INSTITUTO DE ECONOMÍA

TESIS de DOCTORADO

2019

Essays in Applied Microeconometrics

Andrés García Echalar

PONTIFICIA UNIVERSIDAD CATOLICA DE CHILE
INSTITUTO DE ECONOMIA
DOCTORADO EN ECONOMIA

TESIS DE GRADO
DOCTORADO EN ECONOMIA

García Echalar Andrés

Enero, 2019

PONTIFICIA UNIVERSIDAD CATOLICA DE CHILE
INSTITUTO DE ECONOMIA
DOCTORADO EN ECONOMIA

ESSAYS IN APPLIED MICROECONOMETRICS

Andrés García Echalar

Comisión

Tomás Rau
Felipe Gonzáles
Claudia Martínez
Sergio Urzúa

Santiago, enero de 2019

Contents

Part I

A Reweighting Approach for Regression Discontinuity Designs with Discontinuous Density of the Running Variable

A Reweighting Approach for Regression Discontinuity Designs with Discontinuous Density of the Running Variable

Andrés García E. *

This version: December 2018

Abstract

In the Regression Discontinuity (RD) design, discontinuities in the density function of the running variable may harm identification and bias estimations as in the manipulation and heaping cases. This paper proposes a new robust approach that consistently estimates the Average Treatment Effect near the cutoff under discontinuities in the distribution function of the running variable. The approach consists of sample-reweighting the outcome prior to the estimation of the causal effect. This paper also discusses the limitations of existing indirect tests about the validity of the RD estimation results and presents sufficient conditions for identification that are directly linked to those tests. Simulated examples are presented to assess the finite sample performance of the proposed approach, while non-simulated examples use real data and compare to existing correction methods that partially identify the effect. The Reweighted-RD design applies to any setting with or without discontinuities in the conditional or marginal distributions while maintaining all the distinctive features of the conventional RD design.

1 Introduction

Since its introduction in 1960 (Thistlethwaite and Campbell, 1960) and later rebirth in the late '1990s (Angrist and Lavy, 1999; Black, 1999; Van der Klaauw, 2002), the Regression Discontinuity (RD) design has become widely used in economics and other social sciences as a non-experimental evaluation strategy. In this design, the probability of receiving treatment changes discontinuously when an observed "running variable" crosses a known threshold.[1] This discontinuity allows researchers to identify and estimate a causal effect by comparing observations above and below the cutoff point (Imbens and Lemieux, 2008).

---

*Pontificia Universidad Católica de Chile, Vicuña Mackenna 4860, Macul, Santiago, Chile, (email: agarcia8@uc.cl).
I am indebted to the guidance of my thesis advisor Tomás Rau. All errors, however, are my own. My thanks to Felipe González, Claudia Martínez, Sergio Urzúa and Francisco Gallego for their useful comments and suggestions. I also thank the doctoral fellowship from CONICyT (2013-63130177). Powered@NLHPC: This research was partially supported by the supercomputing infrastructure of the NLHPC (ECM-02).

[1] The running variable is an observed continuous variable that fully/partially determines treatment assignment based on a cutoff point. It is also known as assignment or forcing variable.

One of the characteristics of the RD design is that identification conditions are considered mild compared to those required by other non-experimental methods (Lee and Lemieux, 2010). Hahn et al. (2001) formally established assumptions for identification — involving continuity of conditional expectations of potential outcomes on the running variable — and defined the Average Treatment Effect at the threshold as the parameter of interest. Among empiricists, identification with an RD design is considered "as good as" identification from a local (i.e. around the threshold) randomized experiment as discussed by Lee (2008).

Unfortunately, the validity of an RD design cannot be directly assessed due to the impossibility of observing design elements such as the potential outcomes, and violations of identification assumptions bias estimations and lead to an inconsistent RD estimator. The most common threat to identification analyzed in the literature is that of discontinuity of the running variable's distribution at the cutoff. Such a discontinuity may arise if individuals manipulate their running variable to self-select into treatment (Lee, 2008; McCrary, 2008). Discontinuities may also appear in the presence of heaping, which occurs when the running variable is self-reported, measured with limited precision, rounded or discretized (Barreca et al., 2015). For both manipulation and heaping, partial identification methods are provided by Gerard et al. (2018) and Barreca et al. (2015) respectively.[2]

This paper proposes a Reweighting approach to Regression Discontinuity (RRD) design that, under regularity conditions, achieves identification regardless of whether or not there are discontinuities in the density function of the running variable or in any other related distribution function. In other words, the proposed approach is robust to distributional discontinuities in a sense that, unlike the standard (conventional) RD design, it consistently estimates the parameter of interest in presence of discontinuities in the conditional and/or marginal distributions; and furthermore, it is also consistent, same as the standard RD estimator, when conventional assumptions are met. Following the idea discussed in DiNardo et al. (1996), the proposed approach reweights the outcome before estimation to smooth out any possible discontinuity by replacing the observed conditional density function with its counterfactual continuous density.

As a second contribution, this paper discusses in detail the limitations of indirect tests of identification assumptions; discussion for which I revisit Hahn et al. (2001)'s necessary conditions and propose a sufficient condition directly linked to these tests. This condition is that of continuity in the conditional density function of both unobservables and observable covariates at the cutoff. I provide evidence that these indirect tests for manipulation and heaping present a limitation that has been somehow ignored by practitioners; which is that these are informative tests rather than conclusive and are not sufficient nor necessary for identification. If any of these tests do not reject the hypothesis of continuity, then we would judge in favor of the RD estimation's results. However, conditions for identification may not be satisfied despite the success of the indirect test. On the other hand, conditions for identification might be well satisfied even when any of the indirect tests reject the continuity hypothesis.

Also, this article proposes Local Linear Regression (LLR) for the estimation of the weighting scheme,

[2]To the best of my knowledge, with the exception of these two papers, recent literature has focused on other topics and developments such as Regression Kink Design (e.g. Card et al. (2015)), multi-cutoff RD design (e.g. Cattaneo et al. (2016)), multi-running variable RD design (e.g. Keele and Titiunik (2015)), effects away of the cutoff (e.g. Angrist and Rokkanen (2015)) and RD design with covariates (e.g. Calonico et al. (2016)) to name only a few examples and authors. Other recent theoretical and empirical contributions can be found in Cattaneo and Escanciano (2017).

presents asymptotics for the RRD estimator and assesses its finite sample performance with Monte Carlo simulations. The RRD estimator dramatically reduces bias in comparison to the conventional estimator and others related and outperforms them in a Mean Squared Error sense. For completeness, I also present two real data examples where the running variable's discontinuous distribution might be biasing estimation results. The first one illustrates the manipulation case with Bravo and Rau (2012)'s data and results are compared to the RD bounds approach from Gerard et al. (2018). The results obtained by the RRD estimator imply that manipulation is biasing conventional results and that treatment effects found by Bravo and Rau (2012) are still positive but smaller in magnitude. The heaping case is illustrated in the second example with Almond et al. (2010)'s data and results are contrasted to Barreca et al. (2015)'s Donut-RD estimator. The RRD results imply that Almond et al. (2010)'s conclusions are correct and that heaping in this scenario does not pose a threat to the conventional estimation.

This article complements the literature on identification of treatment effects in RD designs when the running variable's density is discontinuous. Gerard et al. (2018) identify sharp bounds for the estimand under manipulation of the running variable. Imposing some structure on how individuals might manipulate, they identify the proportion of "always-assigned" units (i.e. "manipulators") and then use it to derive bounds by considering the worst-case scenarios for the manipulated outcomes. Meanwhile, Barreca et al. (2015) analyze identification in the heaping case. After proposing tests to detect non-random heaping, they discuss how omitting "heaped observations" can achieve identification for non-heaped units. My approach differs from theirs since it consistently estimates the parameter of interest by reweighting the outcome, smoothing out (potential) discontinuities of the relevant distribution function.

The rest of the paper is organized as follows. Section 2 revisits identification conditions for the conventional design and their relation to indirect tests. Section 3 presents the RRD design and discusses its identification capability in comparison to the conventional RD. Estimation and asymptotic properties are shown in Section 4. Performance assessment through simulated examples are presented in Section 5, while Section 6 presents the non-simulated examples. Section 7 concludes.

2 Regression Discontinuity design: Identification

2.1 Revisiting Identification in Sharp RD design

To better understand the identification issues that may arise from a discontinuously distributed running variable, this paper begins by revisiting the conventional identification theorem in RD design as developed by Hahn et al. (2001). The paper also builds on Lee (2008)'s treatment assignment selection model and generalizes his definition of an individual's type to unobservable and observable covariates in order to have a unified framework. To simplify the exposition, consider the Sharp design where treatment status is solely determined by a running variable. In other words, treatment assignment is a deterministic function of the running variable.

Let $(R, \mathbf{X}, U)$ be a random vector with joint distribution $F_{R,\mathbf{X},U}(r, \mathbf{x}, u)$ and support $\Theta \subseteq \mathbb{R}^{J+2}$; where R represents the running variable, $\mathbf{X}$ is a J-length random vector of observable covariates and U is the random variable that represents the unobservable determinants of the outcome variable.

$f_R(r)$ and $f_{\mathbf{X},U|R}(\mathbf{x},u|R=r)$ are the induced marginal and conditional densities respectively.[3] To simplify notation, I omit the conditioning variable while maintaining its value in any conditional moment and/or density function (e.g. $f_{\mathbf{X},U|R}(\mathbf{x},u|R=r) \equiv f_{\mathbf{X},U|R}(\mathbf{x},u|r)$).

Consider the following general model with potential outcomes defined as: $y_1 = g_1(r,\mathbf{x},u)$ under treatment and $y_0 = g_0(r,\mathbf{x},u)$ in absence of treatment, where $g_1(\cdot)$, $g_0(\cdot)$ are continuous real-valued bounded functions. The treatment variable is $t = \mathbf{1}(r \geq c)$, with $\mathbf{1}(\cdot)$ the indicator function and c the threshold or cutoff point. Then, the observed outcome is expressed as $y = y_1 t + y_0(1-t)$. Define the individual treatment effect as $\tau(r,\mathbf{x},u) = y_1 - y_0 = g_1(r,\mathbf{x},u) - g_0(r,\mathbf{x},u)$ and the parameter of interest as the Average Treatment Effect near the cutoff: $E\left[\tau(c,\mathbf{x},u)|c\right]$.[4]

Following Hahn et al. (2001) whom assume: *(RD)* $\lim_{r\downarrow c} E\left[t|r\right]$, $\lim_{r\uparrow c} E\left[t|r\right]$ exist and are different, and *(A1, A2)* both $E\left[y_0|r\right]$, $E\left[y_1|r\right]$ are continuous at $r=c$, we can identify the $E\left[\tau|c\right]$ by:[5]

$$\begin{aligned}
\lim_{r\downarrow c} E\left[y|r\right] - \lim_{r\uparrow c} E\left[y|r\right] &= \lim_{r\downarrow c} E\left[y_1 t + y_0(1-t)|r\right] - \lim_{r\uparrow c} E\left[y_1 t + y_0(1-t)|r\right] \\
&= \lim_{r\downarrow c} E\left[y_1 \mathbf{1}(r\geq c) + y_0(1-\mathbf{1}(r\geq c))|r\right] \\
&\quad - \lim_{r\uparrow c} E\left[y_1 \mathbf{1}(r\geq c) + y_0(1-\mathbf{1}(r\geq c))|r\right] \\
&= \lim_{r\downarrow c} E\left[y_1|r\right] - \lim_{r\uparrow c} E\left[y_0|r\right] \\
&= \lim_{r\downarrow c} E\left[y_1 - y_0|r\right] + \underbrace{\lim_{r\downarrow c} E\left[y_0|r\right] - \lim_{r\uparrow c} E\left[y_0|r\right]}_{=0} \\
&= E\left[\tau|c\right]
\end{aligned}$$

where the last two expressions result from the continuity at $r=c$ assumption. Throughout the paper, continuity will always refer to continuity in r at the threshold c for ease of exposition. It is important to remark that, in this setting, these conditions are necessary and sufficient for identification.

[3]Throughout the paper, it is assumed that the joint and any other induced density functions are strictly positive for every $(r,\mathbf{x},u)$ within the support Θ.

[4]Hahn et al. (2001) define the ATE *at* the cutoff, while Lee (2008) declares "a *weighted* ATE for the entire population", where weights represent the likelihood of having a running variable near the threshold. Given that all RD estimations include observations around the threshold, the common practice to say the estimand of interest is the ATE *near* the cutoff is adopted in this paper.

[5]For simplicity, $E\left[\tau|c\right]$ stands for $E\left[\tau(c,\mathbf{x},u)|c\right]$ from now on.

2.2 The Implications of a Discontinuously Distributed Running Variable

To link the marginal density function $f_R(r)$ with the identification result, using $y_t = g_t(r, \mathbf{x}, u)$ for $t = 0, 1$ we write the conditional expectation of the potential outcome as:

$$
\begin{aligned}
E[y_t|r] &= E[g_t(r,\mathbf{x},u)|r] \\
&= \int_\Omega g_t(r,\mathbf{x},u) f_{\mathbf{X},U|R}(\mathbf{x},u|r)\,d\mathbf{x}du \qquad (1)\\
&= \int_\Omega g_t(r,\mathbf{x},u)\frac{f_{R,\mathbf{X},U}(r,\mathbf{x},u)}{f_R(r)}\,d\mathbf{x}du
\end{aligned}
$$

where $\Omega \subseteq \mathbb{R}^{J+1}$ is the corresponding support for $(\mathbf{X}, U)$. Now, if any of the two conditional expectations $E[y_t|r]$ is discontinuous, the difference $\lim_{r\downarrow c} E[y|r] - \lim_{r\uparrow c} E[y|r]$ does not identify the parameter of interest provided that $\lim_{r\downarrow c} E[y_0|r]$ and $\lim_{r\uparrow c} E[y_0|r]$ do not cancel out and that $\lim_{r\downarrow c} E[y_1 - y_0|r] \neq E[\tau|c]$.

It might seem that (dis)continuity of $f_R(r)$ has an indirect effect on the (dis)continuity of $E[y_t|r]$ through the conditional density function (McCrary, 2008). However, it would be incorrect to make any statement about identification of $E[\tau|c]$ based on $f_R(r)$, given that we have no information about the joint density function $f_{R,\mathbf{X},U}(r, \mathbf{x}, u)$. In consequence, we cannot properly characterize $f_{\mathbf{X},U|R}(\mathbf{x}, u|r)$ or $E[y_t|r]$. Expression (1) suggests that, in distributional terms, continuity of the conditional density function $f_{\mathbf{X},U|R}(\mathbf{x}, u|r)$ is a more relevant condition for identification than continuity of $f_R(r)$. Proposition 1 develops further this claim.

Proposition 1. ***Sufficient Condition for Identification***

Let $(R, \boldsymbol{X}, U)$ be a random vector of length $J + 2$ with joint density function $f_{R,\boldsymbol{X},U}(r, \boldsymbol{x}, u)$ and support $\Theta \subseteq \mathbb{R}^{J+2}$. Potential outcomes are $y_1 = g_1(r, \boldsymbol{x}, u)$ and $y_0 = g_0(r, \boldsymbol{x}, u)$; where $g_1(\cdot)$, $g_0(\cdot)$ are continuous real-valued bounded functions.

If the conditional density function $f_{\boldsymbol{X},U|R}(\boldsymbol{x}, u|r)$ is continuous, then both conditional expectations $E[y_1|r]$ and $E[y_0|r]$ are also continuous; and therefore, following Hahn et al. (2001)'s theorem of identification, $\lim_{r\downarrow c} E[y|r] - \lim_{r\uparrow c} E[y|r]$ identifies $E[\tau|c]$.[6]

Note that this condition is sufficient but not necessary for identification in a sense that both conditional expectations could still be continuous in spite of having a discontinuous conditional density.

Now, from Proposition 1 and expression (1) we can deduce that the marginal density $f_R(r)$ plays no determinant role in identification. It is straightforward to build scenarios that support this claim. Appendix B presents a rather simple example in which (i) $f_R(r)$ is discontinuous but the $E[\tau|c]$ is still identified and (ii) the RD estimand is not identified even though the running variable has a continuous marginal density.

[6]The proof follows Jeffery (1925) and is left to Appendix A.

On one hand, a continuous marginal density does not imply continuity of the conditional density function $f_{\mathbf{X},U|R}(\mathbf{x},u|r)$ and therefore does not guarantee identification. Only if the joint density $f_{R,\mathbf{X},U}(r,\mathbf{x},u)$ were also continuous, then from algebra of continuous functions we would say the conditional density is continuous. This means that continuity of $f_R(r)$ is not necessary for identification; but even more, it is not even a sufficient condition.

On the other hand, a discontinuous marginal density function does not imply the conditional density is also discontinuous. And even if both were discontinuous, it was already pointed that continuity of $f_{\mathbf{X},U|R}(\mathbf{x},u|r)$ is sufficient but not necessary for identification, and therefore the $E[\tau|c]$ could still be identified.

From an empirical perspective the above arguments suggest that indirect tests derived from Lee (2008), McCrary (2008) and Barreca et al. (2015) do not necessarily confirm or reject the validity of the RD estimator given by $\lim_{r\downarrow c}\widehat{E}[y|r]-\lim_{r\uparrow c}\widehat{E}[y|r]$.

Lee (2008) presents a treatment assignment selection model showing that the design's validity could be assessed by analyzing continuity near the threshold of observable covariates. The author argues that if a discontinuity were found, then the density of the running variable conditional on those covariates would presumably be discontinuous and the RD estimates would likely be biased. McCrary (2008) builds on Lee's model and proposes another indirect test which consists of empirically assessing the continuity of the running variable's marginal distribution. Following Lee (2008), the author argues that a rejection of the continuity hypothesis would suggest a discontinuous conditional density function and therefore possible biased estimations. Lee claims that the marginal density function will more likely be discontinuous if manipulation is complete (i.e. individuals have total control over the running variable) while McCrary states that both tests of continuity are informative only if manipulation is monotonic (i.e. all individuals manipulate their running variable in the same direction). Then, it is often interpreted that under complete and monotonic manipulation, the RD estimator's identification capability would be potentially hampered.

Heaping on the other hand, which was first discussed by Barreca et al. (2011) and later formalized in Barreca et al. (2015), could also bias estimations if it occurs at the threshold and if covariates related to the outcome predict the heaping as pointed by the authors. To assess the validity of the RD estimator under heaping, Barreca et al. (2015) present another indirect test which consists of empirically analyzing the difference in covariates between heaped and non-heaped units.

However, in either manipulation, heaping or any other context with a discontinuous marginal density $f_R(r)$, the RD estimator isn't necessarily invalidated. Moreover, continuity of $f_R(r)$ does not guarantee identification. To sum up, care must be taken when assessing validity of an RD estimation through indirect tests given that these are informative rather than conclusive. As McCrary (2008) points out "*a running variable with a continuous density is neither necessary nor sufficient for identification except under auxiliary conditions*".

3 Reweighted Regression Discontinuity design

To the best of my knowledge, there is not a direct or indirect test that definitively assesses the validity of an RD estimation. It is in this context that the main contribution of this paper is the proposal of a new RD approach that uses a reweighting procedure as in DiNardo et al. (1996) to account for any discontinuity in both the marginal $f_R(r)$ and conditional $f_{\mathbf{X},U|R}(\mathbf{x},u|r)$ density functions. As a result, the conditional expectations $E[y_t|r]$ for $t=0,1$ are continuous and identification is then guaranteed.

Moreover, the proposed approach also identifies the estimand of interest — same as the conventional RD design — in the scenario where the conditional density is continuous along with the conditional expectations. These robustness properties allow the estimation results to not hinge on any test about the continuity of the marginal density.

Define the Sharp setting as described in the previous section: $(R,\mathbf{X},U)$ is a $J+2$-length random vector with joint density function $f_{R,\mathbf{X},U}(r,\mathbf{x},u)$, support $\Theta \subseteq \mathbb{R}^{J+2}$ and induced marginal and conditional densities $f_R(r)$ and $f_{\mathbf{X},U|R}(\mathbf{x},u|r)$ respectively. $(R,\mathbf{X})$ with support $\Omega \subseteq \mathbb{R}^{J+1}$ is observed while U represents the unobservable determinants of the outcome. The potential outcomes are defined by $y_t = g_t(r,\mathbf{x},u)$, where $g_t(\cdot)$ are continuous real-valued bounded functions for $t=0,1$. Assignment to treatment is determined by $t=\mathbf{1}(r\geq c)$, so that the observed outcome is written as $y=y_1t+y_0(1-t)$.

Consider the worst case scenario in which Manipulation, Heaping or any other alike (MHO hereinafter), causes both the conditional and marginal densities to be discontinuous in a way that the conventional estimation is no longer consistent for $E[\tau|c]$. In the spirit of McCrary (2008), define R^* as the observed running variable with $f_{R^*}(r)$ and $f_{\mathbf{X},U|R^*}(\mathbf{x},u|r)$ both discontinuous. The true "were there no Manipulation, Heaping, or other" running variable R is no longer observed but it is still conceptually well defined as McCrary points out.

Then, following DiNardo et al. (1996), we can think of counterfactuals such as "what would the continuous conditional expectations, $E[y_1|r]$ and $E[y_0|r]$, have been in absence of MHO?". The authors argue that "*the estimation of such counterfactual[s] ... can be greatly simplified by the judicious choice of a "reweighting" function*"; acknowledging, as is the case in this paper, that "*they can be rewritten in terms of ... [observed realizations] ... with the help of ... [this function]*".

3.1 Identification in the Reweighted-RD design

Within the described framework, define the Reweighting function:

$$w(\mathbf{x},r)=\frac{f_{\mathbf{X}|R}(\mathbf{x}|r)}{f_{\mathbf{X}|R^*}(\mathbf{x}|r)}$$

where $f_{\mathbf{X}|R^*}(\mathbf{x}|r)$ is the conditional density function for the observed running variable and $f_{\mathbf{X}|R}(\mathbf{x}|r)$ is its counterfactual, defined for the true running variable, which is by definition continuous.

Proposition 2. ***Identification with the RRD design***
If, in a neighborhood of $r=c$, $U \perp (R,R^*)|\boldsymbol{X}$ (Local Conditional Independence) then $E[\tau|c]$ is identified by:

$$\lim_{r\downarrow c} E[w(\boldsymbol{x},r)y|r] - \lim_{r\uparrow c} E[w(\boldsymbol{x},r)y|r]$$

Proof.

$$\begin{aligned}
\lim_{r\downarrow c} E\left[w(\mathbf{x},r)y|r\right] - \lim_{r\uparrow c} E\left[w(\mathbf{x},r)y|r\right] =& \lim_{r\downarrow c} E\left[w(\mathbf{x},r)\cdot(y_1 t + y_0(1-t))\,|r\right] \\
&- \lim_{r\uparrow c} E\left[w(\mathbf{x},r)\cdot(y_1 t + y_0(1-t))\,|r\right] \\
=& \lim_{r\downarrow c} E\left[w(\mathbf{x},r)\cdot(y_1\mathbf{1}(r\geq c) + y_0(1-\mathbf{1}(r\geq c)))\,|r\right] \\
&- \lim_{r\uparrow c} E\left[w(\mathbf{x},r)\cdot(y_1\mathbf{1}(r\geq c) + y_0(1-\mathbf{1}(r\geq c)))\,|r\right] \\
=& \lim_{r\downarrow c} E\left[w(\mathbf{x},r)y_1|r\right] - \lim_{r\uparrow c} E\left[w(\mathbf{x},r)y_0|r\right]
\end{aligned}$$

Now, under Conditional Independence and by using the Bayes rule, we have that for $t = 0, 1$:

$$\begin{aligned}
\lim_{r\downarrow\uparrow c} E\left[w(\mathbf{x},r)y_t|r\right] =& \lim_{r\downarrow\uparrow c} E\left[\frac{f_{\mathbf{X}|R}(\mathbf{x}|r)}{f_{\mathbf{X}|R^*}(\mathbf{x}|r)} g_t(r,\mathbf{x},u)|r\right] \\
=& \lim_{r\downarrow\uparrow c} \int_\Omega \frac{f_{\mathbf{X}|R}(\mathbf{x}|r)}{f_{\mathbf{X}|R^*}(\mathbf{x}|r)} g_t(r,\mathbf{x},u) f_{\mathbf{X},U|R^*}(\mathbf{x},u|r) d\mathbf{x} du \\
=& \lim_{r\downarrow\uparrow c} \int_\Omega \frac{f_{\mathbf{X}|R}(\mathbf{x}|r)}{f_{\mathbf{X}|R^*}(\mathbf{x}|r)} g_t(r,\mathbf{x},u) \frac{f_{U,R^*|\mathbf{X}}(u,r|\mathbf{x}) f_{\mathbf{X}}(\mathbf{x})}{f_{R^*}(r)} d\mathbf{x} du \\
=& \lim_{r\downarrow\uparrow c} \int_\Omega \frac{f_{\mathbf{X}|R}(\mathbf{x}|r)}{f_{\mathbf{X}|R^*}(\mathbf{x}|r)} g_t(r,\mathbf{x},u) \frac{f_{U|\mathbf{X}}(u|\mathbf{x}) f_{R^*|\mathbf{X}}(r|\mathbf{x}) f_{\mathbf{X}}(\mathbf{x})}{f_{R^*}(r)} d\mathbf{x} du \\
=& \lim_{r\downarrow\uparrow c} \int_\Omega \frac{f_{\mathbf{X}|R}(\mathbf{x}|r)}{f_{\mathbf{X}|R^*}(\mathbf{x}|r)} g_t(r,\mathbf{x},u) f_{U|\mathbf{X}}(u|\mathbf{x}) f_{\mathbf{X}|R^*}(\mathbf{x}|r) d\mathbf{x} du \\
=& \lim_{r\downarrow\uparrow c} \int_\Omega g_t(r,\mathbf{x},u) f_{U|\mathbf{X}}(u|\mathbf{x}) f_{\mathbf{X}|R}(\mathbf{x}|r) d\mathbf{x} du \\
=& \lim_{r\downarrow\uparrow c} \int_\Omega g_t(r,\mathbf{x},u) \frac{f_{U|\mathbf{X}}(u|\mathbf{x}) f_{R|\mathbf{X}}(r|\mathbf{x}) f_{\mathbf{X}}(\mathbf{x})}{f_R(r)} d\mathbf{x} du \\
=& \lim_{r\downarrow\uparrow c} \int_\Omega g_t(r,\mathbf{x},u) \frac{f_{R,U|\mathbf{X}}(r,u|\mathbf{x}) f_{\mathbf{X}}(\mathbf{x})}{f_R(r)} d\mathbf{x} du \\
=& \lim_{r\downarrow\uparrow c} \int_\Omega g_t(r,\mathbf{x},u) f_{\mathbf{X},U|R}(\mathbf{x},u|r) d\mathbf{x} du \\
=& \int_\Omega g_t(c,\mathbf{x},u) f_{\mathbf{X},U|R}(\mathbf{x},u|c) d\mathbf{x} du \\
=& E\left[y_t|c\right]
\end{aligned}$$

is continuous at $r = c$. Then we have:

$$\begin{aligned}\lim_{r \downarrow c} E\left[w(\mathbf{x}, r) y | r\right] - \lim_{r \uparrow c} E\left[w(\mathbf{x}, r) y | r\right] &= \lim_{r \downarrow c} E\left[w(\mathbf{x}, r) y_1 | r\right] - \lim_{r \uparrow c} E\left[w(\mathbf{x}, r) y_0 | r\right] \\ &= E\left[y_1 - y_0 | c\right] \\ &= E\left[\tau | c\right]\end{aligned}$$

$\square$

Local Conditional Independence requires that around the threshold and conditional on observable covariates, the unobservable determinant of potential outcomes is independent from the running variable(s) that determine treatment status. In other words, the condition requires that MHO is only related to observable covariates, in a sense that we can use them to correct the discontinuity problem. This assumption is equivalent to that of selection-on-observables approaches such as multivariate regression and matching. To identify a causal effect through matching for example, one requires that conditional on observable covariates, potential outcomes are independent of treatment status so that counterfactuals can be estimated by matching individuals with as close as possible observable characteristics. In this setting then, I would equivalently say that, Manipulation, Heaping or any other alike is on observables.

Conditional Independence assumption has already been imposed in RD designs by Angrist and Rokkanen (2015), but for the entire support of the running variable instead of the local version in this paper that imposes it only in a neighborhood of the threshold. In order to identify treatment effects away of the cutoff point, the authors' approach requires to assume that conditional on observable covariates, the running variable is ignorable, i.e. that it is independent of potential outcomes. Within the context of discontinuously distributed running variables, Gerard et al. (2018) impose even further structure to the RD design in order to achieve some form of identification. Specifically, to identify bounds on $E\left[\tau|c\right]$ under Manipulation of the running variable they assume: (i) continuity of conditional expectations of potential outcomes for non-manipulated units, (ii) continuity of the running variable's marginal density function for non-manipulated units, and (iii) monotonic manipulation along the impossibility for manipulated units to have a running variable equal to the cutoff point.

Under MHO where the validity of the conventional RD design might be questioned, the RRD design requires some additional structure — i.e. Local Conditional Independence — to recover the estimand of interest by using the reweighting function $w(\mathbf{x}, r)$. Its main feature is that it identifies $E\left[\tau|c\right]$ regardless the (dis)continuity of the conditional and marginal densities. If $f_{\mathbf{X},U|R^*}(\mathbf{x}, u|r)$ is discontinuous, then $w(\mathbf{x}, r)$ recovers continuity by canceling out $f_{\mathbf{X}|R^*}(\mathbf{x}|r)$, while leaving its continuous counterfactual $f_{\mathbf{X}|R}(\mathbf{x}|r)$. But if $f_{\mathbf{X},U|R^*}(\mathbf{x}, u|r)$ is continuous then $f_{\mathbf{X}|R^*}(\mathbf{x}|r) = f_{\mathbf{X}|R}(\mathbf{x}|r)$ so that $w(\mathbf{x}, r) = 1$ and no change is made to preserve continuity. In either case, the RRD design complies with proposition 1 and identification is therefore guaranteed.

Weighting schemes like the proposed in DiNardo et al. (1996) are widely used in the estimation of counterfactual distribution functions in Data Combination literature or Policy Analysis literature. The survey statistics literature is another example, where observations are reweighted in order to replicate the distribution function that would be in the whole population. See Chernozhukov et al. (2013) and Fan et al. (2014) respectively as examples for both cases.

Finally and for completeness, the following subsection shows the identification problem of the conventional RD design in order to compare it with the proposed approach's identification capability.

3.2 Non-Identification with the conventional RD design

Following Hahn et al. (2001)'s conventional identification theorem and by the Lebesgue's Dominated Convergence theorem, we have that in the worst case scenario described above:

$$\begin{aligned}
\lim_{r\downarrow c} E\left[y|r\right] - \lim_{r\uparrow c} E\left[y|r\right] = & \lim_{r\downarrow c} E\left[g_1(r,\mathbf{x},u)|r\right] - \lim_{r\uparrow c} E\left[g_0(r,\mathbf{x},u)|r\right] \\
= & \lim_{r\downarrow c} E\left[g_1(r,\mathbf{x},u) - g_0(r,\mathbf{x},u)|r\right] \\
& + \lim_{r\downarrow c} E\left[g_0(r,\mathbf{x},u)|r\right] - \lim_{r\uparrow c} E\left[g_0(r,\mathbf{x},u)|r\right] \\
= & \lim_{r\downarrow c} \int_\Omega \tau(r,\mathbf{x},u) f_{\mathbf{X},U|R^*}(\mathbf{x},u|r) d\mathbf{x} du \\
& + \lim_{r\downarrow c} \int_\Omega g_0(r,\mathbf{x},u) f_{\mathbf{X},U|R^*}(\mathbf{x},u|r) d\mathbf{x} du - \lim_{r\uparrow c} \int_\Omega g_0(r,\mathbf{x},u) f_{\mathbf{X},U|R^*}(\mathbf{x},u|r) d\mathbf{x} du \\
= & \int_\Omega \tau(c,\mathbf{x},u) \lim_{r\downarrow c} f_{\mathbf{X},U|R^*}(\mathbf{x},u|r) d\mathbf{x} du \\
& + \int_\Omega g_0(c,\mathbf{x},u) \underbrace{\left[\lim_{r\downarrow c} f_{\mathbf{X},U|R^*}(\mathbf{x},u|r) - \lim_{r\uparrow c} f_{\mathbf{X},U|R^*}(\mathbf{x},u|r)\right]}_{\neq 0} d\mathbf{x} du
\end{aligned}$$

cannot identify $E\left[\tau|c\right]$ because, in the last expression, the second integral does not cancel out and the first one does not correspond to the estimand of interest.

4 Estimation

The RRD design requires reweighting the outcome variable y prior to the RD estimation. Given that $w(\mathbf{x},r)$ is not observable, we need an estimator of it and proceed to a feasible RRD estimation provided that $\hat{w}(\mathbf{x},r) \xrightarrow{p} w(\mathbf{x},r)$. Given an *i.i.d.* random sample $(y_i, r_i, \mathbf{x}_i)$ for $i = 1,\ldots,n$, consider the following two-step estimation by using Local Linear regressions (LLR), the standard technique for RD designs and widely used in other nonparametric settings due to its well behaved boundary properties.

4.1 Weighting Scheme Estimation

The estimator of $w(\mathbf{x}_i, r_i)$ is given by:

$$\widehat{w}(\mathbf{x}_i, r_i) = \frac{\widehat{f}_{\mathbf{X}|R}(\mathbf{x}_i|r_i)}{\widehat{f}_{\mathbf{X}|R^*}(\mathbf{x}_i|r_i)}$$

with $\widehat{f}_{\mathbf{X}|R^*}(\mathbf{x}_i|r_i) > 0$ for all $i = 1, \ldots, n$, and where both densities are estimated separately.

$f_{\mathbf{X}|R^*}(\mathbf{x}_i|r_i)$ is estimated by dividing the sample around the threshold in order to capture a possible discontinuity at $r = c$. On the other hand, $f_{\mathbf{X}|R}(\mathbf{x}_i|r_i)$ is estimated with the entire sample to smooth out any discontinuity and to remain similar to the former density. This idea of empirically capturing the discontinuity and then smoothing it out is inspired by the bunching literature in Empirical Economics, introduced in Saez (2010) and further developed by others like Chetty et al. (2011), Kleven and Waseem (2013) and Ito and Sallee (2014) to name a few.

Following Fan et al. (1996), that propose Local Linear nonparametric estimators for conditional densities, we have that $\widehat{f}_{\mathbf{X}|R^*}(\mathbf{x}|r)$ and $\widehat{f}_{\mathbf{X}|R}(\mathbf{x}|r)$ are the intercept solutions to:

$$\widehat{f}_{\mathbf{X}|R^*}(\mathbf{x}_i|r_i) = \begin{cases} \text{argmin}_{\{\alpha_L,\beta_L\}} \sum_{j:r_j<c} \left[\mathbf{K}_H(\mathbf{x}_j - \mathbf{x}_i) - \alpha_L - \beta_L(r_j - r_i)\right]^2 K_{h_r}(r_j - r_i) & \text{if } r_i < c \\ \text{argmin}_{\{\alpha_R,\beta_R\}} \sum_{j:r_j\geq c} \left[\mathbf{K}_H(\mathbf{x}_j - \mathbf{x}_i) - \alpha_R - \beta_R(r_j - r_i)\right]^2 K_{h_r}(r_j - r_i) & \text{if } r_i \geq c \end{cases}$$

$$\widehat{f}_{\mathbf{X}|R}(\mathbf{x}_i|r_i) = \text{argmin}_{\{\alpha_T,\beta_T\}} \sum_{j=1}^{n} \left[\mathbf{K}_H(\mathbf{x}_j - \mathbf{x}_i) - \alpha_T - \beta_T(r_j - r_i)\right]^2 K_{h_r}(r_j - r_i)$$

where $K_{h_r}(z) \equiv K(z/h_r)/h_r$ and $\mathbf{K}_H(\mathbf{z}) \equiv |H|^{-1/2}\mathbf{K}(|H|^{-1/2}\mathbf{z})$. $K(\cdot)$ is a univariate kernel function, with h_r the scalar bandwidth that smooths over the running variable. $\mathbf{K}(\bullet)$ is a multivariate kernel function with H the bandwidth matrix that smooths over the observed covariates vector. The most commonly used multivariate kernel function is the product kernel so that $\mathbf{K}_H(\mathbf{z}) = 1/(h_1 \times \cdots \times h_J)K(z_1/h_1) \times \cdots \times K(z_J/h_J)$.

Then, $\widehat{f}_{\mathbf{X}|R^*}(\mathbf{x}_i|r_i) = \widehat{\alpha}_L(\mathbf{x}_i, r_i)^{\mathbf{1}[r_i<c]} \cdot \widehat{\alpha}_R(\mathbf{x}_i, r_i)^{\mathbf{1}[r_i\geq c]}$ and $\widehat{f}_{\mathbf{X}|R}(\mathbf{x}_i|r_i) = \widehat{\alpha}_T(\mathbf{x}_i, r_i)$, so that

$$\widehat{w}(\mathbf{x}_i, r_i) = \frac{\widehat{\alpha}_T(\mathbf{x}_i, r_i)}{\widehat{\alpha}_L(\mathbf{x}_i, r_i)^{\mathbf{1}[r_i<c]} \cdot \widehat{\alpha}_R(\mathbf{x}_i, r_i)^{\mathbf{1}[r_i\geq c]}}$$

for $i = 1, \ldots, n$.

4.2 $E[\tau|c]$ Estimation

For the weighted outcome $\widehat{w}(\mathbf{x}_i, r_i)y_i$, consider the conventional RD estimation, where the conditional expectation above and below the threshold are, respectively, the intercept solutions to:

$$\lim_{r\downarrow c} \widehat{E}\left[w(\mathbf{x}, r)y|r\right] = \underset{\{a_R,b_R\}}{\text{argmin}} \sum_{i:r_i\geq c} \left[\widehat{w}(\mathbf{x}_i, r_i)y_i - a_R - b_R(r_i - c)\right]^2 K_{h_r}(r_i - c)$$

$$\lim_{r\uparrow c} \widehat{E}\left[w(\mathbf{x}, r)y|r\right] = \underset{\{a_L,b_L\}}{\text{argmin}} \sum_{i:r_i<c} \left[\widehat{w}(\mathbf{x}_i, r_i)y_i - a_L - b_L(r_i - c)\right]^2 K_{h_r}(r_i - c)$$

So that,

$$\widehat{E}\left[\tau|c\right] = \widehat{a}_R - \widehat{a}_L$$

is the estimator of the Average Treatment Effect near the cutoff point.[7]

4.3 Asymptotic Properties

Following conventional asymptotic theorems developed in Fan et al. (1996) and Calonico et al. (2014) for the Local Linear estimation of conditional densities and conditional expectations respectively, this paper presents the following two propositions for the weighting scheme and $E\left[\tau|c\right]$.

Proposition 3. ***Asymptotic Normality for*** $\widehat{w}(\boldsymbol{x}, r)$
For

$$\widehat{w}(\boldsymbol{x}, r) = \frac{\widehat{f}_{\boldsymbol{X}|R}(\boldsymbol{x}|r)}{\widehat{f}_{\boldsymbol{X}|R^*}(\boldsymbol{x}|r)}$$

estimated as discussed in section 4.1, where both conditional densities are estimated separately through LLR we have:

$$\sqrt{nh_1 \dots h_J h_r}\left(\widehat{w}(\boldsymbol{x}, r) - w(\boldsymbol{x}, r) - e.b.^w(\boldsymbol{x}, r; h_1^2, \dots, h_J, h_r^2)\right) \xrightarrow{d} N\left(0, \frac{w(\boldsymbol{x}, r)}{f_{\boldsymbol{X},R^*}(\boldsymbol{x}, r)} R_K^2\right)$$

Proposition 4. ***Asymptotic Normality for*** $\widehat{E}\left[\tau|c\right]$
For

$$\widehat{E}\left[\tau|c\right] = \widehat{a}_R - \widehat{a}_L$$

estimated as discussed in section 4.2, where both reweighted outcome conditional expectations are estimated separately through LLR we have:

$$\sqrt{nh_r}\left(\widehat{E}\left[\tau|c\right] - E\left[\tau|c\right] - e.b.^E(h_r^2; c)\right) \xrightarrow{d} N\left(0, \left(z_+(c)\frac{\sigma_+^2(c)}{f_R(c)}R_K + z_-(c)\frac{\sigma_-^2(c)}{f_R(c)}R_K\right)\right)$$

where $z(c) = E\left[w^2(x, c)|c\right]$ *and* $\sigma^2(c) = E\left[\epsilon^2|c\right]$ *on each side of the threshold.*

e.b.$^w(x, r; h_x^2, h_r^2)$ and e.b.$^E(h_r^2; c)$ are estimation biases introduced by LLR; a common feature of all nonparametric methods. See Appendix A for the proof of both propositions.

Finally, given that the variances of both conditional expectations' estimators can be written in the traditional Huber–Eicker–White way, this paper adopts Calonico et al. (2014)'s approach and constructs consistent estimators by replacing the estimated residuals into the corresponding expression of the variance.

[7] Bandwidths h_r and H should be optimally chosen for this particular estimator. However, MSE-optimal bandwidth selection is out of the scope of this paper.

4.4 Comparison to the conventional RD estimator

This last subsection compares both the proposed RRD with the conventional RD estimators in terms of expectation and variance in order to shed some light on their performance.

If Hahn et al. (2001)'s condition of continuity does not hold due to a distributional discontinuity the comparison between both estimators is trivial: while the RD estimator is severely biased, the RRD estimator consistently estimates the estimand of interest. Now, consider the following expressions for the expectation of both estimators in the scenario where both approaches identify $E[\tau|c]$ provided that Hahn et al. (2001)'s identification assumptions are met:

$$\widehat{E}_{RRD}[\tau|c] = E[\tau|c] + \text{e.b.}^{E}_{RRD}(h_r^2; c)$$

$$\widehat{E}_{RD}[\tau|c] = E[\tau|c] + \text{e.b.}^{E}_{RD}(h_r^2; c)$$

Both estimators are consistent and they only differ in the estimation bias. It seems reasonable to say that $\text{e.b.}^{E}_{RRD}(h_r^2; c) \geq \text{e.b.}^{E}_{RD}(h_r^2; c)$, given that the proposed RRD involves a two-stage nonparametric estimation and therefore, first stage's estimation bias carries onto the second stage estimation. However, we would expect this difference to vanish asymptotically provided that both estimation biases also vanish.

In terms of efficiency, the usual bias-variance trade-off suggests that, in the case of discontinuity in the conditional expectations, the bias reduction of the RRD estimator comes at the cost of a higher variance in comparison to the RD estimator. However, as will be illustrated with the Monte Carlo experiments, we could expect the bias reduction of the RRD design to more than compensate the higher variance and in consequence, the proposed approach would be preferred over the conventional one. For the case in which both estimators are consistent, the variance for the RRD and the RD estimators are given by, respectively:

$$V\left[\widehat{E}_{RRD}[\tau|c]\right] = \frac{1}{nh_r}\left(z_+(c)\frac{\sigma^2_{+RRD}(c)}{f(c)}R_K + z_-(c)\frac{\sigma^2_{-RRD}(c)}{f(c)}R_K\right)$$

$$V\left[\widehat{E}_{RD}[\tau|c]\right] = \frac{1}{nh_r}\left(\frac{\sigma^2_{+RD}(c)}{f(c)}R_K + \frac{\sigma^2_{-RD}(c)}{f(c)}R_K\right)$$

It is not evident which of both variance expressions is larger. Intuitively, the RD estimator would be more efficient than the RRD estimator due to the additional noise introduced in the latter by the estimation of the weighting scheme $w(\mathbf{x}, r)$. For this reason, we would expect $\widehat{\sigma}^2_{RRD}(c) \geq \widehat{\sigma}^2_{RD}(c)$. The other difference between both expressions is the presence of $z(c) = E\left[w^2(\mathbf{x}, c)|c\right]$ in the RRD case. Theoretically, the weighting scheme is not reweighting the outcome to preserve continuity in this case, which would imply that in average $w(\mathbf{x}, c)$ is close to one and symmetrically distributed. Now, if $w(\mathbf{x}, c) < 1$ then $w(\mathbf{x}, c)^2 \leq w(\mathbf{x}, c)$. Conversely, $w(\mathbf{x}, c)^2 \geq w(\mathbf{x}, c)$ if $w(\mathbf{x}, c) > 1$. We could argue then, that in average $w(\mathbf{x}, c)^2 \geq 1$ and infer that $E\left[w^2(\mathbf{x}, c)|c\right] \geq 1$. With both arguments in mind, it is likely the case that the RD estimator is more efficient than the proposed

RRD estimator. However, we would expect both expressions to converge asymptotically given that $E\left[w^2(x,c)|c\right] \to 1$ and $\widehat{\sigma}^2_{RRD}(c) \to \widehat{\sigma}^2_{RD}(c)$ as the sample size increases.

All of these differences between both estimators are further analyzed in the next section.

5 Simulated examples

This section presents the results of four Monte Carlo (MC) experiments to assess the finite-sample performance of the proposed RRD design in comparison to the conventional RD design. In addition, comparison to Matching and Inverse Probability Weighting approaches is also presented given their similarities in identification conditions and estimation techniques.

For ease of exposition, let the unobservable $U \perp (R, R^*, \mathbf{X}) \sim \mathcal{N}(0, 0.25)$ so that Local Conditional Independence in Proposition 2 is trivially met and therefore the RRD estimator guarantees identification. Define $\mathbf{X}$ as a scalar continuous covariate, R^* the observed running variable with cutoff point $c = 0$ and consider the following $g(\cdot)$ functions for the potential outcomes: $y_0 = r + x + u$ and $y_1 = 1 + 0.5r + x + u$. The individual treatment effect is then $\tau = 1 - 0.5r$ and the estimand is $E\left[\tau|0\right] = 1$. This paper uses very simple linear functions for both potential outcomes to have the same estimand in all four MC experiments, and therefore to better compare the performance of both estimators between them. In addition, note that the focus is on the behavior of the distribution functions and not on how the potential outcomes are determined.

5.1 Discontinuous Conditional Density Function

Consider the following two scenarios for the joint distribution of (X, R^*), where the conditional density $f_{X|R^*}(x|r)$ is discontinuous in a way that both conditional expectations for the potential outcomes are also discontinuous. RD estimates will be biased given that assumptions in the conventional identification theorem are no longer met.

The difference between cases 1 and 2 is on the (dis)continuity of the marginal density function of the running variable:

MC-1: $f_{R^*}(r)$ continuous

Let $f^{[1]}_{X,R^*}(x,r) = f^{[1]}_{X|R^*}(x|r) \cdot f^{[1]}_{R^*}(r)$ the joint density function with induced conditional and marginal densities respectively:

$$f^{[1]}_{X|R^*}(x|r) = \left[\frac{1}{\sqrt{2\pi \cdot 0.75}} \exp\left(-\frac{1}{2} \cdot \frac{(x-0.5r)^2}{0.75}\right)\right]^{\mathbf{1}[r<0]} \cdot \left[\frac{1}{\sqrt{2\pi \cdot 0.25}} \exp\left(-\frac{1}{2} \cdot \frac{(x-0.5-0.5r)^2}{0.25}\right)\right]^{\mathbf{1}[r\geq 0]}$$

$$f^{[1]}_{R^*}(r) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2} r^2\right)$$

where the conditional density of $X|R^*$ is discontinuous and the marginal density of R^* is continuous.

MC-2: $f_{R^*}(r)$ discontinuous

Let $f^{[2]}_{X,R^*}(x,r) = f^{[2]}_{X|R^*}(x|r) \cdot f^{[2]}_{R^*}(r)$ the joint density function with induced conditional and marginal densities respectively:

$$f^{[2]}_{X|R^*}(x|r) = \left[\frac{1}{\sqrt{2\pi \cdot (1 - 0.5^2/V_R)}} \exp\left(-\frac{1}{2} \cdot \frac{(x + 0.5E_R/V_R - 0.5r/V_R)^2}{1 - 0.5^2/V_R} \right) \right]^{\mathbf{1}[r<0]}$$
$$\cdot \left[\frac{1}{\sqrt{2\pi \cdot (1 - 0.5^2/0.5V_R)}} \exp\left(-\frac{1}{2} \cdot \frac{(x - 0.5 + 0.5E_R/V_R - 0.5r/V_R)^2}{1 - 0.5^2/0.5V_R} \right) \right]^{\mathbf{1}[r\geq 0]}$$

$$f^{[2]}_{R^*}(r) = \left[\frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2} r^2 \right) \right]^{\mathbf{1}[r<0]} \cdot \left[\frac{1}{\sqrt{2\pi \cdot 0.5}} \exp\left(-\frac{1}{2} r^2 \right) \right]^{\mathbf{1}[r\geq 0]}$$

with $E_R \equiv E[R^*] = \phi(0) \cdot (\sqrt{0.5} - 1)$ and $V_R \equiv \mathrm{Var}(R^*) = 1.5(0.5 - \phi(0)^2) + 2\phi(0)^2\sqrt{0.5}$. In this case we have both the conditional density of $X|R^*$ and the marginal density of R^* discontinuous.

In both scenarios, the conventional RD design will no longer consistently estimate $[\tau|0] = 1$ while the RRD estimator will; given that the outcome will be reweighted following Proposition 2 to correct the discontinuity in the conditional density.

Every replication of the 1,000 conducted for both scenarios, draws realizations (y_i, x_i, r_i) for $i = 1, \ldots, n$ with $n = \{500, 750, 1000, 2000\}$; where $y_i = y_{1i}\mathbf{1}[r_i \geq 0] + y_{0i}(1 - \mathbf{1}[r_i \geq 0])$. Table 1 presents the results for the estimations through RRD and conventional RD in terms of absolute bias, empirical standard deviation (S.D.) and mean square error (M.S.E.). [8] Bandwidth selection for the weights estimation (i.e. H and h_{r_w}) is conducted by Likelihood Cross Validation following Hall et al. (2004) and the Epanechnikov kernel is used.

Four different bandwidths are considered for the running variable to analyze the sensitivity of the results given that the selection criterion is not optimal for this particular setting: the optimally chosen bandwidth h_{r_w}, 0.75 and 1.25 times the bandwidth; and the one chosen optimally for the conventional estimator (i.e. $h_{r_w} = h_{r_\tau}$) in order to have only one bandwidth for the running variable in both stages.

Estimation of $E[\tau|c]$ uses triangular kernel and bandwidth selection as in Calonico et al. (2014) except for the fourth case where the first stage bandwidth is imposed. Figure 1 presents the first stage's results for a typical replication along the running variable's support. Left panel shows the weighting scheme for MC-1 where $f_{R^*}(r)$ is continuous while the right panel depicts it for MC-2 with discontinuous marginal density $f_{R^*}(r)$. Note that, around the cutoff point of zero, there are weights way above and below one, since continuity of the conditional density $f_{X|R^*}(x|r)$ needs to be restored.

[8] In order to avoid a possible zero in the estimation of the denominator in the weighting scheme, a minimum value of 10e-5 is imposed in this and the non-simulated examples section.

Table 1: RRD estimator unbiased, RD estimator biased

RD vs RRD	$f(\mathbf{x}\|r)$ discontinuous									
	MC-1: $f(r)$ continuous					MC-2: $f(r)$ discontinuous				
	RD	RRD				RD	RRD			
		0.75*h_{r_w}	h_{r_w}	1.25*h_{r_w}	$h_{r_w} = h_{r_\tau}$		0.75*h_{r_w}	h_{r_w}	1.25*h_{r_w}	$h_{r_w} = h_{r_\tau}$
$n = 500$										
Bias	0.50	0.01	0.02	0.03	0.03	0.48	0.03	0.01	0.01	0.00
S.D.	0.27	0.19	0.20	0.27	0.43	0.27	0.21	0.22	0.35	0.42
M.S.E	0.32	0.04	0.04	0.07	0.18	0.31	0.04	0.05	0.12	0.17
Eff.Obs. for										
$E[\tau\|c]$	238	231	235	240	238	237	233	238	242	237
$w(x,r)$	-	208	262	306	417	-	217	272	317	411
$n = 750$										
Bias	0.49	0.01	0.01	0.02	0.02	0.49	0.04	0.02	0.01	0.01
S.D.	0.24	0.17	0.18	0.27	0.33	0.22	0.15	0.19	0.30	0.38
M.S.E	0.30	0.03	0.03	0.07	0.11	0.29	0.02	0.03	0.09	0.15
Eff.Obs. for										
$E[\tau\|c]$	362	348	356	363	362	363	350	358	364	363
$w(x,r)$	-	312	391	457	630	-	323	404	470	625
$n = 1000$										
Bias	0.48	0.02	0.01	0.02	0.03	0.49	0.05	0.02	0.03	0.02
S.D.	0.19	0.18	0.22	0.21	0.32	0.19	0.16	0.16	0.26	0.31
M.S.E	0.27	0.03	0.05	0.04	0.10	0.28	0.03	0.03	0.07	0.10
Eff.Obs. for										
$E[\tau\|c]$	489	456	472	485	489	486	464	479	489	486
$w(x,r)$	-	419	525	613	846	-	420	526	614	836
$n = 2000$										
Bias	0.49	0.03	0.01	0.00	0.00	0.50	0.06	0.04	0.04	0.04
S.D.	0.15	0.11	0.13	0.17	0.28	0.13	0.11	0.14	0.17	0.25
M.S.E	0.26	0.01	0.02	0.03	0.08	0.26	0.02	0.02	0.03	0.06
Eff.Obs. for										
$E[\tau\|c]$	995	891	933	970	995	996	900	941	974	996
$w(x,r)$	-	826	1039	1216	1710	-	808	1016	1188	1695

[1] RRD stands for Reweighted Regression Discontinuity and RD for Conventional Regression Discontinuity.
[2] S.D. stands for (empirical) Standard Deviation and M.S.E. for Mean Squared Errors.
[3] Eff.Obs. stands for Effective Observations used in the estimation; which are selected by the kernel function.
[4] LLR for the estimation of the weighting scheme use the Epanechnikov Kernel and bandwidth selection for H and h_{r_w} as in Hall et al. (2004)
[5] LLR for $E[\tau|c]$ estimation in both the RD and RRD use the Triangle Kernel and bandwidth selection for h_{r_τ} as in Calonico et al. (2014), except for the fourth case where the first stage bandwidth is imposed.

The benefits of implementing the RRD design are substantial: while the conventional estimation is severely biased of up to 50%, the proposed one reduces it almost entirely regardless of the chosen bandwidth. For the first MC experiment, where $f_{R^*}(r)$ is continuous, the bias is reduced between 94 percent and 100 percent; while in the second case with discontinuous marginal density, there are reductions that vary from 88 percent to 100 percent.

The bias-variance trade-off emerges as expected, with slightly higher S.D. in some cases; as a consequence of the first stage variability which is passed onto the second stage. The case where the same bandwidth is imposed for both stages shows larger S.D. and reflects the importance of having optimally chosen bandwidths for the proposed estimator. Nonetheless, the M.S.E. for the

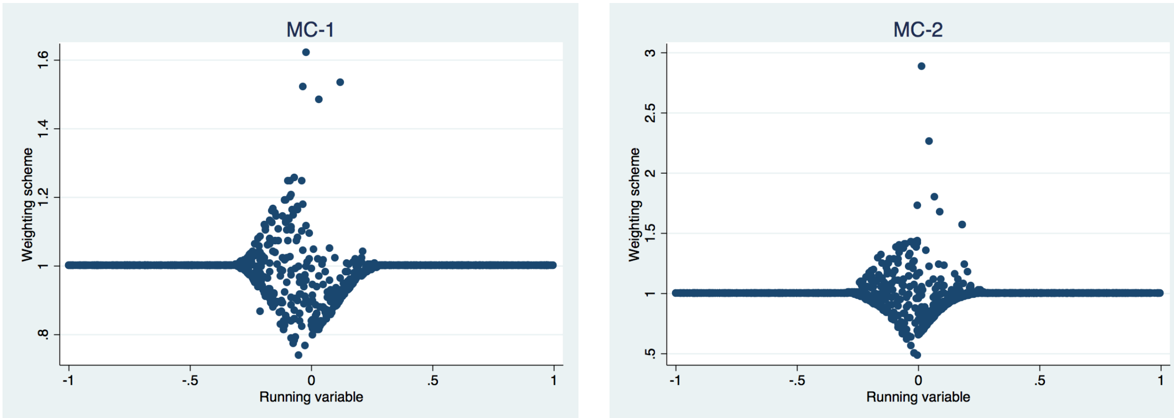

Figure 1: Weighting scheme estimation

RRD estimator is always significantly smaller than its conventional counterpart; implying that the higher variability of the RRD estimator is more than compensated by its bias reduction.

5.2 Continuous Conditional Density Function

In order to assess the robustness of the RRD design, two additional MC experiments are conducted, where the conditional density $f_{X|R^*}(x|r)$ is now continuous in $r = 0$ so that conditions in Proposition 1 are met and therefore, both approaches are consistent for the estimand. As in the previous simulations, the first scenario considers a continuous marginal density while the second one displays a discontinuity in the cutoff:

MC-3: $f_{R^*}(r)$ continuous

Let $f^{[3]}_{X,R^*}(x,r) = f^{[3]}_{X|R^*}(x|r) \cdot f^{[3]}_{R^*}(r)$ the joint density function with induced conditional and marginal densities respectively:

$$f^{[3]}_{X|R^*}(x|r) = \frac{1}{\sqrt{2\pi \cdot 0.75}} \exp\left(-\frac{1}{2} \cdot \frac{(x - 0.5r)^2}{0.75}\right)$$

$$f^{[3]}_{R^*}(r) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2} r^2\right)$$

where both the conditional density of $X|R^*$ and the marginal density of R^* are continuous.

MC-4: $f_{R^*}(r)$ discontinuous

Let $f^{[4]}_{X,R^*}(x,r) = f^{[4]}_{X|R^*}(x|r) \cdot f^{[4]}_{R^*}(r)$ the joint density function with induced conditional and marginal densities respectively:

$$f_{X|R^*}^{[4]}(x|r) = \frac{1}{\sqrt{2\pi \cdot (1-0.5^2/V_R)}} \exp\left(-\frac{1}{2} \cdot \frac{(x+0.5E_R/V_R - 0.5r/V_R)^2}{1-0.5^2/V_R}\right)$$

$$f_{R^*}^{[4]}(r) = \left[\frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}r^2\right)\right]^{\mathbf{1}[r<0]} \cdot \left[\frac{1}{\sqrt{2\pi \cdot 0.5}} \exp\left(-\frac{1}{2}r^2\right)\right]^{\mathbf{1}[r\geq 0]}$$

with $E_R \equiv E[R^*] = \phi(0) \cdot (\sqrt{0.5} - 1)$ and $V_R \equiv \text{Var}(R^*) = 1.5(0.5 - \phi(0)^2) + 2\phi(0)^2\sqrt{0.5}$ as in MC-2. In this case, only the conditional density of $X|R^*$ is continuous while the marginal density of R^* is not.

In these two settings, since $f_{X|R^*}(x|r)$ is continuous then $E[\tau|0] = 1$ will be consistently estimated by both designs even though R^* is discontinuously distributed in the second case. In particular, the RRD design won't reweight the outcome provided that $f_{X|R}(x|r) = f_{X|R^*}(x|r)$. Figure 2 presents the weighting scheme estimation results for both Monte Carlo experiments, where weights vary between 0.5 and 1.5 suggesting some kind of correction in the conditional density function. However, note that weights away from 1 are relatively scarce and are explained by noise estimation, while most of them are concentrated close to 1. In fact, for the graphed replications, deviation from 1 is on average 0.055 and 0.034 for MC-3 and MC-4 respectively while for MC-1 and MC-2 it is 0.116 and 0.103 which more than doubles the former ones.

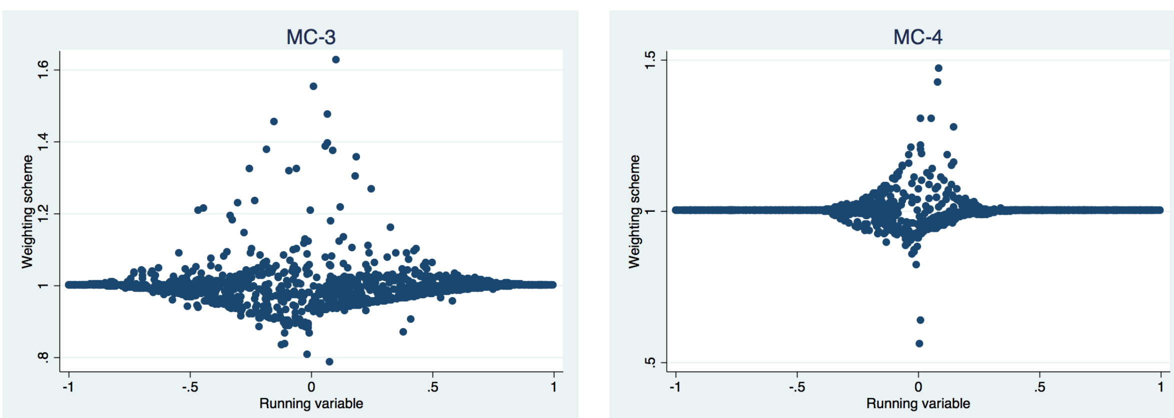

Figure 2: Weighting scheme estimation

Table 2 presents the results for both MC experiments. As expected, the conventional RD estimator outperforms the proposed one in some cases because of efficiency losses and noise introduction of the latter. However, this superiority of the conventional estimator is rather small compared to MC 1 and 2, and differences between both estimators are negligible in general.

This robustness of the RRD design makes it a widely applicable method. Unlike other approaches found in the empirical literature that solve Manipulation or Heaping once the potential problem has been identified, the RRD estimator guarantees identification of the effect regardless of the existence or not of a discontinuity in the conditional and/or marginal distribution functions.

Table 2: Both RRD and RD estimators unbiased, with RD estimator efficient

RD vs RRD	$f(\mathbf{x}\|r)$ continuous									
	MC-3: $f(r)$ continuous					MC-4: $f(r)$ discontinuous				
	RD	RRD				RD	RRD			
		0.75*h_{r_w}	h_{r_w}	1.25*h_{r_w}	$h_{r_w} = h_{r_\tau}$		0.75*h_{r_w}	h_{r_w}	1.25*h_{r_w}	$h_{r_w} = h_{r_\tau}$
$n = 500$										
Bias	0.02	0.03	0.03	0.02	0.01	0.00	0.04	0.03	0.01	0.01
S.D.	0.27	0.18	0.20	0.25	0.35	0.26	0.17	0.27	0.39	0.53
M.S.E	0.07	0.03	0.04	0.06	0.12	0.07	0.03	0.07	0.15	0.28
Eff.Obs. for										
$E[\tau\|c]$	248	251	250	252	248	233	236	236	237	233
$w(x,r)$	-	280	344	391	428	-	255	314	361	409
$n = 750$										
Bias	0.00	0.03	0.03	0.03	0.02	0.00	0.03	0.03	0.02	0.00
S.D.	0.22	0.13	0.17	0.20	0.28	0.21	0.13	0.21	0.20	0.36
M.S.E	0.05	0.02	0.03	0.04	0.08	0.04	0.02	0.04	0.04	0.13
Eff.Obs. for										
$E[\tau\|c]$	380	388	386	385	380	359	365	367	368	359
$w(x,r)$	-	408	503	575	648	-	372	461	532	621
$n = 1000$										
Bias	0.00	0.02	0.02	0.02	0.00	0.00	0.02	0.01	0.01	0.00
S.D.	0.19	0.12	0.15	0.15	0.25	0.18	0.12	0.21	0.22	0.26
M.S.E	0.04	0.01	0.02	0.02	0.06	0.03	0.02	0.04	0.05	0.07
Eff.Obs. for										
$E[\tau\|c]$	515	524	524	524	515	481	490	493	496	481
$w(x,r)$	-	537	663	759	872	-	495	614	707	831
$n = 2000$										
Bias	0.00	0.01	0.01	0.00	0.01	0.00	0.01	0.00	0.00	0.03
S.D.	0.13	0.07	0.10	0.12	0.17	0.13	0.08	0.11	0.16	0.20
M.S.E	0.02	0.01	0.01	0.01	0.03	0.02	0.01	0.01	0.03	0.04
Eff.Obs. for										
$E[\tau\|c]$	1050	1081	1077	1079	1050	989	1020	1024	1026	989
$w(x,r)$	-	1033	1280	1472	1760	-	950	1183	1371	1686

[1] RRD stands for Reweighted Regression Discontinuity and RD for Conventional Regression Discontinuity.
[2] S.D. stands for (empirical) Standard Deviation and M.S.E. for Mean Squared Errors.
[3] Eff.Obs. stands for Effective Observations used in the estimation; which are selected by the kernel function.
[4] LLR for the estimation of the weighting scheme use the Epanechnikov Kernel and bandwidth selection for H and h_{r_w} as in Hall et al. (2004)
[5] LLR for $E[\tau|c]$ estimation in both the RD and RRD use the Triangle Kernel and bandwidth selection for h_{r_τ} as in Calonico et al. (2014), except for the fourth case where the first stage bandwidth is imposed.

5.3 Comparison to other Selection-on-observables Approaches

As a final assessment of the Reweighted RD design, Table 3 presents the comparison to other identification strategies that rely on a similar Conditional Independence assumption. Columns 1, 4, 7 and 10 present the results for the proposed approach along the four Monte Carlo experiments, with optimal bandwidth selection as in Hall et al. (2004) for H and h_{r_w} and as in Calonico et al. (2014) for h_{r_τ}. Columns 2, 5, 8 and 11 display the results for the Nearest Neighbor Matching on Covariates estimator (Abadie and Imbens, 2011; Heckman et al., 1998). One neighbor is used and the matching procedure is performed in a neighborhood around the threshold defined following Calonico et al. (2014) for the conventional RD design. Finally and within the same neighborhood, columns 3, 6, 9 and 12 present the results for the Inverse Probability Weighting estimator (Busso et al., 2014; Tan, 2010). For this, weights are built on the covariate balance propensity score which is estimated following Imai and Ratkovic (2014).

Table 3: RRD compared to Matching and Inverse Probability Weighting

	$f(\mathbf{x}\|r)$ discontinuous						$f(\mathbf{x}\|r)$ continuous					
	MC-1: $f(r)$ continuous			MC-2: $f(r)$ discontinuous			MC-3: $f(r)$ continuous			MC-4: $f(r)$ discontinuous		
	RRD	NNM	IPW	RRD	NNM	IPW	RRD	NNM	IPW	RRD	NNM	IPW
$n = 500$												
Bias	0.02	0.22	0.01	0.01	0.22	0.05	0.03	0.21	0.09	0.03	0.17	0.04
S.D.	0.20	0.14	0.23	0.22	0.14	0.23	0.20	0.11	0.16	0.27	0.12	0.17
M.S.E	0.04	0.07	0.05	0.05	0.07	0.06	0.04	0.06	0.03	0.07	0.05	0.03
$n = 750$												
Bias	0.01	0.22	0.02	0.02	0.20	0.08	0.03	0.20	0.09	0.03	0.16	0.04
S.D.	0.18	0.13	0.20	0.19	0.13	0.23	0.17	0.10	0.15	0.21	0.11	0.15
M.S.E	0.03	0.07	0.04	0.03	0.06	0.06	0.03	0.05	0.03	0.04	0.04	0.03
$n = 1000$												
Bias	0.01	0.20	0.03	0.02	0.19	0.08	0.02	0.20	0.11	0.01	0.16	0.04
S.D.	0.22	0.12	0.18	0.16	0.12	0.19	0.15	0.10	0.14	0.21	0.10	0.15
M.S.E	0.05	0.05	0.03	0.03	0.05	0.04	0.02	0.05	0.03	0.04	0.04	0.02
$n = 2000$												
Bias	0.01	0.19	0.03	0.04	0.17	0.11	0.01	0.19	0.11	0.00	0.15	0.03
S.D.	0.13	0.11	0.16	0.14	0.11	0.17	0.10	0.08	0.13	0.11	0.09	0.13
M.S.E	0.02	0.05	0.03	0.02	0.04	0.04	0.01	0.04	0.03	0.01	0.03	0.02

[1] RRD stands for Reweighted Regression Discontinuity, NNM for Nearest Neighbor Matching and IPW for Inverse Probability Weighting.
[2] S.D. stands for (empirical) Standard Deviation and M.S.E. for Mean Squared Errors.
[4] LLR for the estimation of the weighting scheme use the Epanechnikov Kernel and bandwidth selection for H and h_{r_w} as in Hall et al. (2004)
[5] LLR for $E[\tau|c]$ estimation in both the RRD use the Triangle Kernel and bandwidth selection for h_{r_τ} as in Calonico et al. (2014)

In terms of bias, the Reweighted RD estimator is superior to both the Nearest Neighbor Matching (NNM) and the Inverse Probability Weighting (IPW) estimators in all 4 scenarios and irrespective of the sample size. While bias reduction in MC-1 and MC-2 varies between 56 and 66 percent with NNM and between 78 and 98 percent with IPW, it varies between 88 and 100 percent with the RRD design as already discussed. Moreover, in MC-3 and MC-4 both the NNM and IPW approaches are persistently biased. This superiority is explained by the fact that the RRD design specifically corrects for discontinuities at the conditional distribution level while the other two do so only at the conditional expectation level. Moreover, NNM estimator suffers from additional bias due to the matching discrepancy that does not vanish asymptotically (Abadie and Imbens, 2006). In the case of the IPW estimator, an additional concern relates to the potential misspecification of the propensity score estimation (Drake, 1993). Regarding efficiency, the proposed estimator is in general outperformed by the other two, although this disadvantage disappears as the sample size increases. This follows, same as with the comparison to the conventional RD design, from the additional noise introduced in the first stage of the RRD approach. In terms of mean squared error, that combines bias and variance, we can see that the RRD design is superior to both NNM and IPW estimators when $f(x|r)$ is discontinuous and that all three are similar when $f(x|r)$ is continuous. These results show the advantages of the proposed RRD approach over two other techniques commonly used in the literature that rely on similar identification conditions.

6 Non-simulated examples

This section applies the RRD design to real data where literature has proposed different answers to circumvent a discontinuity in the density of the running variable which could harm the validity of conventional estimations. To this purpose, a manipulation and a heaping examples are presented. The first case is based on Bravo and Rau (2012) that analyze the labor market effects of a youth-employment subsidy and where a concern is present regarding the discontinuity in the marginal distribution of the running variable used by the authors. For this case, the results are compared with Gerard et al. (2018) which propose a way to partially identify a causal effect – through bounds – under manipulation. The second example builds on Almond et al. (2010) which analyze marginal returns of medical spending based on newborn's mortality rates and Barreca et al. (2011) who argue that heaping at the 1,500-gram birth weight might be biasing the results. In this case, the RRD design will be compared to Barreca et al. (2015)'s Donut-RD approach.

6.1 A Manipulation Example

Bravo and Rau (2012) analyze with a conventional RD design the labor market effects of a youth-employment subsidy program implemented in Chile in July 2009. By using administrative records from various sources between July 2009 and December 2010 they exploit the discontinuity in program eligibility to identify treatment effects. The running variable is FPS (*Ficha de Protección Social*) — a continuous proxy-means index that measures economic vulnerability — with 11,734 points as cutoff point so that individuals in the 18-24 age range that have FPS equal or below this threshold are eligible for the subsidy.

The estimation procedure consists of the following LLR:

$$\min_{\{a_L,b_L\}} \sum_{i=1}^{n} [y_i - a_L - b_L(r_i - 11734)]^2 \cdot K_h(r_i - 11734) \cdot \mathbf{1}\,[r_i \leq 11734]$$

$$\min_{\{a_R,b_R\}} \sum_{i=1}^{n} [y_i - a_R - b_R(r_i - 11734)]^2 \cdot K_h(r_i - 11734) \cdot \mathbf{1}[r_i > 11734]$$

where r_i is the running variable FPS and y_i is a binary outcome for formal employment. The ATE near the cutoff is then estimated by $\widehat{a}_L - \widehat{a}_R$. The authors follow Calonico et al. (2014) for bandwidth selection, use triangle kernel and present results for 8 selected months between July 2009 and December 2010. They find significant initial effects that decrease overtime and conclude that "*the program is effective in increasing formal employment*".

Concern regarding identification arises when analyzing the distribution of the running variable given that the histogram — presented in Figure 3 — and McCrary (2008)'s indirect test clearly show a discontinuity at 11734 points. Bravo and Rau (2012) discuss that "*some manipulation of the FPS is possible if vulnerable families [achieve] to decrease their score below 11,734 points*" but argue — in favor of the RD's validity — (i) that the marginal density presents a saw-tooth shape with more than one discontinuity, (ii) that the mass point is found to the right of the threshold which is not consistent with the sorting hypothesis and (iii) that manipulation is limited given that the FPS is "*constructed with several indicators and its formula is unknown*".

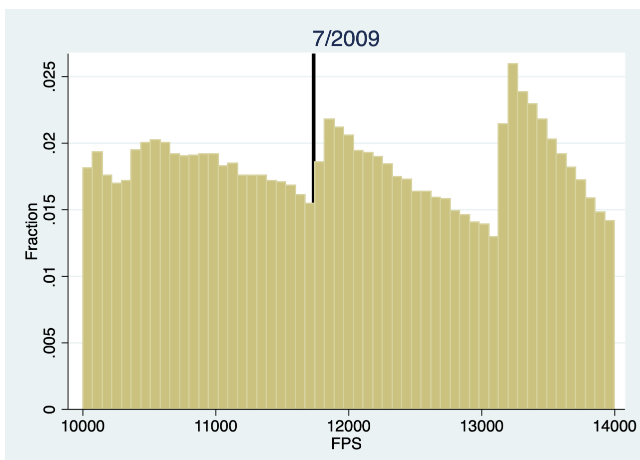

Figure 3: Distribution of the running variable

Is manipulation in this setting harming identification assumptions such that the authors' findings are biased? The proposed RRD design can provide evidence to answer this question given its capability of identifying the estimand of interest under Manipulation of the running variable. In addition and for comparison purposes, this paper also implements the proposal of Gerard et al. (2018) that partially identifies the treatment effect in RD designs with manipulated running variables. Their approach consists of estimating the discontinuity in the density of the running variable as a measure of the proportion of manipulated units and then use it to estimate bounds on the treatment effect considering the worst case scenarios for the values that the observed outcome could take. By doing so, the authors identify bounds for the ATE near the cutoff for non-manipulated units and propose estimators for those bounds and confidence intervals for the aforementioned treatment effect.

Table 4 presents the comparison between the three approaches: in Panel A the conventional RD as estimated by Bravo and Rau (2012), in Panel B the RD bounds as presented by Gerard et al. (2018) and finally in Panel C the proposed RRD design. For ease of exposition, results are presented for 4 out of the 8 months originally used by Bravo and Rau (2012). Replication of the conventional RD finds similar results to that of the authors: an initial significant positive effect on employment of 5.3 percentage points (pp.) (18 percent) that decreases one year later to 2.6 pp. (8 percent) and becomes non-significant after 18 months.

Table 4: Manipulation Example: Comparison between estimators

	$E(\tau\|11734)$	95% C.I.		Eff.Obs.
Panel A: Conventional RD estimations as in Bravo & Rau (2016)				
(A.1) Jul, 2009: LLR with 629-points bandwidth	0.053***	0.038	0.067	820,405
(A.2) Dec, 2009: LLR with 636-points bandwidth	0.049***	0.034	0.064	790,197
(A.3) Jul, 2010: LLR with 640-points bandwidth	0.026***	0.011	0.041	775,943
(A.4) Dec, 2010: LLR with 622-points bandwidth	0.005	-0.012	0.022	712,260
Panel B: RD bounds estimations as in Gerard et al (2018)				
(B.1) Jul, 2009: LLR with 16% manipulation	[-0.004; 0.193]	-0.023	0.217	820,405
(B.2) Dec, 2009: LLR with 18% manipulation	[-0.026; 0.189]	-0.032	0.196	790,197
(B.3) Jul, 2010: LLR with 21% manipulation	[-0.067; 0.192]	-0.086	0.216	775,943
(B.4) Dec, 2010: LLR with 24% manipulation	[-0.134; 0.182]	-0.180	0.215	712,260
Panel C: Reweighted Regression Discontinuity estimations				
(C.1) Jul, 2009: LLR with 626-points bandwidth	0.048***	0.034	0.062	820,331
(C.2) Dec, 2009: LLR with 632-points bandwidth	0.045***	0.030	0.060	790,130
(C.3) Jul, 2010: LLR with 640-points bandwidth	0.024***	0.008	0.039	775,880
(C.4) Dec, 2010: LLR with 615-points bandwidth	0.004	-0.013	0.021	712,206

[1] C.I. stands for Confidence Intervals. Eff. Obs. are observations effectively used within the selected bandwidth.
[2] LLR use Triangle kernel for the treatment effect estimation and for the weighting scheme estimation use Epanechnikov Kernel for continuous covariates and Aitchison-Aitken Kernel for discrete covariates.
[3] Bandwidth selection as in Calonico et al. (2014) is used for the conventional RD and RRD treatment effects estimators. A fixed bandwidth of 623-points is used for the estimation of the RD bounds.
[4] Bandwidth selection as in Hall et al. (2004) is used for the weighting scheme estimation, except for FPS where a 629-points bandwidth is fixed for computational restrictions.
[5] *** significant at 1%, ** significant at 5% and * significant at 10%.

The first stage of Gerard et al. (2018)'s approach estimates a proportion of manipulated units that varies between 16% in July 2009 and 24% in December 2010. Estimated bounds, based on these proportions, are large with ranges of up to 31.6 pp. In July 2009 for example, I find that the employment effect near the cutoff — for non-manipulated units — might be of 19.3 pp. but it might well also be null or even negative of 0.4 pp.; and this without taking into account the estimated confidence intervals that are even wider. In this particular setting, given the large discontinuity found in the marginal density of the running variable, the estimated RD bounds seem of impractical use.

The Reweighted RD approach on the other hand estimates in initial effect of 4.8 pp. that decreases to 2.4 pp. one year later and also becomes non-significant after 18 months. These results imply that manipulation of the running variable in this setting is effectively biasing the conventional results in about 0.5 pp., which accounts for a 10.4 percent bias. First stage results of the RRD approach provide some insight on the correction process. Figure 4 displays all four weighting schemes along the running variable's support. In all cases weights vary between 0.5 and 2 approximately, suggesting important corrections to the outcome variable. Consistent with the increase of the proportion of manipulated units over time found with Gerard et al. (2018)'s approach, the standard deviation of the weighting scheme also increases from 0.0386 in July 2009 to 0.0404 in December 2010. Gender, education, age, labor market participation and an indicator of rural area are used as covariates for the estimation of the weighting scheme following Bravo and Rau (2012) that discuss the process determining the FPS score and describe them as the most relevant ones. It is concluded then that in this particular setting, the conventional approach as used in Bravo and Rau (2012) is biased. The short-run effect of the subsidy on youth employment is actually smaller, although it remains positive and significant. We now turn to the heaping case as the second example in this paper.

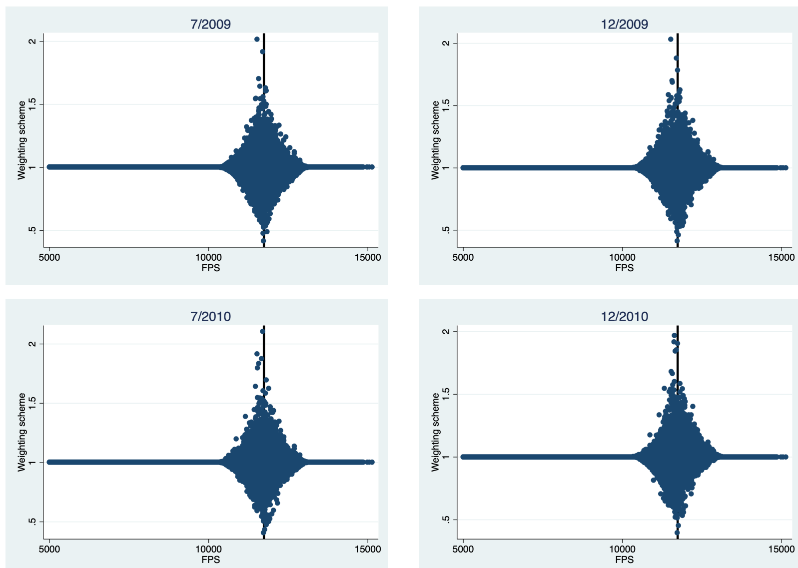

Figure 4: Weighting scheme estimation

6.2 A Heaping Example

Almond et al. (2010) use an RD design to compare the mortality rate of newborns around the 1500-gram birth weight threshold and infer about the marginal returns of medical spending. To this purpose they use the U.S. births' census between 1983 and 2002 (except for the period 1992-1994 where data is not available) and find that "*newborns with birth weights just below 1500 grams have lower one-year mortality rates than do newborns with birth weights just above this cutoff, even though mortality risk tends to decrease with birth weight*". Specifically, they point out that "*One-year mortality falls by approximately one percentage point as birth weight crosses 1500 grams from above, which is large relative to mean infant mortality of 5.5% just above 1500 grams*", suggesting an economically and statistically significant positive effect of medical spending on newborn mortality rates.

To obtain these results they perform four different estimations, where the outcome is 1-year mortality and birth weight is the running variable. First a LLR with Triangle Kernel and 85-gram bandwidth:

$$\min_{\{a_L,b_L\}} \sum_{i=1}^{n} [y_i - a_L - b_L(r_i - 1500)]^2 \cdot K_h(r_i - 1500) \cdot \mathbf{1}[r_i < 1500]$$

$$\min_{\{a_R,b_R\}} \sum_{i=1}^{n} [y_i - a_R - b_R(r_i - 1500)]^2 \cdot K_h(r_i - 1500) \cdot \mathbf{1}[r_i \geq 1500]$$

where the treatment effect is given by $\widehat{a}_L - \widehat{a}_R$ and then 3 OLS regression for observations in the $[1415; 1585]$-gram interval only; first without controls, then adding year effects and finally also adding variables related to parental characteristics and the mother's pregnancy:

$$\min_{\{\alpha\}} \sum_{i=1}^{n} [y_i - \alpha_0 - \alpha_1 t_i - \alpha_2 t_i(r_i - 1500) - \alpha_3(1 - t_i)(r_i - 1500)]^2 \cdot \mathbf{1}[1415 \leq r_i \leq 1585]$$

where $t_i = \mathbf{1}[r_i < 1500]$ is the treatment variable and $\widehat{\alpha}_1$ corresponds to the treatment effect.

Barreca et al. (2011) revisit the results of Almond et al. (2010), suggesting they may be upward biased due to a large heap found at the 1500-gram threshold in particular and at 100-gram and 1-oz multiples in general — as shown in Figure 5. They argue that a reason for potentially biased results is that "*those at the 1500-gr heap appear to be of a particularly disadvantaged sort, ..., a signal that poor-quality hospitals have relatively high propensities to round birth weights*". In particular, they claim that "*newborns at exactly 1500 grams are ... substantially less likely to be white and more likely to have a mother with less than high school education*". To circumvent this potential issue, Barreca et al. (2011) and Barreca et al. (2015) suggest ignoring the heap by dropping observations near the 1500-gram threshold; an approach commonly known among empiricists as a Donut-RD estimator. The authors argue that by doing this, we can obtain an unbiased estimation of the treatment effect for non-heaped individuals. Specifically, Barreca et al. (2011) estimate four OLS regressions within the $[1415; 1585]$-gram interval. All models include year effects and other controls as in Almond et al. (2010)'s OLS, differing in the way they drop observations: from those with 1500 grams up to those within 3 grams around 1500. Then, Barreca et al. (2011) report that, in the first case, the estimated effect drops by more than 50 percent and that in the last model "*the point estimate falls further such that it is now, ..., statistically indistinguishable from 0*".

This setting with mixed results from the literature is ideal for the RRD design to provide more reliable evidence. The central point of the discussion is, in terms of this proposal, if the 1500-gram heaping is causing the conventional RD estimates to be biased. I use the mother's race and education as covariates for the weighting scheme estimation, since they are considered by Barreca et al. (2011) as the most important ones in determining the non-random-heaping.[9] Figure 6 illustrates the weighting scheme estimation results along the running variable's support. Note

[9] Information about both covariates is not available for the entire sample used in Almond et al. (2010), which restricts the estimation of the weighting scheme to a subsample. Nonetheless, it is important to note that this information loss accounts only for a 9.6% of the total data and that both the subsample and the entire sample remain similar in all the relevant variables (i.e. mortality and birth weight).

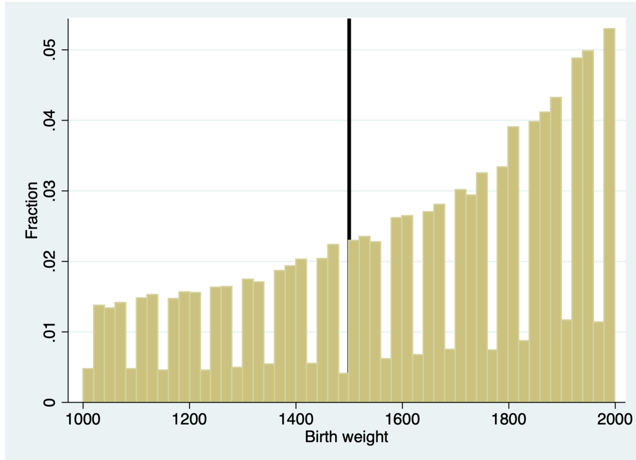

Figure 5: Distribution of the running variable

that for this example we have all weights very close to 1 with variation between 0.944 and 1.030. In fact, the standard deviation for this weighting scheme is 0.002 compared to approximately 0.040 in the manipulation example. As it was discussed in the simulated examples, this small variation could be explained by estimation noise suggesting that the conditional density may well be continuous with conventional RD results unbiased.

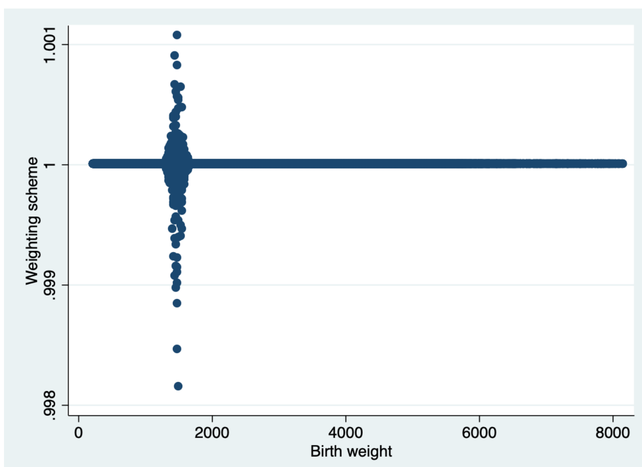

Figure 6: Weighting scheme estimation

Table 5 presents the results for all three approaches: the conventional RD as in Almond et al. (2010) in Panel A, the Donut-hole RD as in Barreca et al. (2015) in Panel B and the proposed RRD in Panel C. Replication results are nearly identical to those of both papers with a significant effect on mortality of between 0.5 and 1 percentage points in Panel A and a precise null effect in Panel B. To mimic as close as possible the specifications used by Almond et al. (2010) and Barreca et al. (2011), four different estimations for the second stage of the RRD approach are performed: a LLR with Triangle Kernel and 85-gram bandwidth and three OLS regressions within the $[1415; 1585]$-gram interval, progressively adding year effects and other controls. The RRD estimation results are nearly identical to those of Almond et al. (2010) implying that dropping observations around the cutoff point after finding evidence of nonrandom heaping would not benefit the analysis because, even if present, heaping is not biasing the RD estimates. With this example, it is concluded then that one-year mortality drops for newborns with birth-weight just below 1500-grams in comparison to those with birth-weight just above that threshold as a result of additional medical spending.

Table 5: Heaping Example: Comparison between estimators

	$E(\tau\|1500)$	S.E.	Eff.Obs.
Panel A: Conventional RD estimations as in Almond et al (2010)			
(A.1) LLR with 85-gram bandwidth	-0.00862***	0.00253	182,696
(A.2) OLS in [1415; 1585]	-0.00669***	0.00222	184,016
(A.3) OLS in [1415; 1585], with year effects (y.e.)	-0.00489**	0.00221	184,016
(A.4) OLS in [1415; 1585], with (y.e.) and controls	-0.00429*	0.00220	184,016
Panel B: Donut-hole RD estimations as in Barreca et al (2015)			
(B.1) OLS as in (A.4) dropping those at 1,500	-0.00135	0.00222	180,960
(B.2) OLS as in (A.4) dropping those at [1,499, 1,501]	-0.00149	0.00223	180,767
(B.3) OLS as in (A.4) dropping those at [1,498, 1,502]	-0.00103	0.00225	179,625
(B.4) OLS as in (A.4) dropping those at [1,497, 1,503]	-0.00038	0.00276	159,586
Panel C: Reweighted Regression Discontinuity estimations			
(C.1) LLR with 85-gram bandwidth	-0.00858***	0.00252	182,696
(C.2) OLS in [1415; 1585]	-0.00666***	0.00222	184,016
(C.3) OLS in [1415; 1585], with (y.e.)	-0.00485**	0.00221	184,016
(C.4) OLS in [1415; 1585], with (y.e.) and controls	-0.00425*	0.00220	184,016

[1] S.E. stands for Standard Error. Eff. Obs. are observations effectively used within the selected bandwidth.
[2] See Almond et al. (2010) for the list of controls included.
[3] LLR for the estimation of the weighting scheme use the Epanechnikov Kernel for continuous variables and Aitchison-Aitken Kernel for race. Triangle Kernel is used in LLR for the estimation of the treatment effects.
[4] Bandwidth selection as in Hall et al. (2004) is used for the weighting scheme estimation, except for birth weight where the 85-gr bandwidth is fixed for computational restrictions.
[5] *** significant at 1%, ** significant at 5% and * significant at 10%.

7 Conclusions

Regression Discontinuity design has been more frequently used in recent years, to estimate causal effects when assignment to treatment is based on a running variable and its cutoff point; due to its high internal validity and ease of implementation. Hahn et al. (2001) formally presented conditions for identification from which continuity at $r = c$ of conditional expectations for both potential outcomes is the most important one. From an empirical perspective, testing continuity of the running variable's distribution function has been a key factor for judging the design's validity.

This paper discussed the importance of interpreting indirect tests in RD designs as informative rather than conclusive by arguing, one hand, that discontinuity of $f_R(r)$ does not imply invalidation of the design; and on the other hand, that a continuous marginal density function does not assure unbiased estimation results. Continuity of the conditional distribution function $f_{\mathbf{X},U|R}(\mathbf{x}, u|r)$ is then presented as a sufficient condition for identification; condition on which this paper's main contribution is based.

Combining Hahn et al. (2001) and Lee (2008)'s treatment assignment models a general model is developed, where potential outcomes are expressed as continuous functions of the running variable, observed covariates and unobservables. Then, following DiNardo et al. (1996), this paper proposed a new RD approach which consists of properly reweighting the observed outcome prior to the conventional estimation.

In the probable scenario in which Manipulation, Heaping or any other setting with distributional discontinuities harm identification, the RRD design is able to consistently estimate the estimand of interest by having the weighting scheme cancel out the discontinuous conditional density and replace it with its counterfactual counterpart; which is, in the spirit of McCrary (2008), the continuous covariates' distribution function conditional on the "were there no Manipulation, Heaping, or other alike" running variable.

Moreover, if a distributional discontinuity does not harm identification, then the RRD design is capable of achieving unbiased estimation results, same as the conventional RD design. Both of these properties make the proposed approach widely applicable to any setting with (or without) Manipulation, Heaping or any other alike, given that identification is guaranteed.

Estimation of the weighting scheme is inspired by the bunching literature introduced by Saez (2010) and performed by local linear regressions to exploit its well-behaved-at-boundaries property. Montecarlo exercises were conducted to analyze the RRD performance in comparison to the conventional RD and other similar approaches. Two non-simulated examples with real data were also presented to compare the proposed approach with current existing methods of partial identification. The RRD design has been proved to be robust to any distributional discontinuities and superior to the conventional design in a MSE sense.

Further research should point to develop an MSE-optimal bandwidth selector suitable to this setting; and to the possibility of introducing a new (indirect) test based on the difference around the threshold between the observed $f_{\mathbf{X}|R^*}(\mathbf{x}|r)$ and its counterfactual part $f_{\mathbf{X}|R}(\mathbf{x}|r)$.

Part II

Effects of a Reduction in Credit Constraints on Educational Attainment: Evidence from Chile

Effects of a Reduction in Credit Constraints on Educational Attainment: Evidence from Chile

Andrés García E., Pinjas Albagli

This version: December, 2018

Abstract

This paper analyzes the enrollment and retention effects of a student loan reform that loosened credit constraints in Chile. The reform reduced the interest rate from an average of 6 percent to a fixed rate of 2 percent. The identification strategy follows a Difference-in-difference approach that compares the effects of this policy change among eligible and ineligible students. We find a precise null effect of the reform on the overall immediate enrollment, along with a diversion effect that increased enrollment to universities in 2.5 percentage points (pp.) (7 percent), while enrollment to vocational institutions dropped in 2.5 pp. (14 percent). Moreover, we find that for female students the decrease in enrollment to vocational institutions is not fully offset by the increase in enrollment to universities. We also find that the reform increased university retention by 3 percent in two-year enrollment and reduced by 10 percent the dropout rate. Our findings are mainly driven by medium-income-family students.

1 Introduction

There is a growing interest in the economics literature about the effects of financial aid on educational attainment as a result of the ongoing debate regarding the relative importance of long-run and short-run constraints.[10] On one hand, some researchers argue in favor of early-stage investments as the main drivers of long-run educational and labor outcomes (Cameron and Heckman, 2001; Carneiro and Heckman, 2002; Heckman et al., 2006). And on the other hand, some others focus on short-run credit constrains as the main obstacles for higher educational attainment, especially among low-income families (Lochner and Monge-Naranjo (2011) provide an excellent detailed review for the latter).

Empirically evaluating how credit constraints affect tertiary educational decisions is a difficult challenge for several reasons. For instance, (i) the impossibility of directly observing credit constraints, (ii) the potential endogeneity in enrollment-based regressions, and (iii) the fact that most tertiary education systems have admission processes that are highly determined by unobserved measures such as alumni status of parents and recommendation letters.[11]

[10]See Dynarski and Scott-Clayton (2013) for a review on the economics-of-education literature on financial aid.

[11]See Riegg (2008) for a discussion about causal inference and selection bias on the financial aid literature.

In recent years, a branch of literature has focused on studying the Chilean Higher Education system (CHES hereinafter) thanks to the similarities with the U.S. system but mostly because its institutional setting allows overcoming several of the aforementioned issues. The admission process, for example, is entirely determined by observable academic variables such as high-school GPA and the national admission test score. Moreover, the CHES has a highly centralized and standardized grants system, allowing researchers to come in hand with a rich set of administrative records. Finally, in the last decades, Chile has introduced and modified different aid programs to boost access to post-secondary education. Empiricists have exploited these programs as quasi-natural experiments to identify the effects of credit access on educational attainment and labor market outcomes (e.g. Rau et al. (2013), Solis (2017), Bucarey et al. (2018), Montoya et al. (2018)).

The main goal of this paper is to analyze post-secondary enrollment and retention effects of a Chilean reform to state-guaranteed loans that took place in 2012, loosening credit constraints for high school graduates. The reform consisted of the following changes to repayment conditions: (i) a decrease of the interest rate from approximately 6 percent average to a fixed interest rate of 2 percent, (ii) repayments were made contingent on income with a cap of 10 percent, and (iii) the possibility to delay repayments in case of unemployment. From these, the interest rate drop is the most relevant change since it is automatically applicable to all loans, while the two others are available upon request and only a small fraction of debtors apply for them.[12]

Our identification strategy is based in a Difference-in-differences (DiD) approach, exploiting the differences between students who were exposed and not exposed to the 2012 changes (post and pre 2012 high school graduation cohorts) and between students who were eligible and ineligible for the loan. We combine a rich set of administrative records at the individual level. Our data covers the entire population of high school graduates in Chile between 2006 and 2014 who faced their first enrollment decisions in the 2007-2015 period. We have detailed information about their enrollment and permanence choices, about the academic variables that determine loan eligibility, and other individual, school, and educational program characteristics that we use as control variables.

We contribute to the empirical literature on the effects of financial aid on educational attainment in two ways. First, while most of the research focuses on the effects of having access to student loans (i.e. on the extensive margin), this paper is the first one, to the best of our knowledge, to evaluate the effects of a reform that introduces changes in the intensive margin by loosening constraints and in a context where those changes were designed to affect repayment behavior only.[13] Analyzing the effects of these intensive margin changes to credit access is of important relevance for policymakers, specially in countries with similar student-loan systems where a reform that loosens credit constraints might be part of the future agenda.[14] Second, and in comparison to other related research, our data has the advantages of having (i) complete non-missing information about the entire population of high school graduates, (ii) large sample sizes that improve efficiency of our estimates, and (iii) a considerable number of cohorts.

Our results suggest that the loosening of credit constraints had no effect on overall immediate

[12] In 2015, only 8% and 4% of the debtors were beneficiary of the 10%-cap and delayed repayments respectively. See Ingresa (2015) for details.

[13] See Nielsen et al. (2010) and Dynarski (2003) as examples of papers that study the effects of reforms to other types of financial aid different from student loans.

[14] Colombia, Mexico, U.S., Canada, U.K. and Australia are examples of countries with student loans as mechanisms of financial aid.

enrollment (i.e., enrollment in any CHES institution during the year immediately following high school graduation). However, we find a diversion effect: enrollment to universities increased by 2.5 percentage points (pp.) — which amounts to a 7 percent increase relative to enrollment of non-exposed eligible individuals — in detriment of enrollment to vocational institutions, which fell by 2.5 pp. — equivalent to a 14 percent decrease of enrollment relative to the same group. This effect is stable over time except for a decrease in 2015 when a new free-tuition program was announced by the government. This shift in institutional choice is explained — in line with Angrist et al. (2016) — by the implicit subsidy the reform creates toward universities relative to vocational institutions given that the former are more expensive in terms of tuition fees and program length. Moreover, the diversion effect implies welfare effects since some individuals that diverted their decision toward universities would be likely better-off had they pursued a vocational degree instead (Rodriguez et al., 2016).

Our findings are consistent with the evidence on the enrollment effects of financial aid in general (Cornwell et al. (2006), Fack and Grenet (2015), Perna and Titus (2004), Van der Klaauw (2002)), and particularly with the Chilean evidence on the effects of having access to the CAE loan. With a Regression Discontinuity design, Solis (2017) and Montoya et al. (2018) find that loan eligibility increases university immediate enrollment by 18 pp. and 15.2 pp. respectively; although these results apply only for individuals with PSU score near 475 points. In addition, our results are smaller due to the fact that we analyze a reform on the intensive rather than the extensive margin.

Regarding retention, we find that as a result of the 2012 changes, the diversion from vocational institutions to universities also encouraged enrollment in universities for a second consecutive year, increasing it in almost 1 pp. — a percent increase relative to non-exposed eligible individuals — while two-year enrollment in vocational institutions fell by 0.5 pp. — equivalent to a 4 percent decrease relative to the same group. Conditional on being enrolled, we also estimate that the dropout rate from universities decreases in almost 2 pp. (a 10 percent decrease), while the effect in vocational institutions is statistically non-significant. This improvement in university retention along with a small deterioration in vocational persistence results from two mechanisms: a sorting effect in ability caused by the diversion effect in enrollment that reduces the likelihood of dropping out in universities while increasing it in vocational institutions (Rodriguez et al., 2016); and a perverse incentive from the CAE loan itself that encourages all institutions to reduce dropout rates given their guarantors role (Rau et al., 2013).

The international literature on the persistence effects of financial aid has found mixed results so far. For example, Glocker (2011) and Chatterjee and Ionescu (2012) discuss the importance of financial aid on retention and completion; while Herzog (2005), Stinebrickner and Stinebrickner (2008) and Stinebrickner and Stinebrickner (2012) find that there are other factors that are more relevant than credit constraints for persistence and graduation. Our results are consistent with recent Chilean evidence. Solis (2017) finds an increase of 16 pp. in university two-year enrollment. Our result is smaller but, again, his finding applies for selected individuals only and considering access to the loan instead of a loosening in credit constraints. Rau et al. (2013) build a structural model for sequential schooling decisions and find that access to this particular loan reduces dropout rates in both universities and vocational institutions.

Finally, we also examine the possibility of heterogeneous effects across two dimensions, namely gender and family income. Regarding gender, we estimate that the only significant difference between men and women is on immediate enrollment such that for females the diversion from vocational institutions (-3.2 pp.) is not fully compensated by the increase in university enrollment (2.3 pp.). As a result, the interest rate drop has a negative impact on immediate enrollment for women (-0.9 pp.), which is explained by female students delaying their enrollment decision to better prepare for the challenge of being accepted at the university. Along the family-income dimension, which we proxy by high school financing scheme, we find that all of our results in enrollment and persistence are entirely driven by students graduating from voucher high schools (middle-income family) with no effects whatsoever on students from public schools (low-income family). This results from the fact that the student loan under analysis does not cover the full tuition costs so that students still need to finance the remaining difference along with other expenses. Then, the 2012 reform is not large enough to have an impact on low socioeconomic status high school graduates.

From a public policy perspective, our findings suggest that a reform to student loans that loosens credit constraints might have null overall effects on access to education, but it could instead have unintended consequences in the institutional composition of students in dimensions such as ability, gender and socioeconomic status; which in turn translate into nontrivial welfare effects.

The remainder of the paper is organized as follows. Section 2 describes the institutional background of the CHES, the changes introduced to the loan in 2012, and the data. The empirical strategy for identification of the effects on educational outcomes is presented in Section 3, while Section 4 presents the results for all of our outcomes, analyzes the plausibility of the identification strategy, and studies heterogeneous effects. Section 5 concludes.

2 Background and Data

The Chilean Higher Education System (CHES) comprises two types of institutions: universities and vocational institutions (*Institutos Profesionales* and *Centros de Formación Técnica*). Universities offer professional programs and are the only institutions entitled to confer academic degrees. Programs at universities are usually between 5 and 6 years of length. Vocational institutions on the other hand, offer technical programs which are mainly between 3 and 4 years of duration. Both types of institutions are financed primarily through tuition fees, with the state providing complementary funding by direct and indirect mechanisms assigned almost entirely to universities.

Tuition fees imply an important financial burden for high school graduates that decide to enroll, since they represent a large fraction of family income. Between 2007 and 2015, the period of analysis in this paper, the mean tuition fee in the 62 Chilean universities was roughly about $CLP 2.1 million ($USD 2,970), which represents 41% of the median family income in 2015.[15] For the more than 100 vocational institutions, the mean tuition fee was around $CLP 1.1 million ($USD 1,556), representing 21% of the 2015 median family income.

This is of special relevance for students graduating from state-funded public schools and from voucher schools. In the same period of time, 39% of the students came from public schools and the mean tuition fee represented 42% and 22% of the median family income for universities and vocational institutions, respectively. Similarly, 53% of the students graduated from a voucher school and the mean tuition fee for universities represented 34% of the median family income and 18%

[15] Median family income is calculated in all cases using the household survey *Caracterización Socioeconómica Nacional* CASEN 2015. Conversion from $CLP to $USD uses the official exchange rate of 12/31/2015.

for the vocational institutions case. Finally, for the remainder 8% of students graduating from private high schools, the mean tuition fee represented 10% and 5% of the median family income for universities and vocational institutions, respectively. In the results section, we will assess how the reform heterogeneously impacts graduates from public schools versus graduates from voucher schools.

Students have few options to finance tertiary education. To work-and-study or work-and-save are usually very demanding alternatives and access to the conventional financial market is typically limited by restrictive conditions on income and job formality. That is why students rely on government grants as their principal source of funding, where eligibility is mostly determined by academic performance and socioeconomic characteristics such as family income. In 2015 for example, from a total of 1,165,654 students enrolled in the CHES, 723,216 (58%) had some form of government financial aid. That same year, the government granted 443,299 loans (38%) and 397,386 scholarships (34%) (Ministry of Education, 2016).

Scholarships cover tuition and, in some cases, enrollment fees and others such as transportation and food expenses. Student loans, on the other hand, cover tuition fees only.[16] Students have access to two types of loan: the traditional university loan or FSCU (*Fondo Solidario de Crédito Universitario*) and the state guaranteed loan or CAE (*Crédito con Aval del Estado*). The FSCU loan is granted by the state only to students who enroll in the so called "traditional" universities, has an annual interest rate of 2% with payments that begin two years after graduation, and contemplates a maximum of 15 years of payments with a cap of 5% of total income.[17] The CAE loan is provided, administered, and collected by private banks and guaranteed by the state and the higher education institution where the student is enrolled. Payment conditions, such as the interest rate, changed in the 2012 reform and are described in detail below.

From all the types of financial aid the government grants to students, the CAE loan is the most important, both in number of beneficiaries and amount granted, as shown in Table 6. In fact, one in every three tertiary education students has a CAE loan to pay for tuition fees. These figures hint at the public policy relevance of analyzing the effects of the 2012 reform to CAE.

2.1 The CAE Loan and the 2012 reform

The CAE loan was introduced in 2006 as an alternative to the conventional FSCU loan that was granted only to students enrolled in traditional universities. The main goal of the policy was to broaden access to the CHES regardless of the chosen institution (i.e., university or vocational institution). Participants in the CAE system are: (i) private banks lending the money, (ii) the government and educational institutions as guarantors absorbing the default and dropout risks respectively and (iii) the students/debtors that borrow and make repayments accordingly.

The process of CAE loan applications and CHES enrollment is structured as follows. Students graduating from high school register for the PSU (*Prueba de Selección Universitaria*), a national college admission test that highly determines admission to the CHES and access to grants.[18] During

[16] Moreover, loans only cover tuition fees up to a maximum amount called "referenctial tuition fee" which is annually determined by the Ministry of Education for each program based on its quality.

[17] "Traditional" universities, or more formally *Universidades del Consejo de Rectores*, is a group of the 27 universities created before 1980.

[18] The PSU is administered once per academic year and consists of two mandatory (language and mathematics)

Table 6: Government Grants in 2015

	Quantity		Total Amount	
Scholarships	**397,386**	**47.27%**	**483,597**	**49.80%**
Beca Centenario	99,930	11.89%	240,974	24.81%
Beca Nuevo Milenio	171,576	20.41%	96,362	9.92%
Beca de Articulación	5,557	0.66%	3,892	0.40%
Beca Juan Gómez Millas	63,474	7.55%	70,545	7.26%
Beca Excelencia Académica y PSU	24,946	2.97%	26,859	2.77%
Beca de Nivelación Académica	3,466	0.41%	2,850	0.29%
Beca Hijos de Profesionales de la Educación	10,360	1.23%	5,104	0.53%
Beca Vocación de Profesor	9,555	1.14%	21,715	2.24%
Beca de Reparación	3,858	0.46%	6,222	0.64%
Beca de Reubicación U. del Mar	4,664	0.55%	9,074	0.93%
Loans	**443,299**	**52.73%**	**487,494**	**50.20%**
CAE	369,253	43.92%	415,951	42.83%
FSCU	74,046	8.81%	71,543	7.37%
Total	**840,685**	**100.00%**	**971,091**	**100.00%**

Notes: Ministry of Education, *Memoria Financiamiento Estudiantil 2016*. Quantity refers to the number of grants. Total Amount in CLP $MM.

the PSU registration process, individuals planning to apply for the CAE loan (or other grants) must fill a socioeconomic form which is used to determine income eligibility. Once test results are published, academic eligibility is determined, loans are granted and students decide to either enroll or not to their respective programs.

To become a beneficiary of the CAE loan, a high school graduate must fulfill both the academic and the family income eligibility criteria. Only students with a PSU score greater or equal to 475 or high school GPA greater or equal to 5.3 are eligible.[19][20] The socioeconomic criterion is the least relevant of the two since it has changed overtime and students do not ex-ante know what the cutoff is because the state sorts applicants by income and grants the loans up to the available budget. In 2007, the first year of analysis in this paper, the CAE loan covered up to the fourth income quintile and since 2014 it has been granted based on the academic criteria only, covering applicants from all socioeconomic conditions.

Initially the CAE loan was granted with conditions similar to those of a conventional loan in the financial sector with market interest rates, payments not contingent on income, and banks legally

and two optional tests (science and history/social science; one must be chosen). PSU scores range from 150 to 850 points and are normalized to have a mean of 500 and standard deviation of 110 points. The average score of the mandatory tests is typically used to assess eligibility for grants.

[19] GPA ranges from 1 to 7.

[20] If a student wishes to enroll to an university then she has to comply with the PSU cutoff, while if she wants to enroll to a vocational institution then she has to comply with either of the thresholds.

entitled to use mechanisms to collect debts. CAE loans have maturity up to 20 years, payments begin 18 months after graduation, and between 2006 and 2011 had an average annual interest rate of 5.6 percent. In middle 2011, the government announced a reform to the CAE loan that came into effect in 2012. The changes introduced were (i) a new fixed annual interest rate of 2 percent, similar to that of the FSCU and with the government subsidizing the difference with the market interest rate; (ii) repayments contingent on income upon request, with a cap of 10% and the government subsidizing the remaining difference; and (iii) the possibility, upon request, to delay payments in case of unemployment. With these changes the government intended to level up the conditions between the two loans and expected to improve repayments following a report that estimated a default rate of 36% and predicted a possible increase to up to 50% (World Bank, 2011).

From a theoretical perspective, this reform accounts for a loosening of credit constraints since individuals initially faced tighter repayment conditions that were relaxed in 2012 and implied a reduction of educational costs (in present value). Of the three changes introduced, the interest rate drop is the most relevant one given that the subsidized reduction is automatically applicable to all loans; while the 10%-of-income cap subsidy to repayments and the option to delay them in case of unemployment are available upon request and only a small fraction of debtors has applied for them since its implementation. In 2015 for example, 8% and 4% of the 242,604 CAE debtors were beneficiary of the 10%-cap and delayed repayments respectively (Ingresa, 2015). Moreover, the decrease in the interest rate is considerable in terms of the present value of repayment flows. To illustrate its potential implications, consider the following scenario. A student applying for a CLP$ 2.1 million annual loan at the former 5.6% interest rate would owe a total of CLP$ 15.7 million at the end of a 6-year program and after the 18-month period of grace. With a 20-year maturity loan, this is equivalent to an annuity of CLP$ 1.3 million. With the new interest rate of 2%, she would instead owe a total of CLP$ 13.6 million (a 13 percent drop) with an annuity of CLP$ 0.8 million, which represents a non trivial decrease of 37 percent.[21]

This loosening of constraints constitutes a change in the intensive margin of credit access rather than an extensive margin change such as the introduction of the CAE loan itself. It is important to analyze the potential effects of such intensive margin changes on educational attainment, especially when these changes are substantial as in the 2012 reform.

2.2 Data and Sample

The application process for financial aid is highly centralized in Chile, allowing us to use nation-wide administrative records that contain information about the entire population of high school graduates, along with their eligibility status and enrollment choices in any given year. We obtained information from three sources.

The first is the student performance database from the Ministry of Education that comprises records of all students enrolled in primary and secondary education, from which we built our universe of high school graduates. This source contains relevant information about the student and her high school. Our second source of information is DEMRE (*Departamento de Evaluación, Medición y Registro Educacional*), the institution in charge of the PSU process. They provided us with the PSU score

[21] Several assumptions are implicit in this example for the sake of simplicity. To name a few, we assume that the student requests the same amount every year, that the loan is granted on an annual basis along with the future repayments, that there is no inflation, that the debtor does not request contingent payments nor a delay of them, etc.

for all test takers in our period of analysis. Our third data source from the Ministry of Education provides individual information about enrollment decisions in all universities and vocational institutions. Merging all the data through an individual identifier, we built a dataset consisting of every yearly cohort of high school graduates and information on their eligibility, enrollment and persistence.

We limit our analysis to the 2007-2015 cohorts (i.e., high school students graduating between 2006 and 2014) for two reasons. In 2006, the first year of implementation, the government missassigned the CAE loan due to an error in the income sorting of applicants, granting loans in the opposite order (Ingresa, 2010). And secondly, the government introduced a new program in 2016 that made available tuition-free tertiary education for some individuals. The 2016 reform entirely changed the scenario for students regarding financial restrictions, which in turn could introduce a confounding factor into our analysis of the 2012 reform.[22]

In addition, care must be taken in using the entire population of high school graduates. As already discussed, income eligibility changes over time and its threshold is not observed by the researchers nor by the applicants. To overcome this issue, we drop from our sample all graduates from private high schools in order to resemble as close as possible income eligibility compliance. By doing so — i.e. conditional on being socioeconomically eligible — we exploit eligibility on the academic dimension only. A second concern is related to high school graduates who do not register to take the PSU test, impeding us to determine their eligibility through the PSU score channel. For this reason we additionally restrict our sample to registered students only.

3 Empirical Strategy

Following a simple model of human capital accumulation with imperfect credit markets, state-funded programs such as scholarships and loans increase the net present value of investment in the education project by reducing the associated costs and, in consequence, increasing the probabilities of enrollment, persistence, and graduation.

Although the changes introduced in the 2012 reform affected the intensive margin and focused on the repayment period, the drop in the interest rate is substantial enough to motivate the investigation of the educational effects of this loosening in constraints, given that it reduced the cost associated to the investment in education. To identify these causal effects, we use a Difference-in-differences (DiD) approach exploiting the timing of the reform and the loan's academic eligibility conditions.

3.1 Immediate Enrollment

Our first and main outcome of interest is immediate enrollment, defined as the choice of enrollment the year immediately following high school graduation. Given that the CAE loan is constrained to eligible individuals only, our treatment group is the sample of eligible individuals from cohorts 2012-2015 since they are the only ones exposed to the reform.

Our first difference is the comparison between the treatment group and those eligible students

[22] See Espinoza and Urzúa (2015) for an initial evaluation of the new tuition free program and Bucarey (2017) for an analysis on other educational effects.

unexposed to the reform (i.e. eligible individuals from the 2007-2011 cohorts). The difference in enrollment between these two groups cannot be uniquely attributed to the reform since other confounding factors could also potentially be explaining part of the difference.

In order to solve this issue, our second difference in enrollment is the one between the two groups of cohorts of ineligible individuals (2012-2015 and 2007-2011). Given that these individuals' decision is not affected by the reform, any difference between the 2007-11 and the 2012-15 cohorts will capture those potential confounders.

With this DiD model we are implicitly assuming that the average remaining difference in unobservables between eligible and ineligible individuals is the same before and after the 2012 changes; this assumption is commonly known as the parallel trends condition. In the results section we present evidence of the plausibility of this assumption.

Following standard practice, our estimation base model is:

$$y_{it} = \beta_0 + \beta_1 \text{eligible}_{it} + \beta_2 \text{exposed}_{it} + \beta_3 \text{eligible}_{it} \times \text{exposed}_{it} + \epsilon_{it} \tag{2}$$

where $eligible_{it}$ is an indicator of CAE-eligibility for high school graduate i of cohort t and $exposed_{it}$ indicates exposure to the reform (i.e., $t \geq 2012$).[23]

Immediate enrollment, y_{it}, is to be captured by three binary variables. The first variable is overall enrollment which equals 1 when individual i enrolls to the CHES regardless of the type of institution chosen and 0 if she does not enroll. Our second binary variable is university enrollment that takes the value of 1 if the individual enrolls in an university and 0 otherwise (i.e. if she enrolls to a vocational institution or she does not enroll at all). Similarly, our third variable is vocational enrollment, an indicator that activates when the high school graduate enrolls to a vocational institution. We follow this strategy in order to also capture any possible differences in enrollment between these two types of institutions.

In this model, the interaction coefficient for $eligible_{it} \times exposed_{it}$ (i.e., β_3), captures the Intention to Treat effect (ITT) of the reform on the enrollment rate. This model is also to be extended to include cohort fixed effects and other relevant covariates as robustness checks for our model specification.

A second specification of our DiD identification strategy is:

$$y_{it} = \beta_0 + \beta_1 \text{eligible}_{it} + \sum_{j=2007}^{2015} \alpha_j \text{cohort}_{jit} + \sum_{j=2007}^{2015} \beta_j \text{eligible}_{it} \times \text{cohort}_{jit} + \epsilon_{it} \tag{3}$$

where the $exposed_{it}$ variable is replaced by the cohort fixed effects $cohort_{jit}$. This model is useful in that it disaggregates the overall effect into yearly effects, providing information about the dynamics. In this case, coefficients β_j of the interaction $eligible_{it} \times cohort_{jit}$ for $j = 2012, \ldots, 2015$ are those of interest, since they capture the evolution of the effect over time. Moreover, the remainder β_j coefficients (i.e. those for $j = 2007, \ldots, 2011$) are of particular interest as well since they allow us to test for the parallel trends assumption.

[23] Note that our data is not longitudinal, as each individual is considered only in the corresponding year of her immediate enrollment decision.

3.2 Two-year Enrollment

Our second and third outcomes focus on persistence decisions. Here we define two-year enrollment as a binary variable that takes the value of 1 if the high school graduate immediately enrolls for two consecutive years and 0 otherwise, which includes the scenarios of enrollment for one year only or no enrollment at all. Same as with the immediate enrollment outcome, we make use of three variables: (i) overall two-year enrollment, (ii) university two-year enrollment and (iii) vocational two-year enrollment.

The first difference in our DiD setting comes from the comparison between eligible students that were exposed and those who were not exposed to the reform. Note that in this case, the first cohort that was exposed is the 2011 cohort (and not the 2012 one), since those are the first individuals whose decision of a second year of enrollment is made under the new loan conditions. For this reason, exposed cohorts are now those from 2011 to 2014, while unexposed cohorts are those from 2007 to 2010.[24] To isolate the potential confounding differences between these two groups of cohorts we use the difference in enrollment for ineligible students between periods of exposure and non-exposure as our second difference.

An issue arises with the two-year-enrollment outcome. Eligibility to the CAE loan is potentially endogenous in this setting, given that initially ineligible individuals (i.e. those with PSU < 475 and GPA < 5.3) can retake the PSU test one year later and become eligible if they manage to score above the 475 threshold. For this reason, we use an Instrumental Variables approach within our DiD framework.

The endogenous variable is the overall eligibility within two years following high school graduation ($Eligible_{2it}$). It is given by the student's GPA, which does not change overtime, and the first-attempt PSU score in case she does not retake the test, or the second-attempt PSU score in case she retakes it and scores above her first score. We use as instrument the first-attempt eligibility status ($eligible_{1it}$) which is given by the student's GPA and by her first-attempt PSU score.

A similar identification assumption to that in our DiD model in the previous section (i.e., a paralell trends assumption) provides validity of the instrument in this framework. Relevance of the instrument is also straightforward in this setting: the endogenous variable $Eligible_{2it}$ and its instrument $eligible_{1i}$ are highly correlated by construction since they both build on the GPA and the first-attempt PSU score. In fact, they will only differ in the scenario of a formerly ineligible student retaking the test and scoring above 475 points.[25]

Our two stage least squares (2SLS) base model is defined by the structural equation:

$$y_{it} = \beta_0 + \beta_1 \text{Eligible}_{2it} + \beta_2 \text{exposed}_{it} + \beta_3 \text{Eligible}_{2it} \times \text{exposed}_{it} + \epsilon_{it} \tag{4}$$

and by the first-stage equation:

$$\text{Eligible}_{2it} = \gamma_0 + \gamma_1 \text{eligible}_{1it} + \gamma_2 \text{exposed}_{it} + \eta_{it} \tag{5}$$

[24] Note that with this specification we lose one cohort of students, that of 2015, given that we do not observe their second-year decision in 2016.

[25] In this case we would have for that student that $eligible_{1it} = 0$ and $Eligible_{2it} = 1$. Also, note that all first-attempt eligible individuals are also overall eligible as well (i.e. those with $eligible_{1it} = 1$ also have $Eligible_{2it} = 1$).

where $exposed_{it}$ is the dummy variable for students affected by the reform in their second year as already discussed. In this case, the coefficient β_3 of the interaction $Eligible_{2it} \times exposed_{it}$ captures the local average treatment effect (LATE) of the reform on the two-year enrollment rate of compliers.[26] We will also extend our model to include cohort fixed effects and other relevant covariates as robustness checks.

Just as with equation 3, in order to desegregate the effect into yearly effects, we also estimate the following model:

$$y_{it} = \beta_0 + \beta_1 \text{Eligible}_{2it} + \sum_{j=2008}^{2015} \alpha_j \text{cohort}_{jit} + \sum_{j=2008}^{2015} \beta_j \text{Eligible}_{2it} \times \text{cohort}_{jit} + \epsilon_{it} \tag{6}$$

$$\text{Eligible}_{2it} = \gamma_0 + \gamma_1 \text{eligible}_{1it} + \sum_{j=2008}^{2015} \gamma_j \text{cohort}_{jit} + \eta_{it} \tag{7}$$

where Equation 6 corresponds to the 2SLS structural equation,Equation 7 is the 2SLS first-stage equation, and the exposure variable $exposed_{it}$ is replaced and disaggregated by the cohort fixed effects $cohort_{jit}$. In this model, the parameters of interest are the β_j of the interaction $Eligible_{2it} \times cohort_{jit}$ for $j = 2012, \ldots, 2015$ to capture the dynamics of the effect and for $j = 2008, \ldots, 2011$ to test for the parallel trends assumption.

Two-year-enrollment provides a measure of persistence that comprises information about the immediate first year decision to enroll along with information on the decision to continue onto the second year of enrollment. To disentangle this information and know about the marginal effect on the second year decision we make use of our third and last outcome.

3.3 Second-year Dropout

Our third outcome variable is second-year dropout. Analysis of dropout decisions is conditional on being enrolled: our subsample of study comprises all high school graduates that immediately enrolled in the CHES in the 2007-2014 period and we will be interested in their dropout decision for the following year. Any results from this model should be interpreted with caution since there might be an issue of sample selection.

Given that we only have enrollment records at the beginning of each period, we do not observe if a student completed the year or not. For this reason, we define second-year dropout as a binary outcome that takes the value of 1 if we do not observe a student's registration at the beginning of her second year, regardless of whether she completed her first academic period or not.

Because estimation is now conditional on enrollment, we no longer define three binary variables but analyze second-year dropout across types of institution (i.e., universities and vocational institutions) instead. A 2SLS model just like equations 4 and 5 is used to solve for the potential endogeneity of the overall eligibility condition. Again, the interaction coefficient β_3 captures the effect of interest.

[26] In this setting a complier is a student that enrolls for two consecutive years only if initially eligible to the loan. See Angrist et al. (1996) for details.

Equations 6 and 7 are replicated for this outcome as well to analyze dynamics and test for the parallel trends assumption where, again, the parameters of interest are those of the interactions β_j.

4 Results

This section presents and discusses our main results. For completeness and to better understand the Chilean context, Table 7 presents some descriptive information for selected cohorts.[27] Our sample consists of over 1.5 million high school graduates, 40% of which come from a public school and the remaining 60% from a voucher school. The overall female/male ratio is of 1.15. Eligibility to CAE loan has increased its coverage from 75% in 2007 to 82% in 2015.

Enrollment in the CHES has an upward trend overtime with an annual growth rate of 2.3%, mainly explained by growth in vocational enrollment (4.9% vs 0.7%). Overall, one half of our sample of high school graduates immediately enrolls to the CHES. Within our period of study, the gender gap in enrollment decreased by two thirds from 3 pp. to 1.1 pp. A more subtle decrease is found in the enrollment gap between students from public high schools vs students from voucher schools. The gap decreased from close to 9 pp. to nearly 6 pp.

In terms of retention in the CHES, 36% of high school graduates in our sample enrolls for two consecutive years, with an annual growth rate of 2.5% and driven, once again, by vocational permanence (6.3% vs 0.2%). The gender gap was very small in 2007 and not only it disappeared but at the end of the sample period females are more likely than males to enroll for two years.[28] The gap in the trends by type of school is very similar to that of immediate enrollment, with students from public schools close to 6 pp. less likely to enroll for two years than students from voucher schools.

Also regarding retention in tertiary education, one in every four students enrolled in the CHES drops in her second year of studies. While the dropout rate has decreased over time in vocational institutions, it has marginally increased in universities. In every year of our period of study, females are less likely to drop than males by nearly 3 pp. The gap in dropout rates by type of school has remained stable overtime at about 2 pp.

The following subsections present the estimation results of the models discussed in the previous sections. All regressions follow a Linear Probability Model with clustered standard errors at the class level to account for intra-class correlation. In this setting, a class is defined as the corresponding cohort graduating from a specific high school in a specific year.

In order to assess the relative sizes of our estimates, we report the respective number of non-exposed eligible individuals and their outcome mean in most tables. As a robustness check, we add year effects and three types of control variables to our base models. Student level variables include gender, attendance rate, *comuna*, and number of family members at different levels in the education system. School level variables include indicators of financing scheme (public or voucher), rural area, and geographical region. Finally, program level covariates — which are included only in the regressions for second-year dropout — include tuition fee and program duration.

[27]See Appendix A for detailed information on all our cohorts.

[28] See Becker et al. (2015) and Becker et al. (2010) for an analysis of the overtaking of men by women in higher education.

Table 7: Descriptive Statistics

	Cohort				
	2008	2010	2012	2014	Pooled
Immediate Enrollment	0.457	0.466	0.515	0.543	0.496
Two-Year Enrollment	0.330	0.346	0.355	0.398	0.364
Second Year Dropout	0.266	0.241	0.273	0.261	0.257
Eligible	0.779	0.771	0.769	0.794	0.778
PSU	475	473	475	477	475
GPA	5.6	5.6	5.6	5.6	5.6
Females	0.546	0.531	0.534	0.532	0.534
Public School	0.423	0.421	0.362	0.365	0.396
Obervations	147,480	180,306	169,824	174,789	1,527,798

4.1 Effects on Immediate Enrollment

Table 8 presents results for our three immediate enrollment variables: overall, university, and vocational institution enrollment. Columns (1), (4) and (7) show the results for our base model as presented in equation 2. Estimation results from adding cohort fixed effects are displayed in columns (2), (5) and (8), while columns (3), (6) and (9) also include student and high school control variables.

Eligible students are more likely to enroll. This is not only due to CAE's availability, but also because they are potentially eligible for other grants and/or the FSCU loan. Moreover, given that eligibility is determined by academic variables, which are arguably related to ability, results suggest that more able students are more likely to enroll. However, when we disaggregate by type of CHES institution, we find that this result is driven by university enrollment: eligible students are more likely to enroll to universities and slightly less likely to enroll to vocational institutions. This could be explained by higher economic returns associated to college degrees, but could also be understood in a comparative advantage framework in a Roy selection model. The coefficient on the *exposed* variable captures the trend in enrollment over time, as already discussed.

The overall enrollment effect of the reform is neither statistically nor economically significant, suggesting that the loosening of credit constraints had no impact on immediate enrollment. Interestingly, we find a diversion effect when we conduct our analysis separately by type of institution: the reform increased enrollment to universities in detriment of vocational institutions by 2.5 pp. In absolute terms, this result implies that approximately 16,000 out of 636,760 individuals shifted their enrollment decision toward universities instead of vocational institutions. This finding is robust to the inclusion of different sets of covariates and roughly amounts to a 7 percent increase in university enrollment and a 14 percent decrease in vocational enrollment, relative to the enrollment rate of non-exposed eligible individuals.

Our results are consistent with others found in the literature, although of a smaller magnitude. By means of a Regression Discontinuity Design (RDD) Solis (2017) uses cohorts 2007-09 to estimate the effects of crossing the 475-PSU-score threshold, which enables loan eligibility, and finds that immediate enrollment in universities increases by 18 pp., close to a 100 percent increase relative to

Table 8: Immediate Enrollment

	Overall			University			Vocational		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Eligible	0.254***	0.254***	0.237***	0.286***	0.286***	0.267***	-0.032***	-0.032***	-0.030***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)	(0.002)
Exposed	0.063***	0.074***	0.080***	-0.011***	-0.025***	-0.022***	0.074***	0.099***	0.103***
	(0.003)	(0.007)	(0.006)	(0.001)	(0.007)	(0.007)	(0.003)	(0.004)	(0.004)
Eligible × exposed	0.001	0.001	-0.001	0.026***	0.026***	0.025***	-0.025***	-0.026***	-0.025***
	(0.004)	(0.004)	(0.003)	(0.004)	(0.004)	(0.004)	(0.003)	(0.003)	(0.003)
Cohort effects	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Control variables	No	No	Yes	No	No	Yes	No	No	Yes
Observations	1,527,798	1,527,798	1,527,797	1,527,798	1,527,798	1,527,797	1,527,798	1,527,798	1,527,797
Control group size	636,760	636,760	636,760	636,760	636,760	636,760	636,760	636,760	636,760
Outcome mean	0.524	0.524	0.524	0.348	0.348	0.348	0.176	0.176	0.176

Notes: Clustered standard errors at the class level in parentheses. *** p<0.01, ** p<0.05, * p<0.1. School level control variables include indicators of financing institution, rural area and geographical region. Student level control variables include gender, attendance rate, *comuna* and number of family members at different leves in the education system. Control group size accounts for the number of ineligible individuals in the exposure period, while Outcome mean refers to the mean of the dependent variable of those individuals.

ineligibles. Following a similar RDD with the same three cohorts, Montoya et al. (2018) analyze labor market effects and within their model also estimate effects on different measures of enrollment. The authors find that scoring above the 475 cutoff has a positive effect of 9.6 pp. on overall immediate enrollment and 15.2 pp. on university immediate enrollment, arguing that most of this variation is a reflection of a vocational-to-university substitution.

Two reasons explain the difference with our results. First, we focus on a reform that introduced changes in the intensive margin (i.e. an interest rate drop that loosens credit constraints) while others analyze the effects of having access to the CAE loan itself (i.e. the extensive margin). Second, in the RDD framework results are local in the sense that they are interpreted as treatment effects for individuals near the threshold (i.e., those with a PSU-score close to 475 poins), while our results are interpreted as an average for the treated individuals.

The shift in institutional choice from vocational institutions to universities is explained — in line with Angrist et al. (2016) — by the implicit subsidy the interest rate drop creates for universities. Given that enrolling in this type of institutions entitles more costs both in pecuniary (i.e. tuition fees) and timely (i.e. program length) terms, the loosening of credit constraints is of a bigger scale for the choice of attending universities; which in turn, further increases the relative incentive to enroll in an university in comparison to a vocational institution.

In addition, this diversion effect implies a welfare effect that depends on the characteristics of the individuals that shifted their enrollment decision toward universities as a result of the 2012 reform. Rodriguez et al. (2016) propose a structural schooling decision model to simulate the effects of a reduction of tuition costs in Chile — which can be interpreted as a loosening in credit constraints and therefore similar to the interest rate drop — and find negligible effects on overall enrollment which is consistent with our results. Moreover, the authors find for Chile that (i) more able students obtain lengthy degrees (i.e., pursue degrees at universities instead of vocational institutions); (ii) economic returns (annual earnings) are increasing in ability and are larger for students graduating from university than for those graduating from vocational institutions; (iii) the ability-earnings gradient is steeper for vocational degrees than for university degrees; and, in consequence, (iv) that there is a non trivial likelihood of obtaining negative returns for university graduates since a large fraction of them would have received higher earnings had they chosen a vocational institution instead. In our setting, this means that individuals deciding to enroll in an university instead of a vocational institution as a consequence of the 2012 reform are marginally more able (a sorting effect), but some of them would be likely better off had they pursued a vocational degree instead.

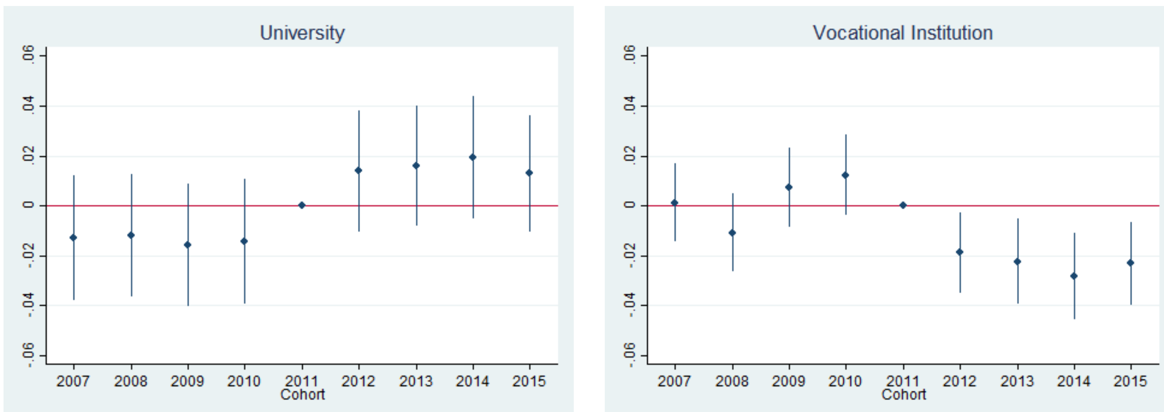

Figure 7: Dynamics of Immediate Enrollment

Figure 7 represents the dynamics of the effect on immediate enrollment by depicting the β_j interaction (i.e., $eligible \times cohort_j$) coefficient estimates described in equation 3, along with their corresponding 99% confidence intervals. Detailed estimation results and robustness checks are presented in Appendix B. The left panel depicts the evolution of the effects on university enrollment while the right panel does the same with vocational enrollment. In both cases we can see a sharp change in the signs of β_j following the 2012 reform: university enrollment increases while vocational enrollment decreases. These effects are stable over time, with a small decrease in magnitude in 2015 when the new tuition-free program was announced for 2016. In addition, the estimated interaction coefficients for cohorts 2007 to 2011 provide a strongly demanding test of the parallel trends assumption: for each year previous to the reform we cannot reject the null hypothesis of non-significance for both the university and the vocational enrollment variables.[29]

4.2 Effects on Retention

We next turn our attention to the effects on retention in tertiary education measured by our two-year enrollment and second-year dropout variables. Table 9 presents the 2SLS results for two-year enrollment. Just as with the immediate enrollment results, columns (1), (4) and (7) display the results for our base model as in equation 5. Columns (2), (5) and (8) add year fixed effects and columns (3), (6) and (9) add further control variables.

[29] Appendix C presents additional evidence in favor of the Parallel Trends assumption for all our outcomes.

Table 9: Two Year Enrollment

	General			University			Vocational		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
				Panel A: Structural Equation					
Eligible_2	0.264***	0.264***	0.246***	0.250***	0.250***	0.235***	0.014***	0.013***	0.011***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)	(0.002)	(0.002)	(0.002)
Exposed	0.034***	0.060***	0.064***	-0.007***	-0.011*	-0.014**	0.040***	0.072***	0.078***
	(0.002)	(0.006)	(0.006)	(0.001)	(0.006)	(0.006)	(0.002)	(0.003)	(0.003)
$\text{Eligible}_2 \times \text{exposed}$	0.005	0.004	0.003	0.009**	0.009**	0.009*	-0.005*	-0.005**	-0.006**
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.003)	(0.003)	(0.003)
				Panel B: First-stage Equation					
eligible_1	0.963***	0.963***	0.963***	0.963***	0.963***	0.963***	0.963***	0.963***	0.963***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Exposed	-0.006***	-0.006***	-0.006***	-0.006***	-0.006***	-0.006***	-0.006***	-0.006***	-0.006***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Cragg-Donald F-stat	1.4E7	1.4E7	1.3E7	1.4E7	1.4E7	1.3E7	1.4E7	1.4E7	1.3E7
Kleibergen-Paap F-stat	7.1E5	7.1E5	7.1E5	7.1E5	7.1E5	7.1E5	7.1E5	7.1E5	7.1E5
Year effects	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Control variables	No	No	Yes	No	No	Yes	No	No	Yes
Observations	1,347,837	1,347,837	1,347,837	1,347,837	1,347,837	1,347,837	1,347,837	1,347,837	1,347,837
Control group size	499,983	499,983	499,983	499,983	499,983	499,983	499,983	499,983	499,983
Outcome mean	0.399	0.399	0.399	0.275	0.275	0.275	0.124	0.124	0.124

Notes: Clustered standard errors at the class level in parentheses. *** p<0.01, ** p<0.05, * p<0.1. School level control variables include indicators of financing institution, rural area and geographical region. Student level control variables include gender, attendance rate and number of family members at different leves in the education system. Control group size accounts for the number of eligible individuals in the before period, while Outcome mean refers to the mean of the dependent variable of those individuals.

Panel B presents first-stage estimation results, showing how strong our initial-eligibility variable is as an instrument for the overall-two-year eligibility. Panel A presents the estimation results for the (LATE) effect of the reform. For this case we also find a null effect in overall persistence. However, there is a statistically significant diversion effect, similar (but smaller in magnitude) to that of immediate enrollment.

The loosening of credit constraints not only leads individuals to be more likely to choose universities but it also encourages them to continue for a second year of studies while diminishing the likelihood of enrolling and staying in a vocational institution. As a result, two-year university enrollment increases in almost 1 pp. (a 3 percent increase relative to non-exposed eligible individuals) while it drops by 0.5 pp. in vocational institutions (a 4 percent relative to the same group). Solis (2017) measures persistence in universities with two-year enrollment within three years after high school graduation and finds an increase of 16 pp., equivalent to a 50 percent increase. Our result is smaller but, again, his finding applies for near-the-cutoff individuals and considering access to the loan instead of a change in repayment conditions.

Results for second-year dropout in Table 10 allow us to further investigate the effects of the reform on

Table 10: Second-year Dropout

	University			Vocational		
	(1)	(2)	(3)	(4)	(5)	(6)
			Panel A: Structural Equation			
Eligible_2	-0.227***	-0.227***	-0.163***	-0.178***	-0.177***	-0.150***
	(0.006)	(0.006)	(0.006)	(0.004)	(0.004)	(0.004)
Exposed	0.048***	0.037***	0.039***	-0.016***	-0.061***	-0.027***
	(0.008)	(0.009)	(0.009)	(0.005)	(0.007)	(0.007)
$\text{Eligible}_2 \times \text{exposed}$	-0.019**	-0.019**	-0.019**	0.005	0.004	0.002
	(0.008)	(0.008)	(0.008)	(0.005)	(0.005)	(0.005)
			Panel B: First-stage Equation			
eligible_1	0.989***	0.989***	0.988***	0.996***	0.996***	0.996***
	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)
Exposed	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001
	(0.002)	(0.002)	(0.002)	(0.000)	(0.000)	(0.000)
Cragg-Donald F-stat	1.6E7	1.6E7	1.6E7	2.8E7	2.8E7	2.7E7
Kleibergen-Paap F-stat	4.0E5	4.0E5	4.9E5	4.1E6	4.1E6	4.2E6
Year effects	No	Yes	Yes	No	Yes	Yes
Control variables	No	No	Yes	No	No	Yes
Observations	386,329	386,329	375,297	273,715	273,715	272,737
Control group size	170,722	170,722	170,722	84,303	84,303	84,303
Outcome mean	0.195	0.195	0.195	0.266	0.266	0.266

Notes: Clustered standard errors at the class level in parentheses. *** p<0.01, ** p<0.05, * p<0.1. School level control variables include indicators of financing institution, rural area and geographical region. Student level control variables include gender, attendance rate and number of family members at different leves in the education system. Control group size accounts for the number of eligible individuals in the before period, while Outcome mean refers to the mean of the dependent variable of those individuals.

retention. Again, columns differ in the inclusion of year effects and other control variables. Panel A presents the effect of the reform and Panel B shows the relevance of our instrument. Among students enrolled in an university, the reform is correlated with a decrease of almost 2 pp. in the dropout rate, which amounts to a 10 percent decrease relative to non-exposed eligible individuals. For vocational institutions on the other hand, we estimate a null effect on the dropout rate.

This difference in retention effects across institutions — i.e. an improvement in university persistence with no changes in vocational institutions — results from two operating mechanisms. The first one is the sorting effect created by the reform that diverts more able individuals to enroll in universities instead of vocational institutions as already discussed in the previous section. As a result, the ability distribution ameliorates in universities while it moves in the opposite direction in vocational institutions. And given that ability is negatively correlated with the dropout probability as documented by Rau et al. (2013) and Rodriguez et al. (2016), then retention measures improve for universities while they worsen for vocational institutions.

The second mechanism comes from a perverse incentive originated by the CAE loan itself. Rau et al. (2013) build a structural model with unobserved heterogeneity for sequential schooling decisions and find that access to this particular loan reduces dropout rates in both universities and vocational institutions; a reduction that the authors discuss is explained by the fact that the CAE loan creates incentives for institutions to reduce dropout rates given their role as guarantors.[30]. In consequence, we find that persistence improves in universities following this perverse incentive which is boosted by the sorting effect, while the two mechanisms operate in opposite directions for vocational institutions.

Figure 8 presents the dynamics of the effects on our retention outcomes.[31] The top panel depicts dynamics of the effect on two-year enrollment and the bottom panel on the second-year dropout. Effects for university are shown in left panels and for vocational institutions in right panels. Out of 16 β_j interaction coefficients for $j = 2008, \dots, 2011$, 13 are not statistically significant, supporting strong evidence of the plausibility of the parallel trends assumption.[32] Regarding the university post-2012 coefficients, we can see that the ones associated to year 2012 are very close to zero for both two-year enrollment and second-year dropout. This reinforces our argument of a sorting effect in enrollment given that the 2012-effects correspond to the 2011 cohort that was exposed to the intervention only after the first year enrollment choices had already been made. The first cohort fully exposed to the reform is that of 2012, and they had to decide in 2013 whether to enroll for a second year or drop out. It is precisely in this year that we observe the biggest effect in dropout and, at the same time, the first non-zero effect in two-year enrollment. In the case of vocational institutions, the trend in the coefficients is very flat with post-2012 coefficients even closer to zero than the pre-2012 coefficients; this suggests, once again, a null effect of the reform on vocational retention.

Finally, the loosening of credit constraints analyzed in this paper will likely have other unintended consequences on outcomes beyond educational attainment, such as labor market outcomes and other long term related variables. Unfortunately, at the moment of writing of this paper there is not enough available information to properly evaluate these effects since cohorts exposed to the reform are just graduating, entering the labor market and beginning to repay their debts. Nevertheless, we can link our results to the literature that focuses on access (extensive margin) to student loans in general and to the CAE loan in particular to anticipate some potential consequences of the reform (intensive margin).

Montoya et al. (2018) suggest that the university-vocational substitution implies longer time to graduation, since programs at universities are of longer length, which in turn translates into less accumulated experience. Moreover, they analyze the long term effects (11 years after high school

[30]As already discussed in Section 2.1, higher education institutions in Chile are guarantors for CAE debtors until graduation and absorb the dropout risk. Rau et al. (2013) argue that " *[CAE loan] creates incentives for [institutions] to reduce dropout rates since they are obliged to repay if the lender drops out. In order to prevent students from dropping out, some [institutions] may lower their standards and shift resources to activities that are less successful at producing human capital but more attractive to students on the margin between continuing their education and dropping out.*"

[31]Detailed estimation results and robustness checks are presented in Appendix B.

[32]Appendix C presents additional evidence in favor of the Parallel Trends assumption for all our outcomes.

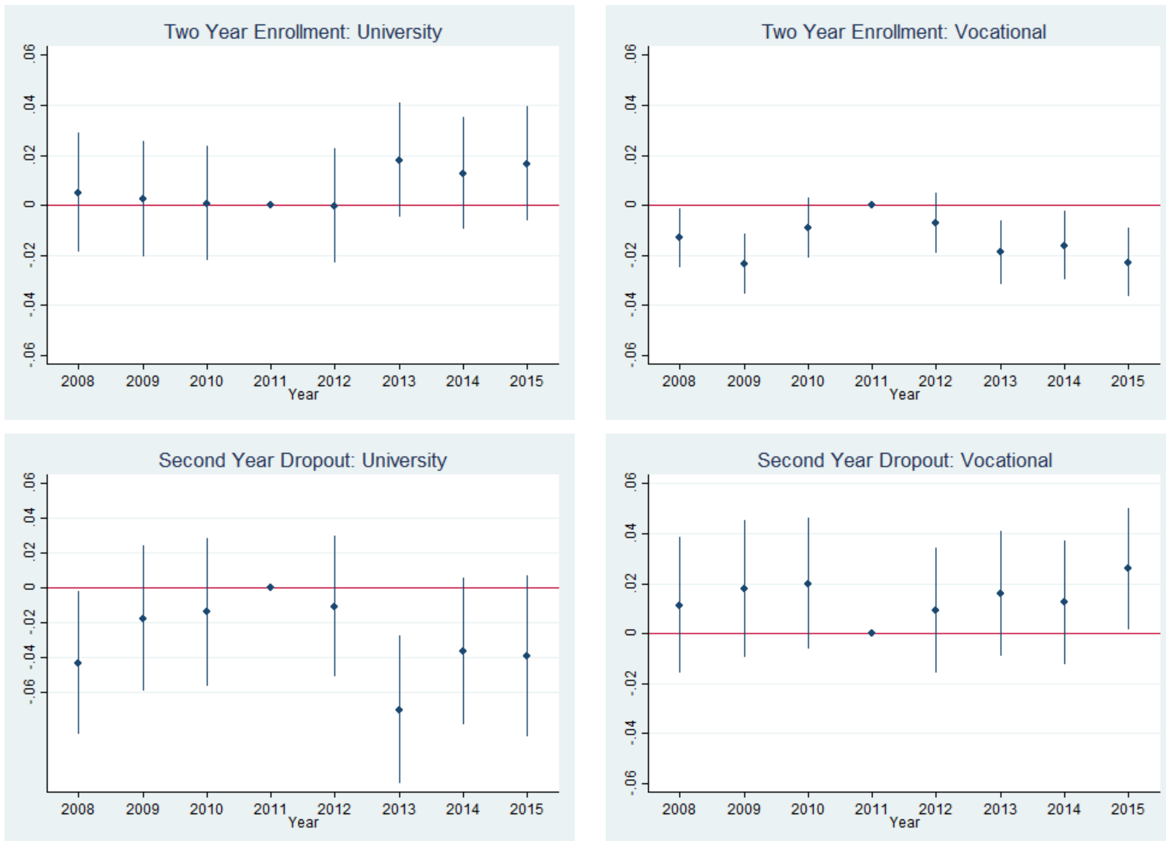

Figure 8: Dynamics of Permanence-related variables

graduation) of loan eligibility on other labor market outcomes and find that there is no difference between graduating from an university or from a vocational institution in terms of annual earnings, participation or job stability. If anything, the authors argue that the investment in university education doubles the monetary cost in comparison to vocational institutions, concluding that "*for individuals up to age 30 in Chile, college does not pay off relative to vocational education*". Bucarey et al. (2018) also follow an RDD approach to analyze the labor market returns to the CAE loan for students enrolled in universities, with similar findings: later completion, larger debt, lower experience, and no difference in wages and employment around the 475-PSU threshold. Recent literature in contexts other than the CHES has focused on several long-term outcomes such as the type of jobs chosen (Rothstein and Rouse, 2011), family planning (Kaufmann et al., 2013), homeownership (Mezza et al., 2016), retirement savings (Elliott et al., 2013), and even intergenerational effects (Kaufmann et al., 2015). Studying the long term effects of loosening credit constraints such as the CAE reform of 2012 or even the free-tuition reform of 2016 will be of great importance in the years to come.

4.3 Heterogeneity

This section analyzes the extent to which enrollment and retention effects of the reform are heterogeneous across two dimensions: gender and high school financing scheme (public versus voucher). We approach this question by estimating equation 2 in the case of enrollment and equations 4 and 5 in the case of persistence outcomes separately for female and male students (Table 11) and public and voucher schools (Table 15). Each table presents the reduced-form estimated effects for each subsample and their difference, along with the corresponding standard errors. We perform seemingly unrelated estimation (SUEST) in order to allow for correlation between subsample estimates.[33] Standard errors are clustered at the class level.

4.3.1 Female versus Male Students

Results in Table 11 suggest significant heterogeneity in enrollment decisions across the gender dimension. While there is no statistically significant difference in immediate university enrollment, the impact of the reform on vocational enrollment is stronger for female students (negative for both), with a difference of -1.3 pp. (significant at the 1% level).

Table 11: Gender Analysis

	General			University			Vocational		
	Female	Male	Difference	Female	Male	Difference	Female	Male	Difference
				Immediate Enrollment					
Eligible × after	-0.009**	0.004	-0.013**	0.023***	0.023***	0.000	-0.032***	-0.019***	-0.013***
	(0.004)	(0.005)	(0.005)	(0.005)	(0.006)	(0.006)	(0.004)	(0.004)	(0.005)
				Two Year Enrollment					
Eligible × after	-0.000	0.004	-0.004	0.010**	0.006	0.005	-0.011***	-0.001	-0.009**
	(0.004)	(0.005)	(0.006)	(0.005)	(0.006)	(0.006)	(0.003)	(0.003)	(0.004)
				Second Year Dropout					
Eligible × after				-0.011	-0.025**	0.014	-0.002	0.005	-0.007
				(0.011)	(0.011)	(0.016)	(0.006)	(0.006)	(0.009)
Cohort effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control variables	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: SUEST clustered standard errors at the class level in parentheses. *** p<0.01, ** p<0.05, * p<0.1. School level control variables include indicators of financing institution, rural area and geographical region. Student level control variables include attendance rate, *comuna* and number of family members at different leves in the education system.

[33] See Weesie (1999) for details.

In the case of males, university enrollment increases by 2.3 pp. in detriment of the vocational enrollment, which decreases by 1.9 pp. (both estimates significant at 1%). Although positive (0.4 pp.), the estimate of the overall effect is not statistically significant. In the case of females, however, the decrease of 3.2 pp. in vocational enrollment (significant at 1%) is not fully compensated by the increase of 2.3 pp. in university enrollment (significant at 1%). The overall estimated effect is a decrease of 0.9 pp. in immediate CHES enrollment for female students (significant at 5%).

The results for two-year enrollment follow a similar pattern, which is reasonable due to the way variables are defined. However, the negative overall effect for females disappears, and none of the results for males are statistically significant. The only significant result for the second-year dropout rate is a 2.5 pp. decline (significant at 5%) for male students enrolled at universities. Nevertheless, we cannot reject equality to the effect for females.

While the negative overall effect on CHES enrollment for female students might seem somewhat counterintuitive, it can be explained in light of the definition of our immediate enrollment outcome. It is possible that, while the reform induced a group of female high school graduates to switch from vocational institution enrollment to university enrollment, some of them did immediately (the year after graduation) and others delayed their enrollment decision. This delay could be an optimal response since eligibility criteria is harder to meet when enrolling in an university.[34] Evidence in Table 12 is consistent with our claim where we analyze the evolution of the proportions of female and male students delaying their PSU assessment in a DiD framework and find that the proportion of women delaying the test increases relative to the corresponding proportion of men after the reform. For this, we estimate the following equation

$$\text{delay}_{it} = \beta_0 + \beta_1 \text{female}_i + \beta_2 \text{after}_t + \beta_3 \text{female}_i \times \text{after}_t + \epsilon_{it}$$

where the outcome variable, delay_{it}, is defined as an indicator that the student sat the PSU test at least once, but not immediately — i.e., not during the corresponding year t of high school graduation.

[34] Recall that the academic criteria for CAE eligibility is more stringent if the student enrolls in an university: passing the PSU cutoff is required, whereas meeting either PSU or GPA cutoff is sufficient if the student enrolls in a vocational institution.

Table 12: PSU Delay

	Delayed PSU
Female	-0.005***
	(0.001)
After	-0.022***
	(0.001)
Female × After	0.002**
	(0.001)
Observations	2,007,043

Notes: Clustered standard errors at the class level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Still, a question remains of why the reform induced women to delay enrollment but not men. There is a vast literature relating the life-cycle production of cognitive and noncognitive skills to a wide range of outcomes such as schooling attainment, labor market outcomes, and even some risky behaviors (Cunha and Heckman, 2007, 2008; Heckman et al., 2006). While cognitive skills seem to be similarly distributed among men and women (Bound et al., 1986), there is evidence that women tend to have higher average, and less dispersed noncognitive skills than men and that this difference can explain the recent boom in higher education of women (Becker et al., 2015, 2010). The decision to delay the timing of CHES enrollment might reasonably require certain levels of noncognitive skills. Following Becker et al. (2015) that argue that grades represent "*a crude but broad measure of noncognitive skills relevant to schooling*" we present in Table 13 evidence that female students have higher noncognitive abilities than males. Here we show yearly differences in average standardized GPA between men and women. We standardize each year's data by the corresponding sample means and standard deviations of female's GPAs. Men's average GPA is systematically below (around 0.3 standard deviations) that of women through time.[35].

Table 13: GPA Gender Gap

	2007	2008	2009	2010	2011	2012	2013	2014	2015
GPA gender gap	-0.31***	-0.31***	-0.30***	-0.31***	-0.30***	-0.28***	-0.29***	-0.30***	-0.35***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)

Notes: Difference of Mean GPA (standardized by corresponding year sample mean and standard devation of female students) between male and female students and the corresponding standard errors (in parentheses). *** p<0.01, ** p<0.05, * p<0.1.

Moreover, if women tend to score lower in the PSU than men, these two factors could induce them to delay their CHES enrollment: in order to enroll in an university and have access to the CAE at the new conditions, female students might need to improve their expected performance

[35]This is consistent with evidence found in the U.S. (Becker et al., 2015; Conger and Long, 2010)

in the test and they have the patience and self-control that this task requires. Table 14 presents evidence suggesting that Chilean female students perform worse in the PSU than male students. We show yearly differences in standardized PSU scores between men and women and the corresponding standard errors. We standardize by the corresponding yearly sample means and standard deviations of women's scores. The first column shows the difference in the average of language and mathematics scores, which determines CAE eligibility. The following columns report the differences in each individual test: mathematics, language, science, and history and social sciences. Men tend to do better than women in every dimension of the test.[36] This result is consistent with other findings in the literature documenting that, while men perform better in mathematics, women tend to do better in language (Fryer and Levitt, 2010; Marks, 2008). In our restricted sample — i.e. registered students from non-private schools — men have higher scores in language, but the gender gap in this dimension is less than half the gaps in the other dimensions. The difference in the math and language average is noteworthy: men's mean score is systematically around ten percent of a standard deviation higher than women's, which makes eligibility for university-CAE harder to achieve for female students. Appendix D presents additional evidence supporting our claim.

[36] Note that this does not necessarily imply that women have lower cognitive skills. For example, the gender gap could be driven by differential effects of a competitive test setting on men and women's performance (Niederle and Vesterlund, 2010).

Table 14: PSU Scores Gender Gap

	Math and Language Average	Math	Language	Science	History and Social Sciences
2007	0.09***	0.13***	0.05***	0.11***	0.12***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)
2008	0.10***	0.16***	0.04***	0.10***	0.15***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)
2009	0.10***	0.14***	0.06***	0.13***	0.15***
	(0.004)	(0.004)	(0.003)	(0.005)	(0.004)
2010	0.09***	0.15***	0.03***	0.11***	0.12***
	(0.004)	(0.004)	(0.003)	(0.005)	(0.004)
2011	0.11***	0.15***	0.05***	0.12***	0.12***
	(0.004)	(0.004)	(0.003)	(0.005)	(0.004)
2012	0.09***	0.16***	0.03***	0.11***	0.15***
	(0.004)	(0.004)	(0.003)	(0.005)	(0.005)
2013	0.11***	0.16***	0.06***	0.12***	0.13***
	(0.004)	(0.004)	(0.003)	(0.005)	(0.005)
2014	0.11***	0.15***	0.07***	0.12***	0.13***
	(0.004)	(0.004)	(0.003)	(0.004)	(0.005)
2015	0.11***	0.16***	0.05***	0.10***	0.15***
	(0.003)	(0.003)	(0.003)	(0.004)	(0.005)

Notes: Difference of Mean PSU scores (standardized by corresponding year sample mean and standard devation of female students) between male and female students and the corresponding standard errors (in parentheses). *** p<0.01, ** p<0.05, * p<0.1.

4.3.2 Public school vs Voucher school students

Table 15 analyzes the differences in the enrollment and persistence effects between individuals graduating from voucher and public schools.[37][38] Results suggest that the diversion effect in immediate enrollment, the small pass-through to two-year enrollment, and the decrease in university dropout rates we found in the pooled sample are entirely driven by students graduating from voucher schools. The loosening of credit constraints had no effect whatsoever on eligible students coming from public schools. There are two differences between public and voucher schools that help us explain why public school graduates did not respond to the reform.

First, students in public schools tend to attain lower scores in standardized tests than students in voucher schools. Literature focused on the effects of voucher systems has found a sorting effect of the

[37] Good descriptions of the Chilean secondary education system can be found, for example, in Anand et al. (2009), Hsieh and Urquiola (2006), Mizala and Romaguera (2000), Sapelli and Vial (2002), and Torche (2005).

[38] As it was discussed in Section 2.2, students from private high schools are dropped from the analysis throughout the paper.

Table 15: High School Financing Scheme Analysis

	General			University			Vocational		
	Public	Voucher	Difference	Public	Voucher	Difference	Public	Voucher	Difference
				Immediate Enrollment					
Eligible × after	0.002	-0.001	0.003	0.009	0.030***	-0.021**	-0.007	-0.031***	0.024***
	(0.006)	(0.004)	(0.007)	(0.008)	(0.005)	(0.009)	(0.005)	(0.004)	(0.006)
				Two Year Enrollment					
Eligible × after	0.001	0.006	-0.005	-0.002	0.014***	-0.016*	0.004	-0.008**	0.012**
	(0.006)	(0.004)	(0.007)	(0.007)	(0.004)	(0.009)	(0.004)	(0.003)	(0.005)
				Second Year Dropout					
Eligible × after				-0.019	-0.022**	0.003	0.009	-0.004	0.012
				(0.014)	(0.010)	(0.017)	(0.007)	(0.006)	(0.009)
Cohort effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control variables	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: SUEST clustered standard errors at the class level in parentheses. *** p<0.01, ** p<0.05, * p<0.1. School level control variables include indicators of rural area and geographical region. Student level control variables include gender, attendance rate, *comuna* and number of family members at different leves in the education system.

introduction of voucher schools in Chile with the "best" public school students moving to voucher schools (Hsieh and Urquiola, 2006; Urquiola, 2016). Table 16 repeats the exercise of Table 14 for voucher and public schools. Graduates of public schools score systematically lower (around 0.9 standard deviations in the math and language average) than graduates of voucher schools. Appendix E presents further evidence for this claim.

Table 16: PSU Scores Voucher versus Public School Gap

	Math and Language Average	Math	Language	Science	History and Social Sciences
2007	0.09***	0.09***	0.08***	0.04***	0.07***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)
2008	0.09***	0.08***	0.08***	0.03***	0.06***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)
2009	0.08***	0.09***	0.07***	0.04***	0.08***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)
2010	0.08***	0.09***	0.08***	0.03***	0.07***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.004)
2011	0.10***	0.10***	0.08***	0.04***	0.07***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)
2012	0.09***	0.10***	0.08***	0.04***	0.07***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)
2013	0.10***	0.12***	0.08***	0.06***	0.07***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)
2014	0.08***	0.10***	0.06***	0.05***	0.06***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)
2015	0.08***	0.11***	0.06***	0.06***	0.06***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)

Notes: Differences of Mean PSU scores (standardized by corresponding year sample mean and standard devation of public school graduates) between graduates of voucher and public schools and the corresponding standard errors (in parentheses). *** p<0.01, ** p<0.05, * p<0.1.

Second, public high school students tend to be poorer than voucher school students. It has been documented that, while public schools serve mostly students coming from low-income households, voucher schools concentrate on the lower-middle and middle-income sectors (Torche, 2005). As discussed in section 2, the CAE only covers up to a "referential tuition fee" which is typically lower than actual tuition fees. This means that even being granted the loan, students (or their families) need sufficient liquidity to cover a non-negligible fraction of the tuition fee. It is arguably harder for poorer households to cover this expense.

Thus, as public school graduates tend to be poor and score low in the PSU, the reform does not actually affect their marginal incentives. Our evidence suggests that the 2012 intensive margin changes in credit access are not big enough to improve educational attainments among low-family-income students.

5 Conclusions

In this paper, we analyze the effects on enrollment and retention in higher education of a reform to student loans that loosened credit constrains by decreasing the interest rate from approximately 6% to a fixed rate of 2%, along with other minor changes that improved repayment conditions. By using a Difference-in-difference approach we exploit these changes to state-guaranteed CAE loans that took place in Chile in 2012.

Our results show that the reform had no effect on overall enrollment in the CHES. Interestingly, we find a diversion effect: enrollment to universities increased by 2.5 p.p. — a 7 percent increase relative to the enrollment rate of eligible students who graduated before the reform — in detriment of enrollment to vocational institutions that fell by 2.5 p.p. — equivalent to a decrease of 14 percent in enrollment relative to the same group. This institutional shift from vocational institutions to universities imply welfare effects given that some diverted individuals would be likely better off had they pursued a vocational degree instead.

In addition we find that retention in universities improves both in two-year enrollment and second-year dropout; while it slightly worsens in vocational institution as a result of a sorting effect in enrollment in conjunction with a perverse incentive to reduce dropout rates by institutions. We also find that for female students the reform had a negative effect on overall enrollment since the decrease in enrollment to vocational institution is not fully offset by the increase in enrollment to universities. We argue this result stems from female students delaying enrollment. Finally, all of our results are entirely driven by students from voucher schools, with null effects for students graduating from public schools.

Our findings constitute important lessons for policymakers — in the CHES and other similar — of the unintended consequences of reforms that loosen constraints. Future research on the long term implications will be of big importance in order to have a complete picture of the short and long run welfare effects.

References

Abadie, A. and Imbens, G. W. (2006). Large Sample Properties of Matching Estimators for Average Treatment Effects. *Econometrica*, 74(1):235–267.

Abadie, A. and Imbens, G. W. (2011). Bias-Corrected Matching Estimators for Average Treatment Effects. *Journal of Business & Economic Statistics*, 29(1):1–11.

Almond, D., Doyle, Jr., J. J., Kowalski, A. E., and Williams, H. (2010). Estimating Marginal Returns to Medical Care: Evidence from At-Risk Newborns *. *Quarterly Journal of Economics*, 125(2):591–634.

Anand, P., Mizala, A., and Repetto, A. (2009). Using School Scholarships to Estimate the Effect of Private Education on the Academic Achievement of Low-Income Students in Chile. *Economics of Education Review*, 28(3):370 – 381.

Angrist, J., Autor, D., Hudson, S., and Pallais, A. (2016). Evaluating Post-Secondary Aid: Enrollment, Persistence, and Projected Completion Effects. *National Bureau of Economic Research*, Working Pa.

Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996). Identification of Causal Effects Using Instrumental Variables. *Journal of the American Statistical Association*, 91(434):444–455.

Angrist, J. D. and Lavy, V. (1999). Using Maimonides' Rule to Estimate the Effect of Class Size on Scholastic Achievement. *The Quarterly Journal of Economics*, 114(2):533–575.

Angrist, J. D. and Rokkanen, M. (2015). Wanna Get Away? Regression Discontinuity Estimation of Exam School Effects Away From the Cutoff. *Journal of the American Statistical Association*, 110(512):1331–1344.

Barreca, A. I., Guldi, M., Lindo, J. M., and Waddell, G. R. (2011). Saving Babies? Revisiting the effect of very low birth weight classification. *The Quarterly Journal of Economics*, 126(4):2117–2123.

Barreca, A. I., Lindo, J. M., and Waddell, G. R. (2015). Heaping-Induced Bias in Regression-Discontinuity Designs. *Economic Inquiry*, pages n/a–n/a.

Becker, G. S., Hubbard, W. H. J., and Murphy, K. M. (2010). New Directions in the Economic Analysis of Human Capital: The Market for College Graduates and the Worldwide Boom in Higher Education of Women. *The American Economic Review: Papers & Proceedeings*, 100:229–233.

Becker, G. S., Hubbard, W. H. J., Murphy, K. M., Journal, S., Fall, N., Becker, G. S., Hubbard, W. H. J., and Murphy, K. M. (2015). Explaining the Worldwide Boom in Higher Education of Women. *Journal of Human Capital*.

Black, S. E. (1999). Do better schools matter? Parental valuation of elementary education. *Quarterly Journal of Economics*.

Bound, J., Griliches, Z., and Hall, B. H. (1986). Wages, Schooling and IQ of Brothers and Sisters: Do the Family Factors Differ? *International Economic Review*, 27(1):77–105.

Bravo, D. and Rau, T. (2012). Effects of Large-scale Youth Employment Subsidies: Evidence from a Regression Discontinuity Design.

Bucarey, A. (2017). Who Pays for Free College? Crowding Out on Campus. *Job Market Paper*, pages 1–71.

Bucarey, A., Contreras, D., and Muñoz, P. (2018). Labor Market Returns to Student Loans.

Busso, M., DiNardo, J., and McCrary, J. (2014). New Evidence on the Finite Sample Properties of Propensity Score Reweighting and Matching Estimators. *The Review of Economics and Statistics*, 96(5):885–897.

Calonico, S., Cattaneo, M. D., Farrell, M. H., Titiunik, R., Bonhomme, S., Canay, I., Drukker, D., Imai, K., Jansson, M., Kilian, L., Kline, P., Ma, X., Santos, A., and Vazquez, G. (2016). Regression Discontinuity Designs Using Covariates.

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica*, 82(6):2295–2326.

Cameron, S. V. and Heckman, J. J. (2001). The Dynamics of Educational Attainment for Black, Hispanic, and White Males. *Journal of Political Economy*.

Card, D., Lee, D. S., Pei, Z., and Weber, A. (2015). Inference on Causal Effects in a Generalized Regression Kink Design. *Econometrica*, 83(6):2453–2483.

Carneiro, P. and Heckman, J. J. (2002). The evidence on credit constraints in post-secondary schooling. *Economic Journal*.

Cattaneo, M. D. and Escanciano, J. C., editors (2017). *Regression Discontinuity Designs (Advances in Econometrics, Volume 38)*. Emerald Publishing Limited.

Cattaneo, M. D., Keele, L., Titiunik, R., and Vazquez-Bare, G. (2016). Interpreting Regression Discontinuity Designs with Multiple Cutoffs. *The Journal of Politics*, 78(4):1229–1248.

Chatterjee, S. and Ionescu, F. (2012). Insuring student loans against the financial risk of failing to complete college. *Quantitative Economics*, 3(3):393–420.

Chernozhukov, V., Fernandez-Val, I., and Melly, B. (2013). Inference on Counterfactual Distributions. *Econometrica*, 81(6):2205–2268.

Chetty, R., Friedman, J. N., Olsen, T., and Pistaferri, L. (2011). Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records. *The Quarterly Journal of Economics*, 126(2):749–804.

Conger, D. and Long, M. C. (2010). Why Are Men Falling Behind? Gender Gaps in College Performance and Persistence. *The Annals of the American Academy of Political and Social Science*, 627(1):184–214.

Cornwell, C., Mustard, D. B., and Sridhar, D. J. (2006). The Enrollment Effects of Merit-Based Financial Aid: Evidence from Georgia's HOPE Program. *Journal of Labor Economics*, 24(4):761–786.

Cunha, F. and Heckman, J. J. (2007). The technology of skill formation. *The American Economic Review*, 97(2):31–47.

Cunha, F. and Heckman, J. J. (2008). Formulating, identifying and estimating the technology of cognitive and noncognitive skill formation. *The Journal of Human Resources*, 43(4):738–782.

DiNardo, J., Fortin, N. M., and Lemieux, T. (1996). Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach. *Econometrica*, 64(5):1001.

Drake, C. (1993). Effects of Misspecification of the Propensity Score on Estimators of Treatment Effect. *Biometrics*, 49(4):1231.

Dynarski, S. and Scott-Clayton, J. (2013). Financial Aid Policy: Lessons from Research. *The Future of Children*, 23(1):67–91.

Dynarski, S. M. (2003). Does Aid Matter? Measuring the Effect of Student Aid on College Attendance and Completion. *American Economic Review*, 93(1):279–288.

Elliott, W., Grinstein-Weiss, M., and Nam, I. (2013). Student Debt and Declining Retirement Savings.

Espinoza, R. and Urzúa, S. (2015). The economic consequences of implementing tuition free tertiary education in Chile. *Revista de Educacion*, pages 10–37.

Fack, G. and Grenet, J. (2015). Improving College Access and Success for Low-Income Students: Evidence from a Large Need-Based Grant Program. *American Economic Journal: Applied Economics*, 7(2):1–34.

Fan, J., Yao, Q., and Tong, H. (1996). Estimation of Conditional Densities and Sensitivity Measures in Nonlinear Dynamical Systems. *Biometrika*, 83(1):189–206.

Fan, Y., Sherman, R., and Shum, M. (2014). Identifying Treatment Effects Under Data Combination. *Econometrica*, 82(2):811–822.

Fryer, R. G. and Levitt, S. D. (2010). An Empirical Analysis of the Gender Gap in Mathematics. *American Economic Journal: Applied Economics*, 2(2):210–240.

Gerard, F., Rokkanen, M., and Rothe, C. (2018). Bounds on Treatment Effects in Regression Discontinuity Designs with a Manipulated Running Variable. *National Bureau of Economic Research*, Working Pa.

Glocker, D. (2011). The effect of student aid on the duration of study. *Economics of Education Review*, 30(1):177–190.

Hahn, J., Todd, P., and Klaauw, W. (2001). Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design. *Econometrica*, 69(1):201–209.

Hall, P., Racine, J., and Li, Q. (2004). Cross-Validation and the Estimation of Conditional Probability Densities. *Journal of the American Statistical Association*, 99(468):1015–1026.

Heckman, J. J., Ichimura, H., and Todd, P. (1998). Matching As An Econometric Evaluation Estimator. *Review of Economic Studies*, 65(2):261–294.

Heckman, J. J., Stixrud, J., and Urzua, S. (2006). The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior. *Journal of Labor Economics*.

Herzog, S. (2005). Measuring Determinants of Student Return VS. Dropout/Stopout VS. Transfer: A First-to-Second Year Analysis of New Freshmen. *Research in Higher Education*, 46(8):883–928.

Hsieh, C.-T. and Urquiola, M. (2006). The Effects of Generalized School Choice on Achievement and Stratification: Evidence from Chile's Voucher Program. *Journal of Public Economics*, 90(8):1477–1503.

Imai, K. and Ratkovic, M. (2014). Covariate balancing propensity score. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*.

Imbens, G. W. and Lemieux, T. (2008). Regression discontinuity designs: A guide to practice. *Journal of Econometrics*, 142(2):615–635.

Ingresa (2010). Balance Anual 2006-2010. Technical report.

Ingresa (2015). Cuenta Pública Año 2015. Technical report.

Ito, K. and Sallee, J. (2014). The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel-Economy Standards. Technical report, National Bureau of Economic Research, Cambridge, MA.

Jeffery, R. L. (1925). The Continuity of a Function Defined by a Definite Integral. *The American Mathematical Monthly*, 32(6):297.

Kaufmann, K., Messner, M., and Solis, A. (2013). Returns to Elite Higher Education in the Marriage Market: Evidence from Chile. *SSRN Electronic Journal*.

Kaufmann, K. M., Messner, M., and Solis, A. (2015). Elite Higher Education, the Marriage Market and the Intergenerational Transmission of Human Capital.

Keele, L. J. and Titiunik, R. (2015). Geographic Boundaries as Regression Discontinuities. *Political Analysis*, 23(1):127–155.

Kleven, H. J. and Waseem, M. (2013). Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan. *The Quarterly Journal of Economics*, 128(2):669–723.

Lee, D. S. (2008). Randomized experiments from non-random selection in U.S. House elections. *Journal of Econometrics*, 142(2):675–697.

Lee, D. S. and Lemieux, T. (2010). Regression Discontinuity Designs in Economics. *Journal of Economic Literature*, 48(2):281–355.

Lochner, L. J. and Monge-Naranjo, A. (2011). The nature of credit constraints and human capital. *American Economic Review*.

Marks, G. N. (2008). Accounting for the Gender Gaps in student Performance in Reading and Mathematics: Evidence from 31 Countries. *Oxford Review of Education*, 34(1):89–109.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142(2):698–714.

Mezza, A. A., Ringo, D. R., Sherlund, S. M., and Sommer, K. (2016). On the Effect of Student Loans on Access to Homeownership. *Finance and Economics Discussion Series*, 2016(010):1–35.

Ministry of Education (2016). Memoria Financiamiento Estudiantil. Technical report.

Mizala, A. and Romaguera, P. (2000). School Performance and Choice: The Chilean Experience. *The Journal of Human Resources*, 35(2):392–417.

Montoya, A. M., Noton, C., and Solis, A. (2018). The Returns to College Choice: Loans, Scholarships and Labor Outcomes.

Niederle, M. and Vesterlund, L. (2010). Explaining the Gender Gap in Math Test Scores: The Role of Competition. *Journal of Economic Perspectives*, 24(2):129–144.

Nielsen, H. S., Sørensen, T., and Taber, C. (2010). Estimating the Effect of Student Aid on College Enrollment: Evidence from a Government Grant Policy Reform. *American Economic Journal: Economic Policy*, 2(2):185–215.

Perna, L. W. and Titus, M. A. (2004). Understanding Differences in the Choice of College Attended: The Role of State Public Policies. *The Review of Higher Education*, 27(4):501–525.

Rau, T., Rojas, E., and Urzúa, S. (2013). Loans for Higher Education: Does the Dream Come True? Technical report, National Bureau of Economic Research, Cambridge, MA.

Riegg, S. K. (2008). Causal Inference and Omitted Variable Bias in Financial Aid Research: Assessing Solutions. *The Review of Higher Education*, 31(3):329–354.

Rodriguez, J., Urzua, S., and Reyes, L. (2016). Heterogeneous Economic Returns to Post-Secondary Degrees: Evidence from Chile. *Journal of Human Resources*, 51(2):416–460.

Rothstein, J. and Rouse, C. E. (2011). Constrained after college: Student loans and early-career occupational choices. *Journal of Public Economics*, 95(1-2):149–163.

Saez, E. (2010). Do Taxpayers Bunch at Kink Points? *American Economic Journal: Economic Policy*, 2(3):180–212.

Sapelli, C. and Vial, B. (2002). The Performance of Private and Public Schools in the Chilean Voucher System. *Cuadernos de economía*, 39:423 – 454.

Solis, A. (2017). Credit Access and College Enrollment. *Journal of Political Economy*, 125(2):562–622.

Stinebrickner, R. and Stinebrickner, T. (2008). The Effect of Credit Constraints on the College Drop-Out Decision: A Direct Approach Using a New Panel Study. *American Economic Review*, 98(5):2163–2184.

Stinebrickner, T. and Stinebrickner, R. (2012). Learning about Academic Ability and the College Dropout Decision. *Journal of Labor Economics*, 30(4):707–748.

Tan, Z. (2010). Bounded, efficient and doubly robust estimation with inverse weighting. *Biometrika*.

Thistlethwaite, D. L. and Campbell, D. T. (1960). Regression-discontinuity analysis: An alternative to the ex post facto experiment. *Journal of Educational Psychology*, 51(6):309–317.

Torche, F. (2005). Privatization Reform and Inequality of Educational Opportunity: The Case of Chile. *Sociology of Education*, 78(4):316–343.

Urquiola, M. (2016). Chapter 4 - Competition Among Schools: Traditional Public and Private Schools. volume 5 of *Handbook of the Economics of Education*, pages 209 – 237. Elsevier.

Van der Klaauw, W. (2002). Estimating the effect of financial aid offers on college enrollment: A regression-discontinuity approach. *International Economic Review*.

Weesie, J. (1999). Seemingly Unrelated Estimation and the Cluster-Adjusted Sandwich Estimator. Technical Report Stata Technical Bulletin 52, Stata Corporation.

World Bank (2011). Chile's State-Guaranteed Student Loan Program (CAE) (English).

Appendix

Appendix I.A

Proofs of Propositions 1, 3 and 4

***Proposition** 1: Sufficient Condition for Identification*

Let $(R, \mathbf{X}, U)$ be a random vector of length $J+2$ with joint density function $f_{R,\mathbf{X},U}(r, \mathbf{x}, u)$ and support $\Theta \subseteq \mathbb{R}^{J+2}$. Potential outcomes are $y_1 = g_1(r, \mathbf{x}, u)$ and $y_0 = g_0(r, \mathbf{x}, u)$; where $g_1(\cdot)$, $g_0(\cdot)$ are continuous real-valued bounded functions.

If the conditional density function $f_{\mathbf{X},U|R}(\mathbf{x}, u|r)$ is continuous, then both conditional expectations $E[y_1|r]$ and $E[y_0|r]$ are also continuous; and therefore, following Hahn et al. (2001)'s theorem of identification, $\lim_{r\downarrow c} E[y|r] - \lim_{r\uparrow c} E[y|r]$ identifies $E[\tau|c]$.

Proof. For $t = 0, 1$ rewrite the conditional expectations as:

$$E[y_t|r] = \int_\Omega g_t(r, \mathbf{x}, u) f_{\mathbf{X},U|R}(\mathbf{x}, u|r) d\mathbf{x} du$$

Define $h_t(r, \mathbf{x}, u) = g_t(r, \mathbf{x}, u) f_{\mathbf{X},U|R}(\mathbf{x}, u|r)$. Probability density function $f(\cdot)$ is real-valued and bounded by construction. Then, under continuity assumption and with $g_t(\cdot)$ continuous, real-valued and bounded we have by algebra of continuous functions that $h_t(\cdot)$ is real-valued, bounded, continuous in r at c for each $(\mathbf{x}, u)$ and measurable in $(\mathbf{x}, u)$ for each r.

Finally, following Jeffery (1925) we have that

$$\int_\Omega h_t(r, \mathbf{x}, u) d\mathbf{x} du$$

is continuous at $r = c$ □

***Proposition** 3: Asymptotic Normality for $\widehat{w}(\boldsymbol{x}, r)$*

For

$$\widehat{w}(\mathbf{x}, r) = \frac{\widehat{f}_{\mathbf{X}|R}(\mathbf{x}|r)}{\widehat{f}_{\mathbf{X}|R^*}(\mathbf{x}|r)}$$

estimated as discussed in section 4.1, where both conditional densities are estimated separately through LLR we have:

$$\sqrt{nh_1 \dots h_J h_r}\left(\widehat{w}(\mathbf{x}, r) - w(\mathbf{x}, r) - \text{e.b.}^w(\mathbf{x}, r; h_1^2, \dots, h_J^2, h_r^2)\right) \xrightarrow{d} N\left(0, \frac{w(\mathbf{x}, r)}{f_{\mathbf{X},R^*}(\mathbf{x}, r)} R_K^2\right)$$

Proof. Following Fan et al. (1996), we have that the LLR estimator of the conditional density has asymptotic normal distribution:

$$\sqrt{nh_1 \dots h_J h_r}\left(\widehat{f}(\mathbf{x}|r) - f(\mathbf{x}|r) - \frac{\kappa_2}{2}\left[(h_1 \dots h_J)^2 f^{(2)}_{\mathbf{xx}}(\mathbf{x}|r) + h_r^2 f^{(2)}_{rr}(\mathbf{x}|r)\right]\right) \xrightarrow{d} N\left(0, \frac{f(\mathbf{x}|r)R_K^2}{f(r)}\right)$$

Since both conditional densities $\widehat{f}_{\mathbf{X}|R}(\mathbf{x}|r)$ and $\widehat{f}_{\mathbf{X}|R^*}(\mathbf{x}|r)$ are estimated separately, we have that:

$$\begin{aligned}\widehat{f}_{\mathbf{X}|R^*}(\mathbf{x}|r) &\xrightarrow{p} f_{\mathbf{X}|R^*}(\mathbf{x}|r) + \frac{\kappa_2}{2}\left[(h_1 \dots h_J)^2 f^{(2)}_{(\mathbf{xx})\mathbf{X}|R^*}(x|r) + h_r^2 f^{(2)}_{(rr)\mathbf{X}|R^*}(x|r)\right] \\ &\xrightarrow{p} f_{\mathbf{X}|R^*}(\mathbf{x}|r) + \text{e.b.}^{f^*}(\mathbf{x}, r; h_1^2, \dots, h_J^2, h_r^2)\end{aligned}$$

while the counterfactual conditional density has:

$$\begin{aligned}E\left[\widehat{f}_{\mathbf{X}|R}(\mathbf{x}|r)\right] &= f_{\mathbf{X}|R}(\mathbf{x}|r) + \frac{\kappa_2}{2}\left[(h_1 \dots h_J)^2 f^{(2)}_{(\mathbf{xx})\mathbf{X}|R}(\mathbf{x}|r) + h_r^2 f^{(2)}_{(rr)\mathbf{X}|R}(\mathbf{x}|r)\right] \\ &= f_{\mathbf{X}|R}(\mathbf{x}|r) + \text{e.b.}^{f}(\mathbf{x}, r; h_1^2, \dots, h_J^2, h_r^2)\end{aligned}$$

and

$$V\left[\widehat{f}_{\mathbf{X}|R}(\mathbf{x}|r)\right] = \frac{1}{nh_1 \dots h_J h_r}\frac{f_{\mathbf{X}|R}(\mathbf{x}|r)}{f_R(r)}R_K^2$$

with $R_K = \int k^2(v)dv$ a known parameter. Then combining both results:

$$\begin{aligned}E\left[\widehat{w}(\mathbf{x}, r)\right] &= \frac{f_{\mathbf{X}|R}(\mathbf{x}|r) + \text{e.b.}^{f}(\mathbf{x}, r; h_1^2, \dots, h_J^2, h_r^2)}{f_{\mathbf{X}|R^*}(\mathbf{x}|r) + \text{e.b.}^{f^*}(\mathbf{x}, r; h_1^2, \dots, h_J^2, h_r^2)} \\ &= w(\mathbf{x}, r) + \frac{\text{e.b.}^{f}(\cdot)f_{\mathbf{X}|R^*}(\mathbf{x}|r) - \text{e.b.}^{f^*}(\cdot)f_{\mathbf{X}|R}(\mathbf{x}|r)}{f_{\mathbf{X}|R^*}(\mathbf{x}|r)(f_{\mathbf{X}|R^*}(\mathbf{x}|r) + \text{e.b.}^{f^*}(\cdot))} \\ &= w(\mathbf{x}, r) + \text{e.b.}^{w}(\mathbf{x}, r; h_1^2, \dots, h_J^2, h_r^2)\end{aligned}$$

and

$$\begin{aligned}V\left[\widehat{w}(\mathbf{x}, r)\right] &= \frac{1}{\left(f_{\mathbf{X}|R^*}(\mathbf{x}|r) + \text{e.b.}^{f^*}(\cdot)\right)^2}V\left[\widehat{f}_{\mathbf{X}|R}(\mathbf{x}|r)\right] \\ &= \frac{1}{\left(f_{\mathbf{X}|R^*}(\mathbf{x}|r) + \text{e.b.}^{f^*}(\cdot)\right)^2}\frac{1}{nh_1 \dots h_J h_r}\frac{f_{\mathbf{X}|R}(\mathbf{x}|r)}{f_R(r)}R_K^2\end{aligned}$$

Ignoring the e.b.$^{f}(\cdot)$ term we then have:

$$V\left[\widehat{w}(\mathbf{x}, r)\right] = \frac{1}{nh_1 \dots h_J h_r}\frac{w(\mathbf{x}, r)}{f_{\mathbf{X},R^*}(\mathbf{x}, r)}R_K^2$$

□

***Proposition** 4: Asymptotic Normality for $\widehat{E}[\tau|c]$*

For

$$\widehat{E}[\tau|c] = \widehat{a}_R - \widehat{a}_L$$

estimated as discussed in section 4.2, where both reweighted outcome conditional expectations are estimated separately through LLR we have:

$$\sqrt{nh_r}\left(\widehat{E}[\tau|c] - E[\tau|c] - \text{e.b.}^E(h_r^2; c)\right) \xrightarrow{d} N\left(0, \left(z_+(c)\frac{\sigma_+^2(c)}{f_R(c)}R_K + z_-(c)\frac{\sigma_-^2(c)}{f_R(c)}R_K\right)\right)$$

where $z(c) = E\left[w^2(\mathbf{x}, c)|c\right]$ and $\sigma^2(c) = E\left[\epsilon^2|c\right]$ on each side of the threshold.

Proof. LLR for the estimation of a regression function has asymptotic normal distribution (**?**):

$$\sqrt{nh_r}\left(\widehat{E}[y|r] - E[y|r] - h_r^2\kappa_2\left[\frac{E^{(2)}[y|r]}{2}\right]\right) \xrightarrow{d} N\left(0, \frac{\sigma^2(r)R_K}{f(r)}\right)$$

with $R_K = \int k^2(v)dv$ and $\kappa_2 = \int v^2 k(v)dv$ known parameters. For ease of exposition write:

$$m_t(c) = E[y_t|c] = \int_\Omega g_t(c, \mathbf{x}, u) f_{\mathbf{X},U|R}(\mathbf{x}, u|c) d\mathbf{x} du$$

for $t = 0, 1$. Then, following Calonico et al. (2014) we have

$$\begin{aligned}
E[\widehat{a}_R - \widehat{a}_L] &= E[\widehat{a}_R] - E[\widehat{a}_L] \\
&= \left(m_1(c) + \frac{h_r^2}{2}\kappa_2 m_1^{(2)}(c)\right) - \left(m_0(c) + \frac{h_r^2}{2}\kappa_2 m_0^{(2)}(c)\right) \\
&= \int_\Omega \left[g_1(c, \mathbf{x}, u) - g_0(c, \mathbf{x}, u)\right] f_{\mathbf{X},U|R}(\mathbf{x}, u|c) d\mathbf{x} du + \frac{h_r^2}{2}\kappa_2\left(m_1^{(2)}(c) - m_0^{(2)}(c)\right) \\
&= E[\tau|c] + \text{e.b.}^E(h_r^2; c)
\end{aligned}$$

and

$$\begin{aligned}
V[\widehat{a}_R - \widehat{a}_L] &= V[\widehat{a}_R] + V[\widehat{a}_L] - 2\text{Cov}[\widehat{a}_R, \widehat{a}_L] \\
&= \frac{1}{nh_r}\left(z_+(c)\frac{\sigma_+^2(c)}{f_R(c)}R_K\right) + \frac{1}{nh_r}\left(z_-(c)\frac{\sigma_-^2(c)}{f_R(c)}R_K\right) - 2\text{Cov}[\widehat{a}_R, \widehat{a}_L]
\end{aligned}$$

Theoretically, $\text{Cov}[\widehat{a}_R, \widehat{a}_L] \neq 0$ since both estimators $\widehat{a}_R$ and $\widehat{a}_L$ use, at some extent, same observations for the estimation of the counterfactual conditional density $f_{\mathbf{X},U|R}(\mathbf{x}, u|r)$ in the first stage. However, this paper's proposal disregards this term given that it is negligible and it was properly verified in the Monte Carlo simulations. □

Appendix I.B

Examples for the implications of a discontinuously distributed running variable

Consider the following two examples to illustrate the arguments in section 2.2. Let $g_0(r, \mathbf{x}, u) = g_1(r, \mathbf{x}, u) = x$ so that $\tau = 0$ and the estimand of interest $E[\tau|c]$ is also zero. In this simple setting, define $c = 0$, x a binary covariate (scalar) and r the running variable with the following marginal $f_R(r)$ and conditional $p_{X|R}(X = 1|r)$ distributions respectively:

Design 1: Discontinuously distributed running variable with continuous conditional probability function

$$f_R(r) = \left[\frac{1}{\sqrt{2\pi}} \exp\left(-\frac{r^2}{2}\right)\right]^{\mathbf{1}(r<0)} \cdot \left[\frac{1}{\sqrt{\pi}} \exp\left(-r^2\right)\right]^{\mathbf{1}(r\geq 0)}$$

a Gaussian-shaped density function and

$$p_{X|R}(X = 1|r) = \begin{cases} 0.75 & \text{if } r < -1 \\ 0.5 & \text{if } r \geq -1 \end{cases}$$

a conditional probability function that is continuous at $r = 0$. Conversely, note that $f_R(r)$ is discontinuous near the threshold: $\lim_{r\downarrow 0} f_R(r) = 1/\sqrt{\pi}$ and $\lim_{r\uparrow 0} f_R(r) = 1/\sqrt{2\pi}$.

Then, following the identification theorem

$$\begin{aligned} \lim_{r\downarrow 0} E[y|r] - \lim_{r\uparrow 0} E[y|r] &= \lim_{r\downarrow 0} E[y_1 - y_0|r] + \lim_{r\downarrow 0} E[y_0|r] - \lim_{r\uparrow 0} E[y_0|r] \\ &= \lim_{r\downarrow 0} E[x - x|r] + \lim_{r\downarrow 0} E[x|r] - \lim_{r\uparrow 0} E[x|r] \\ &= \lim_{r\downarrow 0} E[0|r] + \underbrace{\lim_{r\downarrow 0} p[X = 1|r] - \lim_{r\uparrow 0} p[X = 1|r]}_{0.5-0.5=0} \\ &= 0 \end{aligned}$$

identifies $E[\tau|0] = 0$ despite $f_R(r)$ being discontinuous at the cutoff point.

Design 2: Continuously distributed running variable with discontinuous conditional probability function

$$f_R(r) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{r^2}{2}\right)$$

a standard normal distribution and

$$p_{X|R}(X = 1|r) = \begin{cases} 0.75 & \text{if } r < 0 \\ 0.5 & \text{if } r \geq 0 \end{cases}$$

In this case, $p_{X|R}(X = 1|r)$ is discontinuous exactly at the threshold, while R has a continuous density function. Then,

$$\begin{aligned}
\lim_{r\downarrow 0} E\left[y|r\right] - \lim_{r\uparrow 0} E\left[y|r\right] &= \lim_{r\downarrow 0} E\left[y_1 - y_0|r\right] + \lim_{r\downarrow 0} E\left[y_0|r\right] - \lim_{r\uparrow 0} E\left[y_0|r\right] \\
&= \lim_{r\downarrow 0} E\left[x - x|r\right] + \lim_{r\downarrow 0} E\left[x|r\right] - \lim_{r\uparrow 0} E\left[x|r\right] \\
&= \lim_{r\downarrow 0} E\left[0|r\right] + \underbrace{\lim_{r\downarrow 0} p\left[X = 1|r\right] - \lim_{r\uparrow 0} p\left[X = 1|r\right]}_{0.5-0.75} \\
&= -0.25
\end{aligned}$$

does not identify $E\left[\tau|0\right] = 0$ even though R is continuously distributed.

Figure 9 depicts this example. The upper and lower panel describe designs 1 and 2 respectively. In this simple case, (dis)continuity of the density function of R is irrelevant for identification. Instead, continuity of $E\left[y_1|r\right]$ and $E\left[y_0|r\right]$ hinges on continuity of $p_{X|R}(X = 1|r)$.

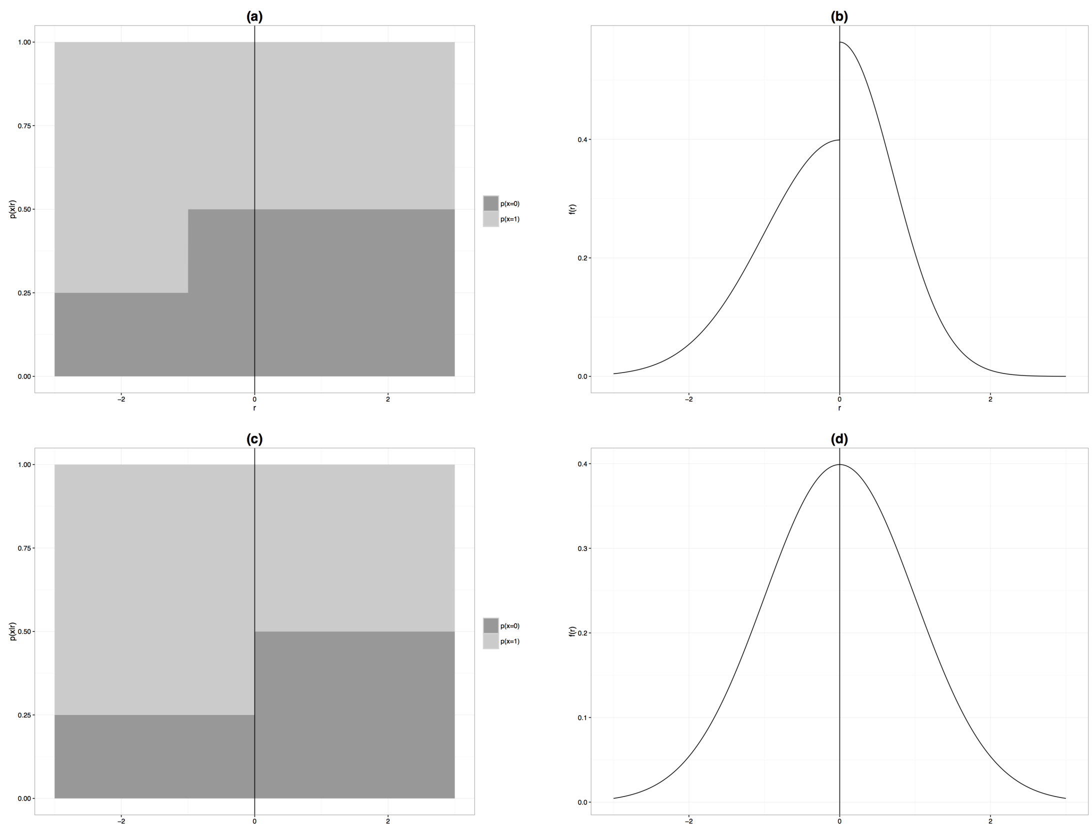

Figure 9: Identification Examples

Note: Panels (a) and (c) show the conditional distributions for design 1 and 2 respectively, while panels (b) and (d) illustrate the marginal distribution.

Appendix II.A

Table 17: Descriptive Statistics

	Cohort									
	2007	2008	2009	2010	2011	2012	2013	2014	2015	Pooled
Immediate Enrollment	0.449	0.457	0.456	0.466	0.488	0.515	0.538	0.543	0.540	0.496
by Institution										
University	0.284	0.292	0.274	0.273	0.285	0.299	0.297	0.298	0.299	0.289
Vocational	0.165	0.165	0.182	0.193	0.203	0.216	0.242	0.245	0.242	0.207
by Gender										
Females	0.436	0.443	0.451	0.463	0.488	0.517	0.539	0.540	0.535	0.492
Males	0.466	0.474	0.462	0.468	0.488	0.513	0.537	0.546	0.546	0.501
by High School										
Public	0.400	0.416	0.418	0.423	0.448	0.462	0.494	0.505	0.503	0.451
Voucher	0.489	0.487	0.483	0.497	0.516	0.545	0.563	0.565	0.562	0.526
Two-Year Enrollment		0.330	0.343	0.346	0.354	0.355	0.379	0.398	0.403	0.364
by Institution										
University		0.223	0.227	0.219	0.220	0.217	0.228	0.226	0.228	0.223
Vocational		0.107	0.115	0.127	0.134	0.138	0.151	0.172	0.175	0.141
by Gender										
Females		0.324	0.338	0.347	0.360	0.363	0.386	0.404	0.408	0.367
Males		0.336	0.347	0.345	0.348	0.346	0.370	0.391	0.398	0.361
by High School										
Public		0.289	0.309	0.311	0.315	0.319	0.333	0.357	0.367	0.324
Voucher		0.362	0.367	0.371	0.383	0.380	0.404	0.421	0.424	0.391
Second Year Dropout		0.266	0.250	0.241	0.239	0.273	0.265	0.261	0.257	0.257
by Institution										
University		0.216	0.221	0.202	0.195	0.239	0.237	0.238	0.232	0.223
Vocational		0.352	0.302	0.300	0.303	0.320	0.304	0.289	0.287	0.305
by Gender										
Females		0.256	0.235	0.231	0.224	0.256	0.254	0.250	0.244	0.244
Males		0.277	0.267	0.252	0.257	0.291	0.277	0.273	0.272	0.271
by High School										
Public		0.277	0.256	0.255	0.255	0.288	0.279	0.277	0.272	0.270
Voucher		0.259	0.246	0.232	0.230	0.264	0.258	0.253	0.249	0.249
Eligible	0.752	0.779	0.767	0.771	0.767	0.769	0.782	0.794	0.815	0.778
PSU	474.827	474.858	474.420	472.508	475.556	474.615	476.212	476.832	478.672	475.402
GPA	5.564	5.600	5.581	5.585	5.581	5.596	5.611	5.643	5.683	5.606
Females	0.541	0.546	0.537	0.531	0.527	0.534	0.531	0.532	0.529	0.534
Public School	0.443	0.423	0.422	0.421	0.405	0.362	0.363	0.365	0.367	0.396
Obervations	146,410	147,480	171,300	180,306	184,636	169,824	174,909	174,789	178,144	1,527,798

Appendix II.B

Table 18: Dynamics of Immediate Enrollment

	University		Vocational	
	(1)	(2)	(3)	(4)
Eligible	0.297***	0.278***	-0.034***	-0.034***
	(0.007)	(0.007)	(0.004)	(0.004)
Cohort 2007	0.013***	0.015***	-0.039***	-0.043***
	(0.003)	(0.004)	(0.005)	(0.005)
Cohort 2008	0.012***	0.011***	-0.029***	-0.032***
	(0.003)	(0.004)	(0.006)	(0.005)
Cohort 2009	0.001	0.003	-0.027***	-0.027***
	(0.003)	(0.004)	(0.006)	(0.005)
Cohort 2010	-0.002	0.001	-0.020***	-0.021***
	(0.003)	(0.004)	(0.006)	(0.005)
Cohort 2012	0.002	0.003	0.028***	0.030***
	(0.003)	(0.004)	(0.005)	(0.005)
Cohort 2013	-0.006**	-0.004	0.056***	0.057***
	(0.002)	(0.004)	(0.006)	(0.005)
Cohort 2014	-0.011***	-0.010***	0.066***	0.067***
	(0.002)	(0.004)	(0.006)	(0.006)
Cohort 2015	-0.011***	-0.010***	0.059***	0.060***
	(0.002)	(0.004)	(0.006)	(0.005)
Eligible × cohort 2007	-0.013	-0.016*	0.001	0.003
	(0.010)	(0.010)	(0.006)	(0.006)
Eligible × cohort 2008	-0.012	-0.012	-0.011*	-0.010
	(0.010)	(0.009)	(0.006)	(0.006)
Eligible × cohort 2009	-0.016*	-0.017*	0.007	0.007
	(0.009)	(0.009)	(0.006)	(0.006)
Eligible × cohort 2010	-0.014	-0.017*	0.012*	0.013**
	(0.010)	(0.010)	(0.006)	(0.006)
Eligible × cohort 2012	0.014	0.010	-0.019***	-0.017***
	(0.009)	(0.009)	(0.006)	(0.006)
Eligible × cohort 2013	0.016*	0.011	-0.022***	-0.021***
	(0.009)	(0.009)	(0.007)	(0.007)
Eligible × cohort 2014	0.019**	0.016*	-0.028***	-0.027***
	(0.010)	(0.010)	(0.007)	(0.007)
Eligible × cohort 2015	0.013	0.010	-0.023***	-0.022***
	(0.009)	(0.009)	(0.006)	(0.006)
Control variables	No	Yes	No	Yes
Observations	1,527,798	1,527,797	1,527,798	1,527,797

Notes: Clustered standard errors at the class level in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$. School level control variables include indicators of financing institution, rural area and geographical region. Student level control variables include gender, attendance rate, *comuna* and number of family members at different leves in the education system.

Table 19: Dynamics of Two Year Enrollment

	University		Vocational	
	(1)	(2)	(3)	(4)
Eligible_2	0.248***	0.233***	0.024***	0.021***
	(0.006)	(0.007)	(0.003)	(0.003)
$\text{Eligible}_2 \times$ year 2008	0.005	0.005	-0.013***	-0.012***
	(0.009)	(0.009)	(0.005)	(0.005)
$\text{Eligible}_2 \times$ year 2009	0.003	0.003	-0.024***	-0.023***
	(0.009)	(0.009)	(0.005)	(0.005)
$\text{Eligible}_2 \times$ year 2010	0.001	0.001	-0.009*	-0.009**
	(0.009)	(0.009)	(0.005)	(0.005)
$\text{Eligible}_2 \times$ year 2012	-0.000	0.002	-0.007	-0.007
	(0.009)	(0.009)	(0.005)	(0.005)
$\text{Eligible}_2 \times$ year 2013	0.018**	0.017*	-0.019***	-0.019***
	(0.009)	(0.009)	(0.005)	(0.005)
$\text{Eligible}_2 \times$ year 2014	0.013	0.010	-0.016***	-0.016***
	(0.009)	(0.009)	(0.005)	(0.005)
$\text{Eligible}_2 \times$ year 2015	0.017*	0.015*	-0.023***	-0.023***
	(0.009)	(0.009)	(0.005)	(0.005)
Year 2008	0.003*	0.004	-0.017***	-0.020***
	(0.002)	(0.003)	(0.004)	(0.004)
Year 2009	0.003*	0.001	-0.001	-0.003
	(0.002)	(0.003)	(0.004)	(0.004)
Year 2010	-0.001	-0.002	0.000	0.001
	(0.002)	(0.003)	(0.004)	(0.004)
Year 2012	-0.001	-0.004	0.009**	0.009**
	(0.002)	(0.003)	(0.004)	(0.004)
Year 2013	-0.006***	-0.005*	0.031***	0.037***
	(0.002)	(0.003)	(0.004)	(0.004)
Year 2014	-0.006***	-0.008***	0.050***	0.052***
	(0.002)	(0.003)	(0.004)	(0.004)
Year 2015	-0.010***	-0.012***	0.059***	0.062***
	(0.002)	(0.003)	(0.004)	(0.004)
Control variables	No	Yes	No	Yes
Observations	1,347,837	1,347,837	1,347,837	1,347,837

Notes: Clustered standard errors at the class level in parentheses. *** p<0.01, ** p<0.05, * p<0.1. School level control variables include indicators of financing institution, rural area and geographical region. Student level control variables include gender, attendance rate and number of family members at different leves in the education system.

Table 20: Dynamics of Second Year Dropout

	University		Vocational	
	(1)	(2)	(3)	(4)
Eligible_2	-0.207***	-0.136***	-0.188***	-0.159***
	(0.011)	(0.011)	(0.007)	(0.007)
$\text{Eligible}_2 \times$ year 2008	-0.043***	-0.063***	0.011	0.005
	(0.016)	(0.016)	(0.011)	(0.010)
$\text{Eligible}_2 \times$ year 2009	-0.017	-0.027*	0.018*	0.011
	(0.016)	(0.016)	(0.011)	(0.010)
$\text{Eligible}_2 \times$ year 2010	-0.014	-0.017	0.020**	0.019*
	(0.016)	(0.016)	(0.010)	(0.010)
$\text{Eligible}_2 \times$ year 2012	-0.010	-0.031**	0.009	0.006
	(0.016)	(0.015)	(0.010)	(0.009)
$\text{Eligible}_2 \times$ year 2013	-0.070***	-0.080***	0.016*	0.012
	(0.016)	(0.016)	(0.010)	(0.009)
$\text{Eligible}_2 \times$ year 2014	-0.036**	-0.041**	0.012	0.004
	(0.016)	(0.016)	(0.010)	(0.009)
$\text{Eligible}_2 \times$ year 2015	-0.039**	-0.026	0.026***	0.019**
	(0.018)	(0.018)	(0.009)	(0.009)
Year 2008	0.059***	0.078***	0.034***	0.028***
	(0.016)	(0.016)	(0.011)	(0.010)
Year 2009	0.041**	0.057***	-0.018	-0.015
	(0.016)	(0.016)	(0.011)	(0.010)
Year 2010	0.017	0.023	-0.020*	-0.024**
	(0.017)	(0.016)	(0.010)	(0.010)
Year 2012	0.052***	0.082***	0.008	0.018*
	(0.016)	(0.015)	(0.010)	(0.010)
Year 2013	0.108***	0.111***	-0.015	-0.012
	(0.016)	(0.016)	(0.010)	(0.010)
Year 2014	0.078***	0.089***	-0.025**	0.002
	(0.017)	(0.016)	(0.010)	(0.009)
Year 2015	0.074***	0.064***	-0.034***	-0.007
	(0.018)	(0.018)	(0.010)	(0.009)
Control variables	No	Yes	No	Yes
Observations	386,329	375,297	273,715	272,737

Notes: Clustered standard errors at the class level in parentheses. *** p<0.01, ** p<0.05, * p<0.1. School level control variables include indicators of financing institution, rural area and geographical region. Student level control variables include gender, attendance rate and number of family members at different leves in the education system.

Appendix II.C

To provide further evidence in favor of the Parallel Trends assumption, Figure 10 shows the time evolution of our three outcomes for both eligible and non-eligible individuals. In turn, each outcome is depicted throughout its three variables: general, university, and vocational institutions. Panel A presents trends for Immediate enrollment, panel B for two-year enrollment and panel C for second-year dropout, respectively. From this visual inspection we can argue that all nine variables have evolved in a parallel way before the 2012 changes between eligible and ineligible individuals, providing evidence in favor of our identification assumption.

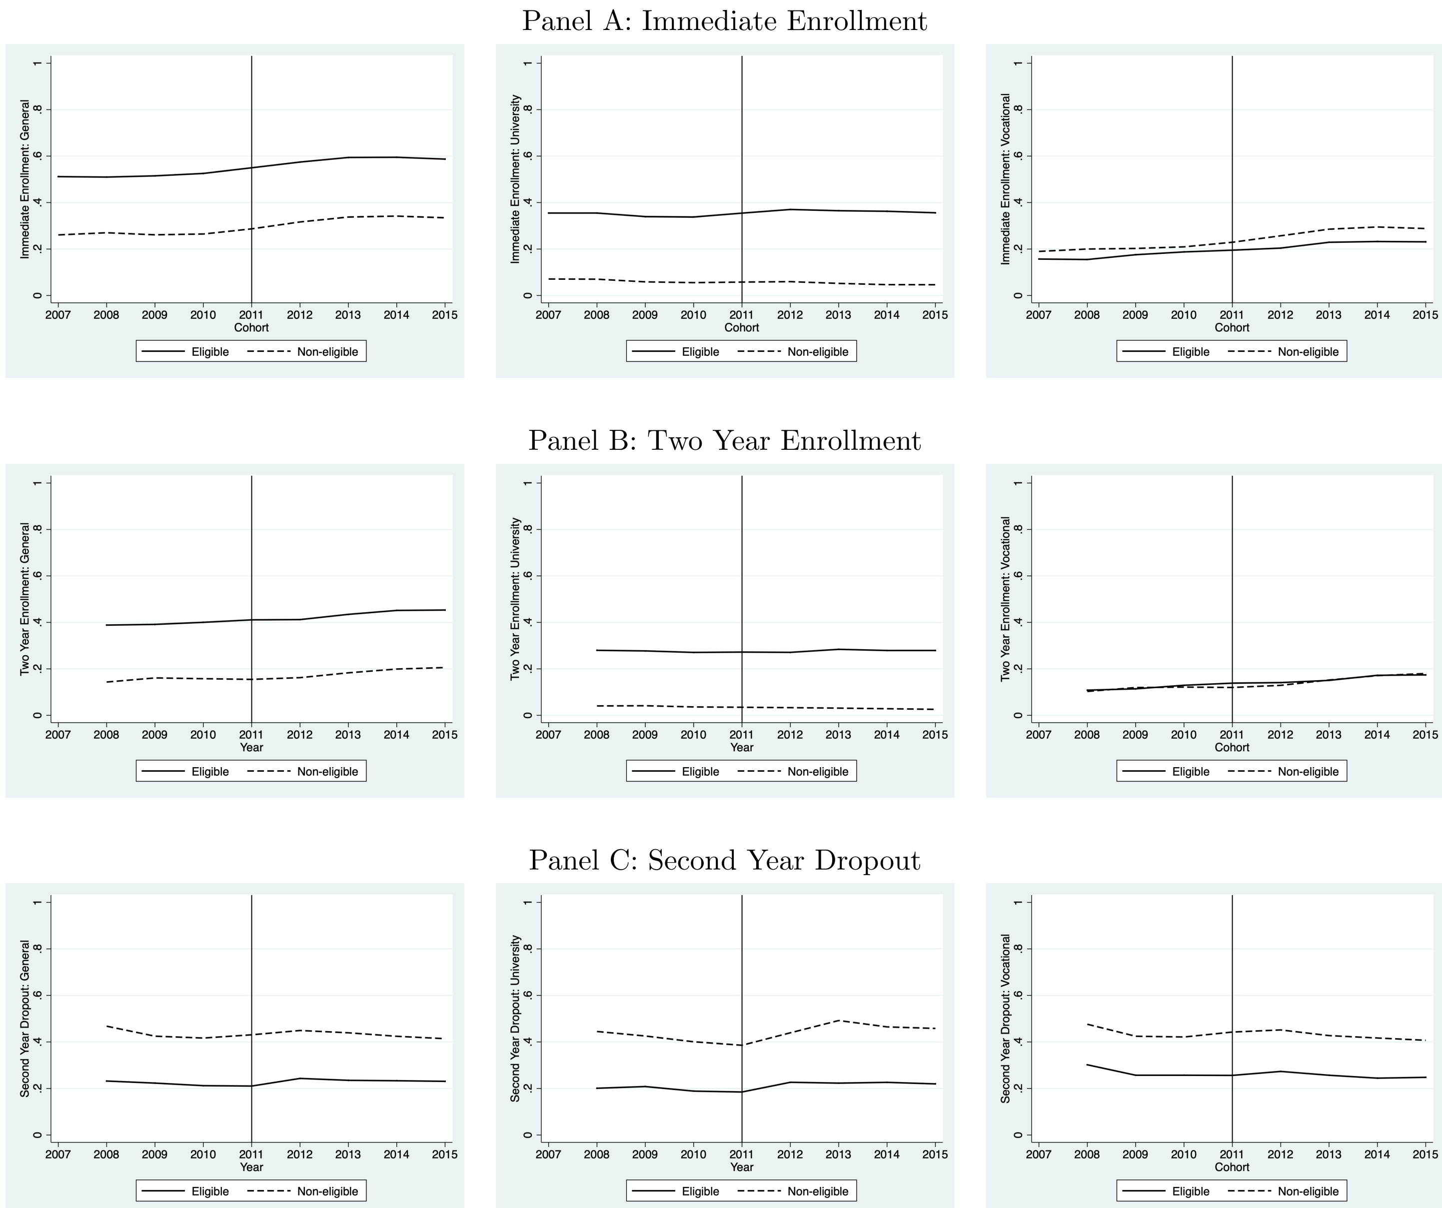

Figure 10: Outcomes over time by eligibility

Appendix II.D

To further elaborate the argument that women perform worse than men in PSU scores, in Table 21 we estimate the yearly differences of the proportions of male and female students scoring under 475 in the PSU. The proportion of women not meeteing the PSU cutoff for CAE eligibility is systematically higher (between 4 and 7 pp.) than the corresponding proportion of men. This evidence is consistent with our explanation for the negative effect of the reform on female overall immediate enrollment.

Table 21: Gender Difference in Proportion of PSU Scores under 475

	2007	2008	2009	2010	2011	2012	2013	2014	2015
Difference of proportions	-0.07***	-0.06***	-0.06***	-0.06***	-0.05***	-0.04***	-0.06***	-0.06***	-0.05***
	(0.003)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)

Notes: Difference of proportions of students scoring less than 475 in PSU between male and female students and the corresponding standard errors (in parentheses). *** p<0.01, ** p<0.05, * p<0.1.

Appendix II.E

Table 22 mimics the exercise in Table 21. The proportion of public school students scoring under 475 in the PSU is persistently higher than the corresponding proportion in voucher schools (between 14 and 17 p.p.).

Table 22: Voucher-Public School Difference in Proportion of PSU Scores under 475

	2007	2008	2009	2010	2011	2012	2013	2014	2015
Difference of proportions	-0.14***	-0.14***	-0.14***	-0.15***	-0.14***	-0.17***	-0.16***	-0.14***	-0.16***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)

Notes: Difference of proportions of students scoring less than 475 in PSU between students of Voucher and Public schools and the corresponding standard errors (in parentheses). *** p<0.01, ** p<0.05, * p<0.1.