Title
Essays in Labor Economics

Author
Serrano, Joaquin Augusto

Publication Date
2024

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Essays in Labor Economics

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

by

Joaquin Augusto Serrano

2024

ABSTRACT OF THE DISSERTATION

Essays in Labor Economics

by

Joaquin Augusto Serrano

Doctor of Philosophy in Economics

University of California, Los Angeles, 2024

Professor Adriana Lleras-Muney, Chair

This dissertation investigates factors influencing economic growth in developing countries through their impact on the labor market. Chapter 1 analyzes the effects of industrial robot adoption on local labor markets in Mexico. Using data on robot imports alongside administrative, survey, and economic census data and leveraging spatial and temporal variations across commuting zones, the study finds that robot adoption does not negatively affect formal employment and average wages, unlike developed countries. Additionally, regions importing robots see increased gross value added per worker, indicating productivity gains offset potential displacement effects. Chapter 2 studies the connection between intergenerational social mobility and economic development, using a novel dataset on educational mobility for 52 Latin American sub-regions. Employing a new weighting procedure that accounts for cohort participation in the economy, the study finds that higher intergenerational mobility is strongly linked to economic growth and improved development indicators. The final chapter examines discrepancies in reported Type I errors in linked datasets, which are crucial for studying historical intergenerational mobility. The study reveals that differences arise from reporting conditional versus unconditional error rates. The higher unconditional error rate, considering all algorithm links, better reflects the true error rate and affects inferences from historical data lacking a ground truth.

The dissertation of Joaquin Augusto Serrano is approved.

Martha Jane Bailey

Natalie Danielle Bau

Daniel Haanwinckel Junqueira

Adriana Lleras-Muney, Committee Chair

University of California, Los Angeles

2024

To my mother, my father, and my brother for their unconditional support

A mi madre, mi padre, y mi hermano por su apoyo incondicional

Contents

List of Figures

List of Tables

ACKNOWLEDGMENTS

I am deeply grateful to my advisors for their support during the completion of this dissertation. I am profoundly thankful to my mentor, Adriana Lleras-Muney, for her exceptional guidance, dedication, and patience throughout my time at UCLA. I am equally grateful to Martha Bailey for her permanent encouragement, generous support, and insightful advice that shaped the way I think about economics and the profession. I am incredibly thankful to Daniel Haanwinckel and Natalie Bau for their outstanding support and extensive discussions on the research that constitutes the first chapter of this dissertation. This project would not be the same without their advice. All of you taught me more than you know. I am also thankful to Michela Giorcelli for her continued support. I am grateful for the insightful comments shared during meetings and seminars by Dora Costa, Felipe Goncalves, Yotam Shem-Tov, Juliana Londono, and Maurizio Mazzocco.

I also thank my coauthors, Leonardo Gasparini, Guido Neidhofer, Matias Ciaschi, Jonas Helguertz, and Connor Cole. I particularly want to acknowledge Leonardo for his support and guidance, which have been invaluable to me since my undergraduate years and throughout this journey. I also thank David Jaume for his time, advice, and discussions about the first chapter of this dissertation.

I am grateful to my friends and colleagues for making my PhD a very happy and rich experience. I could not have finished the Ph.D. without the support of Vicky Barone, Diana Flores Peregrina, Domenico Fabrizi, Tomas Guanziroli, Ariadna Jou, Nano Palleja, Fatih Ozturk, and Fernanda Rojas Ampuero. You made this journey very enjoyable and unforgettable. I also want to thank Sungwoo Cho, Francesco Gabriele, Johnny Huynh, Jonathan Kowarski, Calvin Kuo, Santiago Mosquera, Manu Navjeevan, Daniel Ober-Reynolds, Daniel Perez, Brian Pustilnik, Benjamin Pirie, Estefania Saravia, Vickie Wang, and many others for your conversations, time and help in different stages of the PhD.

This research was made possible by the Dissertation Year Fellowship through the UCLA

Graduate Division. I benefited from resources provided by the California Center for Population Research at UCLA (CCPR), which receives core support (P2C-HD041022, NICHD). I thank CCPR for providing resources and a great research environment during my Ph.D. Chapter two was published in the Journal of Economic Growth. No changes were made to the published version. Appropriate credit is given, and the publication can be accessed here. The material is available under the Creative Commons Attribution 4.0 International License. To view a copy of this license, visit http://creativecommons.org/licenses/by/4.0/.

VITA

Joaquin Augusto Serrano

Education

University of California, Los Angeles

M.A. in Economics 2018-2019

Universidad Nacional de La Plata

M.S. in Economics 2013-2016

B.A. in Economics 2008-2012

Publications

"Social mobility and economic development" (2023) *Journal of Economic Growth* –with Neidhöfer, G., Ciaschi, M., and Gasparini, L.

"Economic Cycle and Deceleration of Female Labor Force Participation in Latin America" (2019) *Journal for Labour Market Research*, Vol. 53. –with Gasparini, L., Marchionni, M. and Gluzmann, P.

"Educational Inequality and Intergenerational Mobility in Latin America: A New Database" (2018) *Journal of Development Economics*, Vol. 134, pp. 329-349. –with Neidhöfer, G., and Gasparini, L.

"The Inequality Possibility Frontier in Latin America" (2017) *El Trimestre Económico*, Fondo de Cultura Económica, México, Vol. 84(2), Núm. 334 –with Benzaquén, I.

"Characterizing Female Participation Changes" (2015) In Gasparini, L. and Marchionni, M., editors, *Bridging Gender Gaps? The Rise and Deceleration of Female Labor Force Participation in Latin America*, volume 1, chapter 4, pp. 199–260. Center for Distributive, Labor and Social Studies, 1st edition. –with Gasparini, L., Marchionni, M., and Badaracco, N.

Research in Progress

"Do Robots Necessarily Displace Workers? Evidence from Mexican Local Labor Markets" (*Job Market Paper*)

"Reconciling the Evidence on False Links in Historical Data" (with Martha Bailey, Connor Cole, and Jonas Helgertz)

"Social Mobility and Economic Development: A Lifetime Analysis of Latin American Cohorts" (with Guido Neidhöfer, Matias Ciaschi, Antonio Ciccone, and Leonardo Gasparini)

"Effectiveness and Consequences of Occupation and Industry-specific Minimum Wages" (with Tomas Guanziroli)

Relevant Positions

CCPR Student Affiliate

California Center for Population Research, UCLA. Since 2019.

Associate Researcher

CEDLAS, Universidad Nacional de La Plata. Since 2018

Research Assistant with Prof. Martha Bailey, LIFE-M Project, UCLA. 2020-2022.

Research Assistant with Prof. Michela Giorcelli, *Project: Developing Down Under: The Case of Italian Prisoners of War in Australia*, UCLA. 2019-2021.

Fellowships, Honors, and Awards

Dissertation Year Fellowship, Graduate Division, UCLA. 2023

Graduate Summer Research Mentorship Program, Department of Economics, UCLA. 2019

Graduate Division Fellowship, Department of Economics, UCLA. 2018-2022

Introduction

The dissertation studies factors that potentially affect economic growth through labor markets in developing countries. Specifically, in the first chapter, I analyze how the recent increase in industrial automation in Mexico is affecting labor markets and how these effects differ from those found in developed countries. In the second chapter, I study the relationship between improvements in intergenerational mobility and development through labor markets in Latin America. This includes analyzing how unequal opportunities to invest in human capital cause talent to be misallocated in the labor market. In the third chapter, I explore issues on the performance of linking methods, essential for creating databases that allow researchers to study labor markets and intergenerational social mobility using historical records.

In Chapter 1, I use the rapid rise of automation in Mexico to examine the labor market effects of industrial robot adoption in industrializing countries. To do so, I combine unique data on robot imports with administrative, survey, and economic census data. I exploit spatial and time variations in robot imports across commuting zones to estimate the impacts of automation on labor market outcomes. I find that industrial robot adoption does not negatively affect formal employment and average wages in Mexico; these results contrast with findings for developed countries. I can rule out negative effects of size as large as -0.7% for formal employment and -0.01% for formal wages. Additionally, I show that after a CZ starts importing robots, its aggregated gross value added per worker increases, consistent with the productivity effect dominating or compensating for any adverse displacement effect. Despite a substantial pool of unskilled labor, automation in industrializing economies such as Mexico can yield nonnegative employment effects because of the significant productivity enhancements associated with a higher price elasticity of demand and the early stages of automation.

In Chapter 2, coauthored with Guido Neidhofer, Leonardo Gasparini, and Matias Ciaschi, I analyze the relationship between social mobility and economic development in Latin American countries. The study creates a sub-national region-level panel data for 10 Latin American countries. It develops a new methodology to connect cohort- and year-level observations by weighting the degree of mobility of a cohort based on its contribution to the overall economic performance of the respective country in each year. The results suggest that higher intergenerational mobility is associated with rising income per capita, income growth, and other development indicators. The study's findings have significant policy implications, indicating that improving social mobility is not an equity-efficiency trade-off, and it can have positive economic returns. The hypothesis behind this pattern is that improving social mobility can lead to more efficient accumulation of human capital, reduce misallocation of talent, and enhance economic performance. Despite concerns that improving social mobility may cause inefficiencies in the short-run, the positive long-run impact on the economy justifies its use. The study also reveals that even when controlling for contemporaneous income inequality, intergenerational mobility remains relevant for explaining economic development.

Chapter 3, coauthored with Martha Bailey, Jonas Helgertz, and Connor Cole, aims to reconcile the contradictory evidence of false links in historical data. Technology to link national surveys, administrative data, and historical records is revolutionizing economic history. Still, recent studies have shown that machine algorithms used to create large-scale longitudinal and intergenerational datasets have significant linking error rates, ranging from 5 to 50 percent. To address this issue, we analyze linking error rates across US censuses and the Oldest Old data. We find that the difference in linking error rates depends on whether Type I error rates are conditioned on links in the ground truth. The study shows that unconditional rates, which consider the total number of potential links, tend to be higher than conditional rates. Given that researchers typically lack access to the ground truth dataset, unconditional rates are more relevant for accurate inference, which reveals limitations of many linking methods used in the literature. Additionally, the study explores how linking errors vary across subpopulations, with disadvantaged groups exhibiting higher error rates. Finally, the study demonstrates the impact of conditional and unconditional linking errors using two large-scale linked datasets (MLP and CLP) and provides recommendations for assessing measurement errors in linked historical data.

Chapter 1

Do Robots Necessarily Displace Workers? Evidence from Mexican Local Labor Markets

1. Introduction

Machines—from assembly lines and electric motors to computers and modern robots—have fundamentally reshaped workplaces. In 2022, a staggering 3.5 million industrial robots were operational, with a remarkable annual growth rate in this number of 31% (IFR, 2022). As robotics, artificial intelligence, and computing power advance, there is a growing fear of significant job losses (Acemoglu and Restrepo, 2020; Frey and Osborne, 2017; Arntz et al., 2019). These concerns prompt essential questions about the societal implications of automation. However, amid the uncertainty, this technological shift also offers the potential to boost productivity. Previous research has predominantly examined industrialized countries, demonstrating that automation reduces employment and wages, particularly among low-skilled workers engaged in routine tasks (Acemoglu and Restrepo, 2020, 2022a; Dauth et al., 2021; Chiacchio et al., 2018).

Theoretically, the expected sign of the net effect of automation on outcomes is ambiguous (Acemoglu and Restrepo, 2019, 2020). On the one hand, automation replaces workers (the displacement effect), leading to declines in employment and wages in the short run. On the other hand, automation reduces production costs and allows firms to lower prices and increase sales (the productivity effect), leading to increases in employment and wages. Productivity effects in industrializing countries may be larger because automation occurs in sectors exposed to a more competitive product market such that demand is more sensitive to price changes.[1] In addition, these economies are at an early stage of automation where there is more room for large marginal gains. However, empirical evidence on the effects of automation in these economies is still lacking because it is a relatively recent process and data on robot adoption are scarce.

This paper overcomes these challenges by providing empirical evidence on the labor market effects of automation in Mexico. Mexico serves as an ideal case study for several reasons. First, the country has witnessed a remarkable surge in industrial robot adoption over the past decade, with an average annual growth rate of 40% since 2011. This transformation is especially significant within the manufacturing sector, which accounts for a large share of employment.[2] Second, Mexico provides administrative data on industrial robot imports, including the location and timing of robot adoption at fine geographic levels. These data measure robot adoption much more precisely in a given area than the samples in previous works, which typically rely on time series data by industry at the national level. I complement

[1]The productivity effect will be larger, especially when product demand is relatively more elastic. In developing countries, robot adoption is more prevalent in export manufacturing, which is exposed to foreign competition, and therefore demand is arguably more elastic. Firms in this sector are more likely to invest in automation because of the relatively large fixed cost and the associated product quality upgrade, particularly when interdependence of components is a critical factor (Artuc et al., 2022). Moreover, domestic consumers have relatively low incomes and are more likely to be price sensitive.

[2]Mexico is also among the most populous Latin American countries, second only to Brazil, with a population of approximately 126 million people, 60 million of whom are of working age, among whom 62% participate in the labor force.

these with publicly available administrative data from the Mexican Social Security Institute (IMSS, for its Spanish acronym) to examine labor market outcomes in the formal sector, supplemented by insights from the National Occupation and Employment Survey (ENOE) on the informal sector and from Mexican economic censuses on changes in productivity and balance sheet information across all Mexican firms. These high-quality data allow a comprehensive examination of the effects of robot adoption at the local labor market level.

Identifying the causal effects of automation is challenging because of a combination of data limitations and endogeneity issues. Endogeneity concerns such as omitted variable and simultaneity biases complicate the analysis. Consider, for example, a scenario where an industry faces a negative labor supply shock and turns to automation in part to mitigate rising labor costs. Neglecting this dynamic would overestimate the negative effects of automation on employment. Conversely, if industries facing positive product-market shocks or relaxation of their financial constraints adopt robots to help them increase production, neglecting this dynamic would lead to overestimation of the positive impact of automation on employment.

To address the endogeneity issue, I employ multiple complementary strategies. My primary approach is to take advantage of timing-based event study methodologies. I identify an event as the first year in which a commuting zone (CZ) receives industrial robot imports, and then, I group CZs into cohorts based on these event years. To mitigate selection effects, I limit the sample to only CZs that import robots at some point. Therefore, the comparison group includes only CZs that import robots but do so later than the CZ under consideration. This method assumes that the timing of events is random and is not influenced by anticipation (pretreatment periods are not affected by future events) and that the outcomes for CZs importing robots and the comparison group follow parallel trends. I explore any systematic link between importation timing and observable factors to validate this assumption. Furthermore, I include granular industry-level trends to mitigate potential biases and account for industry-specific shocks, such as increasing competition. Assuming that the comparison groups and the specific fixed effects and trends account for unobservable shocks, I exploit the remaining exogenous variation arising from a combination of changes in robot supply (cost reduction and improved capabilities) and exogenous variation in exposure across Mexican regions due to supply-side factors, like supply chain disruptions, random events at ports or shipping facilities, or differences in trade networks. For instance, when the cost of robots produced in Korea falls, producers with existing connections to Korean robot producers are likely to benefit first.

I find that the rapid adoption of robots in Mexico in the last decade did not lead to significant negative effects on formal employment or average real formal wages. The effect of robot adoption on employment is zero in the first year of adoption. Then, it stays stable for the first four years, with an average effect of 1.6% that is not statistically significant at the 5% level. I can rule out negative effects larger than -0.7% on formal employment and larger than -0.01% on formal wages for the average of these four years. I also provide evidence that these results do not hide compositional effects of the CZs' cohorts around the event time. This zero effect could hide potential heterogeneity across groups. However, there are, at most, minor negative effects, even among groups expected to be more exposed to the negative effects of robot adoption, based on findings from industrialized countries. These groups include workers in the manufacturing sector, workers in smaller CZs, medium-wage and prime-age workers, and workers in occupations related to industrial robots.

When I estimate the effects over an average of 5 to 8 years after the event, they consistently remain nonnegative, increase in magnitude, and achieve statistical significance. However, it is important to acknowledge that changes in composition might influence these long-term effects. This is because cohorts experience the event in different years and are observed over varying time frames after the event. Consequently, I interpret these results cautiously, concluding that automation may not lead to significant job displacement effects. This highlights that robot adoption does not necessarily harm workers and could hold transformative potential for labor market dynamics.

The main concern with my estimation method is that the adoption of robots is endogenous, and in particular could be related to positive demand shocks. To address this, I use alternative methods, such as event study approaches, triple differences (DDD), and synthetic difference-in-differences (SDD). The DDD method compares manufacturing industries that are high robot adopters with non-manufacturing non-robot adopters industries within each CZ. This allows me to net out specific shocks at the CZ level, such as positive growth trends for all the sectors in regions adopting robots first. This fact is important since, a priori, we might expect that growing CZs would adopt robots first. The SDD enables me to address potential violations of the identifying assumption of parallel trends. The method selects a different set of controls for each treated unit, composed of a weighted average of CZ that matches that unit's pre-trend. Therefore, the control group is essentially individualized, and I can investigate treatment effect heterogeneity across cohorts, relevant to analyze long-term effects avoiding compositional changes. I consistently observe the same qualitative patterns as the baseline estimates across different methods and outcome measures. In addition, I conduct a sensitivity analysis that considers varying degrees of potential violations of parallel trends in post-event periods. The results confirm the robustness of the main findings: a substantial deviation from parallel trends would be required to yield negative employment and wage estimates.

The results are also robust when employing different measures of robot adoption. One might suspect that the nonnegative effect is due to the need for a certain number of industrial robots in the CZ. Therefore, I define the event based on the cumulative value of robot imports surpassing a predefined threshold. Alternatively, I use a measure of intensity based on the cumulative number of years that a CZ receives imports. The analysis is also robust if I use alternative measures for my primary outcomes, formal employment, and average wages.

To explore potential mechanisms underlying industrial robot adoption, I employ Mexican economic censuses covering the universe of Mexican firms and balance sheet information. Within this analysis, I consistently observe that robot adoption enhances gross value added (GVA) per worker across all industries, consistent with productivity increases. The effect on the labor share is negative but imprecisely estimated. When I restrict the sample to the manufacturing sectors more likely to adopt automation, I find an increase in GVA per worker, input cost, payroll, and fixed assets, indicating increased productivity, wages, and investment. Despite these gains, the automation effects on the labor share remain negative but not statistically significant. In nonmanufacturing sectors, adopting robots generally results in statistically nonsignificant effects; however, when statistically significant, these effects tend to be small and lack a systematic pattern. These findings suggest that productivity effects dominate or compensate for potential adverse displacement effects in the case of Mexico in the short to medium run.

This study contributes to the literature on automation's impact on local labor markets, not only in developed countries (Acemoglu and Restrepo, 2020; Dauth et al., 2021; Chen and Frey, 2021; Chiacchio et al., 2018; Graetz and Michaels, 2018) but also in industrializing economies (Brambilla et al., 2023; Calì and Presidente, 2021a; Rodrigo, 2021), for which the evidence is still very limited and not conclusive.[3] Most of these papers rely on industry-level proxies of robot adoption, which assume uniformity in robot adoption within a particular industry across different districts, as implied by Bartik measures, generating potential attenuation bias. In addition, this data do not allow for exploiting variation within industry in the adoption of robots –variation that I can exploit in my setting and which allows me to control for industry–year, CZ–industry, and CZ–year specific shocks in some specifications. Because of this more granular level of variation I can control for additional shocks and therefore obtain potentially different estimates. I complement this literature by

[3]Brambilla et al. (2023) examines the impact of domestic robots on labor markets in Argentina, Brazil, and Mexico, finding increases in unemployment and informality but no statistically significant negative effect on the employment rate and wages. Rodrigo (2021) identifies positive effects on firms' productivity but no aggregate employment impact in Brazil. For Indonesia, Ing and Zhang (2022) and Calì and Presidente (2021b) observe significant productivity and employment gains from automation in manufacturing.

leveraging administrative data with precise information on the timing and location of domestic automation adoption, which allows me not to rely on shift-share measures and to use alternative empirical strategies.

There are several studies analyzing the impact of US automation on developing countries through trade linkages and reshoring (Kugler et al., 2021; Stemmler, 2020; Artuc et al., 2022). Two evaluate automation in Mexico (Artuc et al., 2019; Faber, 2020) during 1990–2015 and include exposure to domestic automation in their analysis, using similar industry-level data, Bartik instruments, and periods. However, they arrive at different conclusions about the wage–employment effects of the increase in domestic robots, ranging from null to negative effects. My paper complements these studies by using district-level data on robot adoption, introducing a novel empirical strategy, and focusing on a more recent period marked by accelerated domestic robot adoption in Mexico. It also offers a context where domestic production of robots is minimal and firms primarily depend on imports, a scenario different from the ones in previous studies centered on countries with active domestic robot production (Acemoglu et al., 2021; Koch et al., 2021; Dixon et al., 2019; Humlum, 2020; Acemoglu et al., 2022). These contributions enrich our understanding of how automation is shaping the future of work and economic development in Mexico and beyond, offering valuable insights for policymakers.

The paper is structured as follows. Section 2. describes the data sources, including industrial robot adoption and labor market outcomes, while offering descriptive evidence that sheds light on the association between these two factors. Section 3. gives an overview of the institutional background of the labor market and the rise of automation in Mexico. Section 4. describes the empirical strategy for identifying the impact of industrial robot adoption. Section 5. presents the main results on labor market outcomes, and Section 6. shows results on local mechanisms that may explain the effects of industrial robots. Section 7. presents a conceptual framework and discusses why the effects of automation could differ in developing countries such as Mexico from those found for developed countries. Section 8. concludes.

2. Data

To study the effect of robot adoption on Mexican labor markets, I compile data from multiple sources: (i) detailed administrative records on the universe of industrial robot imports in Mexico provided by the Ministry of Economy; (ii) data on formal employment, publicly accessible and released by the Mexican Social Security Institute (*Instituto Mexicano del Seguro Social*, IMSS); (iii) the Mexican employment survey (ENOE); (iv) information from the 1990 and 2000 population censuses; and (v) information from economic censuses spanning from 2003 to 2018.

I conduct the analysis at the local labor market level.[4] There are many ways of defining labor markets; generally, they are defined based on local administrative (state or municipality) boundaries, but these only sometimes coincide with economically relevant boundaries. Therefore, I use commuting zones (CZs) as the unit of observation, grouping 2438 municipalities (similar to US counties) into 1806 labor markets following Faber (2020). His algorithm follows a similar procedure designed by Atkin (2016), which clusters municipalities into CZs based on the intensity of commuting between them.[5]

[4]Identifying the relevant units of analysis is crucial for analyzing the effects of robot adoption on different outcomes. Firms that adopt robots can extend their impact to other related firms, within their industry or even within their labor market by sharing the same labor and product suppliers, consumers, and competitors (Acemoglu et al., 2023; Acemoglu and Restrepo, 2022b; Aghion et al., 2022). Making comparisons solely between robot-adopting and nonadopting firms within a single labor market could yield biased estimates of the impact of automation. To prevent this issue, one should conduct the analysis at the labor market level, capturing the direct effect of robot adopters and the indirect impact on other economic agents within the same labor market. Furthermore, supporting evidence indicates that workers do not exhibit perfect mobility across geographic regions in the US (Autor and Dorn, 2013); the case of Mexico is likely to be similar in this regard.

[5]The algorithm operates as follows: (i) clustering all municipalities within a metropolitan area into one larger municipality; (ii) computing the intensity of commuting from each municipality to any other municipality by dividing the number of people who commute by the number of residents in the origin municipality; and (iii) clustering municipalities into CZs if more than 10% of residents of either municipality commute into the other. See Faber (2020) for more details.

2.1 Robot adoption data

To measure robot adoption at the labor market level, I use information on the universe of industrial robot imports from customs data collected by the Mexican Ministry of Economy. I follow other papers that use a similar measure for developed countries (Acemoglu et al., 2023; Bessen et al., 2020) and define robots using the code 847950 from the international trade codes of commodities (Harmonized System, 2012), which is defined as "industrial robots, not elsewhere specified or included."[6] According to the International Standards Organizations, an "industrial robot is an actuated mechanism programmable in two or more axes, with a degree of autonomy, moving within its environment, to perform intended tasks."

The data provide information about the month and municipality in which the robots were internationally purchased, the country of origin, and the import value. In this paper, I use the publicly available version of these data, which are aggregated to the municipality–month level. The aggregated import value is available if at least three firms imported in the given period. As explained previously, I map municipalities to CZs based on the intensity of commuting flows between them. This transition from municipality–month to CZ–year aggregation allows a more comprehensive examination of the effects of automation on local labor markets. The dataset covers the period from 2006 to the present.

My data on robot imports at the municipality level allow me to measure the exposure to robot adoption more accurately than has been done in previous literature. The standard approach in the literature is to use a Bartik-type measure, which relies on national industry variation in robot shipments (based on data from the International Federation of Robots

[6] The Harmonized System (HS) is an internationally recognized product classification system that designates unique six-digit codes for various categories and commodities. HS code 847950 specifically pertains to industrial robots not categorized under any other code. These robots are capable of autonomously performing a range of tasks, such as welding, painting, assembling, or handling material, without the need for human intervention. Industrial robots falling under this HS code encompass a variety of machines, including robotic arms employed for manipulating objects or tools in manufacturing settings and robotic vehicles designed for transporting materials or goods within warehouses or production facilities. For more information, see WCO (2013), and visit the website https://www.wcotradetools.org.

(IFR)) and the industry composition of each commuting zone (Acemoglu and Restrepo, 2020; Graetz and Michaels, 2018; Dauth et al., 2018). This approach assumes that the same number of robots per worker are installed in each industry across all commuting zones, which may not be realistic and could generate biased estimates. In addition, in the case of Mexico, the IFR data start in 2011 and are aggregated along with the US and Canada for previous periods. My data can overcome these limitations, as they capture the actual location and year of robot imports over an extended period.

My measure of robot adoption indicates whether the CZ is receiving industrial robot imports in a given year. From the data, I observe 104 different CZs receiving at least one industrial robot importation during 2006–2022. I exclude from the sample the CZs receiving imports in 2006 since this is the first year with available data and I assume that these CZs were receiving robot imports even before 2006. Figure 1.2 displays the geographic distribution of robot imports across CZs within the restricted sample. The robot imports are distributed across the entire territory, with the Bajío area and the northern region containing 46% of the CZs that receive robot imports, followed by the southern central region (15%) and the northeastern region (13%).[7] Within these regions, the percentage of CZs receiving imports varies from 3% to 13%, providing ample geographical variation for estimating effects.

Using custom records may have limitations; for example, it is possible that firms obtain robots from local manufacturers instead of importing them. However, this concern is insignificant in Mexico, where almost all robot adopters import them from the US, Japan, or Europe.[8] To validate the accuracy of the customs data, I compare them with national-level IFR data, finding a solid 98% correlation between the number of importing municipalities and

[7]The Bajío area includes Aguascalientes, Guanajuato, Querétaro, San Luis Potosí, and Zacatecas. The northern region comprises Baja California, Baja California Sur, Chihuahua, Durango, and Sinaloa. The southern central region encompasses Hidalgo, Puebla, Tlaxcala, and Veracruz de Ignacio de la Llave. The northeastern region includes Coahuila de Zaragoza, Nuevo León, and Tamaulipas.

[8]Global robot manufacturing is highly concentrated, with the top ten suppliers in 2021 being Japan, Germany, Italy, China, the United States, Denmark, France, South Korea, Austria, and Sweden, which accounted for 80% of total robot exports. In contrast, Mexico ranks 24th, contributing just 0.4% to global industrial robot exports, with a trade deficit in this sector.

the stock of robot installations over the period 2011–2020 (see Figure 1.1 for a comparison of robot adoption measures).[9,10] This distinctive feature of my research setting is advantageous because, in the developed countries studied in similar works, robot manufacturing is substantial and robot imports capture only a fraction of total robot adoptions.

Another potential concern with the use of import records is the possible concentration of robots in cities where firms' headquarters are located rather than where their plants are situated. In Mexico, it is relatively uncommon for firms that use robots to have multiple plants. Among the firms that imported robots between 2006 and 2022, a significant majority, approximately 85%, have only one plant, and 88% operate within the same CZ. Additionally, since the automotive sector is responsible for most robot adoption in Mexico, I identify the universe of large automotive plants and verify that most are located in CZs receiving robot imports (see Figure 1.2).

To address the possibility that firms import machinery through domestic robot integrators specializing in installing robots at local plants, I examine the practices in Mexico to understand whether this is a significant issue. My investigation reveals that the leading robot integrators in Mexico do not import robots or any other type of machines.

2.2 Labor market outcomes

Formal labor market: To assess the impact of industrial robot adoption on formal employment, I leverage a publicly accessible administrative dataset released monthly by the IMSS. This dataset provides municipality-level information on the total count of registered (i.e., formal sector) workers and total payroll by firm size, industry (up to the 4-digit sector),

[9] I obtain a similarly strong correlation if I compare the IFR stock of robots with the stock of the number of importing firms or the stock of the importation values.

[10] The IFR defines an industrial robot as automatically controlled (not manually operated through a joystick or pushbuttons), reprogrammable without physical alteration, multipurpose-manipulable, programmable in three or more axes, fixed in place or mobile, and for use in industrial automation applications. The definition does not include machines such as elevators, ATMs, smart washing machines, transportation tools, textile looms, software, or autonomous cars.

gender and wage categories. The IMSS administers healthcare services to formal employees in the private sector, and because of the mandatory nature of contributions from private firms, it holds comprehensive information on their workers.

Using these data, I compute the total number of formal employees in a given CZ and 4-digit sector and the corresponding average wage by dividing the total payroll by the total number of workers. Although the frequency of the data is monthly, I focus on the information in December of each year. For the primary analysis, I do not restrict the sample by age. However, the data allow me to distinguish employee age groups and investigate potential heterogeneity.

One evident caveat of this dataset is that it does not include the public or informal sector, which is relatively large in Mexico. The informality rate in Mexico between 2005 and 2018 varied from 56% to 59% of the employed population, according to INEGI. Nevertheless, drawing from prior research, most firms that import and adopt industrial robots are more likely to be large in terms of both employees and revenues (Koch et al., 2021; Acemoglu et al., 2022) and therefore are less likely to be in the informal sector (Ulyssea, 2018; Porta and Shleifer, 2008, 2014). Another limitation of the dataset is its lack of information on hours worked, part-time or full-time employment, specific occupations or tasks, and other characteristics of workers such as educational attainment, ethnicity, and immigration status. However, the IMSS data allow me to dividing employees into different wage groups, in terms of the Mexican minimum wage level.[11]

Mexican labor force survey (ENOE). To shed light on the effects of robot adoption on transitions between the formal and informal sectors, into unemployment, and in and out

[11] Wages in the IMMS dataset are reported as daily taxable income (*salario base de cotización*), which may include various forms of compensation, including paid vacation and end-of-year bonuses. Nonetheless, they may exclude other benefits or compensation not subject to labor income taxation. Notably, wages are both bottom- and top-coded, with approximately 1.3% and 1.7% of observations falling into these categories, respectively. For further details, refer to Puggioni et al. (2022). To ensure consistency, I convert these daily wages to monthly wages by multiplying them by 365 days and dividing them by 12 months.

of the labor force, I use the Mexican labor force survey, called ENOE. This survey is at quarterly frequency and covers most of the localities in urban centers, which I aggregate into CZs. To maintain comparability with the IMSS data, I focus on the fourth quarter of each year between 2005 and 2020. I calculate the count of formal and informal employees, unemployed individuals, and workers out of the labor force. I define formal workers as those affiliated with IMSS and informal workers as those lacking affiliation with any social security subsystems (IMSS, PEMEX, or ISSSTE). However, given ENOE's different sample of localities, I repeat my primary analysis using the IMSS data but the ENOE sample.[12]

1900, 2000, and 2015 Mexican censuses: To calculate general labor market and socioeconomic characteristics for all the Mexican CZs, I use the 1990, 2000, and 2015 population censuses conducted by the Mexican National Institute of Statistics and Geography (*Instituto Nacional de Estadística, Geografía e Informática*, INEGI). I use this information to characterize which CZ observable characteristics are associated with industrial robot import events. Additionally, for the sake of completeness and to facilitate comparisons with existing literature, I compute Bartik-type measures of robot exposure using IFR data.

Mexican economic censuses 2003–2018: I use the Mexican economic census of all firms in the country to compute measures of firm productivity and capital investment within each CZ and sector. The publicly available data are at the municipality–6-digit industry level and can be aggregated to CZ level. The most important limitation of these data is that the frequency is not yearly but every five years.

[12]The ENOE results could be noisier because the survey and administrative data, despite being highly correlated, differ in focus: the administrative data center on firm locations, while the survey data cover worker addresses. Despite my using representative localities, discrepancies may persist because of potential worker location differences (Puggioni et al., 2022, Gutierrez et al., 2023).

3. Institutional Background

3.1 Mexican labor markets and the manufacturing sector

Mexico provides a compelling case for studying the impact of automation on labor markets, particularly within its important manufacturing sector, a key sector for economic development and structural change. This sector is a vital component of Mexico's economy, contributing approximately 17.4% to its GDP and accounting for 30.8% of its total exports and 25.6% of total employment in 2019. Mexico manufactures many products within this sector, including automobiles, aerospace components, electronics, medical devices, textiles, apparel, footwear, and furniture.

The strong integration of Mexico into global value chains, especially with the United States, helps us understand the automation phenomenon. Firms engaged in global value chains face intense competition, particularly from China. They are also typically part of intricate production networks in which precision and reliability, which robots improve, are highly valuable. This explains why industrializing economies, including Mexico, are adopting automation, even with a labor force with relatively lower wages (Artuc et al., 2022). The competitive dynamics in Mexico's industrial landscape can imply that these firms face a higher demand elasticity in the product market, a factor affecting the size of the productivity effect of automation.

I focus the analysis on the period from 2005 to 2020, during which an increasing number of firms started importing and adopting industrial robots between 2014 and 2017 (see Figure 1.1). The rise of automation coincides with some national economic trends, like decreases in poverty and inequality and increases in income, GDP and foreign direct investment. The unemployment rate in Mexico is low (3% of individuals in the labor force) and remained relatively constant during my period of interest. These patterns already suggest that automation does not have a large negative effect, at least at the aggregate level. But they also suggest caution since automation coincides with potentially positive economic shocks (as noted by Brambilla et al. (2023) for Latin America), which is what my empirical strategy will address.

The sector in which automation has occurred at the largest rate is the manufacturing sector. The most potentially affected workers in this sector are not low–wage workers since 69% of these workers earn 2 to 5 times the minimum wage, and there are fewer low-wage workers than in other sectors (Figure A.2, Panel A). The most common occupation in this sector (and thus most likely affected) is industrial operators (65.6%, 69% if formal), who are highly susceptible to automation and are predominantly middle-wage earners. At the same time, professionals and managers typically earn high wages (see Panel B). Most workers in this industry are young (68% aged 20-44). Note that these groups differ from those most affected in industrialized countries, which are relatively low–wage earners in those countries. In my analysis, I will investigate the effects of automation for each of these more and less affected groups.[13]

3.2 Industrial automation in Mexico

Mexico has experienced a rapid increase in industrial robot adoption in the last decade. Figure 1.1 illustrates this trend, showing that the stock of operational robots surged at an average annual rate of 40% since 2011, reaching over 45,000 units by 2020. Correspondingly, the value of robot imports has surged since 2006. Simultaneously, the number of municipalities embracing industrial robot imports has risen steeply, starting at 68 in 2006 and reaching 108 by 2018.[14] A driving force behind this rapid automation is the automotive sector and other manufacturing industries, including computer, plastics, basic metals, pharmaceutical,

[13]Panel C in Figure A.2 indicates that individuals aged 25 to 45 primarily fall into the medium-wage group, with older workers enjoying higher salaries, especially in formal manufacturing. Additional details are available in Table A.1, which profiles the workers in each wage category, highlighting the prevalence of industrial operators aged 20–44 among medium-wage earners in manufacturing.

[14]This growth is mirrored by the expansion in the number of firms importing robots, which surged from 304 firms in 2006 to 649 in 2018.

and equipment manufacturing.

Figure A.4 shows that, although Mexico's automation growth surpasses that of many nations, its level remains a fraction of that in the US or EU countries. This observation underscores Mexico's early stage of automation, wherein the incremental productivity gains from robot adoption might be comparatively larger. Other Latin American countries, such as Brazil, exhibit substantially lower automation levels that have scarcely changed since 2004.

The contribution of Mexican industry to global value chains increased substantially beginning in the mid to late 2000s, partially explained by the good performance of the automotive sector (Chiquiar, 2019; Iacovone et al., 2022). Sectors tightly integrated into global value chains prioritize quality due to the high costs of errors, driving automation adoption even in countries with lower labor costs (Rodrik, 2021; Artuc et al., 2022; Robles and Foladori, 2019). Sectors that adopted robots may be on different trajectories as a result. To address this, I include industry-year fixed effects, therefore only comparing firms within sectors that adopt and do not adopt robots in a given year.

Then, why do some firms adopt robots within a specific industry while others do not? Aside from firm-specific shocks, most of this variation comes from supply-side shocks in the robot market. This fact could arise from differences in events at ports or shipping facilities (like accidents, strikes, or changes in regulations), exogenous supplier disruptions (like shortages of electrical and electronics components), or the history of trade ties with countries that supply robots. For example, If a new robot is developed in Japan (or if the price of Japanese robots falls), firms in CZ that have a previous trade connection with Japan are likely to benefit first. These benefits could stem from existing distribution networks, better after-sales service, preferential access, better information, etc.[15] Unfortunately, I do not observe supply chains at the firm level –in fact, my data is at the commuting zone level. So, I cannot exploit this variation. Instead, I will employ several empirical strategies to

[15]The importance of networks and granularity for trade has been documented by recent studies (Chaney, 2014; Gaubert and Itskhoki, 2021).

match CZ as closely as possible and assume that the remaining variation is exogenous from the point of the firm and essentially as good as random, arguing that CZ that are on similar trends have had the same demand shocks but different supply side shocks.

Which CZs import robots?

Previous literature studying automation at the firm level finds that importer firms are positively selected vis-à-vis nonimporters (Acemoglu et al., 2022; Bessen et al., 2020; Koch et al., 2021). Similarly, CZs that receive robot imports differ from those that do not. Table 1.1 shows that CZs receiving robots tend to be larger in size (population and working-age individuals) and have a higher number of formal workers who earn higher average wages. They also tend to have a larger share of their population living in urban areas and working in manufacturing. Figure A.1 shows an increasing trends in formal employment and wages based on the IMSS data, which is common to every wage group and more pronounced in CZs receiving robot imports. This is a characteristic that is correlated with economic and manufacturing growth, and motivates the sample of CZs used in the empirical analysis.

Panel A of Figure 1.3 presents the estimated coefficients from a linear probability model where the dependent variable is a binary indicator equal to one if a CZ is importing robots and where I simultaneously include many different pre-2006 CZ characteristics. It shows many statistically significant differences between CZs receiving and not receiving robots. Various factors positively affect the predicted probability of importation: the working-age population, the share of manufacturing employment (level in 2000 and change in the 1990s), the average formal wage in 2005, and exposure to EU robots in the 1990s. Given these differences between importers and nonimporters, using an empirical strategy that uses all nonimporters as a comparison group would subject my estimates to significant biases: it would be difficult to argue that these very different CZs are on similar trends despite the considerable differences in their economic development.

What can explain the timing of robot imports?

An alternative approach to studying automation's effects involves keeping the CZs that receive at least one industrial robot importation during the analysis period and exploiting the variation in the import timing. Panel B of Figure 1.3 presents the estimated coefficients from an ordinary least square (OLS) regression where the dependent variable is the first year in which I observe a CZ receiving a robot importation. In this case, no variable except urban population share has a coefficient statistically different from zero. Despite the potential influence of unobservable factors on the timing of the first import, I find no significant and systematic relationship between import timing and relevant observable variables. Additionally, I estimate a linear probability model using the panel of CZs and sectors used in the primary analysis to investigate whether formal employment and wages affect the predicted probability of starting to import robots. Table A.3 shows that, after including panel unit and year fixed effects, there is no association between the labor market variables and the event's timing. Given this result, I focus the empirical strategy on comparing CZs that receive robot imports with the ones that have yet to start importing.

4. Empirical Strategy: Event Studies

Identifying the causal effect of robot adoption on jobs and wages presents critical challenges. Given that importing and adopting robots is a decision made by the firms themselves, there is an endogeneity puzzle that necessitates a careful analysis.

On the one hand, imagine a situation where the labor supply diminishes for specific tasks, such as manual and routine ones. For instance, if a company that manufactures electronic gadgets faces a shortage of skilled manual workers who solder electronic components onto circuit boards, attach wiring, and perform quality checks, they may adopt industrial robots to handle these tasks previously performed by humans. In this scenario, it may appear that robots lead to job reductions, creating a negative bias.

On the other hand, firms could experience positive shocks in product demand. Consequently, they might adopt more robots while increasing the size of their workforce. Similarly, positive financial developments within firms might increase investments in robots and job creation. In such cases, we might erroneously deduce that robots positively affect employment when, in fact, they do not.

The approach that I employ to address this challenge involves leveraging geographical and temporal variations in robot importation while controlling for diverse unobservable factors and shocks at the CZ, 4-digit sector, and year levels. To enhance robustness, I use multiple comparison groups and various methodologies capable of addressing different potential sources of endogeneity.

The main set of strategies that I use is timing-based event study methods. I define an event as the first year in which I observe a CZ receiving industrial robot imports. Then, I can group the CZs into different cohorts by the year of this event. The strategy is to compare the labor market outcomes of CZs receiving industrial robot imports before and after the first importation, using as a comparison group the CZs that have not yet received robot imports but eventually do during the period of analysis. Following Miller (2023), I define the data structure as *timing-based* because I use only ever-treated units and the event dates vary. In this case, the underlying thought experiment is that the event timing is as good as random, implying the traditional assumptions of parallel trends and no anticipation. In this case, I rely on the fact that I find no systematic association between the timing of the importation event and economically relevant observable CZ factors.

Table 1.2 provides descriptive statistics on the distribution of yearly imports by CZ cohort.[16] There are 104 CZs that receive robots during 2006–2022. Most were already

[16]See Figure A.5 in the appendix, which shows a histogram of the events. In addition, Figure A.6 shows the raw means of the number of formal workers relative to the number one year before the event for different cohorts and CZs.

importing robots in 2006, so I exclude these from the sample. I also exclude from the treated sample the observations corresponding to periods after 2019, which are likely affected by the COVID-19 pandemic. After these sample restrictions, I end up with 71 CZs grouped into fourteen cohorts that, on average, imported four times during the analysis period.[17]

Two-way fixed effects model: Event-study specification. I first estimate a traditional *two-way fixed effect (TWFE) model*. I compare the labor market outcomes of CZs receiving robot imports before and after the event, using the CZs that receive imports earlier or later as the control group. The unit of analysis is a 4-digit sector within a CZ. Therefore, the estimated model is:

$$Y_{smt} = \sum_{\tau=-10,\tau\neq-1}^{10} \beta_\tau d^\tau_{mt} + \gamma_{ms} + \delta_{str(m)} + \varepsilon_{mst}, \tag{1.1}$$

where Y_{mt} is the labor market outcome of industry s in CZ m in year t, d^τ_{mt} is dummies indicating the year since the CZ's first robot importation ($\tau = 0$), γ_{ms} is CZ–industry fixed effects, and $\delta_{str(m)}$ is industry–year–region fixed effects. The fixed effects are essential since they capture part of the unobservable factors causing endogeneity. For example, if possible positive demand shocks that importer firms experience happen at the 4-digit industry level, I net out that factor from the analysis.

The β_τ parameters are identified by comparing differences in Y before and after event τ for treated and not-yet-treated or already-treated units.[18] The identification assumptions underlying this approach are as follows: (i) that there is no systematic difference in how the outcomes evolve between the two groups, except the effect of the treatment (parallel trends assumption), (ii) that units do not anticipate or react to their future treatment

[17] Note that this does not mean that on average they received four imports (i.e., a firm imported four times). It means that I observe that a CZ received robot imports in four years, regardless of whether there were many or there was just one importation within a year.

[18] To ensure identification, at least two parameter restrictions are implemented to prevent perfect multicollinearity (Sun and Abraham, 2021; Borusyak et al., 2021; Miller, 2023). These restrictions involve omitting the event dummy -1 and including end-cap dummy variables at the extremes, such as $d^{\tau\leq 10}mt$ and $d^{\tau\geq 10}mt$.

status before the event, and (iii) that the effects are homogeneous across cohort groups. This last assumption is particularly important for settings where the analyzed treatment is staggered over time, as it prevents biases arising from comparing later-treated units with earlier-treated ones. Additionally, I cluster the standard errors at the CZ level, representing the most conservative option and corresponding to the level at which I observe robot imports.

New event-study methods. TWFE models can produce biased estimates when there are heterogeneous effects across cohorts under a staggered treatment timing, which is the setting of this study. Therefore, I implement two alternative methods—those of Callaway and Sant'Anna (2021) (CS) and Sun and Abraham (2021) (SA)—that circumvent the issue of negative weighting and the so-called bad comparisons of cohorts. Both methods consist of choosing comparison groups that are not contaminated by the treatment.

The CS method computes multiple 2-by-2 difference-in-differences estimations, avoiding bad comparisons (i.e., not using early-treated observations as a comparison group), and then it aggregates them over event time. The building block of the estimator is $ATT(g, t)$, the average treatment effect on the treated at time t for units starting treatment at time g (cohort). Let us define $D_t = 1$ if it is treated by time t and $D_t = 0$ otherwise (not yet treated). When I use the units not yet treated by time t as the comparison group, the estimand is as follows:

$$
\begin{aligned}
\text{If } t \geq g: \qquad & ATT(g,t) = \mathbb{E}[Y_t - Y_{g-1} | G = g] - \mathbb{E}[Y_t - Y_{g-1} | D_t = 0, G \neq g] \\
\text{If } t < g: \qquad & ATT(g,t) = \mathbb{E}[Y_t - Y_{t-1} | G = g] - \mathbb{E}[Y_t - Y_{t-1} | D_t = 0, G \neq g]
\end{aligned}
$$

For post-treatment periods, the method calculates the difference between the outcomes of treated CZs in cohort g and those of not-yet-treated CZs for the period t years after the event compared to the outcomes in the period immediately before the event (period $g - 1$). This procedure implicitly controls for linear trends and accounts for unit and period fixed effects. For pretreatment periods, the method calculates the difference between the outcomes of treated CZs in cohort g and those of not-yet-treated CZs for the period t years compared to those in the period $t-1$ years before the event.

Then, I can aggregate the multiple $ATT(g,t)$ across event-time to obtain the average effect for units that have been exposed to treatment for exactly e time periods ($e = t - g$):

$$\begin{aligned}\theta(e) &= \sum_{g}^{T} \omega_e(g,t) ATT(g, g+e) \\ &= \sum_{g}^{T} \mathbb{1}\{g + e \leq T\} P(G = g | G + e \leq T) ATT(g, g+e) \end{aligned} \tag{1.2}$$

The main difference between the CS and SA methods lies in the comparison group chosen in each. Both methods allow use of the not-yet-treated cohorts as the comparison group. In this context, Sun and Abraham (2021) use the last-treated cohort (cohorts that started importing robots in 2019 or later), while Callaway and Sant'Anna (2021) use all the not-yet-treated cohorts.

Synthetic DD. Synthetic difference-in-differences (SDD) is a hybrid method that combines elements from both the difference-in-differences (DD) and synthetic control (SC) approaches.[19] In an SDD analysis, we use a panel setup, balanced in calendar time, where specific units receive treatment while others remain untreated. The goal is to assess the impact of the treatment by comparing changes in outcomes before and after treatment for both the treated units and a synthetic control group that has the same pretrends. To create the synthetic control group, the method selects units from the never-treated group (CZs that have never received robot imports during the analysis period) and time periods before the event. It assigns unit-specific weights to these untreated units and time-specific weights to the pretreatment periods. SDD capitalizes on the strengths of both the DD and SC methods.

[19] The method described here follows the framework proposed by Arkhangelsky et al. (2021). Refer to their article for detailed information on estimation, inference procedures, and formal proofs of the estimator's consistency and asymptotic normality.

Similar to DD models, SDD allows for pretreatment differences in level outcomes between the treated and control units. Conversely, like the SC method, SDD seeks to create a control unit that mitigates the strict requirement of the parallel trends assumption.

As a result, SDD is advantageous because it can overcome common challenges encountered in implementing standard DD and SC methods. It can estimate causal relationships even when pretrends are not strictly parallel in aggregate data, which would violate a key identifying assumption for DD. Furthermore, it does not necessitate that the treated unit fall within the "convex hull" of control units, a condition imposed by the SC method.

Robustness analysis: Alternative methods. I employ a triple-differences (DDD) approach to complement my previous analysis. This method involves comparing labor market outcomes among industries that are more and less exposed to robot adoption within CZs that either receive or do not receive robot imports before and after the first robot importation event. I define more exposed industries as manufacturing industries observed to have a high rate of robot adoption in the IFR data, and low exposed as the rest of non-manufacturing industries.[20] The critical innovation of the DDD approach lies in its introducing a third difference by considering different industry exposure levels. This additional layer of differentiation allows me to effectively control for idiosyncratic labor market shocks that might otherwise confound the analysis. These shocks are a primary concern in conventional DD approaches.

I incorporate CZ–year fixed effects into the model to implement this DDD approach. These fixed effects help account for various shocks affecting a CZ labor market in a given year, such as economic downturns, local policy changes, or large-scale industry expansions. By including these fixed effects, I can isolate the impact of robot adoption within industries

[20]More exposed industries include Transportation, Electronics, Metal products, Chemicals, Plastics, Rubber Energy and power generation, Food, and Beverages. Less exposed industries include all the non-manufacturing industries, excluding wholesale machinery trade. I exclude from the sample the rest of the manufacturing industries, i.e., Textile, Wood, Furniture, Paper, Printing, and Non-metallic mineral products.

while minimizing the influence of external shocks or trends that might affect the outcomes. This triple-difference framework enhances the robustness of my analysis and strengthens the causal inference regarding the effects of robot adoption on labor markets.

Finally, I replicate the Bartik instrument strategy outlined in Acemoglu and Restrepo (2020). This strategy involves using an instrumental variable based on exposure to European robot adoption to instrument for exposure to domestic robot adoption. This approach provides an additional layer of validation for my analysis, as it helps isolate plausible exogenous variation in robot adoption. It assumes that European trends reflect reductions in robot costs and technological advancements and do not directly impact Mexican labor markets. Replicating this strategy enables me to compare my findings with results in the existing literature and underscores the significance of my utilizing new industrial robot data. I can improve the measurement of robot adoption and exploit variation within industries. This is crucial as it can be challenging to separate the effects of robots from industry-level shocks, particularly in this context where the automotive industry accounts for most of the phenomenon.

5. Results

5.1 Formal labor market

I begin by examining the reduced-form effects of starting to import industrial robots on CZs' formal employment and average wages for the set of CZs that eventually receive robot imports. In Figure 1.4 Panel (a), I visually represent the dynamic effects on the logarithmic count of formal workers, defined as individuals affiliated with IMSS in December of each year. In this figure, each dot corresponds to the estimated $\theta(e)$ event-study coefficients, calculated using equation 1.2, and is accompanied by a 95% confidence interval. Since outcome variables are measured in logs, the point estimates can be interpreted as semi-elasticities.

Consistent with my failing to find observable predictors of the timing of the event for the subset of not-yet-treated CZs, the pre-event coefficients consistently hover close to zero and remain statistically insignificant at the 5% significance level (the uniform confidence interval covers zero for all pre-event years). The average of the coefficients of the periods between one and five years prior to $t = 0$ is very close to zero (0.005 log points, SE = 0.004), suggesting that employment trended similarly across all CZs before the first robot importation. This finding is reassuring, as it is consistent with the validity of the parallel trends and no-anticipation assumptions.

The number of formal workers in CZs remains relatively stable during the first four years after the initial robot importation, showing an average positive effect of 1.6%, although this effect is not statistically significant at the 5% level. This result suggests that I can rule out significant short-term negative employment effects exceeding -0.7%. Subsequently, the estimated coefficients are still nonnegative. In fact, I find an increase in formal employment averaging 6.9% for five to eight years after the event, a result that is statistically significant at 1%. However, this effect might be explained by compositional changes across CZ cohorts.

The concern about compositional changes arises due to the varying duration of exposure of different cohorts. For instance, we can determine the instantaneous average effect ($t = 0$) only for the cohort of CZs that started importing robots in 2019. Conversely, for CZs that began importing in 2017, we can assess the effects at event times $t = 0, 1, 2$. While both cohorts contribute when I compute the coefficient for $t = 0$, this is not true for $t > 0$. This situation can introduce confounding dynamics and selective treatment timing among different cohorts if the impacts of robot importation systematically differ.

To mitigate this concern, I balance the sample by including only cohorts exposed to the treatment for a certain minimum number of years and examine dynamic effects within those years. Table 1.3 presents the results for different balanced samples, ranging from event windows of one to four years around the importation event. Encouragingly, the short-run effects remain nonnegative, allowing me to rule out negative effects within the range of [.4%–3.2%] for formal employment and [.2%–.8%] for average wages, depending on the length of exposure.

On the other hand, it is more difficult to rule out composition effects as an explanation for the positive long-run effects I document. The estimates rely on fewer and fewer units after 4 years have passed. One could in principle restrict the sample to include only units that are observed over a long period of time to estimate these long-run effects. However, there are only a handful of units that are observed for more than 4 years (for example, there are only 2 with a six-year post-period). Thus, the best estimates for these long-run effects come from the synthetic control approach. These estimates show that the effects are, in fact, the largest and most positive for these early adopters (see Figure A.13).

This positive employment impact suggests that some of the productivity effects take time to emerge due to learning. The growth might also reflect that the intensity of robot use is increasing over time (for example, CZs that first adopted robots in 2011 continued to adopt more robots in subsequent years). The estimates so far only consider the impact of the first set of robots that are adopted, but in fact automation continues once the CZ starts importing robots. I investigate the effect of the intensity in the next section, and I find that the effects are more positive when the intensity is larger.

Panel (b) of Figure 1.4 plots the dynamic effects of robot adoption on the logarithmic average real monthly wage for formal workers using the same empirical approach. The patterns mirror those in Panel (a), but the effects seem less pronounced. The real monthly wage does not exhibit differential pretrends before the importation event. Then, there is no statistically significant wage increase during the first four years after the first robot importation, and there is a positive effect after $t = 5$ of 3% for all industries, which is statistically significant at 5%. If we restrict the sample to robot adopter industries within manufacturing, the effect is still nonnegative (see Table 1.4).

Another way of understanding whether this null effect is hiding potential effect heterogeneity is to analyze groups that are more exposed, such as manufacturing sector workers. To explore this hypothesis, Table 1.4 summarizes the effects for various subsamples. Indeed, when I focus the analysis on the manufacturing sector, I find that the short-term effect is still small and nonnegative but the long-term effect is larger than in the baseline, although more imprecisely estimated and potentially driven by compositional changes. The same is true for other subsamples such as workers in the manufacturing industries that have a large stock of robots according to the IFR data or the sample restricted to a balanced panel in calendar time. A consistent pattern emerges across all these different subgroups, revealing a nonnegative employment effect associated with robot adoption.[21]

Additionally, I explore whether the coefficients are more prominent for groups for which we would anticipate a significant negative effect based on the previous literature. If other groups had notably larger results, this would suggest the presence of omitted variable bias. Table 1.5 provides the CS model estimates, segmented by various factors such as age, wage, and firm size. This analysis confirms that all groups have the same nonnegative result in the short run. Notably, the long-term employment effects are concentrated among medium-wage and prime-age workers, both groups where industrial operators are more prevalent (see Figure A.7 for the dynamic effects). Moreover, it becomes apparent that the results are most pronounced in larger firms with over 50 employees, where automation is more likely to occur. Panel D reveals that the employment and wage effects tend to be relatively larger in smaller CZs, which are typically more responsive to labor market shocks (see Figure A.8 for further details).

The main results are consistent with Ing and Zhang's (2022) finding for Indonesia of positive employment effects for the manufacturing sector (0.113, SE = 0.038) and wages

[21] Note that I also consider the inclusion of never-treated CZs in the comparison group. The post-treatment effects are similar to the ones considering only the not-yet-treated sample. However, it exhibits a slightly but statistically significantly differential pretrend.

(0.068, SE = 0.040). These results are similar to this paper's findings, even though these authors consider a more general definition of automation equipment (not only industrial robots). Calì and Presidente's (2021b) employment effect for Indonesia is of the same sign but is larger in magnitude at 0.307 (SE = 0.104). Rodrigo (2021) also does not find a large statistically significant decline on employment for Brazil, and also finds positive effects on wages and productivity. These results suggest that in industrializing economies the effects are positive unlike those found for the US and Europe.

The estimates in this paper differ from those in Brambilla et al. (2023), which finds a negative effect of robot exposure on the number of formal workers in a panel of states in Brazil, Mexico, and Argentina in the periods 2004–2016. A potential source of this difference is that the study uses a shift-share instrument strategy, leveraging industry composition variation across less granular geographic units (states), and exploiting a different source of variation (only within geographic units).

5.2 Robustness checks

Alternative outcome measures and financial crisis

Another concern is related to the measurement of outcomes, which can have measurement errors that would decrease the precision of the estimates since I take the information from December of each year. To explore how sensitive the main results are to this issue, I calculate the labor market outcomes as a simple yearly average of the last month of each quarter (comprising March, June, October, and December). These measures may be a better indicator of what is happening in each year to the extent that they use more available information. Table A.4 repeats the main table but changes how I measure the dependent variable. These results closely mirror those derived from the December data exclusively.

A second concern is related to the fact that the observed effects could be driven by firm responses to the 2009 financial crisis. To rule out this potential explanation, I re-estimate the effects excluding from the sample the cohorts that started importing robots in 2009 and 2010. Table A.5 shows that the results are very similar when I exclude these cohorts.

Alternative methods

Figure A.10 shows the dynamic coefficients of the three event-study models—TWFE, SA, and CS—for the sample of not-yet-but-eventually-treated CZs. In the case of employment, there are no significantly differential pretrends, and the postevent effect is even more pronounced than that estimated under the CS method. In the short term, the number of formal workers increases by between 3.8% and 4.1%, while in the long term, it rises to 10.7% on average. On the other hand, when I focus on average formal wages, the pre-event coefficients for the different approaches also lack statistical significance, and the effects follow the same dynamic pattern as those on employment, with more muted and less precise positive results (I can rule out negative wage effects larger than -0.8%).

While the pre-event coefficients derived from the various approaches lack statistical significance, the possibility of a modest pretrend increase in average wages emerges, particularly when I examine a timeframe extending beyond three years before the event. This possibility merits thorough investigation. To assess the significance of these pretrends, I employ a sensitivity analysis method proposed by Rambachan and Roth (2023). This analysis estimates the post-treatment effect under different assumptions of parallel trend violations, demonstrating the extent to which postevent trend differences would need to deviate from the pretrends to render the estimates negative or statistically insignificant. More precisely, I provide confidence intervals that enable the post-treatment deviation from parallel trends to reach a magnitude of up to M times the maximum pretreatment deviation, with M being a variable that can take different values. Figure A.11 illustrates the results of this analysis, suggesting that if we assume that the maximum size of the pretrend persists post-treatment, the short-run and long-run estimates for employment and wages remain positive but lack statistical significance (the estimated coefficient for the employment effect in period 6 remains statistically significant when I assume that the pretrend persists but at half its size). The crucial point is that a substantial deviation of the parallel trend assumption would be necessary to yield negative employment and wage estimates.

Another way of addressing the potential violation of parallel pretrends is by controlling for CZ–time-specific idiosyncratic shocks using a DDD model. To do so, I compare the labor market outcomes of more and less exposed industries within CZs before and after the import event between treated and never-treated CZs. I define more exposed industries as those in the manufacturing sector that I observed in the IFR data adopting robots. Less exposed industries are the other, nonmanufacturing industries.[22] I further restrict the sample to CZs with a population of fewer than 500 thousand to improve comparability. Figure 1.5 shows that, once I control for possible deviations in the employment and wage pretrends using the less exposed group, I still find nonnegative effects on formal employment and wages of similar magnitudes to those in the previous analysis.[23]

SDD is an alternative method that directly addresses pretrend differences by construction. When I use this method, I find no statistically significant negative employment or wage effect for any industries or manufacturing sectors, as seen in Figure 1.6. Most of the employment and wage point estimates are positive but generally have larger confidence intervals than those for the results under the previous methods. For instance, the simple average of the employment effects across CZs is 0.07 (SD=0.30), while the average impact for wages is 0.01 (SD=0.07). Figure A.13 shows the difference-in-difference graphs for a specific CZ, where there is a statistically significant positive effect in the number of workers and average wage.

To conclude the analysis, I present the results from the two-stage least squares (2SLS)

[22] I also exclude from the less exposed industries those in educational service since it is listed as adopting robots in the IFR data but is outside manufacturing. I further exclude the industry of wholesale trade of machinery, which likely includes robot integrators.

[23] Figure A.12 shows that this conclusion holds if I use 3-digit industry codes.

method with the Bartik instrument of EU automation, using as a measure of domestic robot adoption the change in robot exposure between 2011 and 2015. Figure A.14 plots the estimated coefficients and compares them with the results from Acemoglu and Restrepo (2020). Table 1.8 presents the 2SLS results for multiple outcomes.[24] The findings do not show any clear proof of negative effects on jobs or wages, although the standard errors are larger and not the same as those corresponding to the results under previous DD based methods. The reason might be that the Bartik method estimates the effects for a different group of units (those affected by the instrument) and relies on the assumption that if an industry adopts robots, any CZ with firms in that industry adopts robots, even if this is not the case.

Table A.7 summarizes the results of all the different methods showing the average estimated coefficients during the pre-event period (averaged across four periods) and the postevent phase, further differentiating the short-run (within four years of the event) from the long-run effects (five to eight years after the event). The main finding across specifications is that there are no substantial negative employment effects of automation. After the first robot importation, there are positive but relatively small effects on the number of formal workers in the short run and larger effects in the long run. This might be consistent with the fact that integrating robots and changes in the job and task landscape necessitates an adjustment period. Using event studies based on traditional TWFE, either using double or triple differences, I can rule out negative effects larger than -0.4% to -0.8%. Overall, it is encouraging that these alternative approaches support the baseline results.

Other robot adoption measures

Concerns may arise with respect to the relevance of defining the event of interest as a CZ's first year of industrial robot importation. It is essential to consider whether different effects

[24]Table A.6 presents similar results for the change between 2010 and 2015. This version implies imputing robot adoption for the year 2010, given that the IFR data for Mexico are available for years since 2011.

emerge when alternative measures that account for import intensity are considered. To address this, I develop two alternative measures for analysis.

The first one considers the value of the robot imports and defines the event as the year in which the cumulative value of a CZ's robot imports as a share of its total 2005 manufacturing sector payroll surpasses a given threshold. To obtain this measure, I first need to impute the import value for CZs with missing data because, for confidentiality reasons, this information is observable only when more than three firms in the CZ import robots in a given period. Therefore, I impute the import values using the mean of each year–state–country supplier cell. Panel (a) of Figure A.9 in the appendix shows that the imputed measure has a mean and variance similar to those of the nonimputed measure. Panel (b) compares the distribution of CZs across the time passed since the event as originally defined and the event as defined by means of the cumulative import values with a 50% threshold. Under the new definition, there are fewer CZs treated for fewer periods. Figure A.15 shows the CS estimates for events based on the new definition under increasing thresholds. All the estimated coefficients for formal employment and wages are positive, regardless of the threshold. In fact, the average of the long-term coefficients is statistically significant at 1% when I use thresholds larger than 40% (i.e., the cumulative value of robot imports is larger than 40% of the total manufacturing sector wage bill). Table 1.6 summarizes these results.

A second possibility is to use the cumulative number of yearly imports at the CZ level. Since the Callaway and Sant'Anna (2021) method does not allow me to consider the treatment intensity, I implement this measure using the TWFE model, interacting the event-time dummies with the cumulative number of yearly imports. Table 1.6 shows that the employment and wage effects are similar to the baseline estimates described in the previous section.

5.3 Other margins of adjustment

To evaluate whether firms and workers use alternative margins of adjustment, I conduct an analysis of its effects on labor informality, unemployment, and labor force participation using ENOE data. Table 1.7 presents the results of our primary specification, using formal and informal employment measures obtained from ENOE as dependent variables. These results are compared with point estimates derived from the IMSS sample, restricting the sample to CZs for which the ENOE data are representative. It is worth noting that ENOE is representative for only a limited number of major cities. Consequently, the estimates regarding the number of formal workers, when compared to those calculated from the IMSS data, exhibit differences in precision, with the latter being less statistically robust but qualitatively consistent (showing a positive short-term effect and a more pronounced long-term effect).

Notably, our findings reveal a short-term decrease in the number of workers in the informal sector, followed by a long-term increase, potentially attributable to general-equilibrium effects. A similar pattern emerges in the case of unemployment and labor force participation. Intriguingly, there is a concurrent rise in the urban population, suggesting that there might be rural-to-urban migration.

Ruling out potential explanations: Plant openings

In dissecting the underlying factors behind the observed positive employment effects beyond the anticipated gains in productivity from robot adoption, I delve into the prospect that firms, buoyed by robust demand shocks, opt for expansion and open new facility establishments, potentially internationally. This sequence could subsequently trigger robot incorporation and a surge in the workforce.

To assess this, I analyze the sensitivity of the estimated robot–employment relationship to my factoring in a measure of plant openings. Leveraging the 2018 economic census, I utilize the variable tracking the commencement year of operations. This enables the calculation of establishment openings for surviving firms. I focus this analysis on specific industries and firm sizes, particularly on industries substantially exposed to automation and larger firms.

Importantly, these segments are more likely to have persisted until 2018, rendering the measure more reliable. The core point estimates remain largely unaltered upon my including this measure into the CS model (see Figure A.16). Furthermore, I extend this exploration to directly evaluate the impact of robots on plant openings, finding no discernible association (see Figure 1.7).

In essence, this comprehensive exploration effectively narrows down the plausible drivers of potential positive employment effects, bolstering the conjecture that heightened productivity triggered by robot adoption was a significant catalyst. This fact lends further weight to the argument that the positive employment outcomes observed are rooted primarily in productivity enhancements brought about by adopting robots.

6. Mechanisms and Alternative Explanations

To analyze the effects of industrial robot adoption on productivity, output, the labor share, and capital investment, I use the Mexican economic censuses (2003, 2008, 2013, 2018). These censuses are taken every five years, which limits my ability to explore the dynamic effects of robot adoption. For this analysis, I estimate the following TWFE model:

$$Y_{smt} = \beta D_{mt}^{post} + \gamma_{sm} + \delta_{st} + \epsilon_{smt}, \tag{1.3}$$

where Y_{smt} represents the aggregated outcomes of firms at the six-digit sector s in CZ m in census year t and D_{mt}^{post} is an indicator variable that takes the value of one if CZ m starts or has already started to import robots in census year t. This model allows me to control for idiosyncratic shocks at a much more granular (six-digit) sector level than in the analysis with the IMSS data, where I can control for shocks only at the four-digit sector level.[25] γ_{sm} is CZ $\times$ six-digit sector fixed effects, which capture non–time-varying factors, such as product or labor market conditions. δ_{st} is six-digit sector $\times$ census year fixed effects, which capture the common trends across sectors over time. For example, some sectors may have regional policies or may experience demand shocks due to global competition or trade liberalization that affect all firms in the sector over time. By including these fixed effects, I can control for important confounding factors that may affect the relationship between robot adoption and the aggregated firm outcomes.

Table 1.9 presents the results and is divided into two panels: Panel (a) shows the results for all industries, and Panel (b) shows the results for a subset of manufacturing sectors that are more likely to be affected by automation (metal, machinery and equipment, plastic and rubber, power generation, and transportation).[26] The dependent variables are the number of firms, labor share, capital per worker, ln gross value added (GVA), GVA per worker, input cost, total payroll, total fixed assets, number of workers, number of workers in production and sales, number of workers in administration, hours worked, hours worked in production and sales, and hours worked in administration.

For all industries, industrial robot adoption has a positive and significant effect on GVA and the input cost. This suggests that robot adoption may increase the productivity of firms in Mexico. However, robot adoption has no significant effect on the labor share, capital per worker, payroll, fixed assets, or employment. This suggests that robot adoption does not affect the distribution of income between labor and capital, or firm size. However, this could be result of the lack of power given the level of aggregation and the availability of observations

[25] One example of a six-digit sector within the automotive manufacturing sector is 336111, automobile manufacturing. This sector comprises establishments primarily engaged in one or more of the following: (1) manufacturing complete automobiles (i.e., body and chassis or unibody); (2) producing automobile chassis only; (3) manufacturing automobile bodies only; or (4) manufacturing and/or assembling automobile subassemblies. This sector is part of 3361, motor vehicle manufacturing, which is part of 336, transportation equipment manufacturing.

[26] Tables A.8 and A.9 present the same results for the balanced panel in calendar time. The findings are similar to those from the unbalanced panel.

for only few periods. The sign of the labor share is negative, and the magnitude is -1.1%, consistent with results derived from task-based frameworks for considering automation effects (Acemoglu and Restrepo, 2020).

For the subset of manufacturing sectors that are more likely to be affected by automation, industrial robot adoption has a positive and significant effect on GVA per worker, input cost, payroll, fixed assets, and capital per worker. This suggests that robot adoption increases the productivity, efficiency, wages, and investment of firms in these sectors. However, robot adoption has no significant effect on the labor share (although the coefficient is negative and larger in magnitude than that for all industries). This suggests that robot adoption does not affect the distribution of income between labor and capital.

I find that robot adoption has a nonnegative effect not only on employment but also on hours worked. Across all industries, the effect is generally small and lacks statistical significance. However, when I narrow my focus to the manufacturing sector, the effects become more substantial and statistically significant within the balanced sample: employment shows an increase of 18.1%, while hours worked increase by 19.4%. These effects can be attributed to increased production and sales workers rather than administrative workers.[27].

For the other manufacturing sectors, the effects are either not significant or much smaller than the effects found for the manufacturing sectors that are more likely to be affected by automation. However, Panel (a) of Table 1.10 presents some exceptions, showing evidence for sectors such the wood, paper and chemicals industries of an increase in capital per worker (3.2%) and a reduction in employment of production and administrative workers. This finding could suggest a certain degree of reallocation of workers across sectors, depending on how much they automate and how large the productivity channel is. I find no effects for

[27]The differences in the effects on production and administrative workers have been investigated by Dixon et al. (2019). They used firm-level data from Canada and found that adopting robots significantly reduces the number of managers and administrative staff, and increases productivity and wages. According to their hypothesis, robots enable companies to enhance their product and service quality, which, in turn, requires more flexible and adaptable management practices.

manufacturing industries with very low levels of robot adoption, such as the textile sector (see Panel (b)).

For the nonmanufacturing sectors (Panel (c) in Table 1.10), industrial robot adoption has a positive and significant effect on GVA, input cost, number of workers, and hours worked. However, these effects are much smaller than those for the manufacturing sectors more likely to be affected by automation. These outcomes may indicate the presence of broader, general-equilibrium effects stemming from robot adoption. Such effects can include shifts in demand or supply dynamics across various sectors, including construction, wholesale and retail trade, and manufacturing, ultimately highlighting the interconnected nature of these industries.

These findings align with the results in the literature on industrial robot adoption at the firm level, highlighting its positive impact on productivity and employment, as evidenced by significant increases in GVA per worker and number of workers (for example, Acemoglu et al. (2023) for the Netherlands, Acemoglu et al. (2022) for the US, Acemoglu et al. (2021) for France, Koch et al. (2021) for Spain, Dixon et al. (2019) for Canada, and Humlum (2020) for Denmark). These aggregate estimates might hide effects within industry since robot adopters may expand at the expense of nonadopter competitors (Acemoglu et al., 2023). However, I still find zero or positive effects on aggregate employment.

7. Discussion

7.1 Theoretical framework

To understand how industrial robot adoption may impact labor markets differently in developing countries, I summarize the key findings of a simple, competitive labor markets model in the style of Acemoglu and Restrepo (2018) where automation allows robots to substitute for workers in some tasks. Acemoglu and Restrepo (2018) develop a task-based model in which robots and human labor compete to complete various tasks. In this model, the effects of robots on wages and employment are ambiguous.

Consider a perfectly competitive economy with two sectors j: (i) manufacturing ($j = 1$) and (ii) the rest of the sectors ($j = 2$), potentially including the informal sector in the Mexican context. Production at $j = 2$ uses only workers, such that one unit of labor corresponds to one unit of output.

The representative firm in sector 1 produces a final good Q by combining a continuum unit measure of tasks, which are completed by robots and human workers. I assume that the markets for both inputs are perfectly competitive. While workers have identical productivity in all tasks, robots are as productive as workers on tasks in the range $[0, I]$ but have zero productivity at tasks $(I, 1]$. Robots can be purchased in an international competitive market at constant rental rate r. I am primarily interested in the equilibrium effects of I, that is, increased automation that allows robots to perform tasks that previously required labor.

Good 2 can be produced internally, as noted above, or imported at constant price p_2. I set this good as the numeraire in the economy: $p_2 \equiv 1$.

The manufacturing good is exported, facing isoelastic international demand:

$$Q(P_1) = \kappa P_1^{-\sigma},$$

where the constant κ is exogenous.

Turning to individual preferences, I assume that all individuals i consume only good 2, and the amount they consume is equal to the wage they receive, as the price is normalized to one. Their indirect utility is a function of their choice of employment sector ($j = 1$ or 2) and the corresponding wage, and is given by:

$$U_{ij} = \beta ln(w_j) + \epsilon_{ij},$$

where ϵ_{ij} is i.i.d. extreme value type 1 errors, indicating that workers may have a preference for one sector over the other. I denote the equilibrium wage in each sector by w_j and the cost of robots by r. Wages can be different across sectors because workers have idiosyncratic preferences for where they work. Given the assumption that production of good 2 requires one unit of labor, and that the market for good 2 is perfect competitive, it follows that $w_2 = 1$.

I assume that the cost of robots is lower than the cost of labor $r < w_1$, such that all tasks in $[0, I]$ are performed by robots while tasks in $(I, 1]$ are completed by human workers, where I represents the extent of automation. Therefore, the cost per unit is:

$$Ir + (1 - I)w_1 = P_1.$$

We obtain the labor demand by plugging the expression for P_1 into the aggregate demand for good 1:

$$Q(P_1) = \kappa[Ir + (1 - I)w_1]^{-\sigma}$$

$$LD_1 = (1 - I)\kappa[Ir + (1 - I)w_1]^{-\sigma}.$$

Each worker with wage w_j and extreme value type 1 heterogeneity ϵ_{ij} evaluates the utility for each option and chooses the highest utility. Consequently, the aggregate labor supply of the manufacturing sector is as follows:

$$LS_1(w_1) = \frac{w_1^\beta}{w_1^\beta + w_2^\beta} = \frac{w_1^\beta}{w_1^\beta + 1} = \frac{1}{1 + w_1^{-\beta}},$$

where β corresponds to the elasticity of supply.

Equilibrium in the labor market is given by $LS_1(w_1^*) = LD_1(w_1^*)$, such that:

$$\frac{1}{1 + w_1^{-\beta}} = (1 - I)\kappa[Ir + (1 - I)w_1]^{-\sigma}.$$

In equilibrium, the elasticity of labor supply β affects the slope of the labor supply curve, affecting the size of the employment and wage effect but not its direction. When β is relatively low, i.e., inelastic, the effects of shifts in labor demand on employment are attenuated and wages exacerbated.

I am more interested in the direction of the effect governed by the parameter σ of labor demand, representing the price elasticity of aggregate demand for manufacturing products. Suppose that σ is relatively high. In this case, the labor demand curve shifts to the right, and there is a positive effect on employment and wages. Figure 1.8 graphically shows how the equilibrium depends on the parameters in a standard supply and demand model.

I next analyze how labor demand changes when the extent of automation marginally increases.

Writing the labor demand function in logs:

$$ln(LD_1(w_1)) = ln(1 - I) + ln(\kappa) - \sigma ln[Ir + (1 - I)w_1] \tag{1.4}$$

Differentiating 1.4 with respect to I leads to the following expression;

$$\frac{\partial LD_1(w_1)}{\partial I} = -\frac{1}{1 - I} - \sigma \frac{(r - w_1)}{Ir + (1 - I)w_1}.$$

Given the assumption $r < w_1$, I can rewrite this as:

$$\frac{\partial LD_1(w_1)}{\partial I} = \underbrace{-\frac{1}{1 - I}}_{\text{Displacement effect}} + \underbrace{\sigma \frac{w_1 - r}{Ir + (1 - I)w_1}}_{\text{Productivity effect}}.$$

According to this result, extending automation has an ambiguous effect on labor demand. On the one hand, there is a displacement effect, which is always negative, $-\frac{1}{1-I} < 0$.[28] On

[28] In Acemoglu and Restrepo (2020), the formula for the cost savings from using robots in labor market

the other hand, there is a productivity or scale effect that is positive under the assumption that $r < w_1$. This latter effect crucially depends on the size of the elasticity of labor demand σ. The larger σ is, the larger the productivity effect, which may more than compensate for the negative displacement effect.

Automation drives the substitution of labor with capital when capital can perform specific tasks more efficiently at the margin. The productivity effect rises because of a reduction in production costs, which translates into a reduction in prices for goods, increasing real income for consumers and consumer demand. In turn, this increases labor demand within nonautomated tasks or sectors yet to undergo automation. This is why, if consumers are more price sensitive, there could be a larger increase in demand and the productivity effect might offset any adverse effect.

7.2 Why might the productivity effect be larger in industrializing than in rich countries?

To understand whether the positive productivity effects of automation are more likely to be larger in industrializing economies such as Mexico, it is helpful to analyze where the productivity gains from robot adoption come from. As explained in Section 7.1, the productivity gains from automation do not result from capital and labor becoming more productive in tasks that they were already performing—they are a consequence of firms' ability to use cheaper capital in functions previously performed by human workers. Therefore, the productivity effect that arises is proportional to the cost savings.

However, the productivity effect also depends on the size of the elasticity of labor demand σ. If the elasticity is large enough, with all else held constant, the productivity effects could offset or even more than compensate for the displacement effect and increase labor demand.

c is in terms of relative prices: $\pi_c = 1 - \frac{\gamma_L}{W_c}\frac{R_c^M}{\gamma_M}$, where γ is factor productivity and W and R^M the prices of labor and capital, respectively. The greater the productivity of capital in automated tasks relative to the rental rate of capital, and the smaller the productivity of labor in these tasks relative to wages, the greater are the productivity gains from automation.

Are there reasons why the elasticity of demand for goods in Mexico be different from that in the US?

Automation in Mexico happens mostly in specific manufacturing industries that are generally concentrated on labor-intensive tasks (assembly) and producing tradable goods for export. Firms in these sectors often face foreign competition, leading to price reductions through automation, creating a competitive advantage and boosting output.[29] On the other hand, US firms might face low competition from abroad since they could be more interested in satisfying domestic demand and consumers may prefer US goods. Consumers in the US also have higher earnings, so they might be less price-sensitive than Mexican consumers.

Alternative explanations are related to different type of industrial robots being adopted depending on the automation stage. Acemoglu and Restrepo (2020) maintain that the predominant negative displacement effect in the US aligns with the prevalence of low-productivity automation technology (referred to as *so-so automation*), which is sufficiently productive that firms adopt it in the presence of factors such as tax code distortions but not inherently more cost-effective than the processes that it supplants.[30] Consistent with this idea, Graetz and Michaels (2018) and Calì and Presidente (2021b) suggest that the positive productivity effect could be larger in developing countries because of diminishing marginal gains from automation. This implies that, as robot density increases, the corresponding productivity gains become increasingly smaller. Furthermore, negative displacement effects might be smaller in Mexico than in the US, given that the labor displaced and early-stage automation might have higher productivity.

[29]For example, Harasztosi and Lindner (2019) find that the tradable sector and exporting firms have larger responses to minimum wage increases because they are more exposed to foreign competitors and have a higher price elasticity of demand.

[30]In the context of so-so automation, the technology's productivity only marginally surpasses that of labor ($\frac{\gamma_L}{W_c} \approx \frac{R_c^M}{\gamma_M}$). Nonetheless, in Mexico, the likelihood of $\frac{\gamma_M}{R_c^M} \gg \frac{\gamma_L}{W_c}$ is heightened, particularly considering that so-so automation tends to emerge at later industrialization stages.

8. Conclusion

This study sheds light on the impact of industrial robot adoption on labor markets in developing countries, with Mexico serving as a compelling case study. The findings offer evidence to support an optimistic narrative: the rapid adoption of robots in Mexico over the last decade has not yielded substantial negative effects on formal employment or average real formal wages. Instead, it sets a favorable transitional phase, with far-reaching implications. While the immediate impact may appear modest, there is evidence of a potential positive effect in the long run, underlining the gradual labor market adjustments linked to automation.

These findings suggest possible policy considerations for industrializing countries such as Mexico. Policymakers can explore the potential of promoting automation to increase productivity and growth, especially in the tradable manufacturing sector. However, the equitable distribution of these benefits is a crucial aspect that needs attention. Facilitating technology transfer from automating to non-automating sectors through partnerships and collaborations could be a viable strategy. Policymakers can adopt a forward-thinking approach, anticipating a future where automation plays a significant role in the economy. This approach could involve planning for investments in education and training for jobs in automated industries, developing infrastructure to support automation, and formulating regulations to guide the ethical use of these technologies. Adopting such a long-term perspective could ensure a smooth and beneficial transition to a more automated economy.

In conclusion, this study contributes to the emerging body of evidence on automation's impact on industrializing economies and suggests promising avenues for future research. Further investigations using firm- and establishment-level data can provide a deeper understanding of the mechanisms and dynamics underlying the observed effects. Exploring worker-level effects will enable a more nuanced assessment of reallocation effects and of which workers might be beneficiaries and which might be adversely affected, offering valuable insights for public policy interventions. Comparative studies across regions and countries hold the potential to reveal the nuanced impact of automation across different contexts, facilitating more targeted and effective policy responses. As we navigate the transformative era of automation, continued research in these directions will be essential to inform evidence-based policymaking and ensure inclusive and sustainable economic development.

9. Figures

Figure 1.1: Robot adoption in Mexico

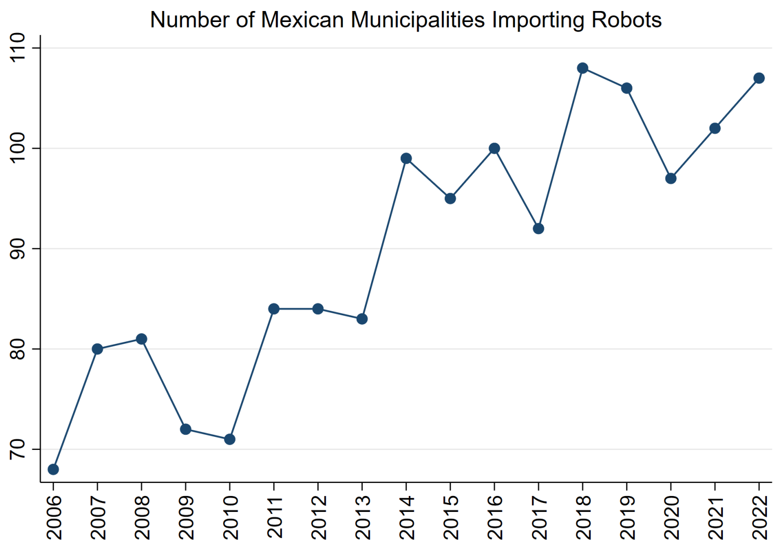

(a) Robot imports from customs records

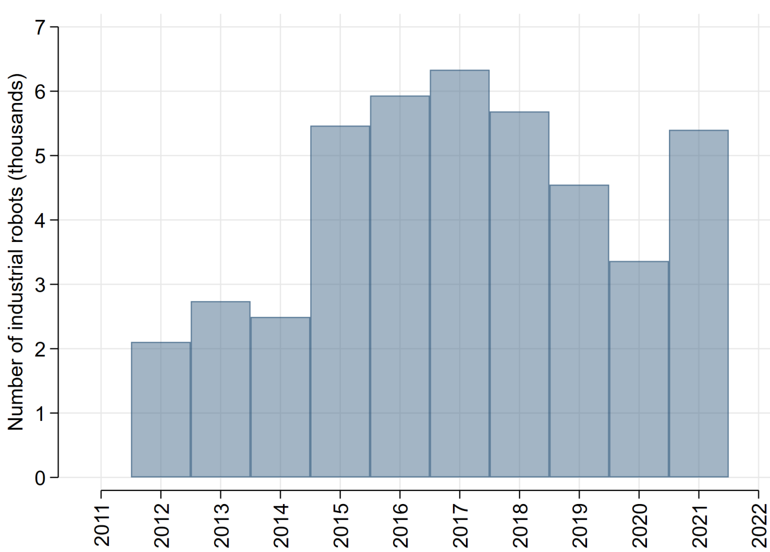

(b) Flow of industrial robots from IFR data

Notes: Panel (a) of Figure 1.1 depicts the flow evolution of the number of Mexican municipalities (similar to US counties) importing robots between 2006 and 2022. Panel (b) shows the evolution of the flow of industrial robots adopted in Mexico according to IFR data. *Data sources*: Mexican customs records from Secretaría de Economía, own calculations based on IFR World Robotics Database.

Figure 1.2: Automation geography: robot imports across commuting zones

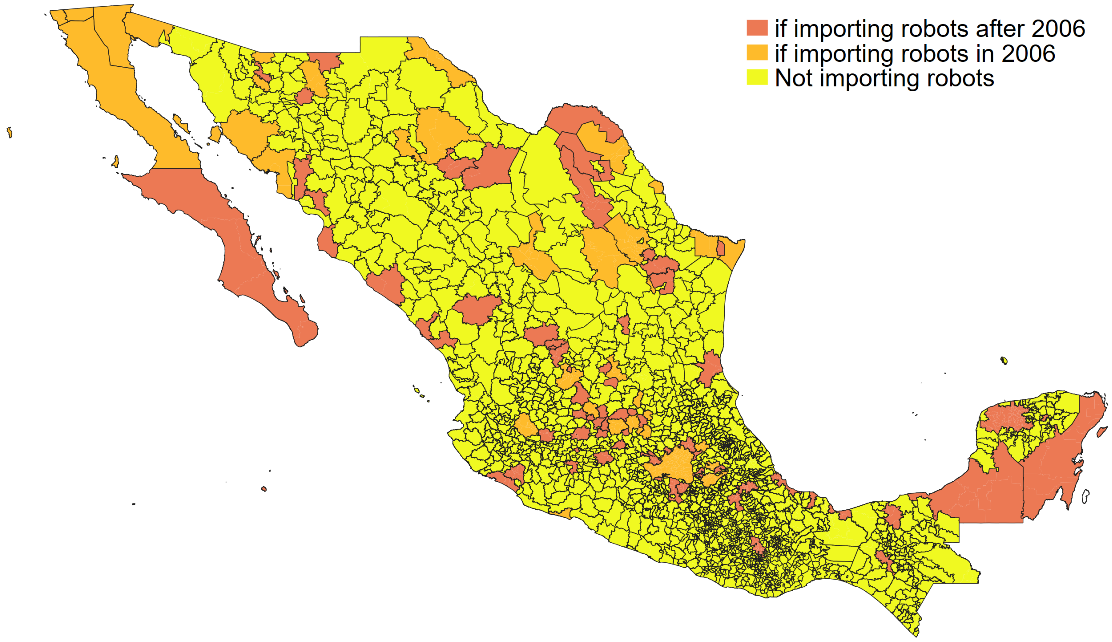

(a) Sample of commuting zones

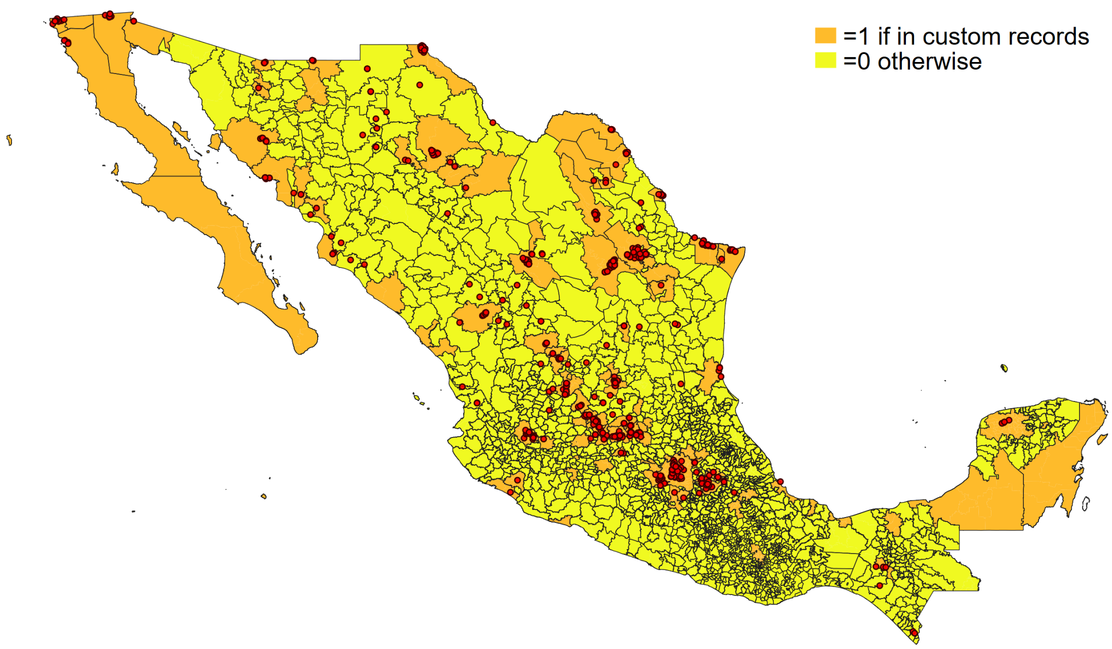

(b) Including large plants in the transportation manufacturing subsector

Notes: Figure 1.2 depicts the geographic distribution of robot imports across commuting zones. Darker shades indicate commuting zones with any robot imports during 2006–2022. The location of importation refers to the exact municipality where purchasing companies are legally registered, as determined by their fiscal address. This registration location may not necessarily coincide with the physical location of the production plant. Red dots correspond to the location of 991 large plants, defined as establishments with more than 250 employees, in transportation equipment manufacturing (NAICS 336). *Data sources*: Mexican customs records from Secretaría de Economía. National Statistical Directory of Economic Units (DENUE, 2020).

Figure 1.3: Predicting who imports robots and when

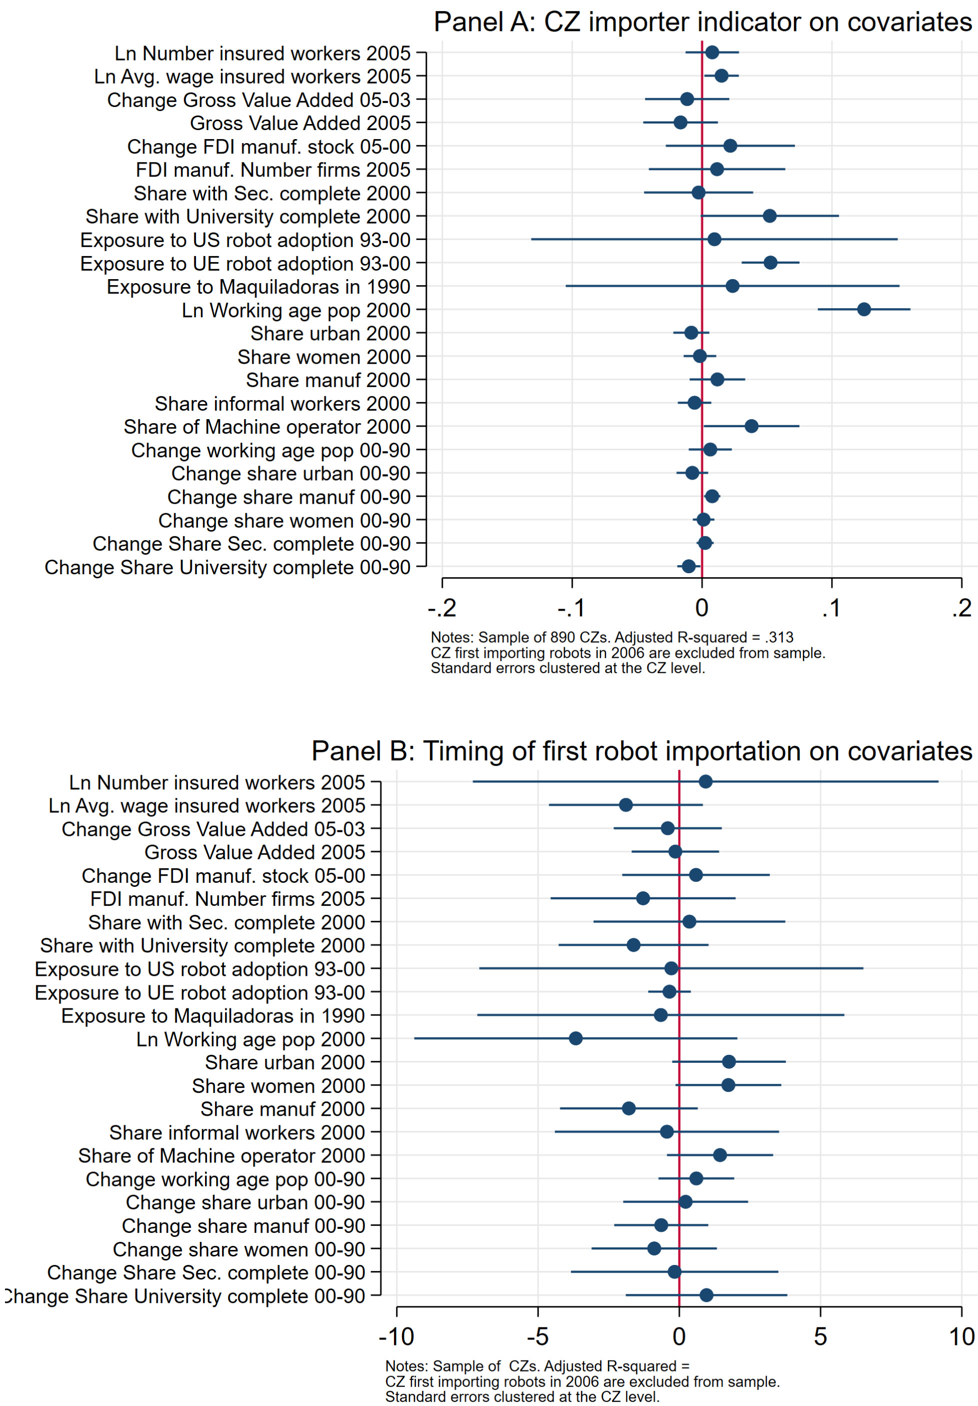

Notes: Panel (a) in Figure 1.3 displays the estimated OLS coefficients along with the corresponding 95% confidence interval of a linear probability model. In this model, the dependent variable indicates whether the CZ received robot imports between 2006 and 2022. Panel (b) illustrates the estimated OLS coefficients of a model with the year of the first robot importation event as the dependent variable. CZs that began importing robots in 2006 are excluded from the sample. The standard errors are clustered at the CZ level. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data. Mexican 1990 and 2000 censuses.

Figure 1.4: Robot adoption and formal labor market

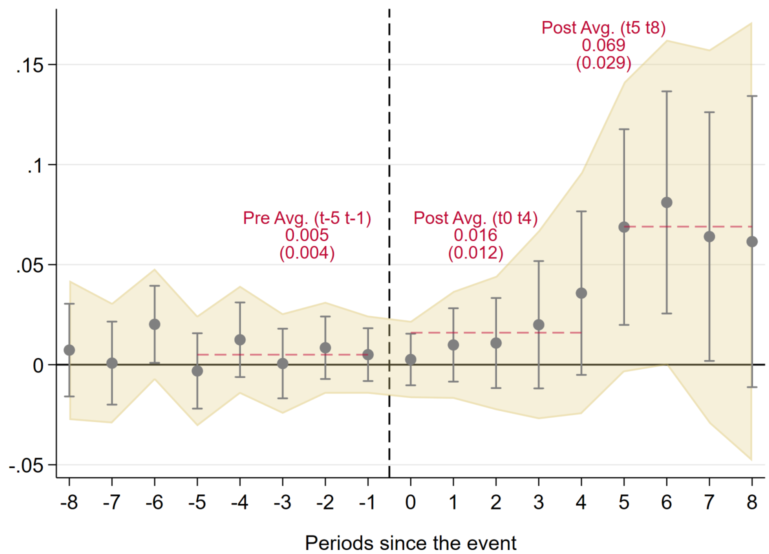

(a) Ln Number of formal workers

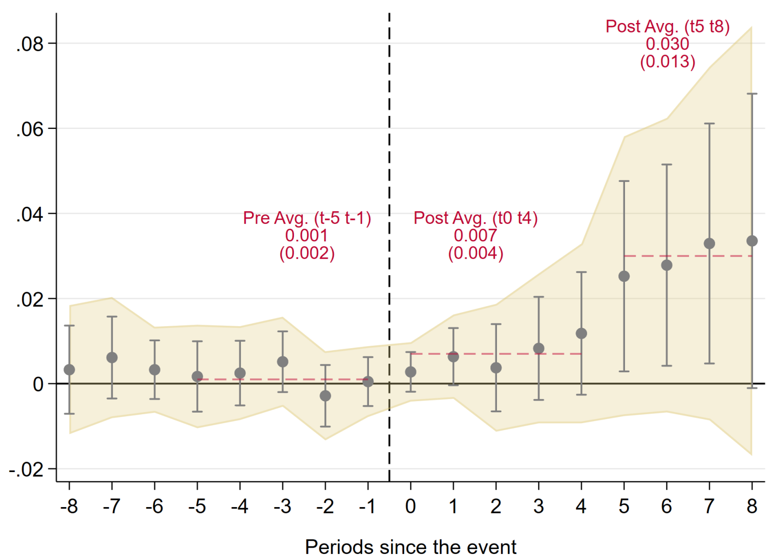

(b) Ln Average monthly wage – formal workers

Notes: Figure 1.4 plots the estimated $\theta(e)$ event-study coefficients using equation 1.2 (Callaway and Sant'Anna (2021) method). The outcome variable is, in turn, the log number of formal workers (Panel (a)) and log average real monthly wage for formal workers, deflated to constant 2010 pesos using Mexican consumer price index (CPI) (Panel (b)). The event is defined as a first-time industrial robot importation. The vertical lines reflect the 95% pointwise confidence intervals, while the shadowed area reflects the 95% uniform confidence intervals, which are robust to multiple hypothesis testing. Horizontal lines show the simple average estimated coefficient for different periods: pre-event refers to years -5 to -1; postevent short run refers to years 0 to 4; long-run refers to years 5 to 8. The clustering of standard errors is at the CZ level. Sample of not-yet-but-eventually-treated CZs for 2005–2019, including cohorts that started importing robots during 2007–2022. The unit of observation is CZ–4-digit industry (N=100,619; 71 CZs). *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Figure 1.5: TWFE models: DD & DDD

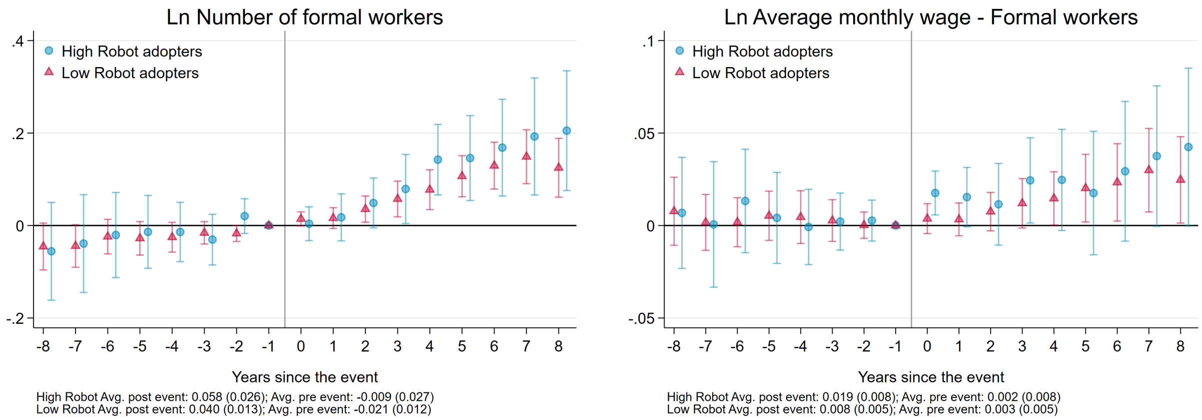

(a) DD estimates. CZ–4-dig. industry level

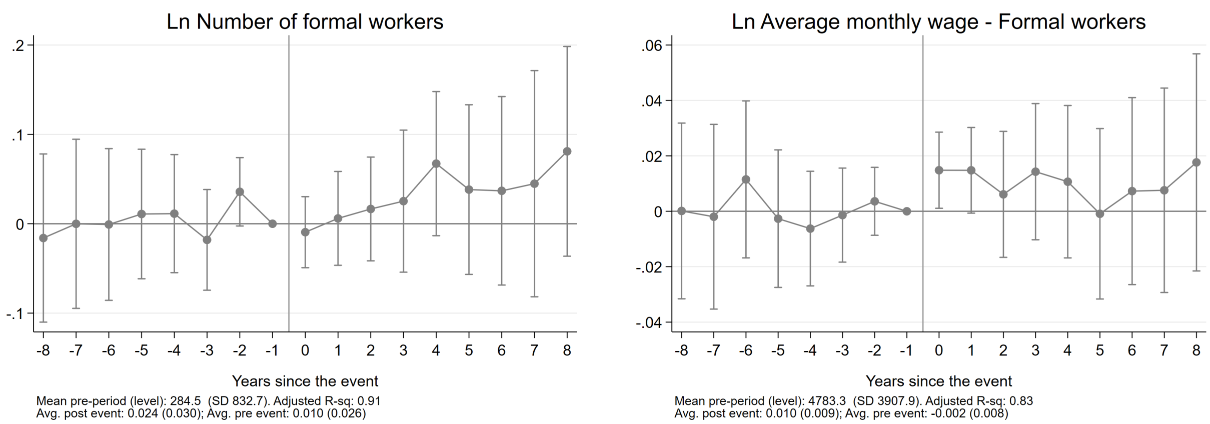

(b) DDD estimates. CZ–4-dig. industry level

Notes: Figure 1.5 displays the estimated event-study coefficients obtained using equation 1.2, a TWFE model. In Panel (a), the results of the difference-in-differences specification are presented, showing overlapping estimates from two different subsamples: "high robot adopters", comprising manufacturing industries with higher rates of robot adoption observed in the IFR data, and "low robot adopters", which include the remaining nonmanufacturing industries. Panel (b) shows estimates from the triple difference-in-differences specification, where the third difference is introduced by contrasting high and low robot adopters. The outcome variable is, in turn, the log number of formal workers and log average real monthly wage for formal workers, deflated to constant 2010 pesos using Mexican CPI. The event is defined as a first-time industrial robot importation. The clustering of standard errors is at the CZ level. Sample of CZs with fewer than 500 thousands inhabitants for the period 2005–2019, including cohorts that started importing robots during 2007–2022. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Figure 1.6: SDD estimates for each CZ from cohorts 2010–2016

(a) Sample: All industries

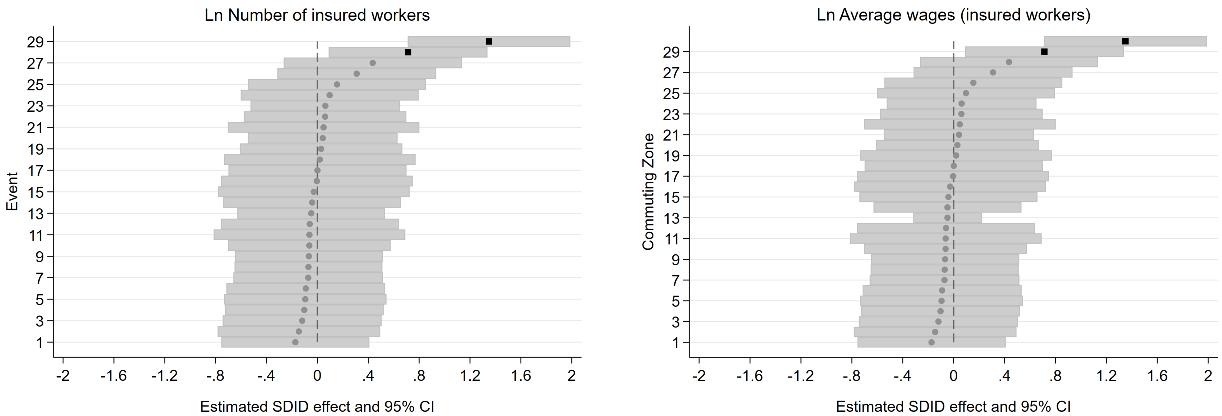

(b) Sample: Manufacturing industries

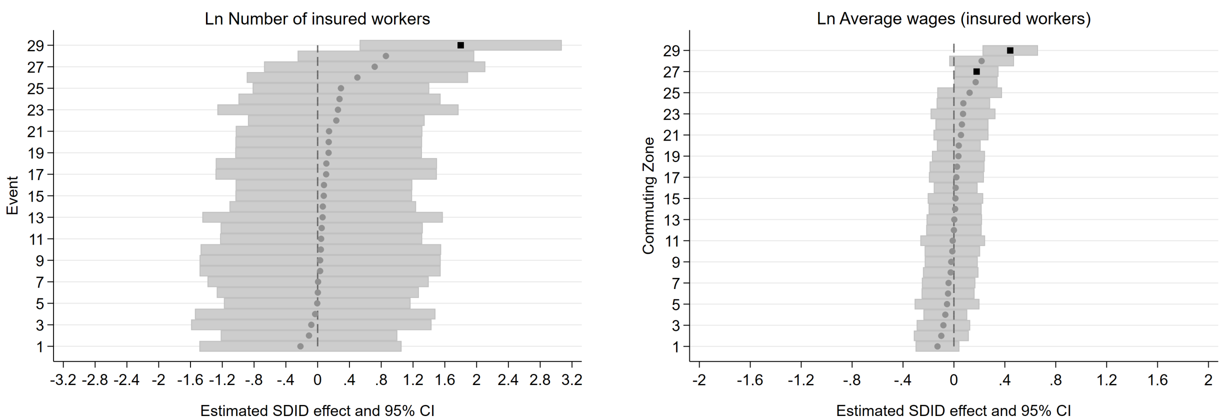

Notes: Figure 1.6 exhibits the estimated treatment effects using the synthetic difference-in-difference model as outlined by Arkhangelsky et al. (2021). In this figure, each dot represents the estimated effect for individual CZs belonging to cohorts from 2010 to 2016, and the gray bars indicate the corresponding 95% confidence intervals. The comparison group for these estimates is derived from the pool of all CZs that have never been treated, and it employs optimal weighting of units and pre-event periods in its construction. The dataset comprises a panel of CZs, balanced in calendar time, spanning from 2005 to 2019. Panel (a) presents the results for the aggregate of all industries, while Panel (b) focuses specifically on the results for the sample of aggregate manufacturing industries. The outcome variables are the logarithm of the number of formal workers and the logarithm of the average real monthly wage for formal workers, adjusted to constant 2010 pesos using Mexican CPI. The event is defined as a first-time industrial robot importation. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Figure 1.7: Role of large plant openings

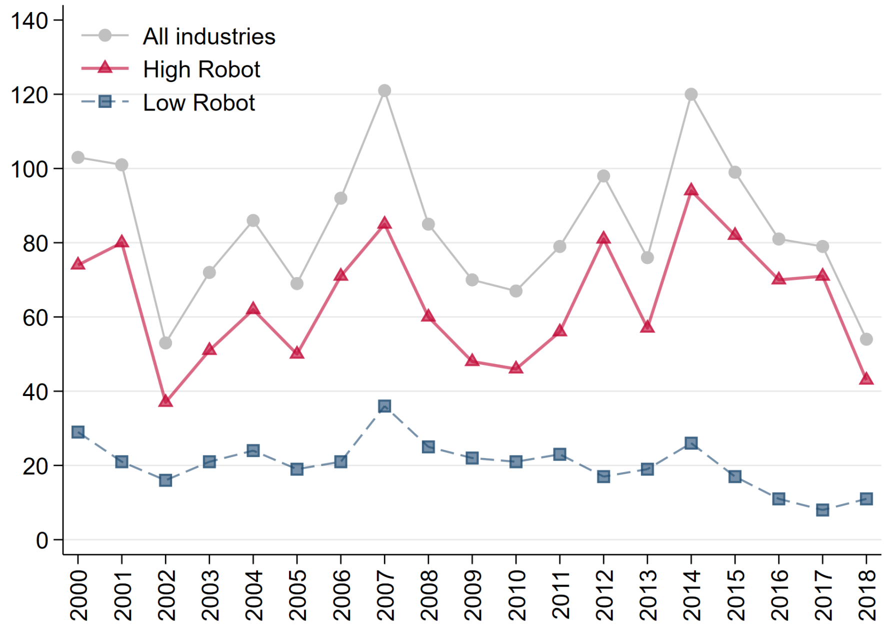

A. Number of large plant openings conditional on the plants' surviving until 2018

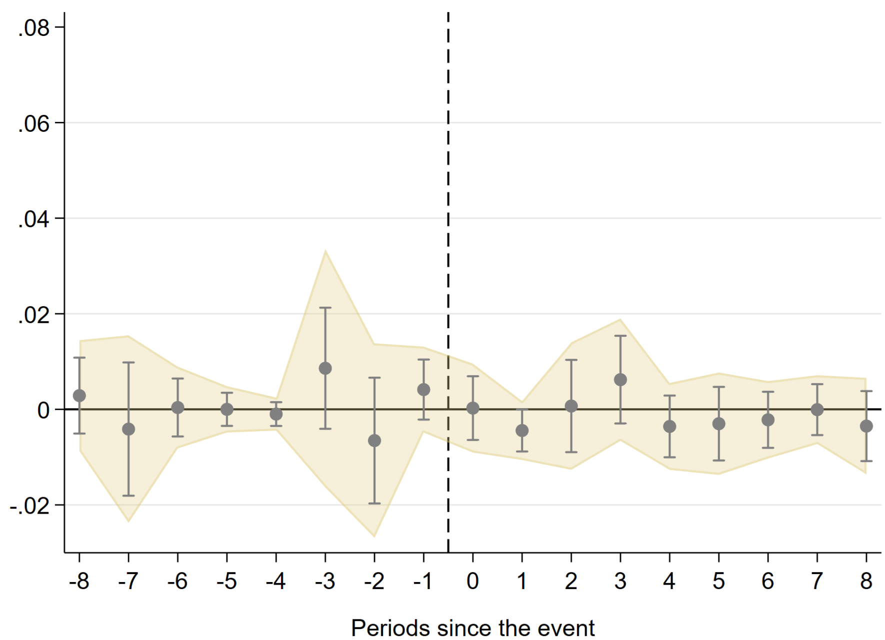

B. Effect of first robot importation on the number of large plant openings

Notes: The first graph in Figure 1.7 illustrates the number of firms in the manufacturing sector (SCIAN codes 31–33) observed in the 2018 census by opening year. The sample is further divided into four groups based on two criteria: firm size (distinguishing between firms with more or fewer than 250 workers) and the extent of robot adoption (classified as high or low robot adopters, as determined by IFR data and described in the notes for Figure 1.5). The second graph displays the estimated OLS coefficient derived from an event-study specification utilizing a TWFE model (equation 1.2). The output is the number of large manufacturing firms that opened in period t across CZs. The clustering of standard errors is at the CZ level. The event is defined as a first-time industrial robot importation. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data. Mexican economic census 2018.

FIGURE 1.8: Equilibrium in the labor market

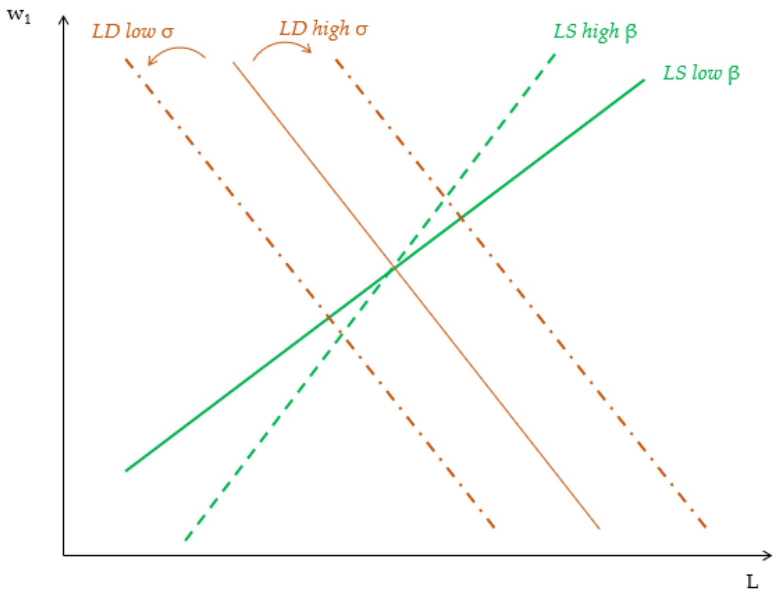

Notes: In Figure 1.8, I visually represent the equilibrium in the labor market and how it depends on given parameters. In equilibrium, the elasticity of labor supply (β) influences the slope of the labor supply curve, affecting the magnitude of employment and wage changes while not altering their direction. When β is relatively low (indicating inelastic labor supply), shifts in labor demand have a muted effect on employment and a heightened impact on wages. The direction of these effects is determined by the parameter σ representing the price elasticity of aggregate demand for manufacturing products. A higher σ results in a rightward shift of the labor demand curve, positively affecting employment and wages.

10. Tables

Table 1.1: Comparison of CZs with and without robot imports

	CZs importing robots				CZs not importing robots			
	All CZ		Excluding 2006 cohort		Balanced sample		Not balanced	
	Mean	s.d.	Mean	s.d.	Mean	s.d.	Mean	s.d.
Number of CZs	104		71		924		1612	
Number of CZs with manufacturing	104		71		730		733	
Number of CZs with manufacturing	82		52		134		134	
IMSS data (levels in 2005)								
Formal workers	1094	(6825)	423	(1468)	65	(509)	65	(508)
Monthly real wage	5350	(3950)	4739	(3697)	3487	(2560)	3487	(2558)
Formal workers in manuf.	879	(3526)	250	(791)	53	(237)	53	(237)
Monthly real wage in manuf.	5191	(3350)	4383	(3041)	3031	(1753)	3031	(1753)
Formal workers in automotive	2059	(4829)	473	(886)	145	(413)	145	(413)
Monthly real wage in automotive	6958	(4273)	5321	(3621)	3130	(1624)	3130	(1624)
Census 2000								
Population (thou.)	591.9	(2004.6)	255.6	(258.1)	33.5	(54.6)	21.9	(43.8)
Working age population (thou.)	373.6	(1307.9)	157.3	(163.4)	19.1	(33.1)	12.4	(26.4)
Share urban	76.9	(18.9)	72.3	(20.2)	42.9	(28.5)	29.7	(31.2)
Share women	51.1	(1.3)	51.2	(1.3)	51.0	(1.8)	51.1	(2.0)
Share manufacturing industry	22.3	(10.6)	19.8	(10.8)	12.1	(8.5)	10.5	(9.4)
Share automotive industry	1.9	(3.1)	1.3	(2.8)	0.3	(1.6)	0.2	(1.3)
Share agriculture industry	14.5	(11.6)	17.6	(12.2)	41.7	(17.3)	51.4	(21.0)
Share service industry	43.3	(9.9)	42.4	(11.2)	29.0	(10.6)	22.8	(12.2)
Share informal workers	26.6	(6.1)	28.1	(6.0)	43.0	(14.7)	52.9	(20.2)
Number in manuf. occupation	12.2	(3.4)	11.9	(3.5)	10.1	(6.6)	9.2	(8.2)
Share secondary complete	14.5	(4.7)	13.5	(5.0)	7.5	(4.1)	5.7	(4.2)
Share superior ed. complete	6.6	(3.2)	6.1	(3.4)	2.8	(2.1)	2.1	(2.0)
Percentage change censuses 1990–2000 (%)								
Population	24.2	(15.9)	21.3	(16.3)	8.7	(15.6)	6.9	(16.8)
Working age population	31.6	(16.9)	29.2	(18.0)	15.9	(17.4)	12.9	(18.8)
Share urban	5.4	(12.1)	6.8	(13.8)	5.8	(22.1)	3.1	(19.7)
Share women	0.5	(1.6)	0.8	(1.5)	1.0	(3.4)	1.3	(4.7)
Share manufacturing industry	18.7	(48.7)	23.3	(56.8)	42.9	(112.8)	57.2	(161.5)
Share automotive industry	189.6	(429.6)	215.0	(499.3)	39.9	(497.1)	20.6	(364.2)
Share agriculture industry	-32.6	(23.2)	-31.8	(22.3)	-24.0	(17.3)	-18.8	(29.3)
Share service industry	14.6	(17.1)	18.6	(18.1)	68.4	(71.4)	102.8	(190.6)
Number in manuf. occupation	73.2	(52.6)	77.6	(60.9)	140.6	(233.7)	158.4	(324.7)
Share secondary complete	46.0	(28.1)	50.5	(31.2)	104.1	(118.8)	114.5	(156.7)
Share superior ed. complete	72.5	(54.8)	80.1	(62.8)	141.0	(208.3)	117.6	(203.7)

Notes: Table 1.1 provides descriptive statistics comparing CZs that received robot imports during 2006–2022 (treated) to those that did not (never treated). For the treated sample, the table includes sample averages and standard deviations for all CZs. It also presents statistics for a subsample obtained by excluding the always-treated CZs, i.e., the 2006 cohort. The table also shows results for the never-treated sample for balanced and unbalanced calendar time samples. Average real monthly wage for formal workers is measured in constant 2010 pesos (adjusted using Mexican CPI). *Data sources*: IMMS public data. Mexican customs records from Secretaría de Economía. Mexican censuses 1990 and 2000.

Table 1.2: Summary statistics of import events by event cohort in CZs

Cohort	# CZs	Mean	SD	Min	Max
2006	33	13.9	5.3	1	17
2007	13	8.4	3.9	1	15
2008	6	7.3	5.3	2	14
2009	2	3.5	0.7	3	4
2010	5	5.0	2.9	1	9
2011	6	3.8	3.2	1	10
2012	5	4.0	3.1	1	8
2013	4	2.5	1.9	1	5
2014	3	2.7	1.5	1	4
2015	5	3.8	1.6	1	5
2016	6	3.0	1.4	1	5
2018	7	2.0	1.2	1	4
2019	3	1.3	0.6	1	2
2020	2	1.0	0.0	1	1
2021	2	1.0	0.0	1	1
2022	2	1.0	0.0	1	1

Notes: Table 1.2 provides summary statistics for robot import events within each treated cohort available in the data. For instance, six different CZs began receiving industrial robot imports in 2011 (2011 cohort). On average, these CZs received imports for approximately 3.8 years, considering CZs that received imports only once (minimum) or for up to ten years (maximum). *Data sources*: Mexican customs records from Secretaría de Economía.

Table 1.3: Ruling out compositional effects

	Pre-event	Postevent
	(1)	(2)
A. Log Number formal workers		
1-year window	0.002	0.010
	(0.008)	(0.007)
2-year window	0.009	0.011
	(0.006)	(0.008)
3-year window	0.007	0.006
	(0.005)	(0.014)
4-year window	0.007	0.011
	(0.006)	(0.023)
B. Log average wage – formal workers		
1-year window	0.001	0.004
	(0.003)	(0.003)
2-year window	-0.000	0.004
	(0.002)	(0.003)
3-year window	-0.001	0.007
	(0.002)	(0.005)
4-year window	-0.001	0.006
	(0.002)	(0.007)

Notes: Table 1.3 shows the simple average of the estimated $\theta(e)$ event-study coefficients using equation 1.2 for different periods. The samples are restricted to cohorts included in samples balanced in event-time for different window lengths and that are observed during those periods. The event is defined as a first-time industrial robot importation. Sample of not-yet-but-eventually-treated CZs for the period 2005–2019, including cohorts that started importing robots during 2007–2022. Clustering of standard errors is at the CZ level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Table 1.4: Main estimates for different samples

	Pre-event	Postevent short run	Postevent long run
	(1)	(2)	(3)
A. Log Number formal workers			
Baseline	0.005	0.016	0.069***
	(0.004)	(0.012)	(0.029)
Balanced sample	0.005	0.005	0.058**
	(0.004)	(0.011)	(0.029)
Manufacturing	0.009	0.012	0.073
	(0.008)	(0.018)	(0.049)
Manufacturing robot adopters	0.011	0.017	0.103*
	(0.009)	(0.023)	(0.059)
Never treated	0.012***	0.041***	0.108***
	(0.004)	(0.010)	(0.021)
B. Log Average wage – formal workers			
Baseline	0.001	0.007	0.030**
	(0.002)	(0.004)	(0.013)
Balanced sample	0.002	0.005	0.027**
	(0.002)	(0.004)	(0.013)
Manufacturing	0.003	0.000	0.024
	(0.002)	(0.008)	(0.020)
Manufacturing robot adopters	0.001	0.007	0.040*
	(0.003)	(0.010)	(0.024)
Never treated	0.001	0.006	0.026***
	(0.001)	(0.004)	(0.008)

Notes: Table 1.4 shows the simple average of the estimated $\theta(e)$ event-study coefficients using equation 1.2 for different periods. Pre-event: years -5 to -1; postevent short run: years 0 to 4; long run: years 5 to 8. The outcome variable is, in turn, log number of formal workers (Panel (a)) and log average real monthly wage for formal workers, deflated to constant 2010 pesos using Mexican CPI (Panel (b)). The event is defined as a first-time industrial robot importation. The baseline sample includes not-yet-but-eventually-treated CZs for 2005–2019, including cohorts that started importing robots during 2007–2022. The manufacturing sample is a subset of the baseline sample, focusing only on CZs within the manufacturing sector (3-digit SCIAN codes ranging from 311 to 339). The robot-adopter sample further narrows the selection to manufacturing sectors that adopt robots in the IFR dataset. Finally, the never-treated sample is an extension of the baseline sample encompassing CZs that did not receive robot imports during the analyzed period, which serves as a comparison group. Clustering of standard errors is at the CZ level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

TABLE 1.5: Heterogeneities

	Log Number Formal Workers			Log Average Wage		
	Pre-event	Postevent short run	Postevent long run	Pre-event	Postevent short run	Postevent long run
	(1)	(2)	(3)	(4)	(5)	(6)
A. By age group						
Aged 15–19	0.006	0.004	0.032	-0.000	-0.004	0.001
	(0.007)	(0.027)	(0.044)	(0.002)	(0.007)	(0.012)
Aged 20–34	0.006	0.011	0.071**	-0.000	0.004	0.029***
	(0.004)	(0.012)	(0.035)	(0.002)	(0.004)	(0.012)
Aged 35–49	0.006*	0.001	0.028	0.000	0.007*	0.026**
	(0.003)	(0.012)	(0.026)	(0.002)	(0.004)	(0.013)
Aged 50–64	0.007***	0.016	0.085***	0.005***	-0.007	0.008
	(0.003)	(0.011)	(0.023)	(0.002)	(0.007)	(0.015)
Age >64	0.009*	0.006	0.007	0.007	-0.003	0.029
	(0.005)	(0.015)	(0.040)	(0.004)	(0.010)	(0.027)
B. By wage group						
Low-wage group	0.002	0.003	0.035	0.002*	-0.002	0.009
	(0.004)	(0.011)	(0.031)	(0.001)	(0.004)	(0.011)
Middle-wage group	0.002	0.019	0.070*	0.001	0.002	0.007
	(0.005)	(0.017)	(0.037)	(0.001)	(0.002)	(0.005)
High-wage group	0.003	-0.006	0.024	-0.002*	0.001	0.002
	(0.006)	(0.018)	(0.036)	(0.001)	(0.003)	(0.009)
C. By firm Size						
1–5 workers	0.006*	0.005	0.037	0.002	0.003	0.027**
	(0.003)	(0.011)	(0.024)	(0.002)	(0.006)	(0.013)
6–50 workers	-0.003	-0.003	0.023	0.001	-0.003	0.017
	(0.003)	(0.011)	(0.021)	(0.001)	(0.005)	(0.012)
51–250 workers	0.007	0.004	0.041	-0.003	-0.003	-0.006
	(0.005)	(0.018)	(0.030)	(0.003)	(0.007)	(0.022)
>250 workers	-0.001	0.028*	0.023	0.002	-0.007	0.014
	(0.005)	(0.015)	(0.038)	(0.003)	(0.008)	(0.020)
D. By CZ size						
Excluding >500k hab.	0.004	0.022	0.096***	0.002	0.012**	0.038***
	(0.003)	(0.014)	(0.034)	(0.002)	(0.005)	(0.016)
Excluding >150k hab.	0.004	0.017	0.133**	0.002	0.023***	0.049***
	(0.005)	(0.022)	(0.057)	(0.002)	(0.008)	(0.020)

Notes: See notes in Table 1.4. Sample of not-yet-but-eventually-treated CZs for the period 2005–2019, including cohorts that started importing robots during 2007–2022. Clustering of standard errors is at the CZ level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Table 1.6: Other robot adoption measures

	Pre-event	Postevent short run	Postevent long run
	(1)	(2)	(3)
A. Log Number formal workers			
Baseline CS	0.005	0.016	0.069***
	(0.004)	(0.012)	(0.029)
Import value event	<0.001	0.028	0.383***
	(0.007)	(0.025)	(0.071)
Baseline DD	-0.003	0.033*	0.100***
	(0.011)	(0.019)	(0.043)
DD with intensity	-0.001	0.031*	0.066***
	(0.016)	(0.016)	(0.026)
B. Log Average wage – formal workers			
Baseline CS	0.001	0.007	0.030**
	(0.002)	(0.004)	(0.013)
Import value event	-0.001	-0.001	0.067**
	(0.002)	(0.010)	(0.030)
Baseline DD	-0.007	0.017	0.046
	(0.008)	(0.013)	(0.030)
DD with intensity	-0.014**	0.017***	0.035***
	(0.007)	(0.006)	(0.010)

Notes: Table 1.6 shows the event-study estimates using different robot adoption measures. *Baseline CS* corresponds to the main specification in Table 1.4. *Import value event* refers to the same model but defines the event as the period when the cumulative total value of robot imports is larger than 50% of the CZ's 2005 total wage bill (refer to Figure A.15 to see estimates under different thresholds). *Baseline DD* refers to the estimation of a TWFE model defining the event as first-time receipt of industrial robot imports. *DD with intensity* refers to *Baseline DD* interacting the event with a measure of robot importation intensity, the cumulative number of years receiving robot imports. Pre-event: years -5 to -1; postevent short run: years 0 to 4; long run: years 5 to 8. The outcome variable is, in turn, log number of formal workers (Panel (a)) and log average real monthly wage for formal workers, deflated to constant 2010 pesos using Mexican CPI (Panel (b)). Sample of not-yet-but-eventually-treated CZs for the period 2005–2019, including cohorts that started importing robots during 2007–2022. Clustering of standard errors is at the CZ level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Table 1.7: Additional labor market outcomes from ENOE data

	Mean	Pre-event	Event year	Postevent short run	Postevent long run
	(1)	(2)	(3)	(4)	(5)
A. Formal employment and earnings					
Log number workers (IMSS)	54122.5	0.014***	0.003	0.044**	0.118
		(0.006)	(0.008)	(0.020)	(0.076)
Log number workers (ENOE)	46452.3	0.042	-0.054	0.036	0.300
		(0.035)	(0.072)	(0.086)	(0.224)
Log average wage (IMSS)	208.6	<0.001	0.008*	0.008	0.012
		(0.003)	(0.004)	(0.008)	(0.016)
Log hours worked	49.4	<0.001	0.016	-0.013	-0.031*
		(0.003)	(0.017)	(0.009)	(0.018)
B. Informal employment					
Log number informal workers	84293.8	0.033*	-0.139**	-0.043	0.169*
		(0.018)	(0.061)	(0.070)	(0.099)
Informality rate	60.2	0.163	-2.111**	-1.016	1.600
		(0.451)	(1.043)	(1.423)	(2.068)
C. Unemployment					
Log number unemployed	6235.5	0.018	-0.497	-0.203	0.313
		(0.095)	(0.308)	(0.246)	(0.329)
Unemployment rate	0.04	-0.001	-0.003	0.000	-0.002
		(0.001)	(0.004)	(0.004)	(0.005)
D. Educational attainment					
Log number less than primary	27473.5	-0.030	0.063	0.054	0.199
		(0.039)	(0.200)	(0.260)	(0.248)
Log number primary complete	121210.9	0.021	-0.082	-0.025	0.094
		(0.018)	(0.064)	(0.074)	(0.108)
Log number secondary complete	51878.8	0.069	-0.067	-0.062	0.209
		(0.053)	(0.064)	(0.093)	(0.142)
Log number tertiary complete	30868.6	0.121	0.018	0.056	0.813*
		(0.076)	(0.197)	(0.218)	(0.455)
E. Other					
Log number working-age pop.	231568.3	0.020	-0.079	-0.005	0.158
		(0.017)	(0.056)	(0.065)	(0.099)
Log number pop. in LFP	149827	0.027	-0.102*	-0.026	0.128
		(0.017)	(0.055)	(0.065)	(0.095)
Log number living in urban area	134449.7	0.001	-0.002	0.011*	0.034*
		(0.002)	(0.003)	(0.006)	(0.020)

Notes: See notes in 1.4. Sample of not-yet-but-eventually-treated CZs for the period 2005–2019, including cohorts that started importing robots during 2007–2022. Sample is restricted to CZs available in ENOE. Clustering of standard errors is at the CZ level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data. ENOE data.

Table 1.8: Effects of robots on changes in labor market outcomes: 2SLS estimates using the change in robot exposure between 2011 and 2015

A: Population Censuses outcomes. Long differences 2010-2015

	All industries											
	Employment-to-pop		Wage-empl-to-pop		Self-empl-to-pop		Labor Force Particip. rate		Unemployment rate		Share Informal	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
MX exposure to robots	-0.4	-0.4	-0.5	-0.5	0.1	0.1	-0.3	-0.4	-0.7	-0.7	0.4	0.4
2011-2015	(1.6)	(2.1)	(1.4)	(1.6)	(1.1)	(1.2)	(1.4)	(1.9)	(1.1)	(1.2)	(1.7)	(1.6)
Observations	1,804	19	1,804	19	1,804	19	1,804	19	1,798	19	1,797	19
Adj. R-squared	0.154		0.0531		0.180		0.137		0.0231		0.155	
Mean dep. 2005	45.47		23.50		21.97		50.39		4.406		49.13	
Lower limit CI	-3.452	-4.502	-3.211	-3.571	-1.974	-2.156	-3.165	-4.099	-2.803	-3.056	-2.86	-2.793
Upper limit CI	2.703	3.713	2.216	2.565	2.22	2.373	2.479	3.388	1.386	1.659	3.699	3.606
First-stage coefficient	0.26		0.26		0.26		0.26		0.266		0.266	
First-stage F-statistic	4.895		4.895		4.895		4.895		5.169		5.17	

B: IMSS data - Insured workers. Long differences 2010-2015

	All industries						Manufacturing industry					
	Ln Number Workers		Ln Payroll		Ln Mean Wage		Ln Number Workers		Ln Payroll		Ln Mean Wage	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
MX exposure to robots	5.0	5.1	4.2	4.2	-0.9	-0.8	-26.3	-26.1	-19.7	-19.6	6.5	6.5
2011-2015	(10.9)	(9.7)	(12.5)	(8.6)	(4.5)	(2.6)	(28.4)	(16.1)	(27.7)	(15.5)	(10.1)	(6.9)
Observations	1,211	19	1,211	19	1,211	19	777	19	777	19	777	19
Adj. R-squared	0.010		0.014		0.008		0.067		0.037		0.031	
Mean dep. 2005	12050		3004000		158.10		3485		954484		134.20	
Lower limit CI	-16.44	-13.98	-20.29	-12.57	-9.682	-5.979	-81.98	-57.54	-74.04	-49.96	-13.22	-7.108
Upper limit CI	26.5	24.12	28.64	21.02	7.97	4.29	29.47	5.344	34.54	10.83	26.23	20.17
First-stage coefficient	0.252		0.252		0.252		0.24		0.24		0.24	
First-stage F-statistic	3.976		3.976		3.976		3.043		3.043		3.043	

Notes: This table presents 2SLS estimates of the effects of exposure to robots on changes in labor market outcomes between 2010 and 2015. Panel (a) presents outcomes based on population censuses and Panel (b) those based on the IMSS dataset. In all models, I instrument Mexican exposure to robots using exposure to robots from European countries. All IV estimates are from regressions weighted by CZ share in national working-age population in 1990. The covariates included in each model are region dummies, exposure to US robots (2011–2015), exposure to *maquiladoras* in 1990, demographic characteristics of CZs in 1990 (log population; share urban; share women; share population over 65; share population with high school complete; share population with college complete; and shares of employment in manufacturing, services and agriculture sectors). For each outcome, the first column presents standard error estimates that are robust to heteroskedasticity and correlation within CZ, and the second column presents robust standard errors computed following Borusyak et al. (2019) in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. *Data sources*: IMMS public data. Mexican 1990 and 2000 censuses. IFR data. Faber (2020).

TABLE 1.9: Mechanisms

	Number of firms	Labor share	Gross value added (log)	GVA per worker (log)	Total fixed assets (log)	Capital per worker	Input cost (log)	Total payroll (log)	Number workers (log)	Workers prod. and sales (log)	Workers admin (log)	Hours worked (log)	Hours prod. and sales (log)	Hours admin (log)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Panel (a). All industries														
D_{mt}^{post}	0.012	-0.011	0.054*	0.028	0.017	-0.019	0.060**	0.025	0.018	0.009	0.000	0.020	0.011	-0.006
	(0.013)	(0.068)	(0.030)	(0.024)	(0.043)	(0.014)	(0.025)	(0.029)	(0.015)	(0.023)	(0.028)	(0.015)	(0.024)	(0.028)
Observations	93,464	32,979	32,352	32,352	32,908	43,054	32,982	30,945	43,054	40,283	27,181	42,913	40,149	27,071
Adj. R-squared	0.932	0.131	0.849	0.626	0.803	-0.0818	0.875	0.852	0.890	0.842	0.751	0.885	0.832	0.737
Mean	27.80	0.330	14.83	0.0642	19.09	0.0508	14.86	4.812	200.5	98.15	24.59	501.9	245.9	60.44
Panel (b). Manufacturing sector: Metal, machinery and equipment, power generation, transportation, plastics and rubber														
D_{mt}^{post}	0.024	-0.025	0.301*	0.169*	0.351**	0.026*	0.298**	0.304**	0.072	0.099	-0.040	0.077	0.105	-0.022
	(0.024)	(0.056)	(0.156)	(0.098)	(0.159)	(0.014)	(0.134)	(0.139)	(0.073)	(0.095)	(0.127)	(0.078)	(0.105)	(0.131)
Observations	5,509	769	761	761	765	983	769	695	983	901	586	983	901	586
Adj. R-squared	0.887	0.0558	0.794	0.608	0.751	0.190	0.797	0.782	0.802	0.755	0.669	0.786	0.738	0.662
Mean	12.50	0.400	15.78	0.0521	11.99	0.0478	21.20	8.296	164.1	98.48	19.18	359.7	224.1	42.48

Notes: Table 1.9 shows the results of estimating the TWFE model, equation 1.3, where the event is defined as a first-time industrial robot importation. Panel (a) presents results for all the industries and Panel (b) for the following manufacturing industries: basic metals, metal products, machinery, equipment, electronics, appliances, power generation, transportation, and other manufacturing not included in Table 1.10. The first dependent variable is the log number of firms at the CZ–five-digit sector level (column (1)). The following seven dependent variables are balance-sheet measures capturing firm productivity and size, aggregated to the CZ–five-digit sector level: labor share, calculated as the total wage bill divided by gross value added (column (2)); capital per worker, calculated as the total fixed assets divided by the number of workers (column (3)); log total gross value added (column (4)) and GVA per worker (column (5)); log input costs (column (6)); log total wage bill (column (7)); and log fixed assets (as a proxy for capital, column (8)). Column (9) presents results for the log number of workers, column (10) for the log number of workers in production and sales, and column (11) for the log number of workers in administrative jobs (including managers). Columns (12)–(14) show similar results for the log number of hours worked in a year. The sample includes not-yet-but-eventually-treated CZs for 2005–2019, including cohorts that started importing robots during 2007–2022. Means are calculated for the period before the event, in levels, and are reported in millions of Mexican pesos (CPI-adjusted to 2005), except for the number of workers, number of hours worked in a year (in thousands), and shares (ratios). Standard errors are clustered at the CZ level and are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. *Data sources*: Mexican economic censuses 2003, 2008, 2013, and 2018.

TABLE 1.10: Mechanisms: Other industries

	Number of firms	Labor share	Gross value added (log)	GVA per worker (log)	Total fixed assets (log)	Capital per worker	Input cost (log)	Total payroll (log)	Number workers (log)	Workers prod. and sales (log)	Workers admin (log)	Hours worked (log)	Hours prod. and sales (log)	Hours admin (log)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Panel (a). Manufacturing: Wood and paper, chemical industries														
D_{mt}^{post}	-0.014	-0.018	0.031	0.014	0.063	0.032**	-0.013	-0.026	-0.084	-0.203**	-0.195*	-0.073	-0.203**	-0.185
	(0.023)	(0.031)	(0.111)	(0.066)	(0.116)	(0.013)	(0.100)	(0.114)	(0.066)	(0.100)	(0.111)	(0.067)	(0.101)	(0.114)
Observations	5,356	1,068	1,036	1,036	1,063	1,433	1,068	1,014	1,433	1,354	929	1,433	1,354	929
R-squared	0.923	0.393	0.799	0.612	0.790	0.581	0.819	0.754	0.843	0.753	0.649	0.842	0.755	0.632
Mean	24.10	0.346	32.84	0.0933	32.55	0.0542	74.43	6.827	285.4	140.3	30.88	731.1	382.7	77.46
Panel (b). Manufacturing: Textiles, apparel, leather goods, food, and beverage industries														
D_{mt}^{post}	-0.028	0.019	0.042	0.051	-0.009	0.005	0.030	0.144	-0.094	-0.017	-0.072	-0.070	0.003	-0.065
	(0.030)	(0.054)	(0.093)	(0.060)	(0.118)	(0.010)	(0.081)	(0.113)	(0.063)	(0.091)	(0.097)	(0.065)	(0.096)	(0.099)
Observations	7,393	1,721	1,715	1,715	1,690	2,235	1,721	1,443	2,235	1,902	969	2,235	1,902	969
R-squared	0.888	0.0699	0.751	0.516	0.729	0.521	0.787	0.739	0.783	0.739	0.727	0.773	0.726	0.715
Mean	15.32	0.388	15.30	0.0419	15.72	0.0378	27.63	7.844	227.6	177.7	33.62	496.6	393.5	75.50
Panel (c). Nonmanufacturing														
D_{mt}^{post}	0.017	-0.012	0.050*	0.025	0.009	-0.023	0.059**	0.015	0.027*	0.016	0.011	0.027*	0.017	0.003
	(0.013)	(0.075)	(0.030)	(0.024)	(0.043)	(0.016)	(0.025)	(0.027)	(0.014)	(0.022)	(0.029)	(0.015)	(0.023)	(0.029)
Observations	75,206	29,421	28,840	28,840	29,390	38,403	29,424	27,793	38,403	36,126	24,697	38,262	35,992	24,587
R-squared	0.935	-0.0462	0.856	0.629	0.807	-0.0718	0.881	0.861	0.899	0.852	0.757	0.894	0.841	0.742
Mean	30.28	0.325	14.11	0.0648	18.96	0.0514	11.72	4.495	196.5	92.28	24.11	496.7	233.2	59.58

Notes: Table 1.10 shows the results of estimating the TWFE model, equation 1.3, where the event is defined as a first-time industrial robot importation. Panel (a) presents results for the following manufacturing industries: food, drinks, tobacco, energy, fossil fuels, chemicals, plastics, and rubber. Panel (b) presents results for the following manufacturing industries: textile, clothing, leather, wood, paper, printing, and nonmetallic mineral products. Panel (c) includes results for all nonmanufacturing industries. The first dependent variable is the log number of firms at the CZ–five-digit sector level (column (1)). The following seven dependent variables are balance-sheet measures capturing firm productivity and size, aggregated to the CZ–five-digit sector level: labor share, calculated as the total wage bill divided by gross value added (column (2)); capital per worker, calculated as the total fixed assets divided by the number of workers (column (3)); log total gross value added (column (4)) and GVA per worker (column (5)); log input costs (column (6)); log total wage bill (column (7)); and log fixed assets (as a proxy for capital, column (8)). Column (9) presents results for the log number of workers, column (10) for the log number of workers in production and sales, and column (11) for the log number of workers in administrative jobs (including managers). Columns (12)–(14) show similar results for the log number of hours worked in a year. The sample includes not-yet-but-eventually-treated CZs for 2005–2019, including cohorts that started importing robots during 2007–2022. Means are calculated for the period before the event, in levels, and are reported in millions of Mexican pesos (CPI-adjusted to 2005), except for the number of workers, number of hours worked in a year (in thousands), and shares (ratios). Standard errors are clustered at the CZ level and are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. *Data sources*: Mexican economic censuses 2003, 2008, 2013, and 2018.

Chapter 2

Social Mobility and Economic Development

1. Introduction

Equality of opportunity and social mobility are values shared by many and are very important policy objectives rooted in the constitution of most countries. From an empirical perspective, it remains an open question whether higher social mobility is also beneficial for economic performance. Establishing the existence of a positive effect of improved social mobility on economic indicators would provide an even greater justification for targeting it as a policy objective, beyond the usual equity argument.

From a theoretical point of view, in a world in which abilities are transmitted perfectly from parents to children, and income inequality is merely the result of returns to individual ability, redistributing opportunities to the children of less able (and hence less affluent) parents at the expense of the children of more able ones might induce distortions causing considerable efficiency loss. However, abilities are not perfectly transmitted across generations, and other factors play an important role in the distribution of income (e.g. Bowles and

Gintis, 2002; Black et al., 2020; Sacerdote, 2011). Under these conditions, creating better opportunities for the less affluent, and thus increasing social intergenerational mobility, should lead to a more efficient accumulation of human capital, reduce the misallocation of talent, and eventually improve the performance of the economy (e.g. Galor and Tsiddon, 1997; Galor and Moav, 2004; Mincer, 1984). Our aim in this study is to test these predictions, analyzing the role of intergenerational mobility for economic development.[1]

Our paper makes a contribution to the literature that studies how inequality in access to resources and in opportunities may affect economic performance (e.g. Galor and Zeira, 1993; Banerjee and Duflo, 2003; Voitchovsky, 2005; Brueckner and Lederman, 2018; Van der Weide and Milanovic, 2018; Marrero and Rodríguez, 2013; Ferreira et al., 2018) providing the first large-scale study on the role of social mobility for economic efficiency. Recent descriptive studies suggest a positive correlation between mobility and economic performance indicators across countries and within single countries across geographical areas (e.g. Chetty et al., 2014; Güell et al., 2018; Neidhöfer et al., 2018; Aghion et al., 2019; Aydemir and Yazici, 2019). In this study, we go one step further by providing estimates based on subnational region-level panel data for multiple countries. Our laboratory of analysis is Latin America, a region with interesting historical similarities and, at the same time, ample heterogeneity in economic development, income inequality, and social mobility trends. For instance, while in most Latin American countries, such as Argentina, Brazil, Chile, and Mexico, educational upward mobility increased substantially from 1940 to 1990, cohorts born in the 1980s in Guatemala and Nicaragua still experience relatively low levels of upward mobility (Neidhöfer et al., 2018). At the same time, Guatemala and Nicaragua show substantially lower GDP per capita levels and trends from 1990 to today compared to the other countries (World

[1]The essay "The Misallocation of Talent" by Rodríguez Mora (2009) motivates the importance of the subject: "A society with low intergenerational mobility is not only unfair, it is inefficient. There is no trade-off between fairness and efficiency when increasing mobility: the more there is, the fairer and more efficient society. (...) It is hard to think about fairness, since what is fair for some is unfair for others. Efficiency is a much more powerful concept; if an allocation is inefficient, it is so for everybody. Society (as a whole) could do better."

Bank: World Development Indicators). To test the relationship between these two phenomena, we construct a unique data set of (sub-national) region-year observations for 10 Latin American countries.The dataset includes information about the intergenerational mobility of education for people born between 1940-1989, and several development indicators, such as average income, poverty rates, labor formality, and luminosity information from satellite data, covering the period from 1981 to 2018. To link social mobility and economic development, we implement a new methodology that connects cohort- and year-level observations by weighting the degree of mobility of a cohort based on its contribution to the overall economic performance of the respective country in each year.

Our results show that intergenerational mobility is consistently associated with economic development. We document strong variation in terms of social mobility and the level of economic development across and within Latin American countries and find that higher intergenerational mobility is associated with rising income per capita, income growth, and other development indicators. These results are robust to different social mobility measures, hold when controlling for unobserved cross-regional heterogeneity and spillover effects, and do not depend on factors related to migration, educational expansions, and initial conditions. Results are also robust to the inclusion of contemporaneous income inequality, meaning that even when controlling for this factor, intergenerational mobility remains relevant for explaining economic development. An interesting picture also emerges when observing the interaction of cross-sectional income inequality and intergenerational mobility: Holding social mobility constant, the association between inequality and economic development is positive. However, the interaction between the two can be particularly detrimental to development when inequality is high and, at the same time, social mobility is low.

These findings have important policy implications. They suggest that there is no equity-efficiency trade-off regarding social mobility. Instead, our results show that improving the opportunities of disadvantaged individuals creates positive economic returns. Hence, even if interventions aimed at improving intergenerational mobility may cause inefficiencies in the short-run, cost-benefit analyzes should also take their positive long-run impact on the economy into account, which may still justify their use.

This paper is organized as follows: Section 2. provides an intuitive conceptual framework about the role of opportunities and social mobility for economic development and reviews the theoretical and empirical literature. Section 3. explains the estimation strategy. Section 4. describes the data, as well as the measurement of social intergenerational mobility and economic development. Section 5. maps the geography of intergenerational mobility in Latin America. Section 6. estimates the impact of social mobility on economic development. Section 7. concludes.

2. Social Mobility and Economic Development: Conceptual Framework and Literature Review

Economic reasoning suggests that human capital promotes development. Hence, improving the opportunities to invest in human capital should enhance its accumulation and allocation, eventually supporting the process of economic development (e.g. Galor and Zeira, 1993; Galor and Tsiddon, 1997; Galor and Moav, 2004; Hassler and Rodríguez Mora, 2000; Maoz and Moav, 1999; Owen and Weil, 1998). In turn, inequality of opportunities harms the accumulation and allocation of human capital and reduces social mobility. In modern economics, the works by Becker and Tomes (1979), Becker and Tomes (1986), Loury (1981), Solon (1992), among others, set the theoretical and conceptual basis of the literature on social intergenerational mobility, modeling the mechanisms and transmission channels that explain the persistence of economic outcomes of families between generations. In these models, intergenerational mobility depends primarily on the inheritance of abilities from parents to children, as well as on private and public investments in human capital.[2] Thus, the persistence of inequality between family lineages over time is an indicator of the opportunities for individuals to achieve economic well-being through their own efforts, independent of circumstances beyond their control, such as the family environment they were born into (Roemer, 1998). Generally, the more equally distributed opportunities for human capital formation are, the higher is intergenerational mobility. These opportunities are directly influenced by under-investments that may exist due to budget constraints, credit market imperfections, or informational asymmetries, among other factors (Heckman and Mosso, 2014).

Social intergenerational mobility is, thus, a measure of how likely it is for people to realize their full potential and make the most of their intrinsic talents and abilities, regardless of the family background they were born into.[3] If the innate abilities of children and their parents are not perfectly correlated, and the distribution of talent in the population has an idiosyncratic component, unequal opportunities to invest in human capital cause talent to be misallocated. Improving social mobility implies that people have better opportunities to take advantage of their potential. This has, in turn, positive repercussions on the accumulation and allocation of human capital, increases the pool of talent in the labor force, and eventually improves economic performance.

Empirical studies at the micro-level find support for the positive relationship between individual opportunities for economic success and economic performance. Bell et al. (2019) highlight the role played by the childhood-environment for innovation and progress. Bandiera et al. (2017) evaluate an intervention that enabled poor women by reducing barriers to taking on better work opportunities and find that the program contributed to sustainable poverty reduction among beneficiaries while not making ineligible households to be worse off.

[2] On the role of genetics and the environment in determining long-run outcomes of children see, among others, Bowles and Gintis (2002), Black et al. (2020), Sacerdote (2011).

[3] Analyzing the mechanisms affecting social mobility—such as territorial segregation across neighborhoods, early childhood policies, educational systems, informational barriers etc- and their relative effectiveness in improving equality of opportunity goes beyond the scope of this work. For a review of the causal evidence on the topic, see Stuhler (2018).

Hsieh et al. (2019) show that improving occupational opportunities for disadvantaged groups causes a better allocation of talent and higher aggregate productivity. Hereby, barriers to forming human capital, such as credit constraints (e.g. Galor and Zeira, 1993) or under-nutrition (e.g. Dasgupta and Ray, 1986), have been argued to be particularly important. Another factor limiting individual opportunities and, hence, harming economic development are inefficiently low aspirations (e.g. Genicot and Ray, 2017; La Ferrara, 2019). Individuals belonging to poor households may have lower aspirations than rich individuals because they anticipate unfair chances in their future. This anticipation can push the poor to choose lower levels of human capital investment, thus perpetuating their economic disadvantage. The resulting non-optimal investment decisions are detrimental to economic performance. All this evidence is consistent with the hypothesis that inequality of opportunity is harmful for growth and that higher social mobility has a positive impact on economic development.

Focusing on the inequality of opportunity, rather than inequality of outcomes, may also shed some light on the so far contrasting findings on the inequality-growth nexus (see e.g. Barro, 2000; Panizza, 2002; Banerjee and Duflo, 2003; Voitchovsky, 2005; Neves and Silva, 2014; Neves et al., 2016; Berg et al., 2018; Brueckner and Lederman, 2018; Van der Weide and Milanovic, 2018). This shift of focus to opportunities, which was already proposed by Rawls (1971), Sen (1980), Roemer (1998), among others, is materialized in the central message of the World Development Report 2006 (Bourguignon et al., 2007). Still, the empirical literature on the topic is rather scant. Ferreira et al. (2018), one of the few studies testing the opportunities-growth relationship, finds evidence that suggests a negative association between inequality of opportunity and growth in a cross-country analysis, though the findings are not robust. Likewise, Marrero and Rodríguez (2013) decompose the level of total inequality at the state-level in the US into inequality due to effort, and inequality due to opportunities. They consistently find that economic growth is positively related to the former, and negatively linked to the latter. Choosing social intergenerational mobility as an indicator of opportunity, some recent studies have shown a positive correlation between mobility and economic indicators, both between countries (e.g Neidhöfer et al., 2018; Aiyar and Ebeke, 2020) and within countries across geographical areas (e.g. Chetty et al., 2014; Fan et al., 2015; Bradbury and Triest, 2016; Güell et al., 2018; Aghion et al., 2019; Aydemir and Yazici, 2019). In this study, we are the first to exhaustively analyze the relationship between social mobility and economic performance going beyond merely describing geographical correlations. We complement the existing evidence by exploiting within-country (and within-region) variation based on a unique panel of Latin American regions, introducing a novel way of linking the intergenerational mobility of cohorts to year-level indicators of economic development.

3. Estimation Strategy

3.1 Social Mobility and Economic Development

To test the association between intergenerational mobility and economic development, we translate the conceptual framework discussed in Section 2. into a linear panel regression. Hereby, the units of analysis are the subnational regions and the time dimension is in years:

$$Y_{jct} = \upsilon_{jc} + \tau_{tc} + \delta M_{jct} + \xi X_{jct} + \theta Y_{jct-1} + \varepsilon_{jct} \tag{2.1}$$

In equation (2.1) Y is the level of economic development, measured for instance by log income per capita, of region j, which is located within the borders of country c, in year $t.M$ is our main variable of interest, which displays the degree of intergenerational mobility. This variable is measured as a weighted average of the degree of intergenerational mobility of people born from 1940 to 1989 living in region j, taking into account their participation in the economy in year t given their age. The exact weighting procedure is explained more exhaustively below in Section 3.2 X is a vector of control variables for regional characteristics in t, including controls for economic conditions, and average characteristics of the cohorts used to estimate social mobility. The model further includes one lag of the dependent variable, fixed effects for regions (v), and country-specific trends (τ), while ε is the error term. In Section 4. we describe the measurement and data sources for each variable more in detail. In different estimations, described and discussed in Section 6., the control variables, fixed effects, and the lag of the dependent variable are included gradually, such that in some estimations their coefficients are restricted to be equal to zero.

3.2 Weighting procedure

One fundamental challenge of linking social mobility to economic development is the temporal association of the two phenomena: while aggregate economic indicators are measured in particular years, an insightful indicator for intergenerational mobility should be measured at the birth cohort level. When the aim is to measure the impact of aggregate indicators - such as growth, income inequality, or public expenditures - on intergenerational mobility, one possible way is to estimate the association between the level of these aggregate outcomes that individuals experienced during their childhood and their future degree of intergenerational mobility (e.g. Mayer and Lopoo, 2008; Neidhöfer et al., 2018). However, this method is not feasible when the aim is to estimate the reverse, namely the impact of intergenerational mobility on aggregate economic outcomes. Indeed, most of the empirical literature overcomes this problem by taking averages of both measures across geographical areas, thus omitting the temporal dimension.

To go one step further in the direction of a proper measurement of the effect of social mobility on economic indicators, the aim is to find a strategy that accounts for the fact that, for reasons related to the life cycle, individuals born in different cohorts are at different stages of their individual contribution to the economy in each year. Neidhöfer et al. (2018) address this issue by arbitrarily choosing time lags of 30, 40, and 50 years to measure economic development when the individuals of each birth cohort were old enough to contribute substantially to the economic activity of the country. In this paper, we develop a novel weighting procedure that enables us to obtain more accurate estimates. The procedure associates the intergenerational mobility of individuals belonging to certain birth cohorts to the economic development of their region of residence by weighting their contribution to the economy in that particular year. This contribution is defined by the wage, experience, and labor market participation associated with the individual's stage of life in a given year.

We compute the weights by estimating cohort-participation profiles for each country in each year. The cohort-weights are constructed such that they sum up to one in every year. The cohort with the highest weight is the one with the highest contribution to the economy in that particular year, while cohorts with a weight equal to zero are not participating in the labor market because they are either too young or too old. In our main specification, these cohort-participation profiles represent the share of total wages earned by all individuals belonging to the respective birth cohort; i.e. $w_{bct} = \frac{\Omega_{bct}}{\sum_{b=1}^{b} \Omega_{bct}}$ where Ω is the sum of wages in year t of individuals residing in country c belonging to cohort b. To avoid potential correlation between the degree of intergenerational mobility of cohorts and their labor market participation affecting the construction of the weights, we define the participation profiles at the national level, rather than at the regional level, and normalize the weights to sum up to one in each year. Reassuringly, we do not observe any consistent pattern of correlation between the degree of mobility of a cohort and its weight across regions and over time.

For illustrative purposes, Figure 2.1 shows these participation profiles for all countries in three different years. To test the robustness of our results, we also compute the weights based on other definitions of cohort-participation rates: i) measured by the average wages of the cohorts w.r.t. the average national wages in each year; ii) defining a minimum share of 10% of contribution to total wages to get a non-zero weight and dividing the weights equally for every cohort satisfying this requirement; iii) defining a minimum share of 10% of contribution to total employment to get a nonzero weight and, again, dividing the weights equally for every cohort satisfying this requirement. Results of these additional exercises are included in the Online Supplementary material.[4] We observe that most cohorts show an active contribution to the economy in each year, while younger and older individuals have the lowest weights.

Following the procedure, M in equation (2.1) results in a weighted average of the intergenerational mobility of people born from 1940 to 1989:

$$M_{jct} = \sum_{b=1}^{B} w_{bct} m_{bcj} \tag{2.2}$$

Here, m_{bcj} is the degree of intergenerational mobility of individuals residing in j and belonging to cohort $b.w_{bct}$ is the weight measuring cohort b 's participation in the economy in t. The variation across years and regions in our estimations is then given by the interaction between the degree of intergenerational mobility and the cohort-participation weight. To measure intergenerational mobility we adopt several indicators, which we describe below in Section 4.2.

4. Data & Measurement

4.1 Data

To obtain our estimates of social mobility and economic development, we rely on 44 nationally representative household surveys from ten Latin American countries (Argentina, Brazil, Chile, Colombia, Ecuador, Guatemala, Mexico, Nicaragua, Panama, and Peru). Hereby, our selection criteria to include a country in our sample is the availability of at least one

[4]To further prove the consistency of our proposed method we also run a series of placebo tests where we use weighting schemes that do not relate at all to the cohort-participation profile in each year. The results of these placebo tests (included in the Online Supplementary material, Section E.8) show, reassuringly, no clear pattern of association between social mobility and economic development across specifications.

representative survey with retrospective questions on parental education and a sufficiently large sample size to enable a subdivision of the country into subnational regions. Using these surveys, we measure intergenerational mobility of people born from 1940 to 1989.

Then, we retrieve the surveys with the highest available quality for each country in our sample usually deriving from national statistical offices and not necessarily the same surveys used before to measure intergenerational mobility-to estimate different measures of economic development for the subnational regions of these countries from 1981 to 2018. We complement our analysis with, firstly, additional information on alternative local development indicators, such as night-time luminosity information from satellite data and, secondly, regional control variables on demographic characteristics. Thirdly, we incorporate historical data on GDP per capita, population size, and weather conditions retrieved from different data sources.

In what follows, we briefly describe the measurement of the two main variables studied in this analysis, social intergenerational mobility and economic development, and of the control variables, as well as the data employed to obtain the estimates. A more detailed description of the data sources for each country is included in the Online Supplementary material.

4.2 Social Mobility

The idea behind the measurement of social intergenerational mobility is to capture the likelihood of changes in the lifetime socioeconomic status of children with respect to their parents.[5]. Measuring socioeconomic status through appropriate proxy measures, such as permanent income, can be challenging, mainly because of data availability (Black et al., 2011; Jäntti and Jenkins, 2015).[6] Instead, information on the completed level of education

[5]Intergenerational mobility measures give meaningful insights on the stratification of societies and are closely related to the notion of equality of opportunity; both empirically and conceptually (Brunori et al., 2013).

[6]For instance, measures of income mobility may suffer from so-called life cycle bias if measured on few income spells for parents and children (e.g. Nybom and Stuhler, 2017).

of parents and children is, firstly, more likely to be available in households surveys, secondly, highly correlated with other measures using income or occupation (Blanden, 2013), and, thirdly, less affected by measurement error (Hertz, 2008). Hence, in our analysis, we focus on the education of individuals and their parents to measure intergenerational associations.

To measure m in equation (2.2), we estimate four different intergenerational mobility indicators: first, the slope coefficient of a linear regression of children's years of education on the years of education of their parents; second, a standardized measure of educational persistence; third, the probability of educational upward mobility; fourth, the relative risk of high school completion. We estimate these measures separately for individuals residing in different subnational regions and who were born in different birth cohorts, spanning 10 -year intervals.[7]

The slope coefficient is the most widely used mobility index in the intergenerational mobility literature. In our application, we regress the years of education y of an individual i on the years of education of the parent with the highest educational qualification y^p:[8]

$$y_i = \alpha + \beta \cdot y_i^p + \vartheta x_i + \varepsilon_i \tag{2.3}$$

x is a set of control variables for age and sex, and ε the error term. The regression coefficient β, the estimated value of which usually lies between zero and one, measures the degree of regression to the population mean between two generations. The higher is β, the stronger the association between parents' and children's education, and, hence, the lower is intergenerational mobility.

This measure of intergenerational mobility has the advantage of comparability between

[7]The correlation between these four measures for social intergenerational mobility is high but not perfect (see Section A. 2 of the Online Supplementary material). Hence, each captures different aspects of social mobility.

[8]Following the so-called "dominance principle", in all mobility measures we define parental education as the education of the parent in the household with the highest qualification. For individuals that indicated only the education of one parent, we use the available information.

countries, regions, and over time. However, it does not account for changes in the marginal distribution of years of education. To consider this, we estimate an indicator for the standardized persistence of education from parents to children:

$$\rho = \beta \frac{\sigma^p}{\sigma} \tag{2.4}$$

Here, σ and σ^p are the standard deviations of children's and parents' years of education, respectively.[9] Intuitively, both are indicators of relative mobility. While β mirrors the degree of association of one year of parental education with the education of their children, ρ measures this association in terms of one standard deviation.

We complement the analysis with two other indicators of social intergenerational mobility that instead of accounting for the entire distribution of years of education focus on an important threshold, namely high school completion. The first indicator, which we define as the probability of upward mobility, measures the likelihood of disadvantaged individuals - i.e. individuals whose parents both did not complete secondary education - to complete high school:

$$UM = \text{Prob}\left(y \geq s \mid y^p < s\right) \tag{2.5}$$

Here, y and y^p are defined as in the equations above and s is the number of regular years of education attached to the completion of secondary schooling. The higher this likelihood, the higher (absolute) intergenerational mobility.

Building on the probability of upward mobility we estimate also our last indicator for intergenerational mobility, namely the relative risk of high school completion:

[9] When no control variables are included in equation (2.3), ρ is equivalent to Pearson's correlation coefficient between y and y^p.

$$RR = \frac{\text{Prob}\,(y \geq s \mid y^p \geq s)}{\text{Prob}\,(y \geq s \mid y^p < s)} \qquad (2.6)$$

The relative risk of high school completion indicates how much more likely it is for the children of high-educated parents (i.e. parents with a completed secondary qualification or more) to complete high school compared to their peers with low-educated parents. The higher RR, the lower intergenerational mobility.

As mentioned before, to avoid co-residency bias we estimate all indicators using surveys that include retrospective information about parental education for each respondent (e.g. Emran et al., 2018).[10] Furthermore, since our aim is to only include individuals who are no longer enrolled in the education system, we restrict the sample to respondents that are older than 22.

Although the inclusion of retrospective questions is not common across Latin American household surveys, and we need enough large sample sizes to subdivide the sample within representative subnational regions and birth cohorts, we were able to obtain suitable data sets for 10 countries: Argentina, Brazil, Chile, Colombia, Ecuador, Guatemala, Mexico, Nicaragua, Panama, and Peru. Pooling all available survey waves we are able to estimate intergenerational mobility for five birth cohorts (1940-49, 1950-59, 1960-69, 1970-79, and 1980-89) in 52 regions. By using similar variable definitions and consistent data processing methods, the resulting statistics are comparable not only across countries and regions but also over time. Our final sample, including all countries and cohorts, comprises almost 1.2 million individuals.[11] In all of our micro-level estimations of intergenerational mobility,

[10]Neidhöfer et al. (2018) discuss potential selectivity issues that derive from non-responses to questions on parental education in household surveys. The analysis shows that although non-response (which is between 2% and 22% of respondents depending on the survey) might be systematic, i.e. respondents with lower education are slightly less likely to report their parents education, this does not affect significantly the sample mean and variance of years of education. Selective non-response could, if substantial, lead to upwardly biased intergenerational mobility estimates. However, the pattern is found to be the same in all countries, such that cross-country comparisons should keep their validity. For more information, see Neidhöfer et al. (2018), Online Appendix, Section 1.3.

[11]The surveys that we use for nine of the ten countries are nationally representative for urban and rural

we weight each observation by the inverse probability of selection provided by the survey, normalizing the weights over the different survey waves.

4.3 Regional Development

We collect data that enables us to estimate the level of economic development Y for each of the subnational regions in our sample. For the final analysis, we were able to construct an unbalanced panel of 52 regions for the period 1981 to 2018. National household surveys are our main data source for retrieving our estimates. When measuring economic development we are not forced to use household surveys that include retrospective questions about parental education. Hence, we use all available sub-nationally representative household surveys for the ten countries in our mobility sample. Since these surveys are not necessarily uniform in terms of geographical coverage and questionnaires across countries and over time, we process the surveys in order to harmonize the variable definitions, the subdivision in subnational units, and the measurement of economic development; i.e. we make the surveys comparable across countries and regions, and over time.[12]

In our baseline specification, the main indicator for the level of regional development is the average of household per capita income measured in purchase power parity (PPP). We estimate this aggregate measure with the household surveys mentioned above, adding up all individual labor and non-labor incomes reported during the last month within a household and dividing by the number of household members. Our second indicator of economic development is the population-weighted night-time luminosity of regions measured with satellite data. This indicator has been shown to be a consistent proxy for economic

areas. The survey that we use to measure intergenerational mobility in Argentina only includes urban areas (defined as localities with more than 2,000 inhabitants) covering 91.1% of the total Argentinian population (see Piovani and Salvia, 2018). More information on the employed surveys is included in Section A of the Online Supplementary material.

[12]We follow the methodology of the Socioeconomic Database for Latin America and the Caribbean (SEDLAC), a project jointly developed by CEDLAS at the Universidad Nacional de La Plata and the World Bank. For more information, see the project website.

growth (Henderson et al., 2012). We retrieve this data from Hodler and Raschky (2014). We also test our findings on a battery of further indicators for economic development: poverty, overall employment, labor formality, and access to water and electricity. All these indicators and their sources are described more exhaustively in the Online Supplementary material, Section B.

4.4 Control Variables

The vector X in equation (2.1) includes a set of control variables such that the uncovered patterns of association between social mobility and economic development are not spurious. The set of controls can be subdivided into three groups: i) year-level controls; ii) cohort-level controls; and iii) cohortspecific initial conditions.

Year-level controls

The first group of covariates includes income inequality in region j and year t, measured by the Gini index of disposable household per capita income, total regional population (polynomial of the second degree), and the urban share of the population. We estimate the first from household survey data and retrieve the two other from census data (their sources are described in the Online Supplementary material, Section C).

Cohort-level controls

The second group of covariates includes the cohort's average years of education and its variance, as well as the share of migrants. The average years of education are included to control for different levels of human capital accumulation, while its variance is used to control for differences in its allocation. These measures also control for the overall geographic sorting by skill level across regions (Diamond, 2016; Moretti, 2012). The share of migrants is included to control for migration from low mobility regions to high mobility regions that may bias our estimates (e.g. Ward, 2020). Including migrants could lead to upward bias of the estimates because of the positive selection of migrants (who, on average, have higher degrees of social intergenerational mobility) into regions with higher levels of economic development. Controlling for the weighted share of migrants in the cohort should correct for this bias. Furthermore, we test the robustness of our results by excluding migrants when estimating our mobility measures and obtain consistent results.[13] All variables are obtained from the surveys that we use to estimate intergenerational mobility, estimated at the cohort level, and weighted by the cohort-participation rate; exactly as the variable m in equation (2.2).

Cohort-specific initial conditions

The inclusion of the last group of controls aims to abstract from the potential effect of so-called initial conditions, i.e. the past development level of the economy that could have had both, an effect on social mobility, as well as on subsequent economic development (e.g. Johnson and Papageorgiou, 2020). In our empirical set-up, we are mostly interested in controlling for the conditions of the economy in the years when the individuals in our social mobility sample were born and grew up. Since historical data on economic conditions is not available at the regional level for Latin America, we approximate the initial conditions for the cohorts measured in each region (i.e. between 1940 and 1989 which are the years of birth of the individuals for whom we estimate social mobility) with five different indicators.

The first indicator is an estimate for regional GDP per capita from 1940 to 1989 that we obtain following three steps: First, using the first available household survey for each country

[13] For the purposes of this paper, an individual is considered a migrant if he or she was born in a different geographic area from his or her geographic area of residence (see Online Supplementary material, Section D). Chetty and Hendren (2018) evaluate the impacts of neighborhoods on intergenerational mobility and find heterogeneous effects depending on the age of children at the time of migration. However, we do not have information on the age of migration, which would allow us to consider this aspect in our analysis. The results excluding migrants, included in the Online Supplementary material, show the pure, but downward biased, local level effect of social mobility on development and can, thus, be considered a lower bound estimate.

we compute the share of regional income over total national income for each sub-national region. Then, we retrieve

country-level data on historical per capita GDP from the Maddison Project database (Bolt and van Zanden, 2020). Finally, assuming that the regional shares computed in the first step are constant over time, we multiply these shares with the historical country-level values for per capita GDP.

The second indicator for initial conditions is the child mortality rate around the year of birth of individuals. This variable controls for both, parental investments in children and the environment in which these investments take place. The idea behind this is inspired by the so-called quantity-quality model of fertility; i.e. the characterization of the trade-off in the choice between the number of children and the amount invested in the education of each child (Becker and Lewis, 1973). Under consideration of the quantity-quality trade-off, the degree of infant mortality mirrors the probability that individuals grow up in households with more or less children, and thus, ceteris paribus, their chances of receiving a higher or lower amount of investment in education. Negative shocks to infant mortality, for instance, due to medical and pharmaceutical advances, could thus lead to an increased number of children per family, and result in a lower investment in the education of each child. Additionally, high levels of infant mortality could also reflect adverse environmental conditions experienced while in-utero or in early childhood, such as natural catastrophes or epidemics, that may have a direct effect on mortality, future health, and cognitive capacities of survivors and, thus, on economic growth (e.g. Almond, 2006; Caruso and Miller, 2015).

The regional population from 1940 to 1989 is our third indicator. The inclusion of this variable is motivated by the literature relating population growth to economic growth (e.g. Headey and Hodge, 2009). The fourth and fifth indicators capture the regional weather conditions from 1940 to 1989 retrieved from National Oceanic and Atmospheric Administration, measured by the average air temperature and the average precipitation. As has been shown by past research, early-life weather conditions may have a persistent effect on future health, schooling, and socioeconomic outcomes (e.g. Maccini and Yang, 2009) as well as on economic development (e.g. Dell et al., 2012). Since all these variables are measured in the years associated with the birth cohorts, the same weighting procedure explained in Section 3. is applied to them. To account for non-linear interactions, the variables for population, temperature, and precipitation are included as a polynomial of the second degree.

5. Geography of Intergenerational Mobility in Latin America

In this section, we characterize the variation of intergenerational social mobility across the 52 subnational regions we constructed for Latin America. Our goal in this section is to provide a first detailed spatial picture of the extent to which children's education is related to their parental educational background. This analysis is relevant since it allows to identify regions with less social progress.[14]

As a first approach, Figure 2.2 maps the geography of social intergenerational mobility in Latin America for three cohorts. Interestingly, two main spatial patterns emerge: First, social mobility varies significantly across countries. The high levels of social mobility found in the south of South America (primarily Chile and Argentina) contrast with lower levels in the Northern part of the region, including Mexico and Central American countries. Second, there is also a substantial variation within countries. For instance, the south of Chile presents

[14]Munoz (2021) estimates intergenerational mobility of education across Latin American provinces using cohabitation samples from census data. Since the estimates are relying on parents and children cohabiting in the same household, and hence a sample of older individuals is likely to suffer from co-residency bias (e.g. Emran et al., 2018), the analysis mostly focuses on the probability to complete primary education of younger individuals, following Alesina et al. (2021). This dimension is, actually, important for older cohorts of Latin American residents, but less relevant for younger cohorts because of the expansion of secondary education in recent decades (e.g. Levy and Schady, 2013). Indeed, changes in returns to education just above and below high school completion are closely related to the changes in inequality experienced in the region (López-Calva and Lustig, 2010).

low upward mobility compared to the north of the country. In turn, the northern regions of Brazil shows considerably lower levels of mobility relative to the south. These findings complement previous country-level studies which show that intergenerational mobility is rising in Latin America (e.g. Neidhöfer et al., 2018). We provide evidence suggesting that this trend reached almost every sub-national region, but with a high degree of heterogeneity between and within countries.[15]

To emphasize the relevance of within-country variation, Figure 2.3 shows the distribution of different measures of social mobility for each country and its regions. The country-level values can reasonably give a general picture of social mobility in Latin America. However, most of the countrylevels estimates are not a sufficient summary of the heterogeneity within countries. For instance, Ecuador, Nicaragua, and Panama have levels of intergenerational persistence above the Latin American average (i.e., lower social mobility), while many of their sub-regions reach substantially lower levels, comparable to the most socially mobile countries (Argentina and Chile). This heterogeneity is also visible in Figure 2.4, which shows the 10% regions with the highest and lowest levels of intergenerational mobility.

Figure 2.5 plots the evolution of social mobility measures for regional level (grey) and countrylevel (black) estimates by comparing individuals belonging to the first two cohorts of our analysis (1940-1949) with people born in the last two (1980-1989). As is evident, Latin Americans benefited differently from the development of social mobility over time, even considering areas within the same country. Estimates over the 45 -degree line imply that intergenerational mobility did not change over the time period considered here. On the other hand, estimates reveal improvements in social mobility when they are on the right of

[15]Note that these estimates are merely descriptive and do not consider, so far, the role of migration in shaping intergenerational mobility patterns. The level of intergenerational mobility of a region is measured on a sample including all residents of that region. Since the intention of this part of the analysis is to give a descriptive overall picture on the geography of intergenerational mobility in Latin America we abstain from excluding migrants here. However, when measuring the impact of intergenerational mobility on economic development in the next sections we do take this important dimension into account, including appropriate control variables and testing the robustness of our results.

the 45 -degree line for intergenerational persistence, standardized persistence, and risk ratio, and on the left for the probability of upward mobility. In general, intergenerational mobility is rising in our sample of Latin American countries at both the regional and national level. For instance, while in all countries the chance of upward mobility for people born in 1940-49 with low-educated parents is less than 50%, the chances of people born in 1980-89 in many regions are significantly higher. However, substantial heterogeneity remains regarding both the degree of mobility as well as its evolution over time. In particular, the dispersion of social mobility across regions for younger cohorts is much less prominent than it was in past.

6. Social Mobility and Economic Development

In this section we report the results of our empirical analysis testing the relationship between social mobility and economic development.[16] First, in 6.1 , we present the results on the relationship between social mobility and economic development measured by regional income per capita in levels. Then, in 6.2, we show the estimates obtained by including lags of the dependent variable, which indicate the relationship between social mobility and economic growth. In 6.3 we then investigate the association between social mobility and other indicators of economic development, such as nighttime luminosity and poverty rates. We discuss the strengths and limitations of our results in 6.4. Finally, we provide additional evidence on the accumulation vs. allocation of human capital hypothesis in 6.5 , and on the mobility-inequality nexus in 6.6.

[16]Throughout this section, we present the results weighting social mobility measures using the aggregated cohortparticipation profiles. All the results presented here are robust to the utilization of the other alternatives of cohort weights described in Section 3.. These additional results are shown in Section E of the Online Supplementary material.

6.1 Economic Development

As a first approximation, Figure 2.6 plots the averages over the entire time period of all four measures of social intergenerational mobility described in Section 4.2 and log average household per capita income. This first stylized analysis shows a clear and robust positive (negative) correlation between intergenerational mobility (persistence) and economic development, both across countries as well as across regions.

Table 2.1 presents the results of estimating equation (2.1) using the slope coefficient to measure intergenerational mobility (M) and average household per capita income as indicator of economic development (Y). So far, these estimates are obtained without including lags of the dependent variable. Recall that the slope coefficient is a measure of persistence; it shows the degree of association of one year of parental schooling with the years of schooling of their children. The higher this coefficient, the lower intergenerational mobility. Hence, a negative regression coefficient of M in Table 2.1 indicates higher intergenerational persistence (i.e. lower intergenerational mobility) is associated with lower average per capita income.[17] To allow a more straightforward interpretation of the coefficients, all variables are included as logarithms in the estimations. Robust standard errors are obtained by clustering at the country-year level to account for serial correlation of the error term within countries.[18] The significance of the point estimates is consistent with the main analysis if we cluster standard errors by countries, or regions.

We gradually include the control variables described in Section 4.4. In column (3) we first include the year-level covariates mentioned in Section 4.4 and in column (4) the cohort-level variables described in Section 4.4, which are aimed to abstract from contemporaneous and past characteristics related to social mobility that could influence economic development.

[17] The same applies for the standardized persistence (ρ) and the relative risk of high school completion (RR). For the probability of upward mobility (UM) a positive coefficient indicates that higher mobility is associated with economic development.

[18] Results with bootstrapped standard errors are included in the Online Supplementary material.

Among these, the second set of controls includes information on migration and human capital accumulation, measured by the share of migrants in the cohort and the average and variance of years of schooling of the cohort. In column (5), results are obtained controlling for cohort-specific initial conditions, i.e. the economic conditions during the formative childhood years of individuals in our social mobility sample. These controls are necessary as said conditions could have had an effect on both social mobility as well as subsequent economic development. These include past GDP per capita, child mortality, population, temperature

and precipitation from 1940 to 1989 (see Section 4.4). All variables at the cohort-level are weighted adopting the cohort-participation profiles explained in Section 3.2. Models including lags of the dependent variable are reported and discussed in Section 6.2.

The results show that in all estimations the coefficient of M, measured by the slope coefficient, is negative and highly significant. Hence, social mobility is consistently associated with economic development. These findings hold when controlling for i) unobserved heterogeneity by including region and time fixed effects, ii) potential mediators (cross-sectional inequality, share of migrants, average education, and cohort-specific initial conditions), iii) spillover effects between regions in the same country, and iv) country-specific time trends (country-by-time fixed effects).[19] On average, a 10% increase in intergenerational mobility, measured by the slope coefficient, is associated with a rise in per capita income by 17%.[20] To give benchmarks for this estimate, intergenerational mobility in education measured by the slope coefficient rose in Latin America, on average, by 4% from one four-year-cohort to the next between 1940 and 1991 , and by 12% for people born at the end of the 70 s with respect to people born at the beginning of the 60 s.[21]

[19]Spillover effects are controlled by including the average degree of intergenerational persistence in year t of all other regions $-j$ in the country (i.e. region j is excluded to estimate this average).

[20]The results obtained using the other measures of mobility described in Section 4.2 confirm these findings. The average effect over all mobility measures is around 12%. In terms of standard deviations, the effect size is also similar across specifications and mobility indicators: a one standard deviation increase in mobility is associated with an income per capita increase of around 0.5-1 standard deviations. All additional results tables, including several robustness checks, can be found in the Online Supplementary material, Section E.

[21]These estimates are obtained from the Mobility-Latam Data at https://mobilitylatam.website (see

Among the covariates included in the models, income inequality deserves a special mention. Its coefficient in most specifications shows that, controlling for the degree of intergenerational mobility, inequality is positively associated with economic development. However, the interaction between social mobility and cross-sectional income inequality in column (8) has a negative sign, meaning that low social mobility is particularly detrimental when income inequality is high. We will analyze the relationship between social mobility and inequality separately in Section 6.6.

6.2 Economic growth

We also analyze the relationship between social mobility and economic growth, rather than economic development measured in income levels. equation (2.1) can be reformulated as a growth equation

$$GY_{jct} = v_{jc} + \tau_{tc} + \delta M_{jct} + \xi X_{jct} + \tilde{\theta} Y_{jct-1} + \varepsilon_{jct} \tag{2.7}$$

where $GY_{jct} = Y_{jct} - Y_{jct-1}$ is the logarithmic growth rate of regional income per capita. The only difference between ((2.7)) and the baseline equation (2.1) is the interpretation of the coefficient of the lagged dependent variable, $\tilde{\theta} = \theta - 1$ (Durlauf et al., 2005). With this in mind, we estimate equation (2.1) including one lag of regional income per capita among the set of covariates. Table 2.2 shows the results of the estimations.

Column (1) of Table 2.2 shows the estimates obtained by OLS regressions omitting the fixed effects v_{jc} and τ_{tc}. The coefficient of M is negative and significantly different from zero, suggesting that social mobility is positively associated with economic growth. Also, the requirement for conditional convergence is fulfilled, which is $\tilde{\theta} < 0$ or, respectively, $\theta < 1$. However, this OLS estimate may be biased because of the potential correlation between

Neidhöfer et al., 2018).

lagged income and the error term. Hence, in the next columns, we gradually include fixed effects in the model. Column (2) includes region fixed effects, column (3) region and year fixed effects, and column (4) region and country-by-year fixed effects. In these models, equivalently to most specifications in Table 2.1, the variation in the degree of social mobility within-regions explains the variation in economic growth. In all fixed effects estimations, the coefficient of M is consistently negative and significant, while conditional convergence still holds.

As shown by Nickell (1981), within-group estimates of dynamic panel data models such as equation (2.1) and (2.7) relying on a low number of observations over time (small T) may be seriously biased. If the number of observations over time included in our panel can be considered to be relatively high ($T = 28$ on average across regions, with a maximum of 36 time periods), the fixed-effects model that we estimate should provide consistent results (Roodman, 2009). However, following Marrero and Rodríguez (2013), among others, we also account for region-specific dynamics by estimating the models using dynamic panel data methods, and implement a system GMM-estimator (Arellano and Bover, 1995; Blundell and Bond, 1998). This approach is based on the use of lagged levels of the regressors as instruments. We employ the one-step system GMM estimator and consider robust standard errors. We use all available lags of income per capita $> t - 2$ and limit the number of instruments by collapsing the instrument set (Roodman, 2009). The results of this application are shown in columns (5) to (8) of Table 2.2. For transparency, we show the estimates for the same specifications as in columns (1)-(4). Below the estimates in these columns, we also report the p-value of the Hansen test of over-identifying restrictions, and the two Arellano-Bond tests of autocorrelation. A significant Hansen statistics suggests that the set of instruments is not valid, while absence of autocorrelation in the Arellano-Bond test requires that the AR(1) test rejects the null hypothesis ($\mathrm{p} < 0.1$) while the AR(2) test does not ($p > 0.1$).

Generally, the results are consistent with the main analysis: Social mobility is negatively associated with economic growth when considering region-specific dynamics. The coefficients obtained by applying System-GMM, shown in columns (5)-(8) of Table 2.2, are in the same order of magnitude and are not significantly different from the OLS estimates, shown in columns (1)-(4) of the same table. Likewise, the coefficient of the income lag consistently points at conditional convergence. In column (5), which is the specification that does not include fixed effects, the validity of the instruments is not rejected by the Hansen test and the coefficient of M is negative, but not statistically significant. In columns (6) and (7), the coefficient of M is negative and significant, but the Hansen test suggests that the set of instruments is not valid. Finally, in column (8), which is our preferred specification because it properly controls for country-specific heterogeneity in income trends, the coefficient of M is statistically significant and the p-value of the Hansen test suggests that the set of instruments is valid. Altogether, the estimates suggest that social mobility is positively associated with future economic growth.[22]

6.3 Different Dimensions of Development

We test whether the positive association between social mobility, income per capita and economic growth also extends to other dimensions of economic development. Table 2.3 presents the estimated coefficient of social mobility M in equation (2.1) for different variables as indicators of economic development Y. These estimations include the full set of control variables described in Section 4.4, region and country-by-time fixed effects, and spillover effects. The results show that the positive relationship between social mobility and economic development is robust to considering different indicators, namely the log of average nighttime lights per pixel (i.e. luminosity), poverty (headcount ratio at 1USD a day), total employment,

[22]The results we obtain with the other social mobility measures are mainly consistent with the baseline analysis, although the statistical significance of the coefficients and the validity of the instrument set sometime varies. However, qualitatively, all estimates suggest the same pattern: social mobility is positively associated with economic development and growth. The additional System GMM estimates obtained with the other indicators are included in the Online Supplementary material, Section E.

labor formality, and houses with access to water and electricity. A 10% decrease in the slope coefficient (i.e. an increase in social intergenerational mobility)

is associated with a 8% stronger luminosity, 25% less poverty, 8% more employment, 5% more labor formality, and 8% and 2% higher share of houses with access to water and electricity, respectively.[23]

6.4 Discussion of the Results

Although the exact identification of the effect of improving social mobility on economic performance is empirically challenging, and we cannot completely exclude that other sources of unobserved heterogeneity not considered here may bias our results, these new estimates allow us to make an important step toward understanding the relationship between social mobility and economic development.

First, the results presented above show that the positive and significant association between social mobility and economic development (and growth) is not explained by confounding factors such as migration, human capital accumulation, contemporaneous income inequality, and the initial conditions of the economy; i.e. the persistent effect of regional economic development in the past (1940 to 1989—which represents the circumstances faced during the formative years of the individuals in our sample) on present economic development.

Second, we perform the analysis within subnational regions over time. The inclusion of region and time fixed effects, and even country-specific time trends in some estimations, warrants that our estimates account for unobserved heterogeneity that could drive the results, for instance due to the role of culture and institutions as drivers of economic development. In addition, we also take into account region-specific dynamics affecting growth and development by estimating dynamic panel data models, which provide results that are consistent

[23] The coefficient of the last parameter is not statistically significant.

with the main analysis.

Third, given the structure of our data and the construction of our variable for social mobility through the weighting procedure explained in Section 3., the association that we measure relates past mobility with future economic development. Due to the applied cohort-participation profiles methodology, at the point in time when economic development is measured the individuals for whom mobility is estimated have already completed their educational careers. Further, we control for the past level of development—i.e. the cohort-specific initial conditions of the economy (including past GDP per capita, child mortality, population, temperature and precipitation in the period 1940-1989—which assures that the uncovered relationship between social mobility and economic development is not spuriously driven by the past level of development that is correlated with both, social mobility and future development. Hence, the estimated correlation is not affected by a feedback effect resulting in reverse causality.

Finally, all results hold when considering different dimensions of economic development, and the significance of the correlation is robust to the consideration of different measures of intergenerational mobility, when measuring the degree of intergenerational mobility of men and women separately, and to the exclusion of migrants (see Additional Results in the Online Supplementary material).[24]

We conclude that these findings allow us to make a step forward toward understanding the relationship between social mobility and economic development. As mentioned in Section 2., the theoretical mechanism behind this relationship is that higher social mobility results in a better allocation of talent, thus, improving the overall productivity of the population

[24]Generally, the intergenerational mobility of men and women, and of migrants and non-migrants, are highly correlated across regions and cohorts. The mobility of each population subgroup is likely to be influenced by the shared overall equality of opportunity-enhancing environment and, since all subgroups participate to the economy, we do not expect substantial differences in the estimated relationship between each subgroup level of mobility and economic development. Interestingly, the point estimates showing the association between the social mobility of men and economic development are stronger then the estimates for women. This is in line with a lower labor market participation-both at the extensive and intensive margin-of women in most Latin American countries.

in the labor force. Less inequality of opportunity in the process of human capital formation enables individuals from households in the lower end of the income distribution to translate their talent and abilities into human capital. As a consequence, assuming a constant distribution of innate abilities, the pool of talent in the labor force increases, and the allocation of individuals to occupations depends more on individual's skills than on socioeconomic background. With low levels of social mobility, economic development is negatively affected by the misallocation of talent which is prevalent in the society. Since the frequency of our data is annual, the estimates obtained by including region fixed effects show that in a year when social mobility (i.e. the weighted average social mobility across all cohorts) is higher than average, economic development indicators show a higher-than-average performance. Hence, the positive association between mobility and development is mainly driven by cohorts of individuals that had better opportunities to develop their talent entering the labor market or gaining more experience and lower employment shares among cohorts of individuals that faced lower equality of opportunity and a stronger misallocation of talent.

6.5 Accumulation vs. Allocation

After having shown that social mobility is consistently and positively associated with economic development, and that this relationship is robust, we further test whether the main driver of this relationship is the accumulation of human capital or its allocation. Generally, a stronger accumulation of human capital and lower social mobility could coexist, for instance when it is mostly the children of higheducated parents who benefit from educational expansions. In the regressions presented thus far, we controlled for the average years of education to avoid bias in our estimates capturing the "trickledown-effect" of this type of accumulation (at the top of the distribution) on economic development, instead of the impact of social mobility and equality of opportunity. Furthermore, the fact that also measures of relative mobility-such as the standardized persistence and the relative risk of high school completion-yield consistent results on the positive relationship between social mobility and economic development provides suggestive evidence in favor of the allocation-hypothesis. In this section, we further test this assumption including both the degree of upward mobility from the bottom, and the degree of persistence at the top. The results of this exercise are shown in Table 2.4.

The regression estimates in column (1) of Table 2.4 are obtained including the full set of control variables with the exception of average years of education. The coefficient of upward mobility, i.e. the likelihood of completing secondary education for the children of low-educated parents, is positively and significantly associated with economic development. The same applies to the degree of top persistence, i.e. the likelihood of completing secondary education for the children of high-educated parents, which is highly correlated with the degree of upward mobility from the bottom since secondary school expansions benefited most of the population in Latin American countries. However, when including the degree of upward mobility in column (4) and (5), the coefficient of top persistence becomes very small in size and statistically indistinguishable from zero. In contrast, the level of upward mobility is consistently, significantly, and substantially associated with economic development.

These estimates confirm that it is not only the overall accumulation of human capital that positively affects economic development, but also in which part of the distribution this accumulation takes place is important. Reduced inequality of opportunity implies a higher level of human capital accumulation for children from disadvantaged families leading to a more efficient allocation of talent, and, hence, to improved aggregate economic performance. A higher level of accumulation that only benefits advantaged families may have no direct effect on development.

6.6 *Mobility and Inequality*

As a final exercise, we estimate the relationship between social mobility and income inequality. This relationship has attracted special attention by researchers and policy makers since descriptive evidence suggests that countries with high levels of income inequality also have low degrees of intergenerational mobility. A graph showing this relationship became very famous under the name Great Gatsby Curve (Corak, 2013).

Economic theory, indeed, suggests the existence of a negative correlation between inequality and social mobility (e.g. Becker and Tomes, 1979; Loury, 1981; Galor and Zeira, 1993; Owen and Weil, 1998; Maoz and Moav, 1999; Hassler et al., 2007). The main mechanism hypothesized to be behind this relationship is inequality in investment in human capital: since parents invest one part of their income in the human capital of their children, a higher degree of income inequality leads to a higher dispersion of parental investments. Hence, the human capital of children from families in the upper part of the distribution rises at a higher rate than the human capital of children from families with less resources and, as a consequence, social mobility decreases. Neidhöfer (2019) tests this side of the relationship and finds that, indeed, higher levels of inequality experienced during childhood and adolescence by children of low-educated parents are associated with lower levels of mobility measured in adulthood. Our data allows us to investigate the other side of the relationship, namely the effect of intergenerational mobility on future cross-sectional income inequality. The proposed mechanism driving this relationship is straightforward: higher inequality of opportunity to invest in human capital, which mirrors a lower degree of social mobility, leads to higher levels of income inequality in the future.

We follow the same approach as before and estimate the partial correlation between social mobility in year t, obtained by weighting the mobility of each cohort by their cohort-participation profile, and income inequality in t. Figure 2.7 plots the unconditional relationship between our four indicators for social mobility and income inequality (i.e. the Great Gatsby Curve), measured by the Gini coefficient of disposable household income per capita. Table 2.5 shows estimates obtained via linear regressions subsequently including the control variables described above. All results are consistent with the hypothesis that lower levels of social mobility, and hence higher inequality of opportunity, are associated with higher levels of future income inequality.

7. Conclusions

In this paper, we explored the relationship between social intergenerational mobility and economic development constructing a new panel data set including 52 regions of 10 Latin American countries. For these regions, we estimate the degree of intergenerational mobility of people born between 1940 and 1989, and aggregate measures of economic development from 1981 to 2018. These are linked using a new weighting procedure that we develop to account for the relative participation of the cohorts in the economy in each year. Our results show a positive, significant, and robust association between increasing social mobility and the economic development of Latin American regions.

To the best of our knowledge, this paper represents the first large scale study on the role of social mobility on economic development and contributes to our understanding of the nexus between inequality and economic growth. Our findings suggest the non-existence of the equity-efficiency trade-off regarding social mobility. Conversely, they suggest that improving equality of opportunity generates positive economic returns. Our analysis provides evidence for the robustness of this positive association and shows that it is not driven by confounders such as migration, human capital accumulation, and initial development conditions. Although a clear causal identification of the relationship is challenging, our empirical set-up makes a decisive step forward. In addition, the cohort-participation profiles methodology that we propose should also be suitable for a more thorough evaluation of the relationship between human capital, measured by education, and growth. This new methodology represents a valuable contribution to this branch of the literature, which thus far has mainly focused on contemporary (or lagged) relationships between the average education of the working age population and economic growth.

Our findings are also relevant for the evaluation of the effectiveness of market interventions. Arguably, interventions aimed at improving equality of opportunity may create distortions and thus lead to inefficiency in the short-run. However, if these interventions are indeed able to contribute to better opportunities and less misallocation of talent, they should simultaneously contribute to increased efficiency in the long run. Consequently, both effects could possibly outweigh each other and change the terms of the trade-off. For the sake of sustainable policy decisions, these long-run considerations should be taken into account to evaluate the effectiveness of policy measures in the future.

Finally, our analysis also contributes to the literature on the geography of intergenerational mobility (e.g. Alesina et al., 2021; Chetty et al., 2014; Corak, 2020; Güell et al., 2018) by providing the first geographical trends for 52 sub-national regions in Latin America. Our findings show that there is considerable variation among sub-national regions in both intergenerational mobility and economic development, even within countries. Since previous country-level estimations showed that Latin America is a region with strong intergenerational persistence (e.g. Torche, 2014; Neidhöfer et al., 2018), these new findings contribute to the overall understanding that country-wide patterns obscure within-country heterogeneity.

8. Figures

Figure 2.1: Cohort-participation profiles. Aggregated cohort participation rate

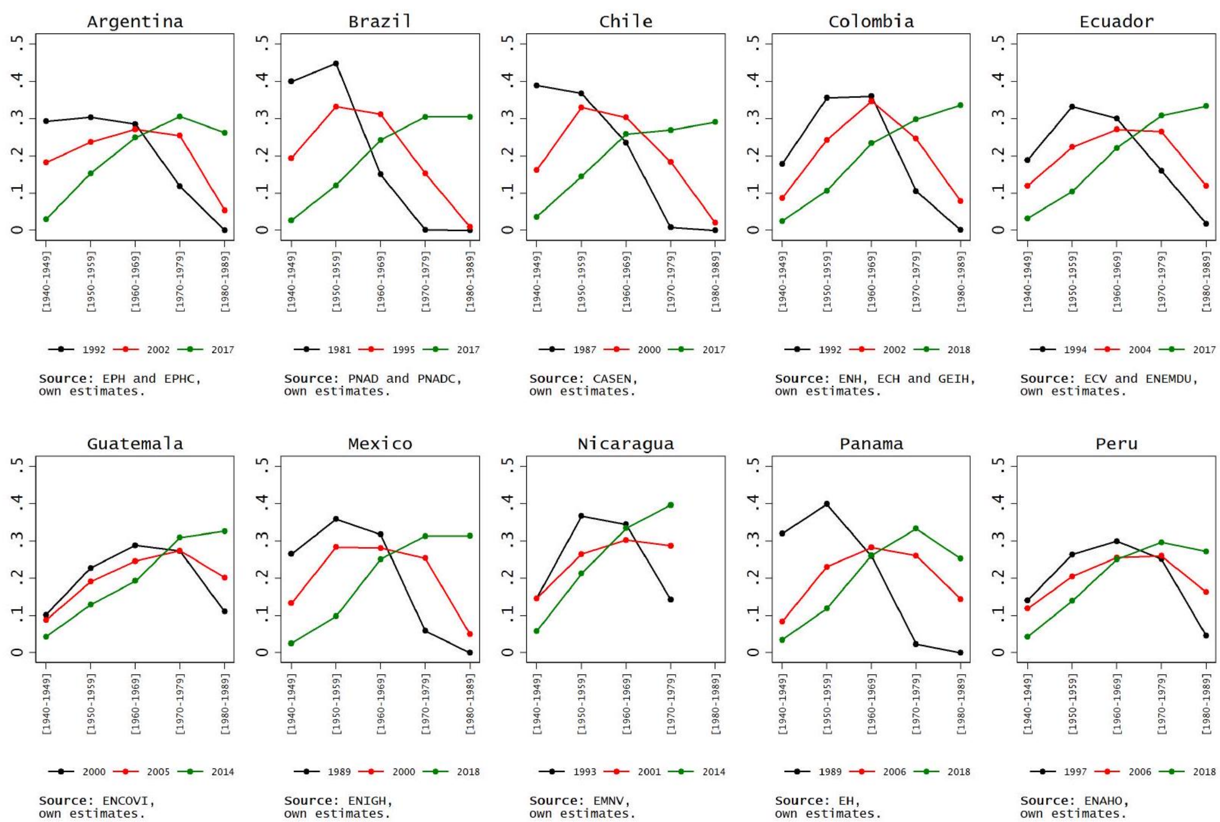

Source: National Household Surveys, own estimates.

Figure 2.2: The geography of social mobility in Latin America

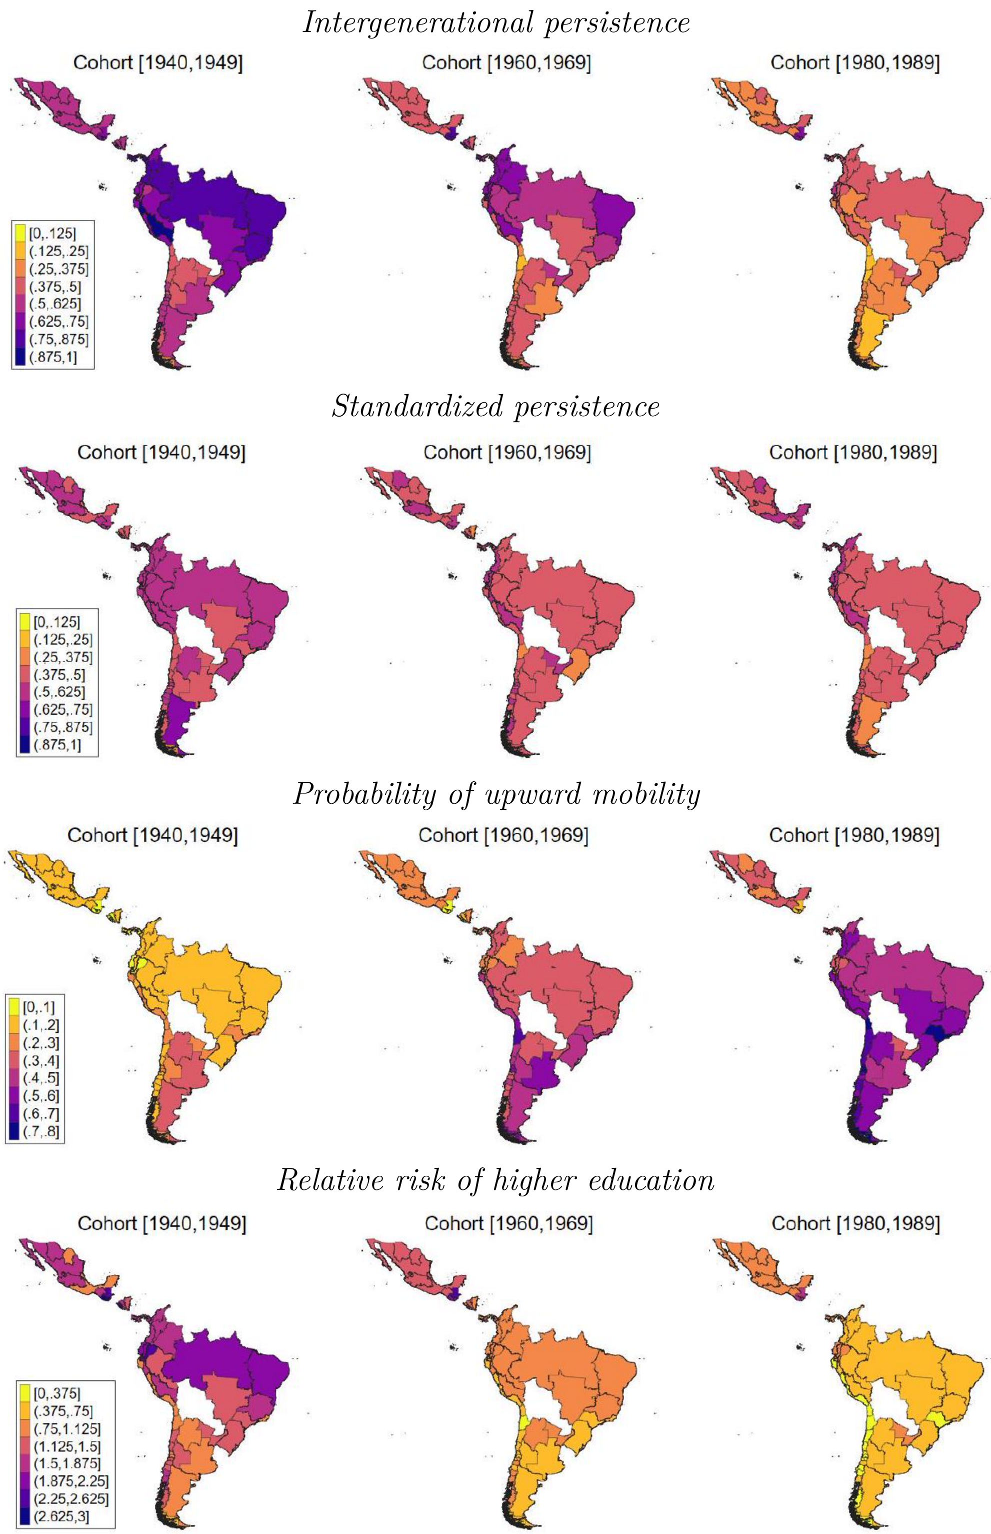

Source: National Household Surveys, own estimates.

Figure 2.3: Comparison of social mobility at national and sub-national level

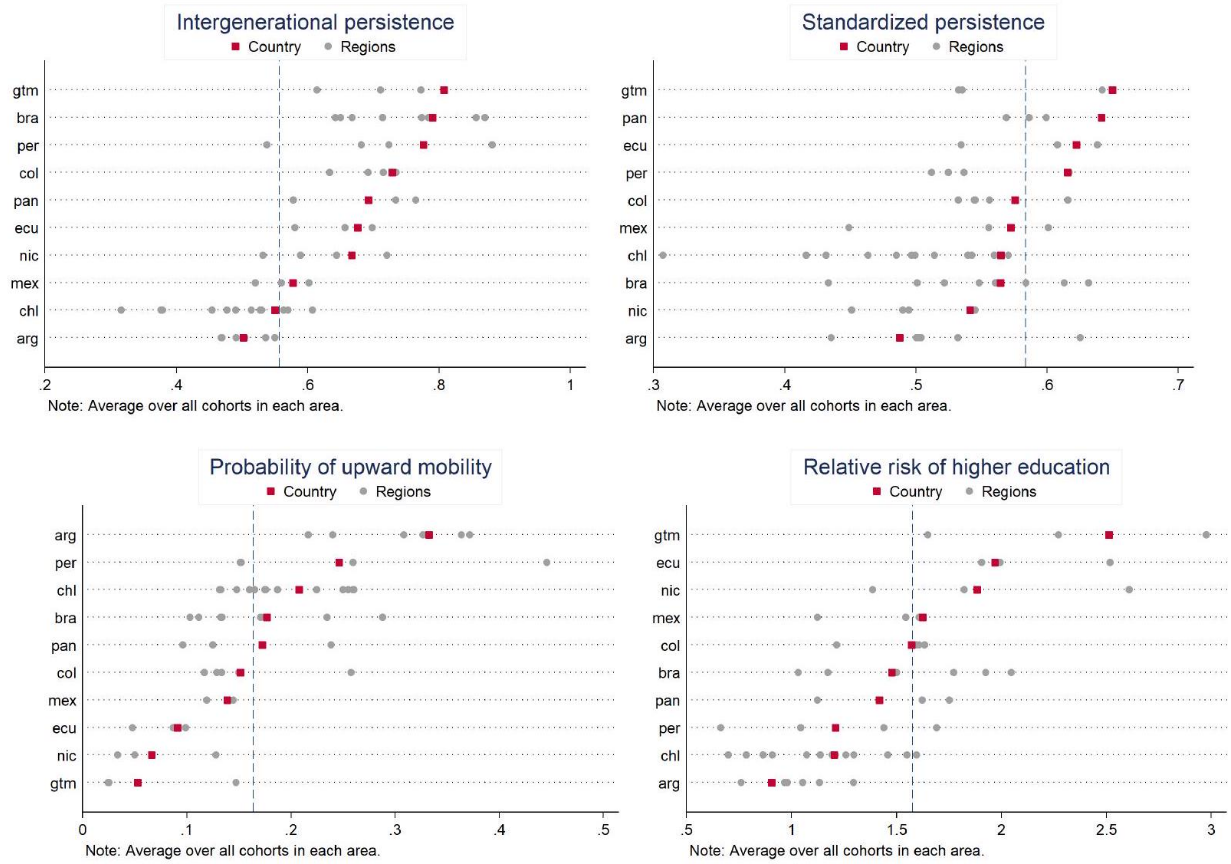

Source: National Household Surveys, own estimates.

Figure 2.4: Rankings of social mobility across Latin American sub-national regions

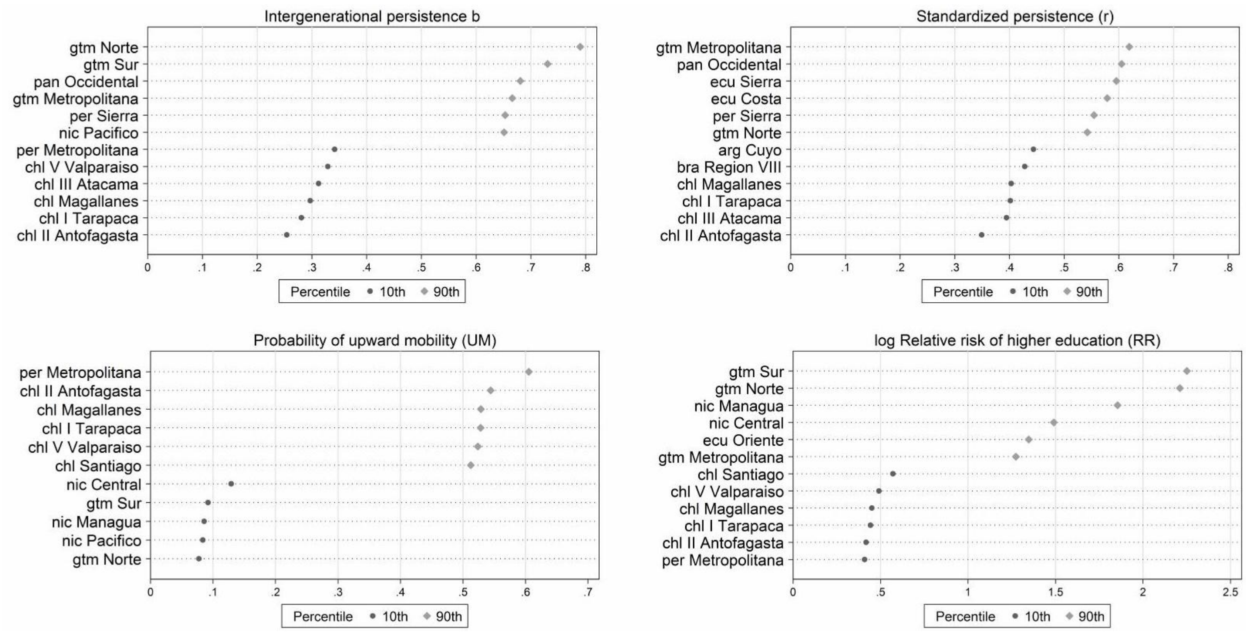

Source: National Household Surveys, own estimates.

Figure 2.5: Evolution of social mobility in Latin American regions

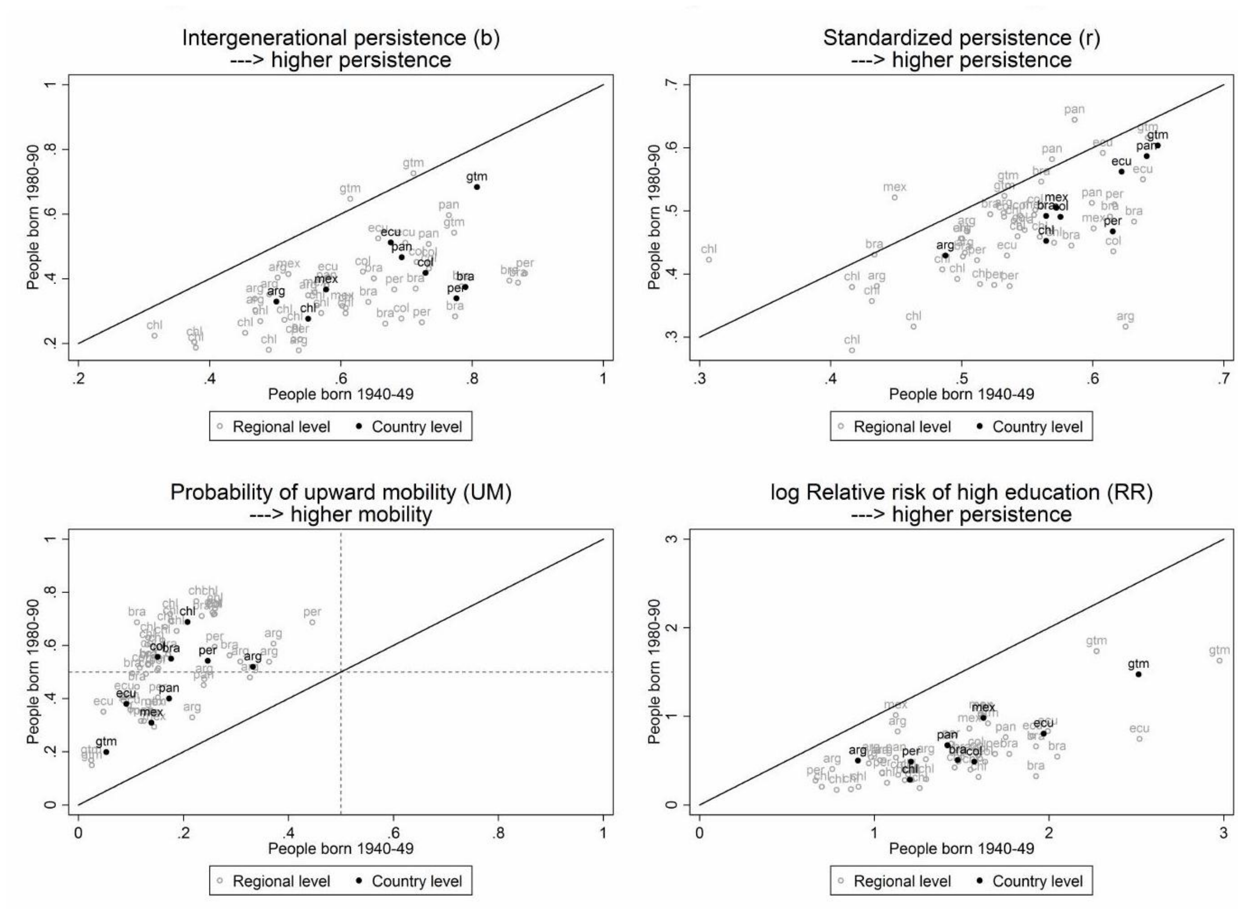

Source: National Household Surveys, own estimates.

Figure 2.6: Social mobility and economic development. Unconditional relationship.

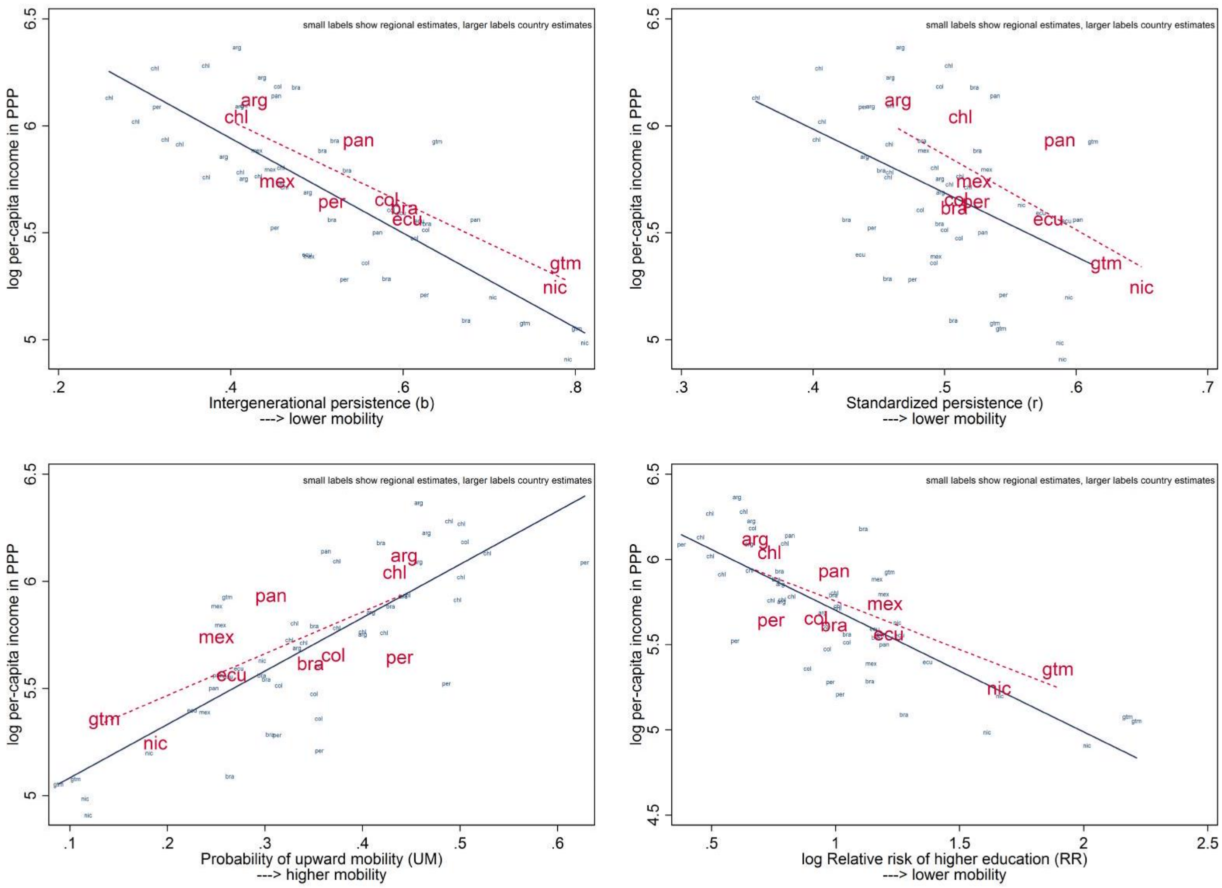

Source: National Household Surveys, own estimates.

Figure 2.7: Social mobility and income inequality. Unconditional relationship.

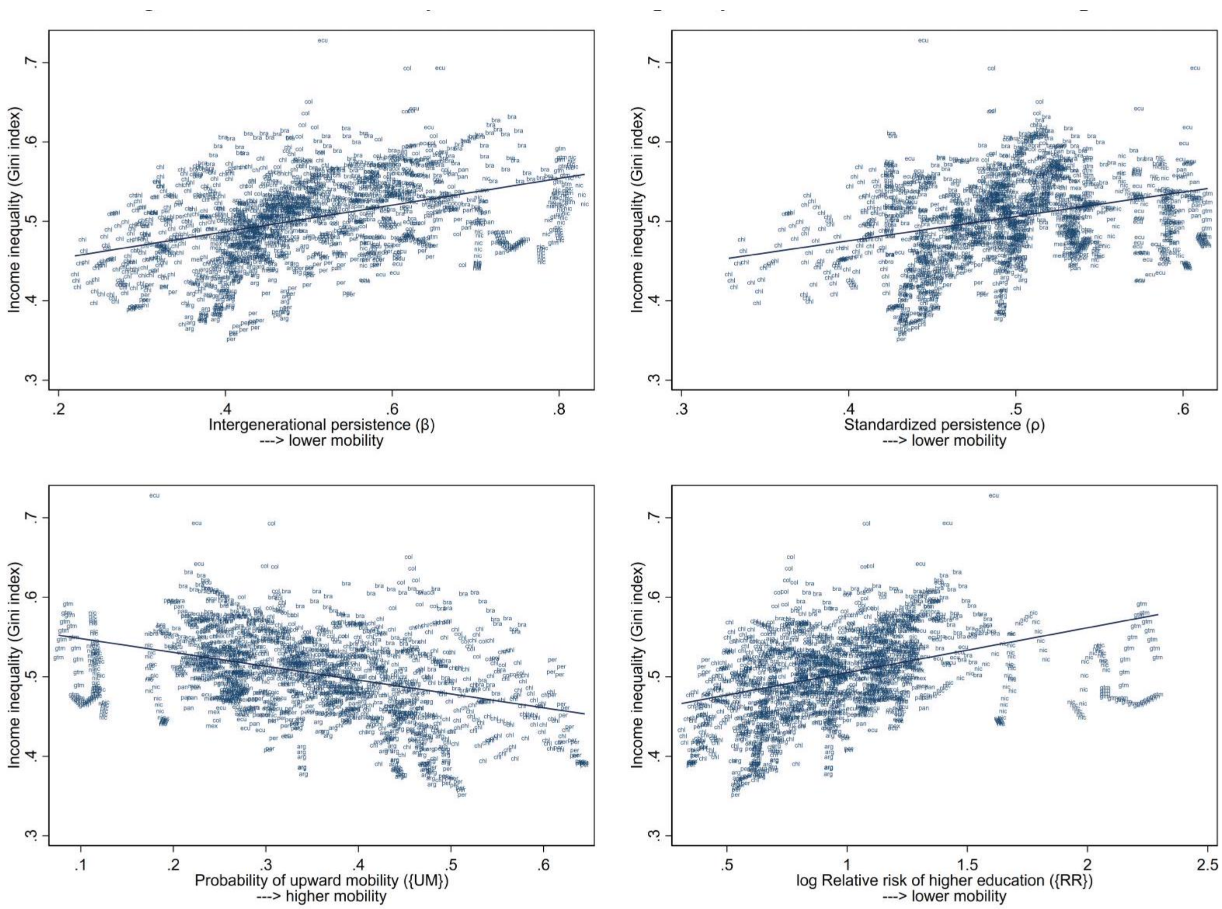

Source: National Household Surveys, own estimates.

9. Tables

Table 2.1: Estimates on social mobility and economic development. Intergenerational persistence β

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
M(w)	-1.268^{***} (0.0638)	-1.292^{***} (0.230)	-1.506^{***} (0.243)	-2.012^{***} (0.268)	-2.032^{***} (0.216)	-1.967^{***} (0.228)	-1.593^{***} (0.305)	-2.645^{***} (0.303)
M (w) × Inequality (Gini)								-1.409^{***} (0.192)
Year-level Controls								
Inequality (Gini)			0.356^{**} (0.158)	0.456^{***} (0.156)	0.498^{***} (0.167)	0.512^{***} (0.155)	0.746^{***} (0.0823)	-0.453^{***} (0.165)
Urban Population			0.187 (0.131)	-0.0155 (0.130)	-0.131 (0.136)	-0.0588 (0.130)	-0.137 (0.0937)	-0.230^{***} (0.0803)
Population			-0.918 (0.647)	-0.329 (0.528)	-0.0659 (0.689)	0.827 (0.635)	0.103 (0.464)	-0.0220 (0.424)
Population × Population			0.0270 (0.0226)	0.00439 (0.0187)	-0.00663 (0.0244)	-0.0370 (0.0226)	-0.0138 (0.0161)	-0.00669 (0.0148)
Cohort-level Controls								
Migrant share (w)				0.633^{***} (0.160)	0.680^{***} (0.159)	0.964^{***} (0.172)	0.0583 (0.161)	0.0528 (0.148)
Average years of education (w)				0.528^{*} (0.295)	0.704^{**} (0.274)	-0.744^{**} (0.288)	0.979^{***} (0.288)	1.005^{***} (0.299)
Variance of education (w)				0.350^{*} (0.178)	0.402^{**} (0.194)	1.079^{***} (0.180)	-0.140 (0.218)	-0.221 (0.228)
Cohort-specific initial conditions								
GDP p.c. 1940-89 (w)					0.131^{**} (0.0526)	0.0565 (0.0481)	0.202^{***} (0.0481)	0.144^{**} (0.0576)
Region fixed effects			Yes	Yes	Yes	Yes	Yes	Yes
Birth-year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	154	154	154	154	154	154	154	154
Adjusted R-squared	0.663	0.667	0.667	0.671	0.673	0.671	0.665	0.687

Notes: This table presents estimates of intergenerational persistence β based on various model specifications. The dependent variable is the intergenerational income elasticity β. Standard errors are in parentheses. The symbols *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 2.2: Estimates on social mobility and economic growth. Intergenerational persistence β

	OLS				System-GMM			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
M(w)	-0.0938^{**}	-0.551^{***}	-0.534^{***}	-0.378^{*}	-0.0603	-0.401^{*}	-0.421^{**}	-0.530^{*}
	(0.0424)	(0.199)	(0.181)	(0.218)	(0.0410)	(0.220)	(0.184)	(0.301)
Income Lag (-1)	0.900^{***}	0.683^{***}	0.710^{***}	0.653^{***}	0.949^{***}	0.715^{***}	0.768^{***}	0.643^{***}
	(0.0280)	(0.0624)	(0.0568)	(0.0390)	(0.0353)	(0.0575)	(0.0436)	(0.0912)
Country-Year F.E.				Yes				Yes
Year F.E.			Yes				Yes	
Region F.E.		Yes	Yes	Yes		Yes	Yes	Yes
Other controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	1319	1319	1319	1319	1319	1319	1319	1319
Hansen J-test (p-value)					0.1554	0.0000	0.0000	0.7089
Arellano-Bond test AR(1)					0.0001	0.0002	0.0004	0.0059
Arellano-Bond test AR(2)					0.6180	0.6109	0.8681	0.2716

Notes: Dependent variable is the log per capita income of a region (between 1981 and 2018). $M(w)$ is the weighted intergenerational persistence (measured by the slope coefficient) of people born between 1940 and 1989. Spillover effects are controlled by including the average degree of intergenerational persistence in year t of all other regions $-j$ in the country (i.e. region j is excluded to estimate this average). For a detailed description of the data and all variables included in the regressions see Section 4.. Robust standard errors clustered at the country-year level. Source: National Household Surveys, own estimates.

Table 2.3: Estimates on social mobility and economic development. Intergenerational persistence β

	Luminosity	Poverty	Employment	Formality	Water	Electricity
M (w)	-0.817^{***}	2.518^{**}	-0.795^{***}	-0.525^{**}	-0.786^{***}	-0.192
	(0.132)	(0.997)	(0.105)	(0.206)	(0.172)	(0.156)
Region and Country-Year F.E.	Yes	Yes	Yes	Yes	Yes	Yes
Year level controls	Yes	Yes	Yes	Yes	Yes	Yes
Cohort level controls	Yes	Yes	Yes	Yes	Yes	Yes
Cohort-specific initial conditions	Yes	Yes	Yes	Yes	Yes	Yes
Spillover effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	999	1368	1368	1223	1278	1128

Notes: Dependent variable is indicated in the column-title. M (w) is the weighted intergenerational persistence (measured by the slope coefficient) of people born between 1940 and 1989. For a detailed description of data and variables see Section 4.. Robust standard errors clustered at the country-year level. Source: National Household Surveys, own estimates.

Table 2.4: Allocation vs. accumulation of human capital and economic development

	(1)	(2)	(3)	(4)	(5)
Upward Mobility (w)	1.185***			1.073***	1.181***
	(0.101)			(0.101)	(0.0955)
Top Persistence (w)		0.494*	0.430*	0.0441	0.0279
		(0.257)	(0.250)	(0.229)	(0.227)
Average years of education (w)			1.462***	0.607***	
			(0.205)	(0.217)	
Region and Country-Year FE	Yes	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes	Yes
Observations	1368	1368	1368	1368	1368

Notes: Dependent variable is the log per capita income of a region (between 1981 and 2018). Estimations include the full set of control variables, region fixed effects and country-by-time fixed effects. For a detailed description of data and variables see Section 4.. Robust standard errors clustered at the country-year level. Source: National Household Surveys, own estimates.

Table 2.5: Estimates on social mobility and income inequality. Intergenerational persistence β

	(1)	(2)	(3)	(4)	(5)
M(w)	0.158**	0.412***	0.577***	0.531***	0.727***
	(0.0741)	(0.0809)	(0.0960)	(0.110)	(0.151)
log income per capita	0.149***	0.171***	0.170***	0.204***	0.208***
	(0.0243)	(0.0228)	(0.0226)	(0.0213)	(0.0213)
Region and Country-Year F.E.	Yes	Yes	Yes	Yes	Yes
Year level controls		Yes	Yes	Yes	Yes
Cohort level controls			Yes	Yes	Yes
Cohort-specific initial conditions				Yes	Yes
Spillover effects					Yes
Observations	1368	1368	1368	1368	1368

Notes: Dependent variable is regional income inequality measured by the Gini coefficient of disposable household per capita income (between 1981 and 2018). $M(w)$ is the weighted intergenerational persistence (measured by the slope coefficient) of people born between 1940 and 1989. For a detailed description of data and variables see Section 4.. Robust standard errors clustered at the country-year level. Source: National Household Surveys, own estimates.

Chapter 3

Reconciling the Evidence on False Links in Historical Data

1. Introduction

New data linking technology is revolutionizing economic history, allowing the creation of large-scale longitudinal and intergenerational datasets to answer questions that have long been out of reach. Ongoing infrastructure projects are linking national surveys, administrative data, and research samples to recently digitized historical records, such as full count U.S. censuses (Ruggles, 2006; Ruggles et al., 2021; Helgertz et al., 2022; Abramitzky et al., 2021), vital records (LIFE-M project: Bailey et al., 2023), and Social Security records (Goldstein et al., 2021; Bailey et al., 2023).[1]

[1] Many on-going initiatives link the 1940 U.S. Census to other datasets. The Census Bureau plans to link the 1940 Census to current administrative and Census data (Census Longitudinal Infrastructure Project, CLIP) and the Minnesota Population Center plans to link it to other historical censuses. It has also been linked to five of the most widely used surveys, including the Panel Survey of Income Dynamics (Warren et al., 2020b), the Health and Retirement Survey (Warren et al., 2020a), and the National Health and Aging Trends Survey (Freedman et al., 2021). Supplementing these public infrastructure projects, entrepreneurial researchers have also combined large datasets. See, for example, Abramitzky et al. (2012, 2013, 2014); Boustan et al. (2012); Hornbeck and Naidu (2014); Mill (2013); Mill and Stein (2016); Aizer et al. (2016); Bleakley and Ferrie (2017, 2016, 2013); Nix and Qian (2015); Collins and Wanamaker (2014, 2015, 2016); Eli et al. (2018). This paper discusses many of the linking approaches used in these papers.

Recent findings regarding the accuracy of the machine methods used to create these "big data" have been mixed. Bailey et al. (2020) report that 15 to 37% of individuals linked in historical contexts are incorrect. Ó Gráda et al. (2023) report that up to 50% of Irish immigrants matched by the methods of Abramitzky et al. (2012, 2014) are incorrect. In addition, Bailey et al. (2020), Ó Gráda et al. (2023), and Ghosh et al. (2024) show that these errors are consequential for inference about social and geographic mobility. On the other hand, Abramitzky et al. (2021) report substantially lower error rates for the same algorithms and data, which they argue has little effect on inference. Similarly, Price et al. (2021) report error rates of around 6-7%.

This paper reconciles these seemingly disparate conclusions and explores variation in error rates across subgroups and linked census datasets. Our first finding is that the seemingly different error rates in the literature reflect the fact that different papers report different parameters. Lower error rates such as those headlined in Abramitzky et al. (2021) are conditional, in the sense that they calculate Type I linking errors within the intersection of algorithm and ground truth datasets. Higher error rates such as those featured in Bailey et al. (2020) are unconditional, in the sense Type I linking errors are computed for all algorithm links rather than the subset of the data in the ground truth. We show that the entire difference between the conclusions in Abramitzky et al. (2021) and Bailey et al. (2020) reflects focusing on different parameters. Abramitzky et al. (2021) main results featuring very low linking errors are based on the analysis of a subsample of data that is not representative of research samples in other contexts.

While these two parameters could be similar or even identical in some instances, our second finding is that the unconditional error rate is typically much higher in common empirical applications. The reason is that links made by both the algorithm and the ground truth are not a representative sample of all algorithm links and are much more likely to be correct. For example, links made by the algorithm that are not made in the Oldest Old data from the Early Indicators project are very likely errors: these links have been scrutinized by genealogists using multiple data sources and discarded as erroneous. Links made by an algorithm (but not by these genealogists) which are based on much less information are likely to be incorrect. This means that the Type I error rate computed for the subset of links in both data sets substantially underestimates the unconditional Type I error rate. Because most historical analyses of linked data do not typically have ground truth data available, the higher unconditional error rate is closer to the true error rate and, therefore, the parameter capturing measurement error with linked samples. Differences between the conditional and unconditional error rates may be lower in other ground truth datasets.

A third set of results uses the Oldest Old data from the Early Indicators Project to examine how Type I linking errors vary across subpopulations and with respect to the time period considered. The results show that Type I linking errors are typically much higher for more disadvantaged or mobile groups, implying significant measurement error to analyses of lower socio-economic status groups, racial minorities, and immigrants. In addition, the results show that the share of individuals unlinkable due to death or emigration raises linking error rates, meaning that error rates are likely higher in datasets like the census and when linking individuals over longer time periods.

A final analysis illustrates how consequential the differences in the conditional and unconditional linking errors can be in practice when using two popular, large-scale linked datasets: Minnesota's Population Center's (MPC) Multigenerational Longitudinal Panel (MLP, see Helgertz et al., 2023) and the Census Linking Project (CLP, see Abramitzky et al., 2022). These projects link millions of individuals to decennial censuses using different automated linking methods for doing so. Whereas the CLP generates links using deterministic methods, the MLP links were generated using a supervised, probabilistic, method and incorporate co-resident links. The subsample of individuals in the 1910-1930 Censuses linked to the 1940 Census made both by MLP and CLP differ only by 1 to 6% of the time, a conditional calculation suggesting low overall rates of linking errors. However, this comparison ignores links made by only one of these two sources: up to 35% of the links in CLP are not made by the MLP project, and between 32 (1910-1940) and 38% (1930-1940) of MLP's links are not found by CLP. To evaluate link quality, we use link probability scores based on name, age, race, and birth state similarity, as well as household co-residents and geography for millions of record pairs. Our analysis using the 1930-1940 links show that CLP-only links are characterized by very low link probability scores. Moreover, for 14 to 28% of the CLP-only links, there are better matches in the 1940 Census according to link scores. We conclude with suggestions for assessing the incidence of measurement error in machine-linked historical data.

2. Conditional and Unconditional Error Rates in Historical Settings

"Ground truth" is defined as data obtained by direct observation of the true link. Ground truth data rarely exist for historical populations, so researchers have instead used high quality, hand-linked data to investigate the performance of various linking methods. For instance, researchers have used genealogies of large sets of people compiled by Family Search (Helgertz et al., 2022; Price et al., 2021) or rich but restricted administrative data unavailable to most researchers (Abowd, 2017; Massey, 2017; Scheuren and Winkler, 1993; Winkler, 2006).[2]

Figure 3.1 illustrates this fairly common situation where the available ground truth overlaps incompletely with the data linked by an algorithm. The square represents the universe of individuals. The solid blue circle contains observations linked in the ground truth. (Ground truth data typically include linked data and are often silent on why links are not made.) The dashed red circle contains all observations linked by the algorithm. A represents links

[2]Large, historical population registers are available for other countries (Christen and Goiser, 2007; Eriksson, 2016; Goeken et al., 2017).

made by the ground truth and the algorithm, where the algorithm and the ground truth links agree. D represents links made by the ground truth and the algorithm, for which the algorithm's link disagrees with the ground truth. We assume the ground truth contains the correct link, which implies that all of the links in D are records of the algorithm linked incorrectly. Additionally, let $C \cup E$ be the set of links made by the algorithm that are not contained in the ground truth, where links in C are correct, and links in E are incorrect.

A common approach to evaluating the performance of an algorithm considers only the data linked by the algorithm and the ground truth. Because this parameter conditions the error rate on being linked by the ground truth, we call this the conditional error rate, which we define as,

$$P(\text{ Wrong Link} \mid A \cup D) = \frac{P(D)}{P(A \cup D)}$$

Note that this parameter is different from the unconditional error rate of an algorithm, which is defined,

$$P(\text{ Wrong Link } \mid A \cup C \cup D \cup E) =$$
$$= \frac{P(D \cup E)}{P(A \cup C \cup D \cup E)} = \theta \frac{P(D)}{P(A \cup D)} + (1-\theta)\frac{P(E)}{P(C \cup E)},$$

where $\theta = \frac{P(A \cup D)}{P(A \cup C U D U E)}$ is the share of records linked by the algorithm and the ground truth. Although the unconditional error rate is the parameter of interest for evaluating the role of measurement error, incomplete information about which records fall in regions C and E means that $\frac{P(E)}{P(C U E)}$ is not observed and cannot be calculated directly.

The unconditional and conditional error rates are only equal in special cases, where $P(D \mid A \cup D) = P(E \mid C \cup E)$. This could be the case if the ground truth was created as a random sample of algorithm links (e.g., the researcher randomly drew a set of algorithm links and evaluated each by hand for errors). In most cases, however, this equality does not hold and the relationship between the unconditional and conditional error rates depends on (i) the incidence of links that the algorithm has in common with the ground truth, θ, as well as (ii) the probability that a link made by the algorithm but not made by the ground truth is incorrect. Two cases illustrate this relationship.

Case 1. The ground truth sample contains all true links

For this case, assume that skilled reviewers considered and rejected all links not made in the ground truth. This implies that C is the null set, $P(C) = 0$, and $P(E \mid C \cup E) = 1$. That is, any link made by an algorithm in $C \cup E$ is incorrect. The unconditional error rate is, therefore,

$$P(\text{ Wrong Link } \mid A \cup C \cup D \cup E) = \theta \frac{P(D)}{P(A \cup D)} + (1 - \theta)$$

Note that if $\theta < 1$, then the conditional error rate, $\frac{P(D)}{P(A \cup D)}$, strictly understates the unconditional error rate. This circumstance could arise in very high quality administrative or genealogical data. One example of this type of data could be the Early Indicators data, which was carefully reviewed by skilled genealogists. Links stated by an algorithm using the Early Indicators data that is not stated by the hand-linked data are highly likely to be incorrect, or $\frac{P(E)}{P(C \cup E)} \approx 1$.

Case 2. The ground truth sample does not contain all true links

For this case, assume that skilled reviewers did not consider or reject all links not in the ground truth, so C is not the empty set and $P(C) > 0$. The unconditional error rate of an algorithm is

$$P(\text{ Wrong Link } \mid A \cup C \cup D \cup E) = \theta \frac{P(D)}{P(A \cup D)} + (1-\theta) \frac{P(E)}{P(C \cup E)}$$

Note that the relationship between the conditional and unconditional error rate is theoretically ambiguous. The conditional error rate could under- or overstate the unconditional error rate, depending on whether $\frac{P(E)}{P(C \cup E)}$ is larger or smaller than $\frac{P(D)}{P(A \cup D)}$.

The relationship between the unconditional and conditional error rates depends upon the relationship between $\frac{P(E)}{P(C \cup E)}$ and $\frac{P(D)}{P(A \cup D)}$, which is not known in most applications. Note that the relationship of these error rates will also vary with the ground truth used (i.e., the dataset used in the conditioning).

However, some theoretical reasons and empirical evidence we will provide below suggest that generally $P(E \mid C \cup E) > P(D \mid A \cup D)$, which implies that the unconditional error rate is typically higher than the conditional error rate. Individuals who have not been linked in a ground truth tend to (i) change their name or its spelling, (ii) live with different people, and (iii) be more geographically or economically mobile (migrated within the U.S. or returned to their home country). In addition, many more may have (iv) died between censuses or points of observation.

One example of a ground truth that does not contain all true links is the FamilySearch.org Family Tree (Price et al., 2021). These data were produced by crowd-sourced linking that would reflect primarily the family lineages of the people who selected to use the platform and are generally found to be highly accurate. However, these data do not contain all possible links, so algorithms linking the same data may make some correct links that were not initially on the Family Tree.

3. Conditional and Unconditional Error Rates in the Early Indicators Data

Our analysis uses a subset of the Union Army Dataset drawn from the Early Indicators Project (Costa et al., 2017). The project created what they call the Oldest Old sample, which consists of over two thousand Union Army veterans who lived to over 95 years of age. Using genealogical methods and a rich set of supplementary materials, the project then determined whether these men were confirmed to be alive in 1900 and, if so, attempted to link them to the complete-count 1900 Census.

The Oldest Old data serves as invaluable ground truth for historical census linking and is most similar to Case 1 in section 4.1. This sample contains only people who were confirmed to live to a certain age and excludes those who died or emigrated. Importantly, genealogists have explored multiple data sources to link the Union Army veterans, making it unlikely that an algorithm (using less information) will identify a true link that genealogists missed. Thus, the probability that a link made by an algorithm but not made by genealogists using more information is very close to zero.

3.1 Conditional Error Rates using the Oldest Old Data

Both Bailey et al. (2020) and Abramitzky et al. (2021) use this sample to evaluate the performance of different linking algorithms, restricting the data to observations that are confirmed to have a "high quality" link to the 1900 complete-count U.S. Census.[3] This makes the error rates reported in Bailey et al. (2020) and Abramitzky et al. (2021) for the Oldest Old data a conditional error rate. Here, we re-analyze this data and show how

[3]Bailey et al. (2020) use Early Indicators matches coded as quality 1 and 2, while Abramitzky et al. (2021) only use matches coded as quality 1. In practice, the further restriction in Abramitzky et al. (2021) only alters TPR and PPV somewhat. To strictly replicate Abramitzky et al. (2021), we use the same limitation here.

consequential this restriction is.

The first step in our analysis is replicating Bailey et al. (2020) and Abramitzky et al. (2021). Figure 3.2 plots the Positive Predictive Value (PPV), or 1 minus the Type I error rate in linking, versus its True Positive Rate (TPR) for each algorithm.[4] Panel A uses the conditional Type I error rate.[5] Abramitzky et al. (2021) evaluate more algorithms than **?**, but both find similar results.

3.2 *Unconditional Error Rates in the Oldest Old Data*

Most researchers are not able to condition an analysis on whether individuals are alive or "linkable," which are two of the implicit restrictions in Bailey et al. (2020) and Abramitzky et al. (2021). We next examine the error rates in the entirety of the Oldest Old data confirmed to live until 1900, rather than restricting our calculations to observations that were linked by the algorithm and had a link to the 1900 Census in the Oldest Old.[6] This broader Oldest Old sample includes (1) lower quality links to 1900 (excluded from the conditional rate calculations), (2) individuals with multiple potential links to 1900, and (3) individuals with no link to 1900 . We present these results in Figure 3.2 Panel B.

One interesting finding is that the TPR is generally unchanged. This is likely a consequence of the small sample size in the Early Indicators data. Most of the observations in the

[4]PPV is the share of correct links made—according to the ground truth—out of the total number of links made by the algorithm. TPR is the share of correct links made out of the total number of potential correct links in the ground truth. Both are standard measures of algorithm performance, and linking algorithms often have a trade-off between the two.

[5]Figure 3.2 Panel A is almost identical to Abramitzky et al. (2021) Figure 1 ("Conditional", panel A) and Figure 3.2 Panel B to Abramitzky et al. (2021) Appendix Figure A1 ("Unconditional", panel B). The exception to this similarity is for Feigenbaum (2016), where the difference in PPV likely stems from the fact that we use different training data. Abramitzky et al. (2021) created their own training data by hand to offer an example of how the Feigenbaum (2016) matching algorithm would work in many practical applications, but we trained the algorithm using the matches stated in the Oldest Old data that were made by genealogists. It is not surprising that the supervised machine learning algorithm in Figure 3.2 performed somewhat better, as we likely trained the algorithm using links closer to the ground truth.

[6]To confirm death, genealogists used sources such as gravestone databases, obituaries, newspaper accounts, veterans associations and pension files to allow multiple cross-validation. The Oldest Old data also have the advantage of containing linked death certificates, which are typically unavailable in census linking.

Early Indicators data have unique names, ages, and birthplaces, and including additional observations does not alter the uniqueness of these combinations of characteristics. In linking settings that use larger datasets (e.g., census-to-census linking), observations that do not have links may look similar in a census to an observation that does not have a link, and including those additional observations may "crowd out" the link that would otherwise be made.

A second finding is that examining all observations (including those unlinked to the 1990 Census in the Oldest Old sample) has a large effect on PPV and Type I error rates. The unconditional PPV rate is around 15 percentage points lower (or the error rate is 15 percentage points higher). The roughly parallel trend lines (see slope parameter on TPR) suggest that the difference between the conditional and unconditional error rates is very similar across all algorithms. The Type I error rate increases by about 5.4 percentage points for every 10 percentage-point increase in TPR. While Abramitzky et al. (2021) imply that linking errors are generally low, this conclusion rests on sample restrictions—conditioning that is not possible for researchers without ground truth. The PPV in Figure 3.2 Panel A does not characterize the Type I error rate relevant to using linked data in regression settings. The relevant error rate is in Figure 3.2 Panel B.

3.3 How Increasing the Incidence of Unlinkable Records Affects Errors

One feature of the Oldest Old data is that researchers know with a high degree of confidence when 71% of the sample dies.[7] This allows our analysis in Figure 3.2 to exclude Union Army veterans who did not survive until 1900 and, consequently, should not link to the 1900 Census. Most research contexts have much less information, and researchers do not know if

[7] In the Union Army data 29.4% of observations have no confirmed death year, and these individuals are more likely to be illiterate and be an immigrant than individuals with a confirmed death year. We exclude individuals without a confirmed death year from this exercise as they cannot be confirmed to be unlinkable, but note that these observations may be especially prone to errors in linking (Ó Gráda et al., 2023). Consequently, PPV and TPR rates we estimate may overstate the levels of recall that would happen in this exercise if we did not have to make this additional restriction on the sample.

someone has died or emigrated.

To evaluate algorithm performance without this information (e.g., census-to-census linking), we progressively increase the sample by 10% (230 veterans), 20% (522 veterans), and 40% (1394 veterans) to include randomly selected individuals from Union Army veterans confirmed to be dead before 1900 . By construction, these veterans do not have a valid link in the 1900 Census.

Figure 3.3 presents the results. Increasing the share of unlinkable individuals from 0% (as in Figures 3.2 Panel A and B) to 40% strengthens the negative relationship between TPR and PPV. Whereas PPV and TPR for the most conservative EM method change little, algorithms with higher TPRs exhibit much lower PPVs, resulting in a clockwise rotation of the regression line. In short, error rates in the Early Indicators data for individuals known to be alive may be substantially lower than error rates in other datasets where unlinkability due to death, emigration, and measurement error is more common. Moreover, linkage projects across longer time periods (when more will have died or emigrated) are also expected to have higher error rates compared to the rates characterized here.

3.4 How Error Rates Vary across Subpopulations

Another feature of the Oldest Old data is that they contain a very diverse group of veterans, including an oversample of Black veterans. This allows us to characterize how algorithms perform for different subpopulations, especially groups of lower socio-economic status (SES) and greater geographic mobility (e.g., younger individuals or immigrants). Figure 3.4 presents the relationship between conditional and unconditional PPV and TPR by literacy (panels A-B), home ownership (panels C-D), race (panels E-F), age (panels G-H), and immigrant status (panels I-J).

Reflecting a pattern noted in Bailey et al. (2020), linking algorithms often perform worse with more disadvantaged groups or mobile groups. Unconditional PPV (Type I error rates) average 74% (26%) for literate veterans versus 52% (48%) for illiterate veterans; 76% (24%) for homeowners versus 72% (28%) for non-owners; 74 % (26%) for White veterans versus 68% (32%) for Black veterans; 80% (20% for older veterans versus 67% (33%) for younger veterans; and 72% (28%) for native-born veterans versus 70% (30%) for immigrants. These findings imply that comparisons across subpopulations could be heavily influenced by linking errors.

Additionally, the difference in error rates between more and less disadvantaged groups is proportionally larger in unconditional data than in conditional data. For example, the average unconditional PPV for illiterate veterans is 30% lower than the PPV for literate people, but this statistic is only 21% lower for the conditional PPV. The drop in TPR for the same algorithms for the same subpopulation is 43% in the unconditional data and 38% in the conditional data. In short, conditional error rates consistently understate linking errors and the understatement is larger for more disadvantaged or mobile populations.

4. Conditional and Unconditional Errors in the Wild: A Case Study of MLP and CLP

Links provided by the Multigenerational Longitudinal Panel (MLP) and the Census Linking Project (CLP) allow us to characterize conditional and unconditional linking errors outside of the laboratory setting—that is, the performance of algorithms "in the wild." Both methods link millions of individuals between every pair of censuses, separated by up to three decades, from 1850 to 1940. In this paper, we evaluate six types of CLP links that are based on variations of the "Iterative Method" by Abramitzky et al. (2012, 2014).[8] This unsupervised linking approach combines data using theoretically time-invariant characteristics, such as

[8]We understand that CLP will soon include links based on Abramitzky et al. (2020) adaptation expectation maximization algorithm (Fellegi and Sunter, 1969; Winkler, 2006; Dempster et al., 1977), but these links are not yet available.

name, age, race, and place of birth information, and different versions use different string-cleaning algorithms and tolerances for age discrepancies.

The MLP method differs from CLP in that it implements a supervised, probabilistic, machine learning algorithm, also incorporating potentially time-varying household characteristics, such as other household members and geography, along with time-invariant characteristics also used in CLP (Helgertz et al., 2022). The benefit of using additional characteristics is that the links may be of higher quality; the cost is that the final linked sample may be less representative (Ruggles, 2006). Although the MLP method links individuals in a two-stage procedure, initially linking men and subsequently linking other household members regardless of sex, we only examine links for the men through the first step for comparability with CLP.

The machine learning model used in MLP is trained using data based on manual searches of genealogical records provided by Ancestry.com. A ten-fold train-test-split procedure is employed, calibrating the model's optimum performance using Matthew's Correlation Coefficient. The software hlink, used to implement the MLP method, can be used to obtain record pairs' probability scores from the model that was used to determine the MLP links. Crucially for this paper, a probability score is assigned to every potential link, implying that we can evaluate the degree of similarity—according to the MLP method—of any record pair that fulfills standard blocking criteria. For potential matches between the 1930 and 1940 Censuses alone, this amounts to 1.25 billion observation pairs.

Importantly, CLP and MLP share many similarities with the linking examples in the previous section. Both datasets contain only individuals that were successfully linked across censuses. Like our examples in Figure 3.3, both datasets try to link individuals who died or emigrated between the census years because death and emigration are not observed. Unlike our examples using the Union Army data, however, researchers do not know whether the links are correct.

4.1 Comparing Links Common to MLP and CLP to Compute Conditional Type I Error Rates

Figure 3.5 illustrates the conceptual overlap between the two sets of links. Different linking methodologies imply that many links are made by both projects ($S \cup T$). However, MLP uniquely links some observations ($X \cup Y$) and different variants of CLP uniquely link some observations ($V \cup W$). The union of all MLP and CLP links ($S \cup T \cup X \cup Y \cup V \cup W$) captures the union of all men linked by either database between the 1910/1920/1930 and 1940 Censuses.

Table 3.1 presents these results for men in the CLP and MLP projects linked between the 1910 and 1940 Censuses (columns 1-3), the 1920 and 1940 Censuses (columns 4-6), and the 1930 and 1940 Censuses (columns 7-9). Panel A shows that the MLP algorithm declared 10.4, 15.8, and 27.8 million links between those censuses, respectively. Corresponding ranges for the variants of CLP are 5.8-9.5, 8.6-13.6, and 12.4-19.1 million. The union of the MLP-CLP links ranges from 14.9 to 33.1 million, depending upon which variant of CLP is used.

Unlike the Early Indicators data, none of the links have been reviewed by genealogists. In this case, we characterize agreement between the two projects when the links overlap in the region $S \cup T$ and when they do not (links in the region $X \cup Y$ or $V \cup W$). Bailey et al. (2020) show that links in the set made by multiple methods tend to have much lower Type I error rates, so we suspect that the rate of disagreement in links in the intersection set understates the Type I error rates in either dataset.

Panel B shows that among individuals in a base year linked by both the MLP and CLP standard methods ($S \cup T$), the disagreement ranges from 3.2% (Exact standard, 1910-1940) to 5.8% (NYSIIS standard, 1920-1940). When we consider links made by both MLP and the conservative CLP methods, the number of disagreements between the two sets of linked samples falls to no greater than 1% between 1910 and 1940, 1.3% between 1920 and 1940, and 1.6% between 1930 and 1940 . Under the assumption that MLP and the CLP are correct when they agree on a match, these results imply a conditional Type I error rate in either dataset is at most between 5.8% using standard CLP methods and 2% using conservative CLP methods.[9] Decreases in the rate of disagreement when comparing the linking results for the more conservative method provide additional evidence that rates of disagreement are correlated with Type I errors.

Taken at face value, these findings are good news: Type I errors in CLP and MLP are lower than in the Early Indicators Oldest Old data (section) and lower than reported in Bailey et al. (2020) or Ó Gráda et al. (2023). These error rates are also similar to estimates reported in Helgertz et al. (2022), which computes the accuracy of both MLP and CLP links based on their overlap with an external genealogical database. In short, the conditional Type I error rates for the intersection set of CLP and MLP appear very low.

4.2 Bounding the Unconditional Error Rates in CLP and MLP

Panel C of Table 3.1 shows that both MLP and CLP produce a large number of links that are not matched by the other algorithm—links falling into the sets $X \cup Y$ or $V \cup W$ in Figure 3.5. For the 1910-1940 links, 32% of MLP's links and 23-35% of CLP's links are unique to each project. For the 1930-1940 period, these shares are 38% for MLP and 11-22% for CLP. The magnitude of these non-overlapping sets means that the error rates in these regions will play a large role in determining the unconditional error rates. In the extreme case scenario where all 1910-1940 MLP-only links are incorrect, MLP would have an unconditional error rate of at most 34.9%. On the other hand, if the 1910-1940 links made only by CLP exact conservative are incorrect, this would result in a 20.4% unconditional error rate for this method, compared to 38.2% for the CLP NYSIIS version. These rates fall when considering links with the census

[9]If both MLP and CLP make the same errors, the actual rates in both datasets may be higher. Investigation of the data using alternative sources suggests this should only occur rarely. See Helgertz et al. (2022) for a discussion of how alternative data and methods highlight much lower error rates in MLP.

years closer to each other, but they imply upper bounds on unconditional Type I error rates that are substantially larger than the conditional error rates. Put succinctly, Type I error rates for 1910-1940 links could be less than 1% for both projects, or as high as 34.9% for MLP (assuming all links not found in CLP are incorrect) or 38.2% for CLP (assuming all links not found in MLP are incorrect).

4.3 Using MLP Link Probabilities to Infer Error Rates

A feature of the software implementing the MLP algorithm is that all potential matches that satisfy the imposed blocking criteria are assigned a probability score. Consequently, these were assigned both to record pairs that are eventually determined to be matches according to MLP but also to record pairs that are not determined to be matches. For example, scores were assigned to 1.25 billion record pairs in the case of the 1930-1940 linking run. These scores allow us to narrow this wide range of Type I errors into a range that is empirically more likely. The MLP probability score is calculated based on all of the characteristics used in the first step of the MLP linking algorithm and ranges from zero to one. All else equal, the score will be low when two records are dissimilar and high when the two records are similar in both observed characteristics of the individual being linked as well as co-resident individuals.

Panel A of Table 3.2 summarizes probability scores for links made by both MLP and CLP. Across linking intervals, the mean probability score for 1930-1940 for links common to both projects is high at around 0.95. The corresponding scores for 1920-1940 and 1910-1940 links are lower at 0.74-0.76 and 0.74-0.78, respectively (see Appendix Table A.10). The decline in scores reflects greater uncertainty characterizing matches made over a longer period and also changing linking thresholds. For example, no record is considered a match in the MLP if the first step linking score is below 0.4 , but these thresholds are lower at 0.16 and 0.15 for the 1920-1940 and 1910-1940 periods, respectively.

Some of the links made by both MLP and the CLP, however, are to different 1940 records. Panel A of Figure 3.6 shows the distribution of these scores. In the case of these discrepant links, no CLP links have scores exceeding the MLP link score (by construction). More telling, however, is that the scores are so different. MLP scores are considerably higher (exceeding 0.4 by construction and typically higher than 0.75), whereas many CLP scores are below 0.25.[10] For example, about 70% of the CLP Exact Standard links have a probability score of 0.15 or less, while about as much of the distribution for competing MLP links obtain scores of at least 0.95 . The mean probability score for the cases linked by both methods, but to different 1940 records, are 0.16 for CLP Exact Standard and 0.91 for MLP. Table S1 (Supplementary Material) illustrates a range of cases in the data, which shows that high probability score MLP links tend to represent matches made in the presence of supporting evidence from additional information, including co-resident household members and geography.[11]

Panel B of Table 3.2 summarizes the probability scores for links outside the intersection set, links only made by MLP (but not by any CLP variant) or only CLP (but not linked by MLP). Notably, there is a slight difference in the probability score between MLP-only links (0.92) and those where both MLP and a CLP version agree (0.95). For 1910-1940 and 1920-1940 (Table A.10, Appendix), the difference is 20.5 and 9.3 percentage points, respectively. This underscores that MLP-only links are made with almost as much confidence as the links made by both CLP and MLP. For a detailed view of the full probability distribution of MLP-only links, refer to Panel B of Figure 3.6. Table S2 (Supplementary Material) provides examples of actual links, along with the next best potential match. These examples illustrate the value of the additional information provided by household co-residents, which MLP uses to construct the linking scores.

Panel B of Table 3.2 also displays the probability scores for CLP-only links, representing

[10]Note that only probability scores between 0-0.25 and 0.75-1 are displayed for illustrative purposes, which includes over 95% of cases.

[11]Tables in the Supplementary Material are available upon request due to containing confidential information, such as individual names.

between 11-22% of the method's 1930-1940 links. For these links, the average linking score is considerably lower than the MLP-only links. The average scores range from a low of 0.21 for both NYSIIS variants to a high of 0.28 for CLP-Exact Conservative. In addition, for a substantial share of the CLP-only links, there exists at least one record pair in the universe of potential matches with a higher score. For example, MLP scored another record as a higher probability link for 27.7% of CLP NYSIIS Standard links. Panel C of Figure 3.6 plots the distribution of scores for the CLP-Exact Standard and CLP-Exact Conservative only matches, where the 1930 record was not matched to a 1940 record by MLP. The dominance of low probability score matches is evident, suggesting that the majority of CLP-only matches are not borderline cases. CLP-ES-only examples are provided in Table S3 (Supplementary Material), along with the next best (or best overall) record match that exists in the 1940 Census.

The low share of very high probability score matches in CLP-only links is another interesting pattern in Panel C of Figure 3.6. Upon further inspection, these links are characterized by several features that led MLP to reject them. For example, many are associated with a potential match with an even higher probability score in the data or with multiple potential matches exceeding the MLP linking thresholds —ambiguity which led MLP to reject these links as correct. Examples are provided in Table S4 (Supplementary Material).

4.4 Variation in Link Probabilities by Race and Nativity

As illustrated in the analysis using the Early Indicators data, the quality of links and Type I error rates differ by subgroups, which also has implications for cross-group comparisons. Table 3.3 investigates patterns in the link scores by subgroups, focusing on Black men and foreign-born U.S. resident men linked between 1930 and 1940.[12] Panel A shows the number of links made by MLP and the different variants of CLP. Overall, the union of these sets

[12]Results for the 1910-1940 and 1920-1940 Census links are in Table A.11 and A.12 in the Appendix.

contains a total of 2,251,536 links for Black men and 3,158,563 for foreign-born men. Panel B breaks down these links according to which project made them (either MLP or CLP), as well as the number of links that were matched to different records (column 2). The share of all links in the intersection that were discordant links ranges from 2.3 to 11.0% for Black men (column 3) and 2.4 to 8.4% for foreign-born men (column 8). Among the links that agree between MLP and CLP, the average link probability hovers around 0.93 for Black men and 0.95 for foreign-born men. For example, the mean probability score for concordant MLP and CLP-Exact Standard 1930-1940 links for the full population is 0.95, dropping to 0.92 when limited to Black men.

Focusing on the 1930-1940 links, a similar story emerges when examining links made by both MLP and CLP versions but linked to different 1940 Census records. For CLP-Exact Standard links, the mean probability score is 0.11 and 0.19 for Black and the foreign-born men, respectively, compared to 0.88 and 0.92 for the competing MLP links. The complete distributions of the discrepant links are displayed in Panel A of Figure 3.7, again illustrating the dominance of very low probability links for CLP ES links for which the MLP algorithm suggests another 1940 match. In contrast, the MLP links are dominated by very high probability scores. (Note: this is true by construction.) Examples of all combinations of links are provided in Tables S5 and S6 (Supplementary Material tables for Black men and foreign-born men).

Turning to Panel C of Table 3.3, MLP-only links have slightly lower average link probabilities (0.88 for Black men and 0.92 for foreign-born men) compared to the intersection of MLP-CLP links. Panel B of Figure 3.7 shows that the probability distribution is heavily skewed towards very high probability score links, with examples of MLP-only links provided in Supplementary Material Tables S5 and S6.

In contrast, links chosen by CLP variants that were not matched by MLP range from 27 to 41% of all CLP links for Black men and 12 to 27% of all CLP links for foreign-born men. The full probability distribution of these links is presented in Panel C of Figure 3.7, illustrating the skew towards very low probability scores, with examples provided in Tables S5 and S6. This is particularly the case for Black men, with CLP-only link mean probability scores ranging from 0.113 to 0.152, and for foreign-born men, ranging from 0.211 to 0.347. In short, additional information used by MLP deems it unlikely that the CLP-only links are correct, in part because 15 to 34% of Black men and 17 to 30% for foreign-born men have plausible links in the 1940 Census with higher probability scores (columns 4 and 9).

5. Conclusion: How to Compute Unconditional Error Rates for Training Samples

Our results suggest that researchers should use caution in interpreting Type I linking error rates when estimated using ground truth data. Using ground truth data typically involves evaluating the quality of a subset of algorithm links. This subset of algorithm links is typically (1) selected and (2) linked with fewer errors and, therefore, (3) likely to provide an optimistic assessment of the measurement error in linked samples. Both theoretical reasons and the empirical evidence in this paper show that this approach underestimates the true error rates overall. The bottom line is that Abramitzky et al. (2021) headline that "a number of automated methods generate very low (less than 5%) false positive rates" is deeply misleading. The 5% false positive rate is a conditional error rate, which our analysis shows—using the same data and algorithms—to be much lower in practice than the true error rate. Had Abramitzky et al. (2021) featured the unconditional error rates in their main paper, their results would have been indistinguishable from those presented in Bailey et al. (2020). A second finding is that Type I linking errors are much higher for racial minorities, foreign-born men, and other disadvantaged groups, which are harder to link.

The Early Indicators data are of exceptional quality and are helpful in clarifying the distinction between conditional and unconditional error rates. Records linked by an algorithm but not by Early Indicators genealogists are very likely incorrect, making clear why the distinction between the two rates is important. Given that Bailey et al. (2020) did not use genealogists in the construction of LIFE-M, one might wonder how this project assessed unconditional error rates. In fact, the genesis of this paper's analysis is closely connected to this question.

At the outset of the LIFE-M project, Bailey et al. (2020) realized that a number of links were made by algorithms that had not been made in clerical review. To evaluate the quality of the links not contained in the hand-linked data, the LIFE-M project staged a "police line up" process. All links that disagreed between the data trainers and automated methods were re-reviewed by two trainers in a blinded process. The links made by the algorithm or ground truth were disguised in a larger set of candidates, so that the data trainers would not favor any link or method. This process created a way of evaluating the quality of links that were not made by both datasets in order to assess the unconditional likelihood of a Type I error. This re-review overturned the first stage LIFE-M hand-links in around 4% of cases, but it overturned links made by prominent automated methods 15 to 37% of the time (see Bailey et al., 2020, Table 1).

This paper also helps resolve another puzzle–the error rates for the Early Indicators data reported by Bailey et al. (2020) seem lower than the LIFE-M ones. The reason for this is that Bailey et al. (2020) only constructed conditional error rates for the Early Indicators data. That is, they only calculated errors for records linked both by the Oldest Old project and the algorithm but they ignored links made by only the algorithm (almost all of which would have been wrong). Thus, their reported error rates for the Oldest Old sample seem lower than the unconditional ones for the LIFE-M data (see Bailey et al., 2020, Figure 4-5). However, this apparent inconsistency reflects differences in the reported parameters rather than method performance. The unconditional error rates for the Oldest Old data reported in this paper are very similar to the unconditional error rates using the LIFE-M and synthetic data.

6. Figures

Figure 3.1: Overlap of Links Made by an Algorithm and Ground Truth

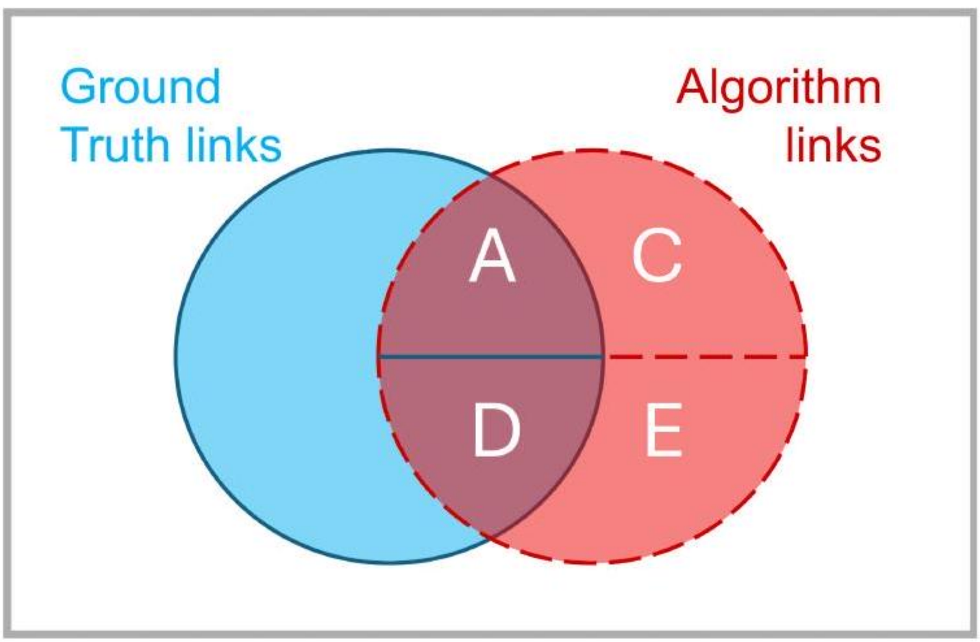

Notes: The square represents the universe of individuals. The left circle contains observations linked in the ground truth (solid blue), while the right circle contains all observations linked by the algorithm (dashed red). Area A represents links made by both the ground truth and the algorithm, where they agree. D represents links made by both, but where the algorithm's link disagrees with the ground truth, it is assumed to be correct. Additionally, C and E are the algorithm's links not in the ground truth, where C denotes correct links and E denotes incorrect ones.

Figure 3.2: Differences in Conditional and Unconditional Error Rates

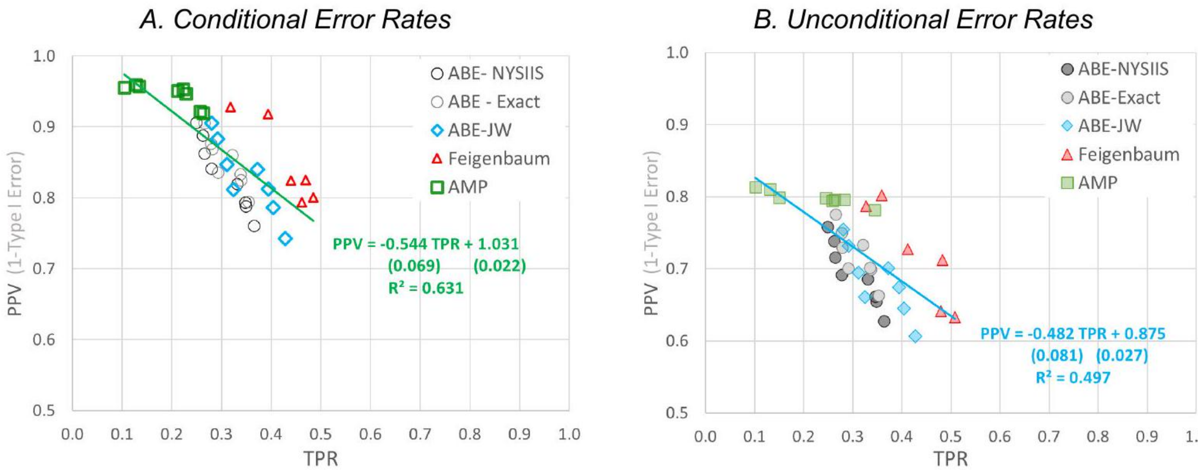

Notes: Figure 3.2 presents the relationship between the conditional PPV (panel A), unconditional PPV (panel B), and TPR for a variety of commonly used machine linking algorithms. The linear regression fit is obtained using OLS, with standard errors shown in parentheses. Source: Authors' calculations from the Oldest Old Early Indicators Data.

Figure 3.3: Changes in Error Rates as the Share of Unlinkable Individuals Rises

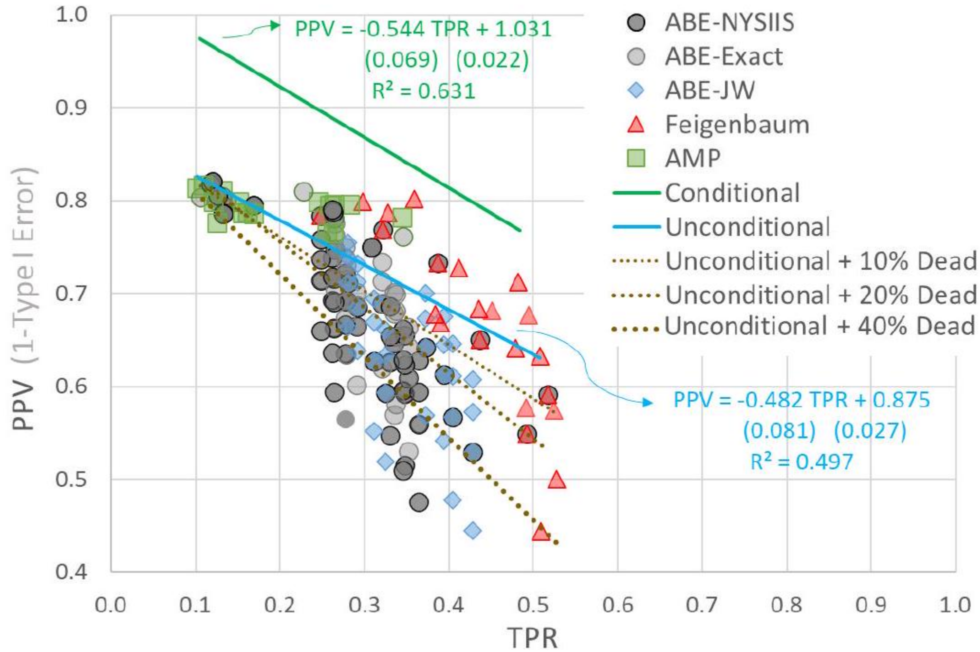

Notes: Figure 3.3 presents the relationship between PPV and TPR as the number of unlinkable Union Army veterans increases from zero (solid blue link) to 40% (large dotted line). See also Figure 3.2 notes. Source: Authors' calculations from the Oldest Old Early Indicators Data.

Figure 3.4: Differences in PPV (1-Type I Error Rates), by Subgroup and Method

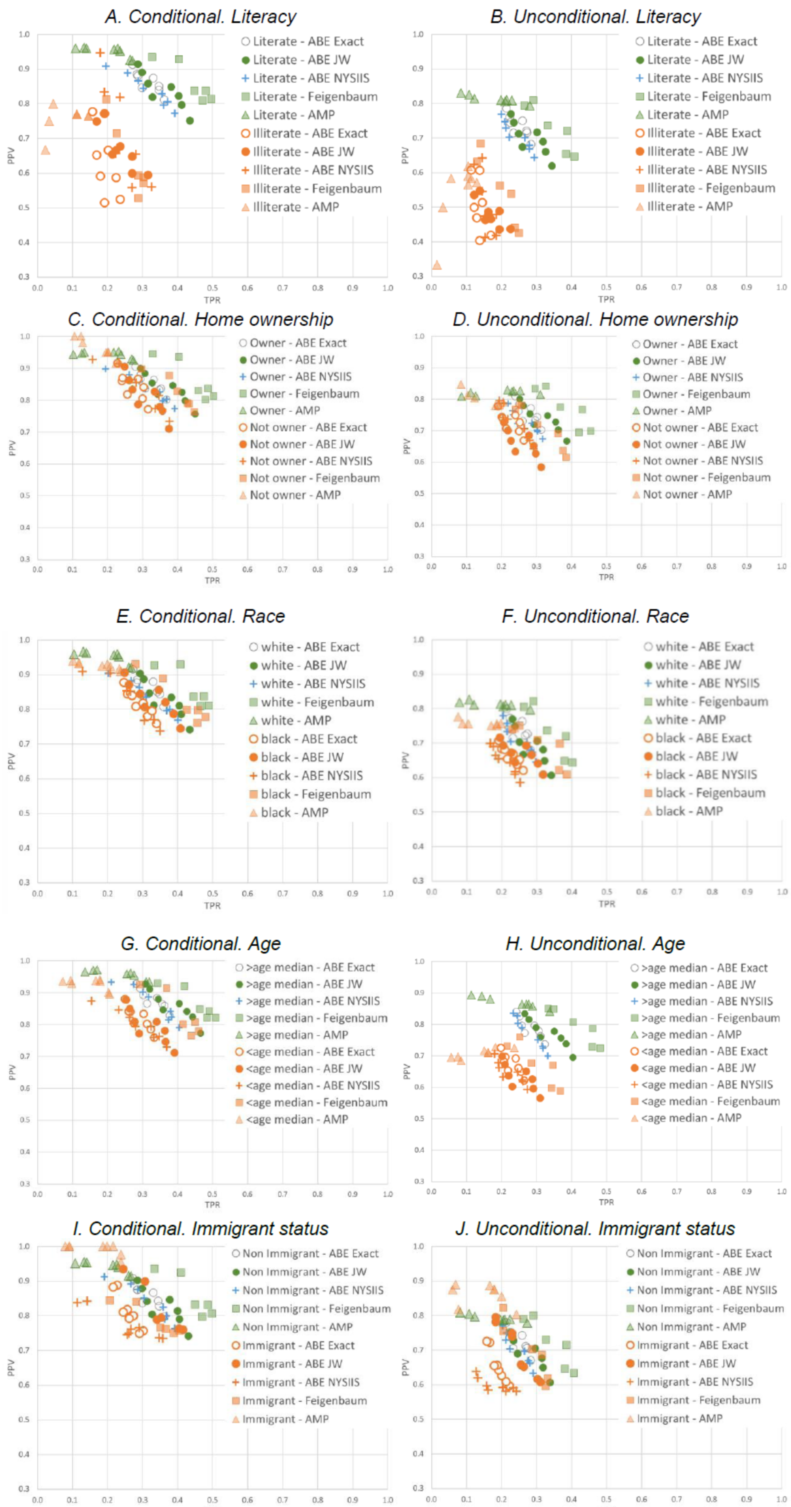

Notes: Figure 3.4 reports results separately for subgroups: panels A-B literate versus illiterate veterans, panels C-D homeowners and renters, panels E-F Black and White veterans, panels G-H above and below median age, and panels I-J native born and immigrants. Source: Authors' calculations from the Oldest Old Early Indicators Data.

Figure 3.5: Links in the CLP and MLP project

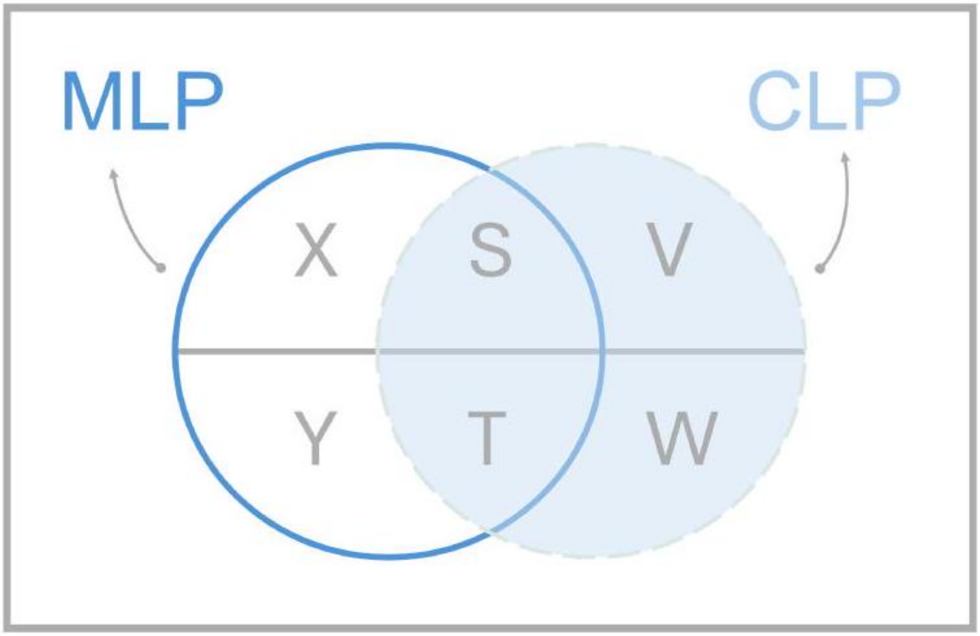

Notes: The square represents the universe of individuals. The square represents the universe of individuals. The left circle contains observations linked by the Multigenerational Longitudinal Panel (MLP), while the right circle contains all observations linked by the Census Linking Project (CLP). Area S represents links made by both projects where they agree. T represents links made by both projects but where they disagree. X and Y denote MLP-only links, with X being correct and Y incorrect. Similarly, V and W denote CLP-only links, where V are correct and W incorrect.

Figure 3.6: Probability Score Distributions for Links Made by Different Algorithms, 1930-1940

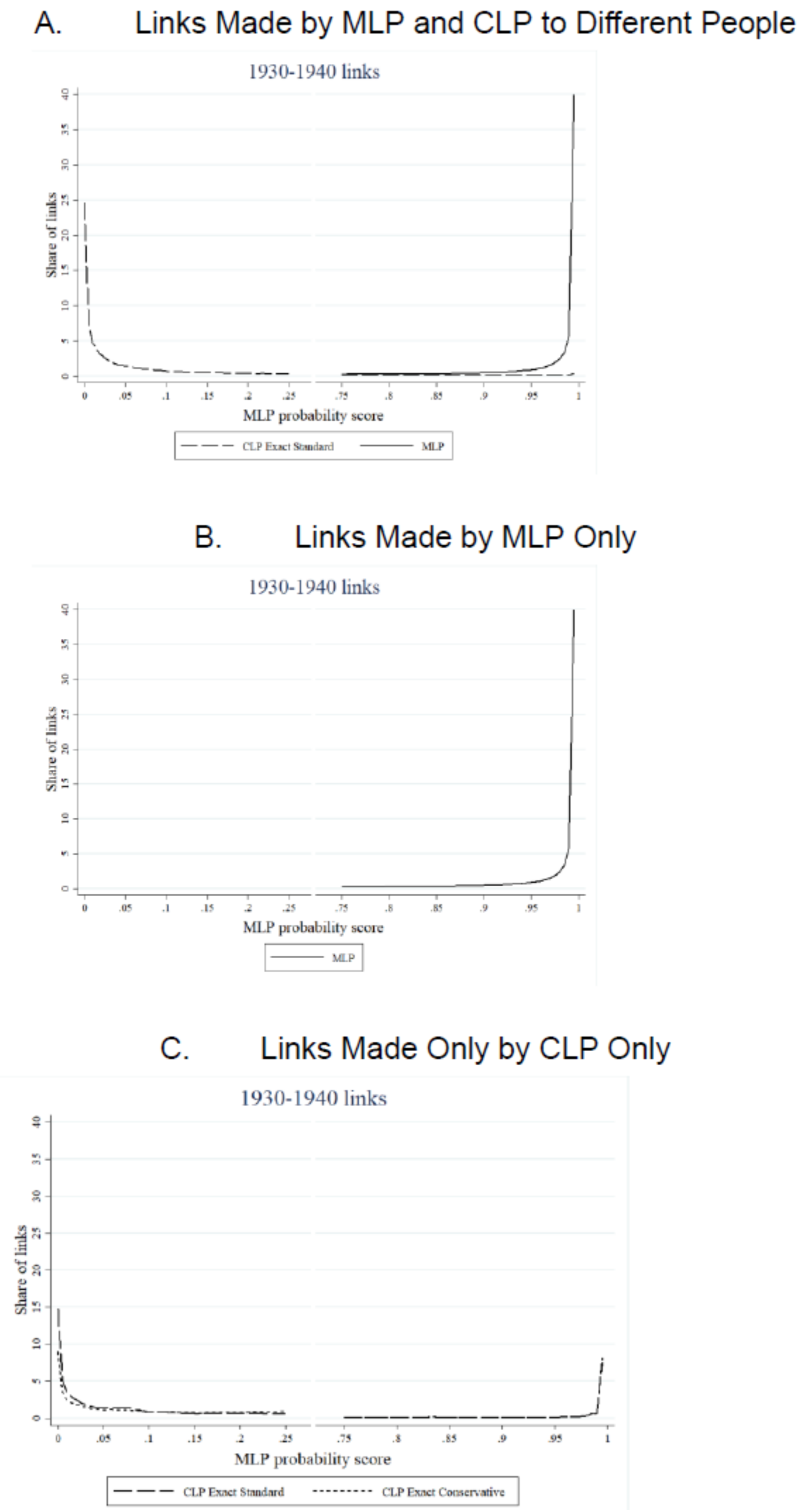

Notes: The x-axis only shows probability scores between $0-0.25$ and $0.75-1$ for illustrative purposes, including over 95% of cases.

Figure 3.7: Probability Score Distribution of MLP and CLP ES Links, by Race and Nativity

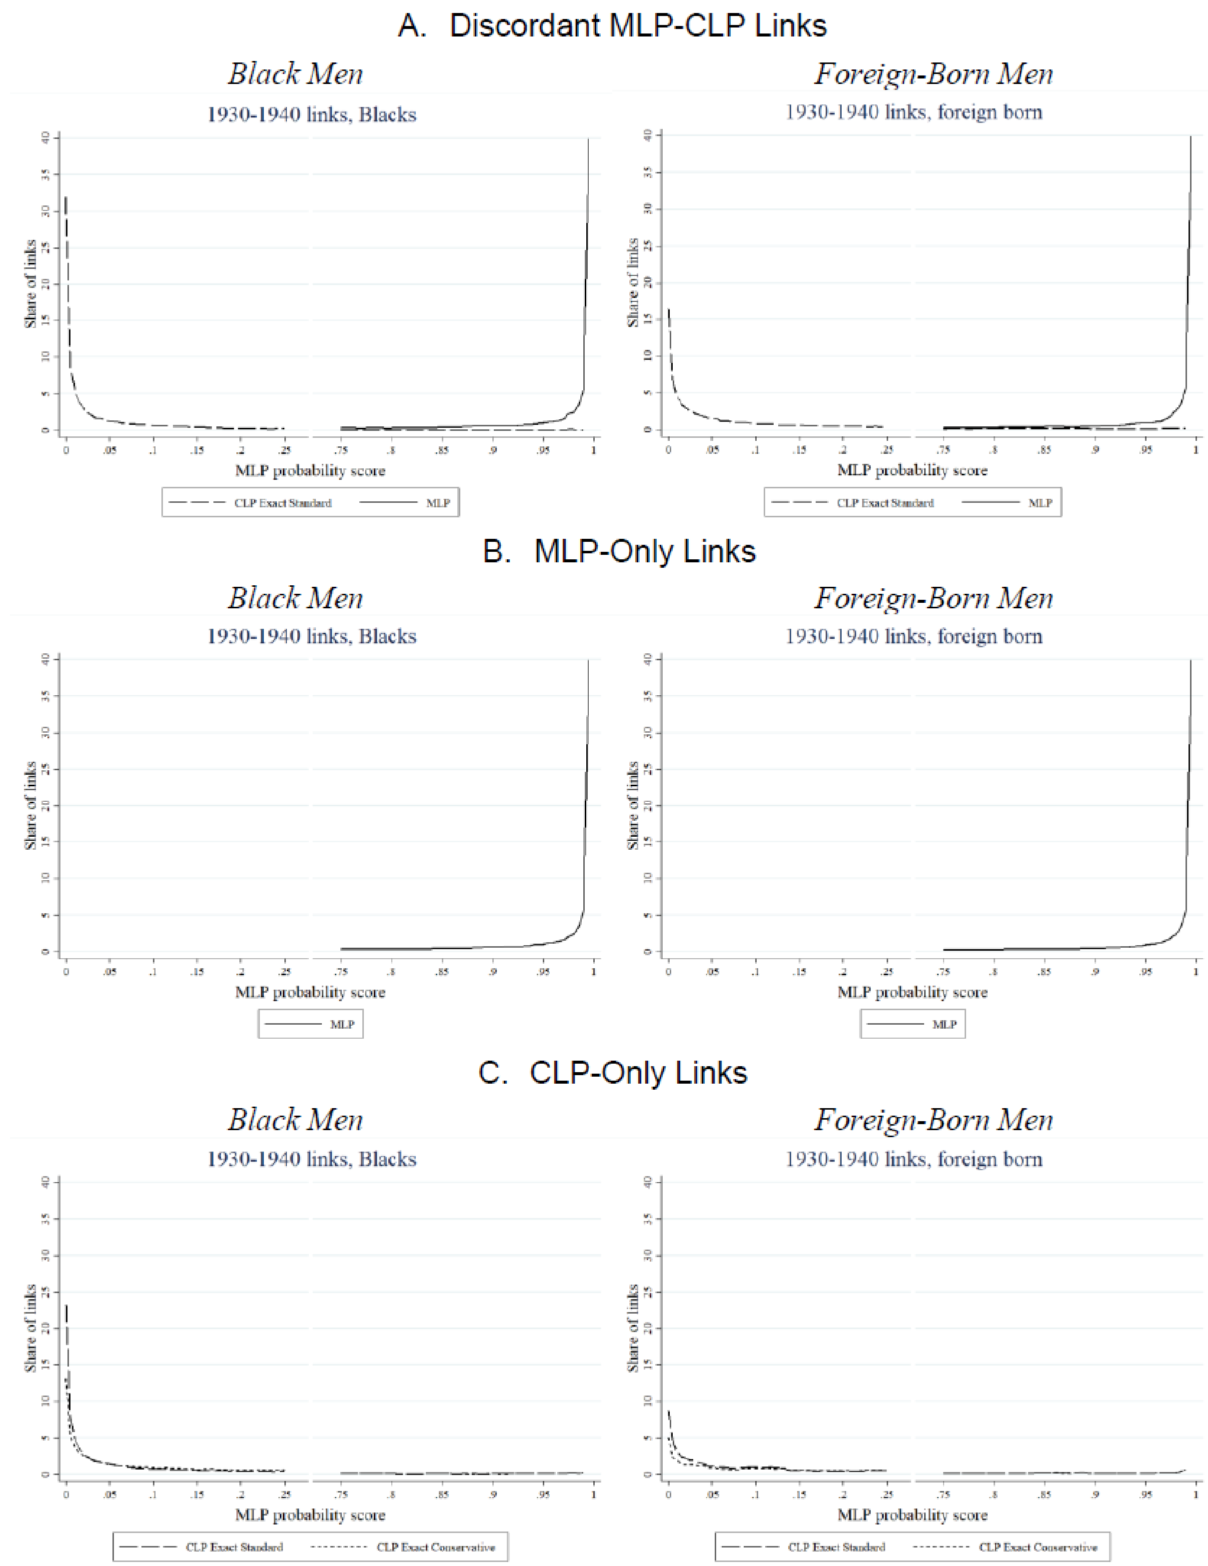

Notes: The x-axis only shows probability scores between $0-0.25$ and $0.75-1$ for illustrative purposes, including over 95% of cases.

7. Tables

Table 3.1: A Comparison of MLP and CLP Links

	1910-1940 Census Links			1920-1940 Census Links			1930-1940 Census Links		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
A. Total links between census years									
MLP links	10,436,600			15,842,252			27,804,653		
CLP links									
Exact Standard	8,253,178			11,917,909			16,995,375		
NYSIIS Standard	9,388,049			13,413,169			18,804,068		
Race NYSIIS Standard	9,512,800			13,602,458			19,112,277		
Exact Conservative	5,820,044			8,621,181			12,563,664		
NYSIIS Conservative	5,910,816			8,661,147			12,373,384		
Race NYSIIS Conservative	6,075,354			8,903,303			12,746,566		
Union of MLP and CLP links	14,850,598			21,780,653			33,112,133		
B. Links common to MLP & CLP									
	Total links	Total linked to different person	Share (2)/(1)	Total links	Total linked to different person	Share (5)/(4)	Total links	Total linked to different person	Share (8)/(7)
Exact Standard	5,683,857	182,334	0.032	8,479,165	376,819	0.044	14,001,333	661,947	0.047
NYSIIS Standard	6,081,590	279,346	0.046	8,801,121	512,409	0.058	14,743,617	811,712	0.055
Race NYSIIS Standard	6,189,919	274,167	0.044	8,967,019	502,602	0.056	15,039,866	794,635	0.053
Exact Conservative	4,677,663	43,930	0.009	6,895,279	96,978	0.014	11,170,597	219,523	0.020
NYSIIS Conservative	4,531,845	46,303	0.010	6,423,764	84,035	0.013	10,511,345	179,479	0.017
Race NYSIIS Conservative	4,661,299	47,015	0.010	6,604,826	82,803	0.013	10,833,420	178,022	0.016
C. Links made by MLP or CLP (but not both)									
		Total linked to different person	Share of all links (panel A)		Total linked to different person	Share of all links (panel A)		Total linked to different person	Share of all links (panel A)
Links in MLP, not found by any CLP link		3,357,926	0.322		5,460,943	0.345		10,414,319	0.375
Exact Standard		2,569,321	0.311		3,438,744	0.289		2,994,042	0.176
NYSIIS Standard		3,306,459	0.352		4,612,048	0.344		4,060,451	0.216
Race NYSIIS Standard		3,322,881	0.349		4,635,439	0.341		4,072,411	0.213
Exact Conservative		1,142,381	0.196		1,725,902	0.200		1,393,067	0.111
NYSIIS Conservative		1,378,971	0.233		2,237,383	0.258		1,862,039	0.150
Race NYSIIS Conservative		1,414,055	0.233		2,298,477	0.258		1,913,146	0.150

Notes: This table presents data for census-linked periods 1910-1940, 1920-1940, and 1930-1940. The dataset is restricted to men. Source: Multigenerational Longitudinal Panel, MLP (Helgertz et al., 2022), and the Census Linking Project, CLP (Abramitzky et al., 2022).

Table 3.2: A Comparison of MLP and CLP Links, by MLP Link Score

1930-1940 Census Links

	(1)	(2)	(3)	(4)
A. Links in MLP & CLP	Total links	Total linked to different person	Average link probability (for MLP-CLP agreements)	Difference in MLP-CLP average link probability (among links in (2))
Exact Standard	14,001,333	661,947	0.949	0.749
NYSIIS Standard	14,743,617	811,712	0.947	0.775
Race NYSIIS Standard	15,039,866	794,635	0.947	0.765
Exact Conservative	11,170,597	219,523	0.951	0.653
NYSIIS Conservative	10,511,345	179,479	0.948	0.715
Race NYSIIS Conservative	10,833,420	178,022	0.948	0.703
B. Links made by MLP or	Total links	Share of all links	MLP link probability	Share of links with higher MLP link probability
Links in MLP, not linked by any CLP variant	10,414,319	0.375	0.920	–
Links in CLP, not linked by MLP				
Exact Standard	2,994,042	0.176	0.238	0.275
NYSIIS Standard	4,060,451	0.216	0.206	0.277
Race NYSIIS Standard	4,072,411	0.213	0.213	0.256
Exact Conservative	1,393,067	0.111	0.278	0.166
NYSIIS Conservative	1,862,039	0.150	0.205	0.158
Race NYSIIS Conservative	1,913,146	0.150	0.210	0.144

Notes: This table presents data for the census-linked period 1930-1940, restricted to men. Link probability scores range from 0 to 1 and are calculated based on the similarity of the record's name, age, race, birth state, household co-residents, and geography. Panel A, column (3) presents the average link probability score among the links made by both the Multigenerational Longitudinal Panel (MLP) and the Census Linking Project (CLP) for the same person in 1940 (i.e., there is agreement). Panel A, column (4) presents the difference in percentage points between the average link probability of the links made by MLP and CLP among the set of disagreements. Panel B, column (2) presents the fraction of links made only by one of the two methods, using as the denominator the total number of MLP links or CLP links by method, not the union of all methods. Source: Multigenerational Longitudinal Panel (MLP) (Helgertz et al., 2022) and Census Linking Project (CLP) (Abramitzky et al., 2022).

Table 3.3: A Comparison of MLP and CLP Links, by Race and Nativity Subgroup

	1930-1940 Census Links									
	Black men					Foreign-born men				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
A. Total links between census years										
MLP links	1,533,946					2,582,182				
CLP links										
Exact Standard	941,689					1,388,583				
NYSIIS Standard	1,102,982					1,726,212				
Race NYSIIS Standard	1,241,548					1,719,835				
Exact Conservative	547,878					1,018,299				
NYSIIS Conservative	610,857					1,103,671				
Race NYSIIS Conservative	791,133					1,099,399				
Union of MLP and CLP links	2,251,536					3,158,563				
B. Links made by MLP & CLP	Total links	Total linked to different person	Share (2)/(1)	Average link probability (for MLP-CLP agreements)	Difference in MLP-CLP average link probability (for links in col. 2)	Total links	Total linked to different person	Share (7)/(6)	Average link probability (for MLP-CLP agreements)	Difference in MLP-CLP average link probability (for links in col. 2)
Exact Standard	592,729	65,067	0.110	0.924	0.777	1,103,766	74,253	0.067	0.952	0.729
NYSIIS Standard	653,521	71,263	0.109	0.925	0.781	1,257,026	105,152	0.084	0.947	0.759
Race NYSIIS Standard	777,374	65,601	0.084	0.922	0.747	1,253,864	105,008	0.084	0.947	0.759
Exact Conservative	416,464	16,576	0.040	0.930	0.715	891,829	23,961	0.027	0.953	0.621
NYSIIS Conservative	433,571	12,942	0.030	0.930	0.731	896,797	21,849	0.024	0.948	0.708
Race NYSIIS Conservative	577,630	13,029	0.023	0.926	0.690	894,461	21,861	0.024	0.948	0.708
C. Links made by MLP or CLP (but not both)	Total links	Share of all links in panel A	Average link probability	Share of links with higher MLP link probability		Total links	Share of all links in panel A	Average link probabilit y	Share of links with higher MLP link probability	
Links in MLP, not linked by any CLP variant	662,634	0.432	0.883	--		1,126,237	0.436	0.915	--	
Links in CLP, not linked by MLP										
Exact Standard	348,960	0.371	0.117	0.344		284,817	0.205	0.289	0.280	
NYSIIS Standard	449,461	0.407	0.113	0.327		469,186	0.272	0.221	0.302	
Race NYSIIS Standard	464,174	0.374	0.131	0.279		465,971	0.271	0.222	0.302	
Exact Conservative	131,414	0.240	0.152	0.187		126,470	0.124	0.347	0.171	
NYSIIS Conservative	177,286	0.290	0.133	0.171		206,874	0.187	0.211	0.182	
Race NYSIIS Conservative	213,503	0.270	0.144	0.147		204,938	0.186	0.213	0.182	

Notes: This table presents data for the census-linked period 1930-1940, restricted to Black men (columns 1 to 5) and foreign-born men (columns 6 to 10). Link probability scores range from 0 to 1 and are calculated based on the similarity of the record's name, age, race, birth state, household co-residents, and geography. Panel B, columns (4) and (9) present the average link probability score among the links made by both the Multigenerational Longitudinal Panel (MLP) and the Census Linking Project (CLP) for the same person in 1940 (i.e., there is agreement). Panel B, columns (5) and (10) present the difference in percentage points between the average link probability of the links made by MLP and CLP among the set of disagreements. Source: Multigenerational Longitudinal Panel (MLP) (Helgertz et al., 2022) and Census Linking Project (CLP) (Abramitzky et al., 2022).

Chapter 4

Appendix and Supplementary Material

1. Appendix to "Do Robots Necessarily Displace Workers? Evidence from Mexican Local Labor Markets"

Figure A.1: Trends in the Mexican formal sector

(a) Labor outcomes by wage group

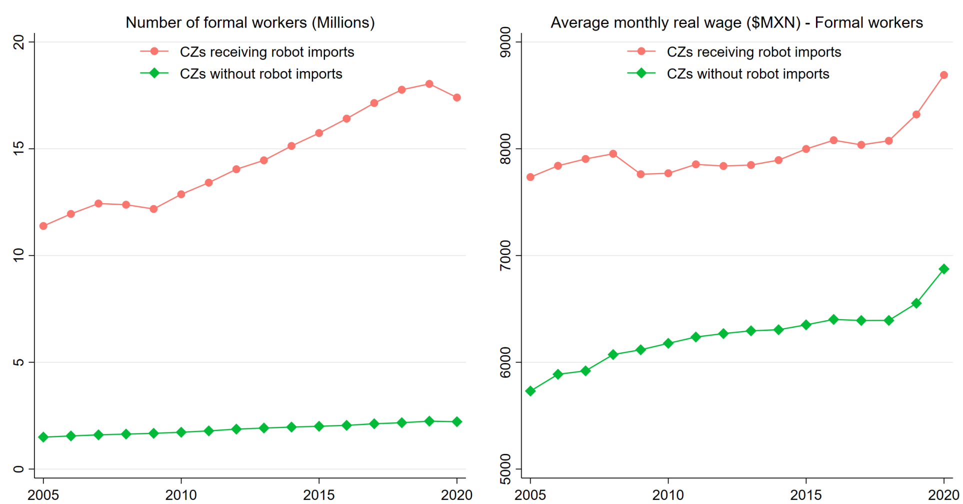

(b) Labor outcomes by treatment status

Notes: Figure A.1 illustrates the trends in the number of formal workers (measured in millions) and average monthly real wages (adjusted to 2010 Mexican pesos) over time. In Panel (a), the sample of formal workers is divided into three distinct wage categories: those earning two minimum wages or less (low wage), those earning between two and five minimum wages (medium wage), and those earning more than five minimum wages (high wage). In Panel (b), the sample is divided into two groups: CZs receiving industrial robot imports at any point during the analysis period and CZs not receiving such imports. *Data sources*: IMMS public data, Mexican customs records from Secretaría de Economía.

Figure A.2: Characterization of Mexican manufacturing sector

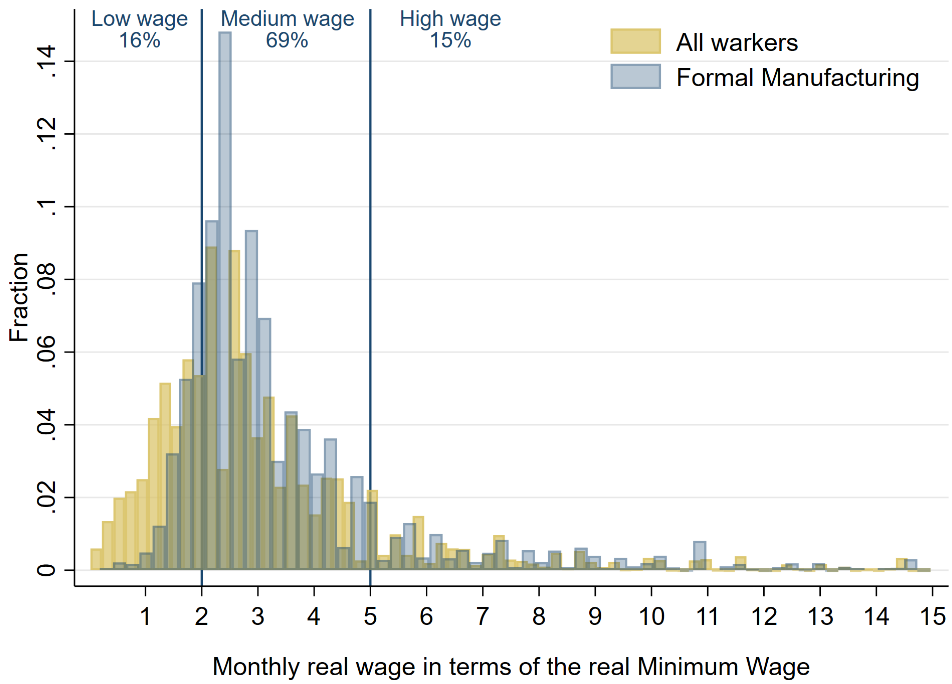

(a) Fraction of workers for each wage bin

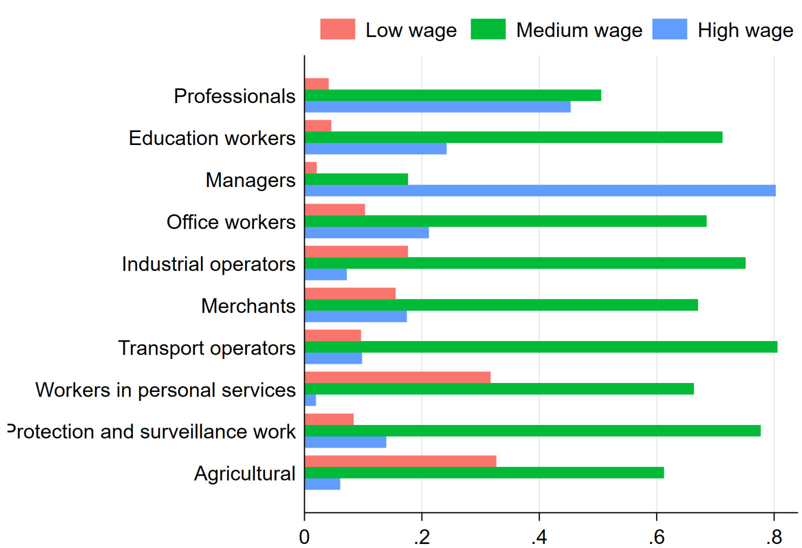

(b) Fraction of wage group for each occupation

Notes: Panel (a) of Figure A.2 shows the histogram of the log average monthly real wages (adjusted to 2005 Mexican pesos) in terms of the monthly minimum wage for either all workers or formal workers. The vertical lines divide formal workers into three distinct wage categories: those earning two minimum wages or less (low wage), those earning between two and five minimum wages (medium wage), and those earning more than five minimum wages (high wage). Panel (b) shows for each occupation the fraction of formal workers in each wage group. For example, 73% of industrial operators fall into the medium-wage group. *Data sources*: Own elaboration based on ENOE 2005, 4th quarter.

Figure A.3: Characterization of Mexican manufacturing sector

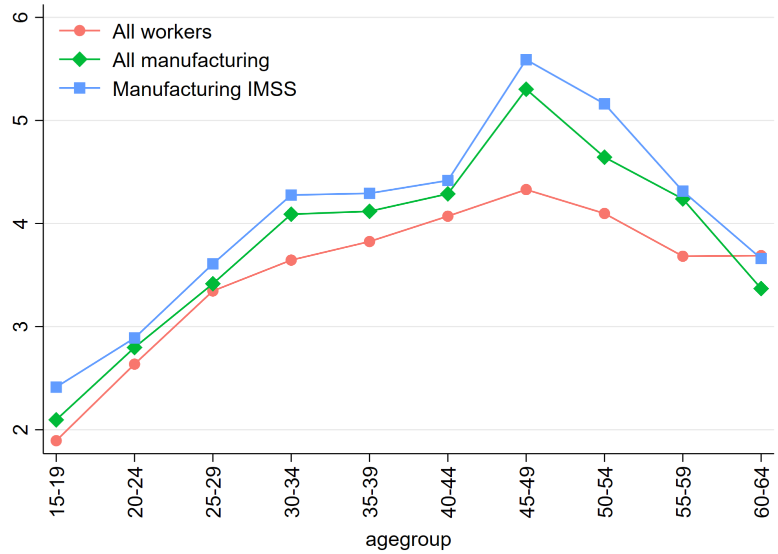

(c) Average monthly wage by age

Notes: Figure A.3 illustrates the average wage trends for different groups of workers across various age brackets. The vertical axis of the figure represents the average monthly wages, measured in terms of the monthly minimum wage, for three distinct groups: all workers (across all industries), all manufacturing workers, and all formal manufacturing workers (those registered in IMSS). *Data sources*: Own elaboration based on ENOE 2005, 4th quarter.

Figure A.4: Industrial robots in Mexico: Large increase but early stage

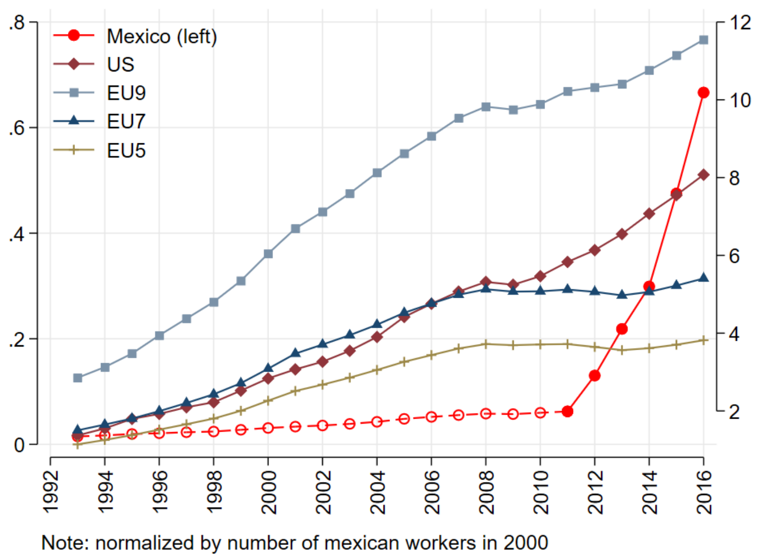

(a) Automation by country

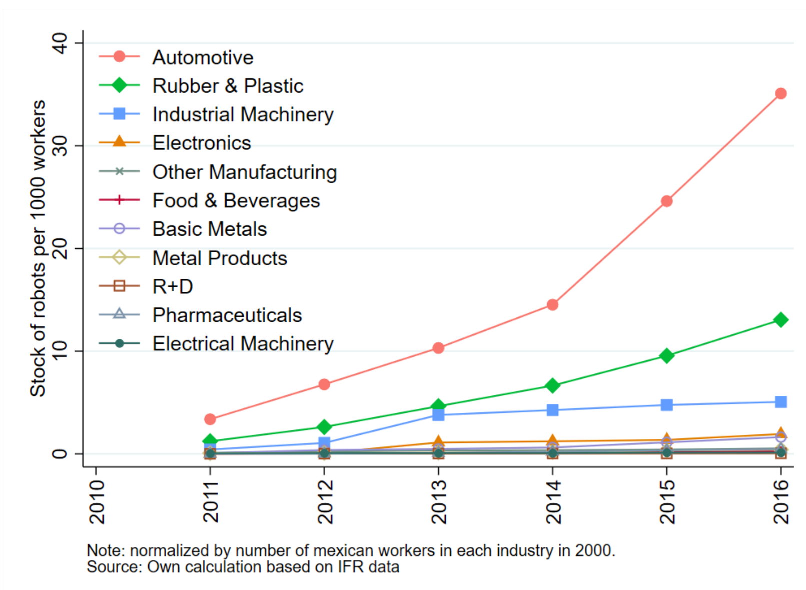

(b) Automation by industry

Notes: Panel (a) of Figure **??** presents a comparison of the cumulative stock of industrial robots in Mexico with that in other developed countries, specifically the United States (US) and European Union (EU) countries. It is important to note that IFR data for Mexico and the US are available only for years from 2011. To create the trend data before 2011, I impute values using the aggregate stock for North America and distributed among Mexico and the US based on their relative sizes in 2011. Panel (b) of the figure depicts the same cumulative robot stock for Mexico, but it further breaks down this stock for different 3-digit industries within the manufacturing sector. *Data source*: IFR.

Figure A.5: Number of CZs that started receiving robot imports by year

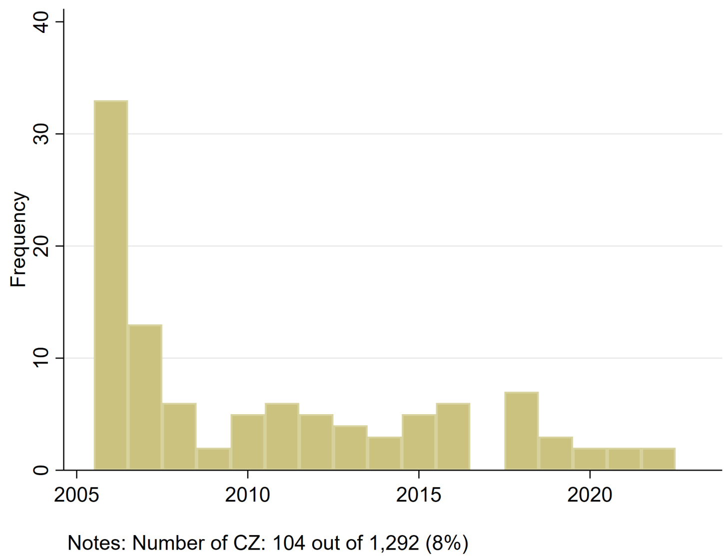

Notes: Figure A.5 shows the count of CZs that started receiving industrial robot imports between 2006 and 2022. *Data sources*: Mexican customs records from Secretaría de Economía.

Figure A.6: Evolution of number of formal workers by treatment cohort

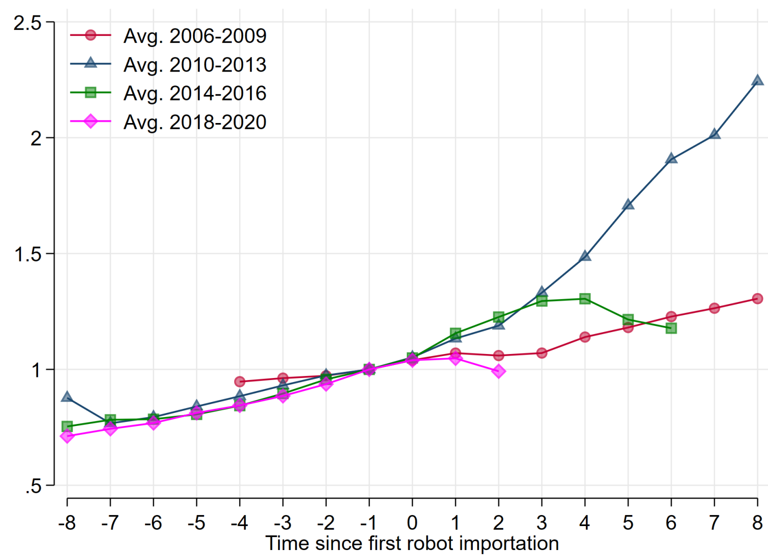

(a) CZ cohort average

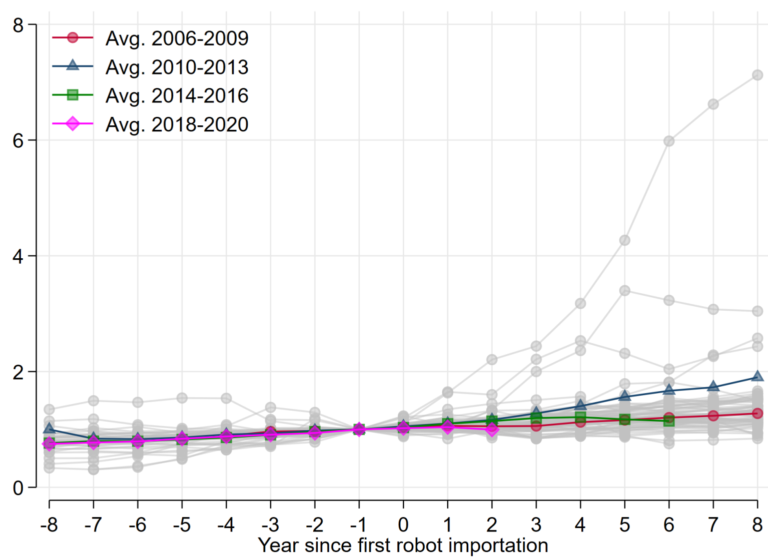

(b) Including every CZ (in gray)

Notes: Figure A.6 shows the raw means for the number of formal workers by treatment cohort over event-time. Estimates are relative to the outcomes one year before the event (period -1). Panel (a) shows cohort averages across CZs. Panel (b) shows the same figures as in Panel (a) and overlaps the raw means for each CZ separately in gray. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Figure A.7: Heterogeneity by wage group

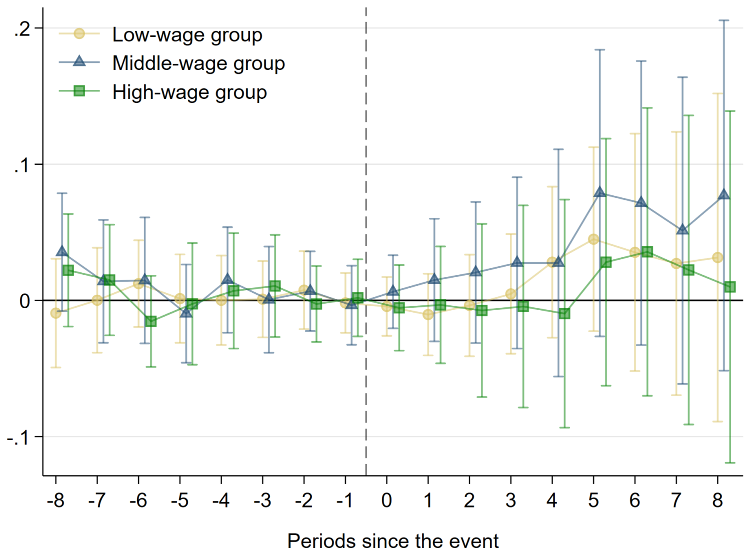

Notes: Figure A.7 plots the estimated $\theta(e)$ event-study coefficients using equation 1.2 for three wage categories of formal workers: those earning two minimum wages or less (low wage), those earning between two and five minimum wages (medium wage), and those earning more than five minimum wages (high wage). The outcome variable is log number of formal workers. The event is defined as a first-time industrial robot importation. The vertical lines reflect the 95% uniform confidence intervals, which are robust to multiple hypothesis testing. Sample of not-yet-but-eventually-treated CZs for the period 2005–2019, including cohorts that started importing robots during 2007–2022. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Figure A.8: Heterogeneity in dynamic effects on formal employment and wages

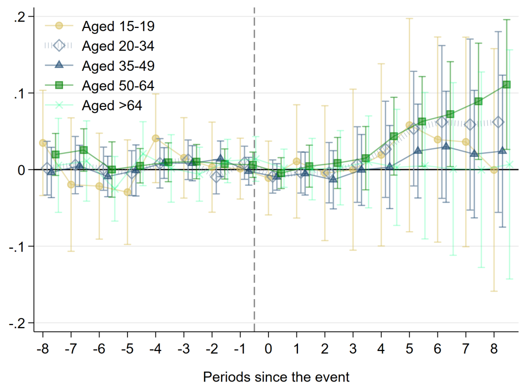

(a) Ln formal workers by **age group**

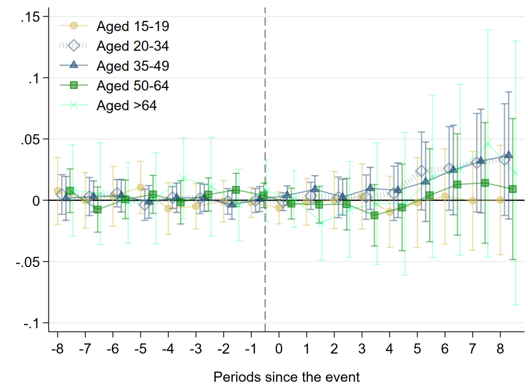

(b) Ln formal wage by **age group**

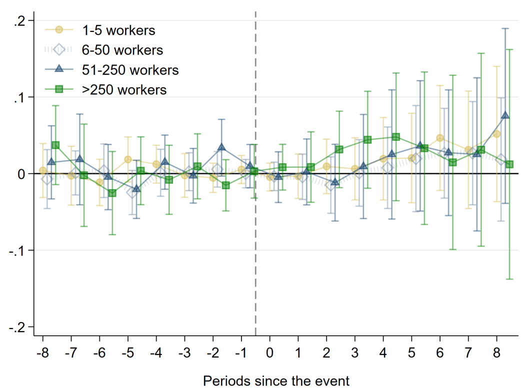

(c) Ln formal workers by **firm size**

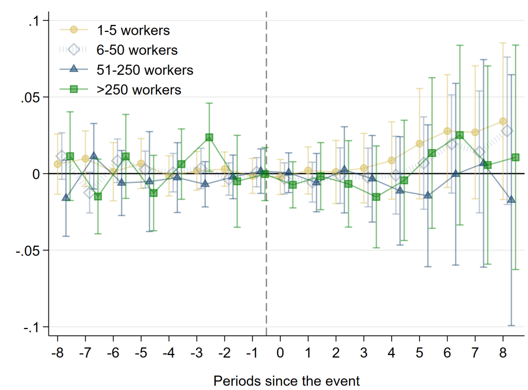

(d) Ln formal wage by **firm size**

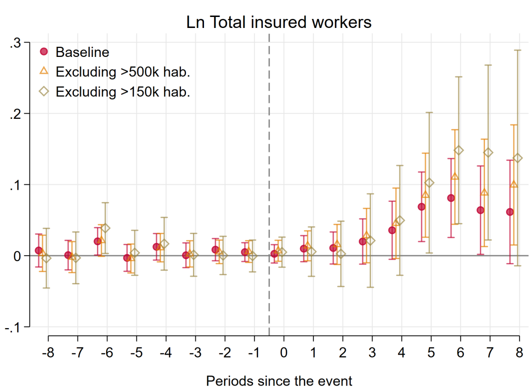

(e) Ln formal workers by **CZ size**

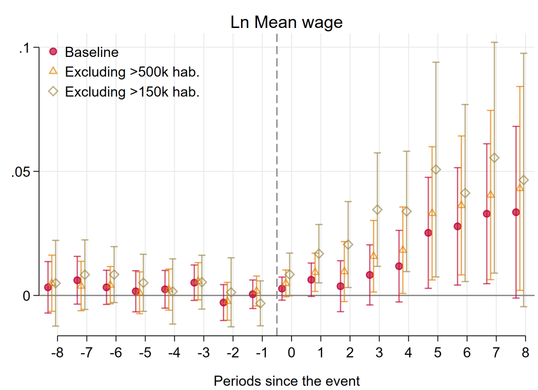

(f) Ln formal wage by **CZ size**

Notes: Figure A.8 plots the estimated $\theta(e)$ event-study coefficients using equation 1.2. The outcome variable is, in turn, the log number of formal workers (Panels (a), (c), and (e)) and the log average real monthly wage for formal workers, deflated to constant pesos using Mexican CPI (Panels (b), (d), and (f)). The event is defined as a first-time industrial robot importation. The vertical lines reflect the 95% uniform confidence intervals, which are robust to multiple hypothesis testing. Sample of not-yet-but-eventually-treated CZs for the period 2005–2019, including cohorts that started importing robots during 2007–2022. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Figure A.9: Alternative measures of robot adoption event

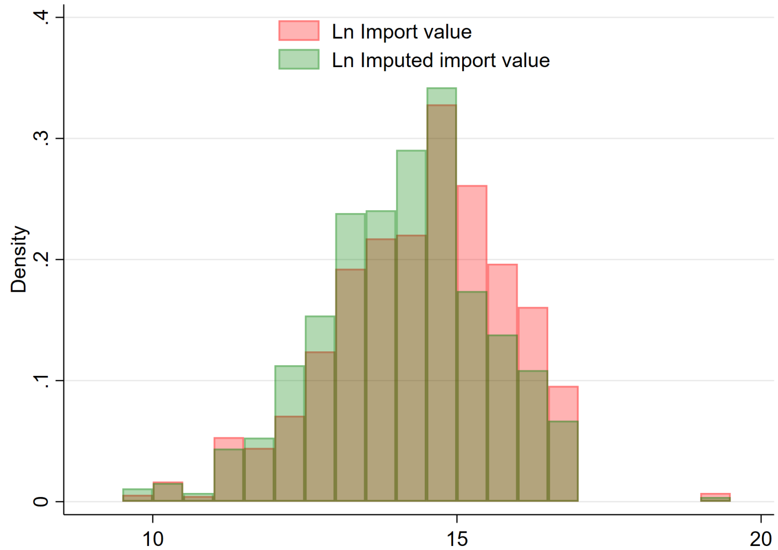

(a) Histogram of imputed and not-imputed values of robot imports

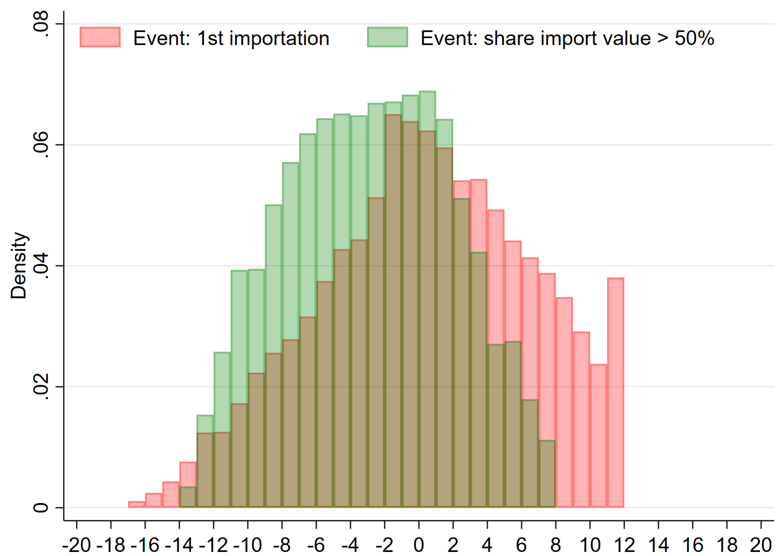

(b) Histogram of original event and alternative events based on import value

Notes: See main text for further explanation. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Figure A.10: Event-study methods

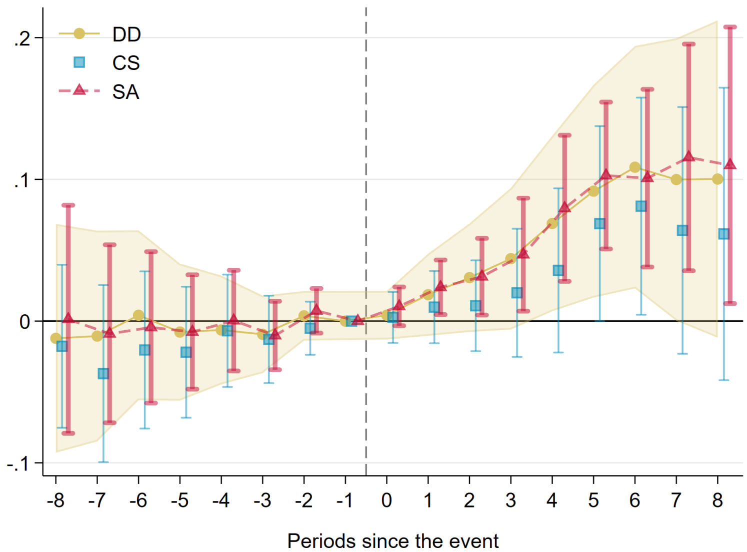

(a) Ln Number of formal workers

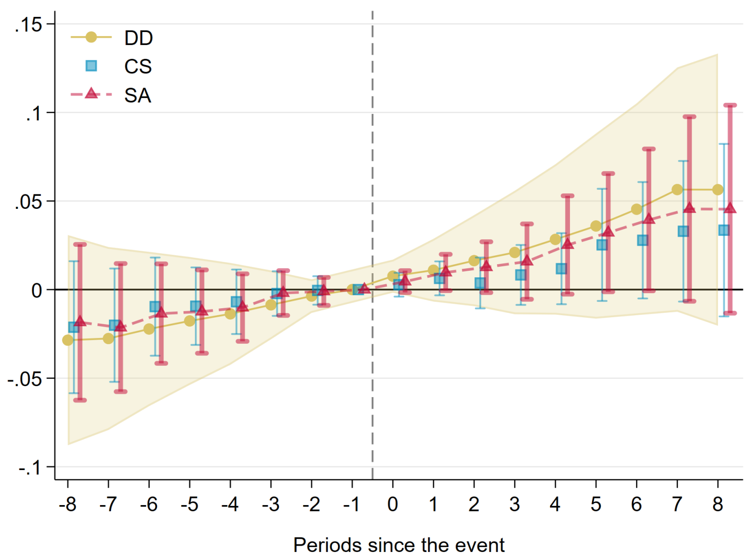

(b) Ln Average monthly wage – formal workers

Notes: Figure A.10 overlaps the estimated event-study coefficients using three different methods: CS (Callaway and Sant'Anna (2021)), SA (Sun and Abraham (2021)), and DD (two-way fixed effect model). Estimates are relative to the outcomes one year before the event (period -1). The outcome variable is, in turn, the log number of formal workers (Panel (a)) and log average real monthly wage for formal workers, deflated to constant 2010 pesos using Mexican CPI (Panel (b)). The event is defined as a first-time industrial robot importation. The shaded area and vertical lines reflect the 95% pointwise confidence intervals for DD and SA and uniform confidence intervals for CS. The clustering of standard errors is at the CZ level. Sample of not-yet-but-eventually-treated CZs for 2005–2019, including cohorts that started importing robots during 2007–2022. The unit of observation is the CZ–4-digit industry (N=100,619; 71 CZs). *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Figure A.11: Pretrends sensitivity analysis

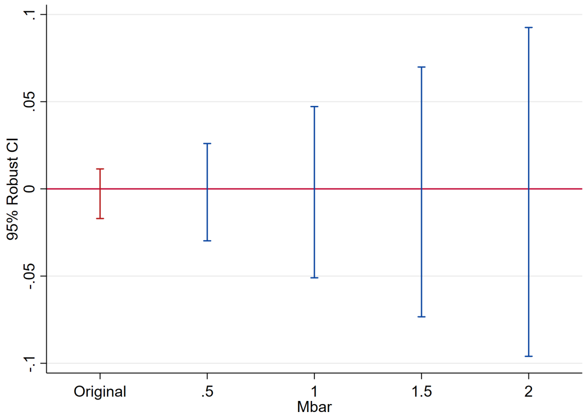

(a) Ln Num. formal workers: Short run

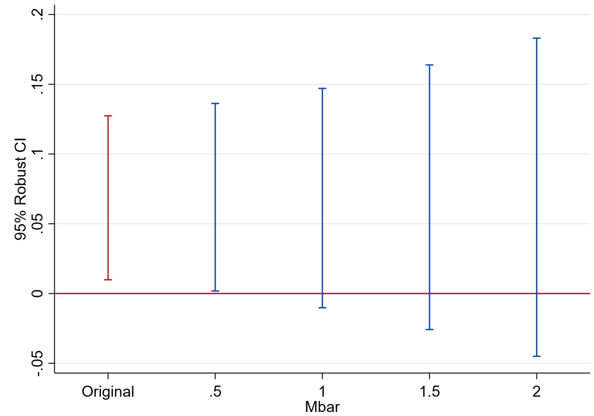

(b) Ln Num. formal workers: Long run

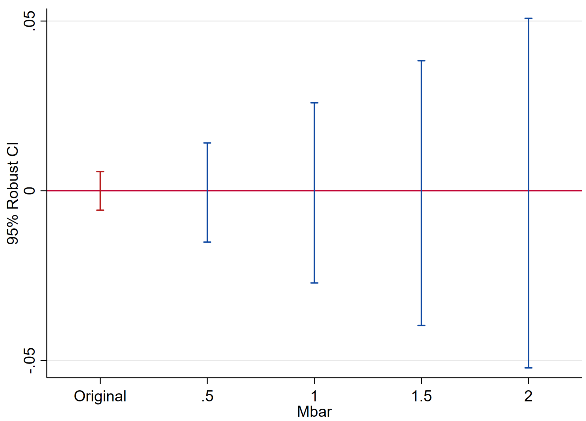

(c) Ln Avg. real formal wages: Short run

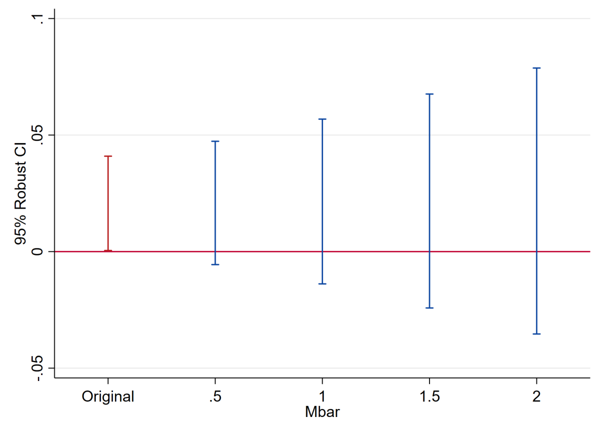

(d) Ln Avg. real formal wages: Long run

Notes: Figure A.11 shows the result of a sensitivity analysis proposed by Rambachan and Roth (2023) over the coefficients estimated under the CS method. This analysis estimates the post-treatment effect under different assumptions about parallel trend violations, demonstrating the extent to which the postevent trend differences would need to deviate from pretrends to render the estimates negative or statistically insignificant. It provides confidence intervals that enable the post-treatment departure from parallel trends to reach a magnitude of up to $Mbar$ times the maximum pretreatment deviation, with M being a variable that can take different values. The outcome variable is, in turn, the log number of formal workers (Panels (a) and (b)) and log average real monthly wage for formal workers, deflated to constant 2010 pesos using Mexican CPI (Panels (c) and (d)). The event is defined as a first-time industrial robot importation. Short term refers to the year of the event, and long term refers to six years after the event occurs. The clustering of standard errors is at the CZ level. Sample of not-yet-but-eventually-treated CZs for 2005–2019, including cohorts that started importing robots during 2007–2022. The unit of observation is the CZ–4-digit industry. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Figure A.12: TWFE models: DD & DDD

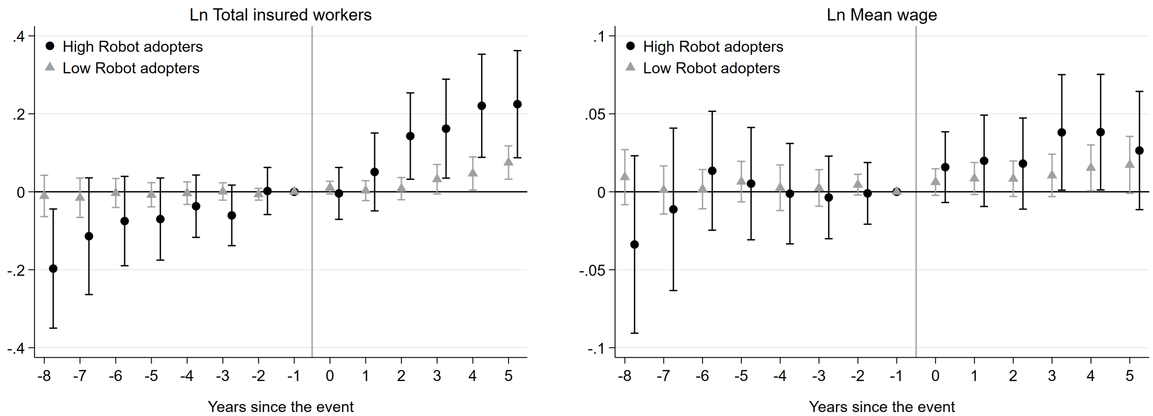

(a) DD estimates: CZ–3-dig. industry level

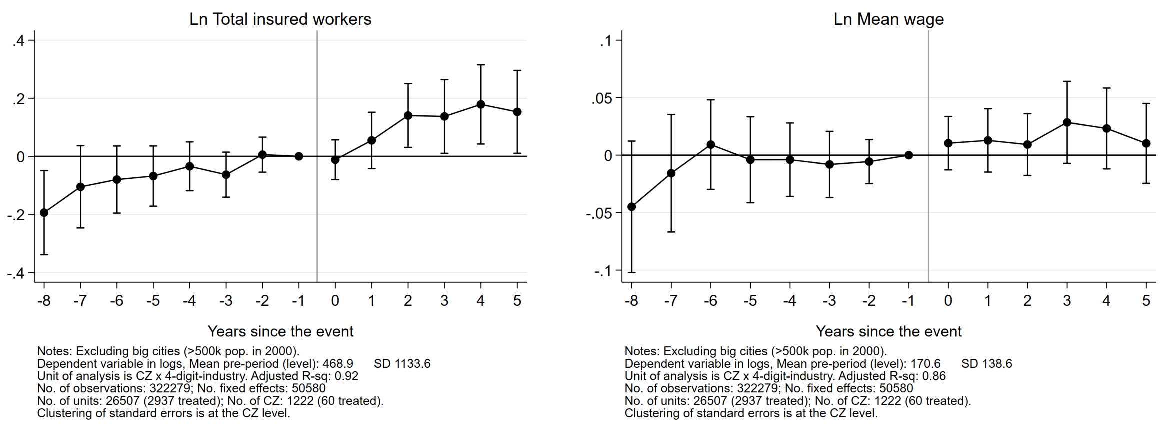

(b) DDD estimates: CZ–3-dig. industry level

Notes: Figure 1.5 displays the estimated event-study coefficients obtained using equation 1.2, a TWFE model. The unit of observation is the CZ–3-digit industry. In Panel (a), the results of the difference-in-differences specification are presented, showing overlapping estimates from two different subsamples: "high robot adopters", comprising manufacturing industries with higher rates of robot adoption observed in the IFR data, and "low robot adopters", which include the remaining nonmanufacturing industries. Panel (b) shows estimates from the triple difference-in-differences specification, where the third difference is introduced by contrasting high and low robot adopters. The outcome variable is, in turn, the log number of formal workers and log average real monthly wage for formal workers, deflated to constant 2010 pesos using Mexican CPI. The event is defined as a first-time industrial robot importation. The clustering of standard errors is at the CZ level. Sample of CZs with population smaller than 500 thousands inhabitants for the period 2005–2019, including cohorts that started importing robots during 2007–2022. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Figure A.13: SDD estimates for CZ #24050 (2011 cohort)

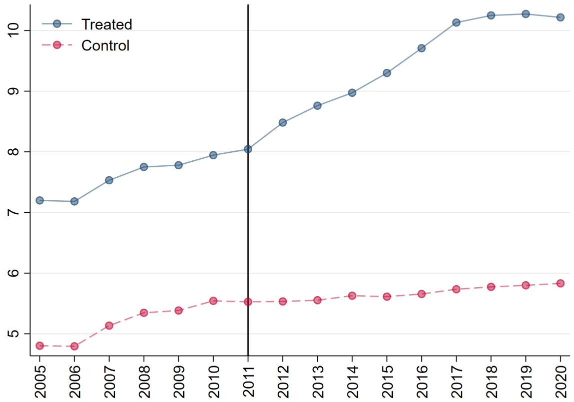

(a) Ln Num. formal workers: Trends

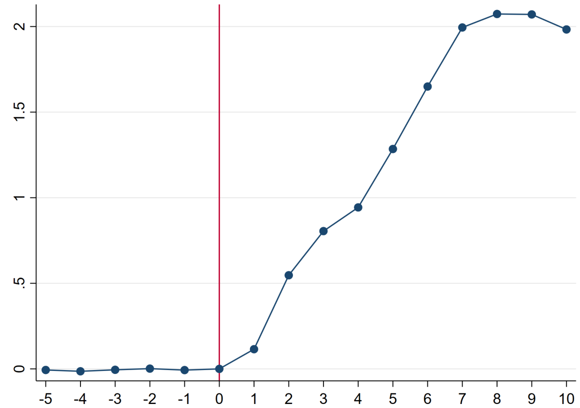

(b) Ln Num. formal workers: DD

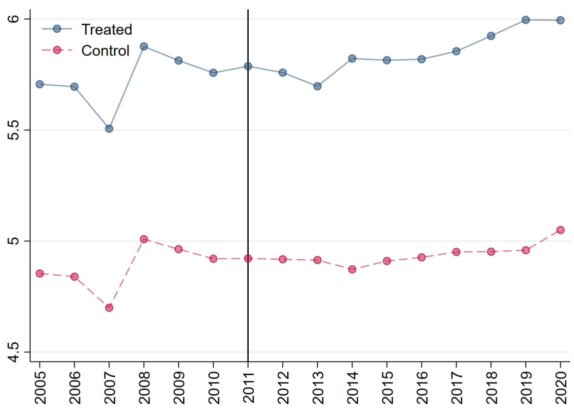

(c) Ln Avg. real formal wage: Trends

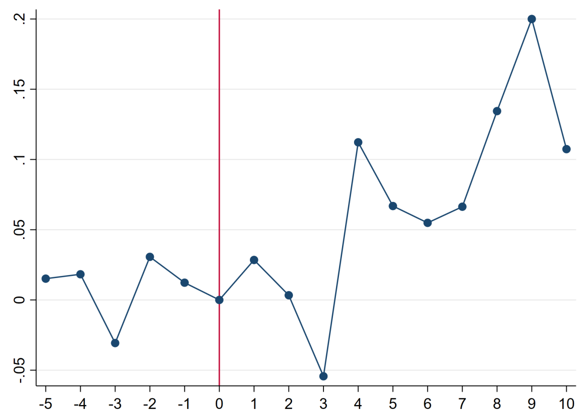

(d) Ln Avg. real formal wage: DD

Notes: Figure A.13 exhibits the estimated treatment effects using the synthetic difference-in-difference model as outlined by Arkhangelsky et al. (2021) for a specific CZ (belonging to 2011 cohort) that experiences a statistically significant positive formal employment and wage effect. The comparison group for these estimates is derived from the pool of all CZs that have never been treated, and it employs optimal weighting of units and pre-event periods in its construction. The dataset comprises a panel of CZs, balanced in calendar time, spanning from 2005 to 2019. The outcome variables are the logarithm of the number of formal workers and the logarithm of the average real monthly wage for formal workers, adjusted to constant 2010 pesos using Mexican CP. The event is defined as a first-time industrial robot importation. Panels (a) and (c) present the raw means of the treated CZ and the synthetic control. Panels (b) and (d) show the corresponding difference between treatment and control over event-time. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Figure A.14: Comparison of 2SLS estimates with results from Acemoglu and Restrepo (2020)

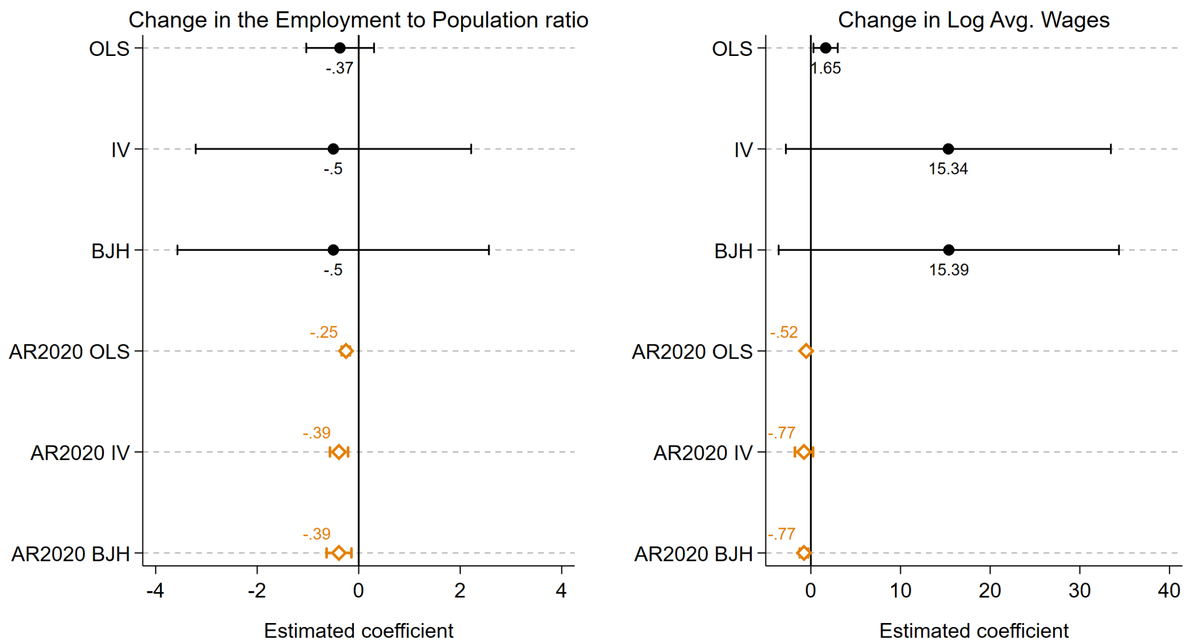

Figure A.14 compares OLS and IV estimates from Acemoglu and Restrepo (2020) for the US and my own replication for Mexico. BJH refers to the standard error estimator proposed by Borusyak et al. (2021). *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data. Mexican censuses 1990 and 200. IFR data. Faber (2020).

Figure A.15: Alternative event definition using robot imports value

(a) Ln Number of formal workers (b) Ln Avg. monthly wage – formal workers

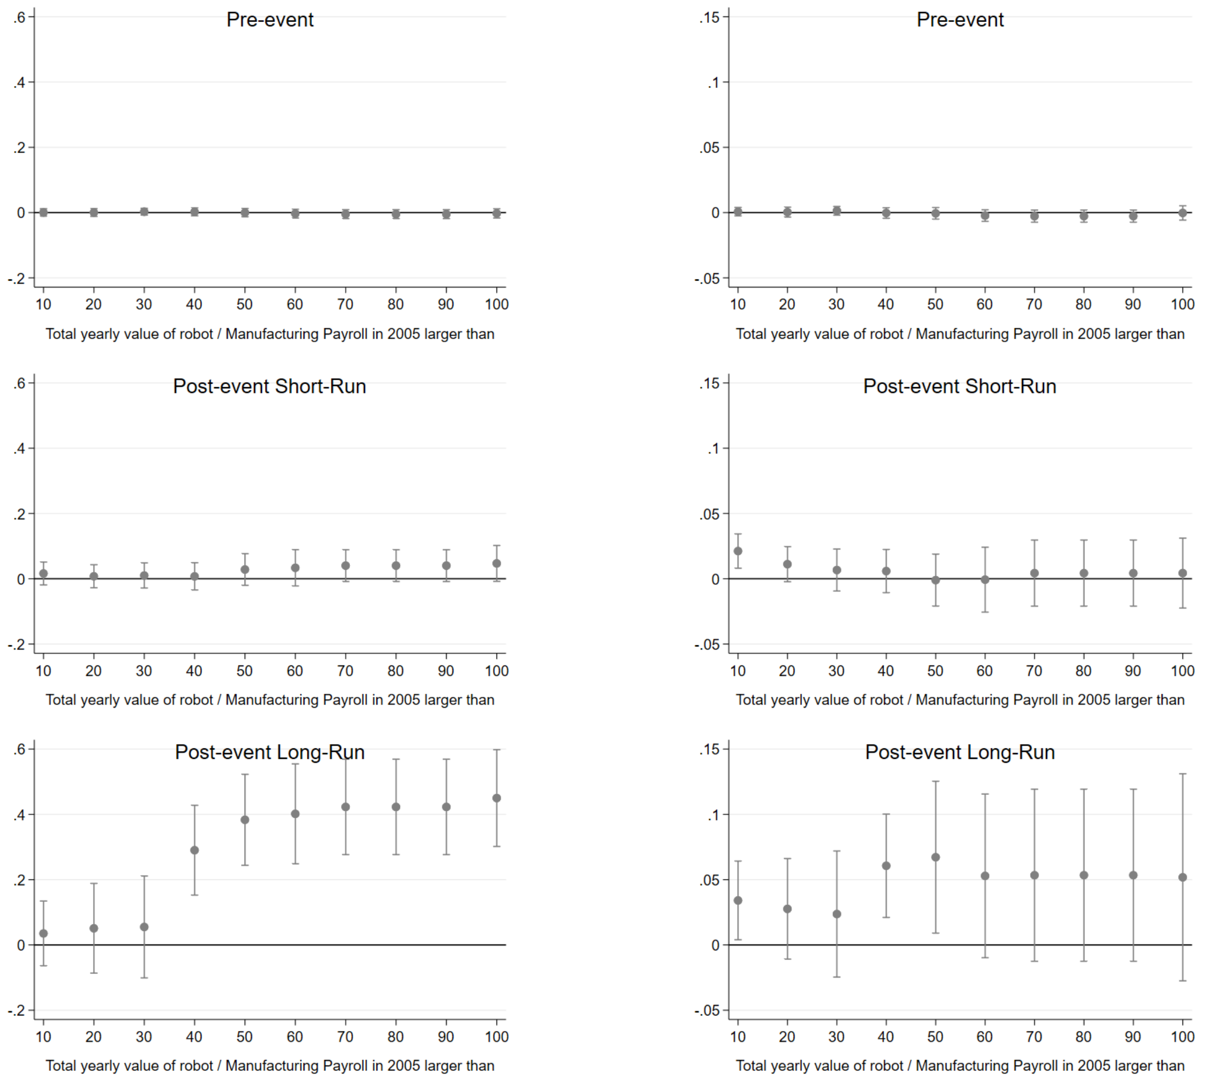

Notes: Figure A.15 plots the simple average of estimated $\theta(e)$ event-study coefficients using equation 1.2 (Callaway and Sant'Anna (2021) method) for different event-periods. The event is defined as the period where the cumulative total value of robot imports is larger than X% of the CZ's 2005 total wage bill (where X% varies on the horizontal axes). Pre-event: years -5 to -1; postevent short run: years 0 to 4; long-run: years 5 to 8. The outcome variable is, in turn, log number of formal workers (Panel (a)) and log average real monthly wage for formal workers, deflated to constant 2010 pesos using Mexican CPI (Panel (b)). Clustering of standard errors is at the CZ level. Sample of not-yet-but-eventually-treated CZs for the period 2005–2019, including cohorts that started importing robots during 2007–2022. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Figure A.16: Controlling for plant openings by large firms

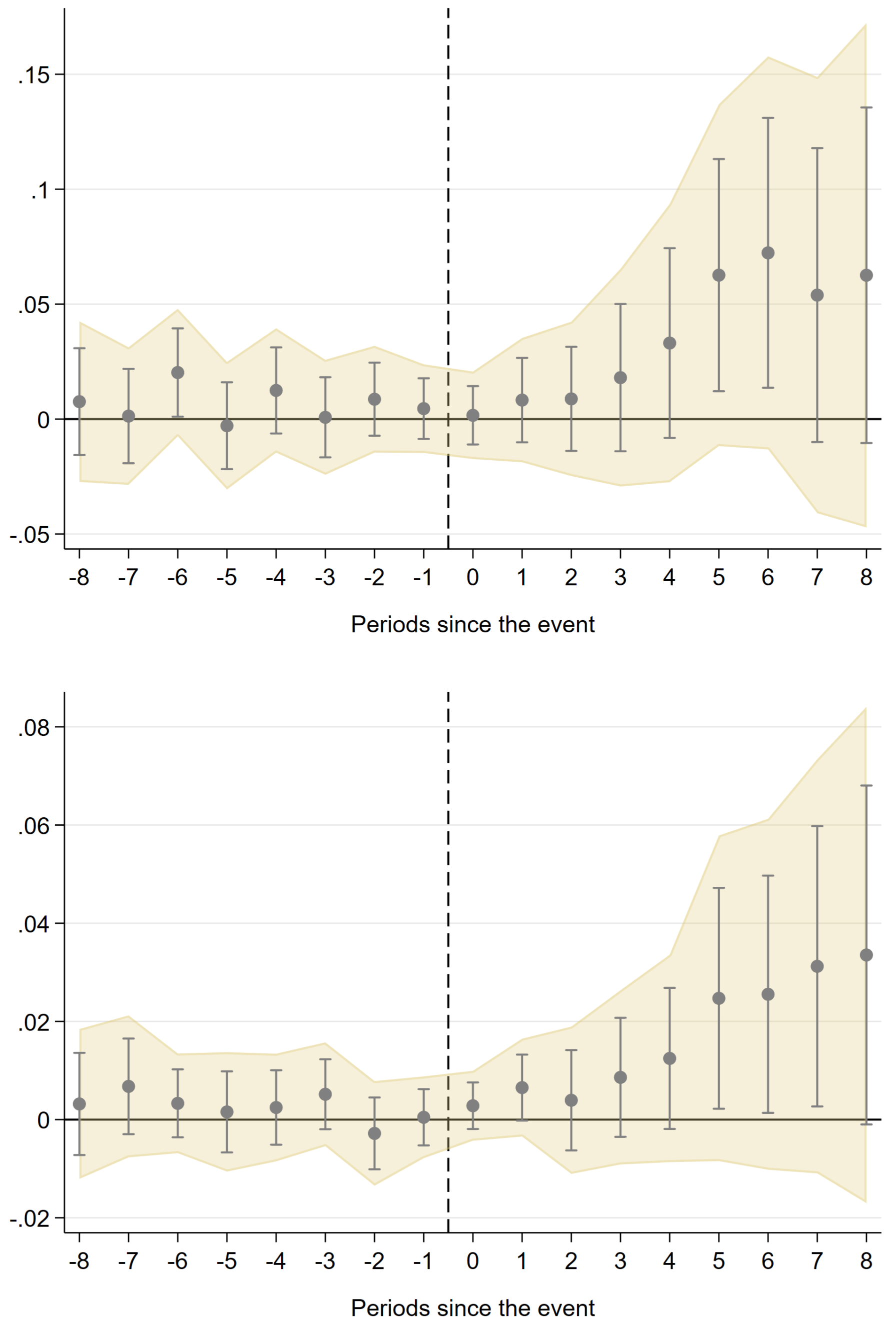

Notes: Figure A.16 shows the estimated event-study coefficients using a two-way fixed effect model, controlling for the opening of large (more than 250 workers) manufacturing firms (which survive up to the 2018 Mexican economic census). In the first graph, the outcome variable is the log number of formal workers, and in the second graph, it is the log average real monthly wage for formal workers, deflated to constant 2010 pesos using Mexican CPI. The clustering of standard errors is at the CZ level. The event is defined as a first-time industrial robot importation. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data. Mexican economic census 2018.

Table A.1: Characterization of mexican manufacturing sector by wage group

	All workers			Manufacturing IMSS workers		
	Low wage	Medium wage	High wage	Low wage	Medium wage	High wage
	(1)	(2)	(3)	(4)	(5)	(6)
A. By occupation						
Professionals	3%	7%	21%	1%	4%	19%
Education workers	2%	5%	17%	0%	0%	1%
Managers	0%	1%	8%	0%	0%	10%
Office workers	7%	14%	20%	6%	8%	19%
Industrial operators	25%	35%	15%	80%	74%	37%
Merchants	15%	10%	7%	4%	5%	8%
Transport operators	3%	7%	6%	1%	4%	3%
Workers in personal services	29%	11%	2%	7%	3%	1%
Protection and surveillance work	1%	5%	3%	0%	1%	1%
Agricultural	17%	5%	1%	0%	0%	1%
Total	100%	100%	100%	100%	100%	100%
N (mill.)	7.84	13.54	3.84	0.59	2.30	0.44
B. By age group						
15–19	19%	8%	1%	13%	7%	0%
20–24	18%	17%	6%	25%	20%	7%
25–29	12%	17%	13%	15%	19%	17%
30–34	11%	15%	17%	15%	16%	20%
35–39	10%	13%	17%	11%	13%	19%
40–44	8%	11%	17%	9%	11%	16%
45–49	6%	8%	14%	5%	6%	10%
50–54	5%	5%	8%	3%	4%	6%
55–59	4%	4%	4%	2%	3%	3%
60–64	3%	1%	3%	1%	1%	1%
65+	3%	1%	1%	0%	0%	0%
Total	100%	100%	100%	100%	100%	100%
N (mill.)	7.62	13.51	3.84	0.58	2.29	0.44

Notes: Panel (a) of Table A.1 shows the share of workers in each occupation within each wage group for the sample of all workers and the subsample of workers in the manufacturing sector who are registered in IMSS. Panel (b) shows similar statistics by age group. Wage categories: those earning two minimum wages or less (low wage), those earning between two and five minimum wages (medium wage), and those earning more than five minimum wages (high wage). *Data sources*: Own elaboration based on ENOE 2005, 4th quarter.

Table A.2: Growth in robot adoption by industry 2015–2010 (%)

	N workers	Change in stock of robots 2010-2015 (%)					
	MEX 2000	MEX	US	EU9	EU7	EU5	LA
All industries	31049.6	694.8	38.2	12.7	2.9	-0.2	96.1
Automotive	460.6	662.3	108.2	3.2	-13.1	-21.6	216.9
Rubber & Plastic	138.5	711.2	92.3	62.8	50.8	45.9	89.5
Industrial Machinery	87.6	1077.1	1485.7	72.9	82.8	79.2	474.2
Electronics	248.1	4714.3	119.3	-6.3	-23.8	-27.4	70.7
Other Manufacturing	739.8	303.9	877.8	-23.2	15.9	45.8	393.3
Food & Beverages	1153.3	2300.0	97.9	49.3	69.3	76.8	344.7
Basic Metals	126.1	1455.6	10162.7	28.8	70.6	65.1	315.0
Metal Products	415.3	322.6	31.9	23.1	8.8	2.8	114.0
R+D	2594.1	489.5	369.9	2.1	22.4	43.6	418.5
Pharmaceuticals	284.3	4850.0	104.6	116.3	130.5	131.9	121.9
Electrical Machinery	204.4	136.4	11554.5	107.5	83.4	81.5	200.0
Other Services	16312.5		279.5	313.1	153.2	590.9	
Minerals	457.1		401.2	-3.4	1.2	2.5	475.0
Paper	236.4		2046.2	12.1	40.9	45.8	300.0
Utilities	159.3			55.6	30.2	152.0	
Wood Products	571.0		1136.4	-25.4	5.5	13.0	150.0
Construction	2767.1		173.9	36.0	32.1	108.7	1200.0
Mining	197.7		1750.0	-29.8	-33.3	-22.7	440.0
Textiles	1474.9		900.0	0.0	-12.1	-3.6	
Agriculture	5678.9		666.7	32.3	15.5	16.5	1600.0
Stock of robots							
in year 2000		963	87999	187480	95756	70249	1300
in year 2010		1855	169550	306760	157492	115105	6514
in year 2015		14743	234245	345770	162070	114883	12775

Notes: Table A.2 shows the number of workers in Mexico 2000 and the change in the stock of robots between 2010 and 2015, calculated as the % change relative to the base year. EU9 includes Germany, Denmark, Spain, Finland, France, UK, Italy, Norway, and Sweden; EU7 excludes Germany and Norway; EU5 excludes the UK and Spain; LA includes Argentina, Brazil, Chile, Colombia, Peru and Venezuela. *Data sources*: IFR data. Mexican population census 2000.

Table A.3: Labor market variables do not predict event timing after including fixed effects

	All industries				Manufacturing industries			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ln Number of formal workers	0.027***	0.0904***	0.005	0.005	0.019***	0.059***	0.002	0.001
	(0.007)	(0.009)	(0.004)	(0.004)	(0.005)	(0.009)	(0.005)	(0.005)
Ln Avg. monthly formal wage	0.046***	0.098***	0.008	0.004	0.090***	0.080***	0.005	0.006
	(0.013)	(0.021)	(0.012)	(0.013)	(0.020)	(0.027)	(0.014)	(0.017)
CZ × 4-dig. sector FE		x	x	x		x	x	x
Year FE			x				x	
Year × 4-dig. sector × Region FE				x				x
Observations	100,619	100,423	100,423	98,988	23,935	23,873	23,873	23,241
Adj. R-squared	0.022	0.334	0.660	0.655	0.025	0.336	0.666	0.658
P-val F-stat	0.00	0.00	0.25	0.40	0.00	0.00	0.77	0.90

Notes: Table A.3 shows the estimates of a linear probability model where the outcome is a binary variable indicating whether the CZ started receiving industrial robot imports. The model includes two explanatory variables at the CZ–4-dig. sector–year level: log number of formal workers and log average real monthly wage for formal workers, deflated to constant 2010 pesos using Mexican CPI. The last row displays the p-value of a joint significance test for the two included covariates. Clustering of standard errors is at the CZ level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. Refer to notes in Table 1.4 for the definitions of the different samples. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Table A.4: Summary of estimates: Annual averages

	Pre-event	Postevent short run	Postevent long run
	(1)	(2)	(3)
A. Log Number formal Workers			
Baseline	0.003	0.008	0.055**
	(0.004)	(0.011)	(0.028)
Balanced sample	0.002	0.001	0.049
	(0.004)	(0.012)	(0.030)
Manufacturing	0.005	0.003	0.073
	(0.007)	(0.019)	(0.048)
Manufacturing robot adopters	0.006	0.000	0.085
	(0.008)	0.023	(0.058)
Never treated	0.011***	0.036***	0.105***
	(0.003)	(0.009)	(0.019)
B. Log Average Wage – formal workers			
Baseline	0.002	0.006	0.028**
	(0.002)	(0.004)	(0.013)
Balanced sample	0.002	0.006	0.026**
	(0.002)	(0.005)	(0.013)
Manufacturing	0.004	0.001	0.024
	(0.002)	(0.007)	(0.019)
Manufacturing robot adopters	0.002	0.004	0.036
	(0.003)	(0.009)	(0.023)
Never treated	0.002	0.006*	0.025***
	(0.001)	(0.004)	(0.008)

Notes: Table A.4 shows the same results as in Table 1.4 but using as outcomes the average of the last month of each quarter for each year instead of the value for December. The table shows the simple average of estimated $\theta(e)$ event-study coefficients using equation 1.2 for different periods. Pre-event: years -5 to -1; postevent short run: years 0 to 4; long-run: years 5 to 8. The outcome variable is, in turn, log number of formal workers (Panel (a)) and log average real monthly wage for formal workers, deflated to constant 2010 pesos using Mexican CPI (Panel (b)). The event is defined as a first-time industrial robot importation. Clustering of standard errors is at the CZ level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. Refer to notes in Table 1.4 for the definition of the different samples. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Table A.5: Excluding 2009 and 2010 cohorts

	Pre-event	Postevent short run	Postevent long run
	(1)	(2)	(3)
A. Log Number formal workers			
Baseline	0.004	0.011	0.051*
	(0.004)	(0.011)	(0.028)
Balanced sample	0.003	0.000	0.045*
	(0.004)	(0.010)	(0.026)
Manufacturing	0.007	0.007	0.078
	(0.009)	(0.018)	(0.049)
Manufacturing robot adopters	0.009	0.010	0.109*
	(0.010)	(0.023)	(0.059)
Never treated	0.012***	0.037***	0.091***
	(0.004)	(0.009)	(0.020)
B. Log Average Wage – formal workers			
Baseline	0.001	0.007*	0.032***
	(0.002)	(0.004)	(0.013)
Balanced sample	0.001	0.007	0.031***
	(0.002)	(0.005)	(0.013)
Manufacturing	0.002	0.000	0.028
	(0.003)	0.008	(0.019)
Manufacturing robot adopters	0.000	0.007	0.047**
	0.003	(0.010)	(0.022)
Never treated	0.001	0.006	0.027***
	(0.001)	(0.004)	(0.009)

Notes: Table A.5 shows the same results as in Table 1.4 but excluding cohorts 2009 and 2010 from the sample. The table shows the simple average of estimated $\theta(e)$ event-study coefficients using equation 1.2 for different periods. Pre-event: years -5 to -1; postevent short run: years 0 to 4; long run: years 5 to 8. The outcome variable is, in turn, log number of formal workers (Panel (a)) and log average real monthly wage for formal workers, deflated to constant 2010 pesos using Mexican CPI (Panel (b)). The event is defined as a first-time industrial robot importation. Clustering of standard errors is at the CZ level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. Refer to the notes of Table 1.4 for the definitions of the different samples. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Table A.6: Effects of robots on changes in labor market outcomes: 2SLS estimates using the change in robot exposure between 2010 and 2015

A: Population Censuses outcomes. Long differences 2010-2015

	All industries											
	Employment-to-pop		Wage-empl-to-pop		Self-empl-to-pop		Labor Force Particip. rate		Unemployment rate		Share Informal	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
MX exposure to robots	-0.2	-0.4	-0.3	-0.5	0.1	0.1	-0.2	-0.3	-0.3	-0.7	0.3	0.4
2011-2015	(0.5)	(2.1)	(0.5)	(1.6)	(0.4)	(1.6)	(0.5)	(1.9)	(0.4)	(1.2)	(0.6)	(1.6)
Observations	1,804	19	1,804	19	1,804	19	1,804	19	1,798	19	1,797	19
Adj. R-squared	0.154		0.05		0.180		0.137		0.026		0.155	
Mean dep. 2005	45.47		23.50		21.97		50.39		4.406		49.13	
Lower limit CI	-1.3	-4.5	-1.2	-3.6	-0.6	-2.1	-1.2	-4.1	-0.9	-3.1	-0.8	-2.8
Upper limit CI	0.9	3.7	0.6	2.6	0.9	2.4	0.7	3.4	0.4	1.7	1.5	3.6
First-stage coefficient	0.721		0.721		0.721		0.721		0.716		0.716	
First-stage F-statistic	39.55		39.55		39.55		39.55		38.75		38.75	

B: IMSS data - Insured workers. Long differences 2010-2015

	All industries						Manufacturing industry					
	Ln Number Workers		Ln Payroll		Ln Mean Wage		Ln Number Workers		Ln Payroll		Ln Mean Wage	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
MX exposure to robots	4.4	5.1	4.7	4.2	0.2	-0.9	-3.4	-25.7	-1.6	-19.2	1.8	6.4
2011-2015	(3.6)	(9.7)	(4.2)	(8.5)	(1.5)	(2.7)	(6.7)	(15.7)	(7.4)	(15.3)	(2.7)	(6.9)
Observations	1,211	19	1,211	19	1,211	19	777	19	777	19	777	19
Adj. R-squared	0.010		0.014		0.007		0.014		0.012		0.011	
Mean dep. 2005	12050		3.004e+06		158.1		3485		954484		134.2	
Lower limit CI	-2.7	-13.9	-3.5	-12.5	-2.7	-6.1	-16.5	-56.6	-16.0	-49.3	-3.5	-7.0
Upper limit CI	11.5	24.2	12.8	20.9	3.1	4.3	9.7	5.2	12.9	10.9	7.1	19.9
First-stage coefficient	0.74		0.74		0.74		0.76		0.76		0.76	
First-stage F-statistic	38.14		38.14		38.14		36.12		36.12		36.12	

Notes: This table presents 2SLS estimates of the effects of exposure to robots on changes in labor market outcomes between 2010 and 2015. Panel (a) presents outcomes based on population censuses and Panel (b) those based on the IMSS dataset. In all models, I instrument Mexico's exposure to robots using exposure to robots from European countries. All IV estimates are from regressions weighted by the CZ share in national working-age population in 1990. The covariates included in each model are region dummies, exposure to US robots (2011–2015), exposure to *maquiladoras* in 1990, demographic characteristics of commuting zones in 1990 (log population; share urban; share women; share population over 65; shares population with high school complete and population with college complete; shares employment in manufacturing, services and agriculture sectors). For each outcome, the first column presents standard error estimates that are robust to heteroskedasticity and correlation within CZ, and in the second column, robust standard errors computed following Borusyak et al. (2019) are given in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. *Data sources*: IMMS public data. Mexican 1990 and 2000 censuses. IFR data. Faber (2020).

Table A.7: Summary of estimates: Different methods

	Formal Workers			Log Average Wage		
Model	Pre-event	Postevent short run	Postevent long run	Pre-event	Postevent short run	Postevent long run
	(1)	(2)	(3)	(4)	(5)	(6)
CA Only not-yet-treated	0.005	0.016	0.069***	0.001	0.007	0.030**
	(0.004)	(0.012)	(0.029)	(0.002)	(0.004)	(0.013)
SA Last-treated cohort	-0.003	0.038***	0.107***	-0.006	0.014*	0.041*
	(0.013)	(0.013)	(0.036)	(0.007)	(0.007)	(0.023)
CA Never-treated	0.012***	0.041***	0.108***	0.001	0.006	0.026***
	(0.004)	(0.010)	(0.021)	(0.001)	(0.004)	(0.008)
SA Never-treated	-0.027***	0.050***	0.141***	-0.000	0.004	0.022***
	(0.008)	(0.008)	(0.019)	(0.003)	(0.003)	(0.008)
DD Only not-yet-treated	-0.005	0.033*	0.100**	-0.011	0.017	0.049
	(0.015)	(0.019)	(0.046)	(0.011)	(0.013)	(0.032)
DDD Only not-yet-treated	-0.018	0.072*	0.180*	-0.013	0.007	0.020
	(0.038)	(0.039)	(0.096)	(0.014)	(0.017)	(0.041)
DD Never-treated	-0.024***	0.046***	0.134***	-0.003	0.005	0.020***
	(0.010)	(0.010)	(0.021)	(0.004)	(0.004)	(0.008)
DDD Never-treated	0.019	0.018	0.041	-0.006	0.006	0.012
	(0.023)	(0.023)	(0.047)	(0.008)	(0.008)	(0.016)
Synthetic DD		0.071			0.068	
		(0.306)			(0.302)	
Bartik IV		0.051			-0.008	
		(0.097)			(0.026)	

Notes: Table A.7 presents a summary of results from different models and methods. Sample including cohorts that started importing robots during 2007–2022. Pre-event: years -5 to -1; postevent short run: years 0 to 4; long run: years 5 to 8. The outcome variables are the log number of formal workers and the log average real monthly wage for formal workers, deflated to constant 2010 pesos using Mexican CPI. The event is defined as a first-time industrial robot importation. Clustering of standard errors is at the CZ level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. *Data sources*: Mexican customs records from Secretaría de Economía. IMMS public data.

Table A.8: Mechanisms: Balanced sample

	Number of firms	Labor share	Gross value added (log)	GVA per worker (log)	Total fixed assets (log)	Capital per worker	Input cost (log)	Total payroll (log)	Number workers (log)	Workers prod. and sales (log)	Workers admin (log)	Hours worked (log)	Hours prod. and sales (log)	Hours admin (log)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Panel (a). All industries														
D_{mt}^{post}	0.027*	0.019	0.042	0.018	0.001	-0.042	0.052*	0.021	0.024	0.022	0.012	0.019	0.013	-0.001
	(0.015)	(0.047)	(0.028)	(0.023)	(0.041)	(0.029)	(0.026)	(0.029)	(0.018)	(0.027)	(0.034)	(0.018)	(0.027)	(0.034)
Observations	30,397	30,397	30,397	30,397	30,397	30,397	30,397	30,397	30,397	30,185	20,039	30,308	30,099	19,963
Adj. R-squared	0.968	0.154	0.843	0.620	0.803	-0.0909	0.874	0.854	0.905	0.860	0.763	0.901	0.851	0.751
Mean	62	0.476	15.92	0.0671	20.49	0.0886	15.65	4.843	207.2	87.95	21.15	489.4	212.5	50.72
Panel (b). Manufacturing sector: Metal, machinery and equipment, power generation, transportation, plastics and rubber														
D_{mt}^{post}	0.056	-0.014	0.286*	0.119	0.324**	0.015	0.242*	0.283*	0.167*	0.225*	0.056	0.178*	0.241*	0.093
	(0.042)	(0.040)	(0.150)	(0.092)	(0.135)	(0.013)	(0.137)	(0.143)	(0.095)	(0.127)	(0.130)	(0.102)	(0.132)	(0.135)
Observations	687	687	687	687	687	687	687	687	687	687	469	687	687	469
Adj. R-squared	0.972	0.0930	0.776	0.574	0.744	0.527	0.778	0.786	0.826	0.789	0.704	0.810	0.770	0.695
Mean	62.95	0.480	18.10	0.0561	12.97	0.0795	23.99	8.328	198.9	107.6	22.65	404.1	226	47.61

Notes: Table A.8 shows the results of estimating the TWFE model, equation 1.3, where the event is defined as a first-time industrial robot importation. Panel (a) presents results for all industries and Panel (b) for the following manufacturing industries: basic metals, metal products, machinery, equipment, electronics, appliances, power generation, transportation, and other manufacturing not included in Table A.9. The first dependent variable is the log number of firms at the CZ–five-digit sector level (column (1)). The first dependent variable is the log number of firms at the CZ–five-digit sector level (column (1)). The following seven dependent variables are balance-sheet measures capturing firm productivity and size, aggregated to the CZ–five-digit sector level: labor share, calculated as the total wage bill divided by gross value added (column (2)); capital per worker, calculated as the total fixed assets divided by the number of workers (column (3)); log total gross value added (column (4)) and GVA per worker (column (5)); log input costs (column (6)); log total wage bill (column (7)); and log fixed assets (as a proxy for capital, column (8)). Column (9) presents results for the log number of workers, column (10) for the log number of workers in production and sales, and column (11) for the log number of workers in administrative jobs (including managers). Columns (12)–(14) show similar results for the log number of hours worked in a year. The sample includes not-yet-but-eventually-treated CZs for 2005–2019, including cohorts that started importing robots during 2007–2022. Means are calculated for the period before the event, in levels, and are reported in millions of Mexican pesos (CPI-adjusted to 2005), except for the number of workers, number of hours worked in a year (in thousands), and shares (ratios). Standard errors are clustered at the CZ level and are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. *Data sources*: Mexican economic censuses 2003, 2008, 2013, and 2018.

Table A.9: Mechanisms: Other industries, balanced sample

	Number of firms	Labor share	Gross Value Added (log)	GVA per worker (log)	Total fixed assets (log)	Capital per worker	Input cost (log)	Total payroll (log)	Number workers (log)	Workers prod. and sales (log)	Workers admin (log)	Hours worked (log)	Hours prod. and sales (log)	Hours admin (log)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Panel (a) Manufacturing: Wood and Paper, Chemical Industry														
D_{mt}^{post}	0.103**	-0.006	-0.022	-0.018	-0.019	0.032*	-0.074	-0.033	-0.004	-0.074	-0.020	-0.015	-0.094	-0.019
	(0.040)	(0.027)	(0.112)	(0.061)	(0.117)	(0.018)	(0.109)	(0.114)	(0.074)	(0.102)	(0.142)	(0.075)	(0.104)	(0.146)
Observations	989	989	989	989	989	989	989	989	989	987	677	989	987	677
R-squared	0.974	0.228	0.784	0.611	0.783	0.604	0.808	0.758	0.846	0.771	0.648	0.844	0.769	0.627
Mean	94.91	0.390	34.43	0.0972	34.37	0.0965	78.99	6.876	299.9	130.6	23.67	726.3	341.2	60.29
Panel (b) Manufacturing: Textiles, Apparel, Leather Goods, Food, and Beverage Industry														
D_{mt}^{post}	0.027	-0.019	0.117	0.056	0.087	-0.006	0.100	0.146	0.061	0.130	0.096	0.074	0.116	0.086
	(0.043)	(0.044)	(0.101)	(0.061)	(0.118)	(0.007)	(0.084)	(0.114)	(0.070)	(0.103)	(0.101)	(0.072)	(0.107)	(0.109)
Observations	1,436	1,436	1,436	1,436	1,436	1,436	1,436	1,436	1,436	1,420	736	1,436	1,420	736
R-squared	0.942	0.0748	0.733	0.489	0.737	0.617	0.772	0.740	0.808	0.754	0.751	0.796	0.738	0.736
Mean	43.76	0.446	17.93	0.0456	17.59	0.0649	32.34	7.873	250.3	166.7	29.38	525.6	360.9	64.90
Panel (c) Nonmanufacturing														
D_{mt}^{post}	0.024	0.022	0.035	0.015	-0.009	-0.047	0.050*	0.012	0.020	0.015	0.009	0.014	0.007	-0.006
	(0.014)	(0.052)	(0.028)	(0.023)	(0.042)	(0.032)	(0.026)	(0.027)	(0.018)	(0.025)	(0.035)	(0.018)	(0.026)	(0.035)
Observations	27,285	27,285	27,285	27,285	27,285	27,285	27,285	27,285	27,285	27,091	18,157	27,196	27,005	18,081
R-squared	0.968	0.162	0.850	0.624	0.806	-0.0817	0.881	0.864	0.913	0.869	0.768	0.909	0.860	0.756
Mean	61.72	0.480	15.05	0.0673	20.29	0.0899	12.14	4.521	201.5	81.56	20.66	480.4	199.1	49.80

Notes: Table A.9 shows the results of estimating the TWFE model, equation 1.3, where the event is defined as a first-time industrial robot importation, and balancing the sample in calendar time. Panel (a) presents results for the following manufacturing industries: food, drinks, tobacco, energy, fossil fuels, chemicals, plastics, rubber. Panel (b) for results for the following manufacturing industries: textile, clothing, leather, wood, paper, printing, and nonmetallic mineral products. Panel (c) includes all nonmanufacturing industries. The first dependent variable is the log number of firms at the CZ–five-digit sector level (column (1)). The first dependent variable is the log number of firms at the CZ–five-digit sector level (column (1)). The following seven dependent variables are balance-sheet measures capturing firm productivity and size, aggregated to the CZ–five-digit sector level: labor share, calculated as the total wage bill divided by gross value added (column (2)); capital per worker, calculated as the total fixed assets divided by the number of workers (column (3)); log total gross value added (column (4)) and GVA per worker (column (5)); log input costs (column (6)); log total wage bill (column (7)); and log fixed assets (as a proxy for capital, column (8)). Column (9) presents results for the log number of workers, column (10) for the log number of workers in production and sales, and column (11) for the log number of workers in administrative jobs (including managers). Columns (12)–(14) show similar results for the log number of hours worked in a year. The sample includes not-yet-but-eventually-treated CZs for 2005–2019, including cohorts that started importing robots during 2007–2022. Means are calculated for the period before the event, in levels, and are reported in millions of Mexican pesos (CPI-adjusted to 2005), except for the number of workers, the number of hours worked in a year (in thousands), and shares (ratios). Standard errors are clustered at the CZ level and are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. *Data sources*: Mexican Economic Censuses (2003, 2008, 2013, 2018).

2. Appendix to "Reconciling the Evidence on False Links in Historical Data"

Table A.10: A Comparison of MLP and CLP Links, by MLP Link Score. 1910-1940, 1920-1940

	1910-1940 Census Links				1920-1940 Census Links			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Links in MLP & CLP	Total links	Total linked to different record	Average link probability (for MLP-CLP agreements)	Difference in MLP-CLP average link probability (for links in col. 2)	Total links	Total linked to different record	Average link probability (for MLP-CLP agreements)	Difference in MLP-CLP average link probability (for links in col. 2)
Exact Standard	5,683,857	182,334	0.761	0.472	8,479,165	376,819	0.756	0.603
NYSIIS Standard	6,081,590	279,346	0.745	0.495	8,801,121	512,409	0.745	0.613
Race NYSIIS Standard	6,189,919	274,167	0.743	0.489	8,967,019	502,602	0.744	0.609
Exact Conservative	4,677,663	43,930	0.785	0.474	6,895,279	96,978	0.764	0.581
NYSIIS Conservative	4,531,845	46,303	0.759	0.443	6,423,764	84,035	0.751	0.570
Race NYSIIS Conservative	4,661,299	47,015	0.757	0.435	6,604,826	82,803	0.749	0.566
B. Links made by MLP or CLP (but not both)	Total links	Share of all links	MLP link probability	Share of links with higher MLP link probability	Total links	Share of all links	MLP link probability	Share of links with higher MLP link probability
Links in MLP, not linked by any CLP variant	3,357,926	0.322	0.553	--	5,460,943	0.345	0.659	--
Links in CLP, not linked by MLP								
Exact Standard	2,569,321	0.311	0.303	0.271	3,438,744	0.289	0.141	0.287
NYSIIS Standard	3,306,459	0.352	0.239	0.278	4,612,048	0.344	0.111	0.279
Race NYSIIS Standard	3,322,881	0.349	0.247	0.258	4,635,439	0.341	0.114	0.249
Exact Conservative	1,142,381	0.196	0.339	0.187	1,725,902	0.200	0.143	0.177
NYSIIS Conservative	1,378,971	0.233	0.187	0.179	2,237,383	0.258	0.094	0.155
Race NYSIIS Conservative	1,414,055	0.233	0.195	0.163	2,298,477	0.258	0.096	0.135

Notes: This table presents data for the census-linked period 1910-1940 and 1920-1940, restricted to men. Link probability scores range from 0 to 1 and are calculated based on the similarity of the record's name, age, race, birth state, household co-residents, and geography. Panel A, columns (3) and (7) present the average link probability score among the links made by both the Multigenerational Longitudinal Panel (MLP) and the Census Linking Project (CLP) for the same person in 1940 (i.e., there is agreement). Panel A, columns (4) and (8) present the difference in percentage points between the average link probability of the links made by MLP and CLP among the set of disagreements. Panel B, columns (2) and (6) present the fraction of links made only by one of the two methods, using as the denominator the total number of MLP links or CLP links by method, not the union of all methods. Source: Multigenerational Longitudinal Panel (MLP) (Helgertz et al., 2022) and Census Linking Project (CLP) (Abramitzky et al., 2022).

Table A.11: A Comparison of MLP and CLP Links, by Race and Nativity Subgroup. 1910-1940.

	1910-1940 Census Links									
	Black men					Foreign born men				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
A. Total links between census years										
MLP links	473,183					823,641				
CLP links										
Exact Standard	482,157					564,929				
NYSIIS Standard	579,254					779,495				
Race NYSIIS Standard	607,378					777,060				
Exact Conservative	253,167					361,756				
NYSIIS Conservative	290,971					422,618				
Race NYSIIS Conservative	352,921					421,707				
B. Links in MLP & CLP	Total links	Total linked to different person	Share (2)/(1)	Average link probability (for MLP-CLP agreements)	Difference in MLP-CLP average link probability (for links in col. 2)	Total links	Total linked to different record	Share (2)/(1)	Average link probability (for MLP-CLP agreements)	Difference in MLP-CLP average link probability (for links in col. 2)
Exact Standard	206,800	15,790	0.076	0.643	0.431	330,927	19,130	0.058	0.780	0.485
NYSIIS Standard	226,752	15,790	0.070	0.633	0.445	390,121	34,481	0.088	0.754	0.472
Race NYSIIS Standard	268,917	15,613	0.058	0.617	0.390	390,126	34,496	0.088	0.754	0.472
Exact Conservative	153,729	3,354	0.022	0.675	0.422	262,634	5,150	0.020	0.809	0.499
NYSIIS Conservative	157,589	2,856	0.018	0.657	0.391	270,329	6,431	0.024	0.773	0.428
Race NYSIIS Conservative	207,664	2,175	0.010	0.639	0.390	270,439	6,443	0.024	0.774	0.428
C, Links made by MLP or CLP (but not both)	Total links	Share of all links	Average link probability	Share of links with higher MLP link probability		Total links	Share of all links	Average link probability	Share of links with higher MLP link probability	
Links in MLP, not linked by any CLP variant	178,738	0.378	0.452	-		367,297	0.446	0.536	-	
Links in CLP, not linked by MLP										
Exact Standard	275,357	0.571	0.111	0.442		234,002	0.414	0.302	0.280	
NYSIIS Standard	352,502	0.609	0.095	0.402		389,374	0.500	0.210	0.307	
Race NYSIIS Standard	338,461	0.557	0.149	0.197		386,934	0.498	0.211	0.307	
Exact Conservative	99,438	0.393	0.104	0.336		99,122	0.274	0.348	0.189	
NYSIIS Conservative	133,382	0.458	0.070	0.267		152,289	0.360	0.151	0.199	
Race NYSIIS Conservative	145,257	0.412	0.140	0.197		151,268	0.359	0.152	0.200	

Notes: This table presents data for the census-linked period 1910-1940, restricted to Black men (columns 1 to 5) and foreign-born men (columns 6 to 10). Link probability scores range from 0 to 1 and are calculated based on the similarity of the record's name, age, race, birth state, household co-residents, and geography. Panel B, columns (4) and (9) present the average link probability score among the links made by both the Multigenerational Longitudinal Panel (MLP) and the Census Linking Project (CLP) for the same person in 1940 (i.e., there is agreement). Panel B, columns (5) and (10) present the difference in percentage points between the average link probability of the links made by MLP and CLP among the set of disagreements. Source: Multigenerational Longitudinal Panel (MLP) (Helgertz et al., 2022) and Census Linking Project (CLP) (Abramitzky et al., 2022).

Table A.12: A Comparison of MLP and CLP Links, by Race and Nativity Subgroup. 1920-1940.

	1920-1940 Census Links									
	Black men					Foreign born men				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
A. Total links between census years										
MLP links	615,189					1,343,439				
CLP links										
Exact Standard	668,173					911,488				
NYSIIS Standard	788,428					1,216,549				
Race NYSIIS Standard	855,619					1,213,076				
Exact Conservative	363,942					629,304				
NYSIIS Conservative	410,829					715,394				
Race NYSIIS Conservative	517,419					713,875				
B. Links in MLP & CLP	Total links	Total linked to different record	Share (2)/(1)	Average link probability (for MLP-CLP agreements)	Difference in MLP-CLP average link probability (for links in col. 2)	Total links	Total linked to different record	Share (2)/(1)	Average link probability (for MLP-CLP agreements)	Difference in MLP-CLP average link probability (for links in col. 2)
Exact Standard	265,780	24,967	0.094	0.647	0.537	579,618	39,987	0.069	0.756	0.622
NYSIIS Standard	284,104	27,936	0.098	0.644	0.541	644,620	61,084	0.095	0.744	0.624
Race NYSIIS Standard	341,233	23,754	0.070	0.647	0.518	644,358	61,021	0.095	0.744	0.624
Exact Conservative	192,300	5,415	0.028	0.659	0.538	464,430	11,773	0.025	0.763	0.600
NYSIIS Conservative	193,539	4,005	0.021	0.653	0.521	449,936	11,324	0.025	0.747	0.585
Race NYSIIS Conservative	259,701	3,659	0.014	0.647	0.518	449,933	11,317	0.025	0.747	0.585
C, Links made by MLP or CLP (but not both)	Total links	Share of all links	Average link probability	Share of links with higher MLP link probability		Total links	Share of all links	Average link probability	Share of links with higher MLP link probability	
Links in MLP, not linked by any CLP variant	240,550	0.391	0.551	-		582,994	0.434	0.653	-	
Links in CLP, not linked by MLP										
Exact Standard	402,393	0.602	0.039	0.404		331,870	0.364	0.096	0.288	
NYSIIS Standard	504,324	0.640	0.037	0.361		571,929	0.470	0.068	0.302	
Race NYSIIS Standard	514,386	0.601	0.051	0.207		568,718	0.469	0.068	0.302	
Exact Conservative	171,642	0.472	0.048	0.247		164,874	0.262	0.106	0.169	
NYSIIS Conservative	217,290	0.529	0.043	0.195		265,458	0.371	0.062	0.169	
Race NYSIIS Conservative	257,718	0.498	0.051	0.207		263,942	0.370	0.063	0.168	

Notes: This table presents data for the census-linked period 1920-1940, restricted to Black men (columns 1 to 5) and foreign-born men (columns 6 to 10). Link probability scores range from 0 to 1 and are calculated based on the similarity of the record's name, age, race, birth state, household co-residents, and geography. Panel B, columns (4) and (9) present the average link probability score among the links made by both the Multigenerational Longitudinal Panel (MLP) and the Census Linking Project (CLP) for the same person in 1940 (i.e., there is agreement). Panel B, columns (5) and (10) present the difference in percentage points between the average link probability of the links made by MLP and CLP among the set of disagreements. Source: Multigenerational Longitudinal Panel (MLP) (Helgertz et al., 2022) and Census Linking Project (CLP) (Abramitzky et al., 2022).

Bibliography

Abowd, J. M. (2017). Large-scale data linkage from multiple sources: Methodology and research challenges. NBER Summer Institute Methods Lecture.

Abramitzky, R., Boustan, L., and Eriksson, K. (2012). Web appendix: Europe's tired, poor, huddled masses: Self-selection and economic outcomes in the age of mass migration.

Abramitzky, R., Boustan, L., and Eriksson, K. (2013). Have the poor always been less likely to migrate? evidence from inheritance practices during the age of mass migration. *Journal of Development Economics*, 102:2–14.

Abramitzky, R., Boustan, L., and Eriksson, K. (2014). A nation of immigrants: Assimilation and economic outcomes in the age of mass migration. *Journal of Political Economy*, 122(3):467–506.

Abramitzky, R., Boustan, L., Eriksson, K., Feigenbaum, J., and Pérez, S. (2021). Automated linking of historical data. *Journal of Economic Literature*, 59(3):865–918.

Abramitzky, R., Boustan, L., Eriksson, K., Rashid, M., and Pérez, S. (2022). Census linking project: 1910-1940, 1920-1940, and 1930-1940 crosswalks.

Abramitzky, R., Mill, R., and Pérez, S. (2020). Linking individuals across historical sources: A fully automated approach. *Historical Methods: A Journal of Quantitative and Interdisciplinary History*, 53(2):94–111.

Acemoglu, D., Anderson, G., Beede, D., Buffington, C., Childress, E., Dinlersoz, E., Foster, L., Goldschlag, N., Haltiwanger, J. C., Kroff, Z., Restrepo, P., and Zolas, N. J. (2022). Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey. *SSRN Electronic Journal*.

Acemoglu, D., Koster, H., and Ozgen, C. (2023). Robots and Workers: Evidence from the Netherlands. *SSRN Electronic Journal*.

Acemoglu, D., Lelarge, C., and Restrepo, P. (2021). Competing with Robots: Firm-Level Evidence from France. *SSRN Electronic Journal*.

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542.

Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30.

Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, 128(6):2188–2244.

Acemoglu, D. and Restrepo, P. (2022a). Tasks, Automation, and the Rise in U.S. Wage Inequality. *Econometrica*, 90(5):1973–2016.

Acemoglu, D. and Restrepo, P. (2022b). Tasks, Automation, and the Rise in U.S. Wage Inequality. *Econometrica*, 90(5):1973–2016.

Aghion, P., Akcigit, U., Bergeaud, A., Blundell, R., and Hémous, D. (2019). Innovation and top income inequality. *The Review of Economic Studies*, 86:1–45.

Aghion, P., Antonin, C., Bunel, S., and Jaravel, X. (2022). Modern manufacturing capital, labor demand, and product market dynamics: Evidence from France Working Paper.

Aiyar, S. and Ebeke, C. (2020). Inequality of opportunity, inequality of income and economic growth. *World Development*, 136:105115.

Aizer, A., Eli, S., Ferrie, J., and Lleras-Muney, A. (2016). The long term impact of cash transfers to poor families. *American Economic Review*, 106(4):935–971.

Alesina, A., Hohmann, S., Michalopoulos, S., and Papaioannou, E. (2021). Intergenerational mobility in africa. *Econometrica*, 89:1–35.

Almond, D. (2006). Is the 1918 influenza pandemic over? long-term effects of in utero influenza exposure in the post-1940 us population. *Journal of political Economy*, 114:672–712.

Arellano, M. and Bover, O. (1995). Another look at the instrumental variable estimation of error-components models. *Journal of econometrics*, 68:29–51.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2021). Synthetic Difference-in-Differences. *American Economic Review*, 111(12):4088–4118.

Arntz, M., Gregory, T., and Zierahn, U. (2019). Digitalization and the future of work: Macroeconomic consequences.

Artuc, E., Bastos, P., Copestake, A., and Rijkers, B. (2022). *Robots and Trade: Implications for Developing Countries*. Number 2018.

Artuc, E., Christiaensen, L., and Winkler, H. (2019). Does Automation in Rich Countries Hurt Developing Ones? Evidence from the U.S. and Mexico. *Does Automation in Rich Countries Hurt Developing Ones? Evidence from the U.S. and Mexico*, (February).

Atkin, D. (2016). Endogenous skill acquisition and export manufacturing in Mexico. *American Economic Review*, 106(8):2046–2085.

Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US Labor Market. *American Economic Review*, 103(5):1553–1597.

Aydemir, A. B. and Yazici, H. (2019). Intergenerational education mobility and the level of development. *European Economic Review*, 116:160–185.

Bailey, M., Lin, P. Z., Mohammed, A. R. S., Mohnen, P., Murray, J., Zhang, M., and Prettyman, A. (2023). The creation of life-m: The longitudinal, intergenerational family electronic micro-database project. *Historical Methods: A Journal of Quantitative and Interdisciplinary History*, 56(3):138–159.

Bailey, M. J., Cole, C., Henderson, M., and Massey, C. G. (2020). How well do automated linking methods perform? evidence from us historical data. *Journal of Economic Literature*, 58(4):997–1044.

Bandiera, O., Burgess, R., Das, N., Gulesci, S., Rasul, I., and Sulaiman, M. (2017). Labor markets and poverty in village economies. *The Quarterly Journal of Economics*, 132:811–870.

Banerjee, A. V. and Duflo, E. (2003). Inequality and growth: What can the data say? *Journal of economic growth*, 8:267–299.

Barro, R. J. (2000). Inequality and growth in a panel of countries. *Journal of economic growth*, 5:5–32.

Becker, G. S. and Lewis, H. G. (1973). On the interaction between the quantity and quality of children. *Journal of political Economy*, 81:S279–S288.

Becker, G. S. and Tomes, N. (1979). An equilibrium theory of the distribution of income and intergenerational mobility. *Journal of political Economy*, 87:1153–1189.

Becker, G. S. and Tomes, N. (1986). Human capital and the rise and fall of families. *Journal of labor economics*, 4:S1–S39.

Bell, A., Chetty, R., Jaravel, X., Petkova, N., and Van Reenen, J. (2019). Who becomes an inventor in america? the importance of exposure to innovation. *The Quarterly Journal of Economics*, 134:647–713.

Berg, A., Ostry, J. D., Tsangarides, C. G., and Yakhshilikov, Y. (2018). Redistribution, inequality, and growth: new evidence. *Journal of Economic Growth*, 23:259–305.

Bessen, J., Goos, M., Salomons, A., and van den Berge, W. (2020). Firm-Level Automation: Evidence from the Netherlands. *AEA Papers and Proceedings*, 110:389–393.

Black, S. E., Devereux, P. J., Lundborg, P., and Majlesi, K. (2020). Poor little rich kids? the role of nature versus nurture in wealth and other economic outcomes and behaviours. *The Review of Economic Studies*, 87:1683–1725.

Black, S. E., Devereux, P. J., and Salvanes, K. G. (2011). Older and wiser? birth order and iq of young men. *CESifo Economic Studies*, 57:103–120.

Blanden, J. (2013). Cross-country rankings in intergenerational mobility: a comparison of approaches from economics and sociology. *Journal of Economic Surveys*, 27:38–73.

Bleakley, H. and Ferrie, J. (2016). Shocking behavior: Random wealth in antebellum georgia and human capital across generations. *Quarterly Journal of Economics*, 131(3):1455–1495.

Bleakley, H. and Ferrie, J. (2017). Land opening on the georgia frontier and the coase theorem in the short- and long- run. University of Michigan Working Papers.

Bleakley, H. and Ferrie, J. P. (2013). Up from poverty? the 1832 cherokee land lottery and the long-run distribution of wealth. (19175).

Blundell, R. and Bond, S. (1998). Initial conditions and moment restrictions in dynamic panel data models. *Journal of econometrics*, 87:115–143.

Bolt, J. and van Zanden, J. L. (2020). Maddison style estimates of the evolution of the world economy. a new 2020 update. Maddison project working paper, University of Groningen, Groningen Growth and Development Centre.

Borusyak, K., Jaravel, X., and Spiess, J. (2021). Revisiting Event Study Designs: Robust and Efficient Estimation. (2014):1–85.

Bourguignon, F., Ferreira, F. H., and Walton, M. (2007). Equity, efficiency and inequality traps: a research agenda. *The Journal of Economic Inequality*, 5:235–256.

Boustan, L. P., Kahn, M. E., and Rhode, P. W. (2012). Moving to higher ground: Migration response to natural disasters in the early twentieth century. *American Economic Review: Papers and Proceedings*, 102(3):238–244.

Bowles, S. and Gintis, H. (2002). The inheritance of inequality. *Journal of economic Perspectives*, 16:3–30.

Bradbury, K. and Triest, R. K. (2016). Inequality of opportunity and aggregate economic performance. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 2:178–201.

Brambilla, I., César, A., Falcone, G., and Gasparini, L. (2023). The impact of robots in Latin America: Evidence from local labor markets. *World Development*, 170:106271.

Brueckner, M. and Lederman, D. (2018). Inequality and economic growth: the role of initial income. *Journal of Economic Growth*, 23:341–366.

Brunori, P., Ferreira, F. H., and Peragine, V. (2013). Inequality of opportunity, income inequality and economic mobility: some international comparisons. In *Getting Development Right: Structural Transformation, Inclusion, and Sustainability in the Post-Crisis Era*, page 85. Palgrave Macmillan.

Calì, M. and Presidente, G. (2021a). Automation and manufacturing performance in a developing country. *Oxford Martin School, Oxford University*, pages 1–58.

Calì, M. and Presidente, G. (2021b). Robots For Economic Development. *IDEAS Working Paper Series from RePEc*, pages 1–78.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Caruso, G. and Miller, S. (2015). Long run effects and intergenerational transmission of natural disasters: A case study on the 1970 ancash earthquake. *Journal of development economics*, 117:134–150.

Chaney, T. (2014). The network structure of international trade. *American Economic Review*, 104(11):3600–3634.

Chen, C. and Frey, C. B. (2021). Oxford Martin Working Paper Series on Economic and Technological Change Automation or Globalization? The Impacts of Robots and Chinese Imports on Jobs in the United Kingdom Automation or Globalization? The Impacts of Robots and Chinese Imports on Jobs in the United Kingdom *.

Chetty, R. and Hendren, N. (2018). The impacts of neighborhoods on intergenerational mobility i: Childhood exposure effects. *The Quarterly Journal of Economics*, 133(02):1107–1162.

Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014). Where is the land of opportunity? the geography of intergenerational mobility in the united states. *The Quarterly Journal of Economics*, 129:1553–1623.

Chiacchio, F., Petropoulos, G., and Pichler, D. (2018). The impact of industrial robots on EU employment and wages.

Chiquiar, D. (2019). Working Papers Global Value Chains in Mexico : A Historical Perspective Global Val ue Chai ns i n M ex ico : A His to r ical Per s p ecti v e *.

Christen, P. and Goiser, K. (2007). *Quality and Complexity Measures for Data Linkage and Deduplication*, volume 43.

Collins, W. J. and Wanamaker, M. H. (2014). Selection and economic gains in the great migration of african americans: New evidence from linked census data. *American Economic Journal: Applied Economics*, 6(1):220–252.

Collins, W. J. and Wanamaker, M. H. (2015). The great migration in black and white: New evidence on the selection and sorting of southern migrants. *Journal of Economic History*, 75(4):947–992.

Collins, W. J. and Wanamaker, M. H. (2016). Up from slavery? african american intergenerational economic mobility since 1880. (23395).

Corak, M. (2013). Income inequality, equality of opportunity, and intergenerational mobility. *Journal of Economic Perspectives*, 27:79–102.

Corak, M. (2020). The canadian geography of intergenerational income mobility. *The Economic Journal*, 130:2134–2174.

Costa, D. L., DeSomer, H., Hanss, E., Roudiez, C., Wilson, S. E., and Yetter, N. (2017). Union army veterans, all grown up. *Historical Methods*, 50(2):79–95.

Dasgupta, P. and Ray, D. (1986). Inequality as a determinant of malnutrition and unemployment: Theory. *The Economic Journal*, 96:1011–1034.

Dauth, W., Findeisen, S., Suedekum, J., and Woessner, N. (2021). The Adjustment of Labor Markets to Robots. *Journal of the European Economic Association*, 19(6):3104–3153.

Dauth, W., S Findeisen, Suedekum, J., and Woessner, N. (2018). Adjusting to robots: Worker-level evidence", opportunity and inclusive growth. *Institute Working Paper Federal Reserve Bank of Minneapolis.*, 13.

Dell, M., Jones, B. F., and Olken, B. A. (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics*, 4:66–95.

Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38.

Diamond, R. (2016). The determinants and welfare implications of us workers' diverging location choices by skill: 1980-2000. *American Economic Review*, 106:479–524.

Dixon, J., Hong, B., and Wu, L. (2019). The Employment Consequences of Robots: Firm-Level Evidence. *SSRN Electronic Journal*.

Durlauf, S. N., Johnson, P. A., and Temple, J. R. (2005). Growth econometrics. *Handbook of economic growth*, 1:555–677.

Eli, S., Salisbury, L., and Shertzer, A. (2018). Ideology and migration after the american civil war. *Journal of Economic History*, 78(3).

Emran, M. S., Greene, W., and Shilpi, F. (2018). When measure matters coresidency, truncation bias, and intergenerational mobility in developing countries. *Journal of Human Resources*, 53:589–607.

Eriksson, B. (2016). The missing links: Data quality and bias to estimates of social mobility. Accessed September 15, 2016.

Faber, M. (2020). Robots and reshoring: Evidence from Mexican labor markets. *Journal of International Economics*, 127:103384.

Fan, Y., Yi, J., and Zhang, J. (2015). The great gatsby curve in china: Cross-sectional inequality and intergenerational mobility. Working paper, CUHK.

Feigenbaum, J. J. (2016). Automated census record linking: A machine learning approach.

Fellegi, I. P. and Sunter, A. B. (1969). A theory for record linkage. *Journal of the American Statistical Association*, 64(328):1183–1210.

Ferreira, F. H., Lakner, C., Lugo, M. A., and Özler, B. (2018). Inequality of opportunity and economic growth: how much can cross-country regressions really tell us? *Review of Income and Wealth*, 64:800–827.

Freedman, V. A., Warren, J. R., Pfeffer, F., Helgertz, J., Xu, D., and Kasper, J. D. (2021). National health and aging trends study 1940 census linkage user guide.

Frey, C. B. and Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114:254–280.

Galor, O. and Moav, O. (2004). From physical to human capital accumulation: Inequality and the process of development. *The Review of Economic Studies*, 71(10):1001–1026.

Galor, O. and Tsiddon, D. (1997). The distribution of human capital and economic growth. *Journal of Economic Growth*, 2:93–124.

Galor, O. and Zeira, J. (1993). Income distribution and macroeconomics. *The review of economic studies*, 60:35–52.

Gaubert, C. and Itskhoki, O. (2021). Granular comparative advantage. *Journal of Political Economy*, 129(3):871–939.

Genicot, G. and Ray, D. (2017). Aspirations and inequality. *Econometrica*, 85:489–519.

Ghosh, A., Hwang, S. I. M., and Squires, M. (2024). Links and legibility: Making sense of historical us census automated linking methods. *Journal of Business and Economic Statistics*, 42(2):579–590.

Goeken, R., Lynch, T., Lee, Y. N., Wellington, J., and Magnuson, D. (2017). Evaluating the accuracy of linked u. s. census data: A household approach.

Goldstein, J. R., Alexander, M. B., Casey, F., Miranda-González, A., Menares, F., Osborne, M., and Yildirim, U. (2021). Censoc mortality file: Version 20.

Graetz, G. and Michaels, G. (2018). The Review of Economics and Statistics. *The Review of Economics and Statistics*, 100(5):753–768.

Güell, M., Pellizzari, M., Pica, G., and Rodríguez Mora, J. V. (2018). Correlating social mobility and economic outcomes. *The Economic Journal*, 128:F353–F403.

Hassler, J. and Rodríguez Mora, J. V. (2000). Intelligence, social mobility, and growth. *American Economic Review*, 90:888–908.

Hassler, J., Rodríguez Mora, J. V., and Zeira, J. (2007). Inequality and mobility. *Journal of Economic Growth*, 12:235–259.

Headey, D. D. and Hodge, A. (2009). The effect of population growth on economic growth: A metaregression analysis of the macroeconomic literature. *Population and Development Review*, 35:221–248.

Heckman, J. J. and Mosso, S. (2014). The economics of human development and social mobility. *Annual Review of Economics*, 6:689–733.

Helgertz, J., Price, J., Wellington, J., Thompson, K. J., Ruggles, S., and Fitch, C. A. (2022). A new strategy for linking us historical censuses: A case study for the ipums multigenerational longitudinal panel. *Historical Methods: A Journal of Quantitative and Interdisciplinary History*, 55(1):12–29.

Helgertz, J., Ruggles, S., Warren, J. R., Fitch, C. A., Hacker, J. D., Nelson, M. A., Price, J. P., Roberts, E., and Sobek, M. (2023). Ipums multigenerational longitudinal panel: Version 1.1 [dataset].

Henderson, J. V., Storeygard, A., and Weil, D. N. (2012). Measuring economic growth from outer space. *American economic review*, 102:994–1028.

Hertz, T. (2008). A group-specific measure of intergenerational persistence. *Economics Letters*, 100:415–417.

Hodler, R. and Raschky, P. A. (2014). Regional favoritism. *The Quarterly Journal of Economics*, 129:995–1033.

Hornbeck, R. and Naidu, S. (2014). When the levee breaks: Black migration and economic development in the american south. *American Economic Review*, 104(3):963–990.

Hsieh, C.-T., Hurst, E., Jones, C. I., and Klenow, P. J. (2019). The allocation of talent and us economic growth. *Econometrica*, 87:1439–1474.

Humlum, A. (2020). Robot Adoption and Labor Market Dynamics. *Unpublished Manuscript*, pages 1–78.

Iacovone, L., Muñoz Moreno, R., Olaberria, E., and De La Paz Pereira López, M. (2022). *Productivity Growth in Mexico*. World Bank Group.

IFR (2022). World Robotics Report: “All-Time High” with Half a Million Robots Installed in one Year.

Ing, L. Y. and Zhang, R. (2022). *Automation in Indonesia: Productivity, Quality, and Employment*.

Johnson, P. and Papageorgiou, C. (2020). What remains of cross-country convergence? *Journal of Economic Literature*, 58:129–75.

Jäntti, M. and Jenkins, S. P. (2015). Income mobility. In *Handbook of income distribution*, volume 2, pages 807–935. Elsevier.

Koch, M., Manuylov, I., and Smolka, M. (2021). Robots and Firms. *Economic Journal*, 131(638):2553–2584.

Kugler, A. D., Kugler, M., Ripani, L., and Rodrigo, R. (2021). U.S. Robots and Their Impacts in the Tropics: Evidence from Colombian Labor Markets. *SSRN Electronic Journal*.

La Ferrara, E. (2019). Aspirations, social norms, and development. *Journal of the European Economic Association*.

Levy, S. and Schady, N. (2013). Latin america's social policy challenge: Education, social insurance, redistribution. *Journal of Economic Perspectives*, 27:193–218.

Loury, G. C. (1981). Intergenerational transfers and the distribution of earnings. *Econometrica: Journal of the Econometric Society*, pages 843–867.

López-Calva, L. F. and Lustig, N. C. (2010). *Declining inequality in Latin America: A decade of progress?* Brookings Institution Press.

Maccini, S. and Yang, D. (2009). Under the weather: Health, schooling, and economic consequences of early-life rainfall. *American Economic Review*, 99:1006–1026.

Maoz, Y. D. and Moav, O. (1999). Intergenerational mobility and the process of development. *The Economic Journal*, 109:677–697.

Marrero, G. A. and Rodríguez, J. G. (2013). Inequality of opportunity and growth. *Journal of Development Economics*, 104:107–122.

Massey, C. G. (2017). Playing with matches: An assessment of accuracy in linked historical data. *Historical Methods: A Journal of Quantitative and Interdisciplinary History*, pages 1–15.

Mayer, S. E. and Lopoo, L. M. (2008). Government spending and intergenerational mobility. *Journal of Public Economics*, 92:139–158.

Mill, R. (2013). Record linkage across historical datasets. inequality and discrimination in historical and modern labor markets.

Mill, R. and Stein, L. C. (2016). Race, skin color, and economic outcomes in early twentieth-century america.

Miller, D. L. (2023). An Introductory Guide to Event Study Models. *Journal of Economic Perspectives*, 37(2):203–230.

Mincer, J. (1984). Human capital and economic growth. *Economics of education review*, 3:195–205.

Moretti, E. (2012). *The new geography of jobs*. Houghton Mifflin Harcourt.

Munoz, E. (2021). The geography of intergenerational mobility in latin america and the caribbean. 3807350, SSRN.

Neidhöfer, G. (2019). Intergenerational mobility and the rise and fall of inequality: Lessons from latin america. *The Journal of Economic Inequality*, 17:499–520.

Neidhöfer, G., Serrano, J., and Gasparini, L. (2018). Educational inequality and intergenerational mobility in latin america: A new database. *Journal of Development Economics*, 134:329–349.

Neves, P. C., Afonso, O., and Silva, S. T. (2016). A meta analytic reassessment of the effects of inequality on growth. *World Development*, 78:386–400.

Neves, P. C. and Silva, S. M. T. (2014). Inequality and growth: Uncovering the main conclusions from the empirics. *Journal of Development Studies*, 50:1–21.

Nickell, S. (1981). Biases in dynamic models with fixed effects. *Econometrica: Journal of the econometric society*, pages 1417–1426.

Nix, E. and Qian, N. (2015). The fluidity of race: "passing" in the united states, 1880-1940. (20828).

Nybom, M. and Stuhler, J. (2017). Biases in standard measures of intergenerational income dependence. *Journal of Human Resources*, 52:800–825.

Ó Gráda, C., Anbinder, T., Connor, D., and Wegge, S. A. (2023). The problem of false positives in automated census linking: Nineteenth-century new york's irish immigrants as a case study. *Historical Methods*, 56(4):240–259.

Owen, A. L. and Weil, D. N. (1998). Intergenerational earnings mobility, inequality and growth. *Journal of Monetary Economics*, 41:71–104.

Panizza, U. (2002). Income inequality and economic growth: Evidence from american data. *Journal of Economic Growth*, 7:25–41.

Piovani, J. I. and Salvia, A. (2018). *La Argentina en el siglo XXI: Como somos, vivimos y convivimos en una sociedad desigual: Encuesta Nacional sobre la Estructura Social*. Siglo Veintiuno Argentina.

Porta, R. L. and Shleifer, A. (2008). The unofficial economy and economic development. *Brookings Papers on Economic Activity*, 2008(2):275–352.

Porta, R. L. and Shleifer, A. (2014). Informality and Development. *Journal of Economic Perspectives*, 28(3):109–126.

Price, J., Buckles, K., Riley, I., and Van Leeuwen, J. (2021). Combining family history and machine learning to link historical records. *Explorations in Economic History*, 80:1–28.

Puggioni, D., Calderón, M., Cebreros Zurita, A., Fernández Bujanda, L., Inguanzo González, J. A., and Jaume, D. (2022). Inequality, income dynamics, and worker transitions: The case of Mexico. *Quantitative Economics*, 13(4):1669–1705.

Rambachan, A. and Roth, J. (2023). A More Credible Approach to Parallel Trends. *Review of Economic Studies*, (February):2555–2591.

Rawls, J. (1971). *A theory of justice*. Harvard University Press.

Robles, R. and Foladori, G. (2019). Una revisión histórica de la automatización de la minería en México. *Revista Problemas del Desarrollo*, 50(197):157–180.

Rodrigo, R. (2021). Robot Adoption, Organizational Capital, and the Productivity Paradox. *Working Paper*, (November):1–50.

Rodrik, D. (2021). New Technologies, Global Value Chains, and the Developing Economies. *SSRN Electronic Journal*.

Rodríguez Mora, J. V. (2009). The misallocation of talent. Technical report, CREI, Centre de Recerca en Economia Internacional.

Roemer, J. E. (1998). *Equality of opportunity*. Harvard University Press.

Roodman, D. (2009). How to do xtabond2: An introduction to difference and system gmm in stata. *The stata journal*, 9:86–136.

Ruggles, S. (2006). Linked historical censuses: A new approach. *History and Computing*, 14:213–224.

Ruggles, S., Fitch, C., Goeken, R., Hacker, J., Nelson, M., Roberts, E., Schouweiler, M., and Sobek, M. (2021). Ipums ancestry full count data: Version 3.0 [dataset].

Sacerdote, B. (2011). Nature and nurture effects on children's outcomes: What have we learned from studies of twins and adoptees? In *Handbook of social economics*, volume 1, pages 1–30. Elsevier.

Scheuren, F. and Winkler, W. E. (1993). Regression analysis of data files that are computer matched. *Survey methodology*, 19(1):39–58.

Sen, A. (1980). Equality of what? The Tanner lecture on human values.

Solon, G. (1992). Intergenerational income mobility in the united states. *The American Economic Review*, pages 393–408.

Stemmler, H. (2020). Automated Deindustrialization: How Global Robotization Affects Emerging Economies - Evidence from Brazil. (382):58.

Stuhler, J. (2018). A review of intergenerational mobility and its drivers. Technical report, Publications Office of the European Union, Luxembourg.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

Torche, F. (2014). Intergenerational mobility and inequality: The latin american case. *Annual Review of Sociology*, 40:619–642.

Ulyssea, G. (2018). Firms, informality, and development: Theory and evidence from Brazil. *American Economic Review*, 108(8):2015–2047.

Van der Weide, R. and Milanovic, B. (2018). Inequality is bad for growth of the poor (but not for that of the rich). *The World Bank Economic Review*, 32:507–530.

Voitchovsky, S. (2005). Does the profile of income inequality matter for economic growth? *Journal of Economic growth*, 10:273–296.

Ward, Z. (2020). Internal migration, education, and intergenerational mobility: Evidence from american history. *Journal of Human Resources*. 0619-10265R2.

Warren, J. R., Pfeffer, F. T., Helgertz, J., and Xu, D. (2020a). Linking 1940 u.s. census data to the health and retirement survey: Technical documentation.

Warren, J. R., Pfeffer, F. T., Helgertz, J., and Xu, D. (2020b). Linking 1940 u.s. census data to the panel study of income dynamics: Technical documentation. page 1–9.

WCO (2013). *HS Classification Handbook*. Number November.

Winkler, W. E. (2006). Overview of record linkage and current research directions. (2006-02).