}.
\end{aligned}$$

The inner minimization problem is simply a quadratic one, which we can explicitly solve to reduce the dual problem to the following.

$$\hat{\theta} = \underset{\theta = (\tilde{\rho}, \lambda) \in \Theta}{\text{argmin}} \left(\frac{1}{4n} \sum_{i=1}^n \frac{G^2(\theta; X_i - c, W_i)}{\sigma_i^2} + \frac{\lambda_1^2}{4nB^2} + \frac{\lambda_2 - \lambda_3}{n} \right), \tag{32}$$

where $\theta = (\tilde{\rho}, \lambda)$ denotes the dual variable that varies in the set

$$\Theta := \left\{ (f, \lambda) \mid f : [-\ell, \ell] \rightarrow \mathbb{R}, f(0) = f'(0) = f''(0) = 0, f'' \text{ is } \lambda_1\text{-Lipschitz}, \lambda \in [0, \infty) \times \mathbb{R}^5 \right\}, \tag{33}$$

and

$$G(\tilde{\rho}, \lambda; x, w) := \tilde{\rho}(x) + \lambda_2 w + \lambda_3 (1 - w) + \lambda_4 w x + \lambda_5 (1 - w) x + \lambda_6 x^2.$$

The function $\tilde{\rho}(\cdot)$ in the above formulation corresponds to the function $\lambda_1 \rho(\cdot)$ in the original problem. Thus, we can recover the optimal variables $(\hat{\gamma}, \hat{t})$ for the primal problem (31) from the dual optimal variable $\hat{\theta}$ (i.e., solution of (32)), as follows.

$$\hat{\gamma}_i = -\frac{1}{2\sigma_i^2} G(\hat{\theta}; X_i - c, W_i), \quad \hat{t} = \frac{1}{2B^2} \hat{\lambda}_1.$$

Since our optimization problem differs from (Imbens and Wager, 2019) essentially in the choice of the function class, (Imbens and Wager, 2019, Proposition 2) provides a basis for approximating (32) using a finite-dimensional quadratic program, in a way that the resulting discrete solutions

converge uniformly to the continuous solution as the grid size becomes small. In our software implementation, we adopt this finite-dimensional approximation approach to compute the optimal weights using `quadprog`, a function for quadratic programming in R (R Core Team, 2023).

F Supporting Lemmas

To set the stage, we recall the key concepts of entropy with and without bracketing for function classes, which will be central to the results presented in this section. Consider a class of real-valued functions $\mathcal{F}$ defined on a set $\mathcal{Z}$, equipped with a norm $\|\cdot\|$. The *covering number* $N(\varepsilon, \mathcal{F}, \|\cdot\|)$ is the minimum number of balls of radius ε needed to cover $\mathcal{F}$. The centers of the covering balls need not belong to $\mathcal{F}$, but they are required to have finite norms. The *entropy (without bracketing)* is defined as the logarithm of the covering number. An ε -bracket is a pair of functions (l, u) such that $\|u - l\| < \varepsilon$ and $l \leq f \leq u$ for all $f \in \mathcal{F}$. The *bracketing number* $N_{[]}(\varepsilon, \mathcal{F}, \|\cdot\|)$ is the minimum number of ε -brackets needed to cover $\mathcal{F}$. The upper and lower bounds u and l of the brackets need not belong to $\mathcal{F}$, but they are required to have finite norms. The *entropy with bracketing* is defined as the logarithm of the bracketing number.

Following (van der Vaart and Wellner, 1996, Section 2.7), we consider the class of functions $f : \mathcal{Z} \rightarrow \mathbb{R}$ defined on a bounded set $\mathcal{Z} \subset \mathbb{R}$ such that the derivatives of f up to order k are uniformly bounded, and $f^{(k)}$ is Lipschitz continuous. We can define a norm on this class as follows:

$$\|f\|_{k+1} := \max_{j \leq k} \sup_x |f^{(j)}(x)| + \sup_{x, y} \frac{|f^{(k)}(x) - f^{(k)}(y)|}{|x - y|},$$

where the last supremum is over all x, y in the interior of $\mathcal{Z}$ with $x \neq y$. Finally, denote by $\mathcal{C}_M^k(\mathcal{Z})$ the set of all continuous functions $f : \mathcal{Z} \rightarrow \mathbb{R}$ with $\|f\|_k \leq M$. Throughout, we use $A \lesssim_\eta B$ to denote that $A \leq CB$ for some arbitrary constant $C > 0$ that can depend on η . For a vector $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_p)$ we denote $\lambda_{-1} := (\lambda_2, \dots, \lambda_p)$. We use $\|x\|_\infty$ to denote $\max_i |x_i|$ when x is a vector, and $\sup_t |x(t)|$ when x is a function.

Lemma 7. *Consider the following class of functions:*

$$\mathcal{F}_{B, \ell, k} = \{\rho \mid \rho : [-\ell, \ell] \rightarrow \mathbb{R}, \rho(0) = \rho'(0) = \dots = \rho^{(k)}(0) = 0, \rho^{(k)} \text{ is } B\text{-Lipschitz}\},$$

where $k \geq 0$ is an integer and $0 < \ell \leq 1$. It holds that

$$\log N(\varepsilon, \mathcal{F}_{B, \ell, k}, \|\cdot\|_\infty) \lesssim \ell(B/\varepsilon)^{1/(k+1)}, \quad \text{and} \quad \log N_{[]}(\varepsilon, \mathcal{F}_{B, \ell, k}, L_r(P)) \lesssim \ell(B/\varepsilon)^{1/(k+1)},$$

for any $r \geq 1$ and probability measure P on $\mathbb{R}$. Thus, $\mathcal{F}_{B, \ell, k}$ is universally Donsker.

Proof of Lemma 7. For any $\rho \in \mathcal{F}_{B, \ell, k}$, the derivatives of ρ upto order k are all bounded by $B\ell$. Thus, $\|\rho\|_{k+1} \leq B(1 + \ell) \leq 2B$. Consequently,

$$\begin{aligned} N(\varepsilon, \mathcal{F}_{B, \ell, k}, \|\cdot\|_\infty) &\leq N(\varepsilon, \mathcal{C}_{2B}^{k+1}([-\ell, \ell]), \|\cdot\|_\infty) \\ &\leq N(\varepsilon/2B, \mathcal{C}_1^{k+1}([-\ell, \ell]), \|\cdot\|_\infty). \end{aligned}$$

For any $\delta > 0$, note that $N(\delta, [-\ell, \ell], |\cdot|) \lesssim \ell/\delta$. We can thus slightly tweak the proof of (van der Vaart and Wellner, 1996, Theorem 2.7.1), with $\alpha = k + 1$ and $\lambda(\mathcal{X}_1) = 2\ell$ (in their notation), to

conclude that

$$\log N(\varepsilon/2B, \mathcal{C}_1^{k+1}([- \ell, \ell]), \|\cdot\|_\infty) \lesssim \ell(\varepsilon/2B)^{-1/(k+1)},$$

as desired to show. This gives us the first entropy bound (without bracketing).

Next, we can adapt the proof technique of (van der Vaart and Wellner, 1996, Corollary 2.7.2) to deduce that the last display also yields the desired bound on entropy with bracketing. Finally, the bound on $\log N_{[]}(\varepsilon, \mathcal{F}_{B, \ell, k}, L_2(P))$ implies that $\mathcal{F}_{B, \ell, k}$ is P -Donsker for any probability distribution P (see, e.g., (van der Vaart, 1998, Theorem 19.5)). $\square$

For the rest of this section, we assume without loss of generality that the treatment threshold is at $c = 0$. Before stating the next few lemmas, we recall some notation. The empirical dual problem is given by

$$\hat{\theta}_n(\sigma^2, B) = \operatorname{argmin}_{\theta=(\rho, \lambda) \in \Theta(\ell)} \hat{L}_n(\theta; B), \quad \hat{L}_n(\theta; B) := \frac{1}{4n} \sum_{i=1}^n \frac{G^2(\theta; X_i, W_i)}{\sigma_i^2} + \frac{\lambda_1^2}{4nB^2} + \frac{\lambda_2 - \lambda_3}{n},$$

where $G(\theta; x, w) = G(\rho, \lambda; x, w) = \rho(x) + \lambda_2 w + \lambda_3(1 - w) + \lambda_4 wx + \lambda_5(1 - w)x + \lambda_6 x^2$, with $w = 1(x > 0)$, and

$$\Theta(\ell) := \{(f, \lambda) \mid f : [-\ell, \ell] \rightarrow \mathbb{R}, f(0) = f'(0) = f''(0) = 0, f'' \text{ is } \lambda_1\text{-Lipschitz}, \lambda \in [0, \infty) \times \mathbb{R}^5\}.$$

Define the dual problem at the population level as

$$\theta_n^*(\sigma^2, B) = \operatorname{argmin}_{\theta \in \Theta(\ell)} L_n(\theta; \sigma^2, B), \quad \text{where} \quad L_n(\theta; \sigma^2, B) := \frac{1}{4\sigma^2} \mathbb{E} [G^2(\theta; X, W)] + \frac{\lambda_1^2}{4nB^2} + \frac{\lambda_2 - \lambda_3}{n}.$$

In Algorithm 1, we sample splitting to estimate the unknown Lipschitz constant B and the homoskedastic variance $\sigma^2 = \operatorname{Var}(Y_i \mid X_i)$. However, for ease of exposition, we assume throughout this proof that an estimate $\hat{B}$ of B and an estimate $\hat{\sigma}^2$ of σ^2 is available independently of the data using which we solve the empirical dual problem. We define

$$\hat{\theta}_n = \hat{\theta}_n(\hat{\sigma}^2, \hat{B}), \quad \text{and} \quad \theta_n^* = \theta_n^*(\hat{\sigma}^2, B).$$

Lemma 8. *Define*

$$\Theta_n(\ell, \delta, \varepsilon) := \{(\rho, \lambda) \mid (\rho, \lambda) \in \Theta(\ell), |\lambda_1 - \lambda_{n,1}^*| \leq \delta \lambda_{n,1}^* \text{ and } \|\lambda_{-1} - \lambda_{n,-1}^*\|_2 \leq \varepsilon \|\lambda_{n,-1}^*\|_2\},$$

where $\ell, \delta, \varepsilon \in (0, 1)$ are fixed. Then, the following maximal inequality holds.

$$\mathbb{E} \sup_{\theta \in \Theta_n(\ell, \delta, \varepsilon)} |\hat{L}_n(\theta; \hat{\sigma}^2, \hat{B}) - L_n(\theta; \hat{\sigma}^2, \hat{B})| \lesssim (\lambda_{n,1}^*)^2 \frac{\ell^6}{\sqrt{n}} + \|\lambda_{n,-1}^*\|_2^2 \frac{\sqrt{\varepsilon}}{\sqrt{n}} + \lambda_{n,1}^* \|\lambda_{n,-1}^*\|_2 \frac{\ell^3}{\sqrt{n}}.$$

Proof of Lemma 8. We begin by noting that

$$\left| \hat{L}_n(\theta; \hat{\sigma}^2, \hat{B}) - L_n(\theta; \hat{\sigma}^2, \hat{B}) \right| = \frac{1}{4\hat{\sigma}^2} |(\mathbb{P}_n - \mathbb{P})G^2(\theta; X, W)|.$$

We prove a maximal inequality for the above RHS using tools from empirical process theory. Write $G(\rho, \lambda; x, w) = \rho(x) + \lambda^\top v(x, w)$ where $v(x, w) := (w, 1 - w, wx, (1 - w)x, x^2)^\top$, and note that

$\|v(x, w)\|_2 \leq \sqrt{1 + \ell + \ell^2} \leq 2$. To reduce notation, we write $\Theta_n = \Theta_n(\ell, \delta, \varepsilon)$ for the rest of this proof. Using triangle inequality,

$$\begin{aligned} \sup_{\theta \in \Theta_n} |(\mathbb{P}_n - \mathbb{P})G^2(\theta; X, W)| &\leq \sup_{(\rho, \lambda) \in \Theta_n} |(\mathbb{P}_n - \mathbb{P})\rho^2(X)| + \sup_{(\rho, \lambda) \in \Theta_n} |(\mathbb{P}_n - \mathbb{P})(\lambda_{-1}^\top v(X, W))^2| \\ &\quad + 2 \sup_{(\rho, \lambda) \in \Theta_n} |(\mathbb{P}_n - \mathbb{P})\rho(X)\lambda_{-1}^\top v(X, W)|. \end{aligned}$$

We will bound each of the above three quantities in order. First, note that for any $(\rho, \lambda) \in \Theta_n$, $\rho^{(j)}(0) = 0$ for $j = 0, 1, 2$, and ρ'' is $B_1 := (1 + \delta)\lambda_{n,1}^*$ -Lipschitz. This gives $|\rho''(x)| = |\rho''(x) - \rho''(0)| \leq B_1 \ell$. Moreover, a Taylor expansion of $\rho(x)$ about $x = 0$ yields the following

$$|\rho(x)| = |\rho(x) - \rho(0) - \rho'(0)x| = \frac{|\rho''(\xi) - \rho''(0)|}{|\xi - x|} |\xi - x| \frac{x^2}{2} \leq B_1 \frac{\ell^3}{2}. \quad (34)$$

A similar argument gives $|\rho'(x)| \leq B_1 \ell^2$. Putting these all together, we can say that the function $r = \rho^2$ satisfies $r(0) = 0$, $r'(0) = 0$, $r''(x) = 2\rho''(x)\rho(x) + 2\rho'(x)^2$, and thus

$$\|r''\|_\infty \lesssim B_1^2 \ell^4 \lesssim (\lambda_{n,1}^*)^2 \ell^4.$$

Consequently, using the notation of Lemma 7, we can say that $\mathcal{A} := \{\rho^2 : (\rho, \lambda) \in \Theta_n \text{ for some } \lambda\}$ is a subset of the class $\mathcal{F}_{\ell, h, 1}$ with Lipschitz constant $B \lesssim (\lambda_{n,1}^*)^2 \ell^4$. Therefore, Lemma 7 implies that

$$\log N_{[]}(\varepsilon', \mathcal{A}, L_2(P)) \leq \log N_{[]}(\varepsilon', \mathcal{F}_{\ell, h, 1}, L_2(P)) \leq \ell(B/\varepsilon')^{1/2} \lesssim \lambda_{n,1}^* \ell^3 (\varepsilon')^{-1/2}.$$

Finally, note that (34) implies that the constant function $F = (\lambda_{n,1}^* \ell^3)^2$ works as an envelope to the class $\mathcal{A}$. We can now apply (van der Vaart, 1998, Corollary 19.35) to conclude that

$$\begin{aligned} \mathbb{E} \sup_{(\rho, \lambda) \in \Theta_n} |(\mathbb{P}_n - \mathbb{P})\rho^2(X)| &= \mathbb{E} \sup_{r \in \mathcal{A}} |(\mathbb{P}_n - \mathbb{P})r| \\ &\lesssim \frac{1}{\sqrt{n}} \int_0^{\|F\|_{P,2}} \sqrt{\log N_{[]}(\varepsilon', \mathcal{A}, L_2(P))} d\varepsilon' \\ &\lesssim \frac{1}{\sqrt{n}} \int_0^F (\varepsilon')^{-1/4} (\lambda_{n,1}^* \ell^3)^{1/2} d\varepsilon' \\ &\lesssim \frac{1}{\sqrt{n}} F^{3/4} F^{1/4} = (\lambda_{n,1}^*)^2 \frac{\ell^6}{\sqrt{n}}. \end{aligned}$$

Next, we deal with

$$\sup_{(\rho, \lambda) \in \Theta_n} |(\mathbb{P}_n - \mathbb{P})(\lambda_{-1}^\top v(X, W))^2| = \sup_{\lambda \in \Lambda_n} |(\mathbb{P}_n - \mathbb{P})(\lambda_{-1}^\top v(X, W))^2|,$$

where $\Lambda_n = \Lambda_n(\varepsilon) = \{v \in \mathbb{R}^5 : \|v - \lambda_{n,-1}^*\| \leq \varepsilon \|\lambda_{n,-1}^*\|\} \subseteq [0, \infty) \times \mathbb{R}^5$. Note, for any $\lambda \in \Lambda_n$,

$$|\lambda_{-1}^\top v(x, w)| \leq \|\lambda_{-1}\|_2 \|v(x, w)\|_2 \leq (1 + \varepsilon) \|\lambda_{n,-1}^*\|_2 \sqrt{1 + \ell + \ell^2} \leq 4 \|\lambda_{n,-1}^*\|_2.$$

Therefore for any fixed x, w , and $\lambda, \tilde{\lambda} \in \Lambda_n$,

$$\begin{aligned} |\langle \lambda, v(x, w) \rangle^2 - \langle \tilde{\lambda}, v(x, w) \rangle^2| &\leq 8\|\lambda_{n,-1}^*\|_2 |\langle \lambda, v(x, w) \rangle - \langle \tilde{\lambda}, v(x, w) \rangle| \\ &\leq 16\|\lambda_{n,-1}^*\|_2 \|\lambda - \tilde{\lambda}\|_2, \end{aligned}$$

showing that the function $\lambda \mapsto (\lambda^\top v(x, w))^2$ (where $\lambda \in \Lambda_n$) is c -Lipschitz, for $c = 16\|\lambda_{n,-1}^*\|$. Consequently,

$$\begin{aligned} \log N_{[]}(\varepsilon', \Lambda_n, L_2(P)) &\leq 5 \log \left(1 + \frac{2c \operatorname{diam}(\Lambda_n)}{\varepsilon'} \right) \\ &\leq \frac{10c \operatorname{diam}(\Lambda_n)}{\varepsilon'} \lesssim \varepsilon \|\lambda_{n,-1}^*\|_2^2 (\varepsilon')^{-1}. \end{aligned}$$

We can now apply (van der Vaart, 1998, Corollary 19.35) with envelop $(4\|\lambda_{n,-1}^*\|_2)^2$ to conclude that

$$\begin{aligned} \mathbb{E} \sup_{(\rho, \lambda) \in \Theta_n} |(\mathbb{P}_n - \mathbb{P})(\lambda_{-1}^\top v(X, W))^2| &\lesssim \frac{1}{\sqrt{n}} \int_0^{16\|\lambda_{n,-1}^*\|^2} \sqrt{\log N_{[]}(\varepsilon', \Lambda_n, L_2(P))} d\varepsilon' \\ &\lesssim \|\lambda_{n,-1}^*\|_2^2 \frac{\sqrt{\varepsilon}}{\sqrt{n}}. \end{aligned}$$

It only remains to prove a maximal inequality for the following supremum:

$$\sup_{(\rho, \lambda) \in \Theta_n} |(\mathbb{P}_n - \mathbb{P})\rho(X)\lambda_{-1}^\top v(X, W)|.$$

Recall from (34) that $|\rho(x)| \leq B_1 \ell^3 / 2 \leq \lambda_{n,1}^* \ell^3$ for any $(\rho, \lambda) \in \Theta_n$. Also recall from above that $|\langle \lambda_{-1}, v(x, w) \rangle| \leq 4\|\lambda_{n,-1}^*\|_2$. Consequently, for any $(\rho, \lambda) \in \Theta_n$, and for every x, w , we have

$$|\rho(x)\lambda_{-1}^\top v(x, w)| \leq 4\lambda_{n,1}^* \|\lambda_{n,-1}\|_2 \ell^3,$$

proving that the constant function $4\lambda_{n,1}^* \|\lambda_{n,-1}\|_2 \ell^3$ works as an envelope for this class.

Next, we fix $\varepsilon' > 0$. Suppose that $\mathcal{R}_\varepsilon := \{\rho_1, \dots, \rho_{N_1}\}$ is a minimal ε_1 -cover for $\pi_1(\Theta_n) = \{\rho : (\rho, \lambda) \in \Theta_n \text{ for some } \lambda\}$ under the sup-norm, and $\mathcal{L}_\varepsilon := \{B_1, \dots, B_{N_2}\}$ is a minimal ε_2 -cover for $\Lambda_n = \{v \in \mathbb{R}^5 : \|v - \lambda_{n,-1}^*\|_2 \leq \varepsilon \|\lambda_{n,-1}^*\|_2\}$ under the ℓ_2 -norm, where $\varepsilon_1 := \varepsilon' / 8 \|\lambda_{n,-1}^*\|_2$ and $\varepsilon_2 = \varepsilon' / 4 \lambda_{n,1}^* \ell^3$. It is well-known (see, *e.g.*, (van der Vaart, 1998)) that,

$$\log N_2 \leq 5 \log \left(1 + \frac{c\varepsilon \|\lambda_{n,-1}^*\|}{\varepsilon' / 4 \lambda_{n,1}^* \ell^3} \right) \lesssim \varepsilon \lambda_{n,1}^* \|\lambda_{n,-1}^*\| \ell^3 (\varepsilon')^{-1}.$$

On the other hand, Lemma 7 yields

$$\log N_1 \lesssim (\varepsilon' / 8 \|\lambda_{n,-1}\|_2)^{-1/3} (2\lambda_{n,1}^*)^{1/3} \ell = (\lambda_{n,1}^* \|\lambda_{n,-1}^*\|_2)^{1/3} \ell (\varepsilon')^{-1/3}.$$

Now, for any $\theta = (\rho, \lambda) \in \Theta_n$, pick $\theta^0 = (\rho^0, \lambda^0) \in \mathcal{R}_\varepsilon \times \mathcal{L}_\varepsilon$ such that $\|\rho - \rho^0\|_\infty \leq \varepsilon_1$ and

$\|\lambda - \lambda^0\|_2 \leq \varepsilon_2$. Note that,

$$\begin{aligned}
& |\rho(x)\langle\lambda_{-1}, v(x, w)\rangle - \rho^0(x)\langle\lambda_{-1}^0, v(x, w)\rangle| \\
& \leq |\rho(x)| \cdot |\langle\lambda_{-1}, v(x, w)\rangle - \langle\lambda_{-1}^0, v(x, w)\rangle| + |\langle\lambda_{-1}^0, v(x, w)\rangle| \cdot |\rho(x) - \rho^0(x)| \\
& \leq \lambda_{n,1}^* \ell^3 \cdot \|\lambda_{-1} - \lambda_{-1}^0\|_2 \|v(x, w)\|_2 + 4\|\lambda_{n,-1}^*\|_2 \cdot |\rho(x) - \rho^0(x)| \\
& \leq \lambda_{n,1}^* \ell^3 \cdot \frac{\varepsilon'}{4\lambda_{n,1}^* \ell^3} \cdot 2 + 4\|\lambda_{n,-1}^*\|_2 \cdot \frac{\varepsilon'}{8\|\lambda_{n,-1}^*\|_2} \leq \varepsilon'.
\end{aligned}$$

Therefore, the class $\mathcal{B}$ of functions $(x, w) \mapsto \rho(x)\langle\lambda_{-1}, v(x, w)\rangle$, where $(\rho, \lambda) \in \Theta_n$, satisfies

$$\log N(\varepsilon', \mathcal{B}, \|\cdot\|_\infty) \leq \log N_1 N_2 \lesssim (\varepsilon')^{-1/3} (\lambda_{n,1}^* \|\lambda_{n,-1}^*\|_2)^{1/3} \ell + (\varepsilon')^{-1} \lambda_{n,1}^* \|\lambda_{n,-1}^*\| \ell^3 \varepsilon.$$

This also gives us a bound for the entropy with bracketing (cf. (van der Vaart and Wellner, 1996, Corollary 2.7.2)), as follows.

$$\begin{aligned}
\log N_{[]}(\varepsilon', G(\Theta_n), L_2(P)) & \leq \log N(2\varepsilon', G(\Theta_n), |\cdot|_\infty) \\
& \lesssim (\varepsilon')^{-1/3} (\lambda_{n,1}^* \|\lambda_{n,-1}^*\|_2 \ell^3)^{1/3} + (\varepsilon')^{-1} \lambda_{n,1}^* \|\lambda_{n,-1}^*\| \ell^3 \varepsilon.
\end{aligned}$$

Recall from above that the constant function $F_{ab} := 4\lambda_{n,1}^* \|\lambda_{n,-1}^*\|_2 \ell^3$ acts as an envelop function for the class $\mathcal{B}$. We now apply (van der Vaart, 1998, Corollary 19.35) once more to conclude that

$$\begin{aligned}
\mathbb{E} \sup_{(\rho, \lambda) \in \Theta_n} |(\mathbb{P}_n - \mathbb{P})\rho(X)\lambda_{-1}^\top v(X, W)| & \lesssim \frac{1}{\sqrt{n}} \int_0^{F_{ab}} \left((\varepsilon')^{-1/6} F_{ab}^{1/6} + (\varepsilon')^{-1/2} \varepsilon^{1/2} F_{ab}^{1/2} \right) d\varepsilon' \\
& \lesssim \frac{1}{\sqrt{n}} \left((F_{ab})^{5/6} (F_{ab})^{1/6} + (F_{ab})^{1/2} \varepsilon^{1/2} (F_{ab})^{1/2} \right) \\
& \lesssim \lambda_{n,1}^* \|\lambda_{n,-1}^*\|_2 \frac{\ell^3}{\sqrt{n}},
\end{aligned}$$

which completes the proof. $\square$

Lemma 9. Let $G(\rho, \lambda; x, w) := \rho(x) + \lambda_{-1}^\top v(x, w)$, where $v(x, w) := (w, 1 - w, xw, x(1 - w), x^2)$. Assume that the assumptions in Theorem 1 hold. Then, there exists $0 < \alpha \leq 1$ such that, for any $u \in \mathbb{R}^5$,

$$u^\top \mathbb{E} \left[v(X, W) v(X, W)^\top \right] u \geq \alpha \|u\|_2^2.$$

Further, $\theta_n^*(\sigma^2, B) = \operatorname{argmin}_{\theta \in \Theta(\ell)} L_n(\theta; \sigma^2, B)$ satisfies the following.

$$\begin{aligned}
L_n(\theta; \sigma^2, B) - L_n(\theta_n^*; \sigma^2, B) & \geq \frac{1}{4\sigma^2} \mathbb{E} \left(\rho(X) - \rho_n^*(X) + \langle \lambda_{-1} - \lambda_{n,-1}^*, v(X, W) \rangle \right)^2 + \frac{(\lambda_1 - \lambda_{n,1}^*)^2}{4nB^2} \\
& \geq \frac{\alpha}{4\sigma^2} \left(\|\rho - \rho_n^*\|_{L^2(P)}^2 + \|\lambda_{-1} - \lambda_{n,-1}^*\|_2^2 \right) + \frac{(\lambda_1 - \lambda_{n,1}^*)^2}{4nB^2} - R_n,
\end{aligned}$$

where $R_n = \frac{2}{4\sigma^2} \mathbb{E} |(\rho(X) - \rho_n^*(X)) \langle \lambda_{-1} - \lambda_{n,-1}^*, v(X, W) \rangle|$.

Proof of Lemma 9. Define $A := \mathbb{E}[v(X, W)v(X, W)^\top]$. Since A is positive semidefinite, it has non-negative eigenvalues. We first show that the smallest eigenvalue of A , which we denote as $\lambda_{\min}(A)$,

is strictly positive. To prove this, assume to the contrary that $\lambda_{\min}(A) = 0$. Then, there exists $u \in \mathbb{R}^5$, $u \neq 0$, such that $u^\top A u = 0$. We can conclude from $u^\top A u = \mathbb{E}(v(X, W)^\top u)^2 = 0$ that $v(X, W)^\top u = 0$ with probability 1. This implies that, either X is a solution of $u_1 + u_3x + u_5x^2 = 0$ (when $w = \mathbf{1}\{x > 0\} = 1$), or a solution of $u_2 + u_4x + u_5x^2 = 0$ (when $w = \mathbf{1}\{x > 0\} = 0$). But, X having a finite support contradicts the assumption that the running variable has a strictly positive density at the cutoff. So we must have $\lambda_{\min}(A) > 0$. Therefore,

$$u^\top \mathbb{E} \left[v(X, W) v(X, W)^\top \right] u \geq \alpha \|u\|_2^2,$$

for any $u \in \mathbb{R}^5$, where $\alpha := \min\{1, \lambda_{\min}(A)\}$.

For proving the second part, we write $L_n = L_n(\cdot; \sigma^2, B)$ for notational convenience. First, we simplify $L_n(\theta) - L_n(\theta_n^*)$ using the fact that gradients of the loss L_n at θ_n^* are zero. To remove the dependence of the domain on the parameter λ_1 , we write

$$L_n(\rho, \lambda) = \frac{1}{4\sigma^2} \mathbb{E} \left(\lambda_1 \tilde{\rho}(X) + \lambda_{-1}^\top v(X, W) \right)^2 + \frac{\lambda_1^2}{4nB^2} + \frac{\lambda_2 - \lambda_3}{n} =: \tilde{L}_n(\tilde{\rho}, \lambda),$$

where $\tilde{\rho}(\cdot) = \rho(\cdot)/\lambda_1 \in \mathcal{R}$. Since $\tilde{L}_n$ is minimized at $(\tilde{\rho}_n^*, \lambda_n^*)$, we can deduce the following.

1. $\nabla_{\lambda_{-1}} \tilde{L}_n(\tilde{\rho}_n^*, \lambda_n^*) = 0$. This implies that,

$$\frac{1}{2\sigma^2} \mathbb{E} \left[(\lambda_{n,1}^* \tilde{\rho}_n^*(X) + \langle \lambda_{n,-1}^*, v(X, W) \rangle) v(X, W) \right] + \frac{e_1 - e_2}{n} = 0. \quad (35)$$

where $e_1 = (1, 0, 0, 0, 0)$ and $e_2 = (0, 1, 0, 0, 0)$.

2. $\nabla_{\lambda_1} \tilde{L}_n(\tilde{\rho}_n^*, \lambda_n^*) = 0$. This implies that,

$$\frac{1}{2\sigma^2} \mathbb{E} \left[(\lambda_{n,1}^* \tilde{\rho}_n^*(X) + \langle \lambda_{n,-1}^*, v(X, W) \rangle) \tilde{\rho}_n^*(X) \right] + \frac{2\lambda_{n,1}^*}{4nB^2} = 0. \quad (36)$$

3. $\nabla_{\tilde{\rho}} \tilde{L}_n(\tilde{\rho}_n^*, \lambda_n^*) = 0$. We can say that the convex function $t \mapsto \tilde{L}_n(t\tilde{\rho} + (1-t)\tilde{\rho}_n^*, \lambda_n^*)$ is minimized at $t = 0$, which gives us the following.

$$\frac{2\lambda_{n,1}^*}{4\sigma^2} \mathbb{E} \left[(\tilde{\rho}(X) - \tilde{\rho}_n^*(X)) (\lambda_{n,1}^* \tilde{\rho}_n^*(X) + \langle \lambda_{n,-1}^*, v(X, W) \rangle) \right] \geq 0. \quad (37)$$

Equipped with Equations (35) to (37), we are ready to simplify $L_n(\theta) - L_n(\theta_n^*) = \tilde{L}_n(\tilde{\rho}, \lambda) - \tilde{L}_n(\tilde{\rho}_n^*, \lambda_n^*)$, as follows.

$$\begin{aligned} L_n(\theta) - L_n(\theta_n^*) &= \frac{1}{4\sigma^2} \mathbb{E} (\rho(X) - \rho_n^*(X) + \langle \lambda_{-1} - \lambda_{n,-1}^*, v(X, W) \rangle)^2 \\ &= \frac{2}{4\sigma^2} \mathbb{E} [(\rho - \rho_n^*)(X) (\rho_n^*(X) + \langle \lambda_{n,-1}^*, v(X, W) \rangle)] + \frac{\lambda_1^2 - (\lambda_{n,1}^*)^2}{4nB^2} \\ &\quad + \frac{2}{4\sigma^2} \mathbb{E} [\langle \lambda_{-1} - \lambda_{n,-1}^*, v(X, W) \rangle (\rho_n^*(X) + \langle \lambda_{n,-1}^*, v(X, W) \rangle)] + \frac{\lambda_2 - \lambda_3 - \lambda_{n,2}^* + \lambda_{n,3}^*}{n}. \end{aligned}$$

Using (35), we deduce that the second line in the above display is zero. Next, set $\tilde{\rho} = \rho/\lambda_1$,

$\tilde{\rho}_n^* = \rho_n^*/\lambda_{n,1}^*$. Continuing from above,

$$\begin{aligned}
&= \frac{2}{4\sigma^2} \mathbb{E} [(\lambda_1 \tilde{\rho}(X) - \lambda_{n,1}^* \tilde{\rho}_n^*(X)) (\rho_n^*(X) + \langle \lambda_{n,-1}^*, v(X, W) \rangle)] + \frac{\lambda_1^2 - (\lambda_{n,1}^*)^2}{4nB^2} \\
&= \frac{2\lambda_1}{4\sigma^2} \mathbb{E} [(\tilde{\rho}(X) - \tilde{\rho}_n^*(X)) (\rho_n^*(X) + \langle \lambda_{n,-1}^*, v(X, W) \rangle)] \\
&\quad + \frac{2}{4\sigma^2} \mathbb{E} [(\lambda_1 - \lambda_{n,1}^*) \tilde{\rho}_n^*(X) (\rho_n^*(X) + \langle \lambda_{n,-1}^*, v(X, W) \rangle)] + \frac{\lambda_1^2 - (\lambda_{n,1}^*)^2}{4nB^2} \\
&\geq 0 - \frac{2\lambda_{n,1}^*}{4nB^2} (\lambda_1 - \lambda_{n,1}^*) + \frac{\lambda_1^2 - (\lambda_{n,1}^*)^2}{4nB^2} \quad (\text{using (37) and (36)}) \\
&\geq \frac{(\lambda_1 - \lambda_{n,1}^*)^2}{4nB^2}.
\end{aligned}$$

Therefore,

$$\begin{aligned}
L_n(\theta) - L_n(\theta_n^*) &\geq \frac{1}{4\sigma^2} \mathbb{E} (\rho(X) - \rho_n^*(X) + \langle \lambda_{-1} - \lambda_{n,-1}^*, v(X, W) \rangle)^2 + \frac{(\lambda_1 - \lambda_{n,1}^*)^2}{4nB^2} \\
&\geq \frac{1}{4\sigma^2} \left[\mathbb{E} (\rho(X) - \rho_n^*(X))^2 + (\lambda_{-1} - \lambda_{n,-1}^*)^\top A (\lambda_{-1} - \lambda_{n,-1}^*) \right] - R_n + \frac{(\lambda_1 - \lambda_{n,1}^*)^2}{4nB^2} \\
&\geq \frac{\alpha}{4\sigma^2} \left(\|\rho - \rho_n^*\|_{L^2(P)}^2 + \|\lambda_{-1} - \lambda_{n,-1}^*\|_2^2 \right) + \frac{(\lambda_1 - \lambda_{n,1}^*)^2}{4nB^2} - R_n,
\end{aligned}$$

as desired to show. □

Lemma 10. *The optimizer $\theta_n^* = (\rho_n^*, \lambda_n^*)$ of the dual population problem (15) satisfies*

$$\|\lambda_{n,-1}^*\|_2 = O(\lambda_{n,1}^* \ell^3 + n^{-1}).$$

Consequently,

$$\sup_{\theta \in \Theta_n(\ell, \delta, \varepsilon)} |L_n(\theta; \hat{\sigma}^2, \hat{B}) - L_n(\theta; \hat{\sigma}^2, B)| \lesssim \frac{(\lambda_{n,1}^*)^2}{n} \left| \frac{1}{\hat{B}} - \frac{1}{B} \right|.$$

Proof of Lemma 10. Recall from (35) in the proof of Lemma 9 that

$$\frac{1}{2\sigma^2} \mathbb{E} [(\lambda_{n,1}^* \tilde{\rho}_n^*(X) + \langle \lambda_{n,-1}^*, v(X, W) \rangle) v(X, W)] + \frac{e_1 - e_2}{n} = 0.$$

Rearranging this equation yields

$$A\lambda_{n,-1}^* = -\lambda_{n,1}^* \mathbb{E} \tilde{\rho}_n^*(X) v(X, W) - 2\sigma^2 \frac{(e_1 - e_2)}{n}$$

where $A := \mathbb{E}[v(X, W)v(X, W)^\top]$, as defined in the proof of Lemma 9. Since $\|v(X, W)\|_2 \leq 2$ and (34) yields $\|\tilde{\rho}_n^*\|_\infty \leq \ell^3/2$, we deduce that the above right hand side is $O(\lambda_{n,1}^* \ell^3 + 1/n)$, which finishes the proof of the first part. The second conclusion follows immediately from the definition once we use $\|\lambda_{n,-1}^*\|_2 = O(\lambda_{n,1}^* \ell^3 + 1/n)$. □

Lemma 11. Consider the neighborhood $\Theta_n(\ell, \delta, \varepsilon)$ of the optimizer $\theta_n^* = \theta_n^*(\hat{\sigma}^2, B)$ of the population dual problem, as defined in Lemma 8. Denote by $\hat{\theta}_n = \hat{\theta}_n(\hat{\sigma}^2, \hat{B})$ the optimizer of the empirical dual problem. Assume that $\hat{B} \rightarrow_p B$, and that $\ell = o(n^{-1/12})$. For any $\delta, \eta \in (0, 1)$, and $\varepsilon = O(n^{-1/4})$, it holds that

$$\mathbb{P}(\hat{\theta}_n \in \Theta_n(\ell, \delta, \varepsilon)) \geq 1 - \eta,$$

for all sufficiently large n .

Proof of Lemma 11. It follows from Lemma 8 that for any $n \geq 1$,

$$\mathbb{E} \sup_{\theta \in \Theta_n(\ell, \delta, \varepsilon)} |\hat{L}_n(\theta; \hat{\sigma}^2, \hat{B}) - L_n(\theta; \hat{\sigma}^2, \hat{B})| \lesssim (\lambda_{n,1}^*)^2 \frac{\ell^6}{\sqrt{n}} + \|\lambda_{n,-1}^*\|_2^2 \frac{\sqrt{\varepsilon}}{\sqrt{n}} + \lambda_{n,1}^* \|\lambda_{n,-1}^*\|_2 \frac{\ell^3}{\sqrt{n}},$$

where $\Theta_n(\ell, \delta, \varepsilon) := \{(\rho, \lambda) \in \Theta(\ell) \text{ such that } |\lambda_1 - \lambda_{n,1}^*| \leq \delta \lambda_{n,1}^* \text{ and } \|\lambda_{-1} - \lambda_{n,-1}^*\|_2 \leq \varepsilon \|\lambda_{n,-1}^*\|_2\}$. Invoking Lemma 10, we reduce the above to the following.

$$\mathbb{E} \sup_{\theta \in \Theta_n(\ell, \delta, \varepsilon)} |\hat{L}_n(\theta; \hat{\sigma}^2, \hat{B}) - L_n(\theta; \hat{\sigma}^2, \hat{B})| \lesssim (\lambda_{n,1}^*)^2 \frac{\ell^6}{\sqrt{n}} + \lambda_{n,1}^* \frac{\ell^3}{n^{3/2}} + \frac{\sqrt{\varepsilon}}{n^{5/2}}.$$

Fix $\eta \in (0, 1)$. It follows from above that for some absolute constant $C = C(\eta) > 0$, we have

$$\mathbb{P} \left(\sup_{\theta \in \Theta_n(\ell, \delta, \varepsilon)} |\hat{L}_n(\theta; \hat{\sigma}^2, \hat{B}) - L_n(\theta; \hat{\sigma}^2, \hat{B})| \leq C \left((\lambda_{n,1}^*)^2 \frac{\ell^6}{\sqrt{n}} + \lambda_{n,1}^* \frac{\ell^3}{n^{3/2}} + \frac{\sqrt{\varepsilon}}{n^{5/2}} \right) \right) \geq 1 - \frac{\eta}{2}, \quad (38)$$

On the other hand, Lemma 10 also tells us that

$$\sup_{\theta \in \Theta_n(\ell, \delta, \varepsilon)} |L_n(\theta; \hat{\sigma}^2, \hat{B}) - L_n(\theta; \hat{\sigma}^2, B)| \leq C \frac{(\lambda_{n,1}^*)^2}{n} \left| \frac{1}{\hat{B}} - \frac{1}{B} \right|.$$

Consequently,

$$\mathbb{P} \left(\sup_{\theta \in \Theta_n(\ell, \delta, \varepsilon)} |L_n(\theta; \hat{\sigma}^2, \hat{B}) - L_n(\theta; \hat{\sigma}^2, B)| \leq \delta^2 \frac{(\lambda_{n,1}^*)^2}{16nB^2} \right) \geq 1 - \frac{\eta}{2},$$

for all sufficiently large n . This, combined with (38) implies that,

$$\mathbb{P} \left(\sup_{\theta \in \Theta_n(\ell, \delta, \varepsilon)} |\hat{L}_n(\theta; \hat{\sigma}^2, \hat{B}) - L_n(\theta; \hat{\sigma}^2, B)| \leq C \left(\frac{(\lambda_{n,1}^*)^2 \ell^6}{\sqrt{n}} + \frac{\lambda_{n,1}^* \ell^3}{n^{3/2}} + \frac{\sqrt{\varepsilon}}{n^{5/2}} \right) + \frac{\delta^2 (\lambda_{n,1}^*)^2}{16nB^2} \right) \geq 1 - \eta, \quad (39)$$

for all sufficiently large n . Next, we recall from Lemma 9 that

$$L_n(\theta; \hat{\sigma}^2, B) - L_n(\theta_n^*; \hat{\sigma}^2, B) \geq \frac{\alpha}{4\hat{\sigma}^2} \left(\|\rho - \rho_n^*\|_{L^2(P)}^2 + \|\lambda_{-1} - \lambda_{n,-1}^*\|_2^2 \right) + \frac{(\lambda_1 - \lambda_{n,1}^*)^2}{4nB^2} - R_n, \quad (40)$$

for some $\alpha \in (0, 1)$, where

$$R_n = \frac{2}{4\hat{\sigma}^2} \mathbb{E} |(\rho(X) - \rho_n^*(X)) \langle \lambda_{-1} - \lambda_{n,-1}^*, v(X, W) \rangle|.$$

We define another neighborhood of $\theta_n^* = (\rho_n^*, \lambda_n^*)$, as follows:

$$\tilde{\Theta}_n := \{(\rho, \lambda) \in \Theta(\ell) : \|\rho - \rho_n^*\|_{L^2(P)}^2 + \|\lambda_{-1} - \lambda_{n,-1}^*\|_2^2 \leq \varepsilon^2 \|\lambda_{n,-1}^*\|_2^2 \text{ and } |\lambda_1 - \lambda_{n,1}^*| \leq \delta \lambda_{n,1}^*\}.$$

For $(\rho, \lambda) \in \tilde{\Theta}_n$, it follows from the Cauchy-Schwarz inequality that

$$R_n \leq \frac{1}{4\hat{\sigma}^2} \left(\|\rho - \rho_n^*\|_{L^2(P)}^2 + \|\langle \lambda_{n,-1} - \lambda_{n,-1}^*, v(X, W) \rangle\|_2^2 \right) \leq \frac{\beta}{4\hat{\sigma}^2} \varepsilon^2 \|\lambda_{n,-1}^*\|_2^2, \quad (41)$$

where $\beta = \max\{1, \lambda_{\max}\}$, where $\lambda_{\max}$ is the largest eigenvalue of $A = \mathbb{E}[v(X, W)v(X, W)^\top]$. It is immediate from definition that $\tilde{\Theta}_n \subseteq \Theta_n(\ell, \delta, \varepsilon)$, and that for every θ on the boundary $\partial\tilde{\Theta}_n$,

$$\|\rho - \rho_n^*\|_{L^2(P)}^2 + \|\lambda_{-1} - \lambda_{n,-1}^*\|_2^2 = \varepsilon^2 \|\lambda_{n,-1}^*\|_2^2, \quad \text{and} \quad |\lambda_1 - \lambda_{n,1}^*| = \delta \lambda_{n,1}^*.$$

We can now combine (39), (40) and (41) to conclude that, for any $\theta \in \partial\tilde{\Theta}_n$,

$$\begin{aligned} & \frac{\alpha}{4\hat{\sigma}^2} \varepsilon^2 \|\lambda_{n,-1}^*\|_2^2 + \frac{\delta^2 (\lambda_{n,1}^*)^2}{4nB^2} - \frac{\beta}{4\hat{\sigma}^2} \varepsilon^2 \|\lambda_{n,-1}^*\|_2^2 \\ & \leq L_n(\theta; \hat{\sigma}^2, B) - L_n(\theta_n^*; \hat{\sigma}^2, B) \quad (\text{using (40) and (41)}) \\ & \leq \hat{L}_n(\theta; \hat{\sigma}^2, \hat{B}) - \hat{L}_n(\theta_n^*; \hat{\sigma}^2, \hat{B}) + 2 \sup_{\theta \in \Theta_n(\ell, \delta, \varepsilon)} |\hat{L}_n(\theta; \hat{\sigma}^2, \hat{B}) - L_n(\theta; \hat{\sigma}^2, B)| \quad (42) \\ & \leq \hat{L}_n(\theta; \hat{\sigma}^2, \hat{B}) - \hat{L}_n(\theta_n^*; \hat{\sigma}^2, \hat{B}) + C \left((\lambda_{n,1}^*)^2 \frac{\ell^6}{\sqrt{n}} + \lambda_{n,1}^* \frac{\ell^3}{n^{3/2}} + \frac{\sqrt{\varepsilon}}{n^{5/2}} \right) + \delta^2 \frac{(\lambda_{n,1}^*)^2}{8nB^2}, \quad (\text{using (39)}) \end{aligned}$$

with probability at least $1 - \eta$, for any $n \geq 1$. We can rewrite the above inequality as $\hat{L}_n(\theta_n^*; \hat{\sigma}^2, \hat{B}) \leq \hat{L}_n(\theta; \hat{\sigma}^2, \hat{B}) + \Delta_n$ and we show below that $\Delta_n < 0$ for all sufficiently large n . This implies, by taking infimum over $\theta \in \partial\tilde{\Theta}_n$, that

$$\mathbb{P} \left(\hat{L}_n(\theta_n^*; \hat{\sigma}^2, \hat{B}) \leq \inf_{\theta \in \partial\tilde{\Theta}_n} \hat{L}_n(\theta; \hat{\sigma}^2, \hat{B}) \right) \geq 1 - \eta,$$

for all sufficiently large n . Finally, note that since the empirical loss function $\hat{L}_n(\theta; \hat{\sigma}^2, \hat{B})$ is convex in θ , the above event implies that $\hat{\theta}_n = \operatorname{argmin}_{\theta} \hat{L}_n(\theta; \hat{\sigma}^2, \hat{B})$ must be inside $\tilde{\Theta}_n$. Thus,

$$\mathbb{P} \left(\hat{\theta}_n \in \tilde{\Theta}_n \right) \geq 1 - \eta,$$

for all sufficiently large n , as desired to show.

To complete the proof, we need to show that $\Delta_n < 0$ for all sufficiently large n . Observe that,

$$\begin{aligned} \Delta_n &= C \left((\lambda_{n,1}^*)^2 \frac{\ell^6}{\sqrt{n}} + \lambda_{n,1}^* \frac{\ell^3}{n^{3/2}} + \frac{\sqrt{\varepsilon}}{n^{5/2}} \right) + \delta^2 \frac{(\lambda_{n,1}^*)^2}{8nB^2} - \left(\frac{\delta^2 (\lambda_{n,1}^*)^2}{4nB^2} - \frac{\beta - \alpha}{4\hat{\sigma}^2} \varepsilon^2 \|\lambda_{n,-1}^*\|_2^2 \right) \\ &\leq C \left((\lambda_{n,1}^*)^2 \frac{\ell^6}{\sqrt{n}} + \lambda_{n,1}^* \frac{\ell^3}{n^{3/2}} + \frac{\sqrt{\varepsilon}}{n^{5/2}} \right) - \delta^2 \frac{(\lambda_{n,1}^*)^2}{8nB^2} + c_0 \varepsilon^2 \left(\lambda_{n,1}^* \ell^3 + \frac{1}{n} \right)^2 \\ &= \frac{(\lambda_{n,1}^*)^2}{n} \left(C \ell^6 \sqrt{n} + c_0 \varepsilon^2 \ell^6 n - \frac{\delta^2}{8B^2} + \frac{C \ell^3}{\lambda_{n,1}^* \sqrt{n}} + 2c_0 \varepsilon^2 \frac{\ell^3}{\lambda_{n,1}^*} + \frac{C \sqrt{\varepsilon}}{n^{3/2} (\lambda_{n,1}^*)^2} + \frac{c_0 \varepsilon^2}{n (\lambda_{n,1}^*)^2} \right). \end{aligned}$$

We show in Lemma 12 that $\lambda_{n,1}^* = O(n^{-3/7})$. In light of this, we can write

$$\limsup_{n \rightarrow \infty} \frac{\Delta_n}{(\lambda_{n,1}^*)^2/n} \leq \lim_{n \rightarrow \infty} \left(C\ell^6 \sqrt{n} + c_0 \varepsilon^2 \ell^6 n - \frac{\delta^2}{8B^2} \right) \leq -\frac{\delta^2}{8B^2} < 0,$$

since $\ell^6 \sqrt{n} = o(1)$ and $\varepsilon^2 \ell^6 n = \varepsilon^2 \sqrt{n} \cdot \ell^6 \sqrt{n} = o(1)$ as well. This completes the proof. $\square$

Lemma 12. *Under the conditions of Theorem 1,*

$$\hat{\lambda}_{n,1}(\hat{\sigma}^2, \hat{B}) = \mathcal{O}_p(n^{-3/7}), \quad \lambda_{n,1}^*(\hat{\sigma}^2, B) = \mathcal{O}_p(n^{-3/7}), \quad \text{and} \quad \sum_{i=1}^n \hat{\gamma}_i^2 = \mathcal{O}_p(n^{-6/7})$$

Proof of Lemma 12. Recall that $\hat{\lambda}_{n,1}(\hat{\sigma}^2, \hat{B}) = 2\hat{B}^2 \hat{t}_n$, where $\hat{B}^2 \hat{t}_n^2$ is bounded above by the worst-case MSE of any linear estimator of the form $\sum_{i=1}^n \gamma_i Y_i$. In particular, there exists a kernel regression estimator that has worst-case MSE (supremum taken over our function class $\mathcal{M}_{\hat{B}}$) of the order of $\mathcal{O}_p(n^{-6/7})$. To make it precise, consider the kernel regression estimator

$$\left(\hat{\tau}(h_n), \hat{\alpha}, \hat{\beta} \right) = \underset{\tau, \alpha, \beta_1, \beta_0 \in \mathbb{R}}{\operatorname{argmin}} \sum_{i=1}^n K\left(\frac{X_i}{h_n}\right) (Y_i - \alpha - \tau W_i - \beta_0 X_i^- - \beta_1 X_i^+ - \beta_2 X_i^2)^2.$$

For simplicity, we use the triangular kernel $K(x) = 1\{|x| \leq 1\}$. It follows along the lines of the proof of (Wager, 2024, Proposition 8.1) that $\hat{\tau}(h_n)$ has worst-case squared curvature bias bounded above by $\hat{B} h_n^6$, and sampling noise of the order of $1/nh_n$. With $h_n = O(n^{-1/7})$, this estimator has MSE of the order of $n^{-6/7}$. Consequently, $\hat{\lambda}_{n,1} = \mathcal{O}_p(n^{-3/7})$ and $\|\hat{\gamma}_n\|_2^2 = \mathcal{O}_p(n^{-6/7})$. The proof for $\lambda_{n,1}^*$ is analogous. $\square$

Lemma 13. *Under the conditions of Theorem 3, it holds that $\hat{b}/s(\hat{\gamma}) = \mathcal{O}_p(1)$ as $n \rightarrow \infty$, where $s^2(\hat{\gamma}) = \sum_{i=1}^n \hat{\gamma}_i^2 \sigma_i^2$, with $\sigma_i^2 = \operatorname{Var}(Y_i | X_i)$.*

Proof of Lemma 13. Assume $c = 0$ without loss of generality. First, we can replace $s^2(\hat{\gamma}) = \sum_{i=1}^n \hat{\gamma}_i^2 \sigma_i^2$ with $\hat{V}_n := \hat{\sigma}^2 \sum_{i=1}^n \hat{\gamma}_i^2$, using the fact that $0 < \sigma_{\min} \leq \sigma_i \leq \sigma_{\max} < \infty$. Recall that, ignoring cross-fitting notation as mentioned in the beginning of this section, we define

$$(\hat{\gamma}_n, \hat{t}_n) := \underset{\gamma, t}{\operatorname{argmin}} \left(\hat{B}^2 t^2 + \hat{\sigma}^2 \sum_{i=1}^n \hat{\gamma}_i^2 \right) \quad \text{subject to} \quad \sup_{\rho \in \mathcal{R}} \sum_{i=1}^n \gamma_i \rho(X_i) \leq t,$$

where $\hat{B}$ and $\hat{\sigma}^2$ are fit externally, independent of the data $\{X_i, Y_i\}_{i=1}^n$. The worst-case bias is given by $\hat{b} = \hat{B} \hat{t}_n$, and thus the goal here is to show that $\hat{t}_n^2 / \hat{V}_n = \mathcal{O}_p(1)$. Recall the population-level problem with externally fit nuisance parameters:

$$(\gamma_n^*, t_n^*) := \underset{\gamma, t}{\operatorname{argmin}} \left\{ \hat{B}^2 t^2 + \frac{1}{n} \hat{\sigma}^2 \mathbb{E}[\gamma^2(X)] \right\} \quad \text{subject to} \quad \sup_{\rho \in \mathcal{R}} \mathbb{E}[\gamma(X) \rho(X)] \leq t.$$

It suffices to prove that $t_n^* / \sqrt{V_n^*} = \mathcal{O}_p(1)$, where $V_n^* = \hat{\sigma}^2 \mathbb{E}[\gamma_n^{*,2}(X)]/n$, since we have consistency (i.e., that $\hat{t}_n/t_n^* \rightarrow_p 1$ and $\hat{V}_n/V_n^* \rightarrow_p 1$). It follows from Lemma 12 that $\hat{B}^2 t_n^{*,2} + V_n^* \leq_p a_+ n^{-6/7}$ for some $a_+ \in (0, \infty)$, for all sufficiently large n . On the other hand, it follows from the minimax

lower bound (see, *e.g.*, (Cheng et al., 1997)) that $\widehat{B}^2 t_n^{*,2} + V_n^* \geq_p a_- n^{-6/7}$ for some $0 < a_- \leq a_+$ for all sufficiently large n . Therefore, to claim that $t_n^*/\sqrt{V_n^*} = \mathcal{O}_p(1)$, it suffices to show that there exists some $c > 0$ such that $V_n^* \geq_p cn^{-6/7}$ for all sufficiently large n .

To prove this by contradiction, assume that for any $c > 0$, we have $V_n^* \leq_p cn^{-6/7}$ infinitely often. We examine the MSE of γ_n^* for the objective

$$R_m(\gamma) := B^2 \left(\sup_{\rho \in \mathcal{R}} \mathbb{E} \gamma(X) \rho(X) \right)^2 + \frac{1}{m} \mathbb{E}[\gamma^2(X)],$$

where $m = n/k$ (assume without loss of generality that n is divisible by k). Note that

$$\begin{aligned} R_m(\gamma_n^*) &= \widehat{B}^2 \left(\sup_{\rho \in \mathcal{R}} \mathbb{E} \gamma_n^*(X) \rho(X) \right)^2 + \frac{1}{m} \mathbb{E}[\gamma_n^{*,2}(X)] \\ &= \widehat{B}^2 (t_n^*)^2 + k V_n^* \leq_p (a_+ + kc) k^{-6/7} m^{-6/7}. \end{aligned}$$

Choose k large enough so that $a_+ k^{-6/7} \leq a_-/2$, and then choose $c > 0$ small enough such that $ck^{1/7} < a_-/2$. This, along with the above display, implies that $R_m(\gamma_m^*) \leq R_m(\gamma_n^*) \leq_p a_- m^{-6/7}$, which contradicts the minimax lower bound (we choose n is large enough so that $m = n/k$ satisfies the minimax lower bound $\widehat{B}^2 t_m^{*,2} + V_m^* \geq_p a_- m^{-6/7}$). This contradiction completes the proof. $\square$

Lemma 14. *Suppose that $Z \sim N(0, 1)$. Then, for any s, h, t and $\delta > 0$, we have*

$$\inf_{|b| \leq t + \delta s} \mathbb{P}(|b + sZ| \leq h) \geq \inf_{|b| \leq t} \mathbb{P}(|b + sZ| \leq h) - \delta.$$

Proof of Lemma 14. Fix any s, h, t and $\delta > 0$. Denote

$$g(b) = \mathbb{P}(|b + sZ| \leq h) = \Phi\left(\frac{h - b}{s}\right) - \Phi\left(\frac{-h - b}{s}\right).$$

Note that,

$$g'(b) = -\frac{1}{s} \left(\phi\left(\frac{h - b}{s}\right) - \phi\left(\frac{-h - b}{s}\right) \right) \implies \|g'\|_\infty \leq \frac{2\phi(0)}{s} \leq \frac{1}{s}.$$

Next, fix any b_1 with $|b_1| \leq t + \delta s$. There exists b_0 such that $|b_0| \leq t$ and $|b_1 - b_0| \leq \delta s$. Consequently, $|g(b_1) - g(b_0)| \leq \|g'\|_\infty |b_1 - b_0| \leq \delta$. This implies that $g(b_1) \geq g(b_0) - \delta \geq \inf_{|b| \leq t} g(b) - \delta$. Taking infimum over b_1 , we get the desired conclusion. $\square$

Lemma 15. *Under the conditions of Theorem 3, it holds that $\widehat{s}(\widehat{\gamma})/s(\widehat{\gamma}) \rightarrow_p 1$ as $n \rightarrow \infty$, where $s^2(\widehat{\gamma}) = \sum_{i=1}^n \widehat{\gamma}_i^2 \sigma_i^2$ with $\sigma_i^2 = \text{Var}(Y_i | X_i)$, and $\widehat{s}^2(\widehat{\gamma}) = \sum_{i=1}^n \widehat{\gamma}_i^2 \widehat{\sigma}_i^2$ with $\widehat{\sigma}_i^2 = (Y_i - \widehat{\mu}_{W_i}(X_i))^2$ as defined in Algorithm 1.*

Proof of Lemma 15. In Algorithm 1, we use sample-splitting and fit regression on either side of the threshold; here we assume for simplicity in notation that we fit the regression of Y_i on X_i globally (our argument applied on either side of the threshold will then complete the proof). Also assume without loss of generality that $c = 0$. Write $\widetilde{X}_i = (1, X_i)$, $\widehat{\varepsilon}_i = Y_i - \widehat{\mu}(X_i) = Y_i - \widetilde{X}_i \widehat{\beta}$ and $\varepsilon_i^* = Y_i - \mu(X_i) = Y_i - \widetilde{X}_i \beta^*$ where $\widetilde{X} \beta^*$ is the best linear projection of $\mu(X) = \mathbb{E}[Y | X]$ onto

$\text{span}\{1, X\}$. It follows from (White, 1980) that $\hat{\beta} \rightarrow_p \beta^*$. Next, write

$$\sum_{i=1}^n \hat{\gamma}_i^2 \varepsilon_i^2 = \sum_{i=1}^n \hat{\gamma}_i^2 \varepsilon_i^{*2} + 2 \sum_{i=1}^n \hat{\gamma}_i^2 \varepsilon_i^* \tilde{X}_i^\top (\beta^* - \hat{\beta}) + \sum_{i=1}^n \hat{\gamma}_i^2 (\tilde{X}_i^\top (\beta^* - \hat{\beta}))^2. \quad (43)$$

On the other hand, note that $\sum_{i=1}^n \hat{\gamma}_i^2 \sigma_i^2 \geq \sigma_{\min}^2 \sum_{i=1}^n \hat{\gamma}_i^2$. Therefore,

$$\frac{\sum_{i=1}^n \hat{\gamma}_i^2 (\tilde{X}_i^\top (\beta^* - \hat{\beta}))^2}{\sum_{i=1}^n \hat{\gamma}_i^2 \sigma_i^2} \lesssim \left\| (\hat{\beta} - \beta^*)^\top \frac{|\hat{\gamma}_i| \tilde{X}_i}{\|\hat{\gamma}_i\|_2} \right\|_2^2 = o_p(1),$$

since $\hat{\gamma}_i = 0$ for $|X_i| > \ell$. For the cross-term, we apply Cauchy-Schwarz inequality to write

$$\frac{\left| \sum_{i=1}^n \hat{\gamma}_i^2 \varepsilon_i^* \tilde{X}_i^\top (\beta^* - \hat{\beta}) \right|}{\sum_{i=1}^n \hat{\gamma}_i^2 \sigma_i^2} \leq \left(\frac{\sum_{i=1}^n \hat{\gamma}_i^2 \varepsilon_i^{*2}}{\sum_{i=1}^n \hat{\gamma}_i^2 \sigma_i^2} \right)^{1/2} \left\| (\hat{\beta} - \beta^*)^\top \frac{|\hat{\gamma}_i| \tilde{X}_i}{\|\hat{\gamma}_i\|_2} \right\|_2.$$

We can handle the second term exactly as above. For the first part, we show below that

$$\sum_{i=1}^n \hat{\gamma}_i^2 \varepsilon_i^{*2} / \sum_{i=1}^n \hat{\gamma}_i^2 \sigma_i^2 \rightarrow_p 1. \quad (44)$$

As a consequence, we can continue from (43) to write

$$\frac{\sum_{i=1}^n \hat{\gamma}_i^2 \varepsilon_i^2}{\sum_{i=1}^n \hat{\gamma}_i^2 \sigma_i^2} = \frac{\sum_{i=1}^n \hat{\gamma}_i^2 \varepsilon_i^{*2}}{\sum_{i=1}^n \hat{\gamma}_i^2 \sigma_i^2} + o_p(1) + o_p(1),$$

which completes the proof. To prove (44), fix any $\varepsilon > 0$ and $\eta > 0$. Define $a_{n,i} = \hat{\gamma}_i^2 / \sum_{i=1}^n \hat{\gamma}_i^2$, $1 \leq i \leq n$, and write $\Delta_i = \varepsilon_i^{*2} - \sigma_i^2$. It follows from the assumption $\sup_x \mathbb{E}[|Y_i - \mu(X_i)|^q | X_i = x] < \infty$ and Jensen's inequality that $\mathbb{E}[|\Delta_i|^{q/2} | X_i = x] \leq C$ for some universal constant $C \in (0, \infty)$. Then,

$$\begin{aligned} & \mathbb{E} \left[\left| \sum_{i=1}^n a_{n,i} \Delta_i \mathbf{1}\{|\Delta_i| > M\} \right| \mid X_1, \dots, X_n \right] \leq \sum_{i=1}^n a_{n,i} \mathbb{E}[|\Delta_i| \mathbf{1}\{|\Delta_i| > M\} \mid X_1, \dots, X_n] \\ & \leq \sum_{i=1}^n a_{n,i} \left(\mathbb{E}[|\Delta_i|^{q/2} \mid X_i] \right)^{2/q} \mathbb{P}(|\Delta_i| > M \mid X_i)^{1-2/q} \quad (\text{H\"older's inequality}) \\ & \leq \sum_{i=1}^n a_{n,i} \frac{1}{M^{q/2-1}} \mathbb{E}[|\Delta_i|^{q/2} \mid X_i] \quad (\text{Markov's inequality}) \\ & \leq \frac{C}{M^{q/2-1}}. \end{aligned}$$

Since $q > 2$, we can pick M large enough so that

$$\mathbb{P} \left(\left| \sum_{i=1}^n a_{n,i} \Delta_i \mathbf{1}\{|\Delta_i| > M\} \right| > \varepsilon/2 \right) < \eta/2,$$

and $|\mathbb{E}[\Delta_i \mathbf{1}\{|\Delta_i| \leq M\}]| \leq \varepsilon/4$ (the latter is possible since $\lim_{M \rightarrow \infty} \mathbb{E}[\Delta_i \mathbf{1}\{|\Delta_i| \leq M\}] = \mathbb{E}[\Delta_i] = 0$

by Dominated Convergence Theorem). Consequently, by Chebyshev's inequality,

$$\mathbb{P}\left(\left|\sum_{i=1}^n a_{n,i} \Delta_i \mathbf{1}_{\{|\Delta_i| \leq M\}}\right| > \varepsilon/2 \mid \{X_i\}\right) \lesssim \sum_{i=1}^n a_{n,i}^2 M^2 \leq M^2 \max_{1 \leq i \leq n} a_{n,i} < \eta/2,$$

for all sufficiently large n , since we have $\max_{i \leq n} a_{n,i} = o_p(1)$ from Corollary 5. Putting these together, we obtain $\sum_{i=1}^n \hat{\gamma}_i^2(\varepsilon_i^{*2} - \sigma_i^2) / \sum_{i=1}^n \hat{\gamma}_i^2 = o_p(1)$. Since $0 \leq \sigma_{\min}^2 \leq \sigma_i^2$, this finishes the proof of (44). $\square$