

High-School Dropouts: Missing out or cutting losses?

Nano Barahona
(Stanford)

Christina Kent
(Stanford)

Christopher Neilson
(Princeton)

Sebastián Otero*
(Stanford)

November 15, 2017

Abstract

This paper examines the extent to which education is a valuable investment for academically marginal students. We ask this question in the context of the Dominican Republic, where schooling is mandatory until eighth grade and many students must decide between continuing on to high school or entering the labor force. For students struggling academically, the potential returns to education may appear small compared to the immediate wages earned in the labor market. We estimate the value of years of secondary education and high-school completion for low-performing students by exploiting the passing cutoffs on the eighth and twelfth grade national exams. We find that students who barely pass the exam in eighth grade complete more years of education, are more likely to enroll in college and experience higher earnings. Similar results are found for those who marginally pass the high-school exit exam. A broad implication of these facts is that education has large and positive returns even for the lowest part of the ability distribution.

1. INTRODUCTION

The United Nations has made universal primary and secondary education a priority¹, and many developing countries are experiencing fast growth in their access to secondary education. Many believe that secondary education can have large private and public benefits, leading to higher earnings in the labor market for individuals and productivity spillover effects of having a more educated workforce (Bils and Klenow, 2000). Nevertheless, the rapid growth in schooling throughout the world in the past 50 years did not translate into higher GDP growth for countries where education grew most rapidly. As a result, some may question whether or not investing in more education is the best option for students in settings where the quality of education is low and youth employment is common (Pritchett, 2001, 2013; Easterly, 2001). This can be particularly relevant for individuals in the lower part of the ability distribution.

*Emails: hbaraho@stanford.edu; cakent@stanford.edu; cneilson@princeton.edu; and soterio@stanford.edu. We are thankful to PACS for financial support.

¹See <http://www.un.org/sustainabledevelopment/education/>

Moreover, large dropout rates in developing countries may seem inconsistent with large return to secondary education. Despite efforts from local authorities, more than 27% of lower secondary school age children in the developing world are actively deciding to leave school (UNICEF Data, 2016). An alternative hypothesis is that returns to education are not low, but different market frictions make it hard for students to finish high school. A lack of information about labor market outcomes may lead some students to underestimate the value of high school education, resulting in higher than optimal dropout rates and inefficient investment in human capital (Jensen, 2010). Students might also be financially constrained, either because school is expensive and they can't afford it, or because they are expected to contribute to their family income by starting working at an early age (Duflo et al., 2017).

To reduce some of these frictions, in the Dominican Republic, the Ministry of Education has been aggressively expanding high school access throughout the country through new school construction. Schooling is mandatory until eighth grade, after which students decide whether or not to continue on to high school. In this paper we focus on ascertaining the private benefits of secondary education for students who are most likely to drop out.² The key problem of assessing the returns to education from observational data is that human capital investment is often correlated with unobservable student ability. We examine the human capital investment decision for low-performing students by exploiting the fact that students must pass standardized exams in eighth grade in order to continue on to high school. The exam creates a sharp passing cutoff that allows us to follow similar students on either side of the discontinuity. Students who fail the exam must retake it in the following year, which leads to an increase in dropout rates and less total years of total education attained among this group.

We are not the first to use exam cutoffs to estimate returns to education in a developing country. Ozier (2016) uses a high school entrance exam threshold in Kenya to compare students that barely passed the test to those that barely failed, finding a large impact in human capital accumulation and labor market conditions. Our setting has several advantages relative to Ozier (2016). First, Ozier (2016) uses survey data on a small sample of students, where both labor outcomes and scores in the test are self reported by the students and measured with error. We use large administrative data that allow us to get more precise estimates. Second, in the case of Kenya the cutoff score varies across schools and gender and are unknown by the researcher, which brings additional difficulties in assessing the validity of the regression discontinuity design. We have a sharp and well establish cutoff that gives us clean identification. Finally, students that barely failed the test in Kenya can attend to private or vocational schools that have different admission standards, violating the exclusion restriction of the research design. In the Dominican Republic, however, students need to pass the exit exam to attend any kind of secondary education.

Our results show that students who barely pass the exam are more likely to complete high school and less likely than their peers to be employed in the formal sector in the short run. In

²We plan to assess the public benefits of education in later work using data on social outcomes such as crime, fertility, and political participation

the long run, we find that they are more likely to attend college, be employed in the formal sector, and experience higher earnings in the formal sector. We attribute these outcomes to differences in educational attainment between the two groups, with those who pass the exam having on average 0.36 more years of education than their peers who fail, and being 8 percent more likely to complete high school. Our analysis cannot disentangle the distinct returns to years of schooling and the signaling value of a high school diploma, however we can leverage an additional cutoff provided by the exit exam in twelfth grade to try to quantify the size of the sheepskin effect. Evidence on those taking the twelfth grade exam suggests that a large component of the returns to education can be attributed to high school graduation. Future analysis will attempt to ascertain the magnitudes of these two effects.

One major caveat to our analysis is that we lack data on employment and earnings in the informal sector for students in our sample. Given that at least 50 percent of workers in the Dominican Republic work in the informal sector ([Central Bank of the Dominican Republic, 2017](#)), the lack of data on the informal sector has important implications for our analysis. For this paper, we assume that informal sector employment and earnings are zero for all students. However, in future work we plan to incorporate data on informal sector earnings that are reported in administrative data from the public health insurance system. With this data, we will be able to generate reasonable bounds for informal labor market earnings for students in our sample.

The rest of the paper proceeds as follows. Section 2 provides background information on education and the labor market in the Dominican Republic. Section 3 describes the various data sources used for our analysis. Section 4 outlines our empirical strategy. Section 5 presents results on the effect of the passing cutoff on educational outcomes in secondary school, and section 6 shows the effect on university enrollment. Section 7 shows results on the returns to education using labor market outcomes, and section 8 analyzes the signaling value of high school graduation. Section 9 attempts to elucidate the mechanisms driving the observed returns to education. Section 10 concludes.

2. SETTING

In this section we describe the education system and labor market in the Dominican Republic. Subsection 2.1 describes the education system in the Dominican Republic, including information on school types and standardized testing. Subsection 2.2 describes labor markets in the Dominican Republic and includes information on minimum wage policy and formal sector employment.

2.1. Education in the Dominican Republic: The public education system in the Dominican Republic consists of eight years of primary education and four years of secondary academic or vocational education. Public primary and secondary education in the Dominican Republic is free, however students can also pay tuition to attend private schools. Despite widespread access to education, however, students in the Dominican Republic continue to perform poorly compared to

international standards. Figure 1 shows the average score on the PISA science and math exams of students in all participating countries.




Figure 1: PISA Scores 2015

The Dominican Republic has historically used high stakes testing to measure student performance and promote national standards for education. Prior to the 2016 school year, students in eighth grade were required to take national standardized exit exams called Pruebas Nacionales (PN) in order to continue on to secondary education. Students take an exam for each core subject (Spanish, mathematics, social studies, and natural sciences) and must pass each exam in order to matriculate. Approximately 150,000 students take these exams each year. For students who fail one or more exams in the first round, they may retake any failed exam up to two times. Given that approximately 20 percent of students fail the exit exam for at least one subject in the first round, exam retakes are not uncommon.

To be promoted to the ninth grade, students must achieve a passing score based on the sum of their GPA and their score on the exit exam for each subject. GPA is based on student performance in class, ranges from 0 to 70 points, and is determined before students take the PN. The score on the PN ranges from 0 to 30 points. The sum of the GPA and PN score must be above a certain confidential cutoff, and students who do not achieve a passing score in all subjects must repeat the eighth grade and retake the exam in the following year to continue on to ninth grade.

For students who continue on to secondary school, they will take another set of exit exams in twelfth grade that are conducted in a similar manner to the eighth grade exit exams. Approximately 100,000 students take the twelfth grade exit exams each year, including students in private and

vocational secondary schools. Students who pass these exams then graduate from high school and can continue on to tertiary education.

Dropout is extremely prevalent in the Dominican Republic: of the students who take the 8th grade exit exam, approximately 34 percent will not complete high school.³ Some students who drop out then join the labor force, and in 2016 the Central Bank of the Dominican Republic estimated that approximately eleven percent of children between the ages of 10 and 19 were employed ([Central Bank of the Dominican Republic, 2017](#)).

Data from our sample of students who take the eighth grade exam show that many students also engage in at least part time work, with only 8.7 percent of girls and 12.3 percent of boys reporting no paid work at all in the week prior to the exam. It is also not uncommon for students to drop out and return to school at a later date: of the students who drop out each year in our sample, approximately 21 percent of these students will return to school within two years.

Grade repetition is also fairly common in the Dominican Republic. By the time they reach eighth grade, many students will have repeated a grade between second and eighth grade. However, grade repetition in earlier grades is not determined in a standardized way as it is in eighth grade thorough the exit exam. Thus the eighth grade has historically been the first time that grade repetition is mandated by a uniform, national, standardized process with strict cutoffs.

Beginning in the 2016 school year, the government of the Dominican Republic implemented changes in the high stakes testing policy. Students no longer take the eighth grade exit exam and instead take standardized exams in grades 3, 6, and 9 that are used purely for measurement purposes and do not affect grade promotion⁴.

2.2. Labor Markets in the Dominican Republic: By 2015 the Dominican Republic had an estimated population of 6,836,621 inhabitants between 15 and 65 years old. From that population, an estimated amount of 3,814,834 were self reported as having worked during that year ([Central Bank of the Dominican Republic, 2017](#)). According to the microdata from *Tesorería de la Seguridad Social*, on the other hand, 1,946,252 workers received a salary through the formal sector during that same year. Among them, 15% report wages below minimum wage, 9.5% exactly at the minimum wage level and 75.5% above it. In the Dominican Republic, minimum wage levels are firm specific and depend on the size and sector of the employer. They are also set at the month level, rather than by the hour.⁵

Short term returns to education may fluctuate based on macroeconomic trends that affect labor markets. After stagnant growth in 2009 following the global financial crisis, the Dominican Republic

³Here we define completing high school as taking the high school exit exam within six years after eighth grade.

⁴Exams are staggered so that only one grade is tested per year

⁵In Figure [A.2](#) we show a histogram of wages by May, 2015. Darker bars represent workers earning exactly the minimum wage. Observations below minimum wage may correspond to part time workers, however, we don't observe that in the data.

has experienced steady GDP growth each year.⁶ In fact, the Dominican Republic has experienced one of the highest GDP growth rates of all countries in Latin America and the Caribbean in the past 25 years.⁷ This has also been accompanied by a decrease in the unemployment rate over the past ten years even as the population continues to grow.⁸ It is important to take these macroeconomic factors into account when interpreting our results on the returns to education, however this is beyond the scope of this paper.

3. DATA

This section outlines the data used for this project and also includes information on the characteristics of our sample. The project relies on four main datasets. Subsection 3.1 describes the dataset on school-level enrollment. Subsection 3.2 describes the dataset on national exam scores from the Ministry of Education. Subsection 3.3 describes the national registry that is used to link education data with labor market outcomes. Subsection 3.4 describes the dataset on formal sector employment and earnings.

3.1. School Data: This project uses administrative school enrollment data from the Dominican Republic for all students from the 2009 to 2016 school years⁹. This administrative data is provided by schools and includes information on student demographic characteristics for all students in public, private, and vocational schools.

Our main outcomes include dropout, educational attainment, and learning for the students in our sample. Dropout is defined as a student not appearing in the administrative enrollment data in the year following eighth grade. We calculate this outcome for all students in our sample. Educational attainment is defined as the final year of education achieved within five or six years of taking the eighth grade exit exam. Because our sample includes students who took the eighth grade exam between 2010 and 2016, we calculate this outcome only for the 2010 and 2011 cohorts. If students do not repeat schooling years between eighth and twelfth grade, students in the 2010 cohort who pass the eighth grade exam should take the twelfth grade exam in 2014, and those who had to repeat a grade should take the twelfth grade exam in 2015. Similarly, students from the 2011 cohort should take the exam in 2015 and 2016, respectively. However, some students do continue to repeat after eighth grade, thus many students may wait an additional year before taking the twelfth grade exit exam.

3.2. Exam Data: We combine the enrollment data with data on student test scores for the eighth grade and twelfth grade PN exams for the entire universe of students who took these exams in the

⁶See <https://tradingeconomics.com/dominican-republic/gdp>

⁷See <http://www.worldbank.org/en/country/dominicanrepublic/overview>

⁸See <https://tradingeconomics.com/dominican-republic/unemployment-rate>

⁹Administrative school data at the student level does not exist prior to the 2009 school year

Dominican Republic from 2010 to 2016. This data was shared with us by *Educación de la Calidad Educativa* (from now on ECE), the organization in the Dominican Republic that is in charge of creating and administering the national exams. The data includes all exam takes for each student in each year and has identifying information allowing us to match exam scores to student enrollment. The data includes students' question-level responses, in addition to their final posted scores. In addition to the exam and enrollment data, we also have survey data for approximately 4,500 eighth grade students and 2,000 twelfth grade students who took their respective national exit exams in the year 2016. This data provides information on students' career and school performance expectations, in addition to information on their parents' education.

3.3. National Registry: We link the education and test data to data on employment and incomes using the national registry. In the Dominican Republic, most citizens have a national ID number, called a *cédula*. This number is similar to the social security number in the United States and is used for employment, voting, health and unemployment insurance, and in the national pension system. Most citizens get this number when they are eighteen,¹⁰ however some request it earlier for employment reasons, to obtain a driver's license, or to enter college. The national registry contains the names of individuals and their *cédulas*, allowing us to link the data on education to income and employment data.

3.4. Income Data: Data on formal sector employment outcomes comes from the Tesorería de la Seguridad Social, which is the Dominican Republic's equivalent to the social security system in the US. This data includes all formal sector employees in the Dominican Republic in 2015 and 2016, with over 41 million monthly income observations for over 2 million citizens employed at 76 thousand firms in the formal sector during this period. This data provides a snapshot of the formal labor market in the Dominican Republic and allows us to calculate statistics on wages by firm, firm size, and sector. We also have panel data on earnings from 2010 to 2017 for 45 thousand students whom we were able to match to our panel,¹¹ which we use to assess labor market outcomes. Since our panel focuses on students or those who recently completed high school, they tend to be younger and less experienced than the rest of the population.

3.5. College Data: The Dominican Republic has a long tradition of college education, hosting the first university in the New World with the foundation of the University Santo Toms de Aquino (current Autonomous University of Santo Domingo) in 1538. During the period of our study, the higher education system consisted of 36 universities and 14 vocational institutions. Having a high school diploma is required to enter to any higher education institution.

¹⁰For this reason, our test score administrative data does not have *cédulas* for each student, and instead uses names and constructed student IDs to identify students.

¹¹The matching process and potential selection concerns are addressed in the section [7.1](#)

Although the sector is regulated by the Ministry of Higher Education (MESCyT), they do not utilize a centralized data system. Instead, since 2012, universities are mandated to declare their students to MESCyT by filling separate spreadsheets with the names, birthdays, and degrees of their enrolled students. Most universities complete all of the required fields, but some refrain from reporting identifiable information in order to protect student’s confidentiality. The ministry, however, keeps close track of diplomas by mandating students to register their degrees in a MESCyT office in Santo Domingo. This dataset has been systematically maintained since 2006. As a result, we have a list of all students enrolled in the higher education system in a given year, however not all students are identifiable. In section 6.1, we describe the process used to match this college data to other datasets.

4. IDENTIFICATION STRATEGY

In this section we describe the regression discontinuity design used for our empirical strategy. Section 4.1 describes the design and why it is effective in this context. Section 4.2 describes our sample on either side of the discontinuity and tests for balance between the two groups.

4.1. Regression Discontinuity Design: To determine the effect of repetition on outcomes of interest, we take advantage of the fact that there is a sharp passing cutoff for the third retake of the exit exams in eighth grade. A common concern in using score cutoffs as a forcing variable is the issue of bunching at the cutoff, which would introduce bias into the regression discontinuity design. It is documented that passing cutoffs in educational settings are particularly susceptible to manipulation from teachers discretion (Diamond and Persson, 2016), outright cheating (Jacob and Levitt, 2003) or bias from local grading of standardize exams (Dee et al., 2016). Non of these threats are present in our setting. First, although passing cutoffs are publicly known before the test, because the exam grading is based on a Rasch model, it is not possible to perfectly map scores to the number of correct answers. Hence even in the presence of cheating it would be impossible to target the minimum passing score. Second, grading is done at the national level by the Ministry of Education and hence there cannot be any systematic bias due to test-based accountability pressure or student traits (e.g. manipulation helping less advanced students).

When looking at the score distribution we do indeed find significant evidence of bunching in reported scores next to the passing cutoff. Figure 2(a) below shows a histogram of final scores on the third exam. The formal cutoff of the exam is 65 points, nevertheless we see in the data that no student achieved a score higher than 60 and lower than 65 points. This is due to an adjustment formula that applies to scores near the cutoff.¹² To correct for this, we access raw test score responses and reconstruct the actual exam scores below. Due to the adjustment process, students are not aware of the actual passing cutoff, and thus we would not expect to see bunching around

¹²Each year, the Ministry of Education calculates a bandwidth around the passing cutoff and adds points to scores within the bandwidth but below the passing cutoff so that these students receive a passing score.

the true cutoff. Figure 2(b) shows the distribution of raw scores on the eighth grade exit exam. As expected, there is no evidence of bunching around the new cutoff of 61 points.




Figure 2: Number of students taking the exam for a third time by final grade

The cutoff determines who is promoted to the ninth grade and who is assigned to repeat the eighth grade. However, one may be concerned that the cutoff is not binding and some students could still be discretionally allowed to continue to the next grade. To test that we use administrative enrollment data and look at the situation of students the year after taking the test. For a student that scores above the passing cutoff, if in the next academic year she is either in ninth grade or has dropped out, we assign her a value of one. If not, we assign her a value of zero. For a student that scores below the passing cutoff, if she is either in eighth grade or dropped out from school, we assign her a zero. If not, we assign her a value of one. Figure 3 illustrates that the true cutoff is indeed binding. All students with a raw score below the cutoff either repeated or dropped out, while all students with a score above the cutoff either continued to the next grade or dropped out.




Figure 3: Probability of passing to the next grade

This gives us a clean setting to implement a sharp regression discontinuity design at the score cutoff that we will use to ascertain the effect of the grade repetition policy on schooling and labor market outcomes. To do this, we evaluate a regression discontinuity model as described by [Lee and Lemieux \(2010\)](#), where the main identification equation is given by

$$Y_i = \alpha + \tau D_i + f_i(X_i - c) + \varepsilon_i \quad (1)$$

where Y_i is our outcome of interest, D_i is a dummy that takes the value of one if student i is above the threshold and $f(X_i - c)$ is a flexible function controlling for the difference between the score of student i (X_i) and the cutoff point (c). We use a linear functional form for $f(X_i - c)$, allowing for different slopes at each side of the threshold, so that $f_i(X - c) = a(X - c) + D_i b(X - c)$.¹³ We include observations within a bandwidth of three points above or below the cutoff.¹⁴ Note that since our running variable is discrete in nature, we do not follow the existing literature in calculating the optimal bandwidth, as those methods are developed for assignment variables with density $g(x)$, where $g(\cdot)$ is continuous and bounded away from zero ([Calonico et al., 2014](#)). Instead, we present detailed results using two different bandwidths, and also show how our coefficient estimates change across six different bandwidth values.

4.2. Sample Characteristics: Because students can take the exit exams a total of three times, we focus on students around the passing cutoff on the third exam, since these scores are binding.¹⁵ This implies that our sample is local to the lowest-performing students in their cohort, since almost all of these students have failed the exam twice before.¹⁶ Of the approximately 1 million students who took the eighth grade PN between 2010 and 2016, 68 thousand participated in the third exam.¹⁷ Of these students, we focus on the 43,612 students who were within 3 points of the passing cutoff.¹⁸ Table 1 shows the difference in means between the two groups on either side of the passing cutoff for the 2010 through 2016 cohorts.¹⁹ We test for differences between means using the regression discontinuity design described in Section 4.

From this table we can see that most of the students in our sample are one to two years older than the expected age for their grade. Of the students who failed, 65 percent are male, 9 percent attend private schools, and 69 percent are from urban areas. Most of the students who failed were retaking the math exam (56 percent). We see similar means for those who passed. There are statistically significant differences in gender and GPA between those on either side of the passing

¹³Our results are also robust to using a quadratic form. The linear form results are shown in table 2, and the quadratic form results are presented in table A.2.

¹⁴We also present results using a bandwidth of four points.

¹⁵We also conducted analyses using the first and second exam cutoffs as instruments, however the prevalence of exam retakes renders these to be very weak instruments and yielded noisy results.

¹⁶It is possible to skip the first two exams and only take the third exam, however this is extremely uncommon.

¹⁷Approximately 94 percent of students who fail the second exam will go on to take the third exam (the second retake).

¹⁸We also include results for the 52,282 students within 4 points of the passing cutoff.

¹⁹In table A.1 we present statistics for the 2010 and 2011 cohorts only

Table 1: Sample Characteristics and Balance

	Mean		RD difference	Std Error	p-value	Obs
	Failed	Passed				
Age	15.45	15.27	-0.028	0.027	0.31	52282
Male	0.65	0.64	-0.016	0.008	0.06	52282
Private School	0.09	0.11	-0.001	0.005	0.85	52282
Urban	0.69	0.70	-0.007	0.008	0.36	52282
Extended Day	0.10	0.10	-0.005	0.005	0.37	52282
Math retake	0.57	0.43	-0.028	0.009	0.00	52282
Spanish retake	0.31	0.24	-0.014	0.008	0.07	52282
Science retake	0.45	0.34	0.005	0.009	0.57	52282
History retake	0.40	0.31	-0.005	0.008	0.54	52282
GPA Spanish	-0.30	-0.02	0.051	0.017	0.00	52282
GPA Math	-0.35	-0.01	0.023	0.016	0.17	52282
GPA History	-0.31	-0.04	0.034	0.017	0.04	52282
GPA Science	-0.34	-0.04	0.036	0.017	0.03	52282
PN Spanish	-0.41	0.00	0.007	0.017	0.67	52272
PN Math	-0.60	-0.02	-0.016	0.016	0.33	52269
PN Science	-0.52	-0.04	0.021	0.017	0.21	52275
PN History	-0.48	-0.03	0.005	0.017	0.78	52273
Mother's education	0.22	0.21	-0.026	0.030	0.38	3,364
Expected education	0.64	0.69	0.022	0.032	0.50	3,777
Outside wage	0.56	0.52	-0.051	0.035	0.14	3,700

Notes: This table shows the balance test results between groups on either side of the discontinuity. The balance test is run using a linear regression discontinuity design and robust standard errors. The first two columns present the mean values for those just below (failed) and those just above (passed) the passing cutoff. The third column reports the standard errors on the regression coefficients for each characteristic. The fourth column reports the p-value associated with each standard error. The fifth column shows the number of observations in each test. The final three characteristics were measured using voluntary survey data for students who took the eighth grade exam in 2016 and thus are drawn from a much smaller sample. The table shows that the groups are mostly balanced, with statistically significant but small differences in gender and GPA between groups. Figure A.1 shows 8th grade GDP in each subject plotted by score.

cutoff, however these are small in magnitude. We also include responses from survey questions included in the 2016 exam asking about student work perceptions and family characteristics. We see no significant differences in these responses between groups.

We can also compare the data on labor market outcomes for our students to that of the rest of the population. Students in our sample, conditional on being employed in the formal sector in October 2016, earn less on average than the rest of those employed in the formal sector. While in the population data the median income is DOM\$12800, among the students from our sample the

median income in 2016 is DOM\$8800. Also, while 75% of the population has wages higher than the minimum wage, in our sample of students only 55% do.²⁰ This is to be expected given the fact that our sample is relatively young compared to the working population and consists of individuals who are still in school or recent graduates.

5. REPETITION AND EDUCATION

In this section we examine the effect of repetition on educational attainment for students close to the passing cutoff. We present results from running the regression discontinuity design on three different outcomes of interest and show heterogeneity in effect size based on student age.

To analyze the effect of repetition on education, we run the regression discontinuity model on three different dependent variables: whether a student drops out one year after taking the eighth grade exit exam, the maximum grade achieved after six years, and whether or not students took the high school exit exam.²¹ Results from these models are shown in Table 2 for students within three or four points of the passing cutoff.²²

Table 2: Passing status in 8th grade and educational outcomes

	(1)	(2)	(3)	(4)	(5)	(6)
	Dropout	Dropout	Maximum Grade	Maximum Grade	Takes PN	Takes PN
Passed	-0.167*** (0.008)	-0.173*** (0.007)	0.455*** (0.066)	0.405*** (0.058)	0.076*** (0.021)	0.063*** (0.019)
Observations	43,612	52,282	7,815	10,074	9,342	11,927
R-squared	0.042	0.040	0.042	0.049	0.018	0.024
Mean	0.193	0.193	10.654	10.654	0.344	0.344
Bandwidth	3	4	3	4	3	4

Notes: This table shows the effect of failing the eighth grade exit exam on dropout, maximum grade achieved, and probability of taking the exit exam in 12th grade. The first two columns are calculated based on all students who took the eighth grade exit exam from 2010 to 2016. Columns 5-8 include only students who took the exam in 2010 or 2011 in order to have sufficient long-run results. We test two different bandwidths and find that the results are not sensitive to choice of bandwidth. From this table we can see that failing the eighth grade exit exam and being assigned to repeat the grade increases dropout by 16 percentage points from a mean of 19 percent, thus almost doubling dropout rates for students who just barely fail. Failure also decreases maximum grade achieved substantially. Failure has a smaller effect on the probability of taking the high school exit exam, reflecting the fact that many students who barely pass the eighth grade exam will still drop out from high school. *** p<0.01, ** p<0.05, * p<0.1.

²⁰See Figure A.2 for the full income distribution in both the whole population and our estimating sample.

²¹In table A.3 we find that the results are robust when we add controls for gender and GPA.

²²In figure A.3 we show that the regression discontinuity coefficient estimates are relatively stable for bandwidths of 1 to 5 points.

From the Table 2 we can see that students above the passing cutoff are 16.7 percentage points less likely to drop out of school in the following year than similar students who just barely fail. We can also look at the effects of failing the exam on final grade achieved. When we focus on the 2010 and 2011 cohorts, we see that students who pass achieve a final grade level that is 0.46 grade levels higher than similar students who failed. While the differences between the two groups are large for dropout rates in the following year and grade level achieved, the effect is smaller on the probability of taking the twelfth grade exit exam. Students who barely pass are 7.6 percentage points more likely to take the twelfth grade exit exam than their peers who fail. This is due to the fact that many students who just barely pass will also repeat grades in secondary school and drop out before taking the twelfth grade exit exam.²³ Below we display the regression discontinuity results in several figures.

Figure 4 shows the effect of failing the exam on dropout in the following year. We see that those who fail are significantly more likely to drop out in the following year, despite being academically similar to those just above the passing cutoff. This effect leads to a disparity in the level of education between those on either side of the passing cutoff.




Figure 4: Effect of repeating on probability of dropping out

Because many students who took the eighth grade exam in 2012 or later are still in school, we focus on the 2010 and 2011 cohorts to assess long run educational outcomes. Figure 4 shows the effect of failing the exam on years of schooling, maximum grade achieved, and probability of completing high school for these cohorts. We see that there is a large gap in years of schooling and grade level achieved between those who pass and those who fail, and a small but significant difference in the probability of completing high school between these groups.

While we do not find heterogeneity in effect size based on gender (see Figure A.4), the effect

²³Of the students who marginally pass, about sixty percent will drop out before taking the twelfth grade exam. Thus for many of these students, passing the exam delays dropout but does not guarantee high school completion. Whether or not this makes students better off depends on if returns to education accrue with each additional year, or if they are mostly due to the sheepskin effect of a high school diploma.




Figure 5: Effect of repeating on educational attainment (sample 2010-2011)

does change based on the age of the student when he or she takes the 8th grade exam. In figure 6, we group the students based on a categorical variable for age. Students who are “on time” are the age that one would expect for eighth grade students in the Dominican Republic (14 years old). Students older or younger than the expected age are labeled “Old” and “Young” respectively. The figures show that the effect of failing the exam is largest for students who are older than their peers, and smallest for those who are younger than their peers.




Figure 6: Heterogeneous effects of repeating on educational attainment by age

One plausible explanation for the treatment effect heterogeneity that we find in Figure 6 is related to the relationship between age and labor market opportunities. Older students may feel they have better chances of getting a job if they drop out from school. This higher opportunity cost can make them more elastic towards dropping out when faced with the prospect of repeating

a grade. We will come back to this in section 7, where we show that held back students do indeed earn more money than those who passed the eighth grade test during the first 3 to 4 years. Moreover, the differences are larger when we focus on older students.²⁴ Another possible explanation for the difference in effect by age is that older students have received more signals suggesting that they have low ability in school, thus giving them a more precise prior about their expected returns to education. Thus not only may they face a higher outside option to school, but they may also expect lower benefits to school than their peers.

We also find no evidence of students that repeated but were able to finish high school doing better in high school exit exams (see Appendix C).

6. REPETITION AND ACCESS TO COLLEGE

In this section we describe the matching process for college data in the Dominican Republic and assess the impact of repetition on college enrollment.

6.1. Matching College Data: We have a list of all students enrolled in the higher education system in a given year, however not all students are identifiable. We are currently coordinating with MESCyT to complete the share of information that is missing. In the meantime, we used our matching algorithm to pair matriculated students with identifiable information to individuals listed in the National Registry. In total, 686,175 of the 1,063,520 students enrolled in the 2012-2016 period were matched, corresponding to a 64.52%. Missing matches are found across all universities, which is consistent with reported information not always being accurate. Most of the missing matches, however, are concentrated in two universities that were particularly reluctant to share their data. We define the probability of matching students to the National Registry as α .

6.2. Assessing Effects on College Enrollment: Our outcome of interest is a dummy variable defined by whether or not the student is going to college. Note however, that because of the matching issues described above, this outcome will suffer from non-classical measurement error. In other words, only students attending to college may be misreported, hence attenuating the impact coefficient.

To formalize this intuition we use the potential outcomes framework where we assume that repeating (treatment status) is randomly assigned across students. Let $y_i \in \{0, 1\}$ denote the true outcome for individual i , and y_i^* denote the observed one. We can rewrite y_i^* as:

$$y_i^* = \begin{cases} y_i & \text{w/p } \alpha \\ 0 & \text{w/p } 1 - \alpha \end{cases}$$

²⁴See Table A.5 for details.

where $\alpha \in [0, 1]$. Define the true impact of repeating as:

$$\tau = \mathbb{E}[y_i(1) - y_i(0)] \quad (2)$$

where (i) indexes the treatment status. Instead, the impact from our biased sample will be:

$$\begin{aligned} \tau^* &= \mathbb{E}[y_i^*(1) - y_i^*(0)] \\ &= \alpha \mathbb{E}[y_i(1) - y_i(0)] \\ &= \alpha \tau \end{aligned}$$

Since we know the number of students enrolled and the share of matched individuals, we can directly calculate α and use it to correct the magnitude estimated coefficients. It is easy to show that $\mathbb{V}(\hat{\tau}^*) = \alpha^2 \mathbb{V}(\hat{\tau})$, which implies that the significance level of the t-statistics will be correct and the p-values will converge to the true coverage rates.²⁵

For the estimation, only students from the 2010 and 2011 cohorts are included. We do this to give time to held back students who mechanically take longer to finish high school. A visual inspection of the results are presented in Figure 7. As expected, the slopes of the curves are increasing with PN score, and there is a visible discontinuity in the repetition threshold. Students passing 8th grade are significantly more likely to enroll in college 5 years after. Note that the coefficient is heavily attenuated because of the unmatched students.




Figure 7: Effect of repeating on probability of going to college (sample 2010-2011)

We formalize these results in Table 3. Columns (1) and (3) present results using a bandwidth of 3 points, while columns (2) and (4) use 4. In columns (1) and (2) we present biased results given by $\hat{\tau}^*$. In columns (3) and (4) we correct our estimates by scaling them by $\frac{1}{\alpha}$ times. We find that the effect of repeating on college attendance is small but significant, with those who pass being between 3 and 5 percent more likely to attend college than those who fail. One may be surprised to see any effect at all on college enrollment, given that these students were some of the lowest-performing students in eighth grade. However, our results suggest that the effect of failing the eighth grade exam effects students at all levels of educational attainment, including tertiary education.

²⁵See Appendix B for a formal derivation.

Table 3: Treatment effect of being promoted in 8th grade on the probability of going to college (sample 2010-2011)

	(1)	(2)	(3)	(4)
Passed	0.033*** (0.011)	0.029** (0.009)	0.051*** (0.017)	0.045** (0.014)
Constant	0.041*** (0.008)	0.046*** (0.006)	0.064*** (0.012)	0.071*** (0.009)
Corrected	No	No	Yes	Yes
Bdw.	3	4	3	4
Obs.	9345	11930	9345	11930

Notes: In this table we present the results from the potential outcomes framework estimation of the effect of repeating on the probability of going to college. Columns 1 and 2 show results using a bandwidth of 3 exam points, while columns 3 and 4 show results using a bandwidth of 4 points. Columns 1 and 3 present biased results using the $\hat{\tau}^*$ estimator, and columns 2 and 4 present results corrected by scaling by $\frac{1}{\alpha}$.

7. REPETITION AND LABOR MARKET OUTCOMES

So far we have shown that students who were not held back in 8th grade end up attaining more years of education, having a higher probability of graduating from high school, and, conditional on making it to 12th grade, performing equally well on the exit exam as their peers who repeated. Now we focus our analysis on non-educational outcomes. To do so, we collected data from the National Registry of every person in the Dominican Republic that has a National ID (a.k.a. cédula²⁶) by May 2017. We used the National Registry data to match students from our panel to their respective ID numbers. Then, based on the ID, we matched the student panel to monthly earnings provided by the social security administration. In this section we start by describing the matching algorithm implemented to merge the databases, showing the main differences between the original sample and our smaller matched sample and replicating some previous results on the new sample. We then look at the effect of repetition on wages for the estimating sample and use it to infer the returns to education for treated and control students.

7.1. Matching: To match our educational information with the National Registry dataset, we mostly rely on names and birthdates. We used an iterative algorithm that uses string comparison methods to allow for typos in names and birthdates in the data. At first the algorithm is very strict and only matches students for which all characteristics are identical in both databases. Then, on later iterations the algorithm aims to match the remaining students even when some of their names

²⁶Cédula is an equivalent of the Social Security Number in the U.S. It is a unique national identifier that is needed to get any job in the formal sector.

differ in some character between the master and using database. 85 percent of out matched sample was matched in the first iteration, while more than 98 percent of the sample was matched in an iteration that forced two out of three names to be identical and the remaining one to have less than two characters of difference²⁷ in the two databases. We had an average matching success rate of 43 percent in our original sample, which was decreasing in time. In Table 4 we present the matching success rate for each year of data. As expected, earlier cohorts have a higher match rate, as the students taking the national exam in earlier years are older and therefore show up in the National Registry with higher probability.

Table 4: Percent of students matched by year

year	unmatched	matched
2010	12%	88%
2011	19%	81%
2012	44%	56%
2013	68%	32%
2014	85%	15%
2015	92%	8%
2016	100%	0%
Total	57%	43%

For the rest of the analysis, we focus on students who took the national exam of 9th grade between years 2010 and 2013. We do so for two reasons: i) matching success rates in that sample are higher and therefore they correspond to a more representative sample, and ii) since we are looking at outcomes that occur after completing one’s education, later years give us fewer observations of such outcomes for our analysis. In this section of the paper we also focus on *male* students only, as in the Dominican Republic they represent a larger fraction of the formal sector and their results are easier to interpret. Within that sample, our overall average matching success rate is 65.4%. Something we might be worried about before jumping to the analysis is whether there are significant differences in the matching rate on either side of the threshold. Even though the discontinuity is exogenous to the students and covariates are balanced, being on either side of the threshold could have a direct impact in the probability a student registers in the National Registry.

In Figure 8 we show matching success rates at both sides of the discontinuity. In Table A.4 we compare unmatched students to matched ones for different sub-samples. Matched students tend to be younger than unmatched students, more likely to come from private schools, less likely to drop out at the next year, but also less likely to be found taking the high school exit exam later in twelfth grade. Even though matching rates are correlated with test scores, we don’t find any significant jump around the relevant threshold. Once we assigned each student to an ID number, we were able to match them to our panel of income data described in section 3. For security reasons

²⁷measured by the Levenshtein Distance




Figure 8: Probability of matching to voter’s registry

and capacity constraints we were allowed only to extract data from the Social Security Department such that students could not be identified. Each student whose characteristics made him unique in the sample was dropped. In Table 5 we present some summary statistics of the three different samples and compare them. Panel A shows average levels of different covariates and outcomes for male students that received between 56 and 65 points in the 8th grade exam. We find that, compared to the original sample of students, in our final sample there are more students coming from public schools and they tend to be older. Since older students are more likely to be registered in the National Registry, older students will be overrepresented in our results.

In Panel B of Table 5 we estimate our main specifications from section 5 for the three samples.²⁸ That is, we estimate the effect of not repeating in dropping out from school the next year and the effect of not repeating in total years of schooling. We find that both effects are larger in magnitude in our final sample, consistent with students being older as shown in Figure 6.²⁹ It is important to keep in mind for the following results that everything will be estimated in this third sample, where students are older and more likely to react to repetition.

7.2. Results: We start by simply looking at the effect of repeating on yearly income. First, we sum all sources of income that each student had in a given year. With that, we were able to create a panel containing information on a student’s score in the 8th grade exam and how much he earned in the formal sector after x years of taking the exam.³⁰ In Figures 9(a) and 9(b) we present results looking at income 4 and 7 years after taking the exam respectively.

When we focus on Figure 9(a), we find that repeating students earn more when measured four years after taking the test. This is consistent with our previous results, where students who passed

²⁸Note that even in column (1) (full sample) results are expected to be different than in section 5, as here we focus on male students only and use exam takers from 2010 to 2013.

²⁹In Figures A.5 and A.6 we present visual version of Panel B from Table 5.

³⁰That is, for $x = 1$ we compute all earnings from the date of the exam until one year after, for $x = 2$ we compute all earnings from year one until year two after the exam was taken, and so on.

Table 5: Comparison of variables among samples

	Full (2010-2013) (1)	National Registry (2)	Social Security (3)
Panel A: Summary statistics			
Private school	0.11 (0.31)	0.09 (0.29)	0.072 (0.258)
Urban area	0.69 (0.46)	0.67 (0.47)	0.681 (0.466)
Fraction of older students	0.56 (0.5)	0.7 (0.46)	0.729 (0.444)
Fraction of younger students	0.17 (0.38)	0.08 (0.28)	0.069 (0.253)
Probability of dropping out	0.23 (0.42)	0.25 (0.44)	0.234 (0.423)
Total years of education	2.51 (1.48)	2.38 (1.5)	2.362 (1.527)
Number of Obs.	21051	13521	11726
Panel B: Treatment effects			
Probability of dropping out	-0.17*** (0.02)	-0.21*** (0.02)	-0.214*** (0.022)
Total years of education	0.44*** (0.05)	0.53*** (0.07)	0.53*** (0.075)
Number of Obs.	15178	9445	8280

Notes: This table show student characteristics in three different samples: i) full sample of men taking the test between 2010-2013, ii) students that were matched to the National Registry dataset, and iii) students matched to the social security data that were non identifiable. Panel A shows summary statistics of characteristics and Panel B shows treatment effects coming from our regression discontinuities.

the test in eighth grade are in school with higher probability, and therefore are not yet working in the formal sector. When we look at earnings seven years after the exam, there is a reversion such that students who passed the exam in the past are the ones with higher income. In figure A.7 we plot the regression discontinuity coefficients for these plots using several different bandwidths and find that they are relatively stable.

To test this previous result formally we run the following econometric specification

$$Y_{it} = \sum_{k=1}^{\tau} \tau_k D_i T_k + \sum_{k=1}^{\tau} f(X_i - c) T_k + \varepsilon_{it}$$

where Y_{it} is the income of student i during the t th year after taking the exam, D_i is a dummy variable that takes the value of 1 if student i passed the exam, T_k is a dummy that takes the value




Figure 9: Effect of repeating on yearly income

of one whenever $k = t$, Υ is the total number of years we observe in the income data after students take the test, $f(X_i - c)$ is a flexible function of score as defined in equation (1) and ε_{it} is a random shock at the student-period level. We cluster our errors at the student level to take into account for the autocorrelation of income across different years for a given student. Our coefficients of interest are given by τ_k and each of them represent the causal effect of passing the 8th grade exam on yearly income k years after taking it. In Figure 10 we plot such coefficients.




Figure 10: Effect of repeating on yearly income by year after taking exam

The first point that captures our attention when looking at Figure 10 is the trend in the coefficients over time. At the very beginning, the effect of repeating in total income is very close to zero. However, 3 years after taking the exam we see that students who originally failed the test are earning more than their counterparts who passed. This result shouldn't be surprising. Very short term returns to education are negative, as students that stayed at school are not working and have zero income during those years. Moreover, during the first years students who dropped out and got

into the labor market gained valuable experience which may help them to earn larger salaries as time goes on. Nevertheless, when we look at the coefficients 6 and 7 years after taking the exam the situation reverts and students who were not forced to repeat (and therefore stayed in school longer) show higher earnings than their counterparts.

There are a few important things to keep in mind when looking at these results. First, the panel is not balanced. While students taking the exam in 2010 can be followed up to 7 years, those taking the exam in 2013 can be followed only 4 years. That makes our estimates for later years less precise. Our results are robust to estimating the coefficients in a balanced panel using only exam takers from 2010 (see Figure A.8). Second, we are looking at unconditional income levels. If we decompose the effect in an extensive margin (how many months are students working in a year) and an intensive margin (conditional on working, how much are they earning), most of the effect is coming from the extensive margin (see Figure A.9).




Figure 11: Probability of being and school and average income by year after taking exam

In Figure 11 we plot the raw levels of school attendance and unconditional yearly income. During the first 4 years after taking the exam approved students are more likely to be found at school. Then, during the 5th year most of the graduated from high school already and the relationship reverts. After 6 years less than 10% of the students are still in school. This is consistent with the timing of Figure 10 and that can be confirmed with Figure 11(b), where we observe a steady increase in average income, which during the first years after taking the exam is higher for repeating students and where approved students catch up after 5 years. 6 and 7 years after taking the exam students that passed it earn substantially more.

7.3. Estimating returns to education: Figure 11(b) shows unconditional yearly earnings of students that barely passed or failed the eight grade exam. A straightforward question to ask is whether those students that passed ended up earning more than their failing classmates that

dropped out. To answer that question we compare net present value of aggregate income during the first 7 years between the two groups

$$\tilde{Y} = \sum_{t=1}^7 \frac{Y_t}{(1+r)^t}$$

where Y_t is the average unconditional income of each group³¹ and r is a discount rate that we assume to be 10%. We find an average net present income during the first seven years for students that barely passed the test of DOM\$52399, equivalent to US\$1434 in 2010. For those students that barely failed, on the other hand, the average net present income is of DOM\$46650, equivalent to US\$1277 in 2010. In other words, ignoring all future income differences after seven years, the monetary value of passing the eighth grade exam is of around US\$157, equivalent to a 12.3% increase in net present income.

8. RETURNS TO HIGH SCHOOL GRADUATION

In section 7, we showed evidence of positive returns to education for students close to the passing cutoff on the eighth grade exam. However, increased education involves many possible mechanisms that could result in higher earnings in the long run. Education not only purportedly improves human capital, but also involves signaling, where employers use the high school diploma as a signal of higher human capital among potential employees. It also opens doors to other educational levels, such as access to university. In this section, we provide evidence suggesting that some of the positive returns to education may be due to this latter diploma effect. The main difference between the discussion in this section and all previous ones, is that on twelfth grade exam takers there will be no effect of repeating on total years of secondary education.

We do this by taking advantage of the fact that students must also pass a national standardized exit exam in twelfth grade. The twelfth grade exam is organized in the same way as the eighth grade exam, with four subjects and two chances to retake exams in any failed subject. Students who fail the final retake of the twelfth grade exam must wait until the following year for another opportunity to receive their high school diploma. However, many students who fail the exam ultimately give up and do not receive their high school diploma. As in the eighth grade exam, the passing cutoff on the twelfth grade exam allows us to run a sharp regression discontinuity design to estimate the impact of failing the exam on outcomes, such as high school graduation, for students close to the cutoff. Table 6 shows the results of the balance test for the two samples within three points on either side of the passing cutoff.

We then run the regression discontinuity on the outcomes of graduating from high school within 4 years and attending college within 7 years for the 2010 through 2012 cohorts. Results from these regressions are reported in table 7.


In Figure 12(a) below we show the effect of failing the exam on the probability of having a high school diploma four years later.

³¹We present unconditional earnings in Table A.5.


Table 6: Sample Characteristics and Balance

	Mean		Std Error	p-value	Obs
	Failed	Passed			
Age	2.87	2.80	0.26	0.11	56139
Male	0.51	0.50	0.01	0.88	56139
Private School	0.10	0.13	0.01	0.57	56139
Urban	0.80	0.82	0.01	0.05	56139
Extended Day	0.06	0.06	0.00	0.20	56139
Math retake	0.73	0.61	0.01	0.64	56139
Spanish retake	0.56	0.45	0.01	0.15	56139
Science retake	0.65	0.58	0.01	0.93	56139
Soc. Stud. retake	0.55	0.42	0.01	0.40	56139

Notes: This table shows the balance test results between groups on either side of the discontinuity. The balance test is run using a linear regression discontinuity design and robust standard errors. The first two columns present the mean values for those just below (failed) and those just above (passed) the passing cutoff. The third column reports the standard errors on the regression coefficients for each characteristic. The fourth column reports the p-value associated with each standard error. The fifth column shows the number of observations in each test. The sample includes students from the 2010 through 2012 cohorts.



(a) After 4 years of taking the exam



(b) After 7 years of taking the exam

Figure 12: Effect of failing the 12th grade exam on educational attainment (sample 2010-2012)

Figure 12(a) suggests that those who barely fail, in addition to not graduating from high school in that same year, are also significantly less likely to ever graduate high school than their peers. This provides us with a first stage for ascertaining the effect of high school graduation on wages. In Figure 12(b) we also find that those students that passed the test in twelfth grade are more likely

Table 7: Regression discontinuity results for 2010 to 2012 cohorts

	(1)	(2)	(3)	(4)
	Pass after 4 years	Pass after 4 years	college	college
Passed	0.175*** (0.004)	0.182*** (0.004)	0.057*** (0.008)	0.057*** (0.007)
Observations	56,139	67,823	56,139	67,823
R-squared	0.175	0.188	0.011	0.013
Mean	1.000	1.000	0.311	0.311
Bandwidth	3	4	3	4

Notes: This table shows the regression discontinuity results on the effect of passing the 12th grade exam in a given year on various outcomes. Columns 1 and 2 show the results of the effect of passing the exam on the probability of graduating high school within 4 years. Columns 3 and 4 show the effect of passing the exam on the probability of attending college within seven years. Odd numbered columns include students within 3 points of the passing cutoff, while even numbered columns include students within 4 points of the passing cutoff. The row "mean" refers to the sample mean of the outcome for students just above the passing cutoff. From this table we can see that students who pass the exam are much more likely to graduate than students who fail, and are also more likely to attend college.

to attend to college during the following years.

We then match data on those who took the twelfth grade exit exam to data on employment and income for up to seven years later. Figure 13 below shows estimates for the difference in wages 7 years after taking the exam for male students³² close to the passing cutoff for the 2010 cohort.




Figure 13: Effect of failing the 12th grade exam on earnings (sample 2010)

We can also pool all cohorts and run the regression discontinuity specification for income ob-

³²Effects for female students are similar but smaller in magnitude.

served in each year after the exam. In Figure 14 we plot the regression coefficients for the effect of passing in year zero on income earned in subsequent years for students in the 2010 through 2016 cohorts. From this figure we see that passing the exam has a positive effect on earnings that increases in later years. However, as in the case of the eighth grade analysis, the effect on income seems to be attributed to more months spent employed rather than higher wages conditional on employment. Figure A.10 shows the discontinuity estimates plotted for the intensive and extensive margins.




Figure 14: Effect of repeating on yearly income by year after taking the 12th grade exam (sample 2010-2012)

While these results are very preliminary, they suggest that a large component of the returns to education may be due to a direct effect of having a high school diploma rather than the human capital or incapacitation effects of school. How to interpret the value of a diploma is however not straightforward. On one hand, it can be used as a signaling device, allowing to better positions in the later labor market. On the other hand, it can be seen as a milestone in the educational path that opens doors to other levels of educational attainment, such as access to college. Here we don't intend to disentangle these two different mechanisms, but rather understand the value of a diploma as a bundle that has value due to the possibilities it opens.

9. DOES EDUCATION LOOK MORE LIKE A MARATHON OR THE ATP TENNIS WORLD TOUR?

In a Marathon, competitors are rewarded based on the time they take to make it to the finish line. Whatever happens before that, however, does not matter for the final result. Therefore, the marginal return of an extra mile is zero is the runner if not able to finish the race. In other competitions, such as the ATP Tennis World Tour, each tournament you play is important by itself, as it allows you to increase your overall points and improve in your ranking. Even if a player gets injured and is not able to play for the rest of the season, any point collected in previous games can be use to improve her position by the end of the year. Whether education looks more like a Marathon or the ATP tournament is not clear and has different consequences for policy. Do we

want to incentivize students to stay in school one extra year, knowing that many are unlikely to complete high school? Or should we focus our efforts on increasing the number of students who ultimately graduate and get a diploma?

In this section we develop an empirical framework that allows us to separate the effects of education from the two mechanisms exposed above: on one hand, more education can affect future earning through more years spent in school. Students have the opportunity to learn more content and spend time in a safe and stimulating environment. On the other hand, it could be that secondary education is important only as long as students are able to ultimately get a diploma. To answer that question we use the exams of eighth and twelfth grade as instruments for the two margins of education. While the twelfth grade exam only increase education through the value of having a high school diploma, the eighth grade exams increases both the probability of graduating from high school and total years spent in school.

We focus on yearly earning two years after the expected year of graduation. That is, for eighth grade exam takers we look at income 6 years after taking the exam, while for twelfth grade exam takers we look at earning 2 years after they took the exam. Let's assume that earnings are given by the following structural equation

$$y_{ie} = \alpha S_{ie} + \beta D_{ie} + f_e(X_{ie}) + \varepsilon_{ie} \quad (3)$$

where y_{ie} are unconditional yearly earnings of student i taking exam $e = \{8, 12\}$ 6 and 2 years after taking the exam, S_{ie} is the maximum grade achieved by student i (note that $S_{ie} = 12$ for $e = 12$) and D_{ie} is a dummy that takes the value of 1 if student i got a high school diploma. We also control for a flexible function of score X_{ie} that can be different for eighth and twelfth grade exam takers. The function includes a constant that also differs between students with $e = 8$ and $e = 12$. It is important to add such flexibility into the model, as marginal students in the eighth and twelfth grade exam are very different from each other, and therefore we want to allow their income to be different, even if S_{ie} and D_{ie} are different. In our estimations we assume $f_e(X_{ie}) = \sum_{\bar{e} \in \{8, 12\}} (\delta_e^0 \mathbb{1}[e = \bar{e}] + \delta_e^1 X_{ie} \mathbb{1}[e = \bar{e}] + \delta_e^3 X_{ie} \mathbb{1}[e = \bar{e}] \mathbb{1}[X_{ie} > c_e])$, where c_e is the cutoff of exam e .

Because S_{ie} and D_{ie} might be correlated with ε_{ie} , we need instruments for them. We construct two instruments z_8 and z_{12} , that correspond to dummies for whether student i passed the exam in eighth grade and twelfth grade respectively. Note that passing the test can only be used as a source of exogenous variation within a bandwidth close to the threshold. Therefore, we can't use information about whether a student passed the exam $e = 12$ if the students was close to the threshold only in $e = 8$. To deal with that issue, we assign a constant value to the instrument $z_{\bar{e}} = \kappa$ for all students for which $e \neq \bar{e}$. The value of κ does not affect our estimates, as it is absorbed by the constant δ_e^0 of our flexible functional form.

In our identification strategy, we need to assume that the parameter β is equal among eighth and twelfth grade takers. In other words, even though we are allowing for students taking the two

different exams to differ in their income levels due to different unobservable characteristics with our flexible polynomial, we are assuming that the treatment effect of getting a high school diploma does not differ between the two groups. This assumption is crucial, as the estimated coefficient $\hat{\beta}$ will be a weighted average of the true value of β from eight grade exam takers compliers and twelfth grade exam takers compliers. Then, α will be consistently estimated for our eight grade exam takers compliers only if our value of $\hat{\beta}$ is consistent for that group as well. This will only be true if the true parameter β is equal among compliers from the two exams.

Table 8: Different mechanisms

	(1)	(2)	(3)
Panel A: Second Stage			
Grade achieved	6,857.5 (5,993.4)	2,964.3 (4,473.4)	1,791.8 (3,970.0)
Diploma	20,761.0** (9,759.6)	18,558.1** (8,274.2)	20,140.0*** (7,663.3)
Panel B: First stage 1 (maximum grade achieved)			
Passing 8th grade	0.561*** (0.107)	0.618*** (0.094)	0.671*** (0.090)
Passing 12th grade	0 (-)	0 (-)	0 (-)
Panel C: First stage 2 (high school diploma)			
Passing 8th grade	0.052* (0.030)	0.074*** (0.026)	0.090*** (0.025)
Passing 12th grade	0.137*** (0.010)	0.139*** (0.009)	0.146*** (0.008)
Bdw.	3	4	5
Obs.	21499	26556	28148

Notes: Here we show the results of estimating equation (3). To run the regressions we pooled data from eight and twelfth grade exam takers and construct the variables as described in the text.

In Table 8 we present results for different selections of bandwidth. We observe large and very significant effects for the value of the diploma, while smaller and not statistically different from zero effects for each extra year of school.

10. CONCLUSIONS

Is high school a good investment for all students? Secondary education may generate both private benefits and positive externalities, leading to higher earnings in the labor market for individuals

and productivity spillover effects from having a more educated workforce. These benefits, however, may depend on the quality of the education system in a particular context, in addition to students' innate abilities.

In this paper we focus on estimating the returns to secondary education for students in the Dominican Republic who are the most likely to drop out. We use the fact that students must pass standardized exams in eighth grade in order to continue on to high school to identify this effect. We find that students who fail the exam are more likely than their peers to drop out in the following year, and are less likely to complete high school. While students who pass initially earn less than their peers in the labor market, over time they begin to earn more than their peers due to higher rates of employment. We find evidence suggesting that while some of this effect may be due to increased time in school and actual human capital accumulation, a large component of this effect may be due to the signaling value of a high school diploma.

The short time horizon of our data prevents us from observing student outcomes at the peak of their careers, however the pattern observed in the data suggests that students with more education will continue to outperform their peers. Also, the nature of our empirical strategy implies that our results are local to the lowest-performing students in the Dominican Republic. We find this to be an encouraging result for proponents of universal secondary education. If we assume that returns to education are increasing with student ability, then we would expect the results of our study to be even more positive for higher-performing students. Ultimately our work suggests that while struggling students may believe that they are cutting their losses by dropping out of school early, they may in the long run be selling themselves short.

10.1. Next steps: In this subsection we describe possible extensions to the current version that we are currently working on and will help us to better understand what are the returns to education for the marginal students in the Dominican Republic, in addition to the mechanisms through which education affects income and other important outcomes. Here we focus on three complementary next steps.

10.1.1. Getting earlier data on scores: We are currently working on an agreement with *ECE* to get exam scores data from 2006 to 2009. In 2010 there was a change in the administration of *ECE* that made data collection from 2010 onwards easier and more direct. The main drawback of the older data is that there is no available school enrollment data for those earlier years. Therefore, it is not possible to estimate a first stage of the effect of repeating in total years of education for that sample. The main and most obvious advantage, however, is that it will allow us to look at students that finished school longer ago. That will allow us to improve our estimates on labor outcomes, as those student have been in the labour market for almost 7 years already. We hope this new access to data will allow us to improve our estimation in section 7 by increasing the precision of the parameters.

We are also working with the government of the Dominican Republic to gain access to data from the health insurance program, which includes data on self-reported earnings in the informal sector. This will help us to put reasonable bounds on the informal labor market earnings of students in our sample who are not in school or working in the formal sector.

10.1.2. Look at more outcomes: Even though looking at earnings is a useful benchmark to understand returns to education, there are other dimensions through which education can improve someone's longer term utility. We are currently working on getting data on different outcomes that we want to make part of our analysis. Outcomes we hope to obtain include: i) voting behavior, ii) fertility rates of women, iii) crime data, and iv) health data.


References

- Angrist, J., E. Bettinger, and M. Kremer (2006). Long-term educational consequences of secondary school vouchers: Evidence from administrative records in colombia. *The American Economic Review* 96(3), 847–862.
- Bils, M. and P. J. Klenow (2000). Does schooling cause growth? *American economic review*, 1160–1183.
- Calonico, S., M. D. Cattaneo, and R. Titiunik (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica* 82(6), 2295–2326.
- Central Bank of the Dominican Republic (2017). Encuesta nacional de fuerza de trabajo. Technical report, Estadísticas Económicas.
- Dee, T. S., W. Dobbie, B. A. Jacob, and J. Rockoff (2016). The causes and consequences of test score manipulation: Evidence from the new york regents examinations. Technical report, National Bureau of Economic Research.
- Diamond, R. and P. Persson (2016). The long-term consequences of teacher discretion in grading of high-stakes tests. Technical report, National Bureau of Economic Research.
- Duflo, E., P. Dupas, and M. Kremer (2017). The impact of free secondary education: Experimental evidence from ghana. Technical report, Working Paper.
- Easterly, W. (2001). *The elusive quest for growth: economists’ adventures and misadventures in the tropics*. MIT press.
- Jacob, B. A. and L. Lefgren (2004). Remedial education and student achievement: A regression-discontinuity analysis. *The review of Economics and Statistics* 86(1), 226–244.
- Jacob, B. A. and S. D. Levitt (2003). Rotten apples: An investigation of the prevalence and predictors of teacher cheating. *The Quarterly Journal of Economics* 118(3), 843–877.
- Jensen, R. (2010). The (perceived) returns to education and the demand for schooling. *The Quarterly Journal of Economics* 125(2), 515–548.
- Lee, D. S. and T. Lemieux (2010). Regression discontinuity designs in economics. *Journal of economic literature* 48(2), 281–355.
- Ozier, O. (2016). The impact of secondary schooling in kenya: A regression discontinuity analysis. *Journal of Human Resources*.
- Pritchett, L. (2001). Where has all the education gone? *The world bank economic review* 15(3), 367–391.


Pritchett, L. (2013). *The rebirth of education: Schooling ain't learning*. CGD Books.

UNICEF Data (2016). Monitoring the situation of children and women. Technical report, UNICEF, New York.


Appendix A: Figures and Tables




(a) Math GPA by Score



(b) Science GPA by Score



(c) History GPA by Score



(d) Spanish GPA by Score

Figure A.1: 8th Grade GPAs by Score




Figure A.2: Wage distribution




Figure A.3: RD Coefficient Estimates by Bandwidth

Table A.1: Sample Characteristics and Balance for 2010 and 2011 cohorts

	Failed	Passed			
	Mean	Mean	Std Error	p-value	Obs
Age	15.48	15.37	0.07	0.62	9,342
Male	0.68	0.65	0.02	0.02	9,342
Private School	0.08	0.09	0.01	0.59	9,342
Urban	0.70	0.69	0.02	0.92	9,342
Extended Day	0.07	0.06	0.01	0.68	9,342
Math retake	0.36	0.32	0.02	0.73	9,342
Spanish retake	0.52	0.45	0.02	0.22	9,342
Science retake	0.39	0.33	0.02	0.96	9,342
History retake	0.47	0.38	0.02	0.74	9,342
GPA Spanish	-0.46	-0.19	0.04	0.62	9,342
GPA Math	-0.18	0.02	0.05	0.63	9,342
GPA History	-0.37	-0.11	0.04	0.08	9,342
GPA Science	-0.33	-0.07	0.04	0.06	9,342
PN Spanish	-0.80	-0.33	0.04	0.87	9,341
PN Math	0.02	0.29	0.05	0.55	9,339
PN Science	-0.22	0.14	0.05	0.11	9,339
PN History	-0.47	-0.03	0.05	0.95	9,339

Notes: This table shows the balance test results between groups on either side of the discontinuity for students from the 2010 and 2011 cohorts. The balance test is run using a linear regression discontinuity design and robust standard errors. The table shows that the groups are mostly balanced, with statistically significant but small differences in gender and GPA between groups.

Table A.2: Passing status in 8th grade and educational outcomes, with a quadratic functional form


	(1)	(2)	(3)	(4)	(5)	(6)
	Dropout	Dropout	Maximum Grade	Maximum Grade	Takes PN	Takes PN
Passed	-0.141*** (0.013)	-0.148*** (0.011)	0.478*** (0.090)	0.474*** (0.077)	0.086*** (0.030)	0.085*** (0.026)
Observations	43,612	52,282	7,815	10,074	9,342	11,927
R-squared	0.042	0.041	0.042	0.049	0.018	0.024
Mean	0.193	0.193	10.654	10.654	0.344	0.344
Bandwidth	3	4	3	4	3	4

Notes: This table shows the effect of failing the eighth grade exit exam on dropout, yeas of education, maximum grade achieved, and probability of taking the exit exam in 12th grade. The RD specification is the same as the one used in table 2, except in this table we allow for $f(X_i - c)$, the function controlling for the difference between a student's score and the cutoff point, to have a quadratic functional form. We find that the results are not qualitatively different from those in table 2. The first two columns are calculated based on all students who took the eighth grade exit exam from 2010 to 2016. Columns 5-8 include only students who took the exam in 2010 or 2011 in order to have sufficient long-run results. We test two different bandwidths and find that the results are not sensitive to choice of bandwidth. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.


Table A.3: Passing status in 8th grade and educational outcomes, with controls for gender and GPA

	(1)	(2)	(3)	(4)	(5)	(6)
	Dropout	Dropout	Maximum Grade	Maximum Grade	Takes PN	Takes PN
Passed	-0.166*** (0.008)	-0.173*** (0.007)	0.440*** (0.066)	0.391*** (0.057)	0.071*** (0.021)	0.058*** (0.019)
Observations	43,612	52,282	7,815	10,074	9,342	11,927
R-squared	0.044	0.042	0.054	0.060	0.027	0.033
Mean	0.193	0.193	10.654	10.654	0.344	0.344
Bandwidth	3	4	3	4	3	4

Notes: This table shows the effect of failing the eighth grade exit exam on dropout, yeas of education, maximum grade achieved, and probability of taking the exit exam in 12th grade. The RD specification is the same as the one used in table 2, except in this table we control for gender and student GPA in each subject. We see that the coefficient estimates are not qualitatively different from those in 2. The first two columns are calculated based on all students who took the eighth grade exit exam from 2010 to 2016. Columns 5-8 include only students who took the exam in 2010 or 2011 in order to have sufficient long-run results. We test two different bandwidths and find that the results are not sensitive to choice of bandwidth. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.




(a) Probability of dropping out




(b) Total years of schooling


Figure A.4: Heterogeneous effects of repeating on educational attainment by gender



(a) Full sample




(b) Matched to voter's registry




(c) Non uniquely identifiable


Figure A.5: Effect of repeating on drop out in different sub samples



(a) Full sample




(b) Matched to voter's registry




(c) Non uniquely identifiable

Figure A.6: Effect of repeating on total years of school in different sub samples



(a) 4 Years



(b) 7 Years

Figure A.7: RD Coefficient Estimates by Bandwidth

Table A.4: Characteristics of matched and unmatched students

	Full sample		Only male		Male close to threshold	
	Unmatched	Matched	Unmatched	Matched	Unmatched	Matched
Age	15.65	14.59	15.76	14.68	15.83	14.67
Male	0.65	0.60	1.00	1.00	1.00	1.00
Private School	0.10	0.13	0.10	0.14	0.09	0.13
Urban	0.70	0.71	0.68	0.71	0.67	0.71
Extended Day	0.08	0.07	0.08	0.07	0.08	0.07
PN Spanish	-0.07	0.13	-0.18	-0.03	-0.34	-0.10
PN Math	0.27	-0.06	0.30	-0.01	0.13	-0.11
PN Science	0.21	-0.06	0.17	-0.10	-0.00	-0.20
PN History	0.11	0.01	0.07	-0.06	-0.10	-0.14
Dropout	0.24	0.18	0.25	0.18	0.25	0.18
Years of Educ.	2.52	2.77	2.46	2.71	2.38	2.73
Maximum Grade	10.59	10.66	10.51	10.57	10.40	10.55
Takes PN	0.32	0.14	0.29	0.13	0.25	0.11
Observations	14066	26563	8425	17315	7530	13521

Notes: This Tables shows summary statistics for the sample of matched and unmatched students to the National Registry as explained in section 7




Figure A.8: Effect of repeating on yearly income by year after taking exam (sample 2010)




Figure A.9: Effect of repeating on labor outcome in the extensive and intensive margin by year after taking the 8th grade exam (2010-2013, men only)

Table A.5: Yearly earning for passing and repeating students after taking the eight grade exam (\$DOM)

years after exam	1	2	3	4	5	6	7
Failed	209	1728	6146	11976	16263	18061	23271
Passed	1020	2366	4948	9285	16134	23495	31523

Notes: notes




Figure A.10: Effect of failing on labor outcome in the extensive and intensive margin by year after taking the 12th grade exam (2010-2012, men only)

Appendix B: Derivation of the variance of the estimator of (2)

The variance of the estimate using an unbiased sample would be given by:

$$\mathbb{V}(\hat{\tau}) = \mathbb{E}[(\bar{y}_{it} - \bar{y}_{ic} - \mathbb{E}[\bar{y}_{it} - \bar{y}_{ic}])^2]$$

We can show that:

$$\begin{aligned}\mathbb{V}(\hat{\tau}^*) &= \mathbb{E}[(\bar{y}_{it}^* - \bar{y}_{ic}^* - \mathbb{E}[\bar{y}_{it}^* - \bar{y}_{ic}^*])^2] \\ &= \mathbb{E}[(\alpha\bar{y}_{it} - \alpha\bar{y}_{ic} - \mathbb{E}[\alpha\bar{y}_{it} - \alpha\bar{y}_{ic}])^2] \\ &= \alpha^2\mathbb{V}(\hat{\tau})\end{aligned}$$

Appendix C: Repetition and Education Quality

We have shown above that grade repetition decreases educational attainment for students on average, but it is possible that students who do not drop out end up with better learning outcomes from repeating a year. Proponents of grade repetition claim that it reinforces weak concepts that students failed to learn in the previous year. [Jacob and Lefgren \(2004\)](#) found some evidence that grade repetition improves learning for students in primary school, however repetition seemed to have no effect on student learning in secondary school. It also may improve tracking, as held back students are placed in classes of students with similar skill levels and thus teachers are better able to tailor lessons to these students.

Here we analyze whether or not the students in our sample who repeat improve their learning outcomes by looking at their performance on the 12th grade exit exam. However, because students who are assigned to repeat are also more likely to drop out than their peers, this leads to a selection problem when comparing students using the 12th grade exam score. Figure [C.1\(a\)](#) illustrates the difference in the probability of taking the 12th grade exam for students on either side of the discontinuity for the 2010 and 2011 cohorts. Students that are held back are 7 percentage points less likely to take the test in twelfth grade. This will introduce source of bias when looking at causal effects of repetition on twelfth grade scores. Let's think about a very simple model in which unobserved ability is equally distributed at the two sides of the threshold and where low ability students are more likely to drop out upon a negative individual specific shock. That is, higher ability students only drop out if the shock is large enough, while low ability students are likely to drop out with any kind of shock. If we assume repeating increases the magnitude of the shock, we would expect students on the left side of the discontinuity to be higher ability than their counterparts on the right side. Therefore, even if repetition has no causal effect on test scores in twelfth grade, we could find an effect using our the regression discontinuity design only due to selection.

Figure [C.1\(b\)](#) shows our main result of the effect of repetition on twelfth grade scores. Despite the presence of the bias mentioned above, we do not find evidence of held back students doing significantly better later on. This means that, among students that were in the margin of failing the test, those who repeated the eighth grade did not learn more during ninth to twelfth grade than their peers than passed the eighth grade exam. This does not mean, however, that learning during these four years was low for the two groups.

Our estimates from Figure [C.1\(b\)](#) are likely to be biased. To see the consequences of differential test-taking rates at each side of the threshold we follow [Angrist et al. \(2006\)](#) and use a strategy where we fit a parametric model to artificially completed score data by censoring observed scores above a particular quantile, with all students bellow the cutoff and non takers assigned the censoring point. The main drawback of this approach is that we will need to assume normality for the uncensored latent score distribution.

Let's assume that the causal effect of repeating the test can be estimated by regressing latent




Figure C.1: Effect of repeating on academic achievement (sample 2010-2011)

scores y_i on D_i and linear controls on score X_i .

$$y_i = \tau D_i + \beta X_i + \eta_i$$

where η_i is a normally distributed error. Our artificially censored dependent variable is given by

$$Y_i(t) = \mathbb{1}[y_i \geq t \text{ and } T_i = 1]y_i + \mathbb{1}[y_i < t \text{ or } T_i = 0]t$$

where T_i is a dummy indicating whether the student took the test in twelfth grade and t is the cutoff score at which we censor the data. Assuming any untested student would have scored below the threshold t if she would have taken the test we can consistently estimate τ by applying Tobit to $Y_i(t)$. We use the percentiles 0.1, 0.05 and 0.01 for the selection of t . Results are presented in Table C.1.

Table C.1: Effect of repetition on academic achievement adjusting for selection

	OLS	Tobit (10%)	Tobit (5%)	Tobit (1%)
	(1)	(2)	(3)	(4)
Passed	-0.080 (0.072)	0.044 (0.027)	0.067** (0.032)	0.426*** (0.139)
Observations	3008	9345	9345	9345

Notes: notes

Appendix D: The Matching Algorithm

Table D.1: Matching rates per year

Year	Number of students	Sharp match	Total match
2006	145,911	77.8%	92.2%
2007	148,855	80.9%	92.7%
2008	153,689	80.3%	91.4%
2009	162,210	80.2%	90.8%
2010	161,446	76.3%	88.9%
2011	147,877	69.9%	78.9%
2012	148,561	34.6%	40.2%
2013	150,031	14.1%	18.4%