

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Essays in Development Economics

Permalink

<https://escholarship.org/uc/item/1hc2s9jt>

Author

Martinez Heredia, Diana Jimena

Publication Date

2024

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Essays in Development Economics

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Economics

by

Diana Jimena Martinez Heredia

Committee in charge:

Professor Prashant Bharadwaj, Co-Chair
Professor Craig Mcintosh, Co-Chair
Professor Julianne Berry Cullen
Professor Sara Rachel Lowes
Professor Gordon Carlos Mccord

2024

Copyright

Diana Jimena Martinez Heredia, 2024

All rights reserved.

The Dissertation of Diana Jimena Martinez Heredia is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2024

DEDICATION

To Adam, Shere and Bagheera, who got me here.

TABLE OF CONTENTS

Dissertation Approval Page	iii
Dedication	iv
Table of Contents	v
List of Figures	vii
List of Tables	ix
Acknowledgements	xii
Vita	xiv
Abstract of the Dissertation	xv
Introduction	1
Chapter 1 Immigrant children and optimal school choice: Evidence from the Venezuelan migration to Peru	2
1.1 Introduction	2
1.2 Background	7
1.2.1 Venezuelan Migration	7
1.2.2 Peru’s Education System	8
1.3 Data	10
1.4 Reduced Form Empirical Strategy	16
1.5 Reduced Form Results	19
1.5.1 Schooling Outcomes and School Switching	19
1.5.2 School Switching Characterization	22
1.5.3 Mechanisms Behind School Switching	28
1.6 Parental Preferences for Schools	32
1.7 Structural Model Results	35
1.7.1 Demand Estimates	36
1.7.2 Simulations	38
1.8 Conclusions	42
Chapter 2 The Effects of Expanding Worker Rights to Children	45
2.1 Introduction	45
2.2 Child labor legislation in Bolivia	50
2.2.1 Child labor legislation prior to 2014	51
2.2.2 Changes in legislation after 2014	52
2.2.3 Enforcement and awareness	55
2.3 Data	58

2.4	Empirical approach	63
2.4.1	Identification	63
2.4.2	Threats to identification	67
2.5	Effects of the 2014 Law	69
2.5.1	Effects of the law on the prevalence and sector of child work	69
2.5.2	The role of enforcement	74
2.5.3	Robustness	77
2.6	Mechanisms	81
2.6.1	Compliance costs	81
2.6.2	Avoidance behavior	84
2.6.3	Costs of work permits	88
2.7	Conclusion	88
Chapter 3	Gender and Transitions into Adulthood	91
3.1	Introduction	91
3.2	Data and Empirical Analysis	94
3.2.1	Descriptive Statistics	96
3.3	Results	98
3.3.1	Economic Outcomes	98
3.3.2	Cultural Outcomes	102
3.4	Concluding Remarks	112
Appendix A	Immigrant children and optimal school choice: Evidence from the Venezuelan migration to Peru	114
A.1	Venezuelan crisis and migration timeline	114
A.2	Figures	116
A.3	Tables	125
Appendix B	The Effects of Expanding Worker Rights to Children	133
B.1	Appendix Figures and Tables	133
B.2	List of Prohibited Tasks under the 1999 and 2014 Laws	154
B.3	Variable Definitions	155
B.4	Measuring driving time to MTEPS offices	157
B.5	Tables	158
Appendix C	Gender and Transitions into Adulthood	159
Bibliography	166

LIST OF FIGURES

Figure 1.1.	Venezuelan Migrants in the Peruvian Education System	12
Figure 1.2.	Venezuelan Migrants Performance in Math by Year	14
Figure 1.3.	Distribution of the Residualized Share of Venezuelan Migrants	18
Figure 1.4.	Academic Achievement Differences for Students who Switch - Primary ..	41
Figure 1.5.	Academic Achievement Differences for Students who Switch - Secondary	41
Figure 1.6.	Academic Achievement Differences for Students Left Behind - Primary ..	42
Figure 1.7.	Academic Achievement Differences for Students Left Behind - Secondary	42
Figure 2.1.	Ministry of Labor Inspections over Time	56
Figure 2.2.	Changes in Work Probability relative to Pre-law Periods at the 14-Year-Old Cutoff	70
Figure 2.3.	Work Probabilities at the 14-Year-Old Cutoff (Before, During, and After the Law)	71
Figure 2.4.	Compliance with Labor Regulations and Travel Time to Inspectors (Pre-Law)	75
Figure 3.1.	Midline Heterogeneity of Independent Work by Schooling and Gender ...	102
Figure A.1.	Venezuelan Migration by Year	116
Figure A.2.	Venezuelan Migrants' Enrollment by Year	116
Figure A.3.	Venezuelan Children Enrolled in Schools in 2019	117
Figure A.4.	Student Turnover by level	118
Figure A.5.	Student Turnover: Incumbents vs Migrants	118
Figure A.6.	Venezuelan Migrants Performance in Language by Year	119
Figure A.7.	Age Distribution by Grade in Primary	120
Figure A.8.	Age Distribution by Grade in Secondary	121
Figure A.9.	Treatment Distribution	121

Figure A.10.	Class Size	122
Figure A.11.	Residualized Share of Venezuelan Migrants and Residualized Class Size .	122
Figure A.12.	Non-linear Effects of Migrant Exposure on the Probability of Retention ..	123
Figure A.13.	Non-linear Effects of Migrant Exposure on the Probability of Dropout ...	123
Figure A.14.	Non-linear Effects of Migrant Exposure on Math Grades	123
Figure A.15.	Non-linear Effects of Migrant Exposure on Language Grades	124
Figure A.16.	Non-linear Effects of Migrant Exposure on the Probability of Switching Schools	124
Figure B.1.	Ministry of Labor Offices	133
Figure B.2.	Articles on the 2014 Law over Time	134
Figure B.3.	Work Probabilities by Age (Pre-law)	136
Figure B.4.	Manipulation Test: Histograms	137
Figure B.5.	Differences in densities: 14 year-old cutoff	138
Figure B.6.	Differences in densities: 12 year-old cutoff	139
Figure B.7.	Differences in densities: 10 year-old cutoff	140
Figure B.8.	Difference in Discontinuity Event Study-style Estimates: Work Probability (12- and 10-Year-Old Cutoffs)	141
Figure B.9.	Work Probabilities at the 12-Year-Old Cutoff (Before, During, and After the Law)	141
Figure B.10.	Work Probabilities at the 10-Year-Old Cutoff (Before, During, and After the Law)	142
Figure B.11.	Job Risks & Work Injuries (Before and During the Law): Stacked Data...	150
Figure B.12.	Work permits and written contracts by per-capita household income	151

LIST OF TABLES

Table 1.1.	Summary Statistics Venezuelan-Receiving Schools 2014	15
Table 1.2.	Effects of Migrant Exposure on Schooling Outcomes	20
Table 1.3.	Effects of Migrant Exposure in the Probability of Switching Schools	21
Table 1.4.	Heterogeneity of Switching Schools in Primary	23
Table 1.5.	Heterogeneity of Switching Schools in Secondary	24
Table 1.6.	Switching Between Public and Private Schools in Primary	25
Table 1.7.	Switching Between Public and Private Schools in Secondary	26
Table 1.8.	Switching Schools Profile in Primary	27
Table 1.9.	Switching Schools Profile in Secondary	28
Table 1.10.	First Stage - Primary	36
Table 1.11.	First Stage - Secondary	37
Table 1.12.	Baseline Demand Parameters	38
Table 1.13.	Taste Heterogeneity Demand Parameters - Primary	39
Table 1.14.	Taste Heterogeneity Demand Parameters - Secondary	39
Table 2.1.	Key Dimensions of Child Labor Legislation	53
Table 2.2.	Descriptive Statistics (Pre-Law)	61
Table 2.3.	Descriptive Statistics by Employer Type	62
Table 2.4.	Difference in Discontinuity Effects of the Law on the Work Probabilities, Hours, and Occupation for the 14-Year-Old Cutoff	72
Table 2.5.	Heterogeneous Effects of the Law by Distance from MTEPS Offices (Diff.- in-Disc.)	77
Table 2.6.	Effects of the Law on Risk, Injuries at Work and Wages	83
Table 2.7.	Effects of the Law on Job Location and Firm Size	86
Table 3.1.	Descriptive Statistics by Gender	97

Table 3.2.	Productive Time Use with Gender Interactions	99
Table 3.3.	Business Characteristics with Gender Interactions	100
Table 3.4.	Marriage and Fertility with Gender Interactions	108
Table 3.5.	Education and Time Use with Gender Interactions	109
Table 3.6.	Creation of New Households and Moving Across Villages with Gender Interactions.....	110
Table 3.7.	Gender Attitudes	111
Table A.1.	Mig. Share Quintiles Range	125
Table A.2.	Switching Schools Location	125
Table A.3.	Effects of Migrant Enrollment in Class Size	126
Table A.4.	Effects of Migrant Exposure on Schooling Outcomes Adding Class Size as a Control.....	127
Table A.5.	Effects of Migrant Exposure in the Probability of Switching Schools Adding Class Size as a Control	127
Table A.6.	Effect of Migrant Exposure in the Probability of Switching Schools - Sample Split by Resource Level	128
Table A.7.	Effect of Migrant Exposure in Standardized Math Std. Grades - Sample Split by Resource Level	129
Table A.8.	Effect of Migrant Exposure in Language Std. Grades - Sample Split by Resource Level	130
Table A.9.	Effect of Migrant Exposure in Math and Language Std. Grades - Sample Split by Switching Status	131
Table A.10.	Effect of Migrant Exposure By Migrant Baseline Performance	132
Table B.1.	Balance Table: Difference in Discontinuity - Household Survey	135
Table B.2.	Effects of the Law and Reversal on Work Probability for 12 and 14 year-olds (Difference-in-Difference)	142
Table B.3.	Effects of the Law on the Work Probabilities, Hours, and Occupation	143
Table B.4.	Effect of the Law on Time Use	144

Table B.5.	Difference in Discontinuity: Household Outcomes for the 14-year-old Cut-off	144
Table B.6.	Heterogeneous Effects of the Law by Driving Time from MTEPS Offices (Difference-in-Discontinuity)	145
Table B.7.	Heterogeneous Effects by Distance from MTEPS Offices, Allowing for Heterogeneity by Urban and Baseline Child Labor Rates	146
Table B.8.	Functional Form Robustness Checks: Difference in Discontinuity for Work Probability (14-Year-Old Cutoff)	147
Table B.9.	Other Robustness Checks: Difference in Discontinuity for Work Probability (14-Year-Old Cutoff)	147
Table B.10.	Difference in Difference Specifications	148
Table B.11.	Balance for 30% of Child Labor Survey Data	149
Table B.12.	Balance for Reweighted Child Labor Survey Data - Full sample	149
Table B.13.	Effects of the Law on Job Risks, and Work Injuries	152
Table B.14.	Robustness Checks: Difference in Discontinuity for Risk Outcomes	153
Table C.1.	Midline Heterogeneity of Productive Time Use by Schooling and Gender	160
Table C.2.	Income by Gender and Educational Attainment	161
Table C.3.	Endline Heterogeneity of Productive Time Use by Schooling and Gender	162
Table C.4.	Marriage and Fertility	163
Table C.5.	Education and Time Use	164
Table C.6.	Creation of New Households and Moving Across Villages	165

ACKNOWLEDGEMENTS

I want to sincerely thank Professors Craig McIntosh and Prashant Bharadwaj, who served as co-chairs of my committee. Their unique guidance and unwavering support have shaped my academic journey.

Thank you to my co-authors: María Adelaida Martínez, Diego Vera-Cossio, Leah Lakdawala, Craig McIntosh, and Andrew Zeitlin. They contributed to my research and my growth as an economist.

My deep appreciation goes to Cristián Sánchez, whose generous mentorship and unwavering support have guided me throughout this journey. I am also grateful to Isabel Hincapié for her invaluable assistance connecting me with Peruvian government officials. I thank the development and applied microeconomics faculty at UCSD for their insightful feedback and comments that have significantly enhanced my work.

I also extend my earnest thanks to the Peruvian government officials who were willing to collaborate with researchers and showed great openness to learning about the Peruvian education system. I especially thank Claudia Lisboa, the head of the statistics office at the Ministry of Education, and her team, Amalia Sevilla and Jhonny Florian. I would also like to express my gratitude to Richard Rubio and his team at Lima's Regional Education Office (DRE) for their warm welcome and openness to our research.

My most profound appreciation goes to my friends and family. My friends have been my community; their company has brought me immense joy. My family has been pivotal in this process. Thank you to my brother, parents, and dog for always supporting me. Thanks to my husband, who brought the best company along -our cats- and cared for me every step of the way. I could not have done this without him.

Chapter 1, in full, is currently being prepared for submission for publication of the material. Martinez Heredia, Diana; Martinez, Maria A. The dissertation author was the primary investigator and author of this paper.

Chapter 2, in full, has been submitted for publication of the material as it may appear

in the Journal of Development Economics, 2024, Martinez Heredia, Diana; Lakdawala, Leah K.; Vera-Cossio, Diego. The dissertation author was the primary investigator and author of this paper.

Chapter 3, in full, is currently being prepared for submission for publication of the material. Martinez Heredia, Diana; McIntosh, Craig; Zeitlin, Andrew. The dissertation author was the primary investigator and author of this material.

VITA

- 2015 Bachelor of Arts in Economics, Universidad de los Andes
- 2016 Bachelor of Arts in Literature, Universidad de los Andes
- 2014–2016 Research and Teaching Assistant, Department of Economics
Universidad de los Andes
- 2017 Master of Arts in Economics, Universidad de los Andes
- 2016–2018 Consultant
Inter-American Development Bank
- 2020 Research Assistant
University of California San Diego
- 2020 Research Assistant
Inter-American Development Bank
- 2018–2024 Teaching Assistant, Department of Economics
University of California San Diego
- 2024 Doctor of Philosophy, University of California San Diego

ABSTRACT OF THE DISSERTATION

Essays in Development Economics

by

Diana Jimena Martinez Heredia

Doctor of Philosophy in Economics

University of California San Diego, 2024

Professor Prashant Bharadwaj, Co-Chair

Professor Craig McIntosh, Co-Chair

Success in the labor market is a key driver of social mobility, and the primary focus of this dissertation is to explore how to ensure young people possess the skills and conditions necessary to achieve such success.

In recent years, millions of children have been displaced, emphasizing the importance of evidence-informed public policy for migrants and recipient communities. Chapter 1 examines the sudden influx of Venezuelan migrant children into the Peruvian school system. Analyzing cross-grade within-school variation, I find that as Venezuelan migrants enter Peruvian schools, parents transfer their children to higher-quality schools with fewer migrants. While native flight

may mitigate the effects of migrant influx for some students, it generally brings no gains and comes at a high cost.

Chapter 2 investigates the impact of a law in Bolivia that temporarily lowered the legal working age from 14 to 10 and introduced benefits and protections for child workers. Using a difference-in-discontinuity approach, I find a decrease in work for children under 14, particularly in areas with a higher threat of inspections. However, there is no evidence of improved work safety, suggesting reductions in visible child labor may be driven by avoiding legal and social sanctions rather than increased safety measures.

Chapter 3 analyzes young people's constraints as they transition into adulthood by studying programs to improve the employment prospects of underemployed Rwandan youth. I find that while cash transfers affect marriage and fertility differently for men and women, they do not significantly impact labor market outcomes differently. There are also no changes in the perception of gender roles. While financial constraints are a significant barrier for young people in Rwanda, addressing these alone does not alter deeply entrenched gender roles and cultural norms.

Introduction

Success in the labor market is a main driver of social mobility, and guaranteeing that people have the skills and the conditions to achieve that success is what motivates my research. Some recurrent themes in this dissertation encompass studying extensive margin choices as a response to shocks or policy changes as well as examining gender asymmetries in labor and education markets.

Chapter 1

Immigrant children and optimal school choice: Evidence from the Venezuelan migration to Peru

1.1 Introduction

Global migration and displacement have surged in the past two decades due to conflict, severe economic and political instability, and extreme weather events. By 2019, there were approximately 272 million international migrants, a figure that had already surpassed the United Nations' 2050 projections, estimated at around 230 million. In 2022, over 40% of the global refugee population comprised school-age children, and 76% was in low and middle-income countries¹. While access to education for these migrant children is critical to ensure access to economic opportunities, their influx can stress the existing educational system, particularly in the developing countries that are more likely to host them. In this paper, we study the effect of an inflow of one million² Venezuelan school-age migrants³ to Peru on their incumbent peers' academic performance and on the likelihood of *native flight*, a phenomenon in which parents

¹Global Trends. Forced Displacement in 2022 (UNHCR)

²UNHRC estimates from 2017 to 2019

³We will use the term 'migrants' throughout the paper. Given the Venezuelan situation, the term 'immigrant' or 'refugee' may more accurately capture some families' current situation. However, we do not have the necessary information to distinguish the various subcategories. We use the word 'migrants' to capture the migrant, immigrant, and refugee populations.

relocate their children to schools with fewer or no migrants.⁴

Peer composition can have a significant influence on academic and behavioral outcomes [Sacerdote, 2014]. Large inflows of immigrants can alter the composition of peers in schools in two ways: directly, due to differences between migrant and native students, and endogenously, as incumbent students may respond to immigrant exposure by opting for native flight. It is crucial to understand the overall effect of migration influxes on native academic achievement and how native flight impacts the academic performance of incumbent students because both factors can shape downstream outcomes related to inequality and segregation in education. Native flight is one case of a broader set of problems where the native population employs extensive margin responses to adapt to a migration shock. The labor markets and education literature studying the impact of migrant influxes has shown that extensive margin choices can be adaptive strategies for natives to navigate migration shocks.⁵ However, studying these extensive margin responses requires data that can capture the intricacies of school turnover. Moreover, identifying who moves because of native flight to isolate the effects of native flight on academic performance requires an empirical model that imposes structure on parental choices. We address these issues with unique administrative data and a structural model that complements our reduced-form strategy.

In the reduced form, we measure the average effects of exposure to migrants on native flight and the academic achievement of native students. Given the magnitude of native flight in this context, we proceed to study its implications of native flight on academic achievement. We estimate a structural model that allows us to identify specific native students induced to move due to migrant presence and study the academic achievement implications in the native population. In this second part of the paper, we model preferences to identify who moves because of migrants

⁴While typically associated with the shift from public to private schools when exposed to migrants, we use the term 'native flight' in a broader sense to include any movement of native students to schools with less migrants.

⁵Card [2001], Borjas [2006], Lewis [2013], Dustmann and Glitz [2011], Cadena and Kovak [2016] and others document how migration shocks can lead to adjustments in spatial mobility patterns and education choices among certain groups of natives. In the education literature, Tumen [2019], Farre et al. [2018], Betts and Fairlie [2003], Cascio and Lewis [2012] and others show that native students are more likely to switch schools as their exposure to migrants increases.

and analyze the academic achievement effects of native flight for two subgroups: native students who switch schools and native students left behind after the native flight.

We leverage time and cross-grade within-school variation to identify the effect of Venezuelan peers on incumbent students' academic performance and likelihood of switching schools. The cross-grade within-school design allows us to compare incumbent students exposed to a different proportion of migrants, fixing the observable and non-observable characteristics of the school. Ultimately, the variation we use comes from the age distribution of Venezuelans within schools, where we see that different grades have different shares of migrants. The effect we identify is the reduced form relative impact of the influx of Venezuelan migrants into Peruvian schools across grades, inclusive of the native children who leave and the ones who stay.

We find that the large influx of Venezuelan migrants into the Peruvian school system has effects on incumbent students' academic achievement and on the probability of transferring to a different school. Having a higher percentage of migrants as classmates decreases language and math grades. The magnitudes of our point estimates are comparable to those found in the literature [Gould et al., 2009, Imberman et al., 2012, Figlio et al., 2021]. In contemporaneous work, Contreras and Gallardo [2022] use a difference-in-difference approach and find that the Venezuelan and Haitian migration decreased sixth-grade incumbent students' standardized test scores in math and language in Chile in 2018. The magnitude of the effects is in the same range as what we find. All these studies are estimated using cross-section data. Our paper improves on the literature using student-level panel data, making our estimates more precise. We also study a context with fewer school resources and a larger migration influx. Thus, our context represents the features of countries more likely to receive migrants. However, our most significant contribution is in the analysis of native flight.

We find that the effect of migration on native flight is large compared to similar studies.⁶ In Peruvian schools, about 8 to 9% of students switch schools yearly before the migration influx. Our estimates suggest that an increase of 10 percentage points in the share of migrants

⁶Figlio and Özek [2019], Tumen [2019]

–eight migrants in an average-size school grade– increases the probability of an incumbent student switching schools by 1.55 percentage points for primary and 1.17 percentage points for secondary. These effects are equivalent to a 10.4% and 10.5% increase in primary and secondary school student turnover, respectively. The effects are non-linear and increasing in the percentage of migrants. The tipping point where migrant concentration starts affecting incumbents’ school switching is around 2.4% and 4.5% of migrants in their cohort for primary and secondary, respectively. In the native flight literature, the switching of local students to other schools is driven mainly by migrant children who do not speak the recipient country’s language, arguing that language differences demand additional school resources [Tumen, 2019, Farre et al., 2018, Betts and Fairlie, 2003, Cascio and Lewis, 2012]. Adding to this literature, this paper explores migration and school choice in a context where incumbents and migrants speak the same language. Thus, our results are more likely to reflect the effects of the perceptions of natives, the resource constraints, and a deeper interaction between native children and migrants. Contreras and Gallardo [2022] explore school switching in the context of the Venezuelan and Haitian migration to Chile but do not find a native flight between public and private schools or cream skimming. The Venezuelans who migrated to Chile are more selected than the ones who arrived in Peru. According to the IOM (2020)⁷, 74% of Venezuelan migrants who arrived in Chile have a college degree or more. This number is 20% of the Venezuelan migrants in Peru. Our context allows us to study a migration influx with less selection.

We characterize the schools to which these students are more likely to switch. We find migrant concentration increases the probability of students switching to higher-quality schools with fewer migrants in primary and secondary schools. Like the native flight widely studied in the US, migrant inflow generates student transfers from public to private schools. However, given the flexibility of the Peruvian school system (school enrollment is not restricted to the neighborhood of residence), we also observe student mobility within private and public schools. We can see whether students move to schools in different cities. We find that the effects of

⁷Chaves-González and Echevarría Estrada [2020]

migration on student turnover are not explained by families moving to different locations.

We follow the reduced-form analysis with a structural model that allows us to isolate the effects of native flight on the academic achievement of the students who opt for native flight and the native students they leave behind. In the structural part of the paper, we follow the literature that models the preferences for schools to study effects in the demand for schools [Allende, 2019, Burgess et al., 2015, Sanchez, 2018, Neilson et al., 2013, Lavy et al., 2009, Hastings et al., 2009]. The purpose of the model is to identify who moves because they were exposed to migrants, in order to understand who gains and who loses when there is native flight. In addition to our panel administrative education data, we use household level data from the national census to study families' school choice decisions and how they are affected by exposure to migrants. Having estimated the preferences for schools, we can make comparisons between school choices made by native parents when they face the presence of migrants in their children's schools and a counterfactual scenario in which they do not. We estimate these counterfactuals for the children that switch schools and for those that are left behind by the native flight.

Before the migration influx, about 8 to 9% of students switch schools every year. In the reduced form, we observe that, on average, native students that are more exposed to migrants are more likely to switch schools. The structural model shows that, after the migration influx, about 9 to 16% of the total turnover is induced by the presence of migrants in Primary and Secondary schools, respectively. Among the students who switch schools because they are exposed to migrants, there are small academic performance gains from the migrant-induced movement. We see that these students experience an increase of 0.02 to 0.05 SD on their math academic performance when their school choice accounts for migrant presence, compared to a counterfactual in which there are no migrants. These gains are small. Due to the effect size, we consider the gains a precisely estimated zero for all groups except for some lower socioeconomic status students. Moreover, moving is costly. Many students move to private schools. We see that, on average, these families' tuition costs increase by 330 to 412 soles (around 89 to 111 USD) per month, which is above to the monthly equivalent of the cash transfer program *Juntos*

(100 soles). On the other hand, the students who are left behind do not seem to be negatively affected by the native flight. The estimate for their loss in academic achievement ranges from -0.005 to 0.005 SD, which we interpret as a precisely estimated zero. The evidence from the model suggests that native flight can be viewed as a strategic adaptation strategy employed by some parents in response to the influx of migrants. However, it is costly and generally brings no gains to students who switch schools. This shows that native flight is driven by factors beyond academic achievement losses, which motivate parents to make costly decisions.

1.2 Background

1.2.1 Venezuelan Migration

The number of migrants leaving Venezuela has increased significantly in the last years, and 20% of the migrants are going to Peru. The UNHRC estimates that around 1.3 million Venezuelans were living in Peru by 2021. Figure A.1 shows the exponential increase of Venezuelan immigrants in Peru after the Venezuelan Government opened the border with Colombia in 2016. The most common migration route is through Colombia and then Ecuador. The data from the Peruvian migration agency shows that around 95% of the migrants travel by bus, in a journey that takes at the very least four days and can last for months. Government records show that around 500,000 Venezuelans have applied for refugee status and that around 18% of them travel with children⁸. They are located mainly in Lima and in other cities along the Peruvian coast, as shown in Figure A.3. They are either unemployed or working in informal jobs. Those migrants who join the formal sector report low wages⁹.

In 2017, the Peruvian government passed a law to establish a temporary permanence permit (PTP for its acronym in Spanish). This permit allowed Venezuelan migrants to stay legally in Peru for a year and gave their children access to public health and education public services. Even if it expired, Venezuelan migrants could present their Venezuelan ID or passport to meet

⁸Standard Operating Procedure for Venezuelan Migrants in Peru by the IOM.

⁹Standard Operating Procedure for Venezuelan Migrants in Peru by the IOM.

the requirements. These somewhat lenient requirements made Peru a more attractive destination for Venezuelan migrants ¹⁰.

The massive inflow came hand in hand with a change in the attitudes of Peruvian citizens. In 2018, the local newspaper *El Comercio* surveyed people in Lima about their attitudes regarding Venezuelan migrants. Around 55% of them disagreed with allowing Venezuelan migrants into the country. In 2019, a new survey by the same newspaper resulted in 67%. In 2019, the Peruvian migration agency launched a campaign against xenophobia. However, the people's perceptions that Venezuelans are taking scarce jobs and services from Peruvians are pervasive. The Universidad Católica in Peru and the Panamerican Development Foundation report widespread concerns¹¹. These organizations collected testimonies of education experts who report that finding schools to enroll Peruvian children is difficult and perceive that the inflow of immigrant students worsens this situation.

1.2.2 Peru's Education System

The Peruvian education system enrolls more than 6 million students in primary and secondary levels each year. In 2019, 74% were in public and 26% in private schools. Education is compulsory for primary and secondary levels. Public schools are free, and there is a wide variety of private schools in terms of tuition costs and quality levels. Unlike other systems, parents do not face restrictions in choosing a school depending on their neighborhood or residence. They can enroll their children in any school if there are slots. However, public schools prioritize enrollment first for children with disabilities and children whose siblings are already enrolled there and second for children who reside in the school area. Yearly enrollment is automatic for children that are already in a school. According to Peruvian law, the access and permanence of the students in public and private schools cannot be denied or conditioned by students' characteristics. Additionally, private schools can not perform evaluations or tests on students as

¹⁰See details on the Venezuelan migration to Peru on Appendix A.1

¹¹<https://data2.unhcr.org/en/documents/download/70863>

part of their admission process.¹²

Regarding enrollment procedures, students need some form of identification to enroll, but an exception is made for Venezuelan students with no identification who can apply for enrollment and defer the document requirement. Last UNESCO's report about the situation of Venezuelan children in Peru documents that this procedure is subject to the discretion of principals. Some reject Venezuelan migrants if they cannot provide evidence of the last grade they passed, while others let children enroll without any documents regarding this matter¹³. The report also mentions that costs to attend offices and lack of knowledge of the Peruvian school system are the main restrictions that Venezuelan parents face in enrolling their children in Peruvian schools. Despite this, enrollment grew; Figure A.2 shows the evolution of enrollment of Venezuelan children in Peruvian primary and secondary schools. Between 2014 and 2019, 75.6% of Venezuelan students enrolled in public schools and 24.4% in private schools.

The Peruvian school system allows school switching at the end of the school year and within the same school year. Parents who want to transfer their children to a different private or public school need to find a spot in the new school and ask for the enrollment transfer between the origin and the new institution. To help parents search for schools for their children, the Ministry of Education developed a webpage with all the schools' characteristics, including quality measures, location, and the number of free slots by grade. As shown in Figure A.4, turnover rates in Peru lie between 9.5 and 10.5 percent in primary and 8 and 8.5 percent in secondary.¹⁴ These turnover rates are close to those of countries such as Chile, which has a turnover rate in primary of 11.5 percent [Zamora Poblete and Moforte Madsen, 2013]. Furthermore, it is somewhat smaller but similar to Florida's 16 percent turnover rate during the Haitian immigration after the earthquake reported by Figlio and Özek [2019].

There are no guidelines about the expenditure per student in the Peruvian law regarding

¹²Ministry of Education Decree N° 005-2021-MINEDU

¹³<https://inee.org/node/9953>

¹⁴Figure A.5 shows that Venezuelan migrants have higher turnover rates than their incumbent peers, consistent with Venezuelan parents having informal and less stable jobs [Morales and Pierola, 2020]

the budget and resource allocation for public education. Saavedra and Suárez [2002] document how the resources allocated to public schools depend on the bargaining power of school principals, who negotiate with local authorities that allocate budgets. Additionally, they are affected by the inertia in old budget structures that have not changed over time. Finally, schools with more complex infrastructure require more resources for upkeep and operation. These elements have resulted in a high inequality in per-student expenditure in Peru's different regions, cities, and neighborhoods. In this same study, the authors mention that parental investments in education are crucial for the operation of schools, even in public schools. The expenditure per student reported by the Peruvian Statistics Institute (INEI) in 2018 is about 835 USD in primary school and 1,180 USD in secondary school. For reference, on average, countries in the OECD spend about 8,700 USD per primary school student and 10,200 per secondary school student (OECD). This disadvantage in resources goes hand in hand with lower education quality. Peru's average score on the PISA tests was about 401 in 2018, the US score was 505, and the average OECD score was 487 [OECD, 2019].

1.3 Data

We use data from four administrative sources: (i) SIAGIE (Sistema de Información de Apoyo a la Gestión Educativa), a student panel from the Ministry of Education; (ii) ECE (Evaluación Censal de Estudiantes) student-level data on Peru's standardized test; (iii) School Roster (Padrón Escolar) and School Census (Censo Escolar) school characteristics panel that includes both private and public schools in Peru. (iv) SiseVe, a platform where schools report school violence cases. For all of them, we have data from 2014 to 2019.

SIAGIE is the system that keeps enrollment records for every student in the education system in Peru. This dataset is a student-level panel from 2014 to 2019 that includes students' school, grade, classroom, nationality, age, sex, and report cards. The student ID allows us to track students across schools and years and merge the information with other Ministry of Education

datasets. The student tracking gives us information on student transfers between schools and dropouts.

The key outcomes we use from SIAGIE are report cards' grades, switching schools, dropout, and retention. We use school report cards' grades standardized at the grade level. For the school switching outcome, a student transfers schools in year t if the school in year t differs from the school in $t - 1$. Since the first year in our data is 2014, we cannot observe which school the students enrolled in the prior year. Hence, we can only construct this variable from 2015 onwards.¹⁵ For dropout, a student drops out of school in year t if they are not present in the school system in $t + 1$ and did not graduate in $t + 1$. Finally, for retention, if we observe a student in the same grade in year t as in $t - 1$, we classify them as they experience retention.

Our second data source is the ECE, the Student Census Evaluation (known as ECE by its Spanish acronym). The ECE is a mandatory test taken by all Peruvian students in the second and fourth grades of elementary school and the second grade of secondary school. The test evaluates two subjects: language and math, and the scores have no impact on students' GPAs or report cards. The ECE includes a short survey to the students or parents (depending on the grade) that provides data on parental education and the household's socioeconomic characteristics. This data includes a wealth index constructed by the MinEduc using principal components analysis over this household survey information. The Peruvian Ministry of Education uses this index as their primary indicator of the socioeconomic level of the school.

We have access to the ECE data from 2014 to 2019, but the data is somewhat sparse. Only the 2nd-grade tests are available starting in 2014. 4th-grade and 8th-grade tests are available starting in 2016 and 2015, respectively. Due to *El Niño* rainy season and the teacher's strike, the test was suspended in 2017¹⁶. Initially, the universe of students in each grade did each test. However, the Ministry of Education modified who took the standardized test in later years. In

¹⁵Some schools offer primary and secondary education, while others do not have continuity and only offer primary. Thus, we do not have information on school switching when students advance from primary to secondary in non-continuous schools. Hence, we cannot construct this variable for 7th graders.

¹⁶<http://umc.minedu.gob.pe/evaluaciones-censales/sus-ece/>

2018, only a sample of 2nd graders took the test, while the universe of all 4th and 8th graders took it. In 2019, the universe of 8th graders took the test, and 2nd and 4th graders' subsamples took it. We standardize the test scores at the grade level, as we do with report cards' grades.

Our third data source is the Education Quality Statistics System ESCALE (by its Spanish acronym). This Ministry of Education tool contains information on all registered public and private educational institutions in Peru. We will use two primary datasets from ESCALE, the School Roster and the School Census. The School Roster has data on the type of school management (public, private, charter), ownership, whether it is coeducational, type of classrooms (single-teacher, multi-teacher, multi-grade, complete multi-teacher), and geocoded location. The School Census includes data on total enrollment (by grade, sex, age, native language), number of classrooms, teachers' experience, education, tenure, and school infrastructure (construction materials, public services, toilets, library, and computers). We use indicators for public and private schools, school location, the district IDs, the teacher-student ratio, and a school wealth index from the school-level data. We use a principal components analysis to construct the school wealth index, which contains school infrastructure (walls, floors, and roof), whether the school has access to essential services (clean water, electricity, trash, and sewage), the number of computers for pedagogical purposes that the school has, and whether the school has a library.

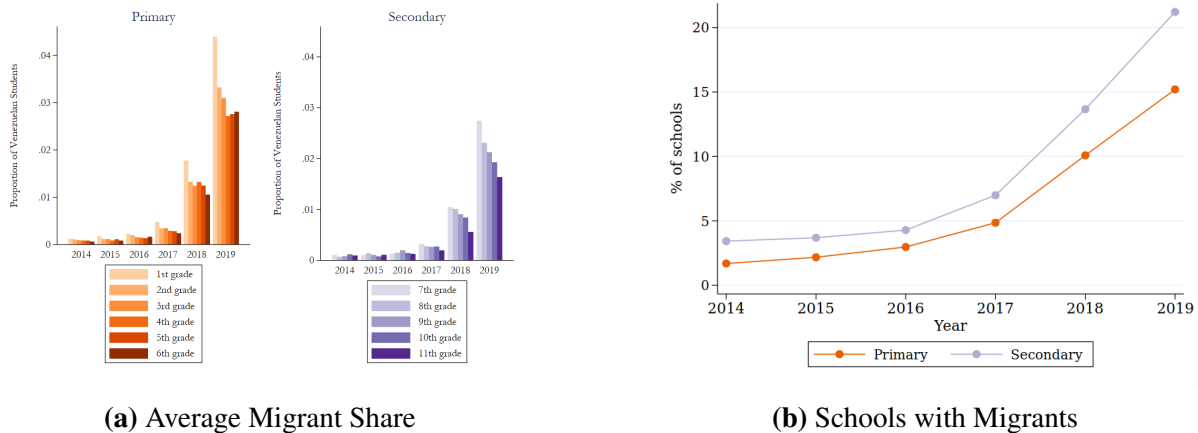


Figure 1.1. Venezuelan Migrants in the Peruvian Education System

Our fourth data source is SiseVe, a Peruvian Ministry of Education platform where

schools, students, and parents can file school violence reports. The list of reports is public, and each report has information on the year, school district, frequency, and motive of the aggression. From the motive, we can count the number of school violence reports related to discrimination in each school district ¹⁷

Finally, our fifth data source is the microdata from the 2017 national census. We are able to match the national census and our enrollment administrative data. This match gives us two crucial pieces of information: the students' proximity to school and their household's socioeconomic status. There is a 73% match between the national census data and the enrollment administrative data in 2019, which accounts for 4,608,866 students. We choose 2019 because that is the year with the most prominent presence of Venezuelan migrants in the school system. The census data has the geolocation of students' residences for around 36% of the sample. This sample is on average more urban and of lower socio economic status than the rest of the sample, however it includes the students who are more affected by the migrant influx. The proximity of students to schools is a key component of model of school preferences. It is rare to find such detailed and comprehensive information in a developing country setting, hence this match between the census data and the enrollment data presents a unique opportunity to study school choice in the context of a migrant influx.

We proceed to provide some descriptive statistics of these data. In Figure 1.1a, we observe that the average migrant share increases exponentially over time. In 2014 it was lower than 1% in primary and secondary. In 2019, the average migrant share by grade was 4 to 5% in primary and 2 to 3% in secondary. These figures can be lower than expected, considering the magnitude of the Venezuelan migratory influx. However, Figure 1.1b shows that the number of schools with Venezuelan migrants is relatively high. In 2019, 15% of primary schools in the country had migrants, while 20% of secondary schools in the country had migrants. Figures 1.1a and 1.1b show that the Venezuelan children migration inflow was large and broadly spread

¹⁷In 2019, the MinEduc included a question on the School Census of whether the school reported or not to the SiseVe and the number of reports. We have this information at the school level only for this year.

among different schools.




Figure 1.2. Venezuelan Migrants Performance in Math by Year

Figures 1.2a and 1.2b show the trends of performance in math for primary and secondary school, respectively. Both figures show the average by year, dividing the sample into two groups: Incumbent students and Venezuelans in the first year they appear on the panel (new migrants). For the second group, we also show the math grades after one year in the system (new migrants t+1) when their grades are more comparable to the ones of their peers. There is considerable heterogeneity in the academic performance of migrants over time, even after a year in the Peruvian school system. This heterogeneity in the performance of Venezuelan migrants is consistent with mixed migration —there are economic migrants, citizens returning to their countries of origin, and refugees. Besides, many highly educated Venezuelans migrated. With surveys, the Peruvian government estimates that 57.9% of the migrants have higher education studies ¹⁸.

Figures 1.2a and 1.2b show that the standardized grades of entering Venezuelan migrants decreased after the migration shock started in 2017. At the beginning of the migration episode, migrants were relatively high achievers, but this tendency reversed as the migration increased. Since most of our variation comes from later years, we expect that the impact of the relatively low-achieving migrants will dominate the effects. Figures A.6a and A.6b show the same pattern

¹⁸Standard Operating Procedure for Venezuelan Migrants in Peru by the IOM.

for language grades.

Table 1.1. Summary Statistics Venezuelan-Receiving Schools 2014

	Primary		Secondary	
	Venezuelan Receiving Schools	Other Schools	Venezuelan Receiving Schools	Other Schools
Public schools = 1	1.000	0.741	0.880	0.569
Total student count	235.438	68.446	353.917	135.697
Proportion of female students	0.491	0.480	0.483	0.463
Student-Teacher ratio	21.941	17.916	15.581	12.461
% of teachers with professional education	0.826	0.722	0.955	0.907
Avg. math std. test score	-0.114	-0.305	-0.260	-0.170
Avg. language std. test score	-0.227	-0.322	-0.289	-0.210
SES index students	-0.433	-0.374	-0.373	-0.302
% of students high SES index	0.062	0.126	0.069	0.140
School violence reported = 1	0.179	0.049	0.338	0.139
Number of schools	5,510	37,494	2,942	13,583

All mean differences are statistically significant at 1% level. The std. errors for the differences are clustered at the district level and include district fixed effects. Math and language test scores are standardized at the grade level and from 2015, the earliest year available. Parents' SES is measured by the socioeconomic index of the ECE surveys on student household characteristics in 2016 (earliest year available). In our sample, the SES index goes from -2.9 to 1.8. The Peruvian MinEduc defines a high SES index as being at the 85th percentile or higher. The school wealth index was constructed using principal components and it includes school infrastructure, essential services, computers and library it ranges from -3.7 to 9.5. The school violence information comes from the school census, which asks the principal for the number of SiseVe reports made during 2019 (there is not school level data for earlier years).

Table 1.1 shows descriptive statistics in 2014 before the migrant influx for schools with and without Venezuelan migrants in 2019. In both primary and secondary, migrants tend to choose larger public schools with a higher proportion of teachers with professional education. However, schools chosen by migrant families were more strained regarding resources, having larger student-teacher ratios and poorer students before most migrants arrived. On average, these schools' 2015 standardized test scores for math and language were lower for both secondary and primary schools. In sum, Venezuelan receiving schools were systematically different from other schools even before the migrants' arrival. This is precisely why a simple difference

between student outcomes in schools with and without migrants will not identify the effect of the Venezuelan migrants' inflow on incumbent students.

1.4 Reduced Form Empirical Strategy

Immigrants are more likely to settle in areas with more immigrants from their country [Stuart and Taylor, 2021, Carrington et al., 1996, Card, 2001]. Then, there is an endogenous placement of immigrants in schools with specific characteristics. A model comparing schools with higher and lower proportions of Venezuelan students will probably generate biased estimates due to selection into schools. The differences between schools will account for all the schools' observable and non-observable characteristics and not only for the immigrant inflow effects. We rely on cross-grade within-school variation in the number of Venezuelan students entering the education system in Peru to address this problem. We implement a school-year fixed effects estimation to study the impact of contemporaneous exposure to Venezuelan migrants.

To identify the effect of a change in the concurrent number of immigrants on incumbent students' outcomes, we compare grades with different proportions of Venezuelan students within the same school and year. The identifying assumption is that the grade placement of Venezuelan students within schools is uncorrelated with what incumbent students' conditional outcomes would have been in the absence of the influx of Venezuelan migrants. In Figures A.7 and A.8 we can see the age distribution of migrants and incumbent students per grade. Grade placement by age is similar for incumbent students and migrants. Migrants are slightly older on average, but their age for grade coincides with the ages for the grade of incumbent students. Hence, our identifying assumption is closely related to the assumption that the age distribution of Venezuelan migrants within schools is uncorrelated to grade-specific educational inputs within a school. Principals can play a role in the selection of migrants into schools. We are assuming that principals will discourage all migrants equally. If they discourage migrants of a specific age, or prefer to enroll migrants into a specific grade, we have to assume that this selection

is not correlated with grade pre-existing characteristics. Our empirical analysis follows this specification:

$$Y_{i,sg,t} = \alpha + \beta V_{sg,t} + \gamma X_i + \theta_{s,t} + \psi_g + \varepsilon_{i,sg,t} \quad (1.1)$$

Where, $Y_{i,sg,t}$ is the achievement measure of the incumbent student i of school-grade sg and year t . X_i is a vector of student characteristics, including sex, age, and the baseline math grade. This is the standardized math grade the first time we observe the incumbent student in the data set. $\theta_{s,t}$ and ψ_g are school by year and grade fixed-effects. Our treatment variable is $V_{sg,t}$, which is the percentage of Venezuelan students of the total student body in school-grade sg and year t . We observe all the outcomes at the end of the school year. The share of migrants, $V_{sg,t}$, corresponds to the peers that the incumbent children had in their grade during the school year. Thus, we calculate the effect of the concurrent share of migrants on the outcomes. This specification only includes incumbent students. Given that the migrant share is at the school-grade level, our standard errors are clustered at that same level.

Incumbent students move to different schools over time. To avoid selection problems induced by parents of incumbent students who choose to move their children to another school (school switchers), we implement an estimation analogous to an intention to treat estimate (ITT). If children move, we assign them their previous school s – the one where they were enrolled before transferring schools after we observe the transfer. We also assign them the share of migrants they would have had if they had not switched to another school. Then, β accounts for the effect of being exposed to a larger share of migrants and the student turnover caused by the exposure.

Students' outcomes reflect the cumulative previous and current investments made to improve their human capital. Including baseline outcomes to control for the earlier investments allows us to focus on the effect of a contemporary input, the share of Venezuelan migrants in the cohort. Effects on test scores often fade out quickly [Bailey et al., 2020]; hence, concurrent

exposure is the key dimension we expect to impact schooling outcomes significantly.

As we mentioned in section 1.3, migrant students are relatively spread out in the education system. Figure A.9 shows the distribution of $V_{sg,t}$. The share ranges from 0 to around 25% of children in primary and secondary schools and is right-skewed. However, we see considerable variability in the migrant share in each grade, even after controlling for school-year, grade, and district fixed effects. Figure 1.3 shows the distribution of the residualized migrant share per grade. The distribution is consistent across grades in both primary and secondary schools, although first grade has a slightly larger range than the other grades in primary schools.

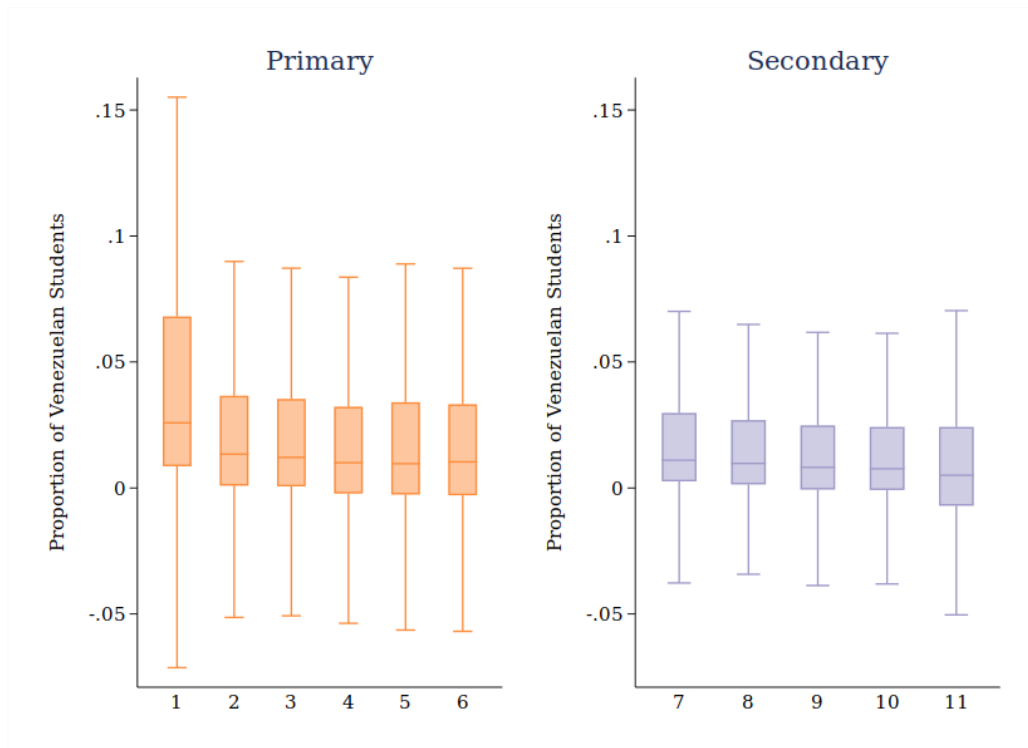


Figure 1.3. Distribution of the Residualized Share of Venezuelan Migrants

The interpretation of our effect could be affected by the reallocation of resources between grades that receive more and fewer migrants within a school. Schools that receive more migrants in one grade likely reallocate resources from other grades to adjust to the changes. This reallocation would affect the outcomes we are using as counterfactuals negatively. If this is the case, our results are a lower bound of the effects of the migration.

1.5 Reduced Form Results

In this section, we first present estimates of the effect of migrant concentration on incumbents' schooling outcomes and the probability of switching schools. Then, we characterize the switching incumbent students and the difference between the origin and destination schools they are being transferred to. Finally, we dig into the mechanisms behind parents' re-optimizing and changing their children to different schools after the Venezuelan migrants' arrival.

1.5.1 Schooling Outcomes and School Switching

Table 1.2 presents the effects of Venezuelan immigrants' concentration on incumbent retention, dropout rates, and language and math grades, estimated using the cross-grade within-school variation on the share of migrants. Panels 1 and 2 present results for incumbent students in primary school grades (1-6) and secondary school grades (7-11), respectively. We find statistically significant results for primary school math and language grades. These results show that an increase of 1 percentage point in the share of Venezuelan migrants in a grade (approximately one migrant) decreases math and language grades by 0.0015 and 0.002 standard deviations, respectively. Similarly, the estimates for secondary school show positive effects on retention and dropout rates and negative effects on math and language grades, both statistically significant. In secondary, a 1 percentage point increase in the share of migrants increases the probability of retention and dropout by 0.009 and 0.023 percentage points, respectively, and reduces math and language grades by 0.007 standard deviations.

These effects are small and comparable with the magnitude of the evacuee effects on math standardized test scores measured by Imberman et al. [2012] (-0.01 standard deviations on math test scores), and the refugee effects measured by Figlio and Özek [2019], (0.003 and 0.006 standard deviations on math and language test scores respectively). A 5 percentage point is a shift in the distribution of the migrant share across school grades, which means going from the 25th to the 75th percentile of the distribution, representing, on average, three more migrant

children in primary and five more migrant children in secondary. Going from the 25th to the 75th percentile of migrant share in primary school will decrease math and language grades by 0.007 and 0.008 standard deviations. In secondary, it will increase the likelihood of retention by 0.045 and the likelihood of dropping out by 0.1 percentage points and reduce math and language grades by 0.03 standard deviations. These effects are plausible and on the lower end of the peer effects range in the literature summarized by Sacerdote [2014].

Table 1.2. Effects of Migrant Exposure on Schooling Outcomes

Primary				
	Retention (1)	Dropout (2)	Math grades (3)	Language grades (4)
Mig. Share ITT	0.00254 (0.00259)	0.00834 (0.00554)	-0.154*** (0.0405)	-0.201*** (0.0385)
R-squared	0.032	0.107	0.182	0.184
Obs.	14,700,335	11,681,011	14,576,843	14,576,961
Mean	.005	.011	-.015	-.012

Secondary				
	Retention (1)	Dropout (2)	Math grades (3)	Language grades (4)
Mig. Share ITT	0.00906** (0.00382)	0.0235* (0.0138)	-0.685*** (0.0875)	-0.701*** (0.0921)
R-squared	0.020	0.105	0.270	0.281
Obs.	12,622,876	10,049,531	12,134,882	12,199,462
Mean	.005	.032	-.044	-.024

Standard errors clustered at school-grade level in parentheses. *** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$. Control variables: sex, age, baseline math grade, school by year fixed, grade, year, and district fixed effects. The sample includes only incumbent students from 2015 to 2019 in primary and secondary schools.

Incumbents' parents might re-optimize and respond to the inflow of migrants to their children's schools. Student turnover is particularly important because children in Peru do not necessarily need to attend the schools in their neighborhoods. Also, parents have the legal right to change their children to a different school at any time during the academic year. This results in a turnover of about 8 to 9% each year, prior to the migrant influx. Table 1.3 examines the effects of Venezuelan immigrant concentration on the likelihood of school switching. Our estimates

suggest that a 1 percentage point increase in the share of migrants increases the probability of an incumbent student switching schools by 0.275 percentage points for primary and 0.174 percentage points for secondary. Given that the average turnover rate for primary schools is 11.6% and for secondary schools is 8.1%. These effects are equivalent to a 2.4% and 2.1% increase in the student turnout for primary and secondary schools, respectively.

Table 1.3. Effects of Migrant Exposure in the Probability of Switching Schools

	Primary (1)	Secondary (2)
Mig. Share ITT	0.275*** (0.0122)	0.174*** (0.0207)
R-squared	0.094	0.085
Obs.	14,700,336	12,622,876
Mean	.116	.081

Standard errors clustered at school-grade level in parentheses. *** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$. Control variables: sex, age, baseline math grade, school by year and grade. The sample includes only incumbent students from 2015 to 2019 in primary and secondary schools.

A 5 percentage point increase in the migrant share will increase the likelihood of switching schools by 1.37 percentage points for a primary student and 0.87 percentage points for a secondary student. In contrast to the effects on student achievement, these effects are larger than similar effects found in the literature. The point estimates are similar in magnitude if we compare them with Figlio and Özek [2019] point estimates of refugees on student mobility¹⁹. However, considering that in Peru in 2019, the student turnover rate was lower than the turnover rate in Florida in Figlio and Özek [2019] paper, which lies between 16% and 17%, effects are more extensive in the context of the influx of Venezuelan migrants to Peru.

As the peer-effects literature points out, the linear-in-means model might be insufficient to understand the mechanisms underlying peer effects [Sacerdote, 2011]. We estimated a non-parametric, non-linear model of the effect of migration on schooling outcomes and school

¹⁹One percentage point increase in refugee concentration increases the probability of student movement by 0.2 percentage points [Figlio and Özek, 2019].

switching. We use the equation 1.1 specification, and instead of having the migrant share $V_{sg,t}$ as our primary explanatory variable, we add five dummy variables that take the value of 1 if the migrant share is on quintiles 1 to 5 of the migrant share distribution in grades where there is at least one migrant. In this way, we ensure that the comparison group is composed of grades with no migrants. Table A.1 shows the range of the migrant share on each quintile. Figures A.12 to A.16 show the coefficient of each dummy and its confidence interval for each outcome. Figures A.12 and A.13 show null results for the incumbent's likelihood of retention and dropout on both primary and secondary schools. Figures A.14 and A.15 show significant negative results for math and language. In primary schools, moving from a zero migrant share to having at least 6 migrants (8.33%) in a school-grade decreases math and language grades by 0.03 standard deviations. There are no significant effects for lower quintiles. In Secondary, we see negative and significant effects that increase between the second and fifth quintiles. Larger effects occur in the fifth quintile, where an 8.33% increase in migrant share reduces incumbents' math and language grades by 0.07 and 0.06 standard deviations, respectively. The effects are between 0.02 and 0.03 standard deviations on the second to fourth quintiles. Figure A.16 shows the point estimates of the non-linear specification on the probability of school switching. As the percentage of migrants increases, the likelihood of incumbents switching schools increases non-linearly. We find that primary incumbents start switching when the migrant share is higher than 2.44%, of their school grade (1.95 migrants), while in secondary, the tipping point is at 4.35% (3.5 migrants).

1.5.2 School Switching Characterization

Parents might re-optimize differently depending on their children's characteristics. We estimate heterogeneity analyses to characterize the students more prone to transfer schools after exposure to a higher share of migrants. Additionally, it generates an indirect change in peer composition that might reinforce or mitigate the migrant effects on schooling outcomes. If the school switchers are low achievers, the positive peer effects will mitigate the adverse effects on achievement. If the high achievers are the ones switching, the negative effect on achievement

will be reinforced by the peer composition changes.

Our heterogeneity analyses follow the main specification in equation 1.1 and include the heterogeneity measure and the interaction between the heterogeneity measure and the migrant share. We explore heterogeneity in three dimensions: gender, baseline math grades, and baseline language grades, and the interaction between them. Table 1.4 shows the results for primary school. We see that boys and girls are equally likely to transfer schools as they are more exposed to migrants. For baseline performance in grades, we find that primary school students with lower grades are more likely to transfer schools when exposed to a higher share of migrants than students with higher grades exposed to the same share of migrants. Increasing 1 percentage point the migrant share in their school grade makes students with one standard deviation higher math grades less likely to move by 0.014 percentage points.

Table 1.4. Heterogeneity of Switching Schools in Primary

	<i>Heterogeneity Measures</i>		
	Girl (1)	Baseline Math Grade (2)	Baseline Lang Grade (3)
Mig. Share \times Heterogeneity	-0.011 (0.013)	-0.014** (0.006)	-0.005 (0.006)
Het. Measure	0.000 (0.000)	-0.010*** (0.000)	-0.010*** (0.000)
Mig. Share ITT	0.282*** (0.014)	0.275*** (0.012)	0.275*** (0.012)
R-squared	0.093	0.094	0.094
Obs.	14,700,292	14,699,846	14,700,292

Standard errors clustered at school-grade level in parentheses. *** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$. Control variables: school by year and grade fixed effects. The specification includes the heterogeneity measure, the share of migrants per grade and the interaction between the migrant share and the heterogeneity measure. The baseline grades correspond to the first grade we observe for every student. The sample includes incumbent students from 2015 to 2019 in primary schools.

Table 1.5 shows the effects on students in secondary schools. Unlike what we found in primary schools, girls are more likely to switch schools when exposed to migrants. With a 1 percentage point increase in migrant share in the school-grade cohort, the likelihood that

girls switch schools increases by 0.22 percentage points, while, for boys, it is 0.138 percentage points. The difference is significant at the 1% level. Additionally, higher-achieving students are more likely to switch secondary schools when exposed to more migrants. An increase of 1 percentage point on the migrant share increases the likelihood of switching by 0.041 percentage points for students with one standard deviation higher math grades and 0.047 percentage points for students with one standard deviation higher language grades. In secondary, girls and students who have higher grades are more likely to change schools as they are exposed to the same share of migrants.

Table 1.5. Heterogeneity of Switching Schools in Secondary

	<i>Heterogeneity Measures</i>		
	Girl (1)	Baseline Math Grade (2)	Baseline Lang Grade (3)
Mig. Share \times Heterogeneity	0.082*** (0.018)	0.041*** (0.009)	0.047*** (0.009)
Het. Measure	0.002*** (0.000)	-0.005*** (0.000)	-0.006*** (0.000)
Mig. Share ITT	0.138*** (0.022)	0.175*** (0.021)	0.174*** (0.021)
R-squared	0.085	0.085	0.085
Obs.	12,640,153	12,610,324	12,640,155

Standard errors clustered at school-grade level in parentheses. *** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$. Control variables: school by year and grade fixed effects. The specification includes the heterogeneity measure, the share of migrants per grade and the interaction between the migrant share and the heterogeneity measure. The baseline grades correspond to the first grade we observe for every student. The sample includes incumbent students from 2015 to 2019 in secondary schools.

We have characterized the students who are more likely to switch schools when exposed to migrants. Now, we describe the schools to which they move. We estimate our main specification from equation 1.1 on the changes of school time-invariant characteristics before and after students switch. First, we focus on whether the movement comes from public or private schools and whether the schools chosen are public or private. Tables 1.6 and 1.7 show the effect of migrants on the likelihood of moving from a public to a private school in column 1, from a private to a

public school in column 2, from a public to a private school in column 3, and from a private to a private school in column 4.

Table 1.6. Switching Between Public and Private Schools in Primary

	Public to Public	Private to Public	Public to Private	Private to Private
	(1)	(2)	(3)	(4)
Mig. Share ITT	0.087*** (0.006)	0.040*** (0.007)	0.049*** (0.004)	0.099*** (0.009)
R-squared	0.065	0.126	0.029	0.165
Obs.	14,698,662	14,698,662	14,698,662	14,698,662
Mean	.051	.02	.013	.031

Standard errors clustered at school-grade level in parentheses. *** p<0.01 ** p<0.05 * p<0.1. Control variables: sex, age, baseline math grade, school by year, and grade fixed effects. By definition the outcome variables take 0 value for all non-switchers. The sample includes only incumbent students from 2015 to 2019 in primary schools.

In Table 1.6, we see that primary school students who are exposed to a higher share of migrants are more likely to move to private schools from both public and private schools. The likelihood of switching from public to public schools increases by 0.087 percentage points as the share of migrants in the school grade increases by one percentage point. The likelihood of switching from private to public schools increases by 0.04 percentage points as the likelihood of switching from private to public schools. The higher effects are on the likelihood of switching from private to private schools, which increases by 0.1. The school mobility rates from public to public, private to public, public to private, and private to private primary schools are 5%, 2%, 1.3%, and 3.1%, respectively. Then, the effect of a one percentage point increase in migrant share increases the probability of switching within public schools by 1.7% and from private to public schools by 2%. In contrast, it increases the probability of switching from a public to a private school by 3.8% and 3.2% from private to private schools.

Table 1.7 shows that secondary school students are most likely to switch from public to public schools but also move to private schools. Increasing the migrant share by 1 percentage point increases the likelihood of switching from public to public schools by 0.024 percentage

Table 1.7. Switching Between Public and Private Schools in Secondary

	Public to Public	Private to Public	Public to Private	Private to Private
	(1)	(2)	(3)	(4)
Mig. Share ITT	0.024** (0.010)	0.039*** (0.011)	0.024*** (0.006)	0.089*** (0.015)
R-squared	0.046	0.094	0.024	0.127
Obs.	12,619,952	12,619,952	12,619,952	12,619,952
Mean	.034	.013	.012	.021
Standard Deviation	.181	.113	.111	.144

Standard errors clustered at school-grade level in parentheses. *** p<0.01 ** p<0.05 * p<0.1. Control variables: sex, age, baseline math grade, school by year and grade fixed effects. By definition the outcome variables take 0 value for all non-switchers. The sample includes only incumbent students from 2015 to 2019 in secondary schools.

points, from private to public by 0.039, from public to private by 0.024 percentage points, and from private to private schools by 0.089 percentage points. The school mobility rates for secondary schools are 3.4% from public to public, 1.3% from private to public, 1.2% from public to private, and 2.1% from private to private. This last group of students incurs new costs to transfer schools after exposure to a higher share of migrants. A 1 percentage point on the migrant share increases the public-to-private and private-to-private switching rates by 2% and 4%, while it changes the public-to-public switching rate by 0.7%.

Second, we characterize schools in different dimensions: the proportion of migrants, test scores, student-teacher ratio, Parents' SES index, and teachers' education. We construct the historical average of these variables by school ²⁰. Then, we construct dummy variables that indicate if students are moving to schools that historically have had fewer migrants, higher test scores, lower student-teacher ratios, higher SES indexes in 2016, and a higher proportion of teachers with professional education. ²¹ Tables 1.8 and 1.9 show the results for primary and secondary schools, respectively.

The school destination characteristics are consistent with the incumbent's avoiding higher

²⁰We use the 2016 Parent's SES index from the household survey made by the Ministry of Education as part of the ECE national standardized test

²¹By construction, these dummy variables equal 0 for all non-switchers

Table 1.8. Switching Schools Profile in Primary

	Fewer Venezuelans (1)	Higher Math Scores (2)	Higher Lang. Scores (3)	Lower Stud/Teach (4)	Higher Parents' SES (5)	Higher Teach Educ (6)
Mig. Share	0.241***	0.152***	0.153***	0.111***	0.123***	0.131***
ITT	(0.012)	(0.011)	(0.010)	(0.008)	(0.009)	(0.010)
R-squared	0.112	0.113	0.104	0.062	0.042	0.138
Obs.	14,700,336	13,872,404	13,873,023	14,685,206	12,806,693	14,698,662
Mean	.039	.061	.06	.054	.05	.052

Standard errors clustered at school-grade level in parentheses. *** p<0.01 ** p<0.05 * p<0.1. Control variables: sex, age, baseline math grade, school by year and grade fixed effects. All the outcome variables are dummies defined by difference in school characteristics after switching. By definition is 0 for all non-switchers. School characteristics are historical averages from 2014 to 2019. Parents' SES is measured by the socioeconomic index of the ECE surveys on student household characteristics in 2016. The sample includes only incumbent students from 2015 to 2019 in primary schools.

concentrations of migrants. In Table 1.8, we can observe that, as primary school students are more exposed to migrants, they tend to go to schools with fewer migrants, higher-income families, and higher quality in all our measures. A 1 percentage point increase in the share of migrants increases the probability of switching to a school with a lower native/migrant proportion by 0.241 percentage points. This same increase in the share of migrants increases the likelihood of switching to a higher-quality school. This effect ranges from 0.111 to 0.153 percentage points for the different quality measures. Table 1.9 shows the same pattern of results for secondary school. Increasing the migrant share by 1 percentage point increases the likelihood of switching to a school with fewer migrants by 0.093 percentage points. It also increases the likelihood of switching to higher-quality schools between 0.068 to 0.125 percentage points, depending on the quality measure we consider. Finally, as we expected from the high likelihood of switching to private schools, primary school and secondary school switchers have a higher probability of switching to schools with wealthier parents.

School switching can result from households moving to neighborhoods with fewer migrants. White flight literature has shown that white households left cities and went to suburban areas in response to the black migration from the rural South [Boustan, 2010]. A natural question in this context is whether the children are not only switching schools but families are also moving

Table 1.9. Switching Schools Profile in Secondary

	Fewer Venezuelans (1)	Higher Math Scores (2)	Higher Lang. Scores (3)	Lower Stud/Teach (4)	Higher Parents' SES (5)	Higher Teach Educ (6)
Mig. Share	0.093***	0.112***	0.112***	0.068***	0.102***	0.125***
ITT	(0.022)	(0.017)	(0.017)	(0.013)	(0.012)	(0.017)
R-squared	0.085	0.072	0.069	0.059	0.040	0.091
Obs.	12,622,876	12,413,962	12,414,030	12,617,787	12,004,923	12,619,952
Mean	.027	.04	.04	.039	.039	.038

Standard errors clustered at school-grade level in parentheses. *** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$. Control variables: sex, age, baseline math grade, school by year and grade fixed effects. All the outcome variables are dummies defined by difference in school characteristics after switching. By definition is 0 for all non-switchers. School characteristics are historical averages from 2014 to 2019. Parents' SES is measured by the socioeconomic index of the ECE surveys on student household characteristics in 2016. The sample includes only incumbent students from 2015 to 2019 in secondary schools.

to different neighborhoods after the migrant's arrival. Given that we do not have data on the student's residence, we do not know if they are changing their neighborhood of residence, but we know the geolocation of the schools. Table A.2 shows the main specification results for dummy variables that take the value of 1 if the incumbent student switches a school in a different region, province, district, and the distance in miles between the origin and the destiny schools²². We find minor significant effects of migrant share on switching school regions, provinces, and districts. Moreover, the effect of a 5 percentage point increase over the distance between the origin and destiny schools is 0.33 miles for primary and non-significantly different from zero in secondary. These estimates suggest that the effects of migration on student turnover are not explained by families moving to different locations.

1.5.3 Mechanisms Behind School Switching

How does the inflow of immigrants into schools translate into higher student turnover rates? This section explores the mechanisms at play in primary and secondary schools. First, we discard this as a mechanical effect of class size changes. Second, we explore resources' role in the parent's decisions and check for evidence of binding resource constraints. Lastly, we examine

²²Peru territory is divided into regions that are subdivided by provinces and provinces are subdivided by districts. There are 25 regions, 196 provinces, and 1869 districts.

the peer composition changes and whether this is a negative peer effect driven by low-achieving migrants or a disruption effect of having new, culturally different children at school. Although this evidence is descriptive, it shows that peer effects are one of the mechanisms at play and evidence of binding resource constraints, especially in low-resourced and public schools.

Class size: The significant influx of migrants may affect class sizes. Figure A.10 shows that class sizes are relatively stable over time, even after the migration increased exponentially in 2017. On average, classrooms in primary schools are smaller and vary between 10 and 14 students; in secondary schools, classrooms range between 18 and 22 students. Although we do not see sharp increases in class size, there still may be a relationship between the number of migrants and class size. As a first approach to check for a correlation between migrant inflow and class size, we plot the residualized class size and number of migrants in the school grade after controlling for all our covariates and fixed effects. In Figure A.11, the solid line shows the correlation between the residualized number of migrants and the residualized class size. The dashed line has a slope of 1 to compare the fitted values to the one-to-one relationship between the X and Y-axis variables. Figure A.11 shows that the relationship between the residualized class size and the number of migrants is not one-to-one. Although class size is increasing and is one of the factors that can explain our effects, we do not find that the migrant influx has increased the probability of children being in classes that exceed the government recommended class size. Table A.3 shows that the likelihood of exceeding the maximum class size is very close to zero, and even negative. When we split the effect by public and private schools, we see that the effect is zero in public schools and negative and small in private schools. This is consistent with private schools having more resources to adapt. Hence, although class sizes increase, we do not find evidence of them being crowded to the point of exceeding the maximum class size established by the Ministry of Education.

The role of school and parent resources: Considering that the Peruvian education system has high inequality in resource availability for students, if the inflow of migrants reduces school resources beyond having mechanical effects on class size, we should see different effects

by resource availability. First, we split the sample between public and private schools. We expect tuition payment in private schools to mitigate resource constraints that the public sector might have experienced after the sudden inflow of Venezuelan migrants. More specifically, in private schools, the migrant inflow is not expected to change per-pupil expenditure. However, in Peru, between 2014 and 2019, 27% of schools were private and had very high variability in prices. Balarin [2015] shows that after its expansion in the late 90s and early 00s, Peruvian private education was no longer a privilege of the wealthy elites. There are low-fee private schools in poor settlements that do not necessarily offer higher quality than the public schools serving the same areas [Balarin, 2015].

For this reason, we added two more resource measures. First, we split the school sample by strictly parent resources using the average parent's socioeconomic index in 2016 before the migration pick. This socioeconomic index is calculated by the ministry of education at the school level using ECE's survey information on parents' education, income, assets, and household characteristics. Finally, we use a cleaner measure of resource availability at the school level: the student/teacher ratio in 2014 before our analysis period starts. Some primary schools in Peru have teachers that simultaneously teach one or more grades²³. In secondary, there are different teachers for different assignments that might teach more than one classroom at a time. Class size does not capture the differences in resources for any of these modalities. In this context student-teacher ratio is a better resource availability measure at the school level. Since there is selection because parents can enroll their children in high and low-resourced schools, this is a descriptive exercise.

Table A.6 shows the resource splitting exercise for school switching. The first panel shows the results for public schools, schools below the 25th percentile of parents' SES index, and schools above the 75th percentile of the student-teacher ratio. The second panel shows the results for private schools, schools above the 75th percentile of parents' SES index, and schools

²³In our sample, these represent 63% of the schools, most of which are in rural areas and 19% of the total primary student population

below the 25th percentile of the student-teacher ratio. Column (1) shows that migrant effects on the likelihood of switching schools are higher in private schools. Given the wide market private schools cover in Peru is not clear if parents' income and willingness to pay or school resource constraints are behind these results. Column (2) shows that in primary, the effect of migrants on the probability of switching schools is higher for schools where the average parent is at the fourth SES index quantile, parents with more resources. However, tables A.7 and A.8 does not show evidence of statistically significant detrimental effects on achievement in this schools. Moreover, the negative effect on math and language in primary is driven by public schools. For high-income parents and private schools, evidence is inconsistent with a mitigating strategy. On the other hand, table A.6 Column (3) shows higher effects of migrant concentration on lower-resourced schools where tables A.7 and A.8 show higher effects on achievement measures. This evidence is consistent with a parents' mitigation strategy and binding school resource constraints in low-resourced and public primary schools.

The results for secondary schools in Table A.6 show a slightly different pattern. Column (4) shows that the effect of migrant concentration on the probability of school switching is positive and significant in private schools. At the same time, it is not statistically significant and is close to zero for public schools. Contrasted by the results shown in Column (4) in tables A.7 and A.8 where the negative effects of migration on achievement are significantly higher for public schools. Again private schools, evidence is inconsistent with a mitigating strategy. On the other hand, when we split the sample by parents' and school resources (Columns (5) and (6) respectively), we observe that the effects of migration on student turnover are higher for low-resourced parents and schools. Tables tables A.7 and A.8 the detrimental effects on grades are higher in low-resource schools. Hence, evidence is consistent with binding resource constraints and parents' using school switching as a mitigation strategy in low-resourced settings.

Changes on peer composition: Following the hypothesis of adverse peer effects driven by low-achieving migrant students, we analyze our main specification and outcomes breaking up the migrant share into two components: migrants that perform above and below the median

performance level. This is a purely descriptive exercise. We construct migrant performance at baseline -the first year we observe the migrants in our data- and calculate the median performance for the baseline grade and year. Column 1 in Table A.10 shows that, in secondary, predominantly low-achieving migrants cause incumbents to move, while in primary, both low and high-achieving migrants cause switching, and high-achieving migrants cause slightly more movements. Results Columns 3 to 5 on Table A.10 show that both higher and lower-performing migrants adversely affect performance measures in primary and secondary schools. According to Hanushek et al. [2004], the disruption caused by new incoming students causes negative peer effects, which are larger in high-turnout schools. The negative point estimates for both types of migrants suggest this might be the mechanism behind our main results and not changes in the skill level peer composition.

1.6 Parental Preferences for Schools

As incumbent students are more exposed to migrants, we observe that there are minor negative effects on their academic performance. However, the likelihood that these students will switch schools is large. We cannot distinguish how much of the effects on academic performance come from the peer re-composition after incumbent students sort.

Our modeling approach allows us to identify which students change schools due to their exposure to migrants. We can then compare the school choices made by families when they face the presence of migrants to a counterfactual scenario in which there is an absence of migrants. We can then study the outcomes of the native students who switch schools and the students they leave behind under both counterfactuals to shed light on who benefits and who is adversely affected by native flight.

In our model, the determinants for school choice are the proximity to the school, the cost of tuition, school quality, school characteristics, and the proportion of migrants in the school. We measure school quality as the value added of the school. We allow preferences for

these determinants to be heterogeneous by gender and baseline achievement of the student after observing that there is heterogeneity by these characteristics in the reduced form native flight results. Student i 's preferences over school j are:

$$U_{ij} = \beta_1 d_{ij} + \beta_2 p_j + \beta_3 X_j + \beta_4 V_j + \beta_5 q_j + \xi_j + \varepsilon_{ij} \quad (1.2)$$

Where p_j is the school's price, d_{ij} is the distance to school, X_j is a vector of school characteristics, V_j is the proportion of Venezuelan students in school j , and q_j is the quality of school j . We allow heterogeneity in the preferences, so for $k \in [2, 5]$, $\beta_{ki} = \beta_k + \sum_r z_{ir} \beta_{kr}$, with z_{ir} being the demographic characteristic r for student i . We use two demographic characteristics based on our findings from the reduced form: gender and whether they are lower or higher achieving. We let $W_{ij} = \beta_1 d_{ij} + \beta_2 p_j + \beta_3 X_j + \beta_4 V_j + \beta_5 q_j + \xi_j$, so that the indirect utility for schools is $U_{ij} = W_{ij} + \varepsilon_{ij}$. We assume ε_{ij} is EV type I. Hence the probability of student i choosing school j is:

$$P_{ij} = \frac{e^{W_{ij}}}{\sum_k e^{W_{ik}}} \quad (1.3)$$

We use a Maximum Likelihood to estimate preference for proximity, taste heterogeneity, and mean utilities or school popularity. The mean utilities absorb the preference components from the indirect utility function that vary only at the school level:

$$LL(\beta) = \sum_i \sum_j C_{ij} \ln \frac{\exp(\beta_1 d_{ij} + (\beta_2 - \beta_2) p_j + (\beta_3 - \beta_3) X_j + (\beta_4 - \beta_4) V_j + (\beta_5 - \beta_5) q_j + \delta_j)}{\sum_k \exp(\beta_1 d_{ik} + (\beta_2 - \beta_2) p_k + (\beta_3 - \beta_3) X_k + (\beta_4 - \beta_4) V_k + (\beta_5 - \beta_5) q_k + \delta_k)} \quad (1.4)$$

Where δ_j is the variation in preferences only at the school level

$$\delta_j = \beta_2 p_j + \beta_3 X_j + \beta_4 V_j + \beta_5 q_j + \xi_j \quad (1.5)$$

From this Maximum Likelihood estimation, we estimate $\hat{\delta}_j$. To estimate \hat{q}_j , we regress $Y_i =$

$Z_i\gamma + q_j + \varepsilon_j$, with Z_i being observable characteristics, and Y_i being test scores. From this process, we obtain \hat{q}_j and $\hat{\delta}_j$. Using a 2SLS estimation, we estimate:

$$\hat{\delta}_j = \beta_2 p_j + \beta_3 X_j + \beta_4 V_j + \beta_5 \hat{q}_j + \xi_j \quad (1.6)$$

Additionally, we recognize that price, quality of the school, and the presence of migrants in the school can be endogenous and we employ a W2SLS strategy to tackle this issue. To account for the endogeneity in price and quality, our first set includes instruments for the price and quality of the schools. Following Allende [2019], we leverage variation from a law reform in Peru that aimed to expand tenured contracts and raise wages for public school teachers. The implementation of the law spanned from 2013 to 2018. We use four instruments from 2018: a teacher wage index for teachers in public schools, teacher job openings in the school, the number of teachers with temporary contracts, and an indicator of whether the school hired teachers under the new regulation (for public schools only). The reform N-29944 regulated the selection process and career advancement of public school teachers in Peru. It established an entrance exam, which is mandated for all candidates to get a tenured teacher contract in a public school. It also established the pay grade scales for each level of experience. The goal of this law was to create better incentives to hire qualified teachers who can guarantee a better quality of education in public schools. Allende [2019] documents how the reform induced variation in the wages and types of contracts through time and space. The assumption behind the exclusion restriction of these instruments is that the variation that the reform introduces on our measures of changes in teacher contracts are unrelated to the unobserved school characteristics that drive parental preferences.

To account for the endogeneity of the presence of migrants in schools, our second set of instruments includes variables related to the geographic settlement of migrants in the previous years. We use the number of migrants by age group in the social security office closest to the school in 2018 and the proportion of migrant students in the three closest schools in

2018. The assumption behind the exclusion restriction of these instruments is that the spatial variation that explains the presence of Venezuelan migrants in 2018 is unrelated to the unobserved school characteristics that drive parental preferences in 2019. The idea behind this exclusion restriction is that the choice of residence happens before school enrollment, so adjusting to current unobservable shocks of school preferences takes a long time. All these elements allow us to construct the preferences of parents for schools.

Each student can choose any school within their market. In our context, each market is a city, except for Lima, which contains four markets (one for each subregion of the city). We use Lima and the following 6 largest cities in Peru, for a total of 10 markets. In secondary school, we estimate the model for grades 8 to 11, since some students attend schools that only offer Primary school up to grade 6, and have to enroll in a different school in grade 7. With these sample constraints, we have a sample of 132,401 students in Secondary.

1.7 Structural Model Results

The preference parameters we have estimated allow us to predict the choices of students. Our modeling approach allows us to identify which students change schools due to their exposure to migrants. We do this by comparing the choices made by families when they consider the presence of migrants in their school selection to a counterfactual scenario in which they do not account for such presence. We see that 9% to 16% of the switching we observe post-migration influx is due to the presence of migrants in the school. Since we can identify which students correspond to this proportion of the sample, we can also predict their outcomes under both counterfactuals. To do this, we follow

With this information, we can study the outcomes of the students who switch schools and the students that they leave behind under both counterfactuals to shed light on who benefits and who is adversely affected by native flight.

1.7.1 Demand Estimates

We start by presenting the first stage of the demand estimates to speak to the relevance of the instruments in Tables 1.10 and 1.11. For the cost and value added instruments in Primary school, we see that the F-statistics are 206.3 and 67.7, respectively. In Secondary, the F-statistics are 57.34 and 73.98, respectively. We see that our measures of vacancies, the teacher wage index, and temporary contracts and teacher test scores in 2018 are positively related to both cost and value-added. The relationship is strong enough to reassure us that we do not have a weak instrument problem. For the instruments of the presence of Venezuelan students in the school, we see that the F statistic is 36.69 in Primary school and 19.74 in Secondary school, supporting the relevance condition for this set of instruments. We see that the geographic location of Venezuelans in nearby schools in 2018 is related to the presence of Venezuelans in 2019.

Table 1.10. First Stage - Primary

	(1) Cost	(2) Value Added	(3) Ven. Students (proportion)
<i>Cost and Value Added Instruments</i>			
Teacher Wage Index × vacancies	1.41e-07*** (2.30e-08)	7.47e-08*** (2.22e-08)	4.92e-09** (2.03e-09)
Teachers under temporary contract	0.0411*** (0.000972)	0.0228*** (0.000938)	-0.000966*** (8.59e-05)
Teachers hired under new regulation	-0.0478* (0.0282)	-0.0342 (0.0272)	0.00918*** (0.00249)
Teacher test scores	0.0567* (0.0303)	0.00577 (0.0292)	0.00627** (0.00267)
<i>Proportion of Ven. Instruments</i>			
Venezuelans in neighboring schools	-0.00891 (0.0239)	-0.00989 (0.0231)	0.00355* (0.00211)
Venezuelans in SS (18+)	-2.33e-05 (1.82e-05)	-1.39e-05 (1.76e-05)	-9.71e-07 (1.61e-06)
Venezuelans in neighboring schools × Venezuelans in SS (18+)	-8.66e-05** (3.53e-05)	-1.90e-05 (3.41e-05)	-9.32e-06*** (3.12e-06)
Venezuelans in SS (13-17)	0.00113* (0.000606)	0.000325 (0.000585)	0.000211*** (5.36e-05)
Venezuelans in neighboring schools × Venezuelans in SS (13-17)	0.00237** (0.000952)	0.000552 (0.000919)	0.000249*** (8.42e-05)
<i>School Characteristics</i>			
School is gendered	0.213*** (0.0332)	0.145*** (0.0320)	-0.00272 (0.00293)
School is public	-0.241*** (0.0141)	0.422*** (0.0136)	0.0265*** (0.00125)
Constant	0.111*** (0.0240)	-0.211*** (0.0232)	0.00743*** (0.00212)
Observations	5,223	5,223	5,223
F-statistic	206.3	67.70	36.69

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 1.11. First Stage - Secondary

	(1) Cost	(2) Value Added	(3) Ven. Students (proportion)
<i>Cost and Value Added Instruments</i>			
Teacher Wage Index × vacancies	2.43e-08** (9.55e-09)	1.61e-08** (8.01e-09)	-5.81e-10 (4.31e-10)
Teachers under temporary contract	0.0204*** (0.00103)	0.0195*** (0.000865)	-0.000395*** (4.65e-05)
Teachers hired under new regulation	-0.0660* (0.0356)	-0.0167 (0.0298)	-0.000894 (0.00160)
Teacher test scores	0.0292 (0.0274)	0.00623 (0.0230)	0.00376*** (0.00124)
<i>Proportion of Ven. Instruments</i>			
Venezuelans in neighboring schools	0.0166 (0.0290)	-0.00277 (0.0243)	0.00275** (0.00131)
Venezuelans in SS (13-17)	-0.000850** (0.000340)	-0.000203 (0.000285)	9.66e-05*** (1.53e-05)
Venezuelans in neighboring schools × Venezuelans in SS (13-17)	5.81e-05 (0.000308)	0.000252 (0.000258)	-2.09e-05 (1.39e-05)
<i>School Characteristics</i>			
School is gendered	0.125*** (0.0309)	0.320*** (0.0259)	-0.00357** (0.00139)
School is public	-0.441*** (0.0178)	-0.0151 (0.0149)	0.00921*** (0.000802)
Constant	0.327*** (0.0284)	0.120*** (0.0238)	0.00671*** (0.00128)
Observations	3,487	3,487	3,487
F-statistic	57.34	73.98	19.74

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

We examine our estimates for the preference parameters for schools in Table 1.12. The baseline parameters have the expected signs. Parents prefer schools closer to their residence, lower-cost schools, schools with better quality, and private schools. We see that the coefficient on the proportion of Venezuelan migrants is negative and large, albeit somewhat noisy. This is also the case for the cost parameters. However, large standard errors are expected, given the number of instrumented variables in the model that introduce noise to our estimates. Tables 1.13 and 1.14 show the taste heterogeneity parameters for girls and high achievers in Secondary schools. We define a high achiever student as a student whose baseline academic achievement is above the median academic achievement. We see that our Secondary school results are in line with what we find in the reduced form. We see that, in Secondary, girls and high-achieving students have a negative estimate for the preference of Venezuelan migrants in their schools. In primary school, the model predictions are different from the reduced form. We see that girls and high achieving

students have a stronger disutility from being exposed to Venezuelan migrants in schools.

Table 1.12. Baseline Demand Parameters

	(1) Primary	(2) Secondary
Distance	-8.188*** (0.021)	-7.3737*** (0.021)
Cost	-2.805 (5.497)	-3.537 (2.894)
Value Added	8.296 (10.63)	6.725** (2.764)
Proportion of Venezuelans in school	-2.188 (18.04)	-7.337 (23.65)
School is gendered	-0.376 (0.437)	-1.490** (0.611)
School is public	-2.434 (6.255)	-0.0106 (1.140)
Constant	28.81*** (2.740)	23.31*** (0.780)
Observations	5,223	3,487
R-squared	0.441	0.649

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

1.7.2 Simulations

With the estimates of the preference parameters of our demand model, we can identify which students change schools due to their exposure to migrants. We can do this by comparing the choices made by families when they consider the presence of migrants in their school selection to a counterfactual scenario in which they do not account for such presence. We observe that, in Secondary school, native flight accounts for 13% of the overall turnover in the education system. With this information, we can study the outcomes of the students who switch schools and the students that they leave behind under both counterfactuals to shed light on who benefits and who is adversely affected by native flight.

Table 1.13. Taste Heterogeneity Demand Parameters - Primary

	Heterogeneity by	
	Girl (1)	Higher Achiever (2)
Cost	-0.020** (0.013)	-0.087*** (0.018)
Proportion of Venezuelan Students	-0.357*** (0.178)	-0.313 (0.230)
School is gendered	0.640*** (0.028)	0.069*** (0.035)
School is public	0.003 (0.015)	-0.028*** (0.019)
\hat{q}	0.081*** (0.015)	0.178*** (0.019)
Observations	114,490,007	114,490,007

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 1.14. Taste Heterogeneity Demand Parameters - Secondary

	Heterogeneity by	
	Girl (1)	Higher Achiever (2)
Cost	-0.034*** (0.013)	0.028*** (0.013)
Proportion of Venezuelan Students	-0.694*** (0.314)	-0.314 (0.332)
School is gendered	0.773*** (0.021)	-0.002 (0.021)
School is public	0.114*** (0.015)	0.048*** (0.016)
\hat{q}	0.152*** (0.017)	0.116*** (0.018)
Observations	54,388,691	54,388,691

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

We estimate the outcomes under both counterfactuals following Dubin and McFadden [1984], who develop a control function approach for a discrete choice model. We can then compare both counterfactual outcomes for two subgroups: the students that produce the native flight, and the students who are left behind by those students who do the native flight.

First, we focus on the students who switch because migrants induce them. In Figures 1.4 and 1.5, the first estimate from the top down is the difference in switchers' achievement in the counterfactual with migrants minus their achievement in the counterfactual in which there are no migrants. In Primary schools, there appear to be no overall gains for the students who move. In Secondary schools, we see that native flight benefits students who move. The effect is close to 0.02 SD and statistically significant. The following four estimates in the figure break down the group of switchers into four subgroups: by academic achievement and by socioeconomic status. In Primary school, high achieving and low SES students show significant gains from movement of 0.05 SD. In Secondary school, we see that there is not any heterogeneity by academic achievement and SES. Overall, the results show that some students experience small benefits in academic achievement from switching schools, but that gain comes at a monetary cost. On average, when accounting for the presence for migrants, students who are induced to switch pay 330 more soles in Secondary and 412 more soles in Primary in tuition (89 and 111 USD) than in the counterfactual in which no migrants are present. For students who pay tuition, the median tuition payment is 255 soles in Secondary and 560 soles in Primary (69-151 USD) under the counterfactual where migrants are present, and 35 soles in Secondary and 60 soles in Primary (10-16 USD) under the counterfactual of no migrants. Although native flight can be an adaptive strategy for some students, it is costly for those who switch to private schools and overall brings no substantial gains in academic achievement to students who switch.

In Figures 1.6 and 1.7, the first estimate from the top down is the difference in achievement in the counterfactual with migrants minus their achievement in the counterfactual in which there are no migrants for the students who are left behind by the native flight. We see that the effect on academic achievement is a precisely estimated zero for both Primary and Secondary schools.




Figure 1.4. Academic Achievement Differences for Students who Switch - Primary

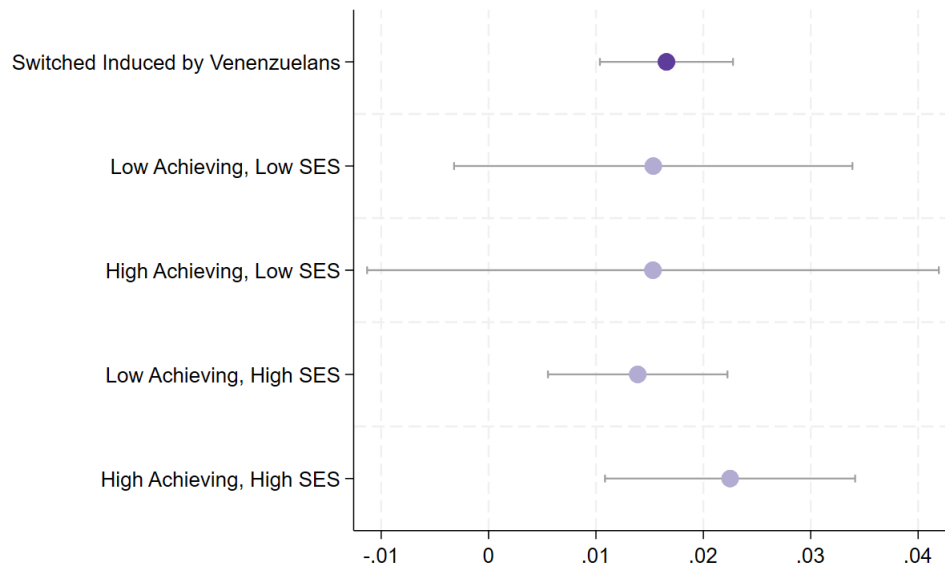


Figure 1.5. Academic Achievement Differences for Students who Switch - Secondary

As before, we look at the four subgroups given by achievement level and socioeconomic status to understand if there are gains or losses that average to zero. We see no distinct patterns for any particular group. Overall, we see that facilitating native flight is not detrimental to this population, but it is costly and not beneficial for the students who switch to private schools.

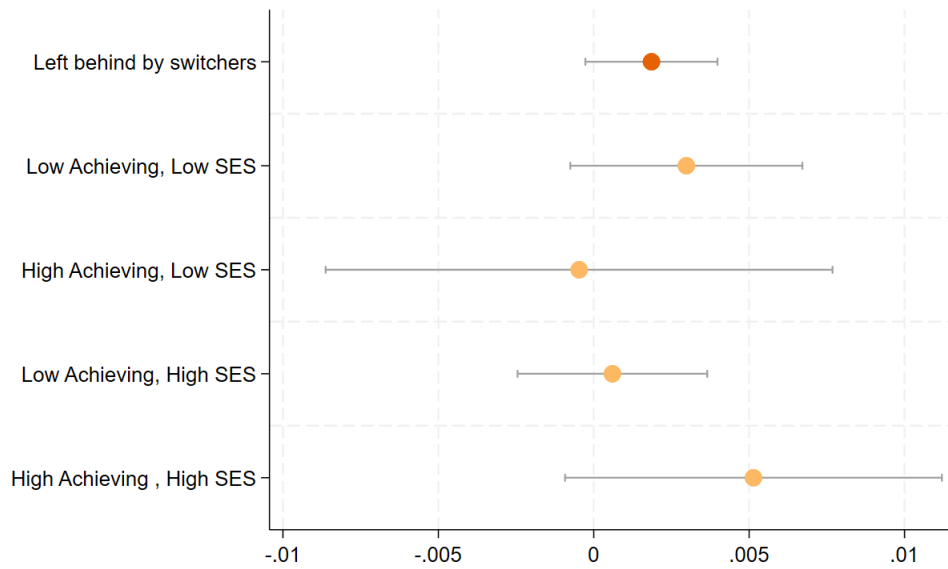


Figure 1.6. Academic Achievement Differences for Students Left Behind - Primary

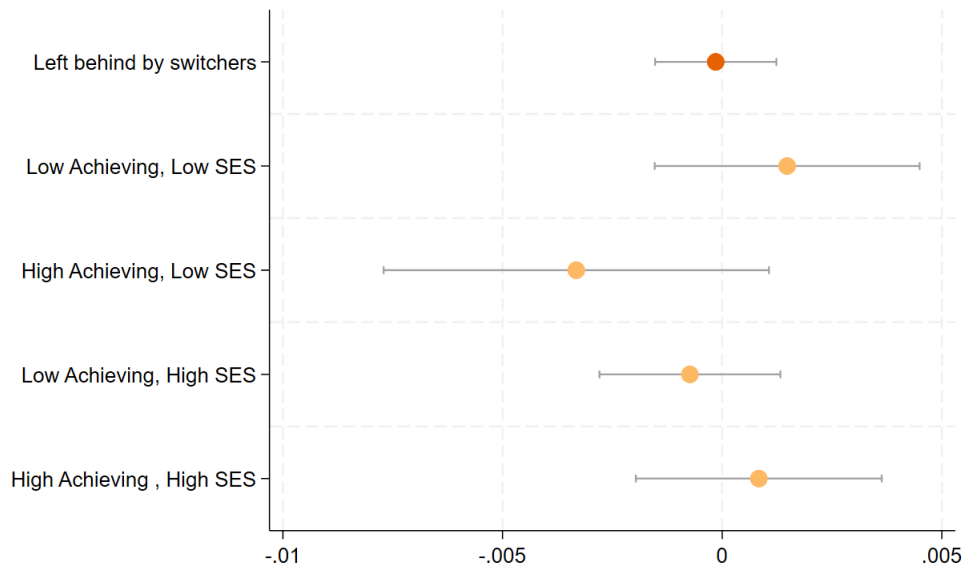


Figure 1.7. Academic Achievement Differences for Students Left Behind - Secondary

1.8 Conclusions

As Venezuelan migrants enter Peruvian schools, incumbent students experience detrimental effects on schooling outcomes. A higher share of migrants increases the likelihood of

dropping out and decreases language and math achievement. The effects are small and comparable to those found in similar studies. However, our estimates are more precise because we use nationwide panel data. We also find that parents re-optimize when their children are more exposed to migrants by sending them to other schools. We characterize the students who move and the schools to which they move. In primary schools, students with lower grades are more likely to switch schools; in secondary schools, students with higher grades and girls are more likely to move. Students transfer predominantly from public schools to private and other public schools. Switching students move to higher-quality schools and schools with fewer migrants. However, we do not find evidence of students moving to schools far away from their original school, suggesting their families are not moving to different neighborhoods.

We discuss potential mechanisms behind the effects. Although larger classes play a role, they are not the main driver of the negative effects on achievement and the rise in school turnout. We find larger effects of migration on the probability of switching schools of high socioeconomic status families and private schools. We also find evidence consistent with binding resource constraints. Parents' are more likely to choose school switching in public and low-resourced schools where migrants have minor but adverse effects on incumbents' schooling outcomes. Nevertheless, we cannot disentangle the sorting effect from the peer effects migrants generate from their lone presence in the classroom unrelated to changes in incumbent composition.

The reduced form provides insights into average effects. we use a structural model to identify specific individuals induced to move due to migrant presence and shed light on the welfare implications. In the Peruvian school system, student turnover is about 8 to 9% per year. In the reduced form, we observe that, on average, native students that are more exposed to migrants are more likely to switch schools. The structural model shows that about 20% of the total turnover is induced by the presence of migrants in Secondary schools. Among the students who switch schools induced by migrants, there are small academic performance gains from the migrant-induced movement. Most of those gains come from the higher socioeconomic status students. However, moving is costly. Many students move to private schools. We see that on

average, the monetary cost of tuition that these families face increases substantially. On the other hand, the students that are left behind do not seem to be negatively affected by the native flight. We interpret the estimate for their loss in academic achievement as a precisely estimated zero. The evidence from the model suggests that native flight can be viewed as a strategic adaptation strategy employed by a few parents in response to the influx of migrants. However, overall the gains are close to zero, and they come at a high cost for the families who switch their children to private school.

Chapter 1, in full, is currently being prepared for submission for publication of the material. Martinez Heredia, Diana; Martinez, Maria A. The dissertation author was the primary investigator and author of this paper.

Chapter 2

The Effects of Expanding Worker Rights to Children

2.1 Introduction

Over 160 million children worldwide are engaged in child labor and roughly 50% (79 million) perform hazardous work [International Labour Organization, 2021]. Most working children are likely hired “off-the-books”, in precarious conditions and under the radar of workplace regulations. Existing evidence focuses on policies to directly prevent children from working such as child labor bans (e.g, Bharadwaj et al. [2020], Basu and Van [1998], Bargain and Boutin [2021], Piza et al. [2023] and Abman et al. [2023]). Much less attention has been paid to policies that aim at improving the working conditions of children or that try to bring child workers “out of the shadows”. Moreover, little is known about the effectiveness of policies to protect workers that are typically hired informally, whose participation in labor markets is often socially condemned and rarely legally recognized. These key characteristics of the labor market for children may alter the incentives of households and employers in ways that run contrary to the policy goals, which could have implications for understanding the effects of regulations in other markets that share similar characteristics.

We leverage a unique setting to better understand the effects of policies regulating the working conditions of children and the mechanisms behind such effects. Specifically, we study a policy change that legally recognized child work with the aim of expanding benefits and

protections to child workers, similar to those granted to adults working formally. In 2014, Bolivia passed legislation that recognized the work of children as young as 10 years old, whose age placed them below the official minimum working age of 14 years old. The law enabled young children aged 10 to 13 to work legally (subject to obtaining a work permit) while simultaneously extending benefits and protections to these workers.¹ For example, the law entitled working children to adult minimum wages and to 2 paid hours per day to devote to school or study; the law also required that employers guarantee safe working conditions for children.

To ensure enforcement, the government tasked local offices of the Ministry of Labor and Social Protection (MTEPS) with adding child labor inspections to their regular labor and workplace inspections. Nationally, child labor inspections doubled between 2013 (the year prior to the law) and 2017 [Ministerio de Trabajo, Empleo y Previsión Social, 2015-2018, U.S. Department of Labor, 2011-2019]. Awareness of the policy change appears to have been widespread, as evidenced by coverage of the law in national and international news outlets and by recorded attendance of official workshops conducted by the MTEPS to educate children, parents, and employers about the law. Indeed, amid high levels of scrutiny, the key features of the law that recognized the work of younger children were reversed in 2018.

We exploit the timing of the changes in legislation and cross-individual variation in the exposure to such changes to empirically estimate the impacts of the law. Specifically, to account for unobserved characteristics of children that vary systematically with age, we employ a difference-in-discontinuity design based on a child's year and month of birth. Using a repeated cross-sectional household survey, we examine differences in work outcomes for children just above and below age thresholds issued by the law across the periods before the law was implemented (2012-2013), during the years in which the law was enforced (2014-2017), and after key components of the law that protected the rights of younger working children were

¹The law allowed children older than 12 to legally work for others and children between the ages of 10 and 12 to work as own-account (self-employed) workers. As detailed in Section 2.2, the law maintained the official minimum working age of 14 but introduced exceptions so that children as young as 10 could work legally. Thus, the law lowered the de-facto minimum working age from 14 to as young as 10.

reversed (2018-2019). This strategy allows other determinants of work to vary (smoothly) with age in months and accounts for any preexisting discontinuities in outcomes prior to 2014.

Even though the law legalized child work and entitled child workers to basic rights and protections, its enactment decreased the prevalence of child labor in terms of the likelihood and hours of work. We find that children under 14 (who were newly able to work legally) were nearly 4 percentage points less likely to work when the law was in effect (roughly 16% of the pre-law mean), relative to children above age 14 (who were always allowed to work legally and whose workers' rights were guaranteed prior to and following the law). These effects dissipate after 2018, when key components of the law were repealed. We find no evidence that the law shifted child labor across regulated and unregulated (prohibited) work or in terms of self versus external employment. Despite the decline in child employment, we find no effects on other measures of child time allocation (schooling and chores) or on household outcomes (adult labor supply and household income).

We find that the effects of the law were strongest in areas with higher probability of inspection by regulators, proxied by the distance to the closest regional offices of the MTEPS — the government agency in charge of conducting labor inspections. These results are robust to excluding the largest urban centers, which suggests that they do not simply capture urban-rural heterogeneity. They also suggests that enforcement, although imperfect, may have increased the perceived threat of inspection. This result is consistent with other studies that analyze how firms respond to regulations and tax-compliance efforts using distance to the regulator as a proxy for enforcement Almeida and Ronconi [2016], McKenzie and Seynabou Sakho [2010].

We analyze several potential mechanisms behind the declines in child work. On the labor-demand side, the law may have increased the relative costs of legally employing younger children among compliant firms, which may have reduced the demand for labor of young workers [Lazear, 1990, Autor et al., 2007]. Thus, one should expect the working conditions for young children to improve. To this end, we study the effects of the law on job characteristics of child workers: namely, job safety and pay, two job attributes specifically targeted by the law. For

job safety, we use two surveys that focused specifically on the nature of child work. We find that the law had no statistically significant impacts on the riskiness of child work or on injuries sustained while at work. We do observe non-statistically significant increases in wages among children that remained employed. However, as few children work for likely compliant, formal firms (3.2%), we believe that the increases in direct costs of complying with the law are unlikely to fully explain the overall decline in child work. Contrary to its key objective, the law does not appear to have improved the working conditions of child workers. Thus, there appears to be a net loss of worker welfare; some child workers lose their jobs and yet this does not lead to increased benefits for child workers who retain their jobs.

Alternatively, the enactment of the law and the high level of scrutiny around it may have induced avoidance behavior from both sides of the market. Among informal employers, the 2014 law may have increased the perceived threat of general labor inspections as both child labor and general labor inspections were carried out by the same government agency (MTEPS), incentivizing them to not hire children who were visible targets of the new legislation in order to remain “under the radar” of regulators. Among parents (households), the law may have intensified scrutiny surrounding child work and thus increased the risk of legal and social sanctions to parents of younger child workers, reducing the supply of child labor. Several pieces of evidence support this mechanism. First, the law’s effects are larger in areas that are closer to regional offices of the enforcement agency, where the threat of inspections and social stigma are likely to be higher. Second, the declines in employment due to the law are driven by declines in the probability of working outside home at fixed establishments, which are more visible and traceable by inspectors; in contrast, we find no changes in employment in less conspicuous and trackable modes, such as at work occurring within the home or in mobile locations. In addition, we find suggestive evidence of substitution from visible to less visible work locations among the children who remained employed. Finally, the fact that the effects dissipate after the law is reversed suggest that the avoidance behavior was driven by employers (as opposed to parents), as the threat of social sanctions and social stigma would have implied a more persistent pattern

of avoidance behavior.

Our results are robust to a battery of robustness checks. First, we show that they are not driven by standard concerns for difference-in-discontinuity designs, such as manipulation of the running variable, changes in sample composition and balance across age thresholds, bandwidth selection, and functional form specifications for the running variable. Second, we show that our results are robust to using an alternative difference-in-difference research design to relax the role of smoothness for identification imposed by our difference-in-discontinuity design. Third, we show that our results are not explained by changes in employment in areas with a higher presence of indigenous communities for which child employment (either for the family or the community) is often conceived as an integral part of engagement with the community and traditions.

Our results also provide novel insights to the literature evaluating the effects of child labor legislation. Previous studies have focused on the effects of child labor bans on work outcomes. We contribute to this literature in two ways. First, we leverage a unique policy change that instead of banning child labor, legally recognizes and regulates the work of children. While some studies find that child labor bans little overall effect on child work [Edmonds and Shrestha, 2012, Bargain and Boutin, 2021] or induce a decline in child work [Piza et al., 2023], our results suggest that legally recognizing the work of younger children does not increase child labor. Instead, our results are consistent with evidence of unintended consequences of child labor bans.² The declines in child labor disappear when the legislation is reversed — which essentially amounts to reinstating the ban on legal work under the age of 14. This finding is consistent with the effects to those in Bharadwaj et al. [2020] and Abman et al. [2023]; removing legal status and worker protections for younger children actually increases the likelihood that they work.

Second, previous studies of the impacts of child labor legislation have mostly focused on the effects on employment and provided little evidence on working conditions for children.

²Other studies explore the impacts of minimum working age laws in the U.S. in the early 20th century. For example, Moehling [1999] finds little effect of minimum ages laws; Manacorda [2006] demonstrates that though minimum age laws reduce work for targeted children, this is often compensated by an increase in the labor supply of siblings such that the overall effect on child work in the household is negligible.

Leveraging novel data, we contribute by providing new evidence on the effects (or lack thereof) of child labor legislation on job safety, a critical dimension of child work and oft-cited rationale for child labor legislation. The results demonstrate that recognizing and regulating child labor does not yield improvements in children’s working conditions and instead appears to have increased the perceived risk of labor inspections and thus the cost of hiring child workers, which ultimately affects child work in ways that can contradict policymakers’ intentions.

Finally, beyond the labor market for children, our results contribute to our understanding of how policy affects the labor market outcomes of vulnerable workers hired informally. In the context of sex workers, legalization and regulation can fail to improve worker safety and well-being [Gertler and Shah, 2011, Ito et al., 2018]; some evidence suggests that this could be due to workers’ reluctance to visibly identify themselves as sex workers by obtaining the necessary certification to work legally [Manian, 2021]. Our contribution is to document a novel channel through which regulation can hamper employment of targeted workers: avoidance behavior driven mostly by informal employers.

2.2 Child labor legislation in Bolivia

Child work is relatively common in Bolivia. From 2012 to 2013, roughly one in five children between the age of 10 and 14 worked despite being younger than the minimum working age of 14 years old.³ The conditions under which children work are also striking. Based on the 2008 Survey of Child Work (Encuesta Nacional sobre Trabajo Infantil, ENTI), more than 65% of child workers worked in occupations that are classified as hazardous by the International Labor Organization and more than one third of working children reported suffering an injury at work. These dramatic patterns were similar even among the 16.5% of children who work for their families.⁴ In comparison, roughly half of working children are engaged in hazardous work

³Authors’ calculations of weighted means based on the 2012-2013 Encuesta de Hogares. This definition does not include participation in household chores.

⁴Specifically, 63% of children working for their families are engaged in hazardous work while 31% reported suffering an injury at work. Authors’ calculations using the 2008 ENTI.

worldwide [International Labour Organization, 2021].

Despite consensus on the importance of protecting the integrity of children, Bolivia has experienced important tensions between policymakers and working children themselves. Setting and enforcing minimum working age requirements that align with compulsory schooling ages are popular policy guidelines recommended by international organizations. However, these policies are often criticized as being at odds with the reality of child work; many argue that child work is often necessary in the face of poverty and that policy should instead focus on regulating child work to ensure safe working conditions and the protection of child rights. In Bolivia, grassroots organizations such as the National Union of Working Children's (*Union Nacional de Niños, Niñas y Adolescentes Trabajadores de Bolivia*, UNATSBO) have been at the forefront of such policy suggestions, demanding the recognition of labor as an integral and unavoidable part of children's development.⁵ In part as a response to this tension, the Child and Adolescents Code of 2014 was implemented to legally recognize some forms of child labor and thus guarantee protections to working children. We describe the main changes induced by the law in the following sections 2.2.1 and 2.2.2.

2.2.1 Child labor legislation prior to 2014

Before 2014, two laws regulated the engagement of children in labor markets: the Child and Adolescents Code (law 206 of 1999), which provided general guidelines about the rights of youths, and the General Labor Law (law 224 of 1943), which regulates overall participation in labor markets.

Title VI of the 1999 Child and Adolescents Code describes the legal framework related to the protection of working children. There are three important dimensions for our analysis. First, the code set a minimum working age of 14 years old (Article 126). Second, the 1999 code put forth regulations for working children between the age of 14 to 18 but did not specify protections for younger children. Third, the code established that the work of adolescents (14

⁵See Chapter 4 in Unión de Niños Niñas y Adolescentes Trabajadores de Bolivia [2010].

years and older) was regulated by the General Labor Law of 1953. Thus, working adolescents were entitled with the same rights and obligations of adult workers.

Specifically, working children were to be paid at least the adult minimum wage and they were to be enrolled in the social security system by their employers. In addition, the 1999 code mandated that employers or parents (in the case of family businesses) offer flexible schedules to working adolescents so that they could attend school and that daily shifts not exceed 8 hours (not more than 40 hours per week). The 1999 code also prohibited child work in occupations deemed hazardous and those that potentially compromised the dignity of working children.⁶

2.2.2 Changes in legislation after 2014

We exploit the enactment of new child labor legislation in 2014 and its subsequent reversal in 2018 as sources of plausibly exogenous variation to estimate the impact of legalizing the work of younger children and increasing worker protections. Law No. 548 of 2014 addressed the general welfare and rights of children and expanded workplace protections to younger children. Specifically, it stated that its objective was “... to recognize, develop, and regulate the exercise of child and adolescent rights ...” (Article 1). Under these broad objectives, the new law changed preexisting child labor regulations in two core dimensions: exceptions that lowered the de facto minimum working age and expansions of worker protections to younger workers.

Table 2.1 summarizes the key changes induced by the law for each age group. The new law set a baseline minimum working age of 14 years, but it also allowed children aged 10 to 13 years to work legally in certain capacities and subject to obtaining work authorizations. Specifically, the new law permitted children aged 10 to 11 to work as self-employed workers⁷ and children aged 12 to 13 to work as both self-employed workers and to work for others. For both age groups, children were required to obtain work authorizations from local child protection

⁶Appendix Section B.2 provides a list of all forbidden activities under Articles 134-135 of Title VI of the 1999 code.

⁷Self-employed – or independent – work is defined by the law as work that is carried out by the child without any employer relationship. It is distinct from work for the family. Examples include street vending and washing vehicle windows at traffic lights.

offices (Defensoría de la Niñez y Adolescencia). This authorization required parental consent and a medical examination of applicants.

Table 2.1. Key Dimensions of Child Labor Legislation

	Before 2014 (Pre-Law)	2014-2018 (During Law)	After 2018 (Post-Reversal)
Age < 10	No legal work	No legal work ¹	No legal work ¹
10 ≤ Age < 12	No legal work	Legal to engage in independent work ²	No legal work ¹
12 ≤ Age < 14	No legal work	Legal to engage in independent work or work for others ² , with worker benefits and protections ³	No legal work ¹
Age ≥ 14	Legal to engage in independent work or work for others ² , with worker benefits and protections ³		

¹ Starting in 2014, children of all ages were allowed to engage in communal work – culturally valued activities taking place in indigenous, Afro-Bolivian, and intercultural communities – as long as it did not infringe on their rights and protections as guaranteed by law.

² In 2014, the list of permitted tasks and sectors for child work was revised to exclude agricultural work for an employer.

³ Prior to 2014, only children age 14 and over were entitled to the same workers’ rights as adults, including minimum wages and social security. After 2014, these rights were extended to working children age 12 and older and the benefits were expanded (for example, to include two paid study hours per day).

By recognizing the work of younger children, the new law also charged the government with regulating work and establishing protections for younger working children that were not accounted for in the previous law. The law explicitly stated “The State at all levels will guarantee the exercise or work performance of adolescents over fourteen (14) years of age, with the same rights enjoyed by adult workers. The protection and guarantees for working adolescents over

fourteen (14) years of age is extended to adolescents under fourteen (14) years of age” (Law 548, Article 130).⁸ Thus, beginning in 2014, working children aged 12 and 13 were entitled to the same benefits and entitlements of adult workers, such as minimum wages and social security. Additionally, the 2014 law required that employers give child employees (age 12 to 17) flexible schedules and at least two paid hours per day to perform their schooling obligations.⁹ It also set a maximum of 30 hours of work per week (6 hours per day) for children between 10 and 14 years old. As was the case prior to 2014, children 14 to 18 years old were allowed to work up to 40 hours per week, with a maximum of 8 hours per day. Finally, the list of prohibited tasks and jobs was updated to include agricultural work occurring outside of family and communal work. Communal work – culturally valued activities taking place in indigenous, Afro-Bolivian, and intercultural communities – was allowed and was subject to separate rules and procedures set and implemented by indigenous jurisdictions. However, children that engaged in family and communal work were still granted the same rights and protections as all child workers.¹⁰

Amid intense debate and scrutiny, some key articles of the law — namely those granting children below the age of 14 the ability to work legally and benefit from the same protections and guarantees as older workers — were reversed in 2018. The 2018 amendment to Article 130 explicitly states the government’s duty to ensure the rights of workers between the ages of 14 to 18 years old and does not establish rights of younger working children, in contrast to the 2014 law. Additionally, the government repealed paragraph IV of Article 132, which regulated weekly work hours for children between 10 and 14 years old. Thus, starting in 2018, the government no longer issued or implemented a program for protecting the rights of working children under the age of 14 [Defensoría del Pueblo, 2022].

⁸Authors’ translation of original document in Spanish.

⁹In the case of self-employed children, the 2014 law required that parents ensure that children can attend school even while working.

¹⁰One unfortunately common example of prohibited exploitative family work relates to children being sent by their parents to beg in the streets. In some cases, parents were accused of family violence for forcing their children to beg [Los Tiempos, 2013b].


2.2.3 Enforcement and awareness

The 2014 law tasked the regional offices of the Ministry of Labor and Social Protection (Ministerio de Trabajo, Empleo y Protección Social, MTEPS) with carrying out inspections and permanent supervision of workplaces to ensure that employers were complying with the regulations under the 2014 law (Article 139). Even prior to the 2014 law, MTEPS offices were in charge of verifying the ownership of valid business registrations, conducting general labor and technical inspections, and carrying out inspections related to preventing forced labor.¹¹ A key component of enforcement was age verification for children. In Bolivia, age verification is relatively straightforward and feasible, due to near universal birth registration and widespread identity cards. According to the 2012 Bolivian Census, 99% of children in the age range of 9-15 years old were registered at birth at the civil registry, and 72.5% of them owned ID cards. Moreover, ID cards are required to obtain government benefits from Bolivia's conditional cash transfer program for school-age children.

There are 25 regional Ministry of Labor and Social Protection offices located in the most populated municipalities of the country (see Appendix Figure B.1). Using data from annual MTEPS reports, Figure 2.1 shows that child labor inspections increased considerably in 2014 and rose thereafter. There were on average around 300 child labor-specific inspections per year conducted during the period following the law's enactment. The total number of inspections (labor and technical) conducted by the MTEPS also increased after 2014, suggesting that the increase in child labor inspections did not crowd out – and perhaps even crowded in – other inspections conducted by the MTEPS.

If any party were found to be in violation of the rights and protections under the law, the MTEPS would turn the case over to the Defensoría de la Niñez y Adolescencia (DNA) for legal restitution. Inspections carried legitimate consequences for employers; in 2018, 17% of child

¹¹Labor inspections verify compliance with national regulations, including being part of the mandatory employer registry (Registro Obligatorio de Empleadores), contributions to social security and health insurance, and compliance with worker protections established in the Labor Law. Technical inspections verify that work facilities comply with safety and sanitary standards.



Note: Data on inspections is obtained from the annual reports by the Ministry of Labor [Ministerio de Trabajo, Empleo y Previsión Social, 2015-2018]. Child labor inspections prior to 2015 are as reported in the US Department of Labor reports [U.S. Department of Labor, 2011-2019].

Figure 2.1. Ministry of Labor Inspections over Time

labor inspections were turned over to the DNAs for resolution [Ministerio de Trabajo, Empleo y Previsión Social, 2015-2018]. Under the 2014 Law, the DNA was allowed to impose penalties such as warnings and reprimands, fines, the removal of children from work, and temporary suspension of business activities.¹² Parents in violation of the code (for example, as employers of their children in family work, but also as guardians of their children more broadly) were also subject to measures ranging from warnings to required attendance of courses and programs and (at the extreme) separation from their children. In the case of repeat offenders, the DNA had the authority to send the proceedings into criminal court.

The threat of an inspection by the MTEPS office is likely to affect employers' compliance with new regulations and their demand for child labor. Formal firms may increase worker protections to avoid sanctions or reduce the demand for younger child workers as they become

¹²As stated in Article 169 of Law 548 and Article 219 of the 1999 code.

relatively more expensive to hire legally. In the case of informal firms—the larger sector in the economy¹³ — the threat of inspection may operate through an additional channel: firms may decide to employ fewer young children in order to avoid being inspected by the Ministry of Labor and continue operating informally. The fact that the MTPES was in charge of both the child labor and regular registration inspections may have increased the perceived risk of inspection among firms hiring young children. A recent survey of Bolivian firms found that the overwhelming majority of firms — even among small and micro-enterprises — perceived costs associated with labor regulations as directly influencing their hiring decisions, suggesting that there is an advantage to remaining “under the radar” of labor inspectors [Muriel and Ferrufino, 2012].¹⁴ Relatedly, prior work has found that firms tend to resist formalization, even when provided information about the registration process and when registration fees are waived, but that firms respond to the increased likelihood of inspections [De Andrade et al., 2016].

The initial enactment of the law was very controversial and highly scrutinized by NGOs, international organizations, and authorities. Several press articles highlight the public support of the legislation by the then-president [Pagina Siete, 2013, Los Tiempos, 2013a], which may have amplified awareness about the policy change.¹⁵ In Appendix Figure B.2 we track articles that mention the 2014 law over time across national and regional Bolivian newspapers. There are clear spikes in the number of published articles around the time that the initial 2014 law was implemented and in the years in which the law amendment was announced and eventually implemented (2018), suggesting that the general public was aware of the policy changes.¹⁶ We also observe coverage of the law in the intervening years – particularly in 2016 and 2017 – indicating that the issue continued to be relevant throughout the period. In addition, the enactment

¹³Informal firms account for almost 80% of employment and 62% of GDP in Bolivia [Elgin et al., 2021].

¹⁴This behavioral response of firms to regulation has been discussed in other settings (see for example, Hsieh and Olken [2014], Tybout [2014]).

¹⁵There is a growing literature documenting how information provided by political leaders can modify citizen’s attitudes and behavior through different media [Ajzenman et al., 2020, Pedemonte, 2020, Jetter and Molina, 2022].

¹⁶Appendix Figure B.2 suggests potential for anticipatory effects because the reversal was announced in February 2018 but not implemented until December 2018. However, this does not represent an issue for our analysis because the survey data we use (described further in Section 2.3) is collected in November and December of each year and we treat 2018 as a post-reversal year.

of the 2014 law was coupled with workshops on worker's rights and protections, delivered by the MTEPS and targeted to employers and children. Over 11,000 workers and employers attended to these child labor workshops between 2015 and 2018, according to MTEPS Annual Reports [Ministerio de Trabajo, Empleo y Previsión Social, 2015-2018].

Throughout the paper, we interpret the enactment of the 2014 law as both a legal recognition of the work of younger children and as an expansion of worker rights for this group. The descriptive evidence on enforcement and awareness suggests that these legislative changes were perceived as important and relevant for firms and families. Accordingly, we interpret the 2018 reversal as an abrupt cessation of both legal recognition and rights for younger working children. Though the implementation and reversal of the law may have led to general changes in overall enforcement on labor regulation (e.g., increased attention paid to safety and protection measures), our identification strategy focuses on the changes for younger children *relative* to those of older children, netting out these general changes over time.

2.3 Data

To measure the effects of the policy change on employment and work hours, we leverage data corresponding to 8 waves of Bolivia's annual household surveys (*Encuesta de Hogares*, henceforth referred to as the household data). Each survey wave contains data from a nationally representative sample of households in Bolivia. We pool survey waves to construct a repeated cross-section covering two years before the policy change (2012 and 2013), 4 post-law years (2014-2017), and 2 post-reversal years (2018-2019). We exclude data preceding 2012 to minimize the potential effects of the rollout of Bolivia's conditional cash transfer (CCT) program targeted at school-age children.¹⁷

As discussed in Section 2.2.2, exposure to different dimensions of the 2014 law (and its

¹⁷The *Bono Juancito Pinto* program was initially delivered to children enrolled in grades 1 to 5 in 2006 and expanded to include children in 8th grade in 2009. In 2012, it was announced that children in 9th grade would also be covered. See Vera-Cossio [2021] for details about the policy. We discuss a further expansion of the program to older children in Section 2.4.

later amendment) is a function of age. Our dataset includes the exact birth date of each household member, which enables us to calculate age at the time of the survey. We compute the number of months elapsed between a child's birth date and the month in which fieldwork of each survey started (typically, November of each year). We then normalize age in months relative to the cutoff of interest—age 10, 12 and 14.

Work is measured by an indicator of whether a child worked at least one hour during the week preceding the interview.¹⁸ We also compute weekly work hours and construct an indicator for overtime work (defined by the 2014 law as working more than 30 hours for children under 14). Further, we separately measure work for self (self-employment), work for others (family or external employer), prohibited work (employment in activities that are prohibited under the law for all children under age 18, such as mining), and allowed activities (those that are allowed and regulated under the law).¹⁹ Finally, we also measure labor force participation, which includes both those who are working and those who are unemployed but actively searching for jobs. We examine the role of enforcement of the law using data on the locations of regional MTEPS offices, which we describe in more detail in Section 2.5.2 and Appendix B.4.

To better understand the mechanisms behind the main results, we use information from the household survey on job attributes (namely, wages and the size of firms children work for). Additionally, we leverage detailed information on the locations where child work takes place (e.g., at home, as a street vendor, or at an establishment with a fixed location) as well as involvement with risky tasks (including, among others, working under extreme temperatures or working in an area exposed to fire, flames, or contaminated dirt and dust) and injuries at work (such as skin injuries, fractures, and respiratory complications) from the 2016 Survey of Children and Adolescents (Encuesta Niño, Niña y Adolescente, ENNA) and the 2008 survey on working children (Encuesta Nacional sobre Trabajo Infantil (ENTI) 2008).²⁰

¹⁸This definition does not include unpaid participation in household chores.

¹⁹See Appendix Sections B.2 and B.3 for a full list of prohibited activities and more detailed variable definitions.

²⁰The sampling frame differs across the two surveys; while the 2016 ENNA is nationally representative, the 2008 ENTI focuses on children who are likely to work. Therefore, in order to pool the two datasets, we reweight the observations in each survey. We discuss this reweighting method in more detail in Section 2.6.1. We also give more

Panel A of Table 2.2 reports summary statistics for children age 9 to 15 years old during the pre-law period (2012-2013). Before the policy change, 14% of children in the sample worked. Among working children, the average number of weekly work hours is 21 and over 19% of working children worked more than 30 hours per week. Self-employment is somewhat rare; less than 2% of working children worked for themselves prior to the 2014 law. Work for others is largely made up of work for a family employer (88%). However, work for a family employer and work for an external employer are similar along many critical dimensions. For example, most employers operate informal firms,²¹ regardless of whether they are family operated or not (see Panel A of Table 2.3); the median firm size (4 workers) is the same across family employers and non-family employers; virtually all jobs are performed outside the household (97%) even in family-operated firms. Family work is largely driven by agriculture and retail, while work for others is more diversified, although still dominated by retail and agriculture. Children tend to work outside home, mostly in fixed establishments, regardless of whether their employer is a household member or not (see Panel A of Table 2.3), although children working for external employers are more likely to work in mobile locations.

Panel B of Table 2.2 shows that roughly 56% of working children are engaged in risky activities and 34% of working children report having experienced a job-related injury in 2008. Children's exposure to risk and injury are high in both work for family and work for employers (Panel B of Table 2.3).

detailed descriptions of variables in Appendix Section B.3.

²¹Formality is defined by whether the firm is formally registered with the national tax authority.

Table 2.2. Descriptive Statistics (Pre-Law)

Panel A: Household Data					
	All Children Ages 10-15 (1)	Working Children Ages 10-15 (2)	All Children Ages 10-11 (3)	All Children Ages 12-13 (4)	All Children Ages 14-15 (5)
<i>Household & Child Characteristics</i>					
HH Head Years of Schooling	8.603	5.685	8.624	8.589	8.595
HH Head is Male	.787	.789	.798	.792	.76
HH Head Age	44.401	45.266	43.403	44.707	45.486
HH Head is Indigenous	.357	.619	.359	.349	.368
Household Size	5.593	6.019	5.635	5.59	5.529
Child is Male	.501	.496	.499	.511	.489
<i>Child Work & Schooling Outcomes</i>					
Any work	.152	-	.125	.158	.19
Hours worked	3.325	21.807	2.306	3.456	4.788
Work for self	.003	.019	.001	.003	.006
Work for others	.15	.981	.123	.154	.184
Work for external employer	.017	.11	.004	.016	.039
Work for family employer	.133	.871	.119	.139	.145
Prohibited work	.006	.042	.002	.006	.014
Allowed work	.146	.958	.122	.152	.175
Work \geq 30 hrs/week	.03	.199	.015	.031	.055
Attends school	.97	.908	.985	.97	.946
Observations	7410	1130	2698	3108	1604
Panel B: Job Attributes (Household Survey)					
	Working Children (1)	Ages 10-11 (2)	Ages 12-13 (3)	Ages 14-15 (4)	
Firm size (median)	4	4	4	4	
Hourly wage (Bolivianos)	7.038	5.44	7.123	7.295	
Firm pays taxes	.035	.021	.046	.034	
Works Outside of Home in Fixed Location	.859	.113	.137	.15	
Works Outside of Home in Mobile Location	.112	.01	.016	.032	
Works at Home	.028	.002	.005	.007	
Observations	1116	336	480	300	
Panel C: Job Attributes (Child Labor Survey)					
	All Children (1)	Working Children (2)	Ages 10-11 (3)	Ages 12-13 (4)	Ages 14-15 (5)
Risk at work	.294	.545	.246	.314	.346
Injured at work	.178	.324	.163	.184	.194
Observations	3477	1749	1343	1389	745

Notes: The table shows the mean of the variables, except for firm size, where the median is displayed. Definitions of the variables appear in Appendix B.3. The list of prohibited tasks appears in Appendix B.2. The sample in both panels includes children from ages 10 to 15. The survey years are 2012-2013 in Panels A and B, and 2008 in Panel C. Observations of the child labor survey are reweighted using the method described in Section 6.1.

Table 2.3. Descriptive Statistics by Employer Type

Panel A: Household Data			
	Work for External Employer (1)	Work for Family Employer (2)	P-value Diff. (3)
Firm size (median)	4	4	.2092
Hourly wage (Bolivianos)	6.291	18.557	0
Formal Firm	.098	.026	.0001
Works Outside of Home in Fixed Location	.64	.899	0
Works Outside of Home in Mobile Location	.36	.071	0
Works at Home	0	.03	.0493
Sector			
Agriculture	.144	.772	0
Sales and retail	.232	.101	0
Other	.624	.127	0
Observations	113	1094	

Panel B: Child Labor Survey Data			
	Work for External Employer (1)	Work for Family Employer (2)	P-value Diff. (3)
Risk at work	.679	.537	.0001
Injured at work	.447	.314	.0006
Observations	186	1741	

Notes: The table shows the mean of the variables, except for firm size, where the median is displayed. Definitions of the variables appear in Appendix B.3. The sample in both panels includes children from ages 9 to 15. The survey years are 2012-2013 in Panel A, and 2008 in Panel B. Observations of the child labor survey are reweighted using the method described in Section 6.1. The third column shows the p-values of the differences between columns 1 and 2.

2.4 Empirical approach

2.4.1 Identification

To identify the causal effects of the exposure to the law, we exploit two sources of variation. First, under the 2014 law, whether and which type of jobs children were allowed to work changed discontinuously at three age thresholds: 10, 12, and 14. Second, we exploit the variation in the timing of the law and its reversal to net out preexisting differences in outcomes across children of different age groups potentially related to the pre-2014 minimum working age of 14.

One key concern is that time varying shocks can differentially affect work outcomes of children based on their age, as there is a steep age-gradient in work probability (see Appendix Figure B.3); for example, 17-year-olds are more than twice as likely to work as 10 year-olds and the probability of working continuously increases by age in months. Thus, it is likely that an aggregate shock to labor markets disproportionately affects the children that are more likely to work. To address this concern, we propose an empirical design that exploits the discontinuous changes in exposure to the law at each age threshold, while allowing for a (continuous) age gradient in outcomes. This strategy compares the work outcomes of children who – based on their age as of data collection – just became eligible to work to the outcomes of children who were only months away from being eligible under the law.

We combine identification at thresholds with temporal variation in the enforcement of the law to account for any preexisting differences in work outcomes that predated the law’s implementation. By relying on *local* comparisons around age thresholds, our empirical strategy helps control for potential time varying shocks with differential effects based on age.²²

²²Imbens and Lemieux [2008] recommend using local linear regressions using observations within a narrow bandwidth of a threshold, as opposed to higher-order polynomials on a wider bandwidth as such estimates tend to be sensitive to observations away from the cutoff. Relative to other studies that focus on a broad bandwidths around the cut-off point (for example, Edmonds and Shrestha [2012] and Edmonds [2014] include children between 8 and 14 years of age in their analysis), our empirical approach exploits information on the month and year of age and survey field work dates to estimate effects around a narrower bandwidth (12 months) using a more granular running variable (age in months). We believe this narrow bandwidth lends internal validity to our estimates by comparing

There are several limitations to this approach. First, it enables us to estimate only short-term responses around a narrow time window after children change exposure status. Specifically, it does not allow us to identify effects of the law on children far away from the age threshold, i.e. children much younger or older than 14. Second, our design estimates a composite policy parameter that reflects the combined effects of different aspects of the policy (such as changes to the legal working age, the introduction of rights and safety regulations for child workers, enhanced resources for monitoring child work). Given that the 2014 Bolivian policy was unlike any prior (or subsequent) child labor policy in recent history, the point estimates we obtain are not directly comparable to those from other contexts. Finally, we are analyzing the context of a developing country which may face substantial barriers to enforcement of labor regulation; Edmonds and Shrestha [2012] argue that lack of enforcement may explain why they do not find evidence of substantial effects of minimum working age restrictions on employment across a wide set of countries. However, our empirical strategy makes two improvements relative to previous (solely) age-based designs. First, we exploit time-series variation in enforcement of the law, by studying changes in work probabilities before, during, and after the enforcement of the law. Second, we exploit spatial variation in the potential enforcement of the law to help validate our estimates. We describe this strategy in detail in Section 2.5.2. Nonetheless, we acknowledge that we do not have data on enforcement by age and, like most other studies, cannot show direct evidence on enforcement (or lack thereof).

More formally, we use a difference-in-discontinuity specification. We model the effect of being exposed to the law on outcome $Y_{i,t}$ corresponding to child i observed in survey wave t as:

$$\begin{aligned}
Y_{i,t} = & \beta_0 + \beta_1 T_i \times Law_t + \beta_2 T_i \times Reversal_t + \beta_3 T_i + \theta_1 (Age_{i,t} - c) \\
& + \theta_2 T_i \times (Age_{i,t} - c) + \gamma x_{i,t} + \delta_{d,t} + \epsilon_{i,t}
\end{aligned} \tag{2.1}$$

where $Y_{i,t}$ is a work outcome for child i in survey year t , $Age_{i,t}$ is the age of child i in months

children who should otherwise be similar, apart from being targeted by the law.

at the beginning of the relevant recall period (which differs by outcome) for survey wave t .²³ We define age relative to the start of the survey recall period because we need to capture the age eligibility for legal work (and thus worker protections) that is relevant for the work outcomes reported in the survey. c is the relevant cutoff age related to the key policy changes induced by the new law (at ages 10, 12, and 14). T_i is an indicator of whether child i is exposed to the policy change associated to each cutoff. In the case of the cutoff at 14 years old, exposure to the law (T_i) is an indicator of whether a child is *younger* than 14 years old. This is because the 2014 law newly allowed children under age 14 to work and do so with protections and benefits; children aged 14 and older were legally allowed to work even under the preexisting law. For the 10- and 12-year-old cutoff, T_i is defined as an indicator of whether a child is 10 years old or older and 12 years old or older, respectively. We define the treatment indicators in this way because at the age 10 threshold, the 2014 law grants children just above the threshold the ability to work legally as self-employed and, at the age 12 threshold, the 2014 law further allows them to work for others. With these definitions, the interpretation of T_i is consistent across all thresholds, in that all treated children have newly expanded working rights under the 2014 law relative to control children. Law_t is an indicator identifying the years in which the law was enforced (2014-2017), while $Reversal_t$ identifies the years after the reversal of the law (2018-2019).²⁴ Finally, $\varepsilon_{i,t}$ is an error term.

For all thresholds, the parameters of interest are β_1 and β_2 , which capture changes in work outcomes of children marginally exposed to each dimension of the law, relative to those just on the control side, between the periods in which the law was enforced and repealed with respect to the pre-law period. While virtually all children have birth registrations and the large majority have ID cards – making age verification by employers straightforward – we recognize that it may be difficult for firms to effectively distinguish between children above and below 14

²³For example, the recall period for employment is the week prior to the survey, so $Age_{i,t}$ reflects the age of the child at the beginning of the prior week when considering employment outcomes.

²⁴The government announced the reversal of the law in mid-2018 and the household surveys are conducted at the end of the year, so we consider 2018 as a post-reversal year.

(especially near to the threshold), introducing fuzziness in the difference-in-discontinuity. Thus we regard β_1 and β_2 as reflecting intent-to-treat effects of the law and reversal.

We also include a vector of demographic household and child characteristics that are unlikely to vary due to the program ($x_{i,t}$).²⁵ These include household head characteristics such as schooling, gender, age, and ethnicity; household characteristics such as number of children in the household in following age categories: 0-6, 7-9, 10-13, and 14-17 and number of adult men and women; whether the household is located in an urban area; and the child's gender. We also include a full set of departamento-by-year fixed effects ($\delta_{d,t}$) to flexibly account for regional time-varying shocks.²⁶

The coverage of Bolivia's flagship CCT program was expanded in 2014 to include children enrolled in grades 9 to 12 (regardless of age). Given that some children in grade 9 are 14 years old, we also control for grade-for-age fixed effects and their interactions with a post-2014 indicator when we estimate equation 2.1 for the 14-year-old cutoff. This helps account for the potential impacts of the CCT on child labor that may also differ for children above and below age 14.²⁷

To account for the age gradient in work outcomes, we use a linear specification of the running variable and allow for different slopes on either side of the cutoff. We show that our results are unchanged when we instead use a second-order polynomial and when we allow the slopes to vary before and after the policy change in Section 2.5.3. We estimate equation (2.1) using triangular kernels that assign a higher weight to observations closer to the eligibility cutoff and conduct inference using standard errors clustered at the age-in-months level (the level at which treatment varies) to account for correlated shocks within age groups.

We estimate equation (2.1) using a twelve-month bandwidth from each age cutoff. This

²⁵We include covariates to increase precision, though we show that our results are robust to specifications without controls in Section 2.5.3.

²⁶Departamento is an administrative/geographic unit roughly comparable to a U.S. state.

²⁷Controlling for CCT exposure is not necessary for younger children (those around the 10- and 12-year-old cutoffs) because by 2009 all children in these age groups were eligible to receive the CCT (regardless of being above or below the thresholds defined in the 2014 law).

bandwidth is narrower than the mean squared error (MSE) optimal bandwidth proposed by Imbens and Kalyanaraman [2012], which ranges from 13 to 25 months for all our main outcomes. However, selecting a narrower bandwidth in our setting avoids classifying observations as part of the treatment group when we analyze one cutoff and as part of the control group in a different cutoff.²⁸ Thus, we compare 9-year-old to 10-year-old children around the 10-year-old cutoff, 11-year-old to 12-year-old children around the 12-year-old cutoff, and 13- to 14-year-old children around the 14-year-old cutoff. In Section 2.5.3, we show that our results are robust to using narrower and wider bandwidths.

2.4.2 Threats to identification

Manipulation. The validity of our diff-in-disc design requires that individuals cannot perfectly manipulate the assignment variable, which in our setting is the age (in months) at the time of data collection. There are two reasons why manipulation is unlikely. First, we study the impact of a law using data that is regularly collected by the government and which was not designed or framed as a tool to measure the impacts of the law; ex-ante, there was no incentive to manipulate child age in order to appear compliant in our analysis. Second, even though age heaping is common, interviewees are asked for the birth date of each household member as opposed to their age.

As we rely on self-reported data, a similar threat to validity is that becoming eligible to work under the law may have caused differential survey response rates of children around each cutoff. Specifically, the concern would be that survey respondents are less likely to truthfully report the age of children younger than 14 years old. Appendix Figure B.4 plots the distribution of observations around the cutoffs, focusing on children with birth dates within a year of each cutoff (the bandwidth of our baseline specifications). It shows no evidence of discontinuous changes at the cutoff either when we compare pre- to post-law periods and post-law to reversal

²⁸For example, a child who is 11.5 years old would be in the treatment group relative to the 10-year-old cutoff, but the same child would be in the control group relative to the 12-year-old cutoff.

periods.

We more formally test for discontinuous changes at each cutoff in the *differences* in densities between the pre-and post-law periods, and the post-law and reversal periods—i.e., the difference-in-discontinuities analogue of the traditional manipulation test in regression discontinuity designs. Appendix Figures B.5 - B.7 plot the differences in densities between periods as a function of the running variable, for each cutoff using household survey data.²⁹ We find no evidence of discontinuous changes in the differences in densities across all cutoffs at a 5% confidence level. In the case of the 14-year old cutoff, there appears to be a significant difference in densities between the post-law and the law reversal period at 10% (p-value=0.098). Such difference is small (-0.0007) and implies that, per 10,000 observations, there were 7 fewer observations corresponding to children just above 14 years old. In Appendix Section 2.5.3, we show that our results are robust to excluding observations very close to the cutoff (and where this small discontinuity could be relevant). In addition, we discuss additional checks for measurement error in Appendix Section 2.5.3.

Changes in sample composition and balance. We test for changes in demographic characteristics around the cutoff before and after the policy change. For this, we estimate (2.1) using demographic characteristics as dependent variables. Appendix Table B.1 shows that, at a 5% significance level, there are no differences across each cutoff. While 2 out of 18 differences are significant at 10% level for the household data, these differences do not reflect a systematic pattern across cutoffs. In addition, for each cutoff, we are unable to reject the null hypothesis that the coefficients in each column are jointly zero.

²⁹Specifically, we follow Grembi et al. [2016] and, for each bin around each cutoff, compute the differences in densities before and after the policy change. We then fit a linear polynomial at each side of the cutoff and compute the difference at each cutoff point.

2.5 Effects of the 2014 Law


2.5.1 Effects of the law on the prevalence and sector of child work

We begin by discussing graphical evidence of the impacts of the law. We focus on the impacts around the 14-year-old threshold, which speak to the combined effects of legalizing and regulating both self-employment and work for others, because there is a substantially higher rate of working children around this cutoff. Figure 2.2 reports flexible difference-in-discontinuity (i.e., event study-style) estimates of the effect of the law around the 14-year-old cutoff using a variation of equation (2.1) that allows the effects of the law to vary over time by grouping observations in two-year bins.³⁰ The work probabilities of 13-year-old (treated) children—whose work was newly regulated by the 2014 law—decline with respect to that of 14-year-old (control) children after 2014. These differences disappear after the law was reversed (2018). Overall, the results suggest that the 2014 law reduced employment for children around the 14-year-old cutoff. We do not observe substantial differences for children around the 12- and 10-year-old cutoffs in Appendix Figure B.8.

Figure 2.3 plots work probabilities as a function of age (in months) relative to the 14-year-old cutoff before, during, and after the implementation of the 2014 law change. During the pre-law period, there is no discontinuous change on work outcomes around the cutoff. This suggests that the preexisting minimum working age was not a binding constraint to child labor. In contrast, we find a discontinuous change around the cutoff after the policy change. Relative to 14-year-old (control) children, marginally younger (treated) children were less likely to work while the 2014 law was in effect.³¹ This difference disappears after the key components of the law recognizing and regulating the work of younger children are reversed in 2018. For children around the 12- and 10-year-old cutoffs (for whom child work is less common), we observe no

³⁰We group observations in two-year bins to gain precision amid the reduced number of observations per survey wave.

³¹We discuss the potential for our estimates to capture a relative *increase* in work for those over 14 rather than a decrease for those under 14 below.




Note: The figure reports changes in work probabilities for 13-year-olds relative to 14-year-olds over time (grouped in two-year bins), with respect to the years preceding the policy change (2012-2013). The specification includes linear splines of the running variable, defined as the difference between the cutoff age and age a week before the survey date in months. Control variables: CCT eligibility indicator, household head characteristics (schooling, gender, age, and indigenous indicator), gender, number of children in the following age categories: 0-6, 7-9, 10-13, and 14-17, number of adult men and women, an urban dummy, and departamento by year fixed effects. The sample includes 2012-2019. The 95% confidence intervals are based on standard errors clustered at the age in months level.

Figure 2.2. Changes in Work Probability relative to Pre-law Periods at the 14-Year-Old Cutoff

discontinuities around the cutoffs during the implementation of the law (see Appendix Figures B.9 and B.10).

We now turn to the regression-based evidence. Table 2.4 reports the effect of the law on work outcomes around the 14-year-old cutoff. We find that the probability of work declines by 3.94 percentage points for 13-year-old children (a 16% decline relative to 14-year-old children; see column 1). Hours of work fall by about an hour per day, averaged across all children (including non-workers). These effects appear to be driven by a decrease in the probability of



The running variable is the difference between age in months and the age cutoff a week before the survey date. The pre sample includes 2012-2013, the post sample includes 2014-2017, and the reversal sample includes 2018-2019. We use a triangular kernel.

Figure 2.3. Work Probabilities at the 14-Year-Old Cutoff (Before, During, and After the Law)

work for others (3.9 percentage points, statistically significant at the 5% level; see column 4) as opposed to self-employment (0.2 percentage points, not statistically significant; see column 3). The decline in work is particularly pronounced in occupations that are legally allowed and regulated under the 2014 law (4.41 percentage points, statistically significant at the 5% level; see column 6). This decline does not coincide with a corresponding increase in work in prohibited occupations (column 5), suggesting that there was no reallocation of child labor across types of work. Finally, these declines in employment translate into similar declines in labor force participation (column 7). We discuss potential mechanisms in detail in section 2.6.

The coefficients associated with the periods following the 2018 reversal of key protections for younger workers under the law validate our empirical approach. Relative to pre-implementation period, there are no substantial differences between marginally exposed and unexposed children when the key protections regulating the work of children under the age of 14 are no longer enforced. The magnitudes of the coefficients associated with the post-reversal period are small and suggest that the changes in work outcomes induced by the enactment of the law fully dissipate after the reversal.

Interestingly, the estimated effects of the reversal of the law (i.e., the removal of legal

Table 2.4. Difference in Discontinuity Effects of the Law on the Work Probabilities, Hours, and Occupation for the 14-Year-Old Cutoff

	Any Work (1)	Hours Worked (2)	Work for Self (3)	Work for Others (4)	Prohibited Work (5)	Allowed Work (6)	Labor Force Participation (7)
Post Law $\times 1\{\text{Age} < 14\}$	-0.039** (0.017)	-0.969* (0.526)	-0.002 (0.004)	-0.037** (0.017)	0.004 (0.006)	-0.043*** (0.015)	-0.040** (0.017)
Post Reversal $\times 1\{\text{Age} < 14\}$	-0.000 (0.019)	0.508 (0.562)	-0.000 (0.005)	-0.000 (0.019)	0.018 (0.012)	-0.019 (0.018)	0.002 (0.019)
Obs.	11991	11991	11991	11991	11991	11991	11991
Mean	0.180	4.397	0.00490	0.175	0.0114	0.169	0.185

Notes: Age in months by year clustered standard errors in parentheses. Significance levels denoted by: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Control variables: CCT eligibility indicator, household head characteristics (schooling, gender, age, and indigenous indicator), gender, number of children in the following age categories: 0-6, 7-9, 10-13, and 14-17, number of adult men and women, an urban dummy, and departamento by year fixed effects. We also include linear splines of the running variable, defined as the difference between the cutoff age and age a week before the survey date in months. We use a bandwidth of 12 months and a triangular kernel. Survey years: 2012-2017.

work status and worker rights) are similar in magnitude to those found in Bharadwaj et al. [2020], who study the effects of a child labor ban in India. We find a 21% increase in the probability of working for those under 14 relative to the pre-reversal average work probability for 13-year-olds³², whereas Bharadwaj et al. [2020] find that the ban results in a 22% increase in work for children under 14 relative to the pre-ban mean. They are also qualitatively similar to estimates of *increases* in child labor when regional trade agreements include child labor bans and *declines* in child labor when regional trade agreements do not include child-labor bans [Abman et al., 2023]. In contrast, our results suggest different effects from a Brazilian law that increased the legal working age from 14 to 16; studies of the Brazilian law found no overall effects [Bargain and Boutin, 2021] or declines in child work [Piza et al., 2023].

We interpret the estimates in Table 2.4 as indicating that the 2014 legislation decreased work for children under the age of 14 (the age group targeted by the law). However, a potential alternative interpretation of our results is that the law simply reallocated work from those under 14 to those over 14, essentially increasing the work of 14-year-olds. Indeed, we observe an apparent increase in the probability of work for 14-year-olds in Figure 2.3 during the law's

³²In column 1, the difference between the Post-law and Post-reversal coefficients suggest that the reversal increased the work probability of 13-year-olds by 3.5 percentage points, relative to the periods in which the law was enforced.

implementation. To show that this increase reflects general trends in the Bolivian labor market (and thus is still an appropriate counterfactual for 13-year-old workers), we first show that the trends in work probabilities for 14 and 15 yearolds are very similar to those of older children (age 16-17). This suggests that the increase in employment among 14-year-olds during the policy implementation period likely reflect general aggregate economic fluctuations. In contrast, the figure also shows that work probabilities among 13-year-olds diverge from the general trend during the period when the policy was implemented. In Appendix Table B.2, we more formally test the extent to which the changes in work probabilities during the implementation of the law are explained by an increase in working probabilities among 14-year-old children by comparing changes in work outcomes before and after the implementation of the law among 14-year-olds and 15-year-olds. If the effects were driven by increased work by 14 year-olds – for example, if the 2014 law simply made clear that age 14 was the acceptable age for work – we would expect to find an increase in work for this group relative to 15-year-olds during the implementation of the law. Column 2 of Appendix Table B.2 shows no evidence that this is the case. In column 3, we allow for more general substitution of older child workers for those under 13 by comparing those age 14-16 to those age 17-18 (around the time children exit secondary school). Again, we find no significant increases in work of children age 14-16, suggesting that a reallocation of work from 13-year-old children to older children does not explain our results.

Appendix Table B.3 corroborates the results from the graphical evidence for younger children; that the law had no statistically discernible effect on the work of 12- and 10-year-old children, respectively. The new law enabled both 11- and 12-year-old children to work; however, only those 12 or older could work for others, subject to obtaining a work permit. Panel A shows that the point estimate of the effect on the likelihood of work for 12-year-old children (column 1) is negative, though not significant at conventional levels. Similarly, we find no statistically significant effects of the 2014 law on work probabilities at the 10-year-old cutoff (Panel B, column 1). We also find that the law does not lead to any changes in the type of work that 10- and 12-year-olds engage in, either in terms of sector of work (allowed versus prohibited), overtime

work, self-employment or work for others.

We also examine the impact of the law on schooling but find no statistically significant effects (see column 1 of Appendix Table B.4).³³ One explanation is that the school day in Bolivia is limited to 4 hours which allows children to combine work and schooling; this aligns with the observation that the overwhelming majority of children in the sample attend school (for example, 93.7% of 13-year-olds attend school). Thus, even if the law had decreased child work (as our results around the 14-year-old cutoff suggest), we expect to find little impacts on school attendance. This finding is bolstered by qualitative evidence that finds that for many working children, work and study are complements; and in many cases, work provides the means to pay for schooling-related expenses [Defensoría del Pueblo, 2022]. Additionally, we estimate the effects of the law on the time children spend performing household chores (in the past week) but we find no evidence that the law impacted children’s time allocation along this dimension; the estimated effect is small and statistically insignificant (column 2).³⁴ We also find that the 2014 law had no significant effects on the labor supply of other household members or on household income per capita (see Appendix Table B.5)

2.5.2 The role of enforcement


The 2014 law highlighted the protections and benefits newly granted to workers under the age of 14 and tasked the MTEPS with ensuring compliance with the law through inspections. These inspections complemented the labor and workplace safety inspections already being conducted by the MTEPS before the law, which verify firms’ formal registration and compliance with general worker regulations. As discussed in Section 2.2.3, MTEPS inspections — both generally, and specifically for child labor — increased after the enactment of the law (see Figure 2.1). However, the threat of enforcement varies across localities; there is substantial variation in

³³Since 2009, schooling has been compulsory for all primary and secondary levels, and free in public schools. Thus, our estimates do not confound any changes in compulsory schooling laws.

³⁴Note that the data on participation in domestic chores comes from the ENTI 2008 and the ENNA 2016, described in more detail in Section 2.6.1. As there is no data beyond 2016, we cannot estimate a post-reversal coefficient for this outcome.

a locality’s proximity to the nearest regional MTEPS office (see Appendix Figure B.1).

We exploit this cross-locality variation to verify whether the effects that we document are driven by children working in areas where inspections are more likely. Previous work finds that distance acts as a deterrent to enforcement of labor regulations [Almeida and Carneiro, 2012, Ponczek and Ulyssea, 2021, McKenzie and Seynabou Sakho, 2010], and evidence from Bolivia suggests that compliance with tax registration is higher among firms closely located to the tax authority [McKenzie and Seynabou Sakho, 2010]. We find corroborating evidence in our data; Figure 2.4 illustrates that adult workers in areas closer to MTEPS offices (based on driving routes optimized to minimize travel time) are more likely to have formal labor contracts and employer-provided health insurance, even after controlling for job and worker characteristics that are likely correlated with distance to MTEPS offices (such as worker education, sector of work, and firm tax registration – a marker of firm formality).



(a) Formal Labor Contracts for Workers

(b) Health Insurance for Workers

This figure presents the proportion of adult workers (age 18+) that have a formal work contract (panel a) and have health insurance through their employer (panel b), by quantiles of driving time to the nearest MTEPS office (50 quantiles) using the 2012-2013 Encuesta de Hogares. The data are residuals after removing variation due to the following controls: age, gender, years of schooling, an urban dummy, a dummy variable denoting department capitals, sector of work fixed effects, and a dummy for firm tax registration.

Figure 2.4. Compliance with Labor Regulations and Travel Time to Inspectors (Pre-Law)

Accordingly, we exploit cross-municipality variation in the driving time to regional MTEPS offices to proxy for variation in the probability of workplace inspections. We compare the effects of the law on work probabilities between municipalities that are “far” and “near”

from the nearest regional Ministry of Labor (MTEPS) office, where “far” is defined as above the median driving time.³⁵ Note that municipality codes are anonymized in the household data starting in 2017, meaning that we cannot link the data to other sources using municipality codes in 2017 and later. Thus, the sample for Table 2.5 does not include data past 2016 and we cannot estimate a “Post-reversal” coefficient. This exercise is similar in spirit to that in Bargain and Boutin [2021], who study heterogeneity in a Brazilian law’s effect using state-level variation in labor inspection rates.

Panel A in Table 2.5 illustrates that the law appears to decrease the likelihood of allowed/regulated work for 13-year-olds relative to 14-year-olds, but only in areas that are located near MTEPS offices, where there was likely to be stronger enforcement. This remains true when we further restrict the sample to municipalities that do not contain an MTEPS office (column 2), illustrating that the result is not being driven only by large, mostly urban municipalities.³⁶ These results are robust to using straight line or “as the crow flies” distance as an alternative measure of distance to MTEPS offices (see Panel B), and are estimated with greater precision. While the effects are not statistically distinguishable across areas near and far from MTEPS offices, the point estimates suggest that the overall declines in child labor are almost exclusively driven by children in localities closer to enforcement offices.³⁷ These results are consistent with those in Bargain and Boutin [2021], who find that the effects of a Brazilian child labor law are detectable only in states with a high potential threat of inspection. We do not find substantially different effects between municipalities that are near and far from the MTEPS regional offices for younger children (see Appendix Table B.6), likely due to the low incidence of overall child labor among younger children. Overall, the results suggest that enforcement was a key driver of the decline in

³⁵We measure the driving time from the municipality capital, typically the most populated locality in the municipality, to the nearest MTEPS office. See Appendix Section B.4 for details.

³⁶These results also help to rule out the concern that the results are driven by family work in subsistence farming, which is more prominent in isolated areas far from MTEPS offices.

³⁷To show that these results are not driven by differences across urban and rural areas or by geographical variation in baseline child labor rates, in Appendix Table B.7, we also report the results after additionally controlling for all possible interactions between the treatment variables, the post-law indicator, and urban status/district-level baseline child labor rates (as measured in 2012). Adding these controls does not change the results in a meaningful way; the estimated effects in areas near and far from MTEPS offices are very similar to those reported in Table 2.5.

Table 2.5. Heterogeneous Effects of the Law by Distance from MTEPS Offices (Diff.-in-Disc.)

Panel A: Driving Time		
	Dependent Variable: Works	
	All (1)	No MTEPS Offices (2)
Post \times 1{Age < 14} for Far	0.002 (0.061)	0.002 (0.058)
Post \times 1{Age < 14} for Near	-0.030 (0.021)	-0.074* (0.043)
Obs.	7650	2984
Mean	0.180	0.317
P-value of difference	0.644	0.338

Panel B: Direct Distance (“as the crow flies”)		
	Dependent Variable: Works	
	All (1)	No MTEPS Offices (2)
Post \times 1{Age < 14} for Far	0.009 (0.060)	0.007 (0.056)
Post \times 1{Age < 14} for Near	-0.037* (0.021)	-0.103** (0.044)
Obs.	7650	2984
Mean	0.180	0.317
P-value of difference	0.496	0.175

Notes: Age in months by year clustered standard errors in parentheses. Significance levels denoted by: *** p<0.01, ** p<0.05, * p<0.1. Municipalities that are classified as Far are above the median distance from a MTEPS office. Control variables: CCT eligibility indicator, urban, HH head characteristics (schooling, gender, age, indigenous indicator), gender, no. of children aged 0-6, 7-9, 10-13, and 14-17, no. of adult men and women, and departamento by year FE. We also include linear splines of the running variable (difference between the cutoff age and age a week before the survey date in months). We use a bandwidth of 12 months and a triangular kernel. Survey years: 2012-2016. We also report the mean of the dependent variable for the pre-law period.

child work due to the law. We discuss the mechanisms behind these results in Section 2.6.

2.5.3 Robustness

Alternative Specifications. We show that our results are robust to alternative specification choices that are common in Regression Discontinuity designs. First, we show that our results are robust to different analysis bandwidths. Our main results on work probabilities

are based on estimates of equation (2.1) using a twelve-month bandwidth around each cutoff. Columns 1 and 3 of Appendix Table B.8 shows that the results are unchanged when we expand the estimation bandwidth to 24 months and when we reduce the bandwidth to six months, albeit with a substantial decline in precision in the latter case. Second, our point estimates are robust to excluding demographic controls from our main specification (see column 4). Third, our results are robust to using a second-order polynomial on each side of the cutoff to flexibly control for the running variable, and to allowing the slopes to vary before and after the policy change on either side of the cutoff, respectively (columns 5-7).

Measurement error. In our main specification we use age in months to determine exposure to the law. However, because we do not have the exact survey interview date, among children born in the same month, there might be children who were exposed to the law at the moment of data collection and others who were not. To ensure that measurement error is not biasing our results, we show that our results are very similar when we exclude observations of children that, according to their age in months, are within a month of exposure and who are more prone to misclassification (column 9). It is worth noting that the fact that the results are robust to this alternative specification also attenuates the concerns that the small difference in densities between the post-law and reversal periods at the 14 year-old cutoff reported in Appendix Figure B.5.

Another potential source of measurement error stems social desirability bias.³⁸ In particular, one might worry that the law changed the stigma surrounding child labor and affected the accuracy of parents' reports of their children's work around the cutoffs. However, we think that this is unlikely for several reasons. First, we observe no discontinuities in either survey responses (Appendix Figure B.4) or in reported work probabilities in the pre-law period (Figure

³⁸The extent to which measurement error in child labor as reported by proxies (e.g., parents) plagues household survey data and whether it is related to social attitudes and norms is debated. Some find that there is no systematic differences across reports by children and proxies when concerning economic activity [Dillon et al., 2012, Dziadula and Guzmán, 2020] while others find differences but no relation to social norms [Dammert and Galdo, 2013]. A recent study from the cocoa industry in Cote d'Ivoire finds that proxies severely under-report work of children attending school and that under-reporting responded to an intervention that potentially signaled support (rather than punishment) for farmers with working children [Lichand and Wolf, 2022].

2.3), when work under 14 was illegal. Second, the 2014 law legalized and legitimized work for those under 14. If anything, we expect that the law reduced pressure for parents to under-report their children's work (i.e., be more likely to report that their children work) after the 2014 law. However, we find that children under 14 become less significantly likely to work after the 2014 law, suggesting that our results may underestimate the true labor-reducing effects of the law.

Alternatively, one might think that the 2014 law increased the salience of the harm caused by work for young children and made parents more reluctant to admit their children were working. If this were the case, the reduction in child work that we document could simply reflect reduced parental reporting of work for children under 14 rather than an effect of the law on work. In this scenario, we would expect the stigma surrounding child work to be especially strong for younger children; however, we find no evidence consistent with this hypothesis. In column 1 of Appendix Table B.2, we compare 12- to 13-year-olds. Both of these age groups are subject to the same legal status and worker protections, but if the law underscored the harm work causes younger children, we should expect work to decline even more for 12-year-olds relative to 13-year-olds; yet, the results indicate that there are no differences in the responses of the two age groups.

Difference in differences approach. As discussed in Section 2.4.1, one limitation of our difference in discontinuities design is that it only enables us to make local comparisons just when children change their treatment status based on their age. If the enforcement of the law varied with how far children are from the cutoff, then our main estimates would be capturing lower bounds. Appendix Table B.10 reports results from three difference-in-difference specifications that allow for comparisons of work outcomes of treated and control children over time, regardless of their proximity to the cutoffs. Column 1 shows results for a simple diff-in-diff model that uses 14-year-old children as controls for 13-year-old children, controlling for a set of demographic attributes as in our main specification. The point estimates are remarkably similar to those in our main specification. Columns 2 and 3 show results of alternative specifications that use two definitions of control groups made up of younger and older children (9 and 14 year old children,

and 7-9 and 14-16 year old children, respectively). Column 2 yields qualitatively similar results than our main specification. However, when we expand our pooled control group to include children as young as 7 and as old as 16 as in Kamei [2020], who analyzes the impact of this law but only using one post-period survey wave, the coefficients drop in magnitude and are not statistically significant, although they remain negative. These changes may reflect potential violations to the identification assumption for this pooled difference-in-difference specification—that in the absence of the policy change, the work outcomes would have evolved similarly for the 7- and 16-year-old control groups and the younger and older treatment groups.³⁹

Accounting for communal work. The 2014 law allowed the participation of children in community activities – culturally valued activities taking place in indigenous, Afro-Bolivian, and intercultural communities – without age restrictions as long as the activities contribute to children’s integration into the community or to the development of skills and only if it did not represent exploitation, contradict a child’s rights, or entail potentially risky activities. Examples of these include working in a communal farm or working for community organizations. While our data do not allow us to identify specific types of community labor (to which the law exceptions apply), work for this purpose appears to be rare; in 2016, only 6% of working children report maintaining family or community customs as the main reason for working.⁴⁰

In column 2 of Appendix Table B.9 we provide evidence that our main results are not driven by changes in these types of activities by excluding municipalities with a high share of residents that identify as indigenous (defined as municipalities with an above-median share of indigenous residents), where communal work related to cultural traditions are more prevalent and

³⁹Our results on child employment are at odds with those of Kamei [2020] who studies the impact of the 2014 Bolivian law and finds that the probability that boys age 12-13 work for their families increases in 2014 relative to the pre-law period. We believe that the differences with our results arise largely from differences along two important dimensions: data and empirical approach. First, we study the effects of the law over a longer horizon (up to 4 years after the introduction of the law), whereas Kamei [2020] restricts attention to the 6 months after the introduction. This longer time span is important if the law’s effects take time to surface — for example, if employers take time to adjust to the new regulations. Second, using children that are much younger and older than the treatment group is prone to violations to the parallel trends assumption stemming from differential responses to labor-market shocks by age.

⁴⁰For children ages 7-17. Authors’ calculations using the 2016 ENNA.

where these exceptions to the law are more likely to apply.⁴¹ The point estimates are remarkably similar to those from our main specification, despite excluding 45% of the sample.

2.6 Mechanisms

The mechanisms behind the negative effects of recognizing work for younger children and extending worker protections to such workers are not *ex ante* obvious. Traditionally, the trade-off between increased worker protections and reductions in labor demand is linked to the idea that as firms comply with new regulations, the cost of hiring increases, which in turn depresses the demand for labor [Lazear, 1990]. However, in markets where most employers are informal and operate under the radar of regulation, firms may also reduce the demand for newly entitled workers to continue avoiding attention from inspectors and regulators. Likewise, in markets associated with a large degree of social stigma, the increased scrutiny amid new regulations may also deter labor supply. We discuss these mechanisms below.

2.6.1 Compliance costs

One key objective of the new law was to improve the working conditions of children. One possible explanation for the overall declines in employment among younger children is that the law increased the safety of child work (at a cost to employers) and subsequently reduced the demand for child workers. We explore this hypothesis by analyzing two child labor surveys on risky tasks and injuries at work: the ETI 2008 and the ENNA 2016.

There are some empirical challenges related to these data. First, the surveys come from different sampling frames. The ETI 2008 samples children that are likely to work while the ENNA 2016 is nationally representative of all children. We combine the two surveys by reweighting the data so that observations that are similar (based on observables) across

⁴¹Municipalities are classified according to the 2012 Census data. Note that municipality codes are anonymized in the household data starting in 2017, meaning that we cannot link the data to other sources using municipality codes in 2017. Thus, the sample for column 2 of Appendix Table B.9 does not include data from 2017.

survey waves are given higher weight.⁴² In Appendix Table B.11, we show balance on these characteristics across the age thresholds and survey rounds (after re-weighting) using random subsamples that were not used in calculating the weights.⁴³ Second, with only two survey waves of these data, we have much smaller samples to assess the effects of the law on job safety outcomes separately at each age threshold. To improve the precision of our estimates, we estimate a stacked difference-in-discontinuity specification, an often-used approach to estimating a common treatment effect across multiple cutoffs (see, for example, Beuermann and Jackson [2020], Pop-Eleches and Urquiola [2013]). Specifically, we pool the samples across age groups but maintain the definitions of treatment variables and running variables to be relative to each specific threshold.⁴⁴ We additionally include cutoff fixed effects, which ensures that our estimates continue to be based on local comparisons around each age cutoff.⁴⁵ Finally, there are no surveys on risky tasks and work injuries after 2016, so we cannot study the effects of the 2018 reversal on these outcomes.

We find neither significant or substantial declines in the incidence of risk (column 1) and injuries at work among treated children (column 3)— who are newly granted worker protections under the 2014 law (see Table 2.6).⁴⁶ We are able to rule out declines in risk larger than 4.3

⁴²To calculate the weights, we pool the observations from a randomly chosen 70% subsample from each survey and then predict the likelihood of appearing in the 2016 nationally representative ENNA using a Probit model based on demographic characteristics of children and their households. We then use these predicted probabilities (propensity scores) to construct weights. Observations from the 2016 survey receive a weight of $\frac{1}{p}$, where p is the predicted probability of being in the 2016 survey. Observations from the 2008 survey receive a weight of $\frac{1}{1-p}$. This reweighting procedure is similar in spirit to the one proposed in Abadie [2005], which aims to minimize bias and maximize balance across the samples.

⁴³We follow this approach to ensure that balance on targeted variables is not simply a consequence of overfitting. We used 70% of the observations to estimate the propensity score p and the remaining 30% to test balance.

⁴⁴Because the treated group are those over the threshold at the 10- and 12-year-old cutoffs but below the threshold at the 14-year-old cutoff, we multiply the running variable by -1 for the observations around the 14-year-old cutoff to maintain consistency across thresholds.

⁴⁵Specifically, we estimate a slightly modified version of the specification in equation 2.1 that includes cutoff fixed effects. In estimating equation 2.1, we use combined weights that reflect both the triangular weights and the constructed sampling weights. For the pre-period (2008), we divide the triangular kernel weights by one minus the inverse probability of being in the post sample in 2016. For the post-period (2016), we divide the triangular kernel weights by the inverse probability of being in the post sample in 2016.

⁴⁶We display graphical evidence in Appendix Figure B.11. There is no evidence of differential changes in sample composition across any of the age cutoffs in the child labor survey (Appendix Table B.12).

percentage points and declines in injuries larger than 4.0 percentage points with 95% confidence. The results are robust to alternative specifications.⁴⁷ The lack of substantial declines in risky activities suggest that compliance with costly safety regulations was not a key driver of the decline in employment among 13-year-olds.

Table 2.6. Effects of the Law on Risk, Injuries at Work and Wages

	Faces Risks at Work (1)	Faces Risks at Work (2)	Has Been Injured at Work (3)	Has Been Injured at Work (4)	Log Hourly Wage (5)
Post Law \times Treated	-0.008 (0.017)	-0.038 (0.035)	-0.015 (0.014)	-0.015 (0.029)	0.103 (0.180)
Post Reversal \times 1{Age < 14}					-0.012 (0.180)
Obs.	8372	2914	8411	3208	712
Mean	0.281	0.536	0.188	0.327	6.656
Sample	All Children	Working Children	All Children	Working Children	Paid Workers

Notes: Age in months by year clustered standard errors in parentheses. Significance levels denoted by: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The sample in columns 1 to 4 comes from the child labor survey, and the sample in column 5 comes from the household survey. Control variables: gender, working indicator (Panel B only), urban indicator, age group fixed effects, household head characteristics (schooling, gender, age, and indigenous indicator), gender, number of children in the following age categories: 0-6, 7-9, 10-13, and 14-17, number of adult men and women, and departamento by year fixed effects. For the risk index regressions, the running variable is the difference between age in months and the age cutoff at the survey date. For the injury index, the running variable is the difference between age in months and the age cutoff a year before the survey date. In columns 1 to 4, we do a stacked difference in discontinuity by multiplying the running variable by -1 for the 13 and 14 year-olds age group for interpretability. For column 5, we do a difference in discontinuity in which the running variable is the difference between age in months and the age cut-off a week before the survey date. The specification includes linear splines of the running variable. The bandwidth for all specifications is 12 months. We use a triangular kernel. Survey years: 2008 and 2016 in columns 1 to 4 and 2012-2019 in column 5. We use a reweighting method for columns 1 to 4 described in Section 2.4.

One concern with our empirical approach is that because the law reduced child work around the 14-year-old cutoff, our reduced-form results do not accurately capture the true impacts of the law on risks outcomes among children who remain working. We offer two pieces of evidence to rule out this concern. Columns 2 and 4 in Table 2.6 show that we are unable to detect significant differences in risk exposure and injuries when we focus only on children who report working. While the point estimate on risk suggests a three percentage point decline (column 2), the point estimate on the probability of suffering an injury at work remains unchanged (column 4). Second, when we replicate our analysis around each age cutoff in Appendix Table B.13, we find relatively small, non-significant effects for younger children—those for which we found

⁴⁷Appendix Table B.14 shows that our results are also robust to changes in bandwidth, excluding controls, including a quadratic polynomial in the running variable, and excluding children within 1 month of the cutoffs (donut-style regressions).

no effects of the law on work probabilities. Reassuringly, we do not find neither substantial nor significant effects around the 14-years-old cutoff either.

Another possibility is that the law directly increased the costs of hiring younger children relative to children age 14 or older, as it established that even children age 13 or younger were entitled to receive the minimum monthly salary. In column 5 of Table 2.6, we use the subsample of working children who report wages to estimate differences in wages induced by the policy change. We find that the hourly wages of working children just under the age of 14 are ten percent larger than those just above age 14 during the period in which the law was enforced, and that these differences vanish when the 2014 worker protections were no longer enforced.⁴⁸ However, these differences are not significant at conventional levels. Moreover, even taking this difference at face value, the potential increase in wages is unlikely to account for the negative effect of the law on child work, as the subset of working children that report receiving a salary is very small (712 children out of roughly 3,600 working children and 18,000 children overall in the sample around the 14-year-old cutoff).

Thus, contrary to previous studies analyzing the impact of increased worker protections [Lazear, 1990, Autor et al., 2007, Almeida and Carneiro, 2012], the evidence from Bolivia suggests that the effects of extending rights to child workers does not seem to be explained by increased worker benefits and hiring costs for complying firms.

2.6.2 Avoidance behavior

Most children work for informal firms. Such firms — by virtue of hiring “off the books” — face different incentives after the introduction of new regulations recognizing the work of younger child workers, who before the policy change were hired illegally. Given the context of high public scrutiny of the 2014 law, hiring younger children — a demographically distinguishable group — may increase the visibility of firms and thus the risk of labor inspections. To the extent

⁴⁸To increase sample size, given the low survey response rates related to child earnings, we estimated the wage equation around a wider bandwidth (18 months) around the 14-year-old cutoff.

that firms internalize this increased risk, they may choose to avoid hiring younger children in order to remain under the radar of regulation. Likewise, as most informal firms are family-owned, the new regulations may have deterred parents from employing their children in their firms. In addition, the high level of scrutiny and social stigma around the new regulations may have deterred other parents from allowing their children to work from others. In equilibrium, with higher perceived risks of regulatory and social sanctions, the new regulations may have triggered avoidance behavior on both sides of the market.

These incentives are consistent with the institutional setting in Bolivia: The entity in charge of child labor inspections (MTEPS) is also in charge of general labor and workplace inspections, and thus firms that draw attention from child labor inspectors will also likely be subject to general inspections. Indeed, as discussed in Section 2.5.2, we find that child employment declined in areas located nearer to MTEPS offices where visibility to inspectors is particularly relevant (see Table 2.5).

One empirical implication of this mechanism is that the declines in child work should be driven by firms with greater visibility. We thus distinguish between children who work outside the home at a fixed establishment and children who work either at home or outside home in non-fixed, mobile locations. The intuition is that inspectors may be better able to track firms operating at fixed external establishments (e.g., a factory, or a shop) as opposed to those operating inside the owner's home with no external visibility or those that frequently change locations and are less traceable (e.g., family farms or street vendors). Panel A in Table 2.7 reports treatment effects of the law on the probability of working at a fixed establishment, on the probability of working at home, and on the probability of working at a mobile work location, around the 14-year-old cutoff. We observe a 5 percentage point decline in the probability of working at a fixed location (column 1); this effect is statistically significant and meaningful in magnitude (about a 33% decline). In contrast, we find no effects on the probability of working at home or in a mobile location (columns 2 and 3). This suggests that the decline in overall employment among 13-year-olds is largely explained by a contraction in the employment of children who

worked in more traceable and visible locations.⁴⁹

Table 2.7. Effects of the Law on Job Location and Firm Size

Panel A: All Children				
	Works in Fixed Location Out of Home (1)	Works in Mobile Location Out of Home (2)	Works at Home (3)	
Post Law \times 1{Age < 14}	-0.051*** (0.014)	0.009 (0.008)	0.003 (0.004)	
Post Reversal \times 1{Age < 14}	-0.009 (0.016)	0.005 (0.013)	0.003 (0.003)	
Obs.	11991	11991	11991	
Mean	0.149	0.0248	0.00588	

Panel B: Working Children				
	Works in Fixed Location Out of Home (4)	Works in Mobile Location Out of Home (5)	Works at Home Size (6)	Firm Size (7)
Post Law \times 1{Age < 14}	-0.098*** (0.035)	0.078** (0.034)	0.021 (0.018)	-0.726 (0.473)
Post Reversal \times 1{Age < 14}	-0.043 (0.050)	0.022 (0.054)	0.021 (0.018)	-0.359 (0.383)
Obs.	2323	2323	2323	2250
Mean	0.829	0.138	0.0327	4.796

Notes: Age in months by year clustered standard errors in parentheses. Significance levels denoted by: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The sample in Panel A includes all children, while the sample in Panel B is restricted to working children only. Control variables: CCT eligibility indicator, household head characteristics (schooling, gender, age, and indigenous indicator), gender, number of children in the following age categories: 0-6, 7-9, 10-13, and 14-17, number of adult men and women, an urban dummy, and departamento by year fixed effects. We also include linear splines of the running variable, defined as the difference between the cutoff age and age a week before the survey date in months. We use a bandwidth of 12 months and a triangular kernel. Survey years: 2012-2019.

In Panel B of Table 2.7, we examine how the law affected the composition of employment. For this, we focus on the subsample of employed children before and after the policy change. We find that relative to the pre-law periods, the share of 13 year old children working in fixed, high-visibility establishments declines by 10 percentage points when the law was enforced, relative to 14 year old workers (column 4). In contrast, we observe a 8 percentage point increase in the share of younger children (under 13 years old) who are employed in mobile, less traceable locations. Even though these estimates are not causal, they suggest a change in the composition of employment that is consistent with younger children moving to less visible forms of employment.

⁴⁹These results complement and strengthen those in Bargain and Boutin [2021], who find “very mild evidence” that the effects of a Brazilian child labor law were concentrated in activities that were more “inspectable.”

Another dimension of visibility of work is firm size. Previous studies find that larger firms are more likely to be targeted by regulators than small firms [Almeida and Carneiro, 2009].⁵⁰ This implies that larger firms have a greater incentive to reduce hiring of young workers targeted by the legislation because of the higher threat of inspection. As a result, young children may end up working for smaller (i.e., less visible) firms; indeed, in column 7, we find that children under 14 work for smaller firms while the law is in place, albeit not significantly so.

The results in Table 2.7 are also consistent with avoidance behavior on the part of families. One possibility is that the law brought more awareness to the potential adverse consequences of work for young children. If the law resulted in a greater perceived stigma related to child work — particularly for those under 14 — then reductions in visible forms of work may also reflect parents' unwillingness to work their children in ways that are visible to others in the community.

To test if the declines in employment are exclusively explained by firm adjustments or also driven by household decisions, we analyze the difference between the effects of the 2014 change in regulation on employment and labor force participation. If the effects were exclusively driven by firm's avoidance behavior, then one should observe more muted effects on labor force participation as laid-off children look for new jobs. In contrast, if both firms and parents react to the new regulations, the effects on employment and on labor force participation should be quantitatively similar. Indeed, in Table 2.4, we see that the effects of the law on work status (column 1) are virtually identical to those on labor force participation (column 7). Combined with the results in Table 2.7, suggests that the reduction in employment we find is consistent with parents removing their children from forms of work that are visible to the community, inspectors, or both.

⁵⁰Almeida and Ronconi [2016] outline a number of reasons why enforcement agencies may target larger firms; for example, larger firms may be less costly to inspect; they may be more visible to media and the public; and they may have more rents to extract if inspectors are corrupt.

2.6.3 Costs of work permits

An alternative explanation for our central finding that younger children are less likely to work during years the law was in place is that the costs and complexity of the application process may have lowered the demand for permits. To qualify for a permit, children first had to be declared fit to work by a doctor following a medical exam, and then visit the closest Child Advocacy office (DNA), often in a different locality. This process has been described as “highly tedious”, characterized by long waits at DNAs and hospitals and requiring considerable time and effort on the part of administrators, parents, and employers [Defensoría del Pueblo, 2022].⁵¹ These transaction costs may deter children from legally entering into the workforce, even when they have the option to do so. Consistent with evidence showing that the complexity of application processes for public services reduces takeup [Banerjee et al., 2021], the probability of having a permit is substantially lower among the children from the poorest households (see Appendix Figure B.12), who are least able to pay the costs of the obtaining a permit. Moreover, the process may discourage employers from hiring children under the age of 14 (who were required to obtain work authorization from DNAs). As one child states, the authorization process is one reason “why they don’t give young people so much work” [Defensoría del Pueblo, 2022, p. 103].

2.7 Conclusion

Overall, we find no evidence that recognizing the work of young children and extending worker protections to them increased child work in Bolivia. In fact, we find that children under age 14 were less likely to work in permitted and regulated activities after the passage of the law (relative to children over age 14). We posit that this is primarily due to the new regulation

⁵¹Low investments in and lack of easily accessible DNA offices exacerbated the cost of obtaining permits. Though the 2014 law mandates that every municipality in the country have a dedicated Child Advocate Office, as of 2016, 20% did not have one and many lack funding, personnel, and materials [U.S. Department of Labor, 2011-2019]. Likewise, in a recent report from a survey to 59 out of 339 municipalities, the People’s Advocate Office (*Defensoría del Pueblo*) found only 12% of surveyed municipalities kept records of child and adolescent labor [Defensoría del Pueblo, 2021].

inducing avoidance behavior by firms and parents. The 2014 new regulations increased the perceived costs of employing younger children— both through increased scrutiny and threat of inspections for firms hiring young children and through the new regulations that granted rights and protections to working children under 14. As some have claimed, “For adolescents, the code frequently had the effect that companies preferred to hire adults rather than jump over bureaucratic hurdles” (Liebel [2019]). Indeed, we find that after the key child labor components of the law (those granting rights and protections to workers under the age of 14) were repealed in 2018, work probabilities and hours of work returned to pre-law levels for children under the age of 14.

Importantly, we find that the law did not significantly affect children’s riskiness of work or injuries on the job. This stands in contrast to one of the purported aims of the policy to make child work safer. Together with the observed decline on employment among younger children due to the law, these results suggests an overall worker welfare loss: the law appears to have reduced employment among the children that it intended to protect without improving the worker conditions of the children who kept their jobs.

The findings are important to the broader discussion of optimal child labor policy. While previous work finds that outright bans are not able to eradicate child labor, our results illustrate that a natural alternative — legal recognition and regulation of child labor — does not necessarily make child work safer. Both bans and legalization/regulation do not address what many consider the root cause of child labor: poverty. Instead, these policies affect employers’ costs of hiring children, and thus affect child labor in nuanced ways that can run contrary to policy aims.

Finally, perhaps a silver lining of the Bolivian case is the powerful role of public scrutiny, which appears to have led to avoidance behavior. Increasing the salience of social issues that are very quite sensitive may be able to achieve what seems challenging for regulation in settings with limited state capacity.

Chapter 2, in full, has been submitted for publication of the material as it may appear in the Journal of Development Economics, 2024, Martinez Heredia, Diana; Lakdawala, Leah K.; Vera-Cossio, Diego. The dissertation author was the primary investigator and author of this paper.

Chapter 3

Gender and Transitions into Adulthood

3.1 Introduction

The productive potential of young people is pivotal in advancing the goals of poverty reduction and economic growth. However, young people face significant challenges transitioning to work. Fares et al. [2006] document that, in most countries, young people face higher unemployment rates and a greater incidence of low-paying or unpaid jobs than adults. Additionally, many young girls are often outside the labor market, with many engaged in home production, and transitioning out of that type of work can positively impact their future labor market outcomes [Azevedo et al., 2012, Fares et al., 2006]. Furthermore, young people, especially women, who actively decide over their fertility and new household formation have better opportunities to participate in the labor market [World Bank, 2006]. Hence, investing in young people and understanding the constraints they face as they join the workforce and form new households during their transition into adulthood is crucial.

In this paper, we evaluate two constraints that young Rwandan people may face as they transition into adulthood: access to training to learn valuable skills for the labor market and access to capital. On one hand, developing skills at a young age is crucial because skills are cumulative and build upon one another over time. By learning general skills, young people enhance their ability to contribute economically, become better prepared to care for their health, and be more effective parents and engaged citizens [World Bank, 2006]. On the other hand, credit

constraints can be a major obstacle for young people to have access to economic opportunities, and this is especially significant for young women, whose role in the household is an obstacle to accessing economic opportunities [Fox et al., 2016, Blattman et al., 2013].

Hence, we evaluate how alternative programs aimed at improving the employment and productive potential of underemployed Rwandan youth affect young women's and men's transitions to adulthood. We focus on gender differences due to the additional barriers that women face in accessing economic opportunities and their role in childbearing. We use data from McIntosh and Zeitlin [2022], who conducted a randomized controlled trial with 1,848 youth, comparing a package of training, soft skills, and networking interventions to cash transfers and a control group. The cash transfer program was implemented by GiveDirectly (GD), a US-based nonprofit specializing in making unconditional household grants via mobile money. The training program, called Huguka Dukore (HD), follows USAID's strategy on workforce readiness and skills training. This program was promoted as a tool for female empowerment. Hence, it is valuable to understand if HD achieved that goal.

These interventions allow us to evaluate how young people respond to receiving financial aid through cash transfers, vocational education training, or both. The financial constraint is a credit constraint, while the skills constraint is the obstacle these young people face in accessing useful training to prepare for the labor market. These young people are past high school age but have not graduated and work mostly in agricultural self-employment. We can assess differential effects by gender on economic outcomes and cultural outcomes, such as the likelihood of marriage, cohabitation, fertility, desired fertility, household formation, gender attitudes, aspirations, financial control that women hold in the household, occupations, being involved in income generating activities, time spent in domestic activities and willingness to relocate for a job. These comparisons help us understand the nature of the constraints young people face as they transition into adulthood and how these constraints vary by gender.

The interventions affect the transitions to adulthood in some subtle but important ways. Overall, there is a remarkable similarity of the final economic impacts of the programs on male

and female beneficiaries, suggesting that lifting financial constraints does not close gender gaps in economic outcomes within the group of individuals who enrolled themselves as eligible for this study. This is surprising for two reasons. First, the literature has found that women experience different effects from receiving cash than men [De Mel et al., 2012, Blattman et al., 2013]. Second, in our sample, women are less likely to be employed, work fewer hours, earn less income, and consume less. Hence, there could be more opportunities for women to catch up as they receive these interventions. The lack of differences between men and women is consistent with the persistence of the roles that women are assigned despite these very involved interventions. We find that schooling can mediate these effects, but this mediation is only present in midline, and not endline.

In Addition, we uncover clear evidence of how lack of income inhibits marriage for men; while no intervention drives cohabitation rates, the larger cash arms have a strong effect on increasing marriage, and only for men. Fertility rises in line with marriage rates, again only for men and only for the larger cash arms. Desired lifetime fertility, on the other hand, shows a sharp decrease from cash transfers, and only for women, indicating that the substantial changes in female entrepreneurship and time use induced by cash transfers do alter the way that they think about family size in line with the predictions of Becker [1965]. These results contrast others in the literature. For example, Baird et al. [2011] shows that cash transfers have decreased the probability of marriage in Malawi, and Baez et al. [2011] shows decreases in the age of marriage in Pakistan. In line with our findings, Baez et al. [2011] and Baird et al. [2011] find reductions in the probability of pregnancy in Pakistan and Malawi, respectively.

We investigate whether these changes co-occur with changes in the perceptions of gender roles. We find that this is not the case. Women do not transition out of domestic chores. We examine measurements of gender attitudes, female financial control, and women's aspirations. Neither HD nor GD have any effects in these dimensions. This is also striking since there is evidence that cash transfers can change women's empowerment. For example, there is evidence that cash transfers increase the probability of women making expenditure decisions in their

household for Bolsa Familia, PROGRESA, and HSNP [De Brauw et al., 2014, Handa et al., 2009, Merttens et al., 2013]. The Rwandan context is relatively progressive for women in the labor market and politics but conservative regarding gender roles. This paper shows that, in these contexts, it takes more than these involved interventions to change perceptions of gender roles.

These results paint an interesting picture. Surprisingly, there are no major gender differences in economic outcomes. Besides, we see differential effects in marriage and fertility that are consistent with the bride price culture, which exerts pressure on men and their families to provide money or assets to formalize marriages, a key social institution. Men use cash to marry, and for them, cash has weak but positive effects on fertility. Women, on the other hand, see decreases in desired fertility as they receive cash. However, none of these changes come with changes in the perception of gender roles. Hence, our results show that young people in Rwanda primarily face financial constraints as they transition into adulthood. Despite having been very involved and having costly interventions, no cultural outcomes have changed other than through the financial constraint channel. In addition, HD, a program promoted as a tool for female empowerment, had no meaningful differential effects for women in any dimension, specifically in perceptions of gender roles.

3.2 Data and Empirical Analysis

We use the data and follow the empirical analyses from McIntosh and Zeitlin [2022]¹. The data analysis follows the design of a three-armed program randomized at the household level. The main specification studies the effects of the interventions with gender interactions. The outcomes are (Y_{ih1}) . The subscript i stands for the individual and h for the household. As in McIntosh and Zeitlin [2022], the specification includes fixed effects for the sector-level assignment blocks within which the randomization was conducted μ_b , and a set of control variables, X_{ih0} , which were selected from the baseline data for their predictive power of the

¹For information about the interventions, enrollment criteria, assignment protocol, balance, and attrition, please reference McIntosh and Zeitlin [2022].

primary outcomes in McIntosh and Zeitlin [2022]. There are three treatment variables: the HD treatment T_{ih}^{HD} , the GD treatment T_{ih}^{GD} , and an indicator for the combined arm T_{ih}^{COMB} . We also include the indicator $female_{ih}$, which is one when the beneficiary is a woman. Finally, we include the interactions between the female indicator and the treatment arm variables:

$$\begin{aligned}
Y_{ih1} = & \beta_0 + \delta^{HD}T_{ih}^{HD} + \delta^{GD}T_{ih}^{GD} + \delta^{COMB}T_{ih}^{COMB} \\
& + \gamma^{HD}T_{ih}^{HD} \times female_{ih} + \gamma^{GD}T_{ih}^{GD} \times female_{ih} + \gamma^{COMB}T_{ih}^{COMB} \times female_{ih} \\
& + \gamma^{FEM}female_{ih} + \beta X_{ih0} + \rho Y_{ih0} + \mu_b + \varepsilon_{ih1}
\end{aligned} \tag{3.1}$$

The standard errors are clustered at the household level. Our main focus is the estimates for γ^{HD} , γ^{GD} and γ^{COMB} , which indicate the differential effects of the treatment arms for women.

We evaluate two sets of outcomes: the economic dimension and the cultural dimension. In the economic dimension, we look at productive time use and business characteristics. Productive time use variables include the probability of doing any work, independent work, being employed by an outside employer, and working on one's own farm. Business characteristics include the likelihood of co-ownership, the amount of debt, the number of employees, the value of assets, and profits. In the cultural dimension, we analyze four main outcomes: the likelihood of being married, the likelihood of cohabiting, the likelihood of having any children, and desired fertility, which is the number of children the beneficiary would like to have, in addition to the ones that they already have. We also look at household formation outcomes, which include the likelihood of forming a new household, moving villages, and moving to urban areas. We also analyze education and time use outcomes, including the highest grade of education attained, hours spent in school, hours spent in domestic chores, and reservation wages. Finally, we analyze variables related to gender roles, such as attitudes about men's and women's roles, women's financial control in the household, and women's aspirations. All these variables are Z-scores. Gender attitudes Z-scores summarize the beneficiary's responses to statements related to the importance of women's education with respect to men's education, marriage, and labor market prospects.

A higher gender attitudes' Z-score means that the beneficiary considers women's education more important. The women's financial control Z-score summarizes women's agency over household financial decisions. A higher financial control Z-score means that the beneficiary has more agency over expenditure decisions in the household. The women's aspirations Z-score summarizes the level of income and assets that women aspire to have. A higher aspirations Z-score means the beneficiary has higher asset and income level aspirations.

3.2.1 Descriptive Statistics

In Table 3.1, we show the demographic characteristics of the beneficiaries in our study by gender. In baseline, on average, beneficiaries were 23 years old, with ages ranging from 16 to 31 and no differential between men and women. Both men and women have, on average, about 7.5 years of education, which means, on average, they have not graduated from high school. Women have, on average, larger households, which is consistent with women having more children. Economic and labor market outcomes show a clear differential between men and women. Men are more likely to be employed, work more hours, earn more income, consume more, and spend less time in domestic work. There is no differential by gender in productive assets, which may be due to the overall lack of productive assets amongst beneficiaries in the baseline. In the cultural dimension, we see no significant differences in the likelihood of marriage or cohabiting between men and women. Women are more likely to have children. Interestingly, women exhibit higher desired fertility. However, desire fertility is measured as the answer to the question, "What is the total number of children you would like to have in your lifetime, including those that you have already?". Hence, a significant factor in the differential in desired fertility is that women have more children.

Table 3.1. Descriptive Statistics by Gender

	All (SD)	Women (SD)	Men (SD)	P-value Of Difference
<i>Baseline</i>				
Age (years)	23.363 (3.50)	23.376 (3.54)	23.344 (3.45)	0.984
Beneficiary years of education	7.579 (2.17)	7.574 (2.17)	7.586 (2.18)	0.827
Household members	4.841 (2.28)	5.007 (2.25)	4.588 (2.29)	0.001
Employed	0.344 (0.48)	0.290 (0.45)	0.424 (0.49)	0.000
Productive hours	11.188 (18.93)	8.709 (15.86)	14.942 (22.30)	0.000
Monthly income	4.452 (4.95)	3.876 (4.81)	5.325 (5.04)	0.000
Productive assets	2.199 (4.19)	2.148 (3.98)	2.277 (4.49)	0.537
Beneficiary-specific consumption	7.512 (2.23)	7.160 (2.45)	8.044 (1.71)	0.000
Married	0.096 (0.29)	0.100 (0.30)	0.090 (0.29)	0.195
Cohabiting	0.057 (0.23)	0.057 (0.23)	0.057 (0.23)	0.805
Observations	1848	1113	735	
<i>Midline</i>				
Desired number of children	2.833 (0.75)	2.954 (0.81)	2.651 (0.62)	0.000
Any children	0.494 (0.50)	0.746 (0.44)	0.215 (0.41)	0.000
Hours Domestic Work	24.553 (16.44)	30.221 (16.14)	15.749 (12.60)	0.000
Observations	457	278	179	

Notes: Table presents baseline means (when available) and standard deviations for all beneficiaries, women and men, and p-value for a test of the hypothesis that compares men and women. Regression-based comparisons and associated hypothesis tests based on a regression with sector indicators and household-level clustered standard errors. *=10%, **=5%, and ***=1% significance. The dummies “Desired number of children”, “Any Children” and “Hours spent in domestic chores” were not measured in baseline, so we present midline figures for the control group. All continuous variables are winsorized at the top and bottom 1 percent. Inverse hyperbolic sine transformations were taken for monthly income, beneficiary consumption, and wealth variables.

3.3 Results

3.3.1 Economic Outcomes

First, we analyze the economic dimension of the differential effects between men and women. We focus on two sets of outcomes: productive time use and business characteristics. Strikingly, we find no differential impacts by gender. Several studies show that women experience different effects from receiving cash than men in economic outcomes [De Mel et al., 2012, Blattman et al., 2013]. However, we see that although gender plays a significant role in the level of all the economic outcomes, it plays no role in the magnitude of the effects of the interventions. Besides, female beneficiaries are less likely to be employed, work fewer hours, earn less income, and consume less. Hence, there could be more opportunities for them to catch up to the men as they receive cash and training.

Productive Time Use

Table 3.2 examines the productive time use outcomes. We analyze the probability of doing any work, independent work, being employed by an outside employer, and working on one's own farm. We see that women are significantly less likely to do any type of work and to be employees. However, the coefficients on the interactions between HD, GD, and the combined arm with the female indicator are small and not statistically significant. Overall, there are few effects on productive time use. GD increases the likelihood of independent work for everyone, independently of gender. HD has no significant effects, which means that the constraint of access to skills is not binding for these economic outcomes.

Business Characteristics

Next, in Table 3.3, we look at business characteristics, including the likelihood of co-ownership, the amount of debt, the number of employees, the value of assets, and profits. For this set of variables, we see that gender is less a determinant of the levels of the outcomes since the coefficient on the female indicator is non-significant across the board. All the interventions

Table 3.2. Productive Time Use with Gender Interactions

	Any Work	Independent	Employed	Farm Work
HD	0.03 (0.05) [1.00]	0.01 (0.03) [1.00]	0.01 (0.05) [1.00]	-0.03 (0.05) [1.00]
GD	-0.03 (0.05) [1.00]	0.11** (0.04) [0.02]	-0.11 (0.05) [0.15]	0.04 (0.05) [1.00]
Combined	0.00 (0.06) [1.00]	0.10 (0.05) [0.32]	-0.09 (0.06) [1.00]	0.01 (0.07) [1.00]
HD × Female	-0.00 (0.06) [1.00]	0.04 (0.04) [1.00]	-0.04 (0.06) [1.00]	0.04 (0.06) [1.00]
GD × Female	0.05 (0.06) [1.00]	-0.03 (0.04) [1.00]	0.06 (0.06) [1.00]	0.01 (0.06) [1.00]
Combined × Female	0.01 (0.08) [1.00]	-0.01 (0.06) [1.00]	0.02 (0.08) [1.00]	0.08 (0.08) [1.00]
Female	-0.29*** (0.04) [0.00]	-0.04 (0.03) [1.00]	-0.27*** (0.05) [0.00]	0.03 (0.05) [1.00]
Control mean	0.50	0.10	0.42	0.41
Observations	1822	1822	1822	1822
R^2	0.08	0.04	0.08	0.05
p -value	0.74	0.51	0.40	0.77

Notes: Table analyzes endline productive time use outcomes at the beneficiary level. Interacted coefficients in rows 4 to 6 give the differential effect of each arm for women, 'Female' gives the difference between women and men in the control group, and the uninteracted treatment terms give the impact of each arm for men. The first column is a dummy for whether the individual does any work at endline. Column 2 is a dummy for whether the individual is self-employed or an entrepreneur at endline. Column 3 analyzes a dummy for whether the beneficiary was an employee at endline. Column 4 is a dummy for whether the individual worked in their own farm at endline. Standard errors are clustered at the household level, *=10%, **=5%, and ***=1% significance.

increase the probability of co-ownership of their business as well as the value of the assets. Cash increases profits, and there are no effects on debt or the number of employees. However, when

we look at the interactions with the female indicator, we see non-differential impacts by gender all across the board. Women do not seem to build businesses in a different way compared to men.

Table 3.3. Business Characteristics with Gender Interactions

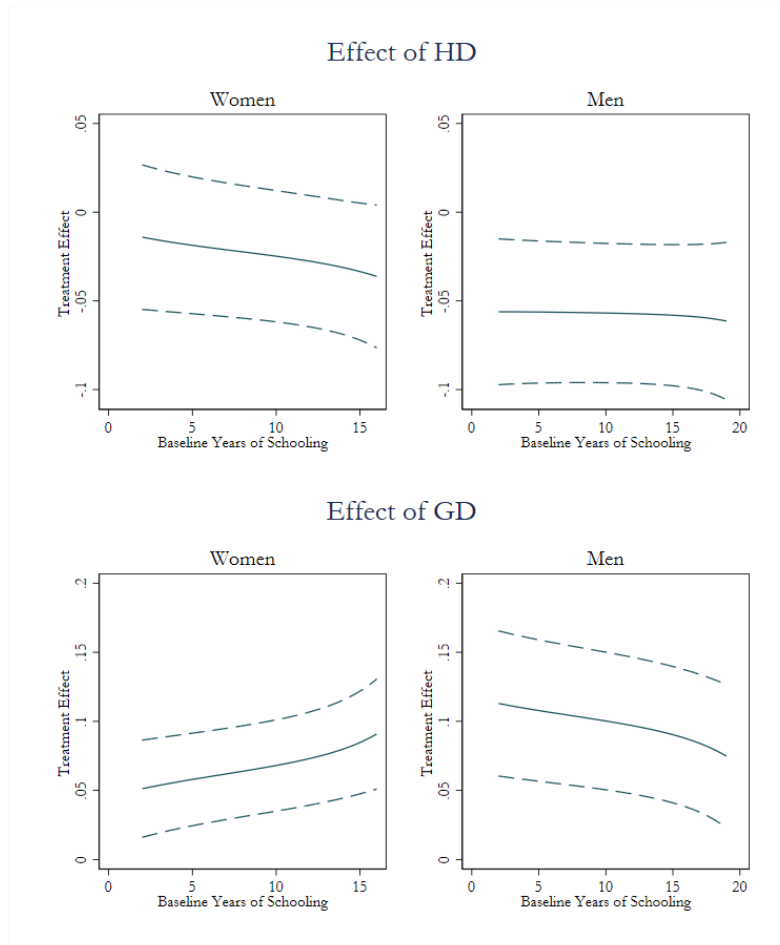
	Co-owned	Debt	Employee Number	Assets IHS	Profits IHS
HD	0.12* (0.05) [0.10]	0.24 (0.47) [0.70]	0.09 (0.08) [0.45]	1.41* (0.61) [0.10]	0.77 (0.39) [0.12]
GD	0.09* (0.04) [0.10]	0.51 (0.44) [0.40]	0.07 (0.05) [0.36]	2.66*** (0.56) [0.00]	0.83* (0.35) [0.10]
Combined	0.18* (0.07) [0.10]	0.90 (0.67) [0.36]	0.01 (0.06) [0.92]	3.53*** (0.80) [0.00]	1.25* (0.54) [0.10]
HD × Female	-0.12 (0.06) [0.12]	0.48 (0.61) [0.56]	-0.05 (0.09) [0.70]	-0.93 (0.76) [0.38]	-0.37 (0.49) [0.56]
GD × Female	-0.08 (0.05) [0.21]	0.13 (0.58) [0.89]	-0.04 (0.06) [0.70]	-0.57 (0.71) [0.56]	-0.42 (0.45) [0.50]
Combined × Female	-0.17 (0.08) [0.11]	-0.07 (0.86) [0.92]	0.04 (0.08) [0.70]	-0.79 (1.02) [0.56]	-0.65 (0.70) [0.50]
Female	0.04 (0.02) [0.21]	-0.02 (0.41) [0.92]	-0.08 (0.04) [0.10]	-0.11 (0.51) [0.89]	0.29 (0.31) [0.50]
Control mean	0.02	2.09	0.09	3.90	1.40
Observations	484	1816	1822	1822	1822
R^2	0.06	0.01	0.02	0.06	0.02
p -value	0.03	0.86	0.72	0.65	0.70

Notes: Table analyzes endline productive time use outcomes at the beneficiary level. Interacted coefficients in rows 4 to 6 give the differential effect of each arm for women, 'Female' gives the difference between women and men in the control group, and the uninteracted treatment terms give the impact of each arm for men. The first column is a dummy for whether the individual does any work at endline. Column 2 is a dummy for whether the individual is self-employed or an entrepreneur at endline. Column 3 analyzes a dummy for whether the beneficiary was an employee at endline. Column 4 is a dummy for whether the individual worked in their own farm at endline. Standard errors are clustered at the household level, *=10%, **=5%, and ***=1% significance.

The Role of Schooling

As part of the Pre-Analysis Plan, we specified triple heterogeneity specifications to understand the education level, marital status, having any children, and age. We find no differential effects by gender and any of these dimensions, except education level. In the first column of Appendix Table C.1, women with higher education levels see larger increases in the likelihood of doing any work than everyone else. In columns 2 and 3, we delve into this result by breaking down the work categories into independent work and wage work. More educated women are more likely to use cash to start their own business than everyone else. This is not true in wage work. To look into the potential reasons for this, we investigate the wage differentials in the wage work labor market. In Appendix Table C.2, we look at how much beneficiaries who do wage work earn in baseline by gender and educational attainment. In general, these men earn more than women. More interestingly, men's earnings do not vary very much across education levels. For women, surprisingly, earnings are lower for those who are more highly educated, than for those who have less education. This could be an indication of more educated women being discriminated against in the labor market. Hence, getting cash makes these women more likely to start their own businesses.

Figure 3.1 summarizes the main result about how schooling interacts with gender. We present how the effects of HD and GD change with the education level, separately for men and women, in the outcome of 'independent work'. We see no differential effects of HD by schooling for men and women. The GD effects are different. While for men, the line declines weakly, for women, we see an increase in the likelihood of independent work as they are more educated. An important caveat is that these results hold only in midline. The endline results in Appendix Table C.3 fade away. However, the endline collection happened during the pandemic, dampened many economic effects of the interventions.



Notes: Figure presents heterogeneity of the probability of independent work by gender and baseline years of schooling. The top panel examines the effects of HD, and the bottom panel delves into the effects of GD. We split the sample between men and women. For each sample split, we present a non-parametric estimation of the effects of additional years of schooling on the likelihood of being entrepreneurs or being self-employed.

Figure 3.1. Midline Heterogeneity of Independent Work by Schooling and Gender

3.3.2 Cultural Outcomes

We capture in our long panel data the period of time where the study subjects are forming their own households, getting married, and having children; understanding how these interventions may advance or retard that process is key to interpreting the effects in the cultural dimension. Given the pivotal role that gender plays in mediating the opportunities that young people have (both because of cultural expectations and also because of the relationship between childbearing and labor supply during this time of life), we are particularly interested in examining

male/female differentials for this set of outcomes.

Specifically, for the cultural outcomes, we analyze the likelihood of being married, the likelihood of cohabiting, the likelihood of having any children, and the desired fertility. To delve deeper into these outcomes, we also look at the likelihood of forming a new household, moving villages, and moving to urban areas. We also analyze the effects on the highest grade of education attained, hours spent in school, hours spent in domestic chores, and reservation wages. Finally, we analyze variables related to gender roles. This includes attitudes about men's and women's roles, women's financial control in the household, and women's aspirations.

Marriage and Fertility

First, we consider a set of variables related to marriage transitions and fertility. To study marriage, we examine whether an individual is married or cohabiting, as well as separate indicators for each of these two statuses. For fertility, we examine whether the individual has ever had any children, and we examine their survey-reported desired fertility (total number of children they hope to have including those they already have). Table 3.4 the gender interactions for this analysis, and Appendix Table C.4 shows the treatment effects. In C.4, we see muted and somewhat confusing results when we pool men and women, but then quite a clear and a clearly differential picture in the interaction analysis in Table 3.4. Money significantly affects allowing men to marry, not whether they cohabit. The sum of the male effect and the female differential is approximately zero, suggesting that money overall has no effect on women's proclivity to marry (and explaining why the pooled results are insignificant). Perhaps unsurprisingly, the story of whether individuals have any children is very similar; money amplifies males' probability of having kids. The effect is significantly less likely to do so for women, resulting in a net effect of zero for women (this is even though the women in our sample overall are 38% more likely to have a child than the men). The cash arm weakly amplifies the desired fertility of males, and here, the significant negative effects are strong enough to mean that on the net, there is a mild overall depressive effect of cash on female desired fertility.

To investigate possible pathways for this latter set of effects, we examine outcomes around education choices and valuation of time, both of which would be key inputs to a Beckerian consideration of fertility choices [Becker, 1965]. First, in Table 3.5 and Appendix Table 3.5) we examine completed schooling. The intent of the program was that it was enrolling individuals who are old enough to have completed their schooling and so neither the cash arm or the HD arm were intended to generate education other than through HD itself. We confirm that this is true, showing that none of the programs have an effect either on the years of completed schooling at endline nor on the time put in to schooling. This null effect is not differential by gender. A different use of time is that spent doing household chores, an activity that we confirm women devote hugely more time to on average than men (30 hours per week for women, 16 for men in the control group). We have already seen that the interventions move productive time use to self-employment (cash arm). Here we see that the interventions have no effect on decreasing overall time in domestic chores. Importantly, we show that HD differentially decreases time in chores for women relative to men by about 4.5 hours per week, suggesting that this program has a labor empowerment effect that closes about a quarter of the gap between male and female time in domestic work. Finally, we use survey questions that asked the beneficiary how much money (as a daily wage) they would need to be paid to accept a job in their village and in the nearest town.² These provide a survey measure, albeit unincentivized, of the opportunity cost of time. If the pathway to the impacts on female desired fertility operated through how she perceives her time on the margin, we would expect to see it here. The results are quite clear that this is not the case; we see no pattern of the treatments increasing the opportunity cost of time overall or for women specifically. This is consistent with the idea that the effective wage rate being achieved in project-created businesses is not out of line with the counterfactual returns they would have achieved in the absence of the program.

Taken together, these results are interesting in a number of dimensions. First, Huguka

²The questions were asked in the form ‘would you be willing to accept 1000 RwF’, ‘if not would you accept 2000 RwF’, ‘if not 4000 RwF’, and if still no then ‘how much would you have to be paid to accept this job’.

Dukore, despite content focusing on family planning and HIV, did not move marriage or observed fertility, although it did have a weak depressive effect on desired fertility, particularly among women. Second, we see evidence that men in Rwanda are income-constrained in marrying, and when this constraint is relaxed they move more quickly to formally marry and to have children. These constraints do not appear to bind in the same way for women (as would be consistent with brideprices from grooms to brides being culturally typical). The bride price culture pressures men and their families to provide money or assets to formalize marriages, which are a key social institution. Finally, while we have not found impacts of these interventions on the economic aspirations of youth or the opportunity costs of time, the considerable effects of cash transfers both on entrepreneurship and on desired fertility for women do suggest a pathway whereby relaxing credit constraints increases young women's economic prospects and thereby alters the way they think about childbearing in a dynamic way.

Household Transitions

Next, we analyze the differences in how individuals form new households as they age out of adolescence and start financially independent adult lives. Do these individuals, who were typically dependents in their parents' households at baseline, become household heads (or spouses of household heads) themselves? A less stringent measure is to examine the share of beneficiaries who have moved away from their baseline households by the time of the ending. The first two columns in Table 3.6 examines these outcomes. HD has no significant effect on any measure of new household formation. The cash arm paints a more complicated picture; overall Appendix Table C.6 there is an elevation of 11 percentage points in the likelihood that the beneficiary lives in a new household at endline. Still, there is no change in the probability of being the head of that household (these effects are significant before controlling for multiple inference).³ We do not see any differential effects by gender.

³Additional analysis (not reported) shows that these new households are no smaller on average (4.7 members in either case), and so it looks more like cash is causing beneficiaries to move to different locations at which opportunities to run a business are improved, rather than actually establishing independent households.

To understand whether the cash-induced movements send youth from that arm into more urban settings, we examine movements from one village to another across rounds. First, we can create a dummy in the midline and endline for being in a different location altogether than the baseline. Then, we match the district, sector, cell, and village of residence in each survey to an official Rwandan government classification of these locations as urban, peri-urban, semi-rural, or rural, and examine treatment effects on the classification of their locations. The third and fourth columns in Table 3.6. Again, we see no differential effects by gender. Appendix Table C.6 shows only a weakly elevated probability from the cash arm of moving across villages, suggesting that about half of the household switching shown previously is within the village. When we examine how urban the locations are, we find that against an overall upward trend in the control (at baseline only 12% of control individuals live outside of rural villages, a rate which rises to 18% in the midline and 22% in the endline) the cash arms have no effect on the movement of people to urban areas. In sum, then, HD plays no strong role in this type of mobility, and while cash encourages people to change households, this is often within village and rarely involves the beneficiary establishing an independent household.

Gender Roles

Finally, we analyze variables related to gender roles, such as attitudes about men's and women's roles, women's financial control in the household, and women's aspirations. Table 3.7 shows that the interventions have no effects on attitudes towards gender or female empowerment. These results are also striking, considering that the literature finds effects on female empowerment, such as financial control. The evidence in this literature shows that cash transfers increase the probability of women making expenditure decisions in their household for Bolsa Familia, PROGRESA, and HSNP [De Brauw et al., 2014, Handa et al., 2009, Merttens et al., 2013].

Table 3.7 shows effects in gender attitudes for men and women separately. The gender attitudes Z-score summarizes the beneficiary's responses to statements related to the importance

of women's education and labor market prospects. A higher gender attitudes' Z-score means that the beneficiary considers women's education more important. The first and third rows show no changes in this dimension. The second row summarizes women's agency over household financial decisions. A higher financial control Z-score means that the beneficiary has more agency over expenditure decisions in the household. Again, we see no changes in women's decision to spend money in the household despite them bringing in a significant amount. Finally, in the last row, the women's aspirations Z-score summarizes the level of income and assets that women aspire to have. A higher aspirations Z-score means the beneficiary has higher asset and income level aspirations. Although these effects are larger, they are noisy and non-significant. Overall, this shows that the changes we see in marriage, fertility, and desired fertility do not stem from changes in the roles, perceptions, and values of beneficiaries towards gender. Rather, the results are more consistent with young people, especially men, facing financial constraints to get married and start a family.

Table 3.4. Marriage and Fertility with Gender Interactions

	Married	Cohabiting	Any Children	Desired Fertility
HD	0.02 (0.04) [0.56]	-0.03 (0.04) [0.45]	0.02 (0.05) [0.65]	-0.00 (0.07) [0.83]
GD	0.11*** (0.04) [0.01]	0.00 (0.04) [0.83]	0.08 (0.05) [0.11]	0.14* (0.06) [0.07]
Combined	0.17** (0.06) [0.01]	0.01 (0.05) [0.77]	0.18** (0.06) [0.02]	0.08 (0.09) [0.37]
HD × Female	-0.05 (0.05) [0.27]	0.06 (0.05) [0.27]	-0.03 (0.06) [0.58]	-0.18* (0.09) [0.09]
GD × Female	-0.10* (0.05) [0.08]	0.05 (0.05) [0.31]	-0.10 (0.06) [0.12]	-0.27*** (0.09) [0.01]
Combined × Female	-0.18** (0.07) [0.03]	-0.02 (0.07) [0.65]	-0.14 (0.08) [0.12]	-0.30** (0.12) [0.03]
Female	0.12*** (0.03) [0.01]	0.01 (0.04) [0.69]	0.38*** (0.04) [0.00]	0.32*** (0.06) [0.00]
Control mean	0.20	0.22	0.58	2.87
Observations	1822	1822	1822	1822
R^2	0.06	0.04	0.12	0.02

Notes: Table analyzes endline marriage and fertility outcomes at the beneficiary level. Interacted coefficients in rows 4 to 6 give the differential effect of each arm for women, ‘Female’ gives the difference between women and men in the control group, and the uninteracted treatment terms give the impact of each arm for men. The first column is a dummy for whether the individual was married at endline. Column 2 is a dummy for whether the individual was cohabiting at endline. Column 3 analyzes a dummy for whether the beneficiary has any children as of the time of the endline. Column 4 uses the response to the question “what is the total number of children you would like to have in your lifetime, including those that you have already”. Standard errors are clustered at the household level, *=10%, **=5%, and ***=1% significance.

Table 3.5. Education and Time Use with Gender Interactions

	Highest Grade	School Hours	Domestic Hours	Wage Village	Wage Town
HD	0.36 (0.40) [1.00]	-0.04 (0.80) [1.00]	2.05 (1.33) [0.98]	-0.09 (0.17) [1.00]	-0.48 (0.25) [0.63]
GD	0.36 (0.37) [1.00]	-0.24 (0.71) [1.00]	0.40 (1.00) [1.00]	1.18 (0.88) [0.98]	1.34 (1.11) [1.00]
Combined	0.66 (0.54) [1.00]	-0.63 (0.81) [1.00]	-0.54 (1.52) [1.00]	0.44 (0.22) [0.63]	0.09 (0.31) [1.00]
HD × Female	-0.14 (0.50) [1.00]	-1.01 (0.94) [1.00]	-4.49 (2.03) [0.63]	0.02 (0.22) [1.00]	0.40 (0.37) [1.00]
GD × Female	-0.35 (0.46) [1.00]	-0.26 (0.93) [1.00]	-1.96 (1.74) [1.00]	-1.32 (0.89) [0.98]	-1.51 (1.14) [0.98]
Combined × Female	-1.15 (0.66) [0.69]	0.27 (1.18) [1.00]	2.03 (2.62) [1.00]	-0.50 (0.25) [0.63]	0.32 (0.45) [1.00]
Female	0.14 (0.35) [1.00]	-0.45 (0.76) [1.00]	19.44*** (1.29) [0.00]	-0.15 (0.15) [1.00]	-0.37 (0.28) [0.98]
Control mean	12.07	1.56	24.31	1.81	2.83
Observations	1822	1822	1822	1810	1804
R^2	0.05	0.01	0.23	0.02	0.02

Notes: Table presents tests for heterogeneity of education and time use effects by gender. Interacted coefficients in rows 4 to 6 give the differential effect of each arm for women, 'Female' gives the difference between women and men in the control group, and the uninteracted treatment terms give the impact of each arm for men. Standard errors are clustered at the household level, *=10%, **=5%, and ***=1% significance.

Table 3.6. Creation of New Households and Moving Across Villages with Gender Interactions

	Formed Household	Head Spouse	New Village	Urban
HD	0.03 (0.05) [1.00]	0.02 (0.05) [1.00]	0.03 (0.05) [1.00]	0.05 (0.03) [1.00]
GD	0.08 (0.05) [1.00]	0.07 (0.05) [1.00]	-0.03 (0.04) [1.00]	0.01 (0.03) [1.00]
Combined	-0.03 (0.06) [1.00]	0.12 (0.06) [1.00]	0.00 (0.06) [1.00]	-0.01 (0.04) [1.00]
Female x HD	0.02 (0.06) [1.00]	0.00 (0.06) [1.00]	-0.03 (0.06) [1.00]	-0.04 (0.04) [1.00]
Female x GD	0.04 (0.06) [1.00]	-0.05 (0.06) [1.00]	0.11 (0.06) [1.00]	-0.00 (0.04) [1.00]
Female x Combined	0.10 (0.08) [1.00]	-0.06 (0.08) [1.00]	-0.01 (0.08) [1.00]	0.02 (0.05) [1.00]
Female	0.03 (0.05) [1.00]	0.03 (0.04) [1.00]	0.14** (0.04) [0.04]	-0.03 (0.03) [1.00]
Control mean	0.40	0.63	0.37	0.08
Observations	1822	1822	1822	1822
R^2	0.03	0.02	0.05	0.05

Notes: Table presents tests for heterogeneity of marriage and fertility effects by Gender. Interacted coefficients in rows 4 to 6 give the differential effect of each arm for women, 'Female' gives the difference between women and men in the control group, and the uninteracted treatment terms give the impact of each arm for men. Standard errors are clustered at the household level, *=10%, **=5%, and ***=1% significance.

Table 3.7. Gender Attitudes

	GiveDirectly			Control Mean	Obs.	R ²
	HD	GD	Combined			
Z-score: women's gender attitudes (midline)	-0.01 (0.09) [1.00]	-0.08 (0.08) [1.00]	-0.16 (0.12) [1.00]	0.01	1064	0.10
Z-score: women's financial control (midline)	-0.01 (0.07) [1.00]	0.02 (0.07) [1.00]	0.08 (0.09) [1.00]	0.00	1064	0.38
Z-score: men's gender attitudes (midline)	-0.03 (0.12) [1.00]	-0.06 (0.10) [1.00]	0.17 (0.14) [1.00]	-0.01	690	0.19
Z-score: women's aspirations	0.09 (0.09) [1.00]	0.10 (0.08) [1.00]	0.15 (0.10) [1.00]	-0.00	1100	0.10

Notes: Table analyzes midline gender attitudes variables, and endline aspirations. Gender attitudes Z-scores in the first and third row summarize the beneficiary's responses to statements related to the importance of women's education with respect to men's education, to marriage, and to their labor market prospects. A higher gender attitudes' Z-score means that the beneficiary considers women's education to be more important. The women's financial control Z-score in the second row summarizes the agency that women have over financial decisions in their household. A higher financial control Z-score means that the beneficiary has more agency over expenditure decisions in the household. The women's aspirations Z-score in the fourth row summarizes the level of income and assets that women aspire to have. A higher aspirations Z-score means that the beneficiary has higher asset and income level aspirations. Standard errors are clustered at the household level, *=10%, **=5%, and ***=1% significance.

3.4 Concluding Remarks

Our study addresses how interventions to enhance the economic prospects of underemployed Rwandan youth impact their transitions to adulthood, focusing on gender dynamics. We analyze the effectiveness of two types of interventions – cash transfers and vocational training programs – in mitigating young people’s financial and skills constraints as they enter the labor market and form new households. We analyze data from a randomized controlled trial to assess the differential effects of these interventions on various economic and cultural outcomes, including employment, marriage, fertility, gender attitudes, and aspirations.

Our findings underscore the complex interplay between economic interventions and social norms in shaping the transitions to adulthood for young people in Rwanda. While the lack of significant gender differences in economic outcomes suggests that both young men and women experience similar effects when provided with access to resources and opportunities, it does not necessarily indicate that they overcome barriers in the same way. Rather, our results reveal that preexisting disparities persist despite the interventions. Besides, the persistence of traditional gender roles, despite these interventions, highlights the deep-rooted nature of cultural norms that cannot be easily altered through economic means alone. HD, a program promoted as a tool of gender empowerment, does not achieve the goal of advancing in women’s economic progress.

Moreover, the gender-specific impacts on marriage and fertility rates provide nuanced insights into the societal dynamics at play. Cash transfers lead to increased marriage and fertility rates among men but not women, reflecting the bride price culture’s influence, which prioritizes men’s financial stability before marriage. In contrast, the decrease in desired fertility among women who receive cash transfers indicates a shift in how young women perceive their future family size, possibly influenced by increased economic opportunities and autonomy. These findings align with the predictions of economic theories that link financial security with fertility decisions. Yet, they also reveal the limitations of economic interventions in transforming social

expectations and behaviors.

Ultimately, our study demonstrates that while financial constraints are a significant barrier for young people in Rwanda, addressing these alone does not alter deeply entrenched gender roles and cultural norms. The persistence of traditional gender attitudes despite substantial economic support suggests that more comprehensive approaches are needed to foster genuine gender equality and social change.

Chapter 3, in full, is currently being prepared for submission for publication of the material. Martinez Heredia, Diana; McIntosh, Craig; Zeitlin, Andrew. The dissertation author was the primary investigator and author of this material.

Appendix A

Immigrant children and optimal school choice: Evidence from the Venezuelan migration to Peru


A.1 Venezuelan crisis and migration timeline

For many years before Nicolás Maduro's presidency in Venezuela in 2013, South Americans migrated to Venezuela looking for better economic opportunities. This tendency has completely reversed, as Venezuela has fallen into one of the greatest economic crises of recent economic history. Hyperinflation and poverty were already concerning when, in May of 2017, Maduro called a Constitutional Assembly. The Venezuelan opposition and the international community rejected this. Regardless, in that Constitutional Assembly, Congress was dissolved. The opposition-held majority in Congress served as a check on Maduro's government, and they lost most of the power they held in 2017. The economic situation only worsened. The IMF reported that hyperinflation reached 65,000% in 2018, and poverty affected about 79% of the population (ENCOVI). Diseases like measles, diphtheria, tuberculosis, and malaria have spread rapidly. The shortage of food and goods for basic needs has been pervasive. Additionally, increasing crime and security issues have forced Venezuelans out of their country.

Corruption and precariousness in Venezuela made it almost impossible for Venezuelans

to emigrate with updated documents. Migration offices in Venezuela could take years to issue a passport or charge large amounts of money to issue them in a reasonable time frame. Migrants also had to present a criminal record that Interpol offices can issue for about 25 USD and legally enter the country (tourists). At first, only migrants who arrived in Peru before December 2016 could apply, but, given the high demand, the Peruvian government expanded the PTP policy several times to allow migrants to legalize their stay even if they came later into the country. As more Venezuelans came, xenophobia proliferated, and the policies for Venezuelan migrants became unpopular. On August 25, 2018, only immigrants with unexpired passports could legally enter the country, increasing illegal immigration. However, after meeting all countries affected by Venezuelan migration in September of that same year, the Peruvian government reversed this change. They allowed Venezuelans with expired passports into the country. The government would reverse this again in June 2019 and allow only Venezuelan migrants with passports.

A.2 Figures



Data between 2014 to 2018 comes from the 2018 nationally representative survey of Venezuelan migrants in Peru ENCEVE. Data in 2019 from the Peruvian Migration authority

Figure A.1. Venezuelan Migration by Year

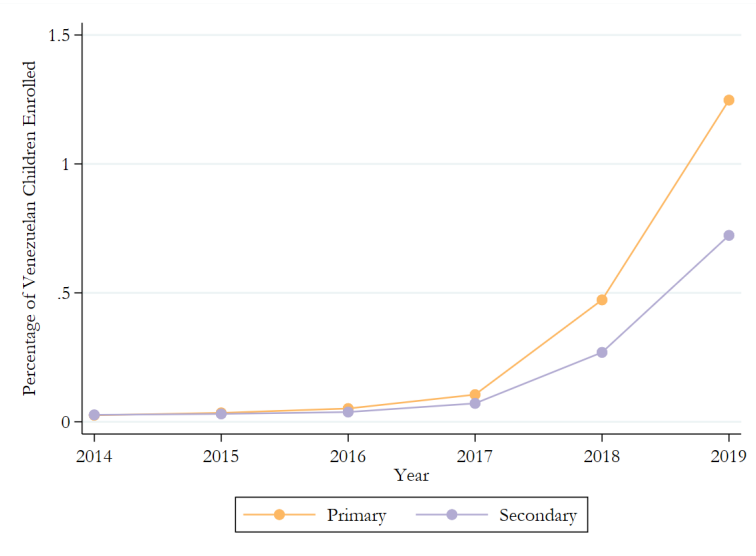


Figure A.2. Venezuelan Migrants' Enrollment by Year



Figure A.3. Venezuelan Children Enrolled in Schools in 2019




Figure A.4. Student Turnover by level


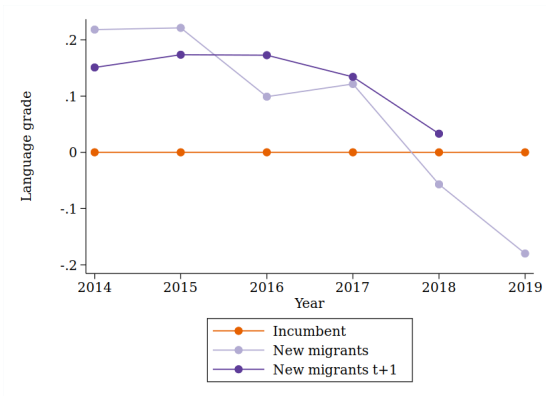
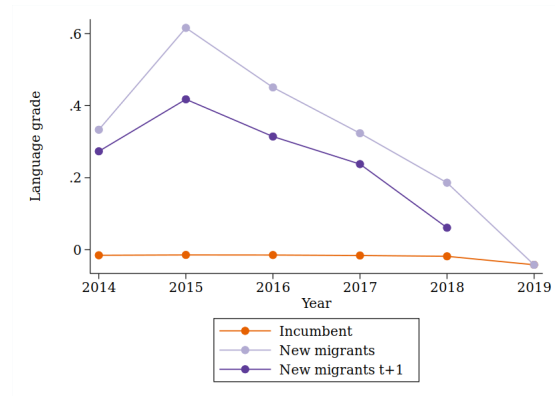


Figure A.5. Student Turnover: Incumbents vs Migrants



(a) Primary



(b) Secondary

Figure A.6. Venezuelan Migrants Performance in Language by Year




Figure A.7. Age Distribution by Grade in Primary

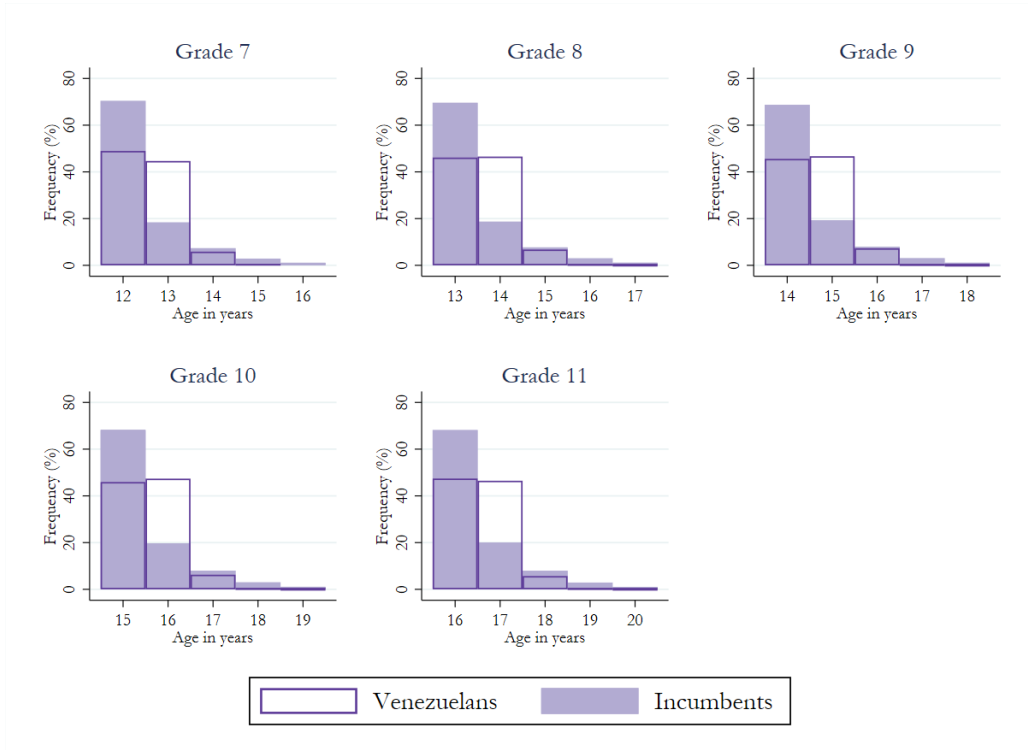


Figure A.8. Age Distribution by Grade in Secondary

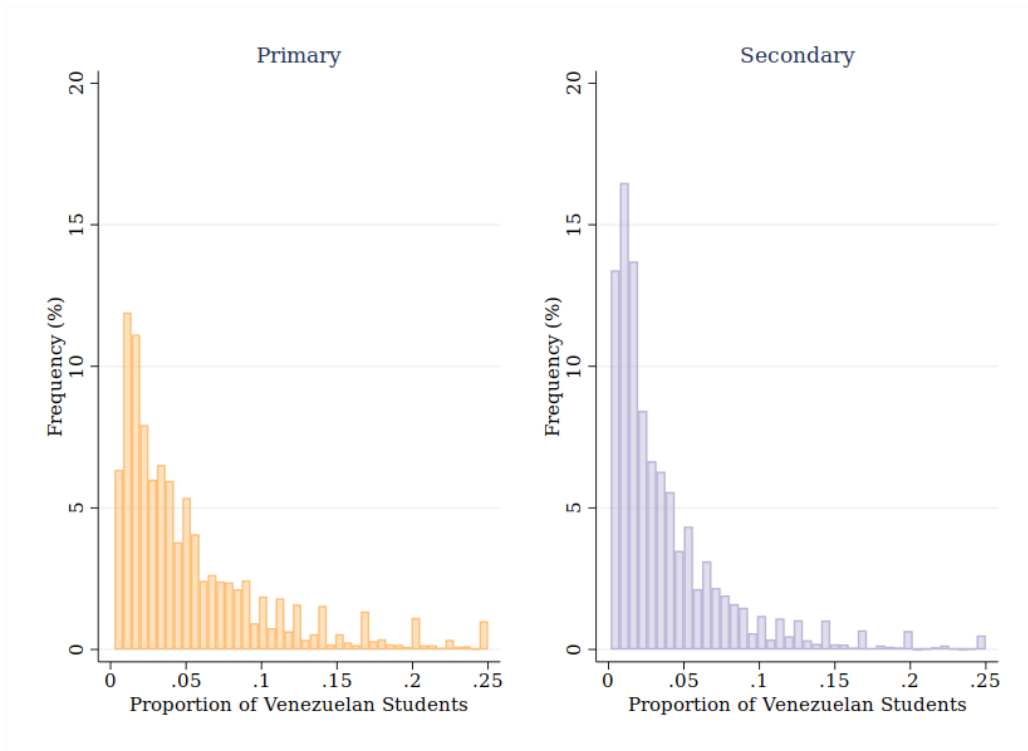


Figure A.9. Treatment Distribution

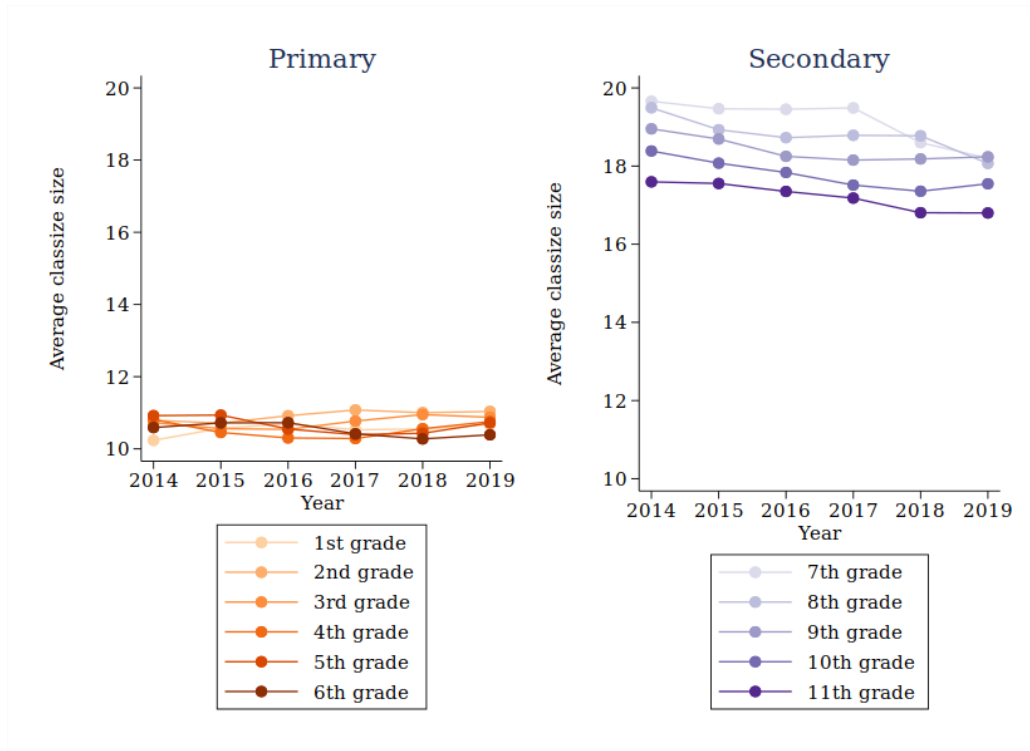


Figure A.10. Class Size

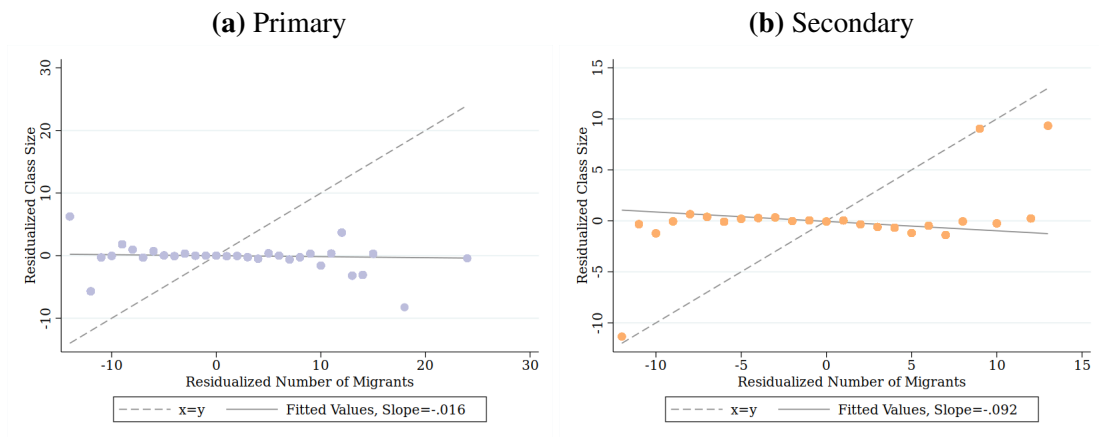


Figure A.11. Residualized Share of Venezuelan Migrants and Residualized Class Size

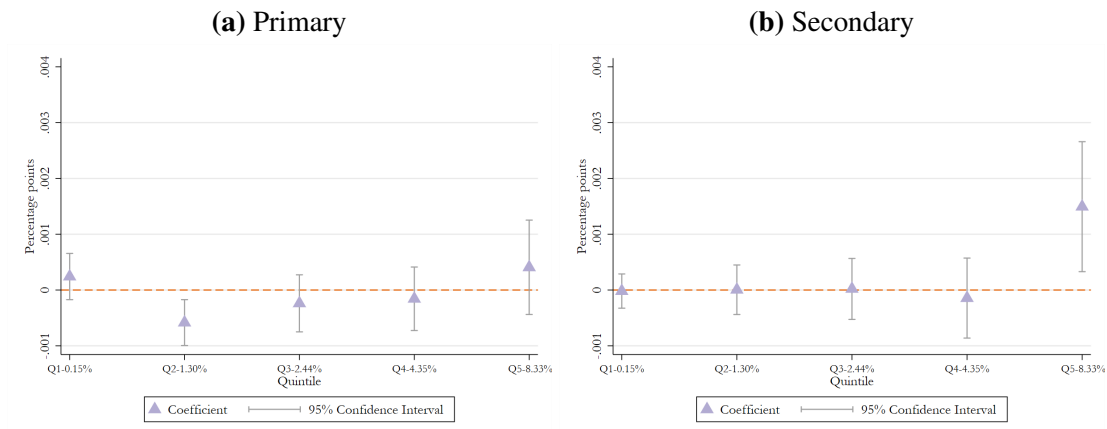


Figure A.12. Non-linear Effects of Migrant Exposure on the Probability of Retention

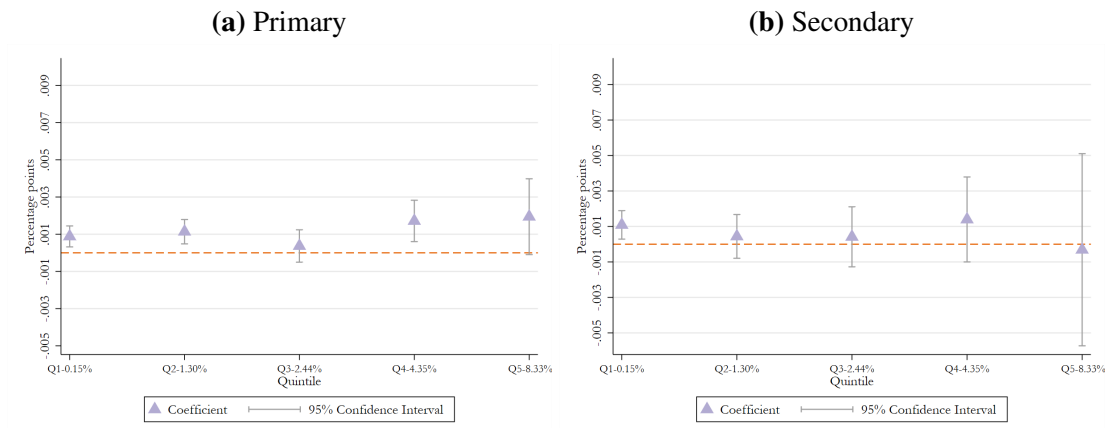


Figure A.13. Non-linear Effects of Migrant Exposure on the Probability of Dropout

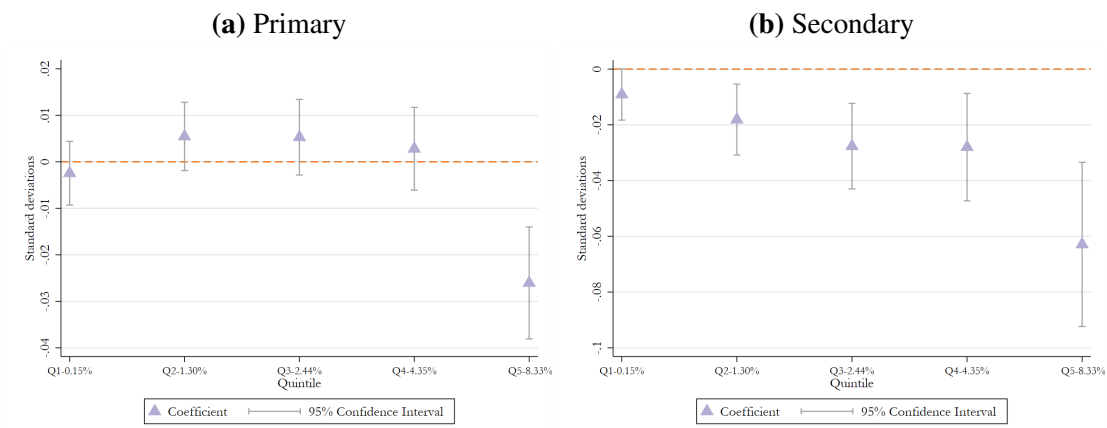


Figure A.14. Non-linear Effects of Migrant Exposure on Math Grades

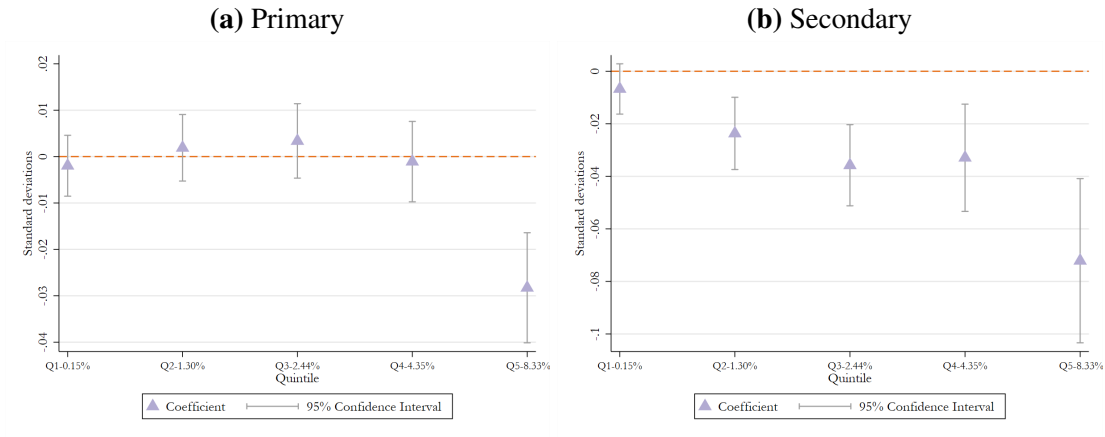


Figure A.15. Non-linear Effects of Migrant Exposure on Language Grades

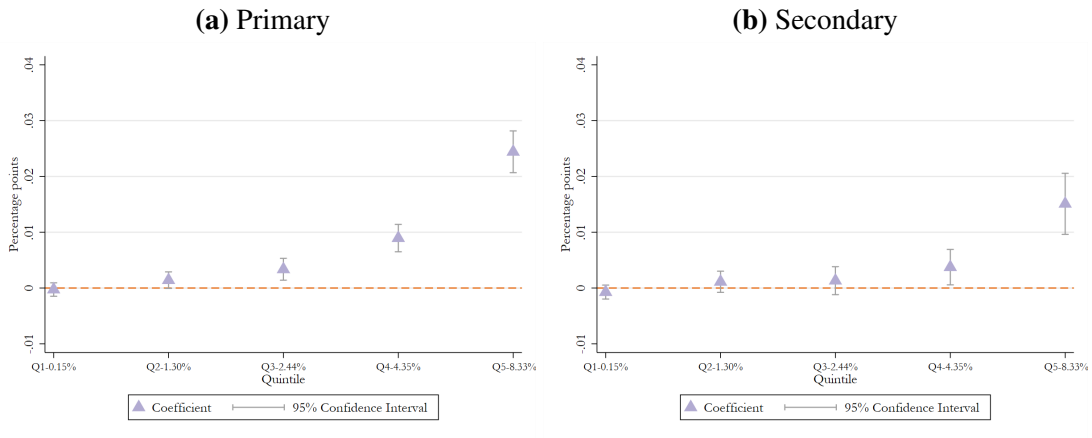


Figure A.16. Non-linear Effects of Migrant Exposure on the Probability of Switching Schools

A.3 Tables

Table A.1. Mig. Share Quintiles Range

Quintile	Range
1	$0 < V_{sg,t} \leq 0.013$
2	$0.13 < V_{sg,t} \leq 0.024$
3	$0.24 < V_{sg,t} \leq 0.043$
4	$0.43 < V_{sg,t} \leq 0.083$
5	$V_{sg,t} > 0.083$

Table A.2. Switching Schools Location

	Primary			
	Region (1)	Province (2)	District (3)	Distance ml (4)
Mig. Share	0.018*** (0.005)	0.023*** (0.005)	0.081*** (0.009)	6.619*** (1.581)
R-squared	0.045	0.053	0.062	0.042
Obs.	14,538,271	14,538,271	14,538,271	14,535,136
Mean	.023	.032	.061	6.204
	Secondary			
	Region (1)	Province (2)	District (3)	Distance ml (4)
Mig. Share	0.019*** (0.007)	0.021*** (0.007)	0.054*** (0.014)	-0.863 (1.895)
R-squared	0.057	0.072	0.078	0.052
Obs.	12,510,851	12,510,851	12,510,851	12,509,229
Mean	.016	.024	.047	4.309

Standard errors clustered at school-grade level in parentheses. *** p<0.01 ** p<0.05 * p<0.1. Control variables: sex, age, baseline math grade, school by year fixed, grade, year, and district fixed effects. All the outcome variables are dummies defined by difference in school location after switching. By definition is 0 for all non-switchers. The geographical distance between schools is calculated using each school latitude and longitude. The sample includes only incumbent students from 2015 to 2019 in Peruvian schools.

Table A.3. Effects of Migrant Enrollment in Class Size

Primary		
	(1)	(2)
	Classize > Max	Classize > Max
Number of Mig.	-0.00469*** (0.00126)	-0.0389*** (0.00155)
Number of Mig. × Public		0.0391*** (0.00127)
R-squared	0.415	0.417
Obs.	3,300,746	3,296,905
Mean	.354	.354

Secondary		
	(1)	(2)
	Classize > Max	Classize > Max
Number of Mig.	-0.00262** (0.00109)	-0.000364 (0.00151)
Number of Mig. × Public		-0.00251* (0.00137)
R-squared	0.406	0.406
Obs.	2,726,419	2,724,032
Mean	.147	.147

Standard errors clustered at school-grade level in parentheses. *** p<0.01 ** p<0.05 * p<0.1. This regressions are run at the classroom level. Class Size ζ Max is a dummy variable that takes the value of one is the classroom size is larger that 30 in primary and larger that 35 in secondary. Control variables: percentage of female students, mean age, school by year fixed, grade, year, and district fixed effects. The sample includes all classrooms from 2015 to 2019 in primary and secondary schools.

Table A.4. Effects of Migrant Exposure on Schooling Outcomes Adding Class Size as a Control

Primary				
	Retention (1)	Dropout (2)	Math Grade (3)	Language Grade (4)
Mig. Share	0.00385 (0.00246)	0.0193*** (0.00567)	-0.141*** (0.0343)	-0.171*** (0.0334)
R-squared	0.031	0.105	0.181	0.183
Obs.	14,543,841	11,568,653	14,423,528	14,423,649
Mean	.005	.013	-.014	-.011

Secondary				
	Retention (1)	Dropout (2)	Math Grade (3)	Language Grade (4)
Mig. Share	0.00858** (0.00367)	0.0179 (0.0125)	-0.643*** (0.0849)	-0.719*** (0.0884)
R-squared	0.019	0.102	0.267	0.275
Obs.	12,513,766	9,970,220	12,034,217	12,098,403
Mean	.005	.033	-.044	-.024

Standard errors clustered at school-grade level in parentheses. *** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$. Control variables: class size, sex, age, baseline math grade, school by year fixed, grade, year, and district fixed effects. The sample includes only incumbent students from 2015 to 2019 in primary and secondary schools.

Table A.5. Effects of Migrant Exposure in the Probability of Switching Schools Adding Class Size as a Control

	Primary (1)	Secondary (2)
Mig. Share	0.154*** (0.0126)	0.111*** (0.0193)
R-squared	0.087	0.095
Obs.	14,543,842	12,513,766
Mean	.115	.08

Standard errors clustered at school-grade level in parentheses. *** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$. Control variables: class size, sex, age, baseline math grade, school by year, grade, year, and district fixed effects. The sample includes only incumbent students from 2015 to 2019 in primary and secondary schools.

Table A.6. Effect of Migrant Exposure in the Probability of Switching Schools - Sample Split by Resource Level

	Low Resources					
	Primary			Secondary		
	Public (1)	Q1 Parents SES (2)	Q4 Stud/Teach (3)	Public (4)	Q1 Parents SES (5)	Q4 Stud/Teach (6)
Mig, Share	0.101*** (0.0162)	0.160*** (0.0244)	0.158*** (0.0248)	0.0278 (0.0232)	0.185*** (0.0465)	0.248*** (0.0529)
R-squared	0.186	0.236	0.332	0.198	0.216	0.484
Obs.	10,869,828	7,828,129	6,474,512	9,365,049	6,632,086	1,824,472
Mean	.094	.096	.102	.062	.054	.101

	High Resources					
	Primary			Secondary		
	Private (1)	Q4 Parents SES (2)	Q1 Stud/Teach (3)	Private (4)	Q4 Parents SES (5)	Q1 Stud/Teach (6)
Mig, Share	0.175*** (0.0185)	0.199*** (0.0232)	0.132*** (0.0240)	0.184*** (0.0252)	0.108*** (0.0322)	0.134*** (0.0254)
R-squared	0.200	0.356	0.483	0.249	0.358	0.435
Obs.	3,658,108	3,318,778	1,280,243	3,139,726	1,882,514	2,420,851
Mean	.174	.118	.168	.132	.116	.096

Standard errors clustered at school-grade level in parentheses. *** p<0.01 ** p<0.05 * p<0.1. Control variables: sex, age, baseline math grade, school by year, grade, year, and district fixed effects. The sample is split in 3 different ways using school-level characteristics. Column(1) splits the sample between private and public schools. Column (2) splits the sample between the first and fourth quintile of the parent's SES index in 2016. Column (3) splits the sample between the first and fourth quintile of the Student-Teacher ratio in 2014. For the Student-Teacher ratio, Q1 is in the high resources panel because this indicator reflects higher resources when it is lower. The sample includes only incumbent students from 2015 to 2019 in primary and secondary schools.

Table A.7. Effect of Migrant Exposure in Standardized Math Std. Grades - Sample Split by Resource Level

	Low Resources					
	Primary			Secondary		
	Public (1)	Q1 Parents SES (2)	Q4 Stud/Teach (3)	Public (4)	Q1 Parents SES (5)	Q4 Stud/Teach (6)
Mig, Share	-0.149*** (0.0506)	0.00125 (0.0704)	-0.0730 (0.0647)	-0.938*** (0.146)	-0.573** (0.254)	-0.524*** (0.187)
R-squared	0.172	0.168	0.185	0.240	0.246	0.312
Obs.	10,781,415	7,763,155	6,419,665	9,006,364	6,369,785	1,750,535
Mean	-.062	-.043	-.071	-.115	-.125	.028

	High Resources					
	Primary			Secondary		
	Private (1)	Q4 Parents SES (2)	Q1 Stud/Teach (3)	Private (4)	Q4 Parents SES (5)	Q1 Stud/Teach (6)
Mig, Share	-0.0534 (0.0439)	0.00493 (0.0599)	-0.0338 (0.0628)	-0.215*** (0.0791)	-0.301*** (0.113)	-0.176* (0.106)
R-squared	0.236	0.236	0.245	0.330	0.350	0.325
Obs.	3,626,323	3,286,682	1,265,721	3,018,870	1,805,606	2,324,540
Mean	.127	-.007	.091	.167	.204	.037

Standard errors clustered at school-grade level in parentheses. *** p<0.01 ** p<0.05 * p<0.1. Control variables: sex, age, baseline math grade, school by year, grade, year, and district fixed effects. The sample is split in 3 different ways using school-level characteristics. Column(1) splits the sample between private and public schools. Column (2) splits the sample between the first and fourth quintile of the parent's SES index in 2016. Column (3) splits the sample between the first and fourth quintile of the Student-Teacher ratio in 2014. For the Student-Teacher ratio, Q1 is in the high resources panel because this indicator reflects higher resources when it is lower. The sample includes only incumbent students from 2015 to 2019 in primary and secondary schools.

Table A.8. Effect of Migrant Exposure in Language Std. Grades - Sample Split by Resource Level

	Low Resources					
	Primary			Secondary		
	Public (1)	Q1 Parents SES (2)	Q4 Stud/Teach (3)	Public (4)	Q1 Parents SES (5)	Q4 Stud/Teach (6)
Mig, Share	-0.185*** (0.0499)	-0.0976 (0.0693)	-0.0803 (0.0635)	-1.101*** (0.154)	-1.163*** (0.264)	-0.438** (0.194)
R-squared	0.172	0.168	0.188	0.249	0.257	0.327
Obs.	10,781,526	7,763,294	6,419,809	9,037,146	6,383,729	1,765,069
Mean	-.066	-.048	-.065	-.102	-.123	.056

	High Resources					
	Primary			Secondary		
	Private (1)	Q4 Parents SES (2)	Q1 Stud/Teach (3)	Private (4)	Q4 Parents SES (5)	Q1 Stud/Teach (6)
Mig, Share	-0.0828** (0.0413)	0.00803 (0.0559)	-0.0460 (0.0615)	-0.170** (0.0756)	-0.226** (0.110)	-0.112 (0.101)
R-squared	0.231	0.244	0.245	0.336	0.352	0.342
Obs.	3,626,332	3,286,680	1,265,716	3,052,282	1,832,999	2,335,004
Mean	.152	-.006	.088	.209	.263	.048

Standard errors clustered at school-grade level in parentheses. *** p<0.01 ** p<0.05 * p<0.1. Control variables: sex, age, baseline math grade, school by year, grade, year, and district fixed effects. The sample is split in 3 different ways using school-level characteristics. Column(1) splits the sample between private and public schools. Column (2) splits the sample between the first and fourth quintile of the parent's SES index in 2016. Column (3) splits the sample between the first and fourth quintile of the Student-Teacher ratio in 2014. For the Student-Teacher ratio, Q1 is in the high resources panel because this indicator reflects higher resources when it is lower. The sample includes only incumbent students from 2015 to 2019 in primary and secondary schools.

Table A.9. Effect of Migrant Exposure in Math and Language Std. Grades - Sample Split by Switching Status

	Primary					
	Non-Switchers		Switchers			
	Math grades (1)	Language grades (2)	Math grades (3)	Language grades (4)	Math grades (5)	Language grades (6)
Mig. Share	-0.218*** (0.0595)	-0.263*** (0.0578)	0.0251 (0.0435)	-0.0473 (0.0419)		
Mig. Share × Before Switch					-0.791*** (0.149)	-0.939*** (0.145)
Mig. Share × After Switch					0.0623 (0.0441)	-0.00660 (0.0424)
R-squared	0.211	0.212	0.166	0.169	0.166	0.169
Obs.	10,019,309	10,019,435	4,555,154	4,555,151	4,555,154	4,555,151
Mean	.004	.003	-.056	-.043	-.056	-.043

	Secondary					
	Non-Switchers		Switchers			
	Math grades (1)	Language grades (2)	Math grades (3)	Language grades (4)	Math grades (5)	Language grades (6)
Mig. Share	-0.607*** (0.110)	-0.799*** (0.114)	-0.646*** (0.0832)	-1.301*** (0.0921)		
Mig. Share × Before Switch					-0.458** (0.221)	-0.355 (0.220)
Mig. Share × After Switch					-0.655*** (0.0847)	-1.344*** (0.0942)
R-squared	0.300	0.308	0.218	0.226	0.218	0.226
Obs.	8,744,722	8,787,929	3,346,142	3,367,426	3,346,142	3,367,426
Mean	-.04	-.017	-.053	-.038	-.053	-.038

Standard errors clustered at school-grade level in parentheses. *** p<0.01 ** p<0.05 * p<0.1. Control variables: sex, age, baseline math grade, school by year, grade, year, and district fixed effects. Columns (1) and (2) run the main specification for students who never switched schools between 2014 and 2019. Columns (3) to (6) for students who switched at least ones in the same period. On columns (5) and (6) the specification includes the concurrent share of migrants per grade (not the baseline as our main specification does) and the interaction between the migrant share and a dummy that indicates the years before and after the first time the student switched. The sample includes incumbent students from 2015 to 2019 in primary and schools.

Table A.10. Effect of Migrant Exposure By Migrant Baseline Performance

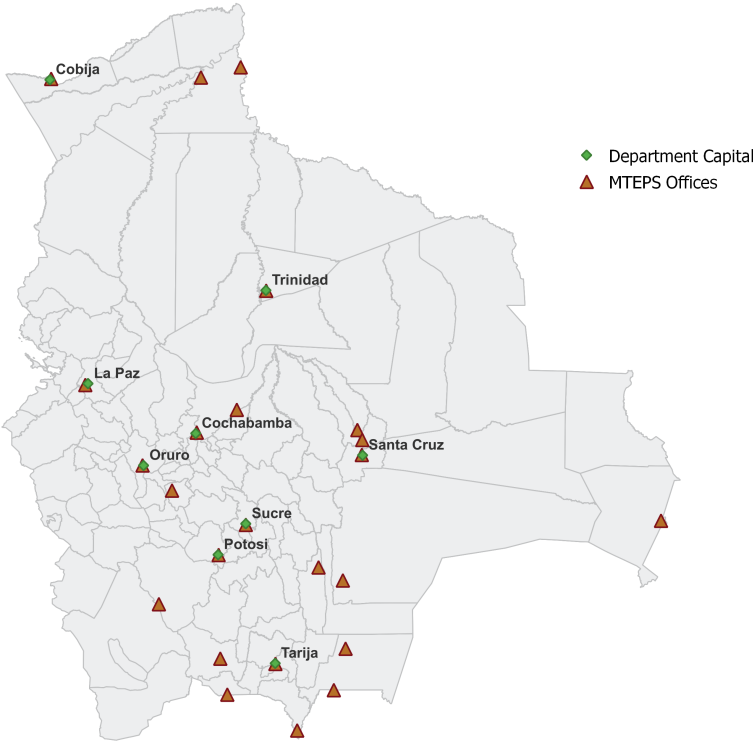
Primary					
	Switching (1)	Retention (2)	Dropout (3)	Math grades (4)	Language grades (5)
Mig. Share BL Achievement > Median	0.164*** (0.0183)	-0.000238 (0.00350)	0.0144** (0.00625)	-0.129*** (0.0482)	-0.143*** (0.0471)
Mig. Share BL Achievement < Median	0.145*** (0.0170)	0.00757** (0.00354)	0.0349*** (0.0126)	-0.141*** (0.0492)	-0.188*** (0.0472)
R-squared	0.087	0.031	0.104	0.180	0.182
Obs.	14,543,842	14,543,841	11,568,653	14,423,528	14,423,649
Mean	.115	.005	.013	-.014	-.011
Secondary					
	Switching (1)	Retention (2)	Dropout (3)	Math grades (4)	Language grades (5)
Mig. Share BL Achievement > Median	0.0398 (0.0307)	0.000775 (0.00526)	0.0316* (0.0165)	-0.943*** (0.138)	-0.862*** (0.142)
Mig. Share BL Achievement < Median	0.160*** (0.0241)	0.0127*** (0.00490)	0.00239 (0.0194)	-0.443*** (0.109)	-0.610*** (0.110)
R-squared	0.095	0.019	0.102	0.266	0.274
Obs.	12,513,766	12,513,766	9,970,220	12,034,217	12,098,403
Mean	.08	.005	.033	-.044	-.024

Standard errors clustered at school-grade level in parentheses. *** p<0.01 ** p<0.05 * p<0.1. Control variables: school by year fixed, grade, year, and district fixed effects. The specification includes the share of migrant students divided between the share of migrants with baseline math GPA above and below the median of their base line year. The sample includes incumbent students from 2015 to 2019 in primary and secondary schools.

Appendix B

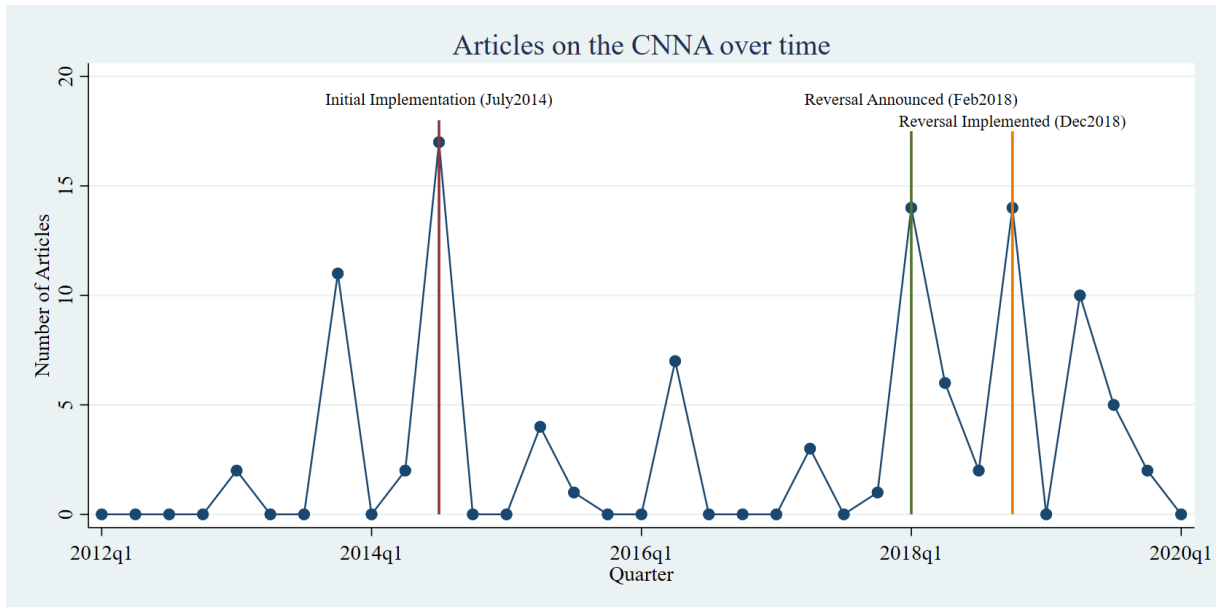
The Effects of Expanding Worker Rights to Children

B.1 Appendix Figures and Tables



The addresses of permanent MTEPS offices can be found here: https://www.mintrabajo.gob.bo/?page_id=2626.

Figure B.1. Ministry of Labor Offices




This figure tracks the number of articles concerning the 2014 law scraped from 43 national and regional Bolivian newspapers between 2012 and 2020. Articles that both mentioned the 2014 law and child labor were included.

Figure B.2. Articles on the 2014 Law over Time

Table B.1. Balance Table: Difference in Discontinuity - Household Survey

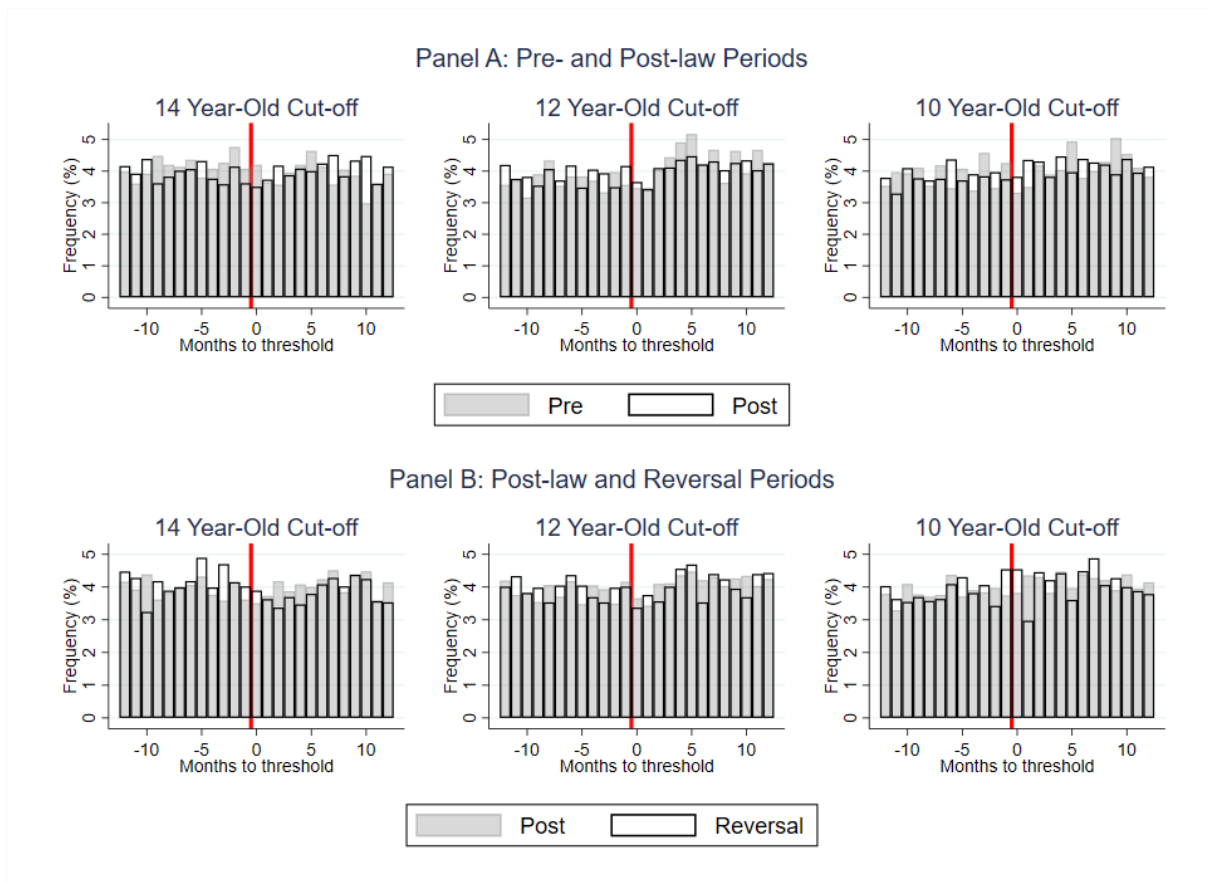
Panel A: 14-Year-Old Cutoff						
	Schooling (HH head)	Male (HH head)	Age (HH head)	Indigenous (HH head)	Male (child)	HH size
	(1)	(2)	(3)	(4)	(5)	(6)
Post Law $\times 1\{\text{Age} < 14\}$	0.197 (0.308)	-0.020 (0.020)	-0.412 (0.530)	0.027 (0.022)	-0.035 (0.028)	-0.072 (0.094)
Post Reversal $\times 1\{\text{Age} < 14\}$	0.310 (0.345)	-0.011 (0.025)	0.370 (0.572)	0.032 (0.026)	-0.009 (0.034)	-0.111 (0.088)
Obs.	11498	11498	11498	11498	11498	11498
Mean Control	8.509	0.798	45.16	0.347	0.499	5.562
Mean Treated	8.595	0.760	45.49	0.366	0.484	5.532
Joint test P-value = .632						
Panel B: 12-Year-Old Cutoff						
	Schooling (HH head)	Male (HH head)	Age (HH head)	Indigenous (HH head)	Male (child)	HH size
	(1)	(2)	(3)	(4)	(5)	(6)
Post Law $\times 1\{\text{Age} \geq 12\}$	-0.184 (0.322)	-0.028 (0.019)	0.064 (0.571)	-0.020 (0.027)	-0.050* (0.030)	0.093 (0.109)
Post Reversal $\times 1\{\text{Age} \geq 12\}$	-0.295 (0.357)	-0.011 (0.021)	-0.133 (0.570)	0.013 (0.029)	-0.043 (0.037)	-0.002 (0.111)
Obs.	11194	11194	11194	11194	11194	11194
Mean Control	8.653	0.790	44.26	0.356	0.522	5.619
Mean Treated	8.574	0.776	43.75	0.354	0.486	5.657
Joint test P-value = .514						
Panel C: 10-Year-Old Cutoff						
	Schooling (HH head)	Male (HH head)	Age (HH head)	Indigenous (HH head)	Male (child)	HH size
	(1)	(2)	(3)	(4)	(5)	(6)
Post Law $\times 1\{\text{Age} \geq 10\}$	-0.115 (0.344)	-0.036* (0.018)	0.577 (0.571)	-0.003 (0.033)	0.036 (0.027)	0.056 (0.121)
Post Reversal $\times 1\{\text{Age} \geq 10\}$	0.129 (0.383)	-0.043** (0.022)	-0.159 (0.628)	0.018 (0.031)	0.022 (0.027)	-0.008 (0.124)
Obs.	11313	11313	11313	11313	11313	11313
Mean Control	8.729	0.813	43.07	0.357	0.504	5.609
Mean Treated	8.848	0.777	42.59	0.369	0.525	5.669
Joint test P-value = .595						

Notes: Age in months by year clustered standard errors in parentheses. Significance levels denoted by: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The running variable is the difference between age in months and the age cut-off at the survey date. The specification includes linear splines of the running variable, an indicator that is one from 2014 to 2017, an indicator equal to one on 2018 and after, and an indicator that is one for the children in the corresponding age group. The bandwidth for all specifications is 12 months. We use a triangular kernel. The sample includes 2012-2019.




This figure plots the average raw work probability by age (in months) as well as a smoothed line for children between the ages of 7 and 17 prior to 2014. Data source: Encuesta de Hogares. Survey years: 2012-13.

Figure B.3. Work Probabilities by Age (Pre-law)



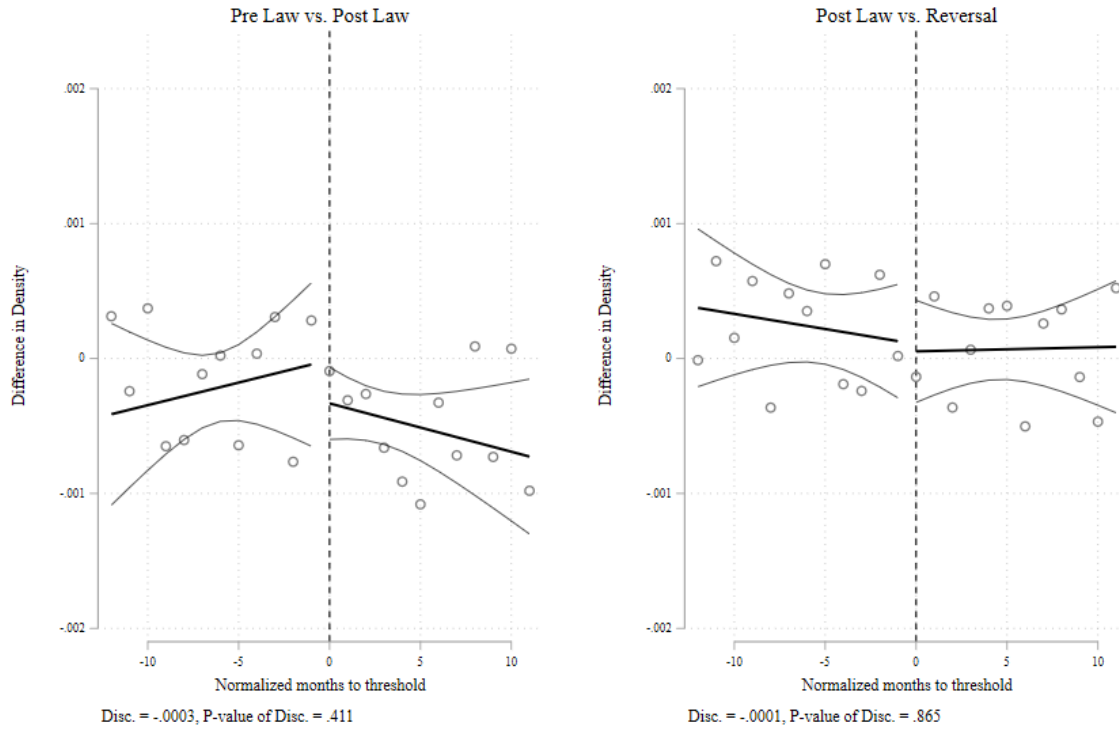
The running variable in both panels is the difference between age in months and the age cutoff at the survey date. In Panel A the pre sample includes 2012-2013 and the post sample includes 2014-2017. In Panel B the post sample includes 2014-2017 and the reversal sample includes 2018-2019. Both panels use data from multiple rounds of household surveys.

Figure B.4. Manipulation Test: Histograms




We use linear splines on each side of the potential density discontinuity to graphically approximate the density of the running variable, age in months, around a 12-month bandwidth. We also use a 12-month bandwidth and a linear polynomial specification interacted with a dummy equal to 1 at the right side of the cutoff to test for manipulation in the running variable. We report the discontinuity and its p-value below the graph. We use robust standard errors for both the graph and the estimate of the potential discontinuity.

Figure B.5. Differences in densities: 14 year-old cutoff




We use linear splines on each side of the potential density discontinuity to graphically approximate the density of the running variable, age in months, around a 12-month bandwidth. We also use a 12-month bandwidth and a linear polynomial specification interacted with a dummy equal to 1 at the right side of the cutoff to test for manipulation in the running variable. We report the discontinuity and its p-value below the graph. We use robust standard errors for both the graph and the estimate of the potential discontinuity.

Figure B.6. Differences in densities: 12 year-old cutoff




We use linear splines on each side of the potential density discontinuity to graphically approximate the density of the running variable, age in months, around a 12-month bandwidth. We also use a 12-month bandwidth and a linear polynomial specification interacted with a dummy equal to 1 at the right side of the cutoff to test for manipulation in the running variable. We report the discontinuity and its p-value below the graph. We use robust standard errors for both the graph and the estimate of the potential discontinuity.

Figure B.7. Differences in densities: 10 year-old cutoff




Household-level clustered standard errors in parentheses. Control variables: household head characteristics (schooling, gender, age, and indigenous indicator), gender, number of children in the following age categories: 0-6, 7-9, 10-13, and 14-17, number of adult men and women, an urban dummy, and departamento by year fixed effects. The sample includes 2012-2019.

Figure B.8. Difference in Discontinuity Event Study-style Estimates: Work Probability (12- and 10-Year-Old Cutoffs)



The running variable is the difference between age in months and the age cutoff a week before the survey date. The pre-law sample includes 2012-2013 and the post sample includes 2014-2017. We use a triangular kernel.

Figure B.9. Work Probabilities at the 12-Year-Old Cutoff (Before, During, and After the Law)



The running variable is the difference between age in months and the age cutoff a week before the survey date. The pre-law sample includes 2012-2013 and the post sample includes 2014-2017. We use a triangular kernel.

Figure B.10. Work Probabilities at the 10-Year-Old Cutoff (Before, During, and After the Law)

Table B.2. Effects of the Law and Reversal on Work Probability for 12 and 14 year-olds (Difference-in-Difference)

	Ages 12 vs. 13 (1)	Ages 14 vs. 15 (2)
Post Law \times 1{Treated}	0.003 (0.014)	-0.005 (0.016)
Post Reversal \times 1{Treated}	-0.001 (0.018)	-0.005 (0.019)
Obs.	12175	12165
Mean	0.160	0.203

Notes: Age in months by year clustered standard errors in parentheses. Significance levels denoted by: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. For Column 1, Treated=1 for 12 year-olds, and =0 for 13 year-olds. For Column 2, Treated=1 for 14 year-olds, and =0 for 15 year-olds. The control variables are: in grade for CCT (only for Column 2), an indicator for urban areas, household head characteristics (schooling, gender, age, and indigenous indicator), gender, number of children in the household in following age categories: 0-6, 7-9, 10-13, and 14-17, number of adult men and women, and departamento by year fixed effects. The specification includes an indicator for the corresponding age group, an indicator equal to one after the law was established, and one equal to one after the law was reversed, and an interaction between the age group indicator and the two indicators post law and reversal. The sample includes 2012-2019.

Table B.3. Effects of the Law on the Work Probabilities, Hours, and Occupation

Panel A: 12-Year-Old Cutoff						
	Any Work (1)	Hours Worked (2)	Work for Self (3)	Work for Others (4)	Prohibited Work (5)	Allowed Work (6)
Post Law $\times 1\{\text{Age} \geq 12\}$	-0.014 (0.015)	-0.339 (0.339)	-0.001 (0.003)	-0.012 (0.016)	-0.003 (0.004)	-0.010 (0.014)
Post Reversal $\times 1\{\text{Age} \geq 12\}$	0.015 (0.019)	0.231 (0.420)	-0.005** (0.003)	0.020 (0.019)	0.003 (0.008)	0.012 (0.016)
Obs.	11719	11719	11719	11719	11719	11719
Mean	0.142	2.846	0.00209	0.140	0.00349	0.138

Panel B: 10-Year-Old Cutoff						
	Any Work (1)	Hours Worked (2)	Work for Self (3)	Work for Others (4)	Prohibited Work (5)	Allowed Work (6)
Post Law $\times 1\{\text{Age} \geq 10\}$	-0.017 (0.014)	-0.199 (0.300)	0.002 (0.002)	-0.018 (0.014)	-0.003 (0.002)	-0.014 (0.014)
Post Reversal $\times 1\{\text{Age} \geq 10\}$	-0.013 (0.015)	-0.316 (0.277)	0.000 (0.002)	-0.014 (0.015)	0.001 (0.008)	-0.014 (0.014)
Obs.	11801	11801	11801	11801	11801	11801
Mean	0.105	1.788	0.000748	0.104	0.00150	0.103

Notes: Age in months by year clustered standard errors in parentheses. Significance levels denoted by: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Control variables: household head characteristics (schooling, gender, age, and indigenous indicator), gender, number of children in the following age categories: 0-6, 7-9, 10-13, and 14-17, number of adult men and women, an urban dummy, and departamento by year fixed effects. We also include linear splines of the running variable, defined as the difference between the cutoff age and age a week before the survey date in months. We use a bandwidth of 12 months and a triangular kernel. Survey years: 2012-2019.

Table B.4. Effect of the Law on Time Use

	Attends School (1)	Minutes Spent on Chores (2)
Post law \times Treated	0.013 (0.012)	-13.843 (18.363)
Post reversal \times Treated	-0.010 (0.013)	
Obs.	11498	8372
Mean	0.955	407.0

Notes: Age in months by year clustered standard errors in parentheses. Significance levels denoted by: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Control variables: CCT eligibility indicator (Column 1 only), household head characteristics (schooling, gender, age, and indigenous indicator), gender, number of children in the following age categories: 0-6, 7-9, 10-13, and 14-17, number of adult men and women, an indicator for urban, and departamento by year fixed effects. For Column 1, we include linear splines of the running variable, defined as the difference between the cutoff age and age at the survey in months. For Column 2, we do a stacked difference in discontinuity by multiplying the running variable by -1 for the 13 and 14 year-olds age group for interpretability. The running variable is the stacked difference between age in months and the age cutoff at the survey date, and the specification includes linear splines of the running variable. We use a bandwidth of 12 months and a triangular kernel for all specifications. Survey years for Column 1: 2012-2019. Survey years for Column 2: 2008 and 2016. We also report the mean of the dependent variable in the pre-law period.

Table B.5. Difference in Discontinuity: Household Outcomes for the 14-year-old Cut-off

	Any Adult in HH Works (1)	Total Hours Worked by Adults (2)	Any Older Sibling Works (3)	Total Hours Worked by Older Siblings (4)
Post law $\times 1\{\text{Age} < 14\}$	-0.001 (0.009)	-3.842* (2.317)	-0.012 (0.037)	-2.225 (1.403)
Post reversal $\times 1\{\text{Age} < 14\}$	-0.007 (0.009)	2.444 (2.524)	-0.010 (0.042)	-1.576 (1.420)
Obs.	10788	10788	3964	3964
Mean	0.969	94.26	0.307	7.993

Notes: Age in months by year clustered standard errors in parentheses. Significance levels denoted by: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The control variables are: an indicator that is one if child in HH is in grade for CCT, an indicator for urban, household head characteristics (schooling, gender, age, and indigenous indicator), number of children in the household in following age categories: 0-6, 7-9, 10-13, and 14-17, number of adult men and women, and departamento by year fixed effects. The income per capita variable in Column 3 is winsorized at the 99th percentile. The running variable is the difference between age in months of the child in the household and the age cut-off a week before the survey date. Hence, we only include households that have only a single child in the corresponding age range. The specification includes linear splines of the running variable, an indicator that is one between 2014 and 2018, an indicator equal to one in 2018 and after, and interaction between the running variable and the indicator for 2014 and after, and an indicator that is one for the children in the corresponding age group. The bandwidth is 12 months. We use a triangular kernel. The sample includes 2012-2019.

Table B.6. Heterogeneous Effects of the Law by Driving Time from MTEPS Offices
(Difference-in-Discontinuity)

Panel A: 12-Year-Old Cutoff		
	Dependent Variable: Works	
	All (1)	No MTEPS Offices (2)
Post \times 1{Age \geq 12} for Far	0.038 (0.037)	0.019 (0.045)
Post \times 1{Age \geq 12} for Near	-0.021 (0.016)	-0.054 (0.037)
Obs.	7313	2938
Mean	0.142	0.257
P-value of difference	0.124	0.128
P-value of difference (urban controls)	0.342	0.180

Panel B: 10-Year-Old Cutoff		
	Dependent Variable: Works	
	All (1)	No MTEPS Offices (2)
Post \times 1{Age \geq 10} for Far	0.046 (0.031)	0.012 (0.038)
Post \times 1{Age \geq 10} for Near	-0.024 (0.016)	-0.052 (0.037)
Obs.	7148	2889
Mean	0.105	0.217
P-value of difference	0.0344	0.146
P-value of difference (urban controls)	0.312	0.627

Notes: Age in months by year clustered standard errors in parentheses. Significance levels denoted by: *** p<0.01, ** p<0.05, * p<0.1. Municipalities that are classified as Far are above the median distance from a MTEPS office. Control variables: CCT eligibility indicator, urban, HH head characteristics (schooling, gender, age, indigenous indicator), gender, no. of children aged 0-6, 7-9, 10-13, and 14-17, no. of adult men and women, and departamento by year FE. We also include linear splines of the running variable (difference between the cutoff age and age a week before the survey date in months). The specification for the p-value with urban controls additionally includes: post \times urban, treatment \times urban, post \times distance \times urban, and treatment \times distance \times urban. We use a bandwidth of 12 months and a triangular kernel. Survey years: 2012-2016. We also report the mean of the dependent variable for the control group.

Table B.7. Heterogeneous Effects by Distance from MTEPS Offices, Allowing for Heterogeneity by Urban and Baseline Child Labor Rates

Panel A: Allowing for Heterogeneity by Urban		
	Dependent Variable: Works	
	All (1)	No MTEPS Offices (2)
Post \times 1{Age < 14} for Far	0.025 (0.069)	-0.016 (0.070)
Post \times 1{Age < 14} for Near	-0.036* (0.021)	-0.094** (0.048)
Obs.	7650	2984
Mean	0.180	0.317
P-value of difference	0.448	0.339

Panel B: Allowing for Heterogeneity by Baseline Child Labor Rates		
	Dependent Variable: Works	
	All (1)	No MTEPS Offices (2)
Post \times 1{Age < 14} for Far	-0.008 (0.066)	-0.019 (0.064)
Post \times 1{Age < 14} for Near	-0.105*** (0.029)	-0.071 (0.052)
Obs.	6874	2210
Mean	0.169	0.308
P-value of difference	0.199	0.565

Notes: Age in months by year clustered standard errors in parentheses. Significance levels denoted by: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Municipalities that are classified as Far are above the median distance from a MTEPS office, where distance is calculated as the driving time from the municipality centroid to the nearest MTEPS office. Control variables: CCT eligibility indicator, urban, HH head characteristics (schooling, gender, age, indigenous indicator), gender, no. of children aged 0-6, 7-9, 10-13, and 14-17, no. of adult men and women, and departamento by year FE. We also include linear splines of the running variable (difference between the cutoff age and age a week before the survey date in months). The specification for Panel A additionally includes: post \times urban, treatment \times urban, post \times distance \times urban, and treatment \times distance \times urban, where urban is normalized to the sample mean. The specification for Panel B additionally includes: post \times baseline CL rates, treatment \times baseline CL rates, post \times distance \times baseline CL rates, and treatment \times distance \times baseline CL rates, where baseline CL rates are defined at the municipality level, are calculated using data from only 2012 (pre-law), and are normalized to the municipality mean. We use a bandwidth of 12 months and a triangular kernel. Survey years: 2012-2016. We also report the mean of the dependent variable for the pre-law period.

Table B.8. Functional Form Robustness Checks: Difference in Discontinuity for Work Probability (14-Year-Old Cutoff)

	Bandwidth (months)			No Controls	Quadratic	Polynomials Pre-Post		
	6	12	24			Linear	Quadratic	Donut
	(Baseline)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post law $\times 1\{\text{Age} < 14\}$	-0.034 (0.021)	-0.039** (0.017)	-0.027** (0.013)	-0.031 (0.024)	-0.040** (0.016)	-0.032 (0.031)	-0.039** (0.016)	-0.030* (0.016)
Post reversal $\times 1\{\text{Age} < 14\}$	0.005 (0.024)	-0.000 (0.019)	-0.000 (0.015)	0.015 (0.032)	-0.001 (0.019)	0.017 (0.034)	-0.000 (0.019)	0.012 (0.020)
Obs.	5983	11991	24340	11991	11991	11991	11991	11057
Mean	0.188	0.180	0.180	0.180	0.180	0.180	0.180	0.183

Notes: Household level clustered standard errors in parentheses. Controls: in grade for CCT, an indicator for urban, household head characteristics (schooling, gender, age, and indigenous indicator), gender, number of children in the household in following age categories: 0-6, 7-9, 10-13, and 14-17, number of adult men and women, and departamento by year fixed effects. The running variable is the difference between age in months and the age cut-off a week before the survey date. We include linear splines of the running variable, an indicator for 2014 and after, and an indicator that is one for the children in the corresponding age group. Column 5 also includes quadratic splines of the running variable. Column 6 includes linear splines that vary across both sides of the cut-off and before and after the law. Column 7 has linear and quadratic splines that vary across both sides of the cut-off and before and after the law. Column 8 omits children within 1 month of the age threshold. We use a triangular kernel. The sample includes 2012-2019.

Table B.9. Other Robustness Checks: Difference in Discontinuity for Work Probability (14-Year-Old Cutoff)

	Baseline Estimation	Excl. Indig.	Excl. CCT control	Excl. 2014	Cluster Age & Region
	(1)	(2)	(3)	(4)	(5)
Post law $\times 1\{\text{Age} < 14\}$	-0.039** (0.017)	-0.034* (0.017)	-0.037** (0.017)	-0.034** (0.017)	-0.030* (0.016)
Post reversal $\times 1\{\text{Age} < 14\}$	-0.000 (0.019)		0.002 (0.019)	0.000 (0.019)	
Obs.	11991	6481	11991	10418	7650
Mean	0.180	0.111	0.180	0.180	0.180

Notes: Household level clustered standard errors in parentheses. Controls: in grade for CCT, an indicator for urban, household head characteristics (schooling, gender, age, and indigenous indicator), gender, number of children in the household in following age categories: 0-6, 7-9, 10-13, and 14-17, number of adult men and women, and departamento by year fixed effects. The running variable is the difference between age in months and the age cut-off a week before the survey date. We include linear splines of the running variable, an indicator for 2014 and after, and an indicator that is one for the children in the corresponding age group. Column 2 excludes municipalities with above median shares of indigenous residents. Because municipality codes are anonymized in the household survey data starting in 2017, we cannot link the data to other sources using municipality codes for the periods after the law was reversed. Column 3 excludes the control that indicates whether the child is eligible for the CCT. Column 4 excludes the year 2014 from the sample. Column 5 clusters by age in months and municipality. We use a triangular kernel. The sample includes 2012-2016 for columns 2 and 5, 2012-2013 and 2015-2019 for column 4, and 2012-2019 for all other columns.

Table B.10. Difference in Difference Specifications

	Dep. Var.: Any Work		
	Control: 14-year-olds (1)	Control: 9- and 14-year-olds (2)	Control: 7-9 and 14-16-year-olds (3)
Post Law \times $1\{\text{Age} < 14\}$	-0.038** (0.015)		
Post Reversal \times $1\{\text{Age} < 14\}$	-0.011 (0.018)		
Post Law \times $1\{10 \leq \text{Age} < 12\}$		-0.012 (0.010)	-0.002 (0.009)
Post Law \times $1\{12 \leq \text{Age} < 14\}$		-0.018* (0.010)	-0.009 (0.009)
Post Reversal \times $1\{10 \leq \text{Age} < 12\}$		-0.008 (0.011)	-0.001 (0.009)
Post Reversal \times $1\{12 \leq \text{Age} < 14\}$		-0.000 (0.012)	0.008 (0.010)
Obs.	11991	35511	53490
Mean	0.180	0.144	0.137

Notes: Age in months by year clustered standard errors in parentheses. Significance levels denoted by: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The control variables are: in grade for CCT (only for 14-year-old cut-off), an indicator for urban areas, household head characteristics (schooling, gender, age, and indigenous indicator), gender, number of children in the household in following age categories: 0-6, 7-9, 10-13, and 14-17, number of adult men and women, and departamento by year fixed effects. The specification includes an indicator for the corresponding age group, an indicator equal to one after the law was established and before it was reversed, an indicator equal to one after the law was reversed, and interactions between the time and the age group indicators. The sample includes 2012-2019.

Table B.11. Balance for 30% of Child Labor Survey Data

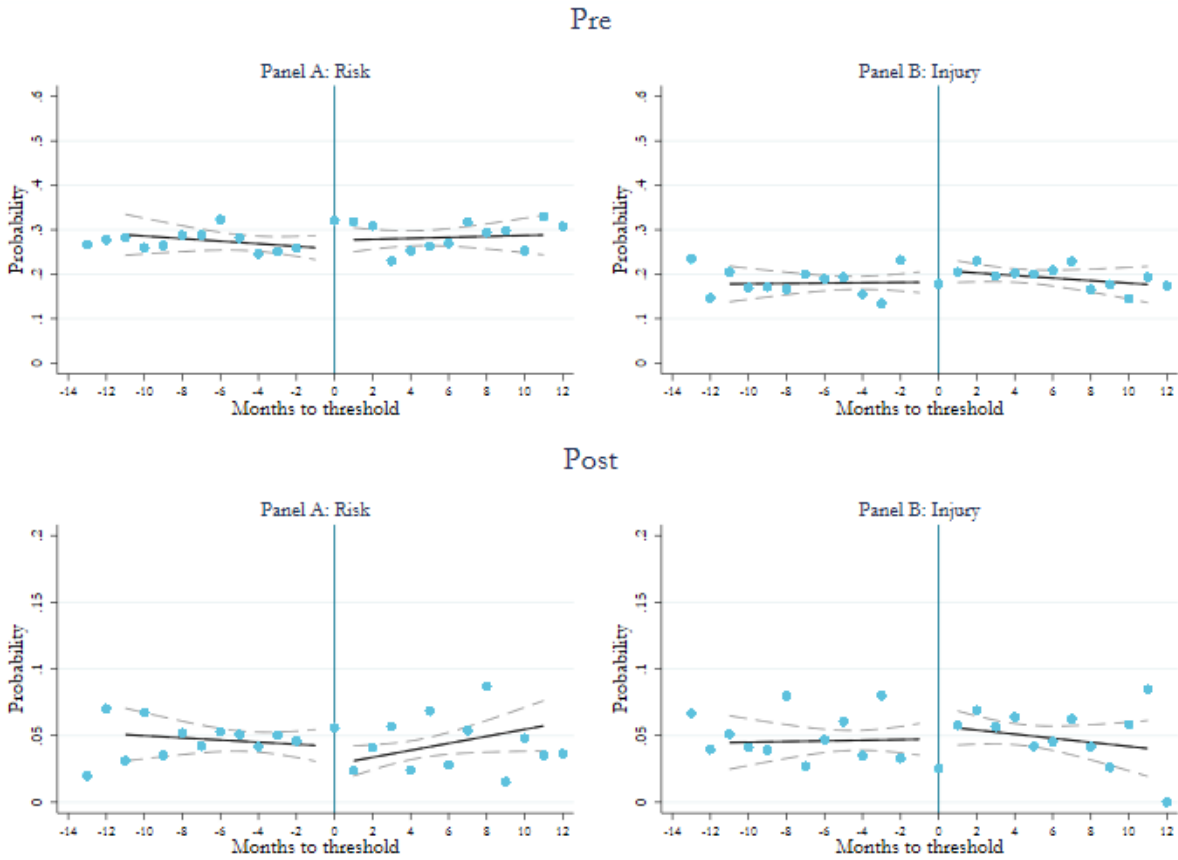
	Male	HH Size	Age HH Head	Education HH Head	Male HH Head	Indigenous HH Head	Urban
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Post	-0.023 (0.020)	0.129 (0.088)	0.696 (0.481)	-0.151 (0.211)	0.013 (0.017)	0.036* (0.020)	0.008 (0.018)
Obs.	2580	2580	2580	2580	2580	2580	2580
Mean	0.510	5.857	42.62	7.888	0.786	0.348	0.742
Joint test P-value = .262							

Notes: Household level clustered standard errors in parentheses. Significance levels denoted by: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The specification includes an indicator that is one in 2016. The running variable is multiplied by -1 for the 13 and 14 year-olds age group for interpretability. The bandwidth for all specifications is 12 months. The sample is 30% of the 2008 and 2016 observations that were not used in the reweighting exercise.

Table B.12. Balance for Reweighted Child Labor Survey Data - Full sample


	Male	HH Size	Age HH Head	Education HH Head	Male HH Head	Indigenous HH Head	Urban
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Post \times Treated	-0.041* (0.022)	-0.039 (0.099)	-0.764 (0.546)	0.237 (0.267)	0.005 (0.019)	0.004 (0.025)	0.012 (0.024)
Obs.	8372	8372	8372	8372	8372	8372	8372
Mean	0.510	5.857	42.62	7.888	0.786	0.348	0.742
Joint test P-value = .604							

Notes: Age in months by year clustered standard errors in parentheses. Significance levels denoted by: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The running variable is the difference between age in months and the age cut-off at the survey date. The specification includes linear splines of the running variable, an indicator that is one in 2016, and an indicator that is one for the children in the corresponding age group. The bandwidth for all specifications is 12 months. We use a triangular kernel. The sample includes 2008 and 2016.



The running variable is the difference between age in months and the age cutoff a week before the survey date, defined separately for each age threshold. We use a triangular kernel and we reweight the observations as described in Section 2.4.

Figure B.11. Job Risks & Work Injuries (Before and During the Law): Stacked Data



The figures present means of the dependent variables by quartiles of per-capita household income using data on children aged 7 to 18 years old. The left hand side figure reports the probability of having a permit using data from the 2016 Child Labor Survey. The right hand side figure reports the probability of having a written contract with an employer on using data from the 2014-2017 household survey waves.

Figure B.12. Work permits and written contracts by per-capita household income

Table B.13. Effects of the Law on Job Risks, and Work Injuries

Panel A: 14-Year-Old Cutoff		
	Faces Risks at Work (1)	Has Been Injured at Work (2)
Post $\times 1\{\text{Age} < 14\}$	-0.007 (0.019)	-0.001 (0.026)
Obs.	2808	2827
Mean	0.349	0.219

Panel B: 12-Year-old Cutoff		
	Faces Risks at Work (1)	Has Been Injured at Work (2)
Post $\times 1\{\text{Age} \geq 12\}$	-0.021 (0.024)	-0.016 (0.018)
Obs.	2733	2767
Mean	0.278	0.183

Panel C: 10-Year-old Cutoff		
	Faces Risks at Work (1)	Has Been Injured at Work (2)
Post $\times 1\{\text{Age} \geq 10\}$	-0.018 (0.020)	-0.025 (0.018)
Obs.	2831	2817
Mean	0.214	0.166

Notes: Age in months by year clustered standard errors in parentheses. Significance levels denoted by: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Control variables: gender, working indicator (Panel B only), urban indicator, age group fixed effects, household head characteristics (schooling, gender, age, and indigenous indicator), gender, number of children in the following age categories: 0-6, 7-9, 10-13, and 14-17, number of adult men and women, and departamento by year fixed effects. For the risk index regression, the running variable is the difference between age in months and the age cutoff at the survey date. For the injury index, the running variable is the difference between age in months and the age cutoff a year before the survey date. The specification includes linear splines of the running variable. The bandwidth for all specifications is 12 months. We use a triangular kernel. Survey years: 2008, 2016. We use a reweighting method described in Section 2.4.

Table B.14. Robustness Checks: Difference in Discontinuity for Risk Outcomes

Panel A: Different Bandwidth Specifications						
	Risk Index			Injury Index		
	<i>Bandwidth (months)</i>					
	Baseline			Baseline		
	6	12	24	6	12	24
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Treated	-0.012 (0.021)	-0.008 (0.017)	-0.009 (0.015)	-0.006 (0.020)	-0.015 (0.014)	-0.010 (0.012)
Obs.	3981	8372	8872	4074	8411	8885
R-squared	0.186	0.179	0.182	0.110	0.107	0.103
Mean	0.277	0.281	0.281	0.194	0.188	0.188

Panel B: Without Controls, Quadratic Splines, and Donut Specification						
	Risk Index			Injury Index		
	No Controls	Quadratic	Donut	No Controls	Quadratic	Donut
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Treated	-0.013 (0.016)	-0.007 (0.016)	-0.003 (0.018)	-0.007 (0.013)	-0.015 (0.014)	-0.036*** (0.013)
Obs.	8372	8372	7325	8411	8411	7351
R-squared	0.109	0.180	0.183	0.0509	0.107	0.109
Mean	0.281	0.281	0.279	0.188	0.188	0.186

Notes: Household level clustered standard errors in parentheses. Significance levels denoted by: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The control variables are: gender, urban indicator, age group fixed effects, household head characteristics (schooling, gender, age, and indigenous indicator), gender, number of children in the following age categories: 0-6, 7-9, 10-13, and 14-17, number of adult men and women, and departamento by year fixed effects. The running variables are the difference between age in months and the age cut-off at the survey date for the risk and hazardous work indices, and the difference between age in months and the age cut-off a year before the survey date for the injury index. The specification includes linear splines of the running variable, an indicator that is one in 2016, and an indicator that is one for the children in the corresponding age group. We use a triangular kernel. The sample includes 2008 and 2016.

B.2 List of Prohibited Tasks under the 1999 and 2014 Laws

Under the 1999 and 2014 laws, children were prohibited from engaging in the following tasks (Authors' translation of original Spanish document):

- Harvesting sugar cane
- Harvesting chestnuts (Brazil nuts)
- Mining
- Fishing in rivers and lakes (other than family or community work activities)
- Brickwork
- Selling alcoholic drinks
- Collecting waste that can affect children's health
- Cleaning hospitals
- Security services
- Live-in domestic work
- Plasterwork
- Agriculture (other than family or community work activities)**This restriction was added in 2014.*
- Large livestock tending (other than family or community work activities)
- Work after hours
- Modeling that has an erotic connotation
- Attending to urinals after hours
- Stone cutting / masonry
- Sound amplification
- Handling heavy machinery
- Construction work (other than family or community work activities)
- Guarding cars after hours

B.3 Variable Definitions

- Any work: Indicator equal to one if the child reports working (or temporarily taking time off from their usual job) in the week prior to the survey. Does not include any unpaid household chores, such as cooking, cleaning, or caring for family members.
- Hours worked: Reported hours worked during the week before the survey; takes the value of zero if children report not working. The survey contains data about the average number of days worked in a week and the average number of hours worked per day for each household member age 7 or older. We compute weekly work hours by multiplying the number of days worked per week by the number of daily hours.
- Prohibited work: Indicator equal to one if the child reports engaging in any work as listed in Appendix B.2.
- Allowed work: Indicator equal to one if the child reports engaging in any other work that is not prohibited as detailed in Appendix B.2.
- Works more than 30 hrs.: Indicator equal to one if the child reports working more than 30 hours in the week before the survey; takes the value of zero if children report not working.
- Work for self: Indicator equal to one if the child reports working as self-employed or as an unpaid business owner in the week before the survey; takes the value of zero if children report not working.
- Work for others: Indicator equal to one if the child reports working for an external employer or for a family employer in the week before the survey; takes the value of zero if children report not working.
- Faces risks at work: Indicator equal to one if the child reports facing any of the following at work in the week prior to the survey:

- Dirt or contaminated dust
- Fire, gas, flames
- Loud noise or vibrations
- Extreme heat or cold
- Dangerous instruments (knives, explosives, etc.)
- Underground work
- Work at height
- Work in water
- Darkness, isolation, or without ventilation
- Chemical products (e.g. pesticides, glue)
- Other risks (given as an option in the survey)

The indicator is zero if children report not working.

- Has been injured at work: Indicator equal to one if the child reports having experienced any of the following injuries at work in the year prior to the survey:
 - Superficial injuries or bites, blisters, etc.
 - Fractures or mutilations
 - Dislocation or distention
 - Burns, scalds, or freezing
 - Respiratory problems
 - Sight problems
 - Skin injuries
 - Stomach problems (diarrhea or chemical poisoning)

- Exhaustion due to task intensity
- Other injuries (given as an option in the survey)

The indicator is zero if children report not working.

- Attends school: Indicator equal to one if children report attending school regularly (or if they report being on vacation but are enrolled in school) at the date of the survey.

B.4 Measuring driving time to MTEPS offices

We describe the process for computing the driving time to the nearest MTEPS office below:

- We obtained addresses and coordinates for MTEPS offices from MTEPS’s website https://www.mintrabajo.gob.bo/?page_id=2626.
- We obtained the coordinates (latitude and longitude) corresponding to the locality where the municipality government is located, typically the locality with the largest population in each municipality. To obtain this information we scraped data from <https://www.municipio.com.bo/>, a website with detailed descriptions of all municipalities in Bolivia. (See, for example, <https://www.municipio.com.bo/municipio-las-carreras.html>)
- For each point (centroid), the travel time to MTEPs offices in the record is calculated (about 8400+ combinations). Then for each municipality, we keep the travel information to the office with the fastest travel by car. Importantly, the algorithm is set to request the API to optimize travel time; therefore, the selected routes are the least time-consuming, although shorter routes (in terms of distance) may be possible. We use two measures to define the closest office to each municipality. First, we estimate the shortest possible distance between each municipality and each MTEPS office (straight line or “as the crow flies” distance). Second, we check for the fastest possible trip by driving. In some cases,

where there was no existing network of routes connecting the points, we were not able to compute distance based on travel time. We avoid this problem by using geocoded centroids (Bing) when the issue arises. Specifically, we feed the algorithm a rough location, typically the name of the municipality (e.g., "Las Carreras, Chuquisaca, Bolivia"), from which we get a precise location that we later use to calculate travel routes.

- As a result, for each municipality, we are able to compute two measures of distance: travel time by road and "as the crow flies" distance.
- Based on each measure of distance, we split municipalities in two groups: Near (minimum distance below the cross-municipality median) and Far (minimum distance above the cross-municipality median).

B.5 Tables

Appendix C

Gender and Transitions into Adulthood

Table C.1. Midline Heterogeneity of Productive Time Use by Schooling and Gender

	Employed	Independent Work	Wage Work
HD	0.09 (0.05) [0.18]	0.03 (0.04) [0.62]	0.09 (0.05) [0.23]
GD	0.04 (0.05) [0.56]	0.15*** (0.04) [0.00]	-0.07 (0.05) [0.36]
Combined	0.04 (0.06) [0.63]	0.11 (0.06) [0.16]	-0.05 (0.06) [0.62]
HD × Female	-0.13 (0.07) [0.16]	0.01 (0.05) [0.88]	-0.11 (0.06) [0.23]
GD × Female	-0.04 (0.06) [0.62]	-0.03 (0.05) [0.67]	-0.04 (0.06) [0.62]
Combined × Female	-0.06 (0.09) [0.62]	0.01 (0.07) [0.88]	-0.04 (0.08) [0.73]
HD × Baseline Years of Schooling	-0.01 (0.02) [0.75]	-0.04* (0.02) [0.08]	0.02 (0.02) [0.41]
GD × Baseline Years of Schooling	-0.02 (0.02) [0.43]	-0.05* (0.02) [0.05]	0.03 (0.02) [0.40]
Combined × Baseline Years of Schooling	0.00 (0.03) [0.88]	-0.04 (0.03) [0.26]	0.03 (0.03) [0.56]
HD × Female × Baseline Years of Schooling	0.04 (0.03) [0.36]	0.05* (0.02) [0.08]	-0.00 (0.03) [0.88]
GD × Female × Baseline Years of Schooling	0.10*** (0.03) [0.01]	0.09*** (0.02) [0.00]	0.02 (0.03) [0.62]
Combined × Female × Baseline Years of Schooling	0.01 (0.04) [0.82]	0.04 (0.03) [0.36]	-0.00 (0.04) [0.88]
Female × Baseline Years of Schooling	-0.05* (0.02) [0.08]	-0.04* (0.02) [0.05]	-0.02 (0.02) [0.58]
Baseline Years of Schooling	0.01 (0.02) [0.56]	0.03* (0.01) [0.08]	-0.02 (0.01) [0.43]
Female	-0.16*** (0.05) [0.01]	-0.05 (0.03) [0.25]	-0.16*** (0.05) [0.01]
Control mean	0.48	0.12	0.37
Observations	1770	1770	1770
R ²	0.07	0.05	0.08
P-value	0.00	0.00	0.82

Notes: Table presents tests for heterogeneity of productive time use effects by Gender. Interacted coefficients in rows 4 to 6 give the differential effect of each arm for women with no education. Interacted coefficients in rows 7 to 9 give the differential effect of each arm for each additional year of education that men acquire. Interacted coefficients in rows 10 to 12 give the differential effect of each arm for each additional year of education that men acquire. 'Female' gives the difference between women and men, and 'Baseline Years of Schooling' gives the differential effect of each arm for each additional year of education. The uninteracted treatment terms give the impact of each arm for men with no education. The first column is an indicator of any type of employment. The second column is an indicator for entrepreneurs and self employed workers. The third and final column is an indicator for employment for others. Standard errors are clustered at the household level, *=10%, **=5%, and ***=1% significance.

Table C.2. Income by Gender and Educational Attainment

	All (SD)	Women (SD)	Men (SD)	P-value Of Difference
	<i>All</i>			
Monthly income (All)	4.452 (4.95)	3.876 (4.81)	5.325 (5.04)	0.000
Observations	1848	1113	735	
	<i>Lower Education</i>			
Monthly income (Working for an Employer)	4.546 (4.91)	4.072 (4.81)	5.286 (4.98)	0.000
Observations	1116	680	436	
	<i>Higher Education</i>			
Monthly income (Working for an Employer)	4.310 (5.02)	3.569 (4.79)	5.383 (5.15)	0.000
Observations	732	433	299	
P-value Diff. by Education Level	0.158	0.049	0.828	

Notes: Table presents baseline means and standard deviations for all beneficiaries, women and men, and p-value for a test of the hypothesis that compares men and women of monthly income. Regression-based comparisons and associated hypothesis tests based on a regression with sector indicators and household-level clustered standard errors. The table includes three panels. The first shows the mean income of employees by gender. The second shows the employees' income for beneficiaries under the median education level by gender. The third panel does the same for beneficiaries above the median education level. The last row shows regression-based comparisons and associated hypothesis tests based on a regression with sector indicators, and household-level clustered standard errors comparing the average income across education levels for all beneficiaries, for women and men. *=10%, **=5%, and ***=1% significance.

Table C.3. Endline Heterogeneity of Productive Time Use by Schooling and Gender

	Employed	Independent Work	Wage Work
HD	0.03 (0.05) [1.00]	0.01 (0.03) [1.00]	0.01 (0.05) [1.00]
GD	-0.03 (0.05) [1.00]	0.11** (0.04) [0.03]	-0.11 (0.05) [0.30]
Combined	0.00 (0.06) [1.00]	0.10 (0.05) [0.75]	-0.09 (0.06) [1.00]
HD × Female	0.00 (0.06) [1.00]	0.04 (0.04) [1.00]	-0.03 (0.06) [1.00]
GD × Female	0.05 (0.06) [1.00]	-0.03 (0.04) [1.00]	0.06 (0.06) [1.00]
Combined × Female	0.01 (0.08) [1.00]	0.01 (0.07) [1.00]	0.01 (0.08) [1.00]
HD × Baseline Years of Schooling	-0.00 (0.02) [1.00]	-0.01 (0.01) [1.00]	-0.00 (0.02) [1.00]
GD × Baseline Years of Schooling	-0.02 (0.02) [1.00]	-0.02 (0.01) [1.00]	-0.01 (0.02) [1.00]
Combined × Baseline Years of Schooling	0.01 (0.03) [1.00]	-0.01 (0.02) [1.00]	0.01 (0.03) [1.00]
HD × Female × Baseline Years of Schooling	-0.02 (0.03) [1.00]	-0.01 (0.02) [1.00]	0.01 (0.03) [1.00]
GD × Female × Baseline Years of Schooling	0.04 (0.03) [1.00]	0.03 (0.02) [1.00]	0.03 (0.03) [1.00]
Combined × Female × Baseline Years of Schooling	0.01 (0.04) [1.00]	0.04 (0.03) [1.00]	-0.01 (0.04) [1.00]
Female × Baseline Years of Schooling	-0.00 (0.02) [1.00]	0.00 (0.01) [1.00]	-0.01 (0.02) [1.00]
Baseline Years of Schooling	-0.00 (0.01) [1.00]	0.00 (0.01) [1.00]	-0.00 (0.02) [1.00]
Female	-0.29*** (0.04) [0.00]	-0.04 (0.03) [1.00]	-0.27*** (0.05) [0.00]
Control mean	0.50	0.10	0.42
Observations	1822	1822	1822
R ²	0.08	0.04	0.08
P-value	0.18	0.18	0.67

Notes: Table presents tests for heterogeneity of productive time use effects by Gender. Interacted coefficients in rows 4 to 6 give the differential effect of each arm for women with no education. Interacted coefficients in rows 7 to 9 give the differential effect of each arm for each additional year of education that men acquire. Interacted coefficients in rows 10 to 12 give the differential effect of each arm for each additional year of education that men acquire. 'Female' gives the difference between women and men, and 'Baseline Years of Schooling' gives the differential effect of each arm for each additional year of education. The uninteracted treatment terms give the impact of each arm for men with no education. The first column is an indicator of any type of employment. The second column is an indicator for entrepreneurs and self employed workers. The third and final column is an indicator for employment for others. Standard errors are clustered at the household level, *=10%, **=5%, and ***=1% significance.

Table C.4. Marriage and Fertility

	GiveDirectly			Control Mean	Obs.	R^2
	HD	GD	Combined			
Married	-0.01 (0.03) [0.73]	0.05 (0.02) [0.18]	0.06 (0.04) [0.24]	0.16	1822	0.06
Cohabiting	0.01 (0.03) [0.73]	0.03 (0.03) [0.34]	-0.01 (0.03) [0.79]	0.18	1822	0.03
Any Children	0.01 (0.03) [0.73]	0.03 (0.03) [0.49]	0.09 (0.04) [0.18]	0.54	1822	0.02
Desired Fertility	-0.11 (0.05) [0.18]	-0.02 (0.04) [0.73]	-0.10 (0.06) [0.24]	2.85	1822	0.01

Notes: Table analyzes endline marriage and fertility outcomes at the beneficiary level. The first row is a dummy for whether the individual was married at endline. Row 2 is a dummy for whether the individual was cohabiting at endline. Row 3 analyzes a dummy for whether the beneficiary has any children as of the time of the endline, and Row 4 uses the response to the question “what is the total number of children you would like to have in your lifetime, including those that you have already”. Standard errors are clustered at the household level, *=10%, **=5%, and ***=1% significance.

Table C.5. Education and Time Use

	GiveDirectly			Control Mean	Obs.	R^2
	HD	GD	Combined			
Highest Grade	0.28 (0.25) [1.00]	0.15 (0.23) [1.00]	0.00 (0.33) [1.00]	12.02	1822	0.04
Hours in School	-0.68 (0.45) [1.00]	-0.40 (0.44) [1.00]	-0.47 (0.60) [1.00]	3.05	1822	0.01
Hours Domestic Work	-0.23 (1.18) [1.00]	-0.56 (1.06) [1.00]	0.09 (1.59) [1.00]	24.49	1822	0.01
Res wage in Village	-0.09 (0.09) [1.00]	0.38 (0.34) [1.00]	0.16 (0.12) [1.00]	1.68	1810	0.01
Res wage in Town	-0.25 (0.17) [1.00]	0.41 (0.45) [1.00]	0.28 (0.23) [1.00]	2.85	1804	0.01

Notes: Table analyzes endline education and time use variables. Highest grade is an ordinal variable measuring completed schooling with the control mean representing one year of post-primary education. ‘Hours in School’ and ‘Hours Domestic Work’ give the number of hours over the seven days prior to the endline that the respondent reports spending in each activity. ‘Reservation wages’ give the survey response to the daily wage the respondent said they would need to be paid to take a job in their village and in the nearest town, respectively (USD). Standard errors are clustered at the household level, *=10%, **=5%, and ***=1% significance.

Table C.6. Creation of New Households and Moving Across Villages

	GiveDirectly			Control Mean	Obs.	R ²
	HD	GD	Combined			
Ever formed new Household	0.04 (0.03) [0.58]	0.11*** (0.03) [0.00]	0.03 (0.04) [0.79]	0.40	1822	0.02
Ever HH Head or Spouse	0.02 (0.03) [0.79]	0.04 (0.03) [0.58]	0.08 (0.04) [0.35]	0.63	1822	0.02
New Village	0.01 (0.03) [0.79]	0.04 (0.03) [0.58]	-0.01 (0.04) [0.93]	0.37	1822	0.02
Urban	0.03 (0.02) [0.58]	0.01 (0.02) [0.79]	0.00 (0.02) [0.94]	0.08	1822	0.05

Notes: Table analyzes the movement and creation of new households by beneficiaries across both survey waves. Row 1 examines a dummy variable for whether the beneficiary was living in a different household than the baseline household at midline or endline. Row 2 is a dummy equal to one if the beneficiary is the head of household or their spouse at midline or endline. ‘New Village’ in row 3 is a dummy for the village in midline or endline being a different one than the baseline village. ‘Urban’ in row 4 is a dummy variable indicating that the village in which the beneficiary resides in midline or endline is classified as semi-urban, peri-urban, or urban (rather than rural). Standard errors are clustered at the household level, *=10%, **=5%, and ***=1% significance.

Bibliography

Alberto Abadie. Semiparametric difference-in-differences estimators. *The Review of Economic Studies*, 72(1):1–19, 2005.

Ryan M Abman, Clark C Lundberg, John McLaren, and Michele Ruta. Child labor standards in regional trade agreements: Theory and evidence. Working Paper 30908, National Bureau of Economic Research, February 2023. URL <http://www.nber.org/papers/w30908>.

N. Ajzenman, T. Cavalcanti, and D. Da Mata. More than Words: Leaders’ Speech and Risky Behavior During a Pandemic. Cambridge Working Papers in Economics 2034, Faculty of Economics, University of Cambridge, April 2020. URL <https://ideas.repec.org/p/cam/camdae/2034.html>.

Claudia Allende. Competition under social interactions and the design of education policies. *Job Market Paper*, 2019.

Rita Almeida and Pedro Carneiro. Enforcement of labor regulation and firm size. *Journal of Comparative Economics*, 37(1):28–46, 2009.

Rita Almeida and Pedro Carneiro. Enforcement of labor regulation and informality. *American Economic Journal: Applied Economics*, 4(3):64–89, 2012.

Rita Almeida and Lucas Ronconi. Labor inspections in the developing world: Stylized facts from the enterprise survey. *Industrial Relations: A Journal of Economy and Society*, 55(3):468–489, 2016.

David H. Autor, William R. Kerr, and Adriana D. Kugler. Does employment protection reduce productivity? evidence from us states*. *The Economic Journal*, 117(521):F189–F217, 2007. doi: <https://doi.org/10.1111/j.1468-0297.2007.02055.x>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0297.2007.02055.x>.

Joao Pedro Azevedo, Marta Favara, Sarah E Haddock, Luis F López-Calva, Miriam Muller, and Elizaveta Perova. Teenage pregnancy and opportunities in latin america and the caribbean: on teenage fertility decisions, poverty and economic achievement. 2012.

- JE Baez, A Alam, and XV Del Carpio. Does cash for school influence young women's behavior in the longer term. *Evidence from Pakistan*, 2011.
- Drew H Bailey, Greg J Duncan, Flávio Cunha, Barbara R Foorman, and David S Yeager. Persistence and fade-out of educational-intervention effects: Mechanisms and potential solutions. *Psychological Science in the Public Interest*, 21(2):55–97, 2020.
- Sarah Baird, Craig McIntosh, and Berk Özler. Cash or condition? evidence from a cash transfer experiment. *The Quarterly journal of economics*, 126(4):1709–1753, 2011.
- María Balarin. The default privatization of peruvian education and the rise of low-fee private schools: Better or worse opportunities for the poor? 2015.
- Abhijit Banerjee, Amy Finkelstein, MIT Rema Hanna, Benjamin A Olken, Arianna Ornaghi, and Sudarno Sumarto. The challenges of universal health insurance in developing countries: Experimental evidence from indonesia's national health insurance. *American Economic Review*, 2021.
- Olivier Bargain and Delphine Boutin. Minimum age regulation and child labor: New evidence from brazil. *World Bank Economic Review*, 35(1):234–260, 2021. doi: 10.1093/wber/lhz047. URL <https://elibrary.worldbank.org/doi/abs/10.1093/wber/lhz047>.
- Kaushik Basu and Pham Hoang Van. The economics of child labor. *American Economic Review*, pages 412–427, 1998.
- Gary S Becker. A theory of the allocation of time. *The economic journal*, 75(299):493–517, 1965.
- Julian R Betts and Robert W Fairlie. Does immigration induce 'native flight' from public schools into private schools? *Journal of Public Economics*, 87(5-6):987–1012, 2003.
- Diether W Beuermann and C Kirabo Jackson. The short and long-run effects of attending the schools that parents prefer. *The Journal of Human Resources*, 2020.
- Prashant Bharadwaj, Leah K Lakdawala, and Nicholas Li. Perverse Consequences of Well Intentioned Regulation: Evidence from India's Child Labor Ban. *Journal of the European Economic Association*, 18(3):1158–1195, 2020. URL <https://ideas.repec.org/a/oup/jeurec/v18y2020i3p1158-1195..html>.
- Christopher Blattman, Nathan Fiala, and Sebastian Martinez. Credit constraints, occupational choice, and the process of development: long run evidence from cash transfers in uganda. 2013.
- George J Borjas. Native internal migration and the labor market impact of immigration. *Journal*

- of Human resources*, 41(2):221–258, 2006.
- Leah Platt Boustan. Was postwar suburbanization “white flight”? evidence from the black migration. *The Quarterly Journal of Economics*, 125(1):417–443, 2010.
- Simon Burgess, Ellen Greaves, Anna Vignoles, and Deborah Wilson. What parents want: School preferences and school choice. *The Economic Journal*, 125(587):1262–1289, 2015.
- Brian C Cadena and Brian K Kovak. Immigrants equilibrate local labor markets: Evidence from the great recession. *American Economic Journal: Applied Economics*, 8(1):257–290, 2016.
- David Card. Immigrant inflows, native outflows, and the local labor market impacts of higher immigration. *Journal of Labor Economics*, 19(1):22–64, 2001.
- William J Carrington, Enrica Detragiache, and Tara Vishwanath. Migration with endogenous moving costs. *The American Economic Review*, pages 909–930, 1996.
- Elizabeth U Cascio and Ethan G Lewis. Cracks in the melting pot: immigration, school choice, and segregation. *American Economic Journal: Economic Policy*, 4(3):91–117, 2012.
- Diego Chaves-González and Carlos Echevarría Estrada. Venezuelan migrants and refugees in latin america and the caribbean: A regional profile. 2020.
- Dante Contreras and Sebastián Gallardo. The effects of mass migration on the academic performance of native students. evidence from chile. *Economics of Education Review*, 91: 102314, 2022.
- Ana C Dammert and Jose Galdo. Child labor variation by type of respondent: Evidence from a large-scale study. *World Development*, 51:207–220, 2013.
- Gustavo Henrique De Andrade, Miriam Bruhn, and David McKenzie. A helping hand or the long arm of the law? experimental evidence on what governments can do to formalize firms. *The World Bank Economic Review*, 30(1):24–54, 2016.
- Alan De Brauw, Daniel O Gilligan, John Hoddinott, and Shalini Roy. The impact of bolsa família on women’s decision-making power. *World Development*, 59:487–504, 2014.
- Suresh De Mel, David McKenzie, and Christopher Woodruff. One-time transfers of cash or capital have long-lasting effects on microenterprises in sri lanka. *Science*, 335(6071):962–966, 2012.
- Defensoría del Pueblo. Trabajo Infantil y Adolescente en Bolivia: Vulneración del Derecho a la Protección de las Niñas, Niños, y Adolescentes con Relación al Trabajo, 2022.

- Defensoría del Pueblo. La defensoría del pueblo alerta que no está funcionando el sistema de protección del trabajo infantil. 2021. URL <https://www.defensoria.gob.bo/noticias/la-defensoria-del-pueblo-alerta-que-no-esta-funcionando-el-sistema-de-proteccion-del-trabajo-infantil>.
- Andrew Dillon, Elena Bardasi, Kathleen Beegle, and Pieter Serneels. Explaining variation in child labor statistics. *Journal of Development Economics*, 98(1):136–147, 2012.
- Jeffrey A Dubin and Daniel L McFadden. An econometric analysis of residential electric appliance holdings and consumption. *Econometrica: Journal of the Econometric Society*, pages 345–362, 1984.
- Christian Dustmann and Albrecht Glitz. Migration and education. In *Handbook of the Economics of Education*, volume 4, pages 327–439. Elsevier, 2011.
- Eva Dziadula and Danice Guzmán. Sweeping It under the Rug: Household Chores and Misreporting of Child Labor. *Economics Bulletin*, 40:901–905, 2020.
- Eric Edmonds and Maheshwor Shrestha. The impact of minimum age of employment regulation on child labor and schooling . *IZA Journal of Labor Policy*, 1(1):1–28, 2012. doi: 10.1186/2193-9004-1-14. URL <https://ideas.repec.org/a/spr/izalpo/v1y2012i1p1-2810.1186-2193-9004-1-14.html>.
- Eric V Edmonds. Does minimum age of employment regulation reduce child labor? *IZA World of Labor*, 2014.
- C Elgin, A. Kose, F. Ohnsorge, and S. Yu. Understanding Informality. Working papers, London, Centre for Economic Policy Research, 2021. URL <https://www.worldbank.org/en/research/brief/informal-economy-database>.
- Jean Fares, Claudio E Montenegro, and Peter F Orazem. How are youth faring in the labor market? evidence from around the world. *Evidence from Around the World (November 1, 2006)*. *World Bank Policy Research Working Paper*, (4071), 2006.
- Lidia Farre, Francesc Ortega, and Ryuichi Tanaka. Immigration and the public–private school choice. *Labour Economics*, 51:184–201, 2018.
- David Figlio and Umut Özek. Unwelcome guests? the effects of refugees on the educational outcomes of incumbent students. *Journal of Labor Economics*, 37(4):1061–1096, 2019.
- David N Figlio, Paola Giuliano, Riccardo Marchingiglio, Umut Özek, and Paola Sapienza. Diversity in schools: Immigrants and the educational performance of us born students. Technical report, National Bureau of Economic Research, 2021.
- Louise Fox, Lemma W Senbet, and Witness Simbanegavi. Youth employment in sub-saharan

- africa: Challenges, constraints and opportunities. *Journal of African Economies*, 25(suppl_1): i3–i15, 2016.
- Paul J Gertler and Manisha Shah. Sex work and infection: what’s law enforcement got to do with it? *The Journal of Law and Economics*, 54(4):811–840, 2011.
- Eric D Gould, Victor Lavy, and M Daniele Paserman. Does immigration affect the long-term educational outcomes of natives? quasi-experimental evidence. *The Economic Journal*, 119 (540):1243–1269, 2009.
- Veronica Grembi, Tommaso Nannicini, and Ugo Troiano. Do fiscal rules matter? *American Economic Journal: Applied Economics*, 8(3):1–30, July 2016. doi: 10.1257/app.20150076. URL <https://www.aeaweb.org/articles?id=10.1257/app.20150076>.
- Sudhanshu Handa, Amber Peterman, Benjamin Davis, and Marco Stampini. Opening up pandora’s box: The effect of gender targeting and conditionality on household spending behavior in mexico’s progresra program. *World Development*, 37(6):1129–1142, 2009.
- Eric A Hanushek, John F Kain, and Steven G Rivkin. Disruption versus tiebout improvement: The costs and benefits of switching schools. *Journal of public Economics*, 88(9-10):1721–1746, 2004.
- Justine Hastings, Thomas J Kane, and Douglas O Staiger. Heterogeneous preferences and the efficacy of public school choice. *NBER working paper*, 2145:1–46, 2009.
- Chang-Tai Hsieh and Benjamin A Olken. The missing” missing middle”. *Journal of Economic Perspectives*, 28(3):89–108, 2014.
- Guido Imbens and Karthik Kalyanaraman. Optimal bandwidth choice for the regression discontinuity estimator. *The Review of economic studies*, 79(3):933–959, 2012.
- Guido W. Imbens and Thomas Lemieux. Regression discontinuity designs: A guide to practice. *Journal of Econometrics*, 142(2):615–635, 2008. ISSN 0304-4076. doi: <https://doi.org/10.1016/j.jeconom.2007.05.001>. URL <https://www.sciencedirect.com/science/article/pii/S0304407607001091>. The regression discontinuity design: Theory and applications.
- Scott A Imberman, Adriana D Kugler, and Bruce I Sacerdote. Katrina’s children: Evidence on the structure of peer effects from hurricane evacuees. *American Economic Review*, 102(5): 2048–82, 2012.
- International Labour Organization. *Child Labour: Global estimates 2020, trends and the road forward*. International Labour Office and United Nations Children’s Fund, New York, 2021.
- Seiro Ito, Aurélia Lépine, and Carole Treibich. The effect of sex work regulation on health and

- well-being of sex workers: Evidence from senegal. *Health economics*, 27(11):1627–1652, 2018.
- Michael Jetter and Teresa Molina. Persuasive agenda-setting: Rodrigo duterte’s inauguration speech and drugs in the philippines. *Journal of Development Economics*, page 102843, 2022.
- Akito Kamei. Lowering the Minimum Age for Child Labor in Bolivia. *Unpublished manuscript*, 2020.
- Victor Lavy, Olmo Silva, and Felix Weinhardt. The good, the bad and the average: Evidence on the scale and nature of ability peer effects in schools. Technical report, National Bureau of Economic Research, 2009.
- Edward P. Lazear. Job Security Provisions and Employment*. *The Quarterly Journal of Economics*, 105(3):699–726, 08 1990. ISSN 0033-5533. doi: 10.2307/2937895. URL <https://doi.org/10.2307/2937895>.
- Ethan Lewis. Immigration and production technology. *Annu. Rev. Econ.*, 5(1):165–191, 2013.
- Guilherme Lichand and Sharon Wolf. Measuring Child Labor: Whom Should Be Asked, and Why It Matters. *Working Paper*, 2022.
- Manfred Liebel. Bolivia bows to international pressure. *Development and Cooperation Op Ed*, 2019.
- Los Tiempos. Presidente no está de acuerdo con eliminar el trabajo infantil, 2013a. URL <https://www.lostiempos.com/actualidad/nacional/20131223/presidente-no-esta-acuerdo-eliminar-trabajo-infantil>.
- Los Tiempos. Presidente no está de acuerdo con eliminar el trabajo infantil, 2013b. URL <https://www.opinion.com.bo/articulo/cochabamba/identifican-ninos-eran-obligados-pedir-limosnas/20201118205333796294.html>.
- Marco Manacorda. Child labor and the labor supply of other household members: Evidence from 1920 america. *American Economic Review*, 96(5):1788–1801, 2006.
- Shanthi Manian. Health Certification in the Market for Sex Work: A Field Experiment in Dakar, Senegal. *Economic Development and Cultural Change*, (Forthcoming), 2021.
- Craig McIntosh and Andrew Zeitlin. Using household grants to benchmark the cost effectiveness of a usaid workforce readiness program. *Journal of Development Economics*, 157:102875, 2022. ISSN 0304-3878. doi: <https://doi.org/10.1016/j.jdeveco.2022.102875>. URL <https://www.sciencedirect.com/science/article/pii/S0304387822000451>.

- David McKenzie and Yaye Seynabou Sakho. Does it pay firms to register for taxes? the impact of formality on firm profitability. *Journal of Development Economics*, 91(1):15–24, 2010. ISSN 0304-3878. doi: <https://doi.org/10.1016/j.jdeveco.2009.02.003>. URL <https://www.sciencedirect.com/science/article/pii/S0304387809000170>.
- Fred Merttens, Alex Hurrell, Marta Marzi, Ramla Attah, Maham Farhat, Andrew Kardan, and Ian MacAuslan. Kenya hunger safety net programme monitoring and evaluation component. *Impact Evaluation Final Report: 2009 to 2012*, 154, 2013.
- Ministerio de Trabajo, Empleo y Previsión Social. *Memoria Institucional*, 2015-2018. Reports for all years are available here: https://www.mintrabajo.gob.bo/?page_id=4387.
- Carolyn M Moehling. State child labor laws and the decline of child labor. *Explorations in Economic History*, 36(1):72–106, 1999.
- Fernando Morales and Martha Denisse Pierola. Venezuelan migration in peru: Short-term adjustments in the labor market. Technical report, IDB Working Paper Series, 2020.
- Beatriz Muriel and Rubén Ferrufino. Regulación Laboral y Mercado De Trabajo: Principales desafíos para Bolivia. *Millenium Foundation Report*, 2012.
- Christopher Neilson et al. Targeted vouchers, competition among schools, and the academic achievement of poor students. *Documento de trabajo*. Yale University. Recuperado de http://economics.sas.upenn.edu/system/files/event_papers/Neilson_2013_JMP_current.pdf, 2013.
- OECD. *PISA 2018 Results (Volume I)*. 2019. doi: <https://doi.org/https://doi.org/10.1787/5f07c754-en>. URL <https://www.oecd-ilibrary.org/content/publication/5f07c754-en>.
- Pagina Siete. Evo morales contrario a prohibir trabajo infantil, 2013. URL <https://www.paginasiete.bo/sociedad/2013/12/23/morales-contrario-prohibir-trabajo-infantil-9390.html>.
- Mathieu Pedemonte. Fireside Chats: Communication and Consumers’ Expectations in the Great Depression. Working Papers 20-30, Federal Reserve Bank of Cleveland, October 2020. URL <https://ideas.repec.org/p/fip/fedcwq/88844.html>.
- Caio Piza, Andr© Portela Souza, Patrick M Emerson, and Vivian Amorim. The Short- and Longer-Term Effects of a Child Labor Ban. *The World Bank Economic Review*, 38(2):351–370, 11 2023. ISSN 0258-6770. doi: 10.1093/wber/lhad036. URL <https://doi.org/10.1093/wber/lhad036>.
- Vladimir Ponczek and Gabriel Ulyssea. Enforcement of Labour Regulation and the Labour Market Effects of Trade: Evidence from Brazil. *The Economic Journal*, 132(641):361–390,

2021.

Cristian Pop-Eleches and Miguel Urquiola. Going to a better school: Effects and behavioral responses. *American Economic Review*, 103(4):1289–1324, 2013.

Jaime Saavedra and Pablo Suárez. El financiamiento de la educación pública en el Perú: el rol de las familias. MISC, 2002.

Bruce Sacerdote. Peer effects in education: How might they work, how big are they and how much do we know thus far? In *Handbook of the Economics of Education*, volume 3, pages 249–277. Elsevier, 2011.

Bruce Sacerdote. Experimental and quasi-experimental analysis of peer effects: two steps forward? *Annu. Rev. Econ.*, 6(1):253–272, 2014.

Cristian Sanchez. Understanding school competition under voucher regimes. *Essays on educational vouchers, Ph. D. Dissertation, University of Maryland*, 2018.

Bryan A Stuart and Evan J Taylor. Migration networks and location decisions: Evidence from us mass migration. *American Economic Journal: Applied Economics*, 13(3):134–75, 2021.

Semih Tumen. Refugees and ‘native flight’ from public to private schools. *Economics Letters*, 181:154–159, 2019.

James Tybout. The missing middle, revisited. *Journal of Economic Perspectives*, 28(4):235–36, 2014.

Unión de Niños Niñas y Adolescentes Trabajadores de Bolivia. *”Mi fortaleza es mi trabajo” de las demandas a la propuesta: niños, niñas y adolescentes trabajadores y la regulación del trabajo infantil y adolescente en Bolivia*. UNATSBO, 2010.

U.S. Department of Labor. *Child Labor and Forced Labor Reports: Bolivia, 2011-2019*. Reports for all years are available here: <https://www.dol.gov/agencies/ilab/resources/reports/child-labor/bolivia>.

Diego Alejandro Vera-Cossio. Dependence or constraints? cash transfers and labor supply. *Economic Development and Cultural Change*, 70 (forthcoming)(4):null, 2021. doi: 10.1086/714010. URL <https://doi.org/10.1086/714010>.

World Bank. *World development report 2007: Development and the next generation*. The World Bank, 2006.

Guillermo Zamora Poblete and Carla Moforte Madsen. Why students change of school? an analysis starting from the family decisions. *Perfiles educativos*, 35(140):48–62, 2013.