Economic Studies 206

Raoul van Maarseveen

Urbanization and Education

Department of Economics, Uppsala University

Visiting address: Kyrkogårdsgatan 10, Uppsala, Sweden
Postal address: Box 513, SE-751 20 Uppsala, Sweden
Telephone: +46 18 471 00 00
Telefax: +46 18 471 14 78
Internet: http://www.nek.uu.se/

ECONOMICS AT UPPSALA UNIVERSITY

The Department of Economics at Uppsala University has a long history. The first chair in Economics in the Nordic countries was instituted at Uppsala University in 1741.

The main focus of research at the department has varied over the years but has typically been oriented towards policy-relevant applied economics, including both theoretical and empirical studies. The currently most active areas of research can be grouped into six categories:

* Labour economics
* Public economics
* Macroeconomics
* Microeconometrics
* Environmental economics
* Housing and urban economics

Raoul van Maarseveen

Urbanization and Education

The Effect of Childhood Urban Residency on Educational Attainment

UPPSALA
UNIVERSITET

Dissertation presented at Uppsala University to be publicly examined in Hall 2, Ekonomikum, Kyrkogårdsgatan 10, Uppsala, Friday, 25 November 2022 at 10:15 for the degree of Doctor of Philosophy. The examination will be conducted in English. Faculty examiner: Daniel Sturm (London School of Economics and Political Science).

Abstract

van Maarseveen, R. 2022. Urbanization and Education. The Effect of Childhood Urban Residency on Educational Attainment. *Economic studies* 206. 210 pp. Uppsala: Department of Economics, Uppsala University. ISBN 978-91-506-2972-9.

Essay I: Does rural to urban migration in developing countries improve the opportunities available to children? And does childhood urban exposure increase educational attainment? Using census data for 14 African countries combined with an age-at-move design, I show that childhood exposure to cities significantly raises primary school completion, school attendance, and literacy rates. The increase in educational attainment is robust to the inclusion of household fixed effects, visible in all subgroups and countries, and particularly large for girls. The paper hence provides evidence of a channel through which urbanization can stimulate economic growth in developing countries, even in the absence of structural transformation.

Essay II: Despite the large urban-rural income gap across the developing world, it remains unclear to what degree this reflects the causal effect of urban residency. This paper presents new evidence by investigating the effect of urban residency during childhood on economic outcomes in adulthood. Causal identification is obtained from an age-at-move design combined with high-quality Brazilian census data. The analysis shows that spending childhood in an environment one log-point denser increases adulthood earnings and wages by 2 - 3 percent. Around half of this effect is due to an increase in educational attainment. The findings suggest that the previous literature, by exclusively focusing on urban exposure during adulthood, has underestimated the causal effect of urban residency on earnings by 50%.

Essay III: Despite a large urban-rural education gap in many countries, little attention has been paid to whether cities enjoy a comparative advantage in the production of human capital. Using Dutch administrative data, this paper finds that conditional on family characteristics and highly predictive measures of cognitive ability, children who grow up in urban regions consistently attain higher levels of human capital compared to children in rural regions. The elasticity of university attendance w.r.t. density is 0.07, which is robust across a wide variety of specifications. Hence, the paper highlights an alternative channel to explain the rise of the city.

Essay IV (with Niklas Bengtsson and Adrian Poignant): Industrialization is a ubiquitous feature of economic development, but the consequences for workers remain poorly understood. In this paper, we study the effects of worker displacement during the second industrial revolution in the Swedish iron industry. Using linked census data and a novel data source on production, we find that displaced ironworkers were 23 pp more likely to exit the industry, 25 pp more likely to migrate internally and lost 10% of their earnings relative to other workers. While the displacement effects persist over time for workers, we find no evidence of spillovers on their children.

Keywords: Urban Residency, Educational Attainment, Urbanization, Human Capital, Urban Migration

Raoul van Maarseveen, Department of Economics, Box 513, Uppsala University, SE-75120 Uppsala, Sweden.

ISSN 0283-7668
ISBN 978-91-506-2972-9
URN urn:nbn:se:uu:diva-485674 (http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-485674)

To Merel

Acknowledgments

Although only my name appears on the cover of the dissertation, this work has benefited from the guidance and support of many individuals. From arriving on the first day in Uppsala to writing the final sentences, I have met an incredible amount of supportive people who have helped with the creation of the dissertation, shaped my love for research, and made this journey a joyful one.

First of all, I want to thank my PhD supervisors, Niklas Bengtsson and Matz Dahlberg, without whose wisdom, encouragement and support throughout the last few years, this project would have been impossible. Niklas has made sure that I always felt comfortable sharing and discussing my ideas at any stage, provided invaluable comments and encouragement, and has been an excellent role model for balancing academic and personal life. Our joint research into the 19th century Swedish iron industry has been greatly rewarding. Matz has always been enthusiastic and supportive of my various research ideas and projects throughout the PhD and made a great effort to connect me with the wider field of urban economics. I have enjoyed our many walks and coffees around Uppsala and I am particularly grateful for Matz his unwavering efforts to expand the urban community in Uppsala, from which both I and the dissertation have benefited greatly.

Apart from my PhD supervisors, this dissertation would not have been possible without the support of the various faculty members at the Department of Economics in Uppsala. I would particularly like to thank Torben Mideksa for his insights into navigating academia and academic life, which is not always clear for a young researcher. Mikael Bask, Mikael Elinder, Eva Mörk and Daniel Spiro have all greatly contributed to making the department a welcoming and productive place for PhD students. I would also like to thank George Graetz and Magdalena Domínguez for their detailed comments on the projects throughout the PhD. Tessa Bold, Johannes Haushofer, Per-Anders Edin and Peter Fredriksson, provided many helpful suggestions during the licentiate seminar and final seminar. Finally, I would also like to thank Jakob Molinder at the Economic History department for his support of our endeavors into the world of economic history during the fourth chapter of the dissertation.

Outside of Uppsala, I have benefited from the generosity of many. Albert Saiz has kindly hosted me for half a year at MIT, which was a great experience. I have never heard of a host as kind as Albert and his family, and I fondly remember our many coffees, the thanksgiving meal and the charity event for

Doctors without Borders, which saw us all freezing in the Atlantic Ocean in the middle of November. I am grateful for the financial support of the Jan Wallander and Tom Hedelius foundation, which enabled this visit. Daniel Sturm has been supportive of my research agenda throughout the dissertation, and I am grateful for his suggestions on the positioning of the various papers and for hosting me at the Center for Economic Performance at LSE in September 2022. Finally, would like to express my gratitude to Edward Glaeser for his encouragement to pursue my research interest in urbanization and education in the developing world at the early stages.

I would also like to thank many of the teachers who have instructed me in economics and research. Anna Tompsett has been one of the most dedicated teachers I have encountered anywhere, and her feedback both during the courses and extending throughout the PhD have been much appreciated. I also would like to thank the various teachers during my studies at Groningen University, in particular Steven Brakman, Tristan Kohl and Robbert Maseland, whose dedication to the undergraduate program have shaped my ambitions and continue to inspire me to this day. Finally, I would like to thank my colleagues at CPB for their support and for starting my journey toward becoming an academic researcher. In particular, I would like to thank Peter Zwaneveld who has been a great mentor throughout the years and whose support and encouragement were invaluable to start the PhD program and Paul Verstraten, who has been a great friend and with whom I have enjoyed exploring and developing our shared interests in the field of urban economics.

The PhD journey would not have been nearly as enjoyable without the support and distractions provided by my fellow PhD students at Uppsala University. First of all, I would like to express my gratitude to the "awesomest PhD cohort from 2010 to forever", consisting of Davide Gandolfi, Alice Hallman, Anna Johansson, Adrian Poignant, Markus Ridder, Josefin Videnord and Yakoslav Yakymovich, for making me feel welcome in Uppsala and for sharing a great deal of fun outside of academia, from skating on the Swedish lakes in winter to exploring the microbreweries of New England. Furthermore, the coffees with Zeynep Atabay, Elin Bostrom, Simon Ek, Edvin Hertegard, Mohammed Sepahvand, Anton Sundberg, Lucas Tilley, and many others have been both productive and very enjoyable. I am also extremely grateful to Akib Khan for our many discussions of new research ideas, and for graciously sharing the latest developments and calls for papers from the world of twitter.

Finally, I would like to thank my family for their encouragement during the PhD and in particular my brother Vincent for the great and lighthearted entertainment in the evenings. My wife Stephanie has been an endless source of support during the writing of the dissertation and in finding the way throughout the PhD process and academic life. I dedicate this dissertation to my proudest accomplishment during the PhD period, which is our Daughter Merel, whose smile in the morning and great sense of fun are the most profound source of happiness.

Uppsala, 30 September 2022
Raoul van Maarseveen

Contents

Introduction

Edward Glaeser famously argued in his 2011 book that cities are man's greatest invention, which made us "richer, smarter, greener, healthier and happier" (Glaeser, 2011). Although reasonable, the empirical support for the "smarter" assertion always felt somewhat tenuous. Certainly, workers in cities had been shown to be more productive (Glaeser and Mare, 2001), but it was unclear whether this was driven by an increase in the accumulation of knowledge and skills due to urban residency or by other factors such as better functioning labor markets that provide higher rewards to existing skills.[1] The lack of evidence has been particularly noteworthy for the period prior to labor market entry, childhood, when the majority of human capital accumulation occurs.

Investigating the relationship between urban residency and educational attainment forms the central topic of this dissertation. The world has rapidly urbanized over the last 30 years, with the majority of children now growing up in an urban environment. This trend will continue for the foreseeable future, with Africa alone expected to add another 950 million urban residents between 2020 and 2050 (OECD, 2020). A substantial share of these urban residents are children, with over 40% of the urban population in sub-Saharan Africa under the age of 18. Despite these trends, we know very little about how urban residency affects the outcomes of children. Does spending childhood in a city improve educational attainment and lead to better learning outcomes? Or do urban disamenities, such as crime, congestion and the distractions that cities provide, mean that children are better off in the countryside? And how does spending childhood in a city eventually affect labor market outcomes in adulthood?

Beyond the level of the individual, studying the relationship between urbanization and education has significant societal and political implications. In recent years much attention has been paid to the growing urban-rural divide in Western countries, particularly in the UK and USA. This divide is visible along both economic dimensions and political attitudes, such as the vote on Brexit and political voting patterns. Implicitly, it is often assumed that these divides reflect spatial selection, with liberal and highly educated households selecting into cities and conservative and less educated households moving to the countryside. Less attention has been paid to how urban and rural residency shapes educational outcomes, preferences and opportunities. This distinction

[1] Follow-up work by Gould (2007) and De La Roca and Puga (2017) showed that urban workers experience faster wage growth which might be evidence of faster learning, although direct evidence on the relationship between learning and urban residency remains scarce.

is important, because if living in a city or a rural area causally shapes the educational outcomes of its residents, then this may contribute to the further divergence in outcomes between residents of urban and rural areas.

The key challenge in studying the relationship between urban residency and educational attainment comes from the fact that households do not randomly select into cities and rural environments. We know from previous research that households who move to cities are typically better educated (Young, 2013), wealthier (Bryan, Chowdhury and Mobarak, 2014) and have higher cognitive ability (Hamory et al., 2020). As such, we would expect the children of households who choose to live in an urban environment to have better educational outcomes, regardless of where they actually live. Thus, we cannot simply compare the educational outcomes of individuals who live in cities and in the countryside and interpret the differences as evidence of any causal effect.

This challenge of spatial selection is not unique to this dissertation, but is a central challenge for most of the work in urban and regional economics. In an ideal context, we would solve this problem by randomly assigning some families and their children to live in cities and others to live in rural areas. Such random assignment in economics often takes the form of a randomized controlled trial (RCT), which are used extensively in the field of development economics. However, such approaches are typically infeasible in the realm of urban economics, since we cannot randomly assign people to locations in virtually all cases.[2] A different method is to use natural experiments where some groups of individuals is forced to move, such as the work by Nakamura, Sigurdsson and Steinsson (2022), who study the effect of the forced displacement in Iceland following a volcanic eruption. However, such natural experiments are rare and typically unavailable in most contexts of interest.[3]

This dissertation instead leverages the strengths and depth of administrative data and census data from a variety of contexts to investigate the effect of urban residency on educational outcomes. The contexts studied in this dissertation have been selected because of some unique features in the context or data that help us answer the central question. Despite the fact that the contexts in this dissertation are very different in their level of development and urbanization, the lessons learned from the various essays are remarkably consistent and together tell us a coherent story. I will now briefly discuss and present the

[2] An important exception here is actually when it comes to international refugees, who are in many cases randomly assigned to a location in their host country. A recent study by Eckert, Hejlesen and Walsh (2022) uses the random assignment of international refugees in Denmark to study the effect of living in large cities.

[3] Some related studies use political shocks which displaced populations to other areas, such as the work by Becker et al. (2020) studying the displacement of Polish households following the annexation of parts of Poland by the USSR after world war II and the work by Sarvimäki, Uusitalo and Jäntti (2022) studying the displacement of Finnish households following the annexation of parts of Finland by the USSR. However, it is difficult to use such displacement shocks to separate the effect of a change of environment from the displacement shock itself.

findings of the individual essays, before discussing the broader lessons which we can learn from this dissertation and the questions that the findings raise for future research.

The Essays

I. The effect of urban migration on educational attainment: evidence from Africa

In the first essay, I study the effect of growing up in a developing world city on educational outcomes in the context of Africa. Africa has rapidly urbanized during the last 40 years, with the share of individuals living in a city almost doubling from 22% in 1980 to 42% in 2021.[4] This growth is set to continue in the coming decades, with Africa's cities expected to add another 950 million residents between 2020 and 2050 (SWAC & OECD, 2020). Many of the urban residents in Africa are young, with over 40% of the urban residents being under the age of 18.[5] At the same time, the living conditions in the cities have substantial room for improvement, with the UN reporting in 2016 that over half of Africa's urban residents live in what they describe as "slum-like" conditions (UN DESA, 2016).[6]

Considering these trends, it is important to understand how urban residency in Africa affects the opportunities of children and how the shift towards cities affects the skills and human capital of future generations. To answer these questions, I use the census data from 14 African countries with a joint population of 550 million individuals. For identification, I bring the methodological innovations of Chetty and Hendren (2018) to the context of urban migration and identify the effect of spending childhood in an urban environment of children who move to and from cities at different ages. Under the assumption that the selection of households into cities does not vary with the age of children at the time of the move, one can compare the outcomes of children who moved to a city at various ages to estimate the effect of spending childhood in an urban environment.

The main finding of the first essay is that spending childhood in a city significantly increases the educational attainment of children. Spending childhood in a large city rather than a fully rural environment increases primary school completion by 8 - 10 percentage points, school attendance at ages 13 - 18 by 4 percentage points and literacy rates by 4 percentage points. The effects of urban residency are concentrated between ages 6 and 12, with little effect of spending additional years in the city prior to age 6 or after age 12 on the educational outcomes.

[4]These statistics refer to Sub-Saharan Africa to be precise.

[5]Statistic based on 27 African countries with census data available on IPUMS international.

[6]Defined as lacking access to water, electricity, improved sanitation or durable housing.

Regarding the robustness of the results, I do not find any evidence that the results are driven by a different selection of households into the cities depending on the age of their children. The results are similar when comparing children between households as well as when comparing children within the same household. I also do not find any evidence that the results are driven by other shocks, such as violence or weather shocks, that may affect both the educational outcomes of children and the migration decisions.
Furthermore, the positive effects of childhood urban residency are visible for all demographic groups and countries under analysis. In particular, the positive effects are visible even among households who live under the aforementioned "slum-like" conditions and lack access to water, electricity or durable housing. As such, it seems to be less about striking it rich in the city as a household, but more about the general improvement of opportunities for children in an urban setting. Furthermore, girls appear to benefit slightly more from childhood urban exposure compared to boys, which suggests that the urbanization of Africa may contribute towards gender equality in educational outcomes. However, the effects of urban residency on educational attainment decline over time, with smaller effects of urban residency on educational outcomes in the later censuses for countries with multiple available censuses. Such decreases might be a general result as countries become wealthier, but may also reflect the push towards universal primary education in the 2000's as part of the UN millennium goals.
The final question studied in the essay is what mechanisms can explain the increase in educational attainment. I find no evidence that the age at which one migrates to the city affects family formation decisions. Instead, I find that children who spend more of their childhood in the city reduce labor force participation at ages 13 - 18. The substitution of employment in favor of schooling can be explained by a variety of channels, such as higher returns to education in cities, better access to schools, changes in peers, changes in family income and changes in the opportunity costs of schooling. Cross-sectional evidence from Guinea suggests that better access to schools in urban regions and the lower opportunity costs of education may play an important role in explaining this trade-off between schooling and working. However, this evidence is from a single context and cannot be interpreted as causal. As such, the decomposition of the main finding into the individual mechanisms within a causal framework forms an interesting challenge for future research.

II. The effect of childhood urban residency on earnings: Evidence from Brazil

An important question raised by the first essay is whether childhood urban residency eventually translates into better labor market outcomes in adulthood or whether the effects dissipate over time. Answering this question is difficult in the African context since information on earnings, wages and other labor market outcomes are largely missing. Therefore, the second essay studies the effect of childhood urban residency in a different setting: Brazil. The Brazilian census data are unique in the world in that they record information on the migration history as well as earnings, wages and hours worked.

Using a similar identification strategy as the first essay, the second essay shows that spending childhood in a denser environment also positively affects adulthood labor market outcomes. Spending childhood in a region which is one log-point denser increases earnings in adulthood by 2 - 3 %. The increase is fully driven by higher wages, with no effects of childhood urban residency on hours worked or labor market participation. Hence, the effects of childhood urban residency persist into adulthood in the form of higher earnings.

The main driver of the higher earnings is the increase in educational attainment. Similar to the African context, I find that children who grow up in an urban environment have better educational outcomes. Spending childhood in an environment one log-point denser increases primary school completion by 2 percentage points and years of schooling by 0.2 years. This increase in educational attainment can explain the majority of the increase in wages. Hence, the second essay confirms and extends the findings of the first essay in a very different setting, a large middle-income country in Latin America.

Regarding the robustness of the results, there is no evidence that the results are driven by a different selection of households into cities by the age of children, the selected urban measure or selective migration during adulthood. Furthermore, I find that the gains of childhood urban residency are visible for all demographic groups, although children who report their race as non-white benefit somewhat less from urban residency compared to children who report their race as white.

Finally, the findings in the Brazilian context allow me to assess how important childhood urban residency is compared to adulthood urban residency. Previous work has estimated an elasticity of earnings with respect to adulthood population density of 0.04 globally (Ahlfeldt and Pietrostefani, 2019) and 0.052 for Brazil specifically (Chauvin et al., 2017). In comparison, I find an elasticity of earnings with respect to childhood density of 0.02 - 0.03. This suggests that previous studies, which have attempted to explain why urban residents are wealthier by only studying urban residency during adulthood, have underestimated the causal effect of urban residency on earnings by around 50%.

III. The urban-rural education gap: do cities indeed make us smarter?

The first two essays show that childhood urban residency positively affects the educational outcomes and labor market outcomes of children in Africa and Brazil. An interesting question is whether such benefits are limited to developing countries or whether similar urban-rural differences exist in developed countries.

In the third essay, I examine the urban-rural gradient in educational investment in The Netherlands using some unique institutional features. At the end of primary school, students select into one of three levels of secondary school (academic, professional or vocational). The three levels of secondary school differ in their difficulty, length and access to tertiary education. Just before making this decision, students participate in the national primary school exit exam, which provides a comprehensive account of students learning outcomes. Using the primary school exit exam scores and a wide range of observed household characteristics, I investigate whether students in urban environments are more likely to enroll in the academic track.

I find that conditional on the detailed test scores and a wide variety of family characteristics, students in urban environments are substantially more likely to enroll in the academic secondary school track. A one log-point increase in density is associated with a 1.68% increase in academic track enrollment, from a base of 23%. I find similar differences between students in urban and rural environments at the next decision point, when students who completed the academic track have to decide whether to enroll in a university. At this stage, a one log-point increase in density is associated with a 0.8 % increase in university enrollment conditional on high-school GPA and family characteristics, from a base of 84%. Hence, students who grow up in an urban environment are significantly more likely to enroll in the academic secondary school track and subsequently to enroll in university compared to students in rural environments, conditional on the observed learning outcomes and family background. Taken together, the results imply an elasticity of university enrollment with respect to childhood population density of 0.07.

In terms of robustness, I do not find evidence that the differences are driven by unobserved heterogeneity between urban and rural students. Furthermore, I use an IV-approach based on the 1840 population density to rule out that the results are driven by factors correlated with contemporary density. Finally, I rule out that the results are driven by broader regional differences between the heavily urbanized Western region and the more peripheral North and East, by showing that the urban-rural gradient in educational investment is also visible within the provinces. Hence, the third essay shows that spending childhood in an urban environment also improves educational outcomes in the setting of a high-income developed country.

IV. The social cost of industrialization - Evidence from the 19th century Swedish iron industry (With Niklas Bengtsson and Adrian Poignant)

Finally, it is important to notice that the migration towards cities and the success of the cities may come at a cost to the rural regions. In many cases, the migration towards the cities and the subsequent structural transformation of the economy can have far-reaching implications for firms and workers operating in rural regions. In the fourth essay, we analyze what happens when the process of industrialization concentrates employment in emerging towns and the closure of industries in the rural environment. We study this question in the context of the Swedish iron industry in the late 19^{th}-century.

During this period, technological advances and economic liberalization favored the concentration of the industry in a small number of large factories. This change resulted in the closure of many smaller iron artisanal production sites in the countryside that had been operating since the early 17^{th} century, whereas the surviving plants expanded and often developed into small towns. The output of the industry increased fourfold in the latter half of the nineteenth century, while the number of production sites declined by $2/3^{th}$.

In the fourth essay, we examine the effect of this industrial transformation on the labor market outcomes of displaced workers. To document the changes on the firm side and to identify the production site closures, we digitalized the annual production reports from the Swedish National Archives covering the universe of iron forges and furnaces between 1860 and 1893. We combine this data with linked census data from the 1880, 1890 and 1900 population censuses, which allows us to track the ironworkers over time. We focus in the analysis on the iron site closures between 1880 and 1890 since the production data and census data overlap for this period.

We find that workers displaced by the site closures are negatively affected. Workers suffering a displacement shock are 25 pp less likely to work in the iron industry in 1890, and are 25 pp more likely to have migrated to a different parish of residence in 1890. Despite the geographical mobility and changes in occupations, we find that the incomes of displaced ironworkers are negatively affected. Displaced ironworkers find employment in occupations that paid on average 10% less compared to ironworkers that are not displaced. Interestingly enough, these negative effects do not dissipate over time. The differences between displaced and non-displaced ironworkers are similar when we study their outcomes in 1890 and outcomes in 1900. However, we do not find evidence that the displacement shock carries over to the next generation. The children of displaced ironworkers are less likely to work in the iron industry and more likely to have migrated during childhood, but their earnings are not significantly different from the earnings of children of non-displaced ironworkers.

Finally, we investigate the effectiveness of the migration response, given that around half of the displaced ironworkers moved to a different parish between 1880 and 1900. We find that the effects of displacement are similar for those who did and did not choose to migrate at the time of the 1890 census. However, at the time of the 1900 census, the ironworkers who migrated had incomes similar to that of non-displaced ironworkers, with the negative effects concentrated among those who remained in their original parish of residence. Thus, the high level of geographic mobility among workers during this era seems to have been an important way in which the workers coped with the displaced shock and the industrial transformation more generally.

Implications of the dissertation for the field of economics

After presenting the findings, an important question how the essays contained in this dissertation contribute to our understanding of the world and the field of economics. In my view, there are four contributions of this dissertation.

The first contribution of this dissertation is to show how urban residency during childhood affects the outcomes of children. Despite a large literature on the effects of urban residency on adults, the effects on children have been largely ignored. In this dissertation, I show that childhood urban residency on average improves the educational and labor market outcomes of children. In the first three essays, I show that spending childhood in an urban environment increases educational attainment. In the second essay, I show that the gains in educational attainment also translate in better labor market outcomes in adulthood. Given that an extensive literature shows that educational attainment also affects a wide host of non-labor market outcomes such as health (Silles, 2009), life satisfaction (Oreopoulos, 2007) and marriage quality (Kaufmann, Messner and Solis, 2013), it seems likely that childhood urban residency improves the opportunities for children also along a range of other dimensions.

This finding stands in contrast to a literature that highlights the potential of developing world cities to act as poverty traps due to the presence of large-scale slums (Marx, Stoker and Suri, 2013). While the findings of this dissertation do not reject the notion that developing world cities may act as poverty traps for some households and their offspring, this does not appear to resemble the experience of the average urban migrant. As such, the three essays also contribute to the newly emerging literature that seeks to measure and understand the geographic variation in the opportunities available to children, following the work by Chetty et al. (2014) and Chetty and Hendren (2018) in the US, Deutscher (2020) in Australia and Alesina et al. (2021) in Africa, by focusing on the urban-rural differences in opportunities.

Second, the dissertation provides a better understanding of the costs and benefits of urban migration in the developing world. A large recent literature has analyzed the effect of urban migration on labor market outcomes (Young, 2013; Bryan, Chowdhury and Mobarak, 2014; Akram, Chowdhury and Mobarak, 2017; Alvarez, 2020; Hamory et al., 2020; Lagakos et al., 2020), but these studies have exclusively focused on urban migration during adulthood. An important reason has been methodological, since the migration literature has traditionally relied on individual-fixed effect models to estimate the gains of rural-to-urban migration. Since childhood migrants typically have no observed income prior to migration, the individual-fixed models could not be estimated to capture the effects of urban residency during childhood. In this dissertation, I use the methodology based on Chetty and Hendren (2018) to expand the world of rural-to-urban migration studies to also encompass the effects of rural-to-urban migration during childhood. The findings highlight that children also benefit from rural-to-urban migration due to their gains in educational attainment. Furthermore, existing studies have used their estimates of the gains of adulthood migration to explain why urban residents in general are wealthier compared to rural residents. The findings of this dissertation suggest that such comparisons have underestimated the causal effect of urban residency by around 50%, by failing to account for the effect of urban residency during childhood.

Third, the dissertation provides a novel contribution on the role of cities in economic development. Economic theory has traditionally approached urbanization and structural transformation as two sides of the same coin, with the migration of households into cities enabling the transformation of an agricultural society into an industrial society (Lewis, 1954; Michaels, Rauch and Redding, 2012). The structural transformation of the Swedish iron industry described in the fourth essay for instance is very much in line with this classical model. Hence, the rapid urbanization of Africa over the recent decades without evidence of subsequent structural transformation has been the source of much debate (Gollin, Jedwab and Vollrath, 2016; Lall, Henderson and Venables, 2017; Henderson and Kriticos, 2018), leading to the question whether urbanization benefits economic development in such context. In this dissertation, I provide clear evidence of a different mechanism through which urbanization can contribute to economic development, namely by increasing the human capital stock of a country. This channel is particularly relevant as Gennaioli et al. (2013) find that human capital is the most important predictor of economic growth across the regions of the world.

Fourth, the dissertation contributes to a long-standing literature in urban economics explaining why cities are more productive (Marshall, 1890; Glaeser, 1999). A large literature shows that cities are more productive in using existing stocks of human capital in the production of goods and services (see Ahlfeldt and Pietrostefani (2019) for a recent review) and in attracting existing stocks of human capital by offering better amenities (Glaeser, Kolko & Saiz, 2001). However, the question of whether cities also have a comparative advantage in the production of human capital itself has received limited attention. Glaeser (1999), Lucas (2004) and Gould (2007) argue that individuals should have higher levels of human capital accumulation in cities, but their contributions were focused on human capital accumulation during working life. More recently De La Roca and Puga (2017) show that this indeed is the case empirically, but the focus has remained on the role of cities in human capital acquisition among workers. This dissertation expands the research to the period prior to labor market entry, when arguably most of the human capital accumulation takes place, and shows that those who grow up in urban environments already enter the labor market with higher stocks of human capital. As such, this dissertation highlights a different explanation for the recent economic success of cities in both the developed and the developing world, which has gone largely unnoticed across the various fields of economics.

Finally, the dissertation naturally has some limitations and certain questions are left for future research. The foremost question raised by the findings of this dissertation is why exactly children in cities have better educational outcomes. Possible channels include the higher income of parents, higher returns to education in cities, better access to education, changes in opportunity costs and changes in peers and role models. The lack of data and the high correlation between the various mechanisms in contexts where data is available makes it empirically difficult to distinguish between the mechanisms without large shocks to each individual mechanism. However, further research into the mechanisms would be helpful to translate the findings of this dissertation into concrete policy implications. Do the better educational outcomes in cities reflect certain factors which could be replicated in a rural setting, such as better access to schools, or are the results driven by factors such as the higher returns to education in the cities or better job prospects of parents, which would be harder to replicate in a rural setting?

Second, the dissertation does not make any claims regarding the welfare of children in urban and rural settings. From the perspective of the individual child, the lower educational attainment of children in rural settings might be efficient, given the local economic conditions. At the same time, it is hard to imagine that attaining basic numerical and literacy skills would be detrimental to the welfare of children in any context. Nonetheless, a holistic assessment of whether urban residency on average improves the welfare of children would require more expansive and interdisciplinary work, particularly on how urban and rural residency affect non-economic outcomes such as (mental) health, social capital and life satisfaction. These interesting questions are beyond the scope of this dissertation and deserve a closer examination in future research.

References

Ahlfeldt, Gabriel M, and Elisabetta Pietrostefani. 2019. "The economic effects of density: A synthesis." *Journal of Urban Economics*, 111: 93–107.

Akram, Agha Ali, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. 2017. "Effects of emigration on rural labor markets." National Bureau of Economic Research.

Alesina, Alberto, Sebastian Hohmann, Stelios Michalopoulos, and Elias Papaioannou. 2021. "Intergenerational mobility in Africa." *Econometrica*, 89(1): 1–35.

Alvarez, Jorge A. 2020. "The agricultural wage gap: Evidence from brazilian micro-data." *American Economic Journal: Macroeconomics*, 12(1): 153–173.

Becker, Sascha, Irena Grosfeld, Pauline Grosjean, Nico Voigtländer, and Ekaterina Zhuravskaya. 2020. "Forced Migration and Human Capital: Evidence from Post-WWII Population Transfers." *American Economic Review*, 110(5): 1430–1463.

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. 2014. "Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh." *Econometrica*, 82(5): 1671–1748.

Chauvin, Juan Pablo, Edward Glaeser, Yueran Ma, and Kristina Tobio. 2017. "What is different about urbanization in rich and poor countries? Cities in Brazil, China, India and the United States." *Journal of Urban Economics*, 98: 17–49.

Chetty, Raj, and Nathaniel Hendren. 2018. "The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects." *The Quarterly Journal of Economics*, 133(3): 1107–1162.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. 2014. "Where is the land of opportunity? The geography of intergenerational mobility in the United States." *The Quarterly Journal of Economics*, 129(4): 1553–1623.

De La Roca, Jorge, and Diego Puga. 2017. "Learning by working in big cities." *The Review of Economic Studies*, 84(1): 106–142.

Deutscher, Nathan. 2020. "Place, Peers, and the Teenage Years: Long-Run Neighborhood Effects in Australia." *American Economic Journal: Applied Economics*.

Eckert, Fabian, Mads Hejlesen, and Conor Walsh. 2022. "The return to big-city experience: Evidence from refugees in Denmark." *Journal of Urban Economics*, 103454.

Gennaioli, Nicola, Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer. 2013. "Human capital and regional development." *The Quarterly journal of economics*, 128(1): 105–164.

Glaeser, Edward. 2011. *Triumph of the city: How urban spaces make us human.* Pan Macmillan.

Glaeser, Edward L. 1999. "Learning in cities." *Journal of urban Economics*, 46(2): 254–277.

Glaeser, Edward L, and David C Mare. 2001. "Cities and skills." *Journal of labor economics*, 19(2): 316–342.

Gollin, Douglas, Remi Jedwab, and Dietrich Vollrath. 2016. "Urbanization with and without industrialization." *Journal of Economic Growth*, 21(1): 35–70.

Gould, Eric D. 2007. "Cities, workers, and wages: A structural analysis of the urban wage premium." *The Review of Economic Studies*, 74(2): 477–506.

Hamory, Joan, Marieke Kleemans, Nicholas Y Li, and Edward Miguel. 2020. "Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata." *Journal of the European Economic Association*.

Henderson, J Vernon, and Sebastian Kriticos. 2018. "The development of the african system of cities." *Annual Review of Economics*, 10: 287–314.

Kaufmann, Katja Maria, Matthias Messner, and Alex Solis. 2013. "Returns to elite higher education in the marriage market: Evidence from Chile." *Available at SSRN 2313369*.

Lagakos, David, Samuel Marshall, Ahmed Mushfiq Mobarak, Corey Vernot, and Michael E Waugh. 2020. "Migration costs and observational returns to migration in the developing world." *Journal of Monetary Economics*, 113: 138–154.

Lall, Somik Vinay, J Vernon Henderson, and Anthony J Venables. 2017. *Africa's cities: Opening doors to the world.* The World Bank.

Lewis, William Arthur. 1954. "Economic development with unlimited supplies of labour." *The Manchester School*, 22(2): 139–191.

Lucas, Robert E Jr. 2004. "Life earnings and rural-urban migration." *Journal of political economy*, 112(S1): S29–S59.

Marshall, Alfred. 1890. "Principles of Economics, 8th edn (1920)." *London, Mcmillan*.

Marx, Benjamin, Thomas Stoker, and Tavneet Suri. 2013. "The economics of slums in the developing world." *Journal of Economic perspectives*, 27(4): 187–210.

Michaels, Guy, Ferdinand Rauch, and Stephen J Redding. 2012. "Urbanization and structural transformation." *The Quarterly Journal of Economics*, 127(2): 535–586.

Nakamura, Emi, Jósef Sigurdsson, and Jón Steinsson. 2022. "The gift of moving: Intergenerational consequences of a mobility shock." *The Review of Economic Studies*, 89(3): 1557–1592.

Oreopoulos, Philip. 2007. "Do dropouts drop out too soon? Wealth, health and happiness from compulsory schooling." *Journal of public Economics*, 91(11-12): 2213–2229.

Sarvimäki, Matti, Roope Uusitalo, and Markus Jäntti. 2022. "Habit Formation and the Misallocation of Labor: Evidence from Forced Migrations." *Journal of the European Economic Association*, jvac037.

Silles, Mary A. 2009. "The causal effect of education on health: Evidence from the United Kingdom." *Economics of Education review*, 28(1): 122–128.

SWAC & OECD. 2020. "Africa's Urbanisation Dynamics 2020: Africapolis, Mapping a New Urban Geography." OECD Publishing, Paris.

UN DESA. 2016. *The Millennium Development Goals Report 2015.* UN.

Young, Alwyn. 2013. "Inequality, the urban-rural gap, and migration." *The*

Quarterly Journal of Economics, 128(4): 1727–1785.

1. The effect of urban migration on educational attainment: evidence from Africa

I would like to express my gratitude to Douglas Almond, Niklas Bengtsson, Tessa Bold, Matz Dahlberg, Magdalena Domínguez, Per-Anders Edin, Peter Fredriksson, Edward Glaeser, Johannes Haushofer, Vernon Henderson, Nathaniel Hendren, Mounir Karadja, Francis Kramarz, David Lagakos, David McKenzie, Torben Mideksa, Adrian Poignant, Albert Saiz, Wayne Sandholtz, Daniel Sturm, Lucas Tilley and Anna Tompsett, as well as the participants to the American UEA conference, European UEA conference, EALE conference, PAA conference, SOLE conference, EEA conference, AEA conference, RES conference, Nordic Conference in Development Economics, 15th Migration and Development Conference and Economic of migration seminar series for their helpful comments and suggestions on previous versions of this paper.

1.1 Introduction

Despite the rapid urbanization of the developing world, little is known about the effects of urban residency during childhood. On the one hand, cities offer access to employment and educational institutions that are typically unavailable in rural areas. On the other hand, developing world cities are characterized by a lack of access to public services, which raises concerns that they may act as poverty traps. This is particularly the case in Sub-Saharan Africa, where more than half of the urban residents live in slum-like conditions (UN DESA, 2016).[1] Does childhood urban residency in such a context improve economic outcomes and is rural to urban migration a viable pathway for families to increase the opportunities available to their offspring?

The effect of childhood urban residency on economic outcomes has received surprisingly little attention in the economics literature. Studying such questions has historically been difficult due to the selection of households into cities, with individual-fixed effect estimations typically not feasible for children. To overcome this challenge I use an age-at-move design based on the work by Chetty and Hendren (2018). The effects of childhood urban exposure are identified of children who move between urban and rural regions at different ages and thus experienced different levels of urban exposure throughout childhood. The data consists of 26 African censuses obtained from IPUMS International, which contain detailed information on individuals place of residence, migration history and educational attainment. The censuses contain information on 57 million individuals and are representative for fourteen countries with a population of 550 million.

The analyses reveal that childhood urban residency significantly increases educational attainment. Spending childhood in a region that is one percent more urbanized increases primary school completion by 0.08 - 0.10 percentage points, school attendance at ages 13 - 18 by 0.04 percentage points, and literacy rates by 0.04 percentage points. Put differently, growing up in a large city instead of a region at the 25th percentile of the density distribution increases primary school completion by 7 - 8 percentage points. The urban exposure effects are concentrated between the ages 6 and 12, with little evidence that urban residency prior to the age of 6 or after the age of 12 affects educational attainment.

The results do not depend on the definition of 'urban' or the selection of countries. The baseline estimates use the population share living in a town with over 10,000 inhabitants as urbanization measure, which yields a consistent definition of urbanization across Africa. However, the estimates are similar when using population density, when using country-specific urban definitions or when restricting attention to capital cities. The urban exposure effects are also visible within the majority of countries, with significant estimates in 13

[1] Defined as lacking access to clean water, improved sanitation, durable housing or sufficient living space (UN DESA, 2016).

out of the 14 countries in the case of primary school completion and 7 (6) out of the 13 (12) countries in the case of school attendance (literacy). Hence, the positive effect of childhood urban exposure on educational outcomes is visible across a wide range of settings.

One concern is that the results might be driven by a different selection of households into urban and rural areas depending on the age of children. However, the estimates are similar when controlling for household characteristics and when including household-fixed effects, suggesting that the results are not driven by selection based on time-invariant household characteristics. A second threat to identification is that shocks in the region of origin may induce (or force) households both to move into cities and to increase educational investment. Becker et al. (2020) for instance find that forced displacement in Poland shifted preferences towards education. However, I find no evidence that the results are driven by violence in the region of origin or agricultural income shocks. More generally, I focus on the years of exceptional migration outflows in the region of origin to identify the subset of households whose timing of migration is least likely to be driven by individual household shocks, following Alesina et al. (2021). The estimates again remain unchanged when restricting attention to this subset of migrations, suggesting that migration in response to household shocks is also unlikely to drive the results.

The heterogeneity analyses reveal three further findings. First of all, childhood urban exposure effects are present for both genders, but are stronger for girls. This stands in contrast to the findings of Chetty and Hendren (2018) in the US and Deutscher (2020) in Australia, who both find stronger neighborhood effects for boys. Second, the estimates are similar across households with different levels of parental education, suggesting that childhood urban exposure neither complements nor substitutes for household characteristics favorable towards education. This stands in contrast to the findings of Chetty and Hendren (2018) who find larger neighborhood exposure effects for the children of wealthier families in the US. Third, the urban exposure effects are significantly larger for families who move to more urbanized regions compared to families who move to less urbanized regions, which might be indicative of efficient spatial sorting by households into cities as previously suggested by Young (2013).

Finally, the question arises which mechanisms explain the increase in educational attainment due to urban residency. Despite large urban-rural differences in teenage births and marriage rates, I find no evidence that the length of urban residency matters for family formation decisions. Instead, the results point towards the importance of local labor market conditions, with urban exposure reducing the probability that a child engages in home production or works outside of the home. A variety of channels can explain this substitution of employment in favor of education, including higher returns to education in cities, better access to schools, household income effects and changes in opportunity costs. Household income effects are unlikely to explain the full results, since the urban exposure effects are visible and similarly sized for children of relatively deprived households, who lack access to electricity, water or durable housing. Evidence from Guinea suggests that better access to educational institutions in cities and the higher opportunity costs of education among rural households may also play an important role.

The two papers most closely related to this study are Nakamura, Sigurdsson and Steinsson (2021), who study the effect of a rural displacement shock following a volcanic eruption in Iceland, and Alesina et al. (2021), who extend the work by Chetty and Hendren (2018) to the African context. Compared to Nakamura, Sigurdsson and Steinsson (2021), I focus on urban migration in the developing world and study all urban migrants, including the majority not driven from their home following a severe displacement shock. I expand on the work by Alesina et al. (2021) by focusing on the causal impact of urban migration and more generally by studying the effects of migration on education attainment, rather than intergenerational mobility.

The findings of this paper contribute to three strands of literature. First of all, the paper contributes to the literature on urban migration and residency in the developing world, following Young (2013), Bryan, Chowdhury and Mobarak (2014), Lagakos et al. (2020) and Hamory et al. (2020). This paper is the first to extend the literature to include the effects on children, who make up 40% of the urban population in Africa. Despite the large literature on the risks of developing world cities as poverty traps for households (Marx, Stoker and Suri, 2013), I find that children on average benefit from urban migration and residency. Furthermore, the findings suggest that existing studies have underestimated the causal effects of urban residency on economic outcomes due to their exclusive focus on urban residency during adulthood, when skills and educational attainment are typically fixed. The paper also contributes to the wider literature on neighborhood effects following Chetty and Hendren (2018) in the US, Deutscher (2020) in Australia and Alesina et al. (2021) in Africa, by studying the effects of one of the most salient features of global migration, the rural to urban migration in the developing world.

Second, the paper contributes to the economic growth literature on the role of urbanization in economic development. A long tradition dating back to Lewis (1954) and Harris and Todaro (1970) views urbanization and structural transformation as two sides of the same coin, with migration into cities driving the growth of the manufacturing and export industry. However, Africa has experienced rapid urbanization over the last decades without any evidence of structural transformation (Gollin, Jedwab and Vollrath, 2016; Lall, Henderson and Venables, 2017; Henderson and Kriticos, 2018), which has led to the question whether and how urbanization benefits economic development. This paper contributes to this literature by highlighting an alternative channel through which urbanization can stimulate economic development, namely by increasing the stock of human capital. This channel is particularly relevant as Gennaioli et al. (2013) find human capital to be the most robust predictor of

economic growth across a wide variety of settings.
Finally, the paper contributes to the urban economics literature by offering a different explanation for the recent success of cities. The focus of the urban economics literature has been on how cities efficiently utilize existing stocks of human capital in the production of goods and services (Combes, Duranton and Gobillon, 2008) and how cities attract existing stocks of human capital by offering superior consumption amenities (Glaeser and Mare, 2001). A smaller literature following Glaeser (1999), Gould (2007) and De La Roca and Puga (2017) highlights the role of cities in the creation of human capital, but the focus has been on human capital accumulation during working life. This paper extends this research to the period prior to labor market entry and shows that cities increase human capital accumulation already in the period before labor market entry.
The remainder of the paper proceeds as follows. Section 1.2 briefly discusses the existing literature on rural to urban migration in developing countries. Section 1.3 discusses the identification strategy and data used in this paper. Section 1.4 presents the main findings and the various robustness checks. Section 1.5 presents the heterogeneity analyses and section 1.6 discusses the potential mechanisms. Section 1.7 concludes.

1.2 Rural to urban migration literature

Despite the rapid urbanization of the developing world in the last decades, measuring the effects of urban residency on individual outcomes has proved challenging. On the one hand, urban consumption levels and incomes are higher across the developing world (Young, 2013; Lagakos et al., 2020) and urban migrants experience large consumption gains compared to those who remain in rural regions (Beegle, De Weerdt and Dercon, 2011). On the other hand, households select into cities based on a variety of characteristics such as education (Young, 2013), wealth (Bryan, Chowdhury and Mobarak, 2014) and cognitive ability (Hamory et al., 2020), which complicates the estimation of the effects of urban residency and migration.
To account for such selection into cities, most papers have relied on individual-fixed effect models. Using both cross-sectional and panel data estimates, Hamory et al. (2020) and Lagakos et al. (2020) find that the individual returns to urban migration are far smaller than the cross-sectional differences between urban and rural households. These panel data estimates have been in line with experimental evidence on the effects of urban migration on wages provided by Bryan, Chowdhury and Mobarak (2014) and Akram, Chowdhury and Mobarak (2017). The difference between cross-sectional returns to urban residency and experimental and individual-fixed effect estimates are typically explained on the basis of Roy selection models, suggesting that individuals

already living in cities are particularly well suited for urban activities.[2]

One limitation of the existing literature has been the focus on urban migration during adulthood. While this has been driven by necessity to some degree, since the individual-fixed effect models require individuals to have worked in the rural areas before moving to the city, this approach also risks underestimating the returns to urban residency. The recent work by Chetty and Hendren (2018) and the Moving To Opportunity program research (Chetty, Hendren and Katz, 2016) highlight that the effect of moving to a region during adulthood may not necessarily be informative of the effect of permanent residency in the same region, since estimates based on adulthood migration do not capture the benefits of spending childhood in the same place. Hence, the existing literature may underestimate the gains of permanent urban residency insofar as childhood residency in urban environments influences individual outcomes. One potential channel through which urban residency during childhood could affect labor market outcomes is through human capital formation. The main use of educational attainment in existing studies has been to correct for observable differences between urban and rural residents in order to 'cleanly' estimate the effect of urban residency on wages and consumption. However, controlling for educational attainment risks underestimating the effect of urban residency insofar as educational attainment itself is a function of density. This concern is explicitly recognized by Henderson and Turner (2020) in a recent literature review. However, the degree to which urbanization affects educational outcomes remains an open question in the literature (Henderson and Turner, 2020; Hamory et al., 2020).

Investigating this question constitutes the central topic of this paper. Cross-sectional evidence reveals that substantial differences in educational attainment exist between urban and rural regions. Figure 1.1 shows primary school completion among 15 and 16 year olds for 27 African countries with available census data. Children in urban communities are on average 25 percentage points more likely to have completed primary school compared to children in rural areas. Similar differences exist in school attendance and literacy rates as shown in appendix figures 1.10 and 1.11. Are such differences in educational attainment between urban and rural areas merely reflective of the spatial selection of households into cities, or does urban residency itself responsible for the better educational attainment of children?

[2]On the other hand, Lagakos et al. (2020) argue that the returns to migration may not be informative of the gains of migration unless we have information about the joint distribution of the returns to migration and the migration costs.

Figure 1.1. Urban-Rural differences in primary school completion in Africa

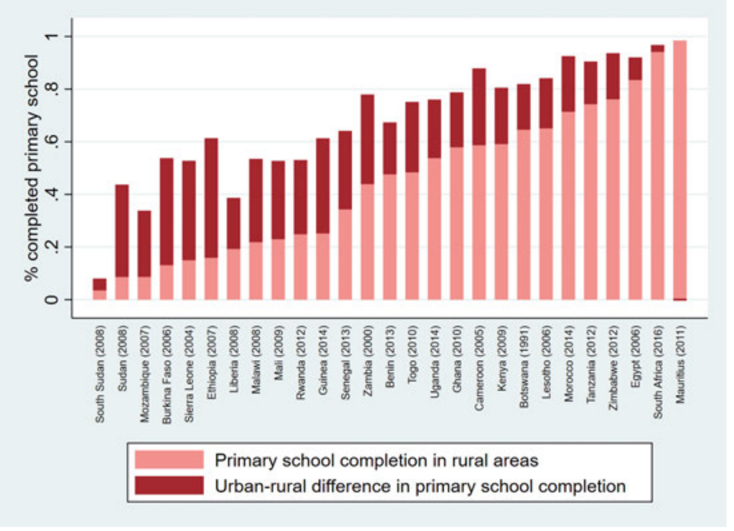

Note: Note: The figure is based on the most recent census available from IPUMS international for all African countries. The sample includes all children between the ages 15 - 16 at the time of the census. The definition of urban and rural is based on the definition of the statistical office in the relevant country. Primary school completion is defined as completion of the sixth year of schooling.

1.3 Data and methodology

1.3.1 Identification strategy

To separate the spatial selection of households from the effects of urban residency, I build on the novel identification strategy of Chetty and Hendren (2018). The key assumption underlying the approach of Chetty and Hendren (2018) is that the selection of households into better or worse neighborhoods does not vary with the age of children at the time of the move. Applied to the context of urban migration, the equivalent assumption is that the selection of households into more or less urban areas does not vary with the age of the children at the time of migration. The effect of urban residency during childhood under this assumption can be estimated by comparing the educational outcomes of children who moved between rural and urban areas at different ages, conditional on a variety of fixed effects. More formally, I use an estimating equation close to the semi-parametric equation of Chetty and Hendren (2018) shown in equation 1.1[3],where i denotes the individual, o the region of

[3]The key difference between the semi-parametric model in equation (5) of Chetty and Hendren (2018) and equation (1.1) in this paper is the separation of the disruption term from the fixed effect. Including the disruption term in the fixed effect has little effect on the baseline estimates, but does not leave sufficient degrees of freedom for estimation when household fixed effects are included. Hence, I separate the disruption term from the fixed effect term throughout the paper to facilitate comparison between models with and without the household fixed effect.

origin, c the census, s the birth cohort, m the age of the child at the time of the move, d the region of residence at the time of the census and h the household.

$$\text{Educational outcome}_i = \alpha_{ocs} + \sum_{m=1}^{15} I(m_i = m)\phi_{mc} + \sum_{m=1}^{15} b_m I(m_i = m)\Delta urbanization_{od} + (\theta_h) + \varepsilon_i \quad (1.1)$$

The first term captures the quality of the region of origin, which is allowed to vary by census year and birth cohort. The second term captures the disruption of moving, which may vary by country and the age of the child at the time of the move. The third term contains the coefficient of interest, b_m, which captures the effect of moving to a place one percent more urbanized at age m. This b_m parameter can be separated into a selection and an exposure effect, as highlighted in equation 1.2:

$$b_m = \beta_m + \delta_m \quad (1.2)$$

The first term (β_m) captures the causal effect of spending the years between migration and the census in an area that is one percent more urbanized, whereas the second term (δ_m) captures the selection of children into urban areas at age m ($\frac{cov(\Delta urbanization, \varepsilon_i)}{var(\Delta urbanization)} \mid m_i = m$). The key assumption underlying the identification strategy is that the selection of households into more and less urbanized regions does not differ by the age of the children at the time of the move, which implies that $\delta_m = \delta$ for all m. Under this assumption it is possible to causally estimate the effect of spending a year at age m in an area one percent more urbanized as $\gamma_m \equiv \beta_m - \beta_{m+1} = b_m - b_{m+1}$.
The baseline estimates of equation 1.1 uses all migrants, but it is possible to also include a household fixed effect term (θ_h) in equation 1.1. When included, the identification of b_m comes from the comparison of children within the same household who experienced different levels of urban exposure during childhood due to their difference in age at the time of migration. The inclusion of the household-fixed effect can serve as a robustness check on the identification strategy, since it controls for the selection based on time-invariant household characteristics. One aspect which the estimation strategy cannot directly control for is potential selection of households into cities in response to time-varying shocks. The potential of such time-varying household shocks to affect the estimates will be discussed in section 1.4.2.

1.3.2 Data

The main data source for this paper are the African census data obtained from IPUMS international (Minnesota Population Center, 2020).[4] The censuses are selected based on two criteria: (1) the census contains information on how long individuals have been living in their current region of residence and (2) the census contains information on the previous region of residence or the region of birth. The two criteria together ensure that the censuses contain sufficient information to construct a migration history, with the first criterion being by far the most stringent. In total 26 censuses across 14 countries satisfy the criteria, which are listed in table 1.4 in the appendix.[5] Figure 1.2 shows a map of the countries included in the sample, as well as the spatial level on which the current and previous region of residence are recorded in the census. The censuses jointly contain information on 57 million individuals and form a representative sample of the 14 countries with a population of 550 million individuals in 2019.[6]

The sample of countries selected based on the two criteria appears fairly representative for the group of African countries with available census data as a whole. The urban-rural gap in primary school completion is 25 percentage points in the 14 countries used in the analyses (25 pp in all countries with census data), 14 percentage points in the case of school attendance (17 pp for all countries) and 19 percentage points in the case of literacy (21 pp for all countries). Furthermore, the population share living in rural areas is also quite similar (63% in the analyses countries versus 65% in all countries with census data). Hence, the group of countries included for the estimation appears reasonably representative for the African countries with available census data.

1.3.3 Sample selection and urbanization measure

For the baseline estimations, I restrict the sample to children who are between the ages of 13 and 18 at the time of the census. Less than 5% of this group reports being the household head or partner to the household head, which somewhat limits concerns regarding the co-habitation bias.[7] At the same time, the majority of children have achieved their final primary school and literacy sta-

[4]While it is possible to extend the design to include countries outside of the African continent, very few non-African countries satisfy the two selection criteria discussed below and those that do are geographically scattered across the globe.

[5]The 1984 Ethiopian census and the 1996 Guinean census only record the time in the current residence in intervals. In these cases, I assume that the household moved in the middle of the interval. Furthermore, South Africa only contains information for individuals who moved in the ten years before the census.

[6]The census extracts available from IPUMS international typically consist of a 10% random draw of the full population census. See table 1.4 in the appendix for details.

[7]Appendix figure 1.15 shows the percentage of individuals who are household head or partner to the household head by age.

Figure 1.2. Selected sample of countries and regions

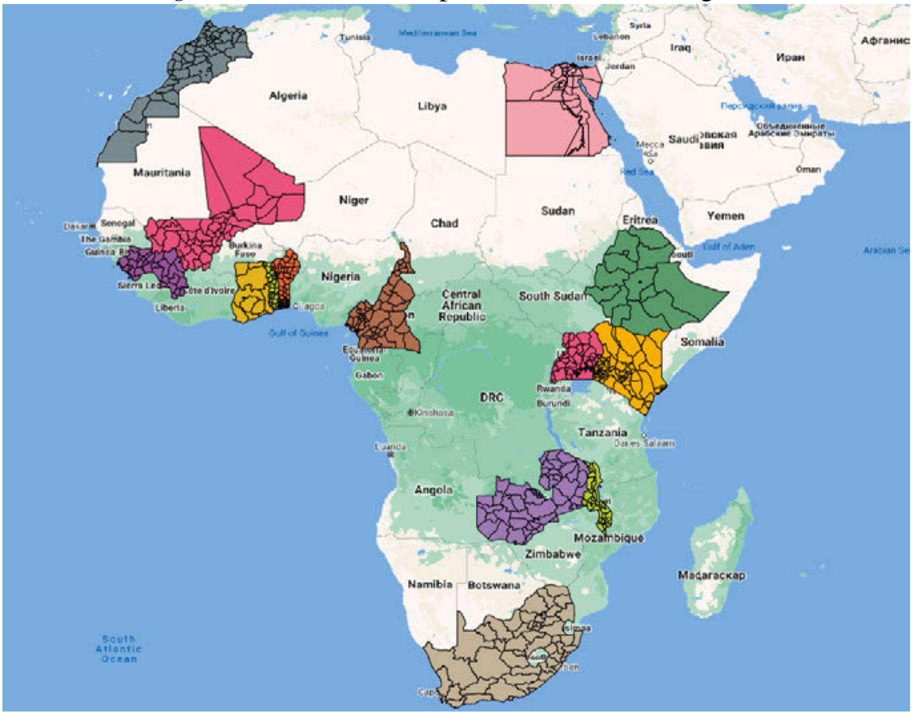

Note: Map displays the 14 countries with census data that satisfied the two selection criteria discussed in section 1.3.2. The regions denote the spatial units on which the information on current and previous region of residence is based, which typically corresponds to a province, district or arrondissement (see table 1.4 for more detail on the spatial units). As the boundaries of administrative regions often shift over time, the spatial units are based on the GEOLEV1 and GEOLEV2 geographic variables of IPUMS International database, which are consistent across time.

tus.[8] However, the age at which outcomes are measured does not drive the results as shown in the robustness section. The restrictions result in a sample of 416,509 individuals who are between the age of 13 and 18 at the time of the census, are not household head or partner to the household head, are citizens of their country of residence and whose previous region of residence differs from their current region of residence.[9] Table 1.1 shows the distribution of the sample by country. The movers are generally well spread across the 14 countries, with only Zambia contributing more than 15% of the observations.

[8] Appendix figure 1.16 shows primary school completion, school attendance and literacy rates by age.

[9] In addition, I exclude children who moved within a year prior to the census, children who are unrelated to the household head and children in households where the household head is below the age of 30.

Table 1.1. *Distribution of core sample of movers across countries*

	Number of observations	Percentage of observations
Benin	21,628	5.19
Cameroon	46,849	11.25
Egypt	56,692	13.61
Ethiopia	9,341	2.24
Ghana	27,185	6.53
Guinea	23,248	5.58
Kenya	49,314	11.84
Malawi	12,288	2.95
Mali	10,427	2.50
Morocco	16,859	4.05
South Africa	12,755	3.06
Togo	8,625	2.07
Uganda	55,467	13.32
Zambia	65,831	15.81
Total	416,509	100.00

Note: Note: Table displays the geographical spread of the core sample of 416,509 movers. See the accompanying text for details on the sample construction.

The final step is to decide how to define the degree of urbanization in both the current and previous region of residence and hence how to compute the $\Delta urbanization_{od}$ term in equation 1.1. The baseline results are based on the urban population data from the Africapolis database developed by the SWAC and OECD (2020), which has the benefit of employing a consistent definition of "urban" across all African countries. The Africapolis database contains information on historical populations and current boundaries for all African towns containing more than 10.000 inhabitants in 2015. For each region, I sum the total population living at the time of the census in the urban areas according to the Africapolis database and divide this by the total population of the region from the census.[10]

The distribution of $\Delta urbanization_{od}$ is shown in figure 1.3. On average, individuals move to regions that are 8.7 percentage points more urbanized compared to their previous region of residence, which is in line with the general migration towards urban areas in Africa. At the same time, there is also sub-

[10] The urban population data in Africapolis is not available yearly, but rather per decade for the years 1990, 2000, 2010 and 2015. I interpolate the annual population assuming an exponential growth function. A small percentage of urban districts are missing information on the urban population in 1990 and 2000. In these cases I extrapolate from the nearest available data point (1990, 2000 or 2010), using the average growth rate of urban regions in the same country. I use nightlight data to distribute the population across regions in case an urban agglomeration spans multiple administrative divisions. All capital city regions, which are typically small and only cover the capital itself, are coded as fully urban.

stantial movement from urban areas into rural areas as previously reported by Young (2013). To highlight that the results do not depend on the selected measure of urbanization, I also present the results when urbanization is based on population density, when using the urban definition of the individual censuses and when restricting attention to exposure to capital cities.

Figure 1.3. Distribution of $\Delta urbanization_{od}$

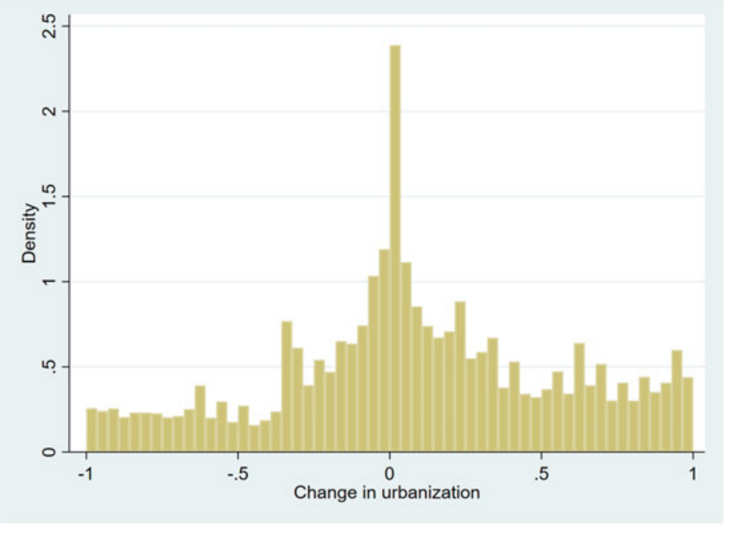

Note: Urbanization is based on the share of the individuals living in towns with more than 10,000 inhabitants. The spike in density at zero is due to households moving between regions classified as either fully rural or fully urban, which is the case in 4% of the migration events.

1.4 Results

1.4.1 Baseline results

Figure 1.4 displays the b_m coefficients obtained by estimating equation 1.1 for primary school completion. The top left panel shows the b_m estimates of the model without household fixed effects. Two things stand out from the figure. First of all, there is clear evidence of selection into cities, with children who move to cities at any age being more likely to have completed primary school. In particular for children who move after the age of 12, the estimates of b_m mainly reflect the spatial selection of households into cities. Second, the b_m estimates provide evidence that childhood urban exposure matters for educational outcomes. Moving to an area that is one percentage point more urbanized before the age of 6 is associated with a 0.20 percentage point (pp) increase in primary school completion, compared to a 0.12 pp increase for children who make the same move after the age of 12.

Under the assumption that the selection of households into more and less urban areas is indeed independent of the age of migration, this 0.08 pp difference can be interpreted as the causal impact of childhood urban exposure on educational outcomes. The magnitude of the urban exposure effect is significant, as it implies that spending childhood in a region that is fully urban instead of fully rural increases the probability of obtaining a primary school degree by 8 percentage points, from a mean of 63%. The urban exposure effects on primary school completion are mostly absent before the age of 6 and after the age of 12.

This interpretation depends on the validity of the Chetty and Hendren (2018)

Figure 1.4. Estimates of b_m for primary school completion at ages 13 - 18

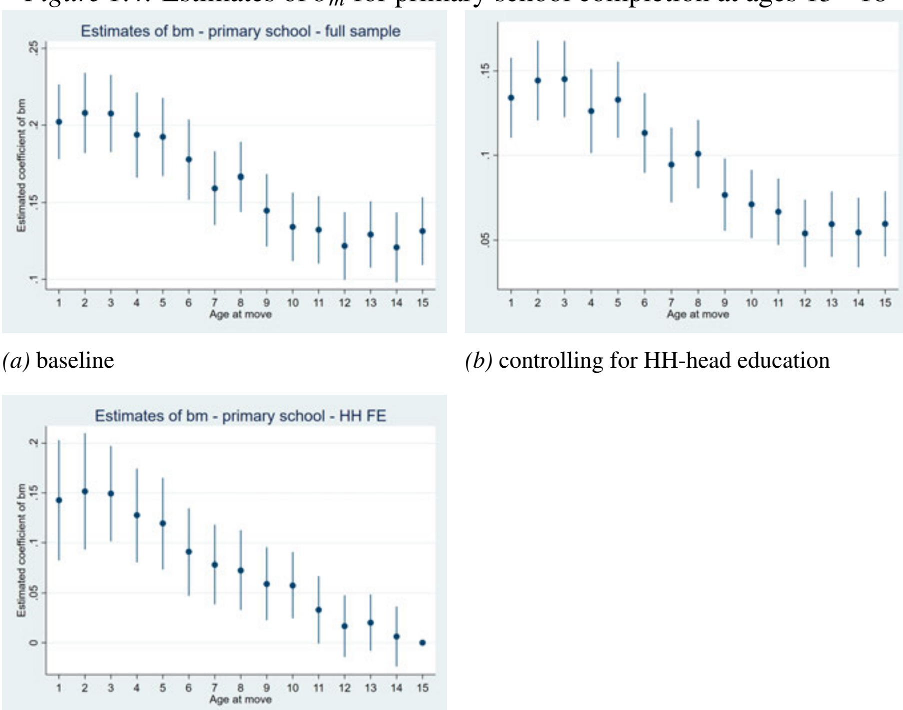

(a) baseline

(b) controlling for HH-head education

(c) including household fixed effects

Note: Figure displays the estimates b_m obtained by estimation equation 1.1 without household fixed effects (top left) and with household fixed effects included (bottom left) for the sample of movers described in section 1.3.3. The estimation in the top right panel shows the results when the education of the household head ("less than primary school", "completed primary school", "completed secondary school" and "completed university") and literacy status ("yes", "no") are included in the fixed effect term of equation 1.1. Primary school completion is defined as successful completion of the sixth grade of schooling. The estimate of b_m at age 15 is normalized to 0 in the household-fixed effect model.

identifying assumption in the context of Africa, which can be assessed in various ways. First of all, equation 1.1 allows for the inclusion of various household characteristics to account for selection into cities. The top right panel of figure 1.4 shows the b_m estimates when the educational attainment and literacy status of the household head are included in the fixed effect term of equation 1.1, thus creating a region of origin by census year by birth cohort by household head education and household head literacy fixed effect. Accounting for the educational attainment and literacy of the household head captures a substantial degree of the selection into cities, as the b_m estimates after the age of 12 drop from around 0.12 to around 0.05. However, the decline is visible through the entire age-at-move distribution with the implied urban exposure effect estimates virtually unchanged. Hence, controlling for the educational attainment status of the household head can account for a substantial part of the selection of households into cities, but not for the observed urban exposure effects.

This approach can be taken one step further with the inclusion of a household fixed effect in equation 1.1, in which case identification comes from comparing children within the same household who moved at different ages. The bottom left panel of figure 1.4 displays the b_m estimates in this case. Since there is no variation within households in $\Delta urbanization_{od}$, the estimate of b_{15} is normalized to zero. The urban exposure estimate is again similar in magnitude to the baseline estimate. To compare the estimates more formally, table 1.5 in the appendix displays the urban exposure estimates of the three specifications in figure 1.4 under two different parameterizations, showing that the inclusion of family characteristics and household-fixed effects indeed does not diminish the estimates.

The next question is whether childhood urban exposure also affects educational investment beyond primary school. To investigate this, figure 1.5 shows

Figure 1.5. Estimates of b_m for school attendance at ages 13 - 18

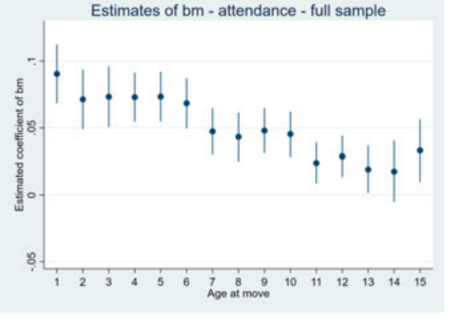

(a) baseline

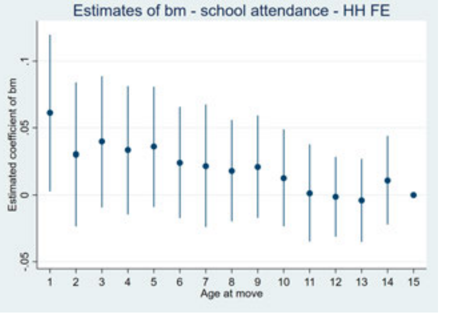

(b) including household fixed effects

Note: Figure displays the estimates b_m obtained by estimation equation 1.1 without (left) and with (right) household fixed effects on school attendance at the time of the census for the sample of movers described in section 1.3.3. The estimate of b_m at age 15 is normalized to 0 in the household-fixed effect model.

the estimates of b_m for school attendance at the time of the census. The pattern is similar to that for primary school completion. Children who move to an area one percentage point more urbanized before the age of 6 are around 0.07 percentage points more likely to attend school at ages 13 - 18, compared to a 0.03 percentage point increase for children who make the same move after age 12. The b_m estimates are similar when including household fixed effects as shown in the right panel. Table 1.6 in the appendix shows the urban exposure estimates of both models under two different parameterizations, with no significant differences between the estimates. Hence, spending childhood in a fully urban region rather than a fully rural region increases the probability of remaining in school at the ages 13 - 18 by around 4 percentage points.
Finally, one concern is that the increase in primary school completion and school attendance may not necessarily translate into better learning outcomes, given the large literature on school quality in Sub-Saharan Africa (see Bold et al. (2017) for an overview). To investigate whether the urban exposure translates into better learning outcomes, figure 1.6 shows the b_m estimates when using literacy as outcome variable. The pattern is similar to the previous two figures, with urban exposure effects clearly visible and concentrated between the ages of 6 and 12. Moving to a region one percentage point more urban before the age of 6 rather than after the age of 12 is associated with a 0.04 percentage point increase in literacy. Table 1.7 in the appendix shows the urban exposure effects under two different parameterizations, with the urban exposure estimates significant at the 0.1% level in all cases and no significant differences in the estimates between the models.

Figure 1.6. Estimates of b_m for literacy at ages 13 - 18

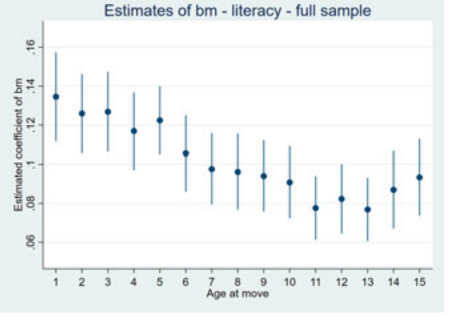

(a) baseline

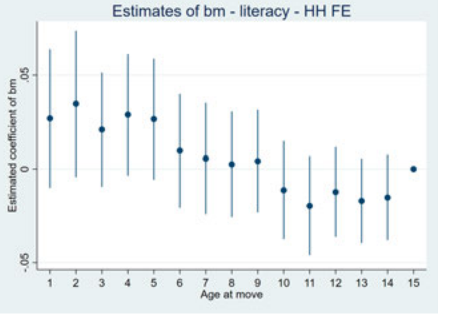

(b) including household fixed effects

Note: Figure displays the estimates b_m obtained by estimation equation 1.1 without (left) and with (right) household fixed effects on literacy for the sample of movers as described in section 1.3.3. The estimate of b_m at age 15 is normalized to 0 in the household-fixed effect model.

1.4.2 Robustness and sensitivity analysis

Although the childhood urban exposure estimates are robust to the inclusion of household-fixed effects, various other potential threats to identification exist. This section discusses four potential concerns: (1) the sensitivity to the definition of urbanization, (2) the sensitivity to the sample of countries and age at which outcomes are measured, (3) the potential impact of time-varying shocks to household conditions and (4) potential selection within the geographical regions.

Urbanization measure

The baseline estimates are based on the urban population data from the Africapolis database, but various other definitions can be used. I verify the results using three alternative measures of urbanization, (1) the urban population data from the censuses[11] (2) population density and (3) exposure to the capital cities. The results under the alternative definitions are shown in appendix 1.D.1. The results are similar under the three alternative definitions, highlighting that the way in which regions are classified does not affect the results.

Country selection and age at census

The next question is whether the results are driven by just a handful of countries or whether the results are broadly representative for the experiences of most African countries. I impose a linear structure on the estimates of b_m as shown in equation 1.3 to compare the estimates across the various countries and to test for their statistical significance. Since the urban exposure effects appear to be largely absent after the age of 12, the linear functional form assumption is restricted to ages 1 - 12.[12]

$$\text{Educational outcome}_i = \alpha_{ocs} + \sum_{m=1}^{12} I(m_i = m)\phi_{mc} + \delta \Delta urbanization_{od} + \gamma \sum_{m=1}^{12} (\frac{12-m}{12}) I(m_i = m) \Delta urbanization_{od} + \varepsilon_i \quad (1.3)$$

The first term again captures the quality of the region of origin with a region of origin by census year by birth cohort fixed effect. The second term captures the disruption effect of moving, which is allowed to vary by the age at

[11] This definition varies greatly across the countries and censuses. For this reason, the baseline results use the Africapolis database, which provides data on the urban population under a definition that is stable across space and time.

[12] An alternative is to impose a piecewise linear function on the estimates of b_m by interacting $\Delta urbanization$ with moving in various age brackets (for instances ages 1 - 5, 6 - 10 and 11 - 15). The results in that case are very similar, with significant urban exposure effects for children who move at ages 1 - 5 rather than at ages 11 - 15 in 9 out of the 14 countries in the case of primary school completion, 8 out of the 13 countries in the case of school attendance and 5 out of the 12 countries in the case of literacy. Furthermore, again only 2 out of the 39 country × education outcomes point estimates are negative.

the move, country and census year. The third term captures the selection by households into more or less urban places. The fourth term captures the urban exposure effect, which consists of the fraction of childhood spent in the new environment before the age of 12, interacted with the change in environment $\Delta urbanization_{od}$. Thus, γ captures the effect of spending childhood before age 12 in a region that is one percentage point more urbanized, identified of children who move between regions at different ages.

The coefficients obtained by estimating equation 1.3 for the individual countries are shown in table 1.8 in the appendix. The urban exposure effects captured by γ are positive and statistically significant in 13 out of the 14 countries, with the exception of South Africa. The point estimates for school attendance and literacy are somewhat smaller as shown in appendix tables 1.9 and 1.10 in line with the baseline estimates, but are nonetheless positive and statistically significant in 7 out of the 13 countries in the case of school attendance and 6 out of the 12 countries in case of literacy.[13] Across the 39 country $\times$ educational outcome estimates in tables 1.8 - 1.10, the point estimate is negative in only two instances.[14] Hence, the results are not driven by any particular region or country, but are instead visible in the majority of African countries across the various educational outcomes.

Furthermore, the results do not depend on the age at which outcomes are measured. The baseline results are estimated jointly on the children between ages of 13 to 18 at the time of the census. Figure 1.23 in the appendix shows the urban exposure estimates for three narrower age groups: 13 - 14 year olds, 15 - 16 year olds and 17 - 18 year olds. The urban exposure effects are visible across all groups, indicating that the estimates are not driven by the age at which educational outcomes are measured. Table 1.11 in the appendix shows the urban exposure effects by age group for the three outcomes, which are positive and statistically significant at the 0.1% level in all cases.

Time-varying shocks to household conditions

Third, one might be concerned that regional shocks in the region of origin may both drive migration into cities as well as increase educational attainment. I analyze two potential shocks that have been shown to affect educational outcomes by prior literature, namely agricultural income shocks stemming from rainfall[15] and civil violence shocks. Tables 1.12 and 1.13 in the appendix show the estimates when removing observations where a substantial rainfall

[13] No data is available on school attendance in the Moroccan census and on literacy in the Kenyan and South African census.

[14] The two cases are south Africa in the case of school attendance and Malawi in the case of literacy. The negative coefficient for school attendance in South Africa is in fact consistent with the cross-sectional evidence in appendix figure 1.10, with South Africa being the only African country where school attendance is lower in urban areas than in rural areas.

[15] Maccini and Yang (2009) for instance show that rainfall shocks in early life significantly affect educational attainment.

shock or civil violence took place in the region of origin around the time of migration. The results remain virtually unchanged in both cases, suggesting that rainfall shocks and violence in the region of origin are not driving the results.[16]

A different concern is that time-varying shocks to household conditions may cause households to relocate to more or less urban areas, in which case the estimates may confound the effect of urban exposure with the effect of the initial shock. One way to test for the influence of time-varying household shocks is to focus on households who migrated from their region of origin during years of exceptionally large migration outflows as suggested by Alesina et al. (2021). During years of large migration outflows, the timing of migration for the individual household is more likely to be driven by exogenous factors and less likely to be driven by idiosyncratic shocks.

Figure 1.7 shows the urban exposure estimates for primary school completion when restricting the sample based on the out-migration in the region of origin. The first estimate is based on all households who moved in the 15 years prior to the census and thus closely corresponds to the baseline estimates, whereas the right-most estimate is based solely on the year during which the migration outflow in the region of origin was highest compared to the predicted migration outflow based on a linear time-trend model. The estimates remain stable as we successively focus on the years with higher migration outflows in the region of the origin, suggesting that also idiosyncratic shocks to household conditions are unlikely to drive the results. The results are similar when analyzing school attendance and literacy, as shown in appendix figures 1.24 and 1.25.

Selection within regions and repeat migration

Finally, one might be concerned regarding the geographic detail of the current and previous region of residence. Figure 1.2 shows that the level of detail on the region of residence in the census can be quite coarse for some countries such as Ethiopia and Ghana, which means that regions may contain very heterogeneous areas. In such cases families might move to very different areas within the broader regions, depending on the age of the children at the time of the move.

While the regions displayed in figure 1.2 are the most detailed level on which both the current and previous region of residence is known, substantially more detailed information is available for some countries regarding the current region of residence. Table 1.14 in the appendix shows the effect of using the more detailed geographical information on the current region of residence for

[16] Alternatively, one might argue that moves in response to violence in fact represent moves which are likely exogenous with respect to the age of children. The urban exposure effects are also clearly visible and statistically significant when the sample is restricted to households who moved from an area where a violent event took place between 3 years prior to the move and 2 years after the move.

Figure 1.7. Migration outflow shocks and urban exposure estimates for primary school completion

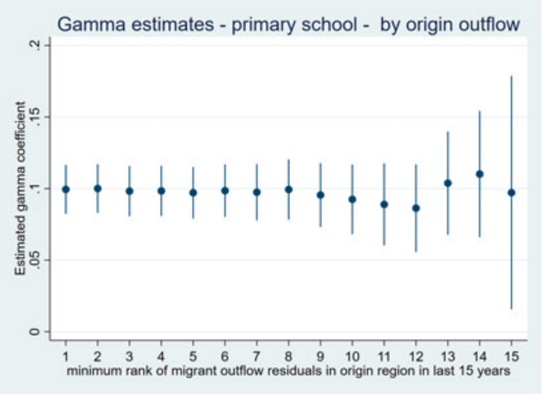

Note: The figure shows the estimates of γ obtained by estimating equation 1.3 on primary school completion for a successively more selective sample. The first estimate is based on all migration events in the 15 years prior to the census, whereas the final estimate is only based on the year with the strongest relative deviation in the migration outflow in the region of origin compared to what would be expected based on a linearly predictive model. The 1996 Guinean census and the 1984 Ethiopian census are excluded as they only record the years since migration in intervals, and the South African census is excluded as it only records migration events in the 10 years prior to the census.

the countries which have a high level of geographic aggregation in the baseline estimates. The urban exposure estimates do not change significantly when using the more detailed geographical information, suggesting that migration to different areas within the current region of residence is not driving the results. Similarly, one might worry about selection within the previous region of residence. Such selection would for instance be the case if children in the urban parts of the region of origin are likely to migrate at younger ages compared to children in the rural areas. To investigate this possibility, table 1.15 in the appendix uses the information contained in eight censuses regarding the rural/urban status of the previous place of residence. Controlling for the urban/rural status of the previous residence also does not affect the urban exposure estimates, suggesting that also differences in selection into migration within the region of origin are not driving the results.
Finally, some children may migrate across regional boundaries multiple times during childhood, in which case the previous place of residence does not accurately describe the area in which they spend their childhood. Thirteen censuses record information on both the previous region of residence and the region of birth, which allows me to test whether individuals who migrated multiple times during childhood affect the results. On average, the place of birth and previous place of residence differs for around 20% of the children. However, excluding the group of children who migrated multiple times during childhood

does not affect the results as shown in appendix table 1.16 and if anything somewhat increases the urban exposure estimates.

1.5 Heterogeneity

One interesting question is whether all children benefit equally from spending childhood in an urban environment, or whether this differs between households. In this section I analyze heterogeneity across three dimensions: the gender of the child, the education of the household head and whether households move to more or less urban areas. Furthermore, I analyze how the urban exposure estimates develop over time using repeated censuses for the same countries.

1.5.1 Gender

Chetty and Hendren (2018) and Deutscher (2020) find that childhood exposure to better neighborhoods matters more for boys than girls in the context of the US and Australia. To investigate this in the African context, I interact the delta term and gamma term in equation 1.3 with the gender of the child.[17] Table 1.2 shows the resulting urban exposure estimates. The exposure effects are larger for girls in all three cases, with the difference statistically significant at the 5% level in the case of school attendance. Hence, both genders benefit from spending childhood in an urban environment, but the effects are somewhat larger for girls compared to boys.

1.5.2 Parental education

A second question is whether childhood urban exposure substitutes or complements for household characteristics favorable towards education. Chetty and Hendren (2018) find that neighborhood quality matters more for children from higher income households, thus suggesting that neighborhood and family characteristics are complementary. Table 1.2 displays the estimates when the urban exposure term is interacted with a dummy variable indicating whether the household head has completed at least secondary school. The interaction term is insignificant across all three educational outcome measures, indicating that the effect of urban exposure on educational attainment does not depend on the parental education in the context of Africa.

[17] For all heterogeneity analyses presented here, the conclusions are the same if a piecewise linear structure is imposed on the estimates of b_m as discussed in appendix 1.C, rather than the linear specification of equation 1.3.

1.5.3 Asymmetric exposure effects

Third, I separate the sample between those households who choose to move to more urbanized areas and those who choose to move to less urbanized areas. If households efficiently sort across space, then the gains of households who move to urban areas should be larger than the losses of households who choose to leave urban areas. Table 1.2 shows that this is indeed the case. The urban exposure estimates are substantially larger for children who move to more urbanized areas compared to children who move to less urbanized areas across all three outcome measures, with the differences statistically significant at the 5% level in the case of primary school completion and school attendance. This stands in contrast to Chetty and Hendren (2018), who somewhat surprisingly find that exposure to better neighborhoods matters more for households who move to worse neighborhoods in the US, and Deutscher (2020), who finds no differences in Australia. The finding supports the suggestion of Young (2013) that African households select into urban and rural regions according to their comparative advantage. Nonetheless, urban exposure effects are still present and statically significant even for households who choose to move to less urban regions.

1.5.4 Development over time

Finally, an interesting question is how the magnitude of the urban exposure effects has developed over time. The censuses included in the estimating sample vary between the years 1984 and 2014, during which Africa developed significantly. In addition, specific attention was paid to improving access to primary education in rural areas as part of the millennium development goals. To investigate how the childhood urban exposure effects have changed over time, I use the fact that 8 out of the 14 countries have data available from two or more census waves. The average gap between the first and the second census wave is 13 years, thus providing a reasonable time frame over which to study the changes.

Table 1.3 shows the urban exposure estimates by census wave. The urban exposure effects based on the second census wave are smaller than the first wave across all the three outcomes measures, with the estimates significantly different at the 5% level. Thus, Africa's development during the 1990's and 2000's seems to have reduced the effect of childhood urban exposure on primary school completion, school attendance and literacy. Nonetheless, the urban exposure estimates based on the second census waves remain statistically significant at the 0.1% level and economically relevant. The results are very similar when focusing on the four countries with three available census waves.

Table 1.2. *Heterogeneity in urban exposure effects by gender, household head education and the change in urbanization status*

Outcome	Primary school completion			School attendance			Literacy		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
gamma	0.0961***	0.109***	0.0687***	0.0496***	0.0745***	0.0339*	0.0495***	0.0637***	0.0441**
	(0.00920)	(0.0104)	(0.0157)	(0.0104)	(0.00945)	(0.0139)	(0.00975)	(0.00919)	(0.0150)
gamma*female	0.0109			0.0297*			0.0195		
	(0.0126)			(0.0140)			(0.0148)		
gamma*education_hh_head		-0.00911			-0.0228			-0.00677	
		(0.0216)			(0.0133)			(0.0159)	
gamma*$\Delta urbanization > 0$			0.0614*			0.0580*			0.0299
			(0.0276)			(0.0254)			(0.0257)
N	332254	332254	332254	267688	267688	267688	283182	283182	283182
R^2	0.225	0.225	0.224	0.153	0.153	0.153	0.143	0.143	0.143

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The table shows the estimates of γ when estimating equation 1.3 with *gamma* and Δ*urbanization* interacted with respectively gender, household head education and the change in urbanization. The HH-head education variable is a dummy indicating whether the household head completed secondary school. The outcome variables are respectively primary school completion (columns 1 - 3), school attendance at ages 13 - 18 (columns 4 - 6) and literacy (columns 7 - 9). The number of observations varies slightly between the three outcomes as school attendance is not observed in Morocco and literacy is not observed in Kenya and South Africa. Standard errors are double clustered at the current and previous region of residence.

Table 1.3. *Urban exposure estimates by census wave*

	(1)	(2)	(3)	(4)	(5)	(6)
	Primary School		School attendance		Literacy	
delta	0.0951***	0.152***	0.0265	0.0294*	0.0626**	0.0980***
	(0.0232)	(0.0176)	(0.0196)	(0.0118)	(0.0229)	(0.0118)
gamma	0.126***	0.0628***	0.115***	0.0529***	0.0952***	0.0562***
	(0.0155)	(0.0114)	(0.0191)	(0.0119)	(0.0177)	(0.0121)
Census wave	First	Second	First	Second	First	Second
N	79230	108170	40977	75635	70598	81029
R^2	0.204	0.219	0.170	0.185	0.156	0.140
Mean census year	1992	2005	1992	2005	1992	2005

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: Note: The table shows the coefficients of γ and δ obtained by estimating equation 1.3 for the eight countries with at least 2 censuses (see table 1.4 for a list of censuses). Standard errors are double clustered at the current and previous region of residence.

1.6 Mechanisms

Finally, it is important to understand what mechanisms can explain the findings of this paper. Moving to a city could increase educational investment due to higher returns to education, improved access to educational institutes or by increasing household income. At the same time, moving to a city may also change the opportunity costs of education and affect family formation decisions. Although data availability and the lack of random variation make it difficult to provide definitive answers, this section provides suggestive evidence which mechanisms are most likely to be at play.

1.6.1 Family formation decisions

One mechanism which could explain the effect of urban exposure on educational outcomes is changes in family formation decisions. Cross-sectional evidence reveals that the percentage of teenage girls who have a child is lower in urban areas across virtually all African countries, as is the percentage of children that are married or report being in a consensual union (see figures 1.12 and 1.13 in the appendix). If growing up in an urban area decreases the probability of teenage marriage or childbirth then this might explain the increase in educational investment in urban regions, particularly among girls.[18]

[18] Households do appear to be forward-looking with respect to educational investment in children. Jayachandran and Lleras-Muney (2009) for instance find an increase in educational attainment among girls in Sri Lanka in response to a reduction in maternal mortality. Hence, if families anticipate that children are more likely to drop out of education due to childbirth or marriage, then this could restrain educational investment in rural areas throughout childhood.

However, there is little evidence that the length of urban exposure matters for childbirth or marriage as shown in figure 1.8. Children who move to a more urbanized area at any age are significantly less likely to have children or to be married, but there is no evidence that the age at which one moves to a city matters for the reduction in childbearing or marriage rates. This is confirmed when estimating equation 1.3 on marriage and childbirth rates, with insignificant estimates of gamma in both cases. Hence, reductions or delays in marriage and childbirth cannot explain the urban exposure effects obtained in section 1.4.

Figure 1.8. Estimates of b_m for having a child (left panel) or being married (right panel) at ages 16/17

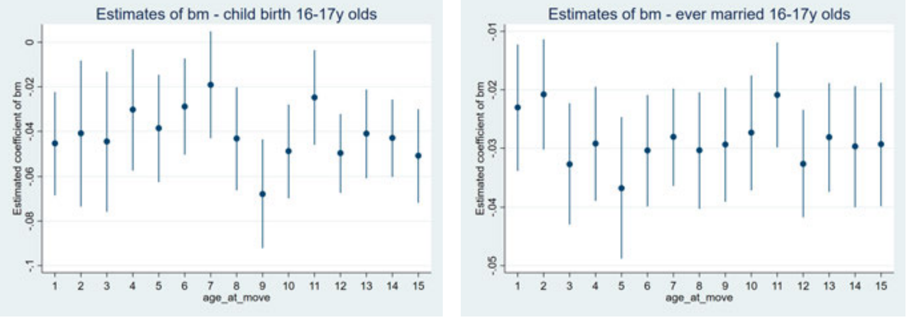

(a) Estimates of b_m for having children

(b) Estimates of b_m for being married

Note: Figure displays the estimates b_m obtained by estimation equation 1.1 for having a child (left panel) or being married (right panel) for children aged 16/17 at the time of the census. Since parenthood is only recorded for females in countries, the left panel is based exclusively on girls.

1.6.2 Labor force participation

A different explanation is that childhood urban exposure may affect labor force participation. The cross-sectional evidence displayed in appendix figure 1.14 shows that children in urban areas are less likely to engage in home production or to work outside the home compared to children in rural areas. The probability of labor force participation is strongly related to the age at which children move into cities as shown in figure 1.9. Children who move to an area one percentage point more urbanized before the age of 6 are around 0.05 percentage points less likely to be engaged in home production or working outside the home at ages 13 - 18 compared to children who move after the age of 12. The decrease in labor force participation is similar in size to the increase in school attendance found in section 1.4.1. Hence, the main margin of adjustment is an increase in educational attainment at the expense of home production and work outside of the home.

There are various factors that could explain the higher school attendance

Figure 1.9. Estimates of b_m for work inside or outside the home at ages 13 - 18

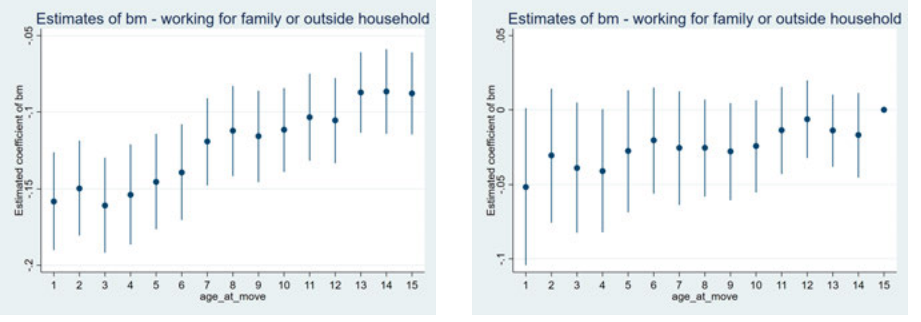

(a) baseline

(b) including household fixed effects

Note: Figure displays the estimates b_m obtained by estimation equation 1.1 without (left) and with (right) household fixed effects on employment status for the sample of movers described in section 1.3.3.

and lower labor participation among children in urban areas, including (1) household income effects, (2) better access to public services and schools, (3) changes in the opportunity costs of education and (4) higher returns to education.

Due to the lack of household income data for most African countries, either from census data or other data sources, it is not possible to directly estimate the regional returns to education, opportunity costs of education or household income effects. However, we can use the information on household conditions from the censuses to make some progress. Most censuses contain information on household access to various utilities, such as piped water and electricity. Furthermore, the censuses also contain information on housing characteristics, such as floor type and whether the walls consist of organic or durable materials. Such information can provide information on whether the increase in educational attainment can be fully explained by the better provision of public utilities in cities and household income effects, or whether children in poor households also benefit from urban exposure, which would suggest that other factors may also play a role.

The results of the interaction of the urban exposure effect with the various measures of utility access and household housing wealth are shown in tables 1.1 and 1.2 in the appendix. The urban exposure effects are very similar for households with and without access to utilities, with the exception of electricity access in the case of primary school completion when the effects are slightly larger for connected households. Nonetheless, the urban exposure effects are positive and statistically significant across all specifications among the comparatively deprived households. The same results hold when splitting households by the type of floor material (dirt versus non-dirt) and by building material of the walls (organic versus brick/stone). As the positive effects

of childhood urban exposure on educational outcomes are also visible among less wealthy citizens and those with poor access to public utilities, it seems somewhat unlikely that household income effects fully explain the findings.
Since data on returns to education, the opportunity costs of education and access to education are unavailable in virtually all countries, it is difficult to assess their importance.[19] One exception is the case of Guinea, where the census records the reason for children to leave school in case they no longer attend. Their responses are summarized in table 1.E.3 in the appendix. Children in rural areas are far more likely to report leaving school due to a lack of access to schooling compared to children in urban areas (15.7% vs 2.0%). Similarly, children in rural areas are more likely to report leaving school due to the need to assist their family with household tasks (17.6% vs 5.0%), with the difference particularly strong when parents are farmers. This suggests that the opportunity costs of education are likely to be higher for households in rural communities. At the same time, there are no substantial differences in the share of students who report dropping out of school due to illness, not wanting to study or academic issues, suggesting that a difference in demand for schooling is unlikely to explain the results. However, the lack of exogenous variation in the various factors and the absence of data makes it difficult to draw any definitive conclusions regarding the importance of the potential mechanisms, which forms an interesting avenue for future research.

1.7 Conclusion

This paper investigates the effect of childhood urban exposure on educational attainment in Africa. Using an age-at-move identification strategy, I find that childhood urban residency increases educational attainment across three educational measures. The childhood urban exposure effects are visible both within and between households and do not appear to be driven by the urbanization measure, time-varying shocks to household conditions, or mismeasurement in the region of residence or origin. Instead, the effects are visible in the majority of African countries and in all subgroups, suggesting that childhood exposure to urban environments increases educational attainment across a variety of settings.
On the individual level, the findings imply that urban migration in Africa improves the opportunities available to children. Despite a large literature highlighting the risks of urban migration in developing countries due to the pres-

[19] In particular the lack of household income data for most African countries makes it difficult to calculate the returns to education and opportunity costs of education on a regional level. One exception is the LSMS-ISA dataset. However, only two of the fourteen countries included in this paper are also included in the most recent version of the LSMS-ISA data. Furthermore, the main focus of the LSMS-ISA data is on income in agricultural regions, which means that information on incomes in urban areas is very limited.

ence of large urban slums without basic public service provision (Marx, Stoker and Suri, 2013) and uncertain property rights (Field, 2007), childhood urban residency on average seems to improve the opportunities available to children. Quantitatively, the effect of spending childhood in a fully urbanized region rather than a fully rural region on educational outcomes is similar to the effect of offering deworming drugs to primary school children (Miguel and Kremer, 2004) or the effect of the school construction program analyzed by Duflo (2001).[20] Since educational attainment has been shown to causally affect a host of other life outcomes, such as health outcomes (Silles, 2009) life satisfaction (Oreopoulos, 2007) and marriage quality (Kaufmann, Messner and Solis, 2013), the increase in educational attainment is likely to also improve opportunities outside of the labor market.

On the national level, the results suggest that urbanization in Africa can contribute to economic development by increasing the stock of human capital. Earlier studies on the role of cities in economic development have mainly focused on structural transformation, which appears less prevalent in the African context compared to other contexts (Gollin, Jedwab and Vollrath, 2016). This paper instead suggests a novel and different way through which urbanization can affect economic development, namely by increasing the human capital stock of the population. This channel is particularly relevant as educational attainment is the most robust predictor of economic development across a wide variety of settings, as shown by Gennaioli et al. (2013).

Finally, an important question for future research is which mechanisms explain the increase in educational attainment. The fact that childhood urban exposure raises educational attainment among comparatively deprived households without access to utilities or durable housing suggests that household income effects may not be the whole story. Evidence from the Guinean census highlights that access to schools and the high opportunity costs of education in rural settings are likely to also play a role in explaining the higher educational attainment in cities. A detailed decomposition of the various potential mechanisms within a causal framework forms an interesting avenue for future research.

Acknowledgements

The author wishes to acknowledge the statistical offices that provided the underlying data making this research possible: National Institute for Statistics and Economic Analysis, Benin; Central Statistics Office, Botswana; National Institute of Statistics and Demography, Burkina Faso; Central Bureau of Cen-

[20]Duflo (2001) finds that the increase of 2 primary schools per 1000 children in Indonesia led to a gain of 0.24 years in years of completed schooling. When analyzing the effect of urban exposure on completed years of schooling, I find that growing up in a fully urban rather than a fully rural area leads to an 0.3 to 0.6 year increase in years of schooling.

sus and Population Studies, Cameroon; Central Agency for Public Mobilization and Statistics, Egypt; Central Statistical Agency, Ethiopia; Ghana Statistical Services, Ghana; National Statistics Directorate, Guinea; National Bureau of Statistics, Kenya; Bureau of Statistics, Lesotho; Institute of Statistics and Geo-Information Systems, Liberia; National Statistical Office, Malawi; National Directorate of Statistics and Informatics, Mali; Statistics Mauritius, Mauritius; Department of Statistics, Morocco; National Institute of Statistics, Mozambique; National Institute of Statistics, Rwanda; National Agency of Statistics and Demography, Senegal; Statistics Sierra Leone, Sierra Leone; Statistics South Africa, South Africa; National Bureau of Statistics, South Sudan; Central Bureau of Statistics, Sudan; Bureau of Statistics, Tanzania; National Institute of Statistics, Togo; Bureau of Statistics, Uganda; Central Statistics Office, Zambia; and National Statistics Agency, Zimbabwe.

References

Akram, Agha Ali, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. 2017. "Effects of emigration on rural labor markets." National Bureau of Economic Research.

Alesina, Alberto, Sebastian Hohmann, Stelios Michalopoulos, and Elias Papaioannou. 2021. "Intergenerational mobility in Africa." *Econometrica*, 89(1): 1–35.

Becker, Sascha, Irena Grosfeld, Pauline Grosjean, Nico Voigtländer, and Ekaterina Zhuravskaya. 2020. "Forced Migration and Human Capital: Evidence from Post-WWII Population Transfers." *American Economic Review*, 110(5): 1430–1463.

Beegle, Kathleen, Joachim De Weerdt, and Stefan Dercon. 2011. "Migration and economic mobility in Tanzania: Evidence from a tracking survey." *Review of Economics and Statistics*, 93(3): 1010–1033.

Bold, Tessa, Deon Filmer, Gayle Martin, Ezequiel Molina, Brian Stacy, Christophe Rockmore, Jakob Svensson, and Waly Wane. 2017. "Enrollment without learning: Teacher effort, knowledge, and skill in primary schools in Africa." *Journal of Economic Perspectives*, 31(4): 185–204.

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. 2014. "Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh." *Econometrica*, 82(5): 1671–1748.

Chetty, Raj, and Nathaniel Hendren. 2018. "The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects." *The Quarterly Journal of Economics*, 133(3): 1107–1162.

Chetty, Raj, Nathaniel Hendren, and Lawrence F Katz. 2016. "The effects of exposure to better neighborhoods on children: New evidence from the Moving to Opportunity experiment." *American Economic Review*, 106(4): 855–902.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon. 2008. "Spatial wage disparities: Sorting matters!" *Journal of urban economics*, 63(2): 723–742.

De La Roca, Jorge, and Diego Puga. 2017. "Learning by working in big cities." *The Review of Economic Studies*, 84(1): 106–142.

Deutscher, Nathan. 2020. "Place, Peers, and the Teenage Years: Long-Run Neighborhood Effects in Australia." *American Economic Journal: Applied Economics*.

Duflo, Esther. 2001. "Schooling and labor market consequences of school construction in Indonesia: Evidence from an unusual policy experiment." *American economic review*, 91(4): 795–813.

Field, Erica. 2007. "Entitled to work: Urban property rights and labor supply in Peru." *The Quarterly Journal of Economics*, 122(4): 1561–1602.

Gennaioli, Nicola, Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer. 2013. "Human capital and regional development." *The Quarterly journal of economics*, 128(1): 105–164.

Glaeser, Edward L. 1999. "Learning in cities." *Journal of urban Economics*, 46(2): 254–277.

Glaeser, Edward L, and David C Mare. 2001. "Cities and skills." *Journal of labor economics*, 19(2): 316–342.

Gollin, Douglas, Remi Jedwab, and Dietrich Vollrath. 2016. "Urbanization with and without industrialization." *Journal of Economic Growth*, 21(1): 35–70.

Gould, Eric D. 2007. "Cities, workers, and wages: A structural analysis of the urban wage premium." *The Review of Economic Studies*, 74(2): 477–506.

Hamory, Joan, Marieke Kleemans, Nicholas Y Li, and Edward Miguel. 2020. "Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata." *Journal of the European Economic Association*.

Harris, John R, and Michael P Todaro. 1970. "Migration, unemployment and development: a two-sector analysis." *The American economic review*, 60(1): 126–142.

Henderson, J Vernon, and Matthew A Turner. 2020. "Urbanization in the developing world: too early or too slow?" *Journal of Economic Perspectives*, 34(3): 150–173.

Henderson, J Vernon, and Sebastian Kriticos. 2018. "The development of the african system of cities." *Annual Review of Economics*, 10: 287–314.

Jayachandran, Seema, and Adriana Lleras-Muney. 2009. "Life expectancy and human capital investments: Evidence from maternal mortality declines." *The Quarterly Journal of Economics*, 124(1): 349–397.

Kaufmann, Katja Maria, Matthias Messner, and Alex Solis. 2013. "Returns to elite higher education in the marriage market: Evidence from Chile." *Available at SSRN 2313369*.

Lagakos, David, Samuel Marshall, Ahmed Mushfiq Mobarak, Corey Vernot, and Michael E Waugh. 2020. "Migration costs and observational returns to migration in the developing world." *Journal of Monetary Economics*, 113: 138–154.

Lall, Somik Vinay, J Vernon Henderson, and Anthony J Venables. 2017. *Africa's cities: Opening doors to the world.* The World Bank.

Lewis, William Arthur. 1954. "Economic development with unlimited supplies of labour." *The Manchester School*, 22(2): 139–191.

Maccini, Sharon, and Dean Yang. 2009. "Under the weather: Health, schooling, and economic consequences of early-life rainfall." *American Economic Review*, 99(3): 1006–1026.

Marx, Benjamin, Thomas Stoker, and Tavneet Suri. 2013. "The economics of slums in the developing world." *Journal of Economic perspectives*, 27(4): 187–210.

Miguel, Edward, and Michael Kremer. 2004. "Worms: identifying impacts on education and health in the presence of treatment externalities." *Econometrica*, 72(1): 159–217.

Minnesota Population Center. 2020. "Integrated Public Use Microdata Series, International: Version 7.3 [dataset]. Minneapolis, MN: IPUMS, https://doi.org/10.18128/D020.V7.2."

Nakamura, Emi, Jósef Sigurdsson, and Jón Steinsson. 2021. "The Gift of Moving: Intergenerational Consequences of a Mobility Shock." *The Review of*

Economic Studies, rdab062.
Oreopoulos, Philip. 2007. “Do dropouts drop out too soon? Wealth, health and happiness from compulsory schooling.” *Journal of public Economics*, 91(11-12): 2213–2229.
Shah, Manisha, and Bryce Millett Steinberg. 2017. “Drought of opportunities: Contemporaneous and long-term impacts of rainfall shocks on human capital.” *Journal of Political Economy*, 125(2): 527–561.
Silles, Mary A. 2009. “The causal effect of education on health: Evidence from the United Kingdom.” *Economics of Education review*, 28(1): 122–128.
SWAC, and OECD. 2020. “Africapolis (database).”
UN DESA. 2016. *The Millennium Development Goals Report 2015.* UN.
Young, Alwyn. 2013. “Inequality, the urban-rural gap, and migration.” *The Quarterly Journal of Economics*, 128(4): 1727–1785.

Appendices

1.A Descriptive statistics on urban-rural differences

Figure 1.10. Urban-Rural differences in school attendance at ages 15 - 16

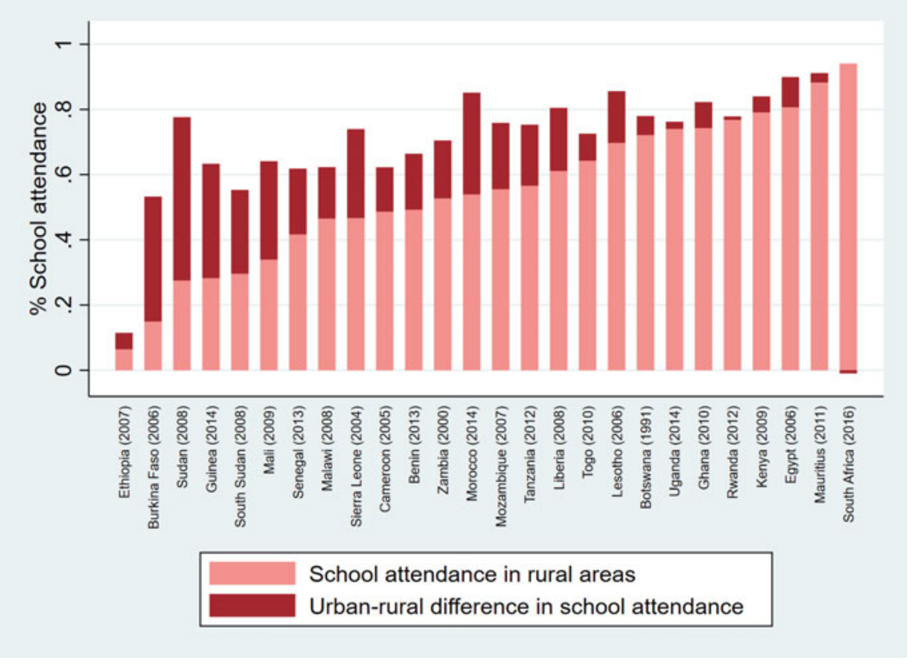

Note: Figure based on the most recently available census available from IPUMS international for all African countries. The sample includes all children between the ages 15 - 16. The definition of urban and rural is based on the definition of the statistical office in the relevant country. Zimbabwe does not gather information on school attendance and is thus excluded.

Figure 1.11. Urban-Rural differences in literacy at ages 15 - 16

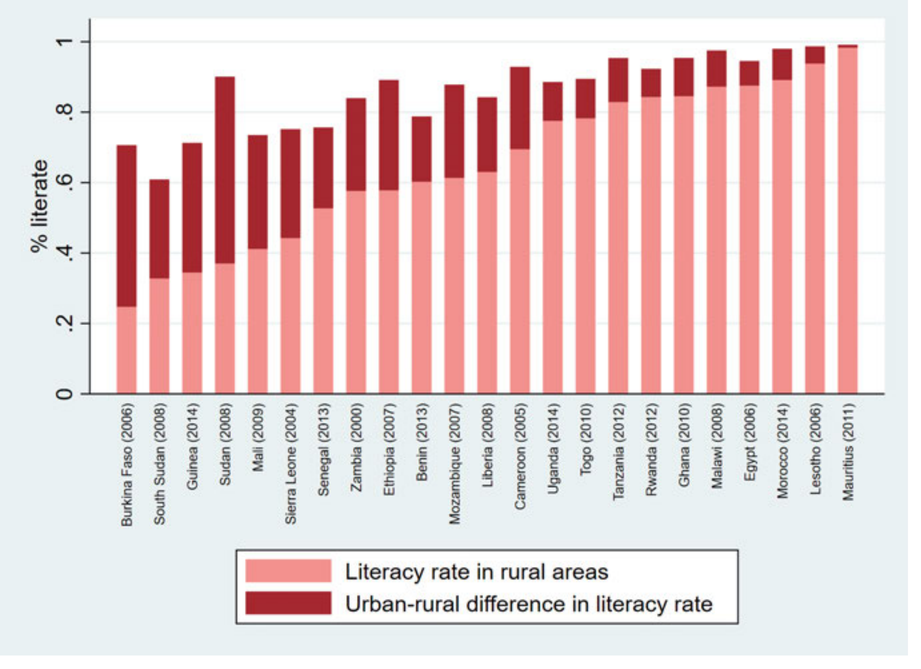

Note: Figure based on the most recently available census available from IPUMS international for all African countries. The sample includes all children between the ages 15 - 16. The definition of urban and rural is based on the definition of the statistical office in the relevant country. Botswana, Kenya, South Africa and Zimbabwe do not gather information on literacy and are thus excluded.

Figure 1.12. Percentage of girls aged 16 - 17 who had a child by urban-rural status

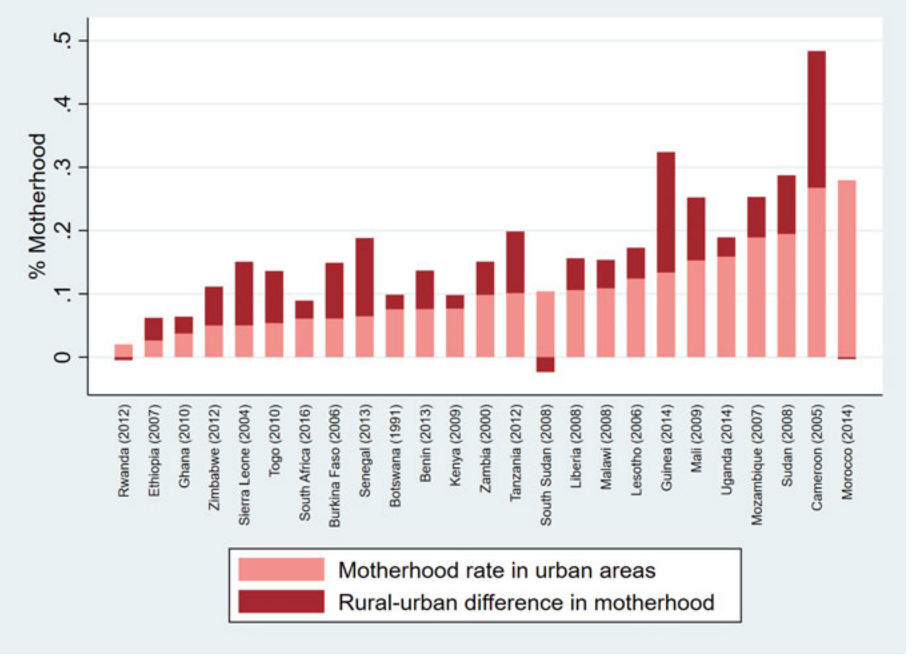

Note: Figure based on the most recently available census available from IPUMS international for all African countries. The sample includes all girls between the ages 16 - 17. The definition of urban and rural is based on the definition of the statistical office in the relevant country. Egypt does not gather information on childbirth and is thus excluded.

Figure 1.13. Percentage of children aged 16 - 17 who are married or in a consensual union by rural-urban status

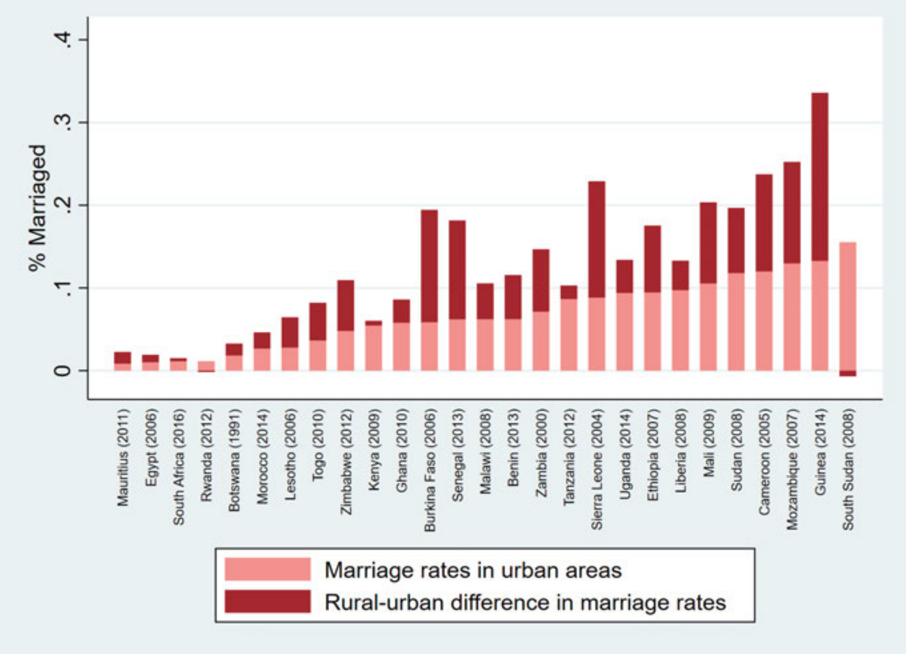

Note: Figure based on the most recently available census available from IPUMS international for all African countries. The sample includes all children between the ages 16 - 17. The definition of urban and rural is based on the definition of the statistical office in the relevant country.

Figure 1.14. Percentage of children aged 15 - 16 whose main activity is working for the family or working outside the family home by urban-rural status

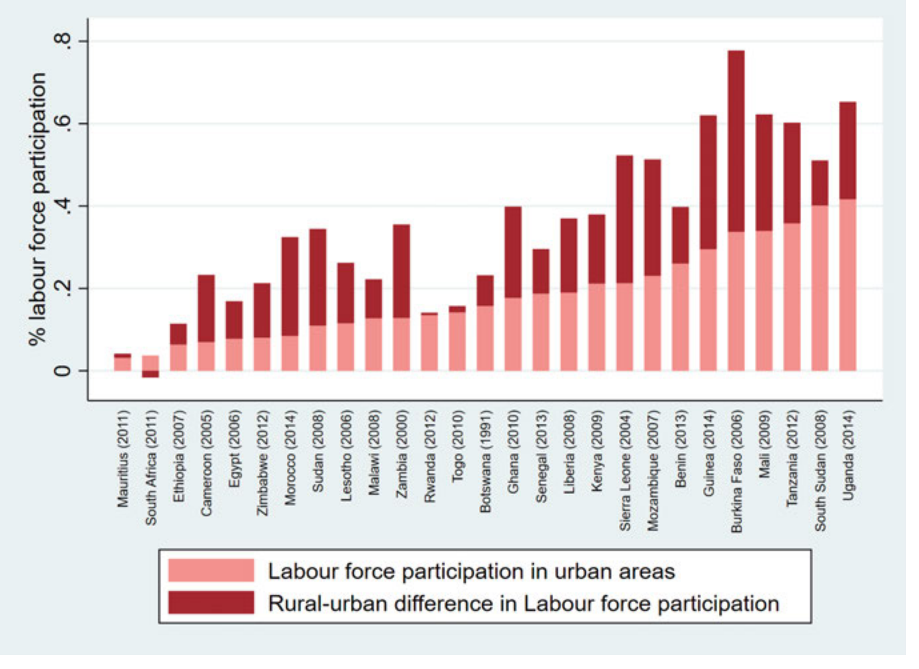

Note: Figure based on the most recently available census available from IPUMS international for all African countries. The sample includes all children between the ages 15 - 16. The definition of urban and rural is based on the definition of the statistical office in the relevant country.

1.B Sample description

Table 1.4. *Overview of included censuses*

Census	Country	Spatial unit	Year	Sample	Spatial level
1	Benin	Commune	2013	10% sample	Geolev2
2		Commune	2002	10% sample	Geolev2
3		Commune	1992	10% sample	Geolev2
4	Cameroon	Department	2005	10% sample	Geolev2
5		Department	1987	10% sample	Geolev2
6	Egypt	Province	2006	10% sample	Geolev1
7		Province	1996	10% sample	Geolev1
8		Province	1986	14% sample	Geolev1
9	Ethiopia	Region	1984	10% sample	Geolev1
10	Ghana	Region	2010	10% sample	Geolev1
11	Guinea	Prefecture	2014	10% sample	Geolev1
12		Prefecture	1996	10% sample	Geolev1
13	Kenya	District	2009	5% sample	Geolev2
14		District	1999	10% sample	Geolev2
15	Malawi	District	2008	10% sample	Geolev1
16	Mali	Circle	2009	10% sample	Geolev2
17		Circle	1998	10% sample	Geolev2
18	Morocco	Province	2004	5% sample	Geolev2
19	South Africa	District	2011	8.5% sample	Geolev2
20	Togo	Prefecture	2010	10% sample	Geolev2
21	Uganda	District	2014	10% sample	Geolev1
22		District	2002	10% sample	Geolev1
23		District	1991	10% sample	Geolev1
24	Zambia	District	2010	10% sample	Geolev2
25		District	2000	10% sample	Geolev2
26		District	1990	10% sample	Geolev2

Note: the boundaries of administrative regions shift over time. The spatial units are based on the GEOLEV1 and GEOLEV2 geographic variables of IPUMS International. The spatial unit in the second and spatial level in the fifth column reflects the level at which previous place of residence is recorded. Note that figure 1.1 and appendix figures 1.10 - 1.14 include censuses not listed above, since those censuses lack the necessary information on the individual migration histories as discussed in section 1.3.2.

Figure 1.15. Percentage of individuals who are household head or partner to household head by age

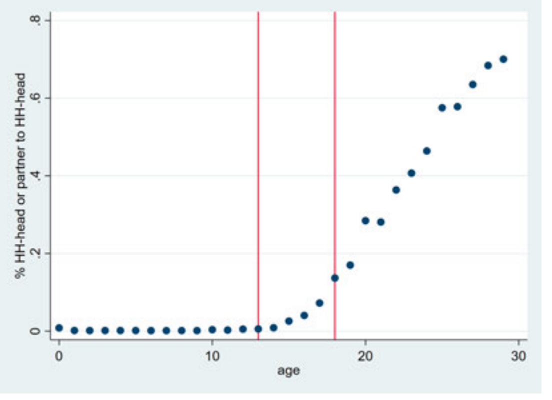

Note: The vertical bars denote ages 13 and 18. Figure based on all individuals included in the censuses listed in table 1.4.

Figure 1.16. Primary school degree completion (top left), school attendance (top right) and literacy (bottom left) by the child's age at the time of the census

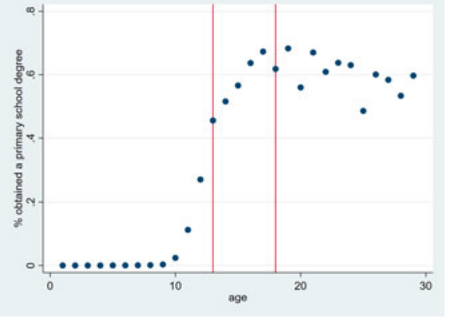

(a) Primary school completion

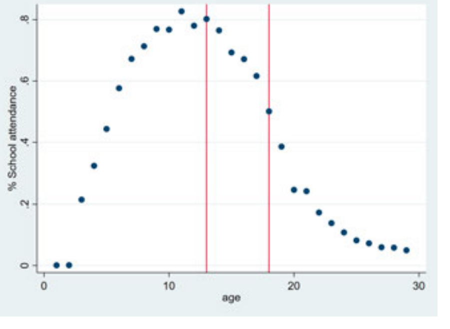

(b) School attendance

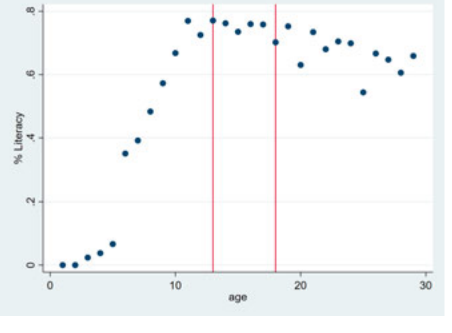

(c) Literacy

Note: The vertical bars denote ages 13 and 18. Figure based on all individuals included in the censuses listed in table 1.4.

1.C Parameterized estimates

To obtain parametric estimates and test the statistical significance of the urban exposure effects in figures 1.4 - 1.6, I impose two different functional form assumptions. The first specification imposes a linear functional form on the urban exposure effects. Equation 1.4 shows the estimating equation, where i denotes the individual, o the region of origin, c the census year, s the birth cohort, m the age of the child at the time of the move, d the region of residence at the time of the census and h the household. The δ parameter captures the selection into more urbanized regions, whereas γ captures the effect of spending the childhood years between ages 1 - 12 in an area that is one unit more urbanized.

$$\text{Educational outcome}_i = \alpha_{ocs} + (\theta_h) + \sum_{m=1}^{15} I(m_i = m)\phi_{mc}$$
$$+ \delta \Delta urbanization_{od} + \gamma \sum_{m=1}^{12} (\frac{12-m}{12}) I(m_i = m)\Delta urbanization_{od} + \varepsilon_i \quad (1.4)$$

Secondly, I estimate a somewhat simpler piecewise linear model which divides the age at moving into three equal groups: those who move between ages 1 - 5, those who move between ages 6 - 10 and those who move between ages 11 - 15. The advantage is that this specification is more flexible and does not rely on the linear functional form assumption for estimation, although it does discard the variation in urban exposure within the age groups for estimation. The estimating equation is provided in equation 1.5, with γ_{6-10} capturing the effect of moving to a region one unit more urbanized between the ages 6 - 10 and γ_{1-5} capturing the effect of moving to a region one unit more urbanized between the ages 1 - 5. The estimates of γ_{1-5} and γ_{6-10} are both relative to δ, which captures the effect of moving to an area one unit more urbanized between the ages of 11 - 15.

$$\text{Educational outcome}_i = \alpha_{ocs} + (\theta_h) + \sum_{m=1}^{15} I(m_i = m)\phi_{mc} + \delta \Delta urbanization_{od}$$
$$+ \gamma_{1-5} I(m_i < 6)\Delta urbanization_{od} + \gamma_{6-10} I(5 < m_i < 11)\Delta urbanization_{od} + \varepsilon_i \quad (1.5)$$

Tables 1.5 - 1.7 display the resulting coefficients of δ and γ obtained by estimating equations 1.4 and 1.5 on respectively primary school completion, school attendance and literacy. The estimates of γ of equation 1.4 and γ_{1-5} and γ_{6-10} of equation 1.5 are statistically significant in all cases, indicating that the urban exposure effects are present and statistically significant across both parameterizations for all three specifications across the three outcome variables. Furthermore, the γ estimates are quite similar across the specifications, indicating that controlling for the educational attainment of the household head and the inclusion of household-fixed effects does not materially alter the urban exposure estimates. At the same time, these household head characteristics can account for a substantial degree of the selection into cities, as indicated by the substantial decline in the δ estimates when household head education and literacy are controlled for.

Table 1.5. *Urban exposure estimates obtained by estimating equation 1.4 (columns 1 - 3) and equation 1.5 (columns 4 - 6) for primary school completion*

	(1) Baseline	(2) Educ.	(3) HH-FE	(4) Baseline	(5) Educ.	(6) HH-FE
delta	0.128***	0.0723***	0	0.130***	0.0714***	0
	(0.0133)	(0.0112)	(.)	(0.0134)	(0.0114)	(.)
gamma	0.0983***	0.0905***	0.156***			
	(0.0120)	(0.0128)	(0.0244)			
gamma_{1-5}				0.0740***	0.0709***	0.0936***
				(0.00861)	(0.00914)	(0.0140)
gamma_{6-10}				0.0283***	0.0301***	0.0420***
				(0.00678)	(0.00723)	(0.00818)
N	168233	144205	168233	168233	144205	168233
R^2	0.230	0.339	0.746	0.230	0.339	0.746

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The table provides the parametarized estimates of the urban exposure effects shown in figure 1.4 in the main text. Columns 1 - 3 display the coefficients obtained by estimating the linear model of equation 1.4, whereas columns 4 - 6 provide the estimates of the piecewise-linear model of equation 1.5, both discussed in appendix 1.C. The sample has been restricted across all specifications to those households with at least 2 children in the age group of 13 - 18 year olds to enhance comparison between the models. The first and fourth columns are the baseline specification, the second and fifth column include the literacy status and education of the household head ("less than primary school", "completed primary school", "completed secondary school" and "completed university") into the fixed effect and the third and sixth column include a household-fixed effect into the estimations. The estimates of δ are omitted in the case of the household-fixed effect models as $\Delta urbanization$ does not vary within the household. Standard errors are double clustered at the current and previous region of residence.

Table 1.6. *Urban exposure estimates obtained by estimating equation 1.4 (columns 1 - 3) and equation 1.5 (columns 4 - 6) for school attendance*

	(1) Baseline	(2) Educ.	(3) HH-FE	(4) Baseline	(5) Educ.	(6) HH-FE
delta	0.0313***	-0.0116	0	0.0330***	-0.0115	0
	(0.00931)	(0.00937)	(.)	(0.00938)	(0.00932)	(.)
gamma	0.0626***	0.0742***	0.0546*			
	(0.0115)	(0.0127)	(0.0214)			
gamma_{1-5}				0.0474***	0.0564***	0.0331**
				(0.00835)	(0.00938)	(0.0126)
gamma_{6-10}				0.0159*	0.0237**	0.0171*
				(0.00708)	(0.00797)	(0.00807)
N	130594	106704	130594	130594	106704	130594
R^2	0.174	0.271	0.722	0.174	0.271	0.722

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The table provides the parametarized estimates of the urban exposure effects shown in figure 1.5 in the main text. Columns 1 - 3 display the coefficients obtained by estimating the linear model of equation 1.4, whereas columns 4 - 6 provide the estimates of the piecewise-linear model of equation 1.5, both discussed in appendix 1.C. The sample has been restricted across all specifications to those households with at least 2 children in the age group of 13 - 18 year olds to enhance comparison between the models. The first and fourth columns are the baseline specification, the second and fifth column include the literacy status and education of the household head ("less than primary school", "completed primary school", "completed secondary school" and "completed university") into the fixed effect and the third and sixth column include a household-fixed effect into the estimations. The estimates of δ are missing in the case of the household-fixed effect models as $\Delta urbanization$ does not vary within the household. Standard errors are double clustered at the current and previous region of residence.

Table 1.7. *Urban exposure estimates obtained by estimating equation 1.4 (columns 1 - 3) and equation 1.5 (columns 4 - 6) for literacy*

	(1) Baseline	(2) Educ.	(3) HH-FE	(4) Baseline	(5) Educ.	(6) HH-FE
delta	0.0714^{***}	0.0312^{***}	0	0.0723^{***}	0.0320^{***}	0
	(0.00943)	(0.00730)	(.)	(0.00952)	(0.00737)	(.)
gamma	0.0679^{***}	0.0669^{***}	0.0565^{***}			
	(0.0111)	(0.0107)	(0.0155)			
gamma_{1-5}				0.0517^{***}	0.0502^{***}	0.0362^{***}
				(0.00791)	(0.00786)	(0.00928)
gamma_{6-10}				0.0194^{**}	0.0200^{**}	0.0141^{*}
				(0.00622)	(0.00606)	(0.00579)
N	145034	143362	145034	145034	143362	145034
R^2	0.153	0.279	0.760	0.153	0.279	0.760

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The table provides the parametarized estimates of the urban exposure effects shown in figure 1.6 in the main text. Columns 1 - 3 display the coefficients obtained by estimating the linear model of equation 1.4, whereas columns 4 - 6 provide the estimates of the piecewise-linear model of equation 1.5, both discussed in appendix 1.C. The sample has been restricted across all specifications to those households with at least 2 children in the age group of 13 - 18 year olds to enhance comparison between the models. The first and fourth columns are the baseline specification, the second and fifth column include the literacy status and education of the household head ("less than primary school", "completed primary school", "completed secondary school" and "completed university") into the fixed effect and the third and sixth column include a household-fixed effect into the estimations. The estimates of δ are missing in the case of the household-fixed effect models as $\Delta urbanization$ does not vary within the household. Standard errors are double clustered at the current and previous region of residence.

1.D Robustness tests

1.D.1 Robustness to changing urbanization measure

Census definition

The majority of the 26 censuses include the urban/rural status of households[21], although the definition of "urban" varies greatly between countries.[22] The correlation with the baseline urbanization measure is relatively high at 0.85, which results in similar urban exposure effects as shown in figure 1.17. Moving to a region one percentage point more urban before the age of 6 instead of after the age of 12 increases the probability of completing primary school by around 0.08 percentage points. This is very close to the baseline result of. The increase in school attendance at ages 13 - 18 (0.05 percentage points) and literacy (0.04 percentage points) are also similar to the baseline results, as are the results of the household fixed effects models shown in figure 1.18. Hence, both the Africapolis and census-based urbanization measures yield very similar results.

Figure 1.17. Estimates of b_m based on the census definitions of urban for primary school completion (left), school attendance (center) and literacy (right)

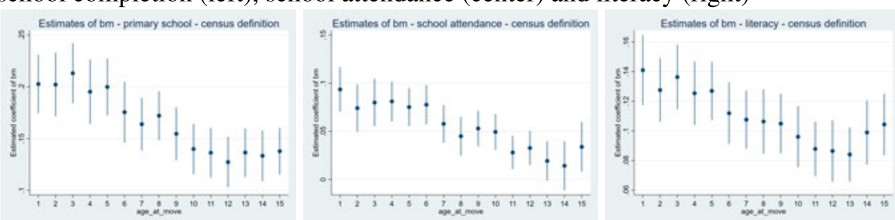

Note: Figure shows estimates of b_m obtained by estimation equation 1.1 with the urbanization measure based on definition of the individual country censuses. See footnote 22 on the definition of urban by country. The left panel shows the coefficients for primary school completion, the center panel school attendance and the right panel literacy.

[21] The 2004 Morocco census and 2010 Zambia census are the exceptions.

[22] In six countries the designation depends on whether the town in which a household resides exceeds either 2000 inhabitants (Ethiopia, Kenya and Uganda) or 5000 inhabitants (Ghana, Mali and Zambia). Three countries designate the municipality or district containing the regional capital as an urban area (Egypt, Guinea and Togo), regardless of its population. Cameroon uses a mixture of both and designates a town as urban if it either serves as an administrative center or if it has a population of more than 5000 inhabitants and various facilities such as a health center and a daily market. Malawi categorizes all townships and district centers as urban areas. Finally, Benin bases the urban definition on the number of individuals living in a census tract, whereas Egypt uses the dominant land use in a census tract.

Figure 1.18. Estimates of b_m based on the census definitions of urban when including household fixed effects for primary school completion (left), school attendance (center) and literacy (right)

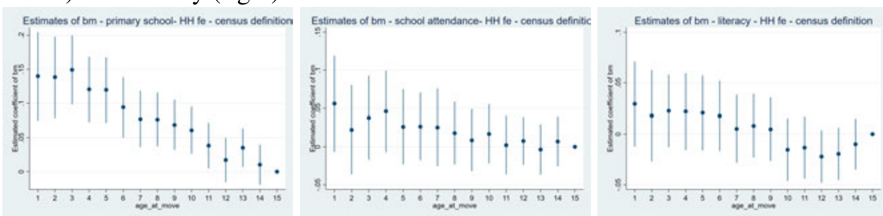

Note: Figure shows estimates of b_m obtained by estimation equation 1.1 when $\Delta urbanization$ is based on the definition of urban of the individual countries (see footnote 22 for details). The left panel shows the coefficients for primary school completion, the central panel for school attendance and the right panel for literacy. The urban exposure effects are statistically significant at the 5% level when applying the linear functional form specification of equation 1.4 discussed in appendix section 1.C.

Population density

Figure 1.19 displays the estimates of b_m when $\Delta urbanization$ is instead based on population density. The coefficient b_m should in this case be interpreted as the effect of moving to an area with a one log-point higher population density at age m. Moving to a region with a one log-point higher population density before the age of 6 increases primary school completion and school attendance by around one percentage point, and literacy by around half a percentage point compared to moving after the age of 12. A regression of log population density on the baseline urbanization measure reveals a coefficient of 0.123, implying that the rescaled coefficients are similar in magnitude to the baseline results. The coefficients are again similar when household fixed effects are included as shown in figure 1.20.

Figure 1.19. Estimates of b_m based on population density for primary school completion (left), school attendance (center) and literacy (right)

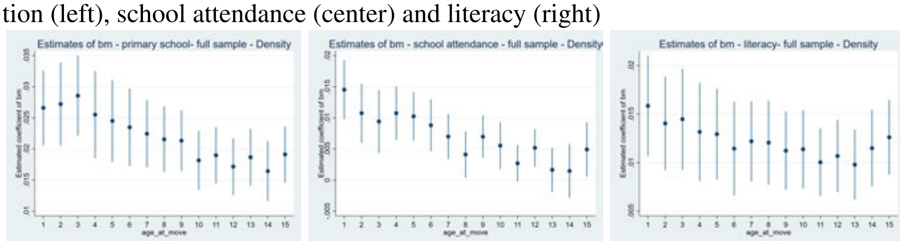

Note: Figure shows estimates of b_m obtained by estimation equation 1.1 when $\Delta urbanization$ is defined as $\log\left(\frac{\text{population density current region}}{\text{population density previous region}}\right)$ for primary school completion (left), school attendance (center) and literacy (right) at ages 13 - 18.

Figure 1.20. Estimates of b_m based on population density when including household fixed effects for primary school completion (left), school attendance (center) and literacy (right)

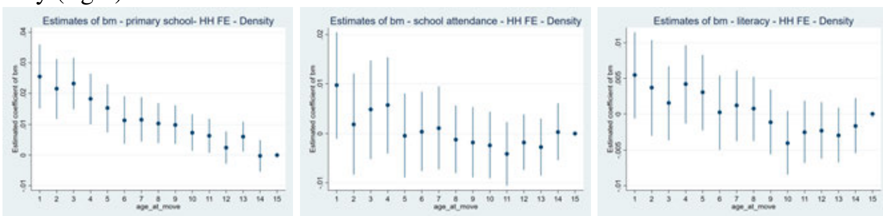

Note: Figure shows estimates of b_m obtained by estimation equation 1.1 including household fixed effects with $\Delta urbanization$ defined as $\log\left(\frac{\text{population density current region}}{\text{population density previous region}}\right)$ for primary school completion (left), school attendance (center) and literacy (right). The urban exposure effects are statistically significant at the 5% level when applying the linear functional form specification of equation 1.4 discussed in appendix section 1.C.

Migration to capital cities.

As the third measure of urbanization, I focus specifically on exposure to capital cities. Migration to and from capital cities accounts for a third of the total moves across regions and is hence of interest to analyze in their own right. To estimate the effect of exposure to capital cities, I re-estimate equation 1.1 with $\Delta urbanization_{od}$ compactly expressed as (I(current region of residence $=$ capital region)$-$ I(previous region of residence $=$ capital region)), thus implicitly assuming symmetric exposure effects for households moving to and from capital cities.[23] As figure 1.21 shows, moving to a capital city before the age of 6 raises primary school completion and school attendance by around 5 percentage points and literacy by around 2.5 percentage points, compared to migrating after the age of 12. Adding household fixed effects does not significantly change the estimates, as shown in figure 1.22. Hence, the urban exposure effects documented in section 1.4.1 do not depend on the selected measure of urbanization, but are also visible when using three other definitions of urbanization.

[23] The heterogeneity tests in section 1.5 show the results when relaxing this assumption.

Figure 1.21. Estimates of b_m based on exposure to capital cities for primary school completion (left), school attendance (center) and literacy (right)

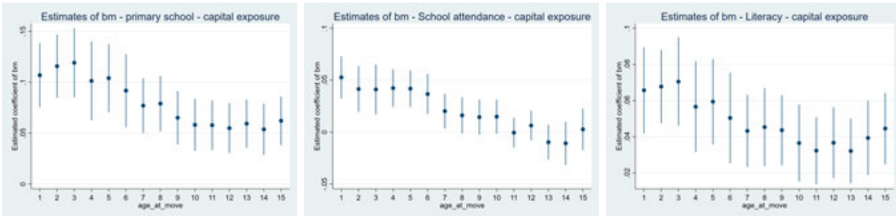

Note: Figure shows estimates of b_m obtained by estimation equation 1.1 when $\Delta urbanization$ is defined as $I(\text{current region of residence} = \text{capital region}) - I(\text{previous region of residence} = \text{capital region})$ and thus captures exposure to capital cities. The left panel shows the coefficients for primary school completion, the center panel for school attendance and the right panel for literacy.

Figure 1.22. Estimates of b_m based on exposure to capital cities when including household fixed effects for primary school completion (left), school attendance (center) and literacy (right)

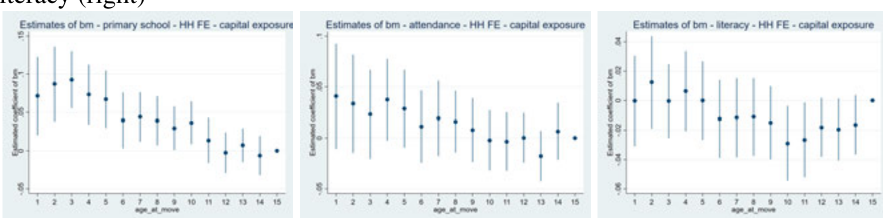

Note: Figure shows estimates of b_m obtained by estimation equation 1.1 including household fixed effects when estimating the impact of exposure to capital cities for primary school completion (left), school attendance (center) and literacy (right). The coefficients are obtained by estimation equation 1.1 with $\Delta urbanization$ defined as $I(\text{current region of residence} = \text{capital region}) - I(\text{previous region of residence} = \text{capital region})$. The urban exposure effects are statistically significant at the 5% level when applying the linear functional form specification of equation 1.4 discussed in appendix section 1.C.

1.D.2 Estimates by country

Table 1.8. *Urban exposure effects by country for primary school completion*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Benin	Cameroon	Egypt	Ethiopia	Ghana	Guinea	Kenya
delta	0.0358	0.0866***	-0.0445	0.109*	0.149***	0.0718*	0.125***
	(0.0341)	(0.0190)	(0.0502)	(0.0505)	(0.0291)	(0.0279)	(0.0127)
gamma	0.123**	0.0714***	0.126***	0.215***	0.172***	0.130***	0.0365*
	(0.0415)	(0.0191)	(0.0322)	(0.0341)	(0.0148)	(0.0343)	(0.0160)
N	16987	34999	51493	7517	21582	17390	38575
R^2	0.229	0.200	0.054	0.197	0.179	0.133	0.256

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Malawi	Mali	Morocco	S. Africa	Togo	Uganda	Zambia
delta	0.107***	0.150***	-0.000798	0.0424*	0.0679*	0.332***	0.123***
	(0.0253)	(0.0403)	(0.0524)	(0.0183)	(0.0285)	(0.0196)	(0.0150)
gamma	0.0735*	0.125***	0.185***	0.0740	0.179**	0.0867***	0.106***
	(0.0355)	(0.0216)	(0.0406)	(0.0755)	(0.0536)	(0.0195)	(0.0166)
N	9841	7455	14208	8545	6691	44861	51645
R^2	0.200	0.155	0.091	0.115	0.148	0.281	0.131

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The table shows the coefficients of gamma and delta obtained by estimating equation 1.3 on primary school completion by country. Standard errors are double clustered at the current and previous region of residence.

Table 1.9. *Urban exposure effects by country for school attendance*

	(1) Benin	(2) Cameroon	(3) Egypt	(4) Ethiopia	(5) Ghana	(6) Guinea	(7) Kenya
delta	-0.0161	-0.0206	-0.0558	0.123**	0.0198	0.0625*	-0.0189
	(0.0331)	(0.0144)	(0.104)	(0.0368)	(0.0251)	(0.0303)	(0.00957)
gamma	0.165***	0.00000573	0.162***	0.0688	0.117*	0.100***	0.0253
	(0.0387)	(0.0191)	(0.0235)	(0.0393)	(0.0443)	(0.0249)	(0.0133)
N	17024	24964	11654	7525	21582	16793	38661
R^2	0.271	0.070	0.063	0.069	0.087	0.061	0.169

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

	(1) Malawi	(2) Mali	(3) South Africa	(4) Togo	(5) Uganda	(6) Zambia
delta	0.0662**	0.163***	0.0343	-0.00210	0.0428*	0.0361**
	(0.0209)	(0.0434)	(0.0182)	(0.0338)	(0.0176)	(0.0129)
gamma	0.0185	0.119***	-0.0571	0.0828	0.101***	0.0560***
	(0.0210)	(0.0277)	(0.0457)	(0.0625)	(0.0212)	(0.0149)
N	9851	7746	8459	6701	44629	51645
R^2	0.078	0.177	0.100	0.107	0.157	0.098

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The table shows the coefficients of gamma and delta obtained by estimating equation 1.3 on school attendance at ages 13 - 18 by country. Morocco is not included as it does not record school attendance in the 2004 census. Standard errors are double clustered at the current and previous region of residence.

Table 1.10. *Urban exposure effects by country for literacy*

	(1) Benin	(2) Cameroon	(3) Egypt	(4) Ethiopia	(5) Ghana	(6) Guinea
delta	0.0301	0.0476***	-0.0207	0.178***	0.0607	0.0982**
	(0.0302)	(0.0127)	(0.0391)	(0.0353)	(0.0269)	(0.0367)
gamma	0.0937**	0.0151	0.109***	0.0531	0.0741	0.100**
	(0.0303)	(0.0102)	(0.0204)	(0.0274)	(0.0357)	(0.0324)
N	17013	35491	51493	7517	21582	17448
R^2	0.220	0.190	0.053	0.126	0.069	0.069

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

	(1) Malawi	(2) Mali	(3) Morocco	(4) Togo	(5) Uganda	(6) Zambia
delta	0.0471***	0.161***	0.0502*	0.0318	0.162***	0.0676***
	(0.0104)	(0.0413)	(0.0203)	(0.0279)	(0.0213)	(0.0107)
gamma	-0.000192	0.159***	0.0720**	0.0213	0.0223	0.0660***
	(0.0110)	(0.0311)	(0.0216)	(0.0348)	(0.0204)	(0.0132)
N	9851	7174	14208	4889	44871	51645
R^2	0.033	0.163	0.076	0.113	0.118	0.094

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The table shows the coefficients of gamma and delta obtained by estimating equation 1.3 on literacy by country. Kenya and South Africa are not included as they do not record literacy in any of their censuses. Standard errors are double clustered at the current and previous region of residence.

1.D.3 Age of children

Figure 1.23. Urban exposure effects by age group

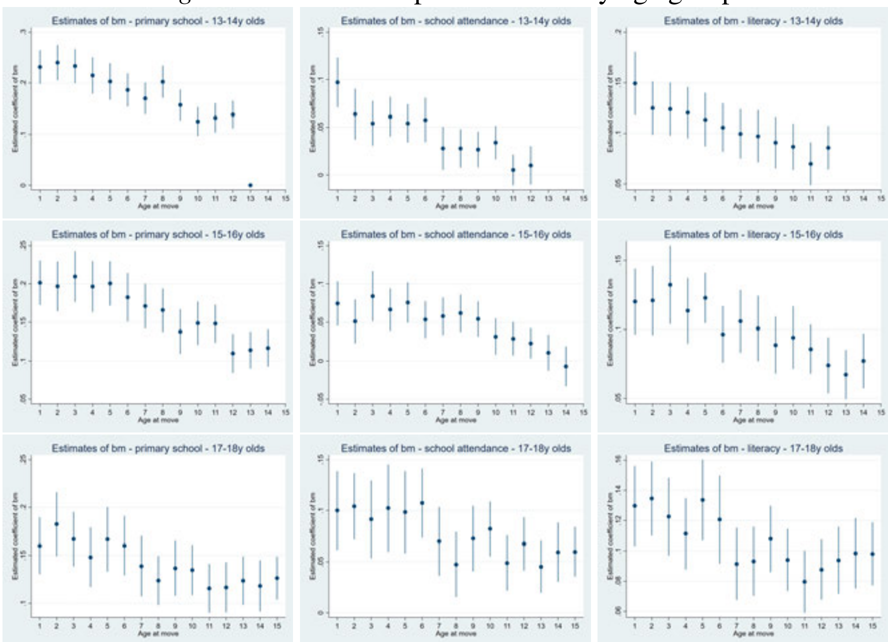

Note: The first row contains the estimates if the sample is restricted to 13 and 14 year olds, the second row the estimates if the sample is restricted to 15 and year olds and the final row the estimates if the sample is restricted to 17 and 18 year olds. The first column displays the estimates of b_m in the case of primary school completion, the second column the estimates of b_m in the case of school attendance and the third column the estimates of b_m in the case of literacy.

Table 1.11. *Urban exposure effects by age group.*

	Primary School completion			School attendance		
	Ages 13/14	Ages 15/16	Ages 17/18	Ages 13/14	Ages 15/16	Ages 17/18
delta	0.126***	0.129***	0.118***	0.00730	0.0311***	0.0635***
	(0.0129)	(0.0122)	(0.0116)	(0.00764)	(0.00905)	(0.0111)
gamma	0.130***	0.100***	0.0642***	0.0790***	0.0586***	0.0524***
	(0.0117)	(0.0102)	(0.0127)	(0.0110)	(0.0102)	(0.0144)
N	137250	111362	83642	113954	89061	64673
R^2	0.205	0.182	0.158	0.120	0.120	0.133

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

	Literacy		
	Ages 13/14	Ages 15/16	Ages 17/18
delta	0.0745***	0.0806***	0.0848***
	(0.0102)	(0.00877)	(0.00848)
gamma	0.0682***	0.0553***	0.0547***
	(0.00999)	(0.00958)	(0.0105)
N	115253	95734	72195
R^2	0.152	0.142	0.135

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The table shows the coefficients of gamma and delta obtained by estimating equation 1.3 seperately for each age group for primary school completion, school attendance, literacy and years of education. Standard errors are double clustered at the current and previous region of residence.

1.D.4 Robustness to potential other shocks

Utilizing exceptionally large migration outflows

Figure 1.24. Migration outflows and urban exposure estimates for school attendance

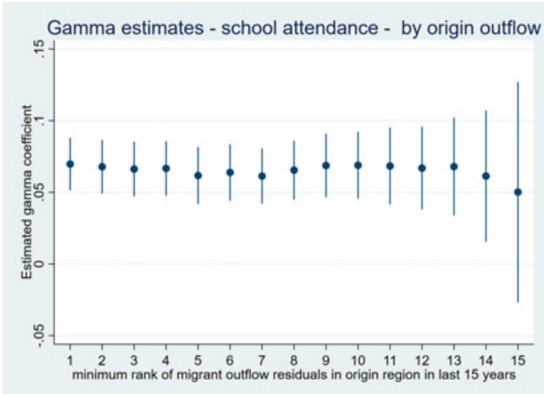

Note: The figure shows the estimates of γ obtained by estimating equation 1.7 on a successively more selective sample. The first estimation is based on all migration events in the 15 years prior to the census, whereas the final estimation is only based on the year with the strongest percentage deviation in the migration outflow compared to what would be expected based on a linearly predictive model. The 1996 Guinean census and the 1984 Ethiopian census are excluded since they only record the years since migration in intervals, and the South African census is excluded as it only records migration events in the 10 years prior to the census.

Figure 1.25. Migration outflows and urban exposure estimates for literacy

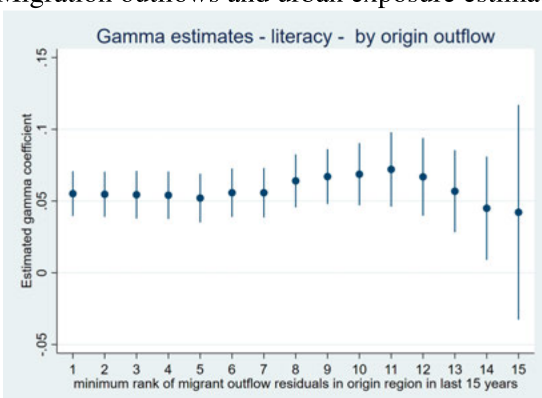

Note: The figure shows the estimates of γ obtained by estimating equation 1.7 on a successively more selective sample. The first estimation is based on all migration events in the 15 years prior to the census, whereas the final estimation is only based on the year with the strongest percentage deviation in the migration outflow compared to what would be expected based on a linearly predictive model. The 1996 Guinean census and the 1984 Ethiopian census are excluded since they only record the years since migration in intervals, and the South African census is excluded as it only records migration events in the 10 years prior to the census.

Rainfall shocks

Table 1.12. *Urban exposure estimates when years with large rainfall shocks in region of origin are excluded.*

	Primary School		School attendance		Literacy	
	(1)	(2)	(3)	(4)	(5)	(6)
delta	0.132***	0.130***	0.0248***	0.0238***	0.0792***	0.0804***
	(0.0115)	(0.0111)	(0.00686)	(0.00722)	(0.00850)	(0.00829)
gamma	0.0930***	0.0973***	0.0537***	0.0498***	0.0518***	0.0487***
	(0.00880)	(0.0106)	(0.00818)	(0.00992)	(0.00887)	(0.0106)
N	271983	168364	248097	154021	222667	137565
R^2	0.240	0.256	0.151	0.163	0.155	0.172
Weather shocks	No	Yes	No	Yes	No	Yes

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The table shows the coefficients of gamma and delta obtained by estimating equation 1.3. The even columns shows the results when migration events are excluded where a large weather shock took place in the year of the migration. Large weather shocks are defined as a rainfall shock above the 80th or below the 20th percentile for the region following Shah and Steinberg (2017). The rainfall data taken from the CHIRPS v2.0 database, which is available since 1981. I drop the censuses taken before 1990 since rainfall data in the region of origin at the time of migration is missing for many migrants. The rainfall is aggregated to the average per year per region, with the percentiles calculated using the rainfall between 1981 and 2017. The estimates remain unchanged when also excluding migration events where an weather shock took place in the year prior to migration (results not shown here)

Conflict shocks

Table 1.13. *Urban exposure estimates when observations with violent events in region of origin are excluded.*

	Primary School		School attendance		Literacy	
	(1)	(2)	(3)	(4)	(5)	(6)
delta	0.126***	0.115***	0.0285***	0.0297***	0.0790***	0.0757***
	(0.0109)	(0.00999)	(0.00698)	(0.00713)	(0.00825)	(0.00810)
gamma	0.103***	0.107***	0.0657***	0.0603***	0.0605***	0.0588***
	(0.00835)	(0.00947)	(0.00811)	(0.00888)	(0.00746)	(0.00784)
N	332254	296772	267688	232407	283182	255846
R^2	0.224	0.222	0.152	0.159	0.143	0.147
Violence shocks	No	Yes	No	Yes	No	Yes

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The table shows the coefficients of gamma and delta obtained by estimating equation 1.3. The even columns show the estimates when migration events are excluded where a violent event resulting in at least one death was reported in the region of origin between three years prior to the move and two years after the move. Geocoded data on violent events is obtained from the GED20.1 database produced by the Uppsala Conflict Data Program (UCDP). Events are only included in the database if they are part of a larger campaign of organized violence, with at least 25 deaths per calendar-year required for inclusion in the database.

1.D.5 Selection within regions by age of migration

Selection within region of arrival

For some of the countries used in the baseline analyses, the current and previous regions of residence are relatively large (see figure 1.2), which means that they likely contain a variety of urban and rural areas within the region. One concern might be that households migrate to more or less urban areas within these regions depending on the age of the children, which might affect the estimates. For 5 countries in particular the regions are of a relatively coarse geographical level (GEOLEV1, see table 1.4) as this is the most detailed level on which the previous place of residence is available. In the baseline specifications, the same regional divisions are used to classify current and previous regions of residence for consistency.

However, we can use the more detailed information from the GEOLEV2 variable available for these five countries to investigate whether the results are driven by heterogeneity in the destinations where migrants move to within the regions of arrival. To do so, I estimate equation 1.3 for these five countries using the more fine-grained current place of residence and calculate the urbanization measure for each of these smaller regions. The number of regions increases from 124 to 721 when using more detailed information on the region of residence. Table 1.14 shows the results for both the baseline results and the more detailed regions for these five countries, with no significant changes to the estimates of gamma. It is worth noting that the region of residence is already based on this finer classification for the countries not included in this table.

Table 1.14. *Urban exposure estimates with the current place of residence based on both the Geolev1 and Geolev2 classification*

	Primary School		School attendance		Literacy	
	(1)	(2)	(3)	(4)	(5)	(6)
delta	0.192***	0.207***	0.0563***	0.0815***	0.117***	0.132***
	(0.0240)	(0.0207)	(0.0115)	(0.0107)	(0.0149)	(0.0130)
gamma	0.131***	0.122***	0.0874***	0.0725***	0.0503***	0.0368**
	(0.0146)	(0.0129)	(0.0144)	(0.0130)	(0.0132)	(0.0119)
N	101191	101191	100380	100380	101269	101269
R^2	0.236	0.246	0.129	0.132	0.151	0.157
Number of regions	124	721	124	721	124	721

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The table shows the coefficients of γ and δ obtained by estimating equation 1.3 for the five countries (Ethiopia, Ghana, Guinea, Malawi, Uganda) for whom the current region of residence in the baseline specifications is at a relatively high aggregation level. The odd columns use the baseline regions of residence for these 5 countries which are also shown in figure 1.2, whereas the even columns use the more detailed Geolev2 regions which are also used for the other countries. Standard errors are double clustered at the current and previous region of residence.

Selection within region of origin

Similarly, one concern might be that children originate from different regions within the region of origin, and that this selection might change with the age of moving. While the baseline specifications already use the most detailed available information on the previous place of residence, eight censuses also record whether individuals originate from a rural or urban area in their previous place of residence. To investigate whether selection within the regions of origin might explain the results, table 1.15 shows the gamma estimates of equation 1.1 with the baseline specification (odd columns) and when the urban/rural status of the previous residence is included in the region of origin by census year by birth cohort fixed effect. The estimates of gamma are very similar across the specifications, suggesting that also selection within the regions of origin by the age of migration is not driving the results.

Table 1.15. *Urban exposure estimates when controlling for the urban/rural status of the previous place of residence*

	Primary School		School attendance		Literacy	
	(1)	(2)	(3)	(4)	(5)	(6)
delta	0.100***	0.0926***	0.0181	0.0165	0.0550***	0.0497***
	(0.0128)	(0.0120)	(0.0103)	(0.00975)	(0.00897)	(0.00813)
gamma	0.0951***	0.0940***	0.0517***	0.0489***	0.0604***	0.0606***
	(0.0124)	(0.0119)	(0.0127)	(0.0127)	(0.0102)	(0.0102)
N	134230	134230	94183	94183	132720	132720
R^2	0.120	0.143	0.091	0.112	0.091	0.115
Urban/rural status?	No	Yes	No	Yes	No	Yes

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The table shows the coefficients of γ and δ obtained by estimating equation 1.3 for the eight censuses (the three Egyptian censuses, the three Zambian censuses, the 2005 Cameroon census and the 2010 Togo census) for whom information is available on the urban/rural status of the previous place of residence of the household. The odd columns show the gamma and delta obtained by estimating equation 1.3 on these eight censuses, whereas the odd columns show the results when the urban/rural status of the previous place of residence are included into the fixed effect term of equation 1.3. Standard errors are double clustered at the current and previous region of residence.

Repeat migration

Finally, one concern is that the baseline regressions base the $urbanization_{od}$ term either on the previous region of residence, or if unavailable, on the region of birth. As such, childhood exposure to regions might not be accurate for those who migrated multiple times across regions during childhood. To assess whether this affects the results, I utilize the fact that 13 censuses separately record the previous region of residence and the region of birth. For those individuals for whom the region of residence is the same as the region of birth, we can be reasonably certain that they have not moved across regions multiple times during childhood.

Table 1.16 below shows the estimates of gamma obtained by estimating equation 1.3 on the 13 censuses with information on both the previous place of residence as well as the region of birth. The uneven columns include all observations, whereas the even columns restrict the sample to those individuals for whom the place of birth and previous place of residence is the same. As can be seen, the estimates of gamma are statistically significant in all cases and actually somewhat larger when excluding individuals who moved multiple times. Hence, the results are also not driven by individuals who migrated repeatedly across regions.

Table 1.16. *Urban exposure estimates when repeat movers are excluded.*

	Primary School		School attendance		Literacy	
	(1)	(2)	(3)	(4)	(5)	(6)
delta	0.123***	0.122***	0.0241	0.00825	0.0773***	0.0751***
	(0.0222)	(0.0256)	(0.0130)	(0.0128)	(0.0139)	(0.0155)
gamma	0.105***	0.128***	0.0961***	0.129***	0.0635***	0.0780***
	(0.0135)	(0.0139)	(0.0158)	(0.0178)	(0.0127)	(0.0157)
N	120790	94927	100706	79051	121113	95127
R^2	0.259	0.278	0.147	0.170	0.187	0.202
Restricted to previous region same as birth region	No	Yes	No	Yes	No	Yes

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The table shows the coefficients of γ and δ obtained by estimating equation 1.3 on the 13 censuses which record information on both the previous region of residence as well as the region of birth. The uneven columns include all observations, whereas the even columns restrict the sample ot the observations where the previous region of residence is the same as the region of birth. Standard errors are double clustered at the current and previous region of residence.

1.E Mechanisms

1.E.1 Access to public services

Table 1.1. *Urban exposure heterogeneity by access to public services (electricity and piped water)*

	Primary School		School attendance		Literacy	
	(1)	(2)	(3)	(4)	(5)	(6)
gamma	0.112***	0.0857***	0.0737***	0.0604***	0.0647***	0.0629***
	(0.0153)	(0.0116)	(0.0142)	(0.0118)	(0.0147)	(0.0133)
gamma $*$ access	0.00396		-0.0142		-0.0187	
piped water	(0.0184)		(0.0164)		(0.0150)	
gamma $*$ access		0.0388*		0.00262		-0.00843
electricity		(0.0157)		(0.0137)		(0.0151)
N	183964	214883	169068	200314	162929	165280
R^2	0.234	0.241	0.136	0.152	0.135	0.144

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table estimates the standard model of equation 1.3, with both gamma and delta interacted with the access of households to piped water (odd columns) and household access to electricity (even columns). The information on access to piped water is available for 15 censuses and information on access to electricity for 17 censuses. Across the censuses, on average 71% of the urban households has access to piped water against 17% of the rural households. For access to electricity the differences is 64% against 17%. Standard errors are double clustered at the current and previous region of residence.

1.E.2 Housing wealth

Table 1.2. *Urban exposure heterogeneity by house material type (organic material versus brick/concrete)*

	Primary School		School attendance		Literacy	
	(1)	(2)	(3)	(4)	(5)	(6)
gamma	0.0952***	0.0807***	0.0535***	0.0398**	0.0556**	0.0485**
	(0.0146)	(0.0144)	(0.0140)	(0.0135)	(0.0173)	(0.0152)
gamma $*$	-0.00222		0.00618		-0.00880	
non-organic walls	(0.0170)		(0.0158)		(0.0198)	
gamma $*$		0.0120		0.0247		-0.000766
non-dirt floor		(0.0175)		(0.0156)		(0.0161)
N	243522	237330	233752	227552	194674	197295
R^2	0.251	0.245	0.160	0.157	0.159	0.161

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table estimates the standard model of equation 1.3, with both gamma and delta interacted with the type of wall construction material (organic versus non-organic) in the odd columns and the floor type (dirt versus non-dirt) in the even columns. The information on wall and floor construction material is available for 19 and 18 censuses respectively. Across the censuses, on average 78% of the urban households has walls made out of non-organic materials, against 47% of the rural households. For the floor type, 86% of the urban households has a non-dirt floor against 42% of the rural households. Standard errors are double clustered at the current and previous region of residence.

1.E.3 Descriptive evidence from Guinean 2014 census

This section presents the descriptive evidence on the mechanisms from the Guinean 2014 census, which specifically asks individuals between the ages 6 and 18 who left school for their motivations. Among the individuals aged 13 - 18, 65% of the children living in urban areas reports still attending school at the time of the census, against 29% in the rural areas. Table 1.3 below shows the school activity among students in both urban and rural areas, as well as the reasons why they left school in case they are no longer enrolled.

Table 1.3. *School participation and reasons for school leaving according to 2014 Guinean census*

Reason for leaving school	Urban	Rural	Difference
Attends school	64.61%	29.06%	-35,55%
Family required for household tasks	4.98%	17.58%	12,60%
Access to a school	2.04%	15.66%	13,62%
Learning outside of school	11.67%	11.73%	0.06%
Insufficient economic means	6.46%	12.22%	5,76%
Did not want to attend school or to study	4.52%	5.77%	1.25%
Marriage or pregnancy	1.78%	4.14%	2.36%
Academic issues	2.96%	2.75%	-0.21%
Illness/disability	0.86%	1.05%	0.15%
Unknown/missing	0.03%	0.02%	-0.01%
Total	100%	100%	0%
N	57,236	83,116	140,352

Note: Data taken from the Guinean 2014 census. The school attendance question is recorded for all children aged 6 - 18, with the follow-up question regarding the reasons for leaving school asked of those who are no longer attending school. The frequencies reported in the table are based on the children between ages 13 and 18.

2. The effect of childhood urban residency on earnings: Evidence from Brazil

I would like to express my gratitude to Douglas Almond, Guillermo Alves, Niklas Bengtsson, Matz Dahlberg, Magdalena Dominguez, Alice Hallman, Johannes Haushofer, Ana Moreno-Maldonado, Jorge Perez Perez, Adrian Poignant, Markus Ridder, Albert Saiz, Daniel Sturm and Gabriel Ulyssea as well as participants to the Uppsala Urban Lab seminar, PAA Conference, SOLE conference, EALE Conference and Urban Economics Association conference for their helpful comments and suggestions on previous versions of this paper.

2.1 Introduction

Are urban dwellers in developing countries wealthier due to spatial selection or by the virtue of living in the city? Despite a long history of economic research dating back to Harris and Todaro (1970), the origins of the urban-rural income gap remain unclear. Recent studies find relatively small income gains from rural to urban migration compared to the cross-sectional income gaps, which has lead to the consensus that the difference mainly reflect the spatial selection of household. However, existing research is largely based on the study of adult migrants, thus ignoring the potential effects of urban residency during childhood. This omission is surprising, since the work by Chetty, Hendren and Katz (2016) and Chetty and Hendren (2018) highlights that childhood is a key period during which the environment affects labor market outcomes.

In this paper, I advance the literature on urbanization in the developing world by studying the impact of childhood urban residency on labor market outcomes. To account for the spatial selection of households, I use an age-at-move identification strategy combined with high quality Brazilian census data. The Brazilian census data are unique in the sense that they provide detailed information on the migration history and labor market outcomes for a large group of households, including income, earnings and hours worked. The censuses contain information on 58 million individuals, which results in a baseline estimating sample of 267,751 individuals who migrated across state boundaries during their childhood and are between ages 25 - 29 at the time of the census.

The analyses show that spending the childhood in an urban environment significantly improves labor market outcomes in adulthood. Spending the childhood prior to age 13 in an environment one log-point denser increases adulthood earnings by 2 - 3 percent. This increase is fully explained by an increase in hourly wages, with little effect of childhood urban exposure on labor market participation or hours worked. Further analysis reveals an important role for human capital accumulation in explaining the wage increase. Spending the childhood before age 13 in an environment one log point denser increases primary and secondary school completion rates by 2 percentage points and years of schooling by 0.2. The increase in educational attainment can account for over half of the effect of childhood urban exposure on earnings, suggesting that others factors play a more limited role or operate through educational attainment.

One concern is that selective mobility during early adulthood might drive the results, since the censuses only record information on the most recent migration event. Such selection may occur if the most ambitious childhood migrants to rural areas subsequently move to urban areas during early adulthood, in which case the observed rural childhood migrants would be negatively selected. However, the childhood urban exposure effects are statistically significant and similar in size for all ages at which outcomes are measured between ages 17 and 29, suggesting that selective migration in adulthood is not driving the results.

A second concern is that the selection of households into more and less urban environments may vary with the age of migration, which would violate the identifying assumption. To investigate this possibility, I focus on the group of individuals aged 17 - 20 who are still living with their parents and for whom we observe household characteristics. Controlling for a variety of individual and household characteristics, including the educational attainment of the parents, does not change the childhood urban exposure estimates, suggesting that selection into migration is unlikely to drive the results. To alleviate concerns over selection based on unobserved characteristics, I use the methodology Oster (2019). Using the selection based on observed characteristics as a guide for selection on unobserved characteristics, I show that the estimates remain similar and statistically significant when correcting for potential selection based on unobserved characteristics.

The heterogeneity analyses reveal two additional findings of interest. First of all, the urban exposure estimates are similar for both genders, but slighly smaller for non-white children, suggesting that the effects of childhood urban exposure somewhat varies between demographic groups. Second, the urban exposure estimates change significantly over time, with the estimates based on the 2010 being around $2/3^{th}$ smaller compared to the 1991 census. Whether this decline is due to the expansion of public service provision in rural Brazil between 1991 and 2010 or whether it reflects a general trend as countries become wealthier is an interesting question for future research. Nonetheless, the urban exposure estimates remain economically relevant and statistically significant for labor market outcomes throughout all periods.

The findings of this paper contribute to three strands of literature. First of all, the paper contributes to the literature on the effect of rural to urban migration in developing countries. Recent studies have focused on the effect of urban migration during adulthood (Alvarez, 2020; Hamory et al., 2020; Lagakos et al., 2020), thus missing the potential gains of urban exposure during the childhood. The findings of this paper highlight that this omission is in fact quite relevant. Quantitatively, the results suggest that spending the childhood in a city rather than a rural area raises adulthood wages by around 5 percent[1], which is similar in magnitude and in addition to the gains of adulthood urban migration reported by Alvarez (2020) for Brazil over the same period.

Second, the paper contributes to an emerging literature on the effects of childhood urban exposure on educational attainment in developing countries, following the work of Adukia, Asher and Novosad (2020) in India, van Maarseveen (2022) in Africa and Chiovelli et al. (2021) in Mozambique. All three paper find that spending the childhood in an urban environment raises educa-

[1]Based on the average population density within urban/rural areas to translate the findings of this article based on the continuous density measure into the urban-rural dichotomy.

tional attainment, but have not been able to extend the findings to labor market outcomes. This paper is the first to extend this literature to labor market outcomes and to provide estimates of the effect of childhood urban residency on income and earnings. Furthermore, the paper is the first to expand this literature beyond Africa and India and shows that similar effects are visible in a middle-income country.

Third, the paper contributes to the the urban wage premium literature following on Glaeser (1999) and Glaeser, Kolko and Saiz (2001). Most of the subsequent research has utilized individual-fixed effect models to study the effect of adulthood urban residency on earnings (Combes, Duranton and Gobillon, 2008; D'Costa and Overman, 2014; De La Roca and Puga, 2017). I complement this literature with novel estimates of the impact of childhood urban residency on adulthood wages. The estimated elasticity of childhood population density on earnings of 0.02 - 0.03 is non-trivial compared to the average elasticity of adulthood population density on earnings of 0.04 across all countries (Ahlfeldt and Pietrostefani, 2019) and the Brazil specific estimate of 0.052 (Chauvin et al., 2017), thus highlighting that childhood is an important period during which urban residency shapes labor market outcomes. Given that the previous literature has not accounted for the effects of childhood urban residency on labor market outcomes, the findings of this paper suggest that the causal effect of urban residency on earnings in developing countries might be around 50% larger than previously assumed.

The remainder of this paper continues as follows. Section 2.2 provides an overview of the literature on urban migration in the developing world. Section 2.3 outlines the empirical approach and the Brazilian census data. Section 2.4 presents the main results as well as a wide variety of robustness tests. Section 2.5 concludes.

2.2 Related literature

Economic research dating back to Harris and Todaro (1970) highlights the differences in earnings and living standards between urban and rural communities across the developing world. Although household data on income and consumption are scarce, Young (2013) using the DHS survey finds that consumption is on average 42% higher in cities compared to the countryside across 65 developing countries. More recently, Lagakos et al. (2020) and Hamory et al. (2020) using country-specific household surveys report similar differences in consumption and income for a select group of developing countries.

A key question is whether these differences in urban and rural incomes merely reflect the spatial selection of households or whether they also reflect the causal impact of urban residency. Answering this question has proven difficult, since households across the developing world select into cities based on a variety of characteristics including educational attainment (Young, 2013), household wealth (Bryan, Chowdhury and Mobarak, 2014) and cognitive ability (Hamory et al., 2020). To account for such selection, researchers have typically turned to individual fixed effect models which control for the time-invariant characteristics of individuals. Using such individual fixed effect models, Lagakos et al. (2020) and Hamory et al. (2020) find that the effect of urban residency on consumption and income declines by 64% - 91% compared to the cross-sectional estimates. The difference between the cross-sectional estimates and the individual-fixed effect estimates is usually interpreted by researchers as reflecting the spatial selection of households within the context of a Roy selection model (Young, 2013; Lagakos et al., 2020; Hamory et al., 2020).

An important shortcoming of individual-fixed effect models for earnings is that they can only be estimated for adulthood migration. To separate the change in the environment from the spatial selection of individuals, individual-fixed effect models require individuals to have worked in their region of origin prior to migration. This requirement has precluded the inclusion of childhood migrants, since they typically have no observed wages in their region of origin. Consequently, the impact of childhood urban residency on labor market outcomes has received little attention in the empirical literature.

However, this omission might be quite relevant. Previous work by Chetty, Hendren and Katz (2016) on the Moving to Opportunity program and Chetty and Hendren (2018) on migration across labor markets highlights that the benefits of migration strongly depend on the length of exposure to the new environment. Children who migrate to a better environment at a younger age seem to benefit more compared to those who move to the same environment during the teenage years or early adulthood.[2] Hence, to the degree that childhood urban residency affects adulthood labor market outcomes, the existing literature underestimates the causal effect of urban residency on labor market outcomes.[3]

Although no literature exists on the effect of childhood urban residency on adulthood labor market outcomes in developing countries, three recent studies have found effects of childhood urban exposure on educational attainment. Adukia, Asher and Novosad (2020) find that educational attainment and school attendance increased in rural Indian villages when they were connected by road to nearby cities, with the effect increasing with the returns to education in the connected cities. van Maarseveen (2022) uses an age-at-move design for urban migrants in 14 African countries and finds that childhood exposure to cities significantly raises educational attainment, school attendance

[2]Nakamura, Sigurdsson and Steinsson (2021) find a similar result in the Icelandic context following a displacement shock due to a volcanic eruption.

[3]Or phrased differently, the previous literature correctly estimated the gains of urban migration during adulthood, but incorrectly attributed the difference between the individual fixed effect estimates and the cross-sectional estimates to spatial selection insofar as childhood urban exposure matters.

and literacy rates. Finally, Giorgio et al. (2021) use the Mozambiquan civil war as an exogenous displacement shock and find increased educational attainment for children who were displaced to urban environments instead of rural environments or refugee camps.
However, the degree to which these improvements in educational outcomes lead to better labor market outcomes for childhood urban migrants remains unclear. In particular, it remains unclear how important the gains of childhood urban residency are compared to the existing estimates in the literature of effects of adulthood urban residency. Studying these two questions forms the central topic of this paper.

2.3 Methodology and data

2.3.1 Methodology

The key challenge in estimating the effect of childhood urban residency on labor market outcomes in adulthood is the spatial selection of households. To account for such selection, I utilize an age-at-move identification strategy based on Chetty and Hendren (2018). The key assumption underlying the identification strategy is that families move across regions for reasons such as the job opportunities of parents or to reduce the distance to other family members, and as such that the selection of households into more and less urbanized areas does not differ by the age of children at the time of the move. Under this assumption it is possible to estimate the effect of childhood urban exposure on labor market outcomes by comparing children who moved to and from urban areas at different ages. Empirically, the urban exposure effects are estimated as follows:[4]

$$\begin{aligned}\text{Labor market outcome}_i = \alpha_{ocs} &+ \sum_{m=1}^{16} I(m_i = m)\phi_{ocm} \\ &+ \sum_{m=1}^{16} b_m I(m_i = m)\Delta urbanization_{od} + \varepsilon_i \quad (2.1)\end{aligned}$$

Where i denotes the individual, o the region of birth, c the census year, s the year of birth, m the age at migration and d the region of residence at the time of the census. The first term of equation 2.1 captures the quality of the region of origin, which is allowed to vary by the census year and the birth cohort. The second term captures the disruption effect of moving, which may vary with the region of origin, the age at the time of the move and the census year. Finally, the third term contains the coefficient of interest, b_m, which captures the effect of moving to an area one unit more urbanized at age m.

[4]Equation 2.1 is very close to estimating equation 5 in Chetty and Hendren (2018), which they refer to as their "semi-parametric" specification.

The b_m coefficient can be separated into a selection effect and the urban exposure effect as highlighted in equation 2.2. The first term (β_m) captures the effect of spending the childhood from age m onwards in a region which is one unit more urbanized, whereas δ_m captures the selection of children into more and less urbanized regions at age m ($\frac{cov(\Delta urbanization_{od}, \varepsilon_i)}{Var(\Delta urbanization_{od})} \mid m_i = m$). The key assumption underlying the identification strategy is that the second term, δ_m which captures the selection of households into cities at age m, is uncorrelated with the age of children at the time of the move (i.e. $\delta_m = \delta$). Under this assumption the effect of spending a year at age m in an area one unit more urbanized can be inferred from equation 2.1 as $\gamma_m \equiv \beta_{m+1} - \beta_m = b_{m+1} - b_m$.

$$b_m = \beta_m + \delta_m \tag{2.2}$$

Whether this identifying assumption is also appropriate in the Brazilian context is an important question. Recent work by Chetty and Hendren (2018), Deutscher (2020) and Alesina et al. (2021) provides strong evidence that the identifying assumption holds well in the contexts of respectively the USA, Australia and Africa. Section 2.4.3 provides a variety of robustness tests to assess whether the identifying assumption also appears appropriate in the Brazilian context and whether a violation of the identifying assumption might reasonably drive the results.

2.3.2 Data

The estimation of equation 2.1 requires detailed data on the migration history and labor market outcomes for a large group of individuals. Such data is rare for developing countries since large data sets such as censuses usually do not include information on earnings and household surveys typically contain too few observations to precisely estimate equation 2.1.[5] The Brazilian census data are uniquely suited to the task since they contain detailed information on the migration history and labor force outcomes for a large sample (10%) of the population. The census data are obtained from the IPUMS International database (Minnesota Population Center, 2020) and include information on 58 million individuals spread across the 1991, 2000 and 2010 census waves. Each census contains detailed information on monthly earnings, income, hours worked, educational attainment and the migration history, in addition to a wide variety of demographic variables.

I restrict the baseline sample to individuals between the age of 25 and 29 who

[5]The inclusion of the large number of fixed effects in equation 2.1 means that households surveys such as the DHS have insufficient observations to estimate the b_m coefficients precisely. Censuses typically do have sufficient observations, although migration histories and wage outcomes are rarely included in the census questions. Brazil is only one of the five developing countries with public use census data that includes questions regarding earnings, and among these five it is the only country that also collects information on the migration history.

are born in Brazil. The majority of this group has finished their education with relatively stable labor market participation rates after the age of 25 as shown in appendix figure 2.8. At the same time, the group is relatively young which limits concerns over imprecise reporting regarding the age at migration. Nonetheless, the results are independent of the age at which outcomes are measured as shown in the robustness section. Since the censuses only include information on the most recent move across state and municipality boundaries, I restrict the sample to those who moved to their current municipality of residence during childhood. Migration events that took place within the same state are excluded to avoid including moves over very short distances.[6] The selection criteria result in a baseline sample of 267,751 individuals who moved to their current municipality of residence during childhood and for whom we observe labor market outcomes between the age of 25 - 29.

The next step is to characterize the degree of urbanization for both the current and previous place of residence. For the baseline estimates, I follow the approach of De La Roca and Puga (2017) and calculate the number of individuals living within a 25km radius for each of the 2,553 municipalities in Brazil.[7] Appendix 2.B provides more detail on the construction of the urbanization variable as well as a map of the urbanization measure. Consistent with the general pattern of migration towards urban areas across the developing world, the childhood migrants on average live in denser environments compared to the overall population of 25 - 29 year olds. The density distribution for both groups is shown in appendix figure 2.9.

2.4 Results

2.4.1 Baseline results

Figure 2.1 shows the b_m estimates obtained by estimating equation 2.1 on log earnings for the sample of childhood migrants. Two things stand out. First of all, moving to a denser environment at any point during the childhood positively affects adulthood earnings. Moving to an environment one log-point denser between the age of 12 - 16 is associated with an 8.5% gain in earnings between the ages 25 - 29. This estimate likely reflects a combination of the positive selection of households into urban environments as well as the benefit of access to a denser labor market. Second, the gains of migrating to an urban environment increases with the duration of the childhood exposure. Children who migrate to an environment one log-point denser before the age of 6 earn

[6]This does of course leave the possibility that some short-distances moves across state boundaries are included in the sample. However, the results are unchanged when excluding migration events between neighboring states.

[7]There are officially around 5,500 municipalities in Brazil. To preserve anonymity, the publicly released census data merge neighboring municipalities with fewer than 20.000 inhabitants to ensure that each municipality reported in the census contains at least 20.000 individuals.

around 2 - 3 percent more at ages 25 - 29 compared to children who make the same move after the age of 12. Under the assumption that the selection of households into more and less urban environments is indeed uncorrelated with the age of children at the time of migration, this difference can be interpreted as the causal effect of spending the childhood before age 12 in an environment one log-point denser. The effects are similar when using log income instead of log earnings as shown in appendix figure 2.12.

The next question is whether the increase in earnings is driven by a change in

Figure 2.1. Estimates of b_m for log earnings

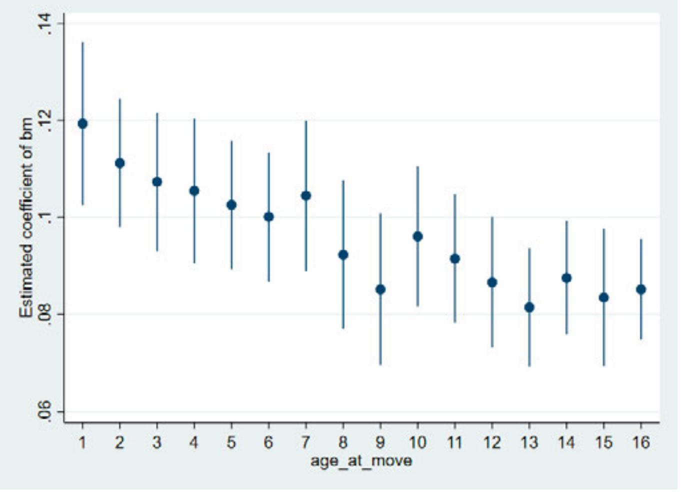

Note:Figure displays the estimates of b_m obtained by estimating equation 2.1 on log earnings at ages 25 - 29 for the sample of 267.751 movers described in section 2.3.2. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

labor force participation or whether it is driven by an increase in hourly wages. The left panel of figure 2.2 shows the b_m estimates for log hourly wage and the right panel the estimates for hours worked. The effect of childhood urban exposure on the hours worked is relatively small, with children who move to an urban environment before the age of 7 working around 0.1 hours per week more compared to children who move after the age of 7. Given the mean hours worked of 44.2, this implies a 0.23% increase in hours worked, which is an order of magnitude too small to explain the earnings increase. There is a similarly small (and statistically insignificant) effect of childhood urban exposure on the probability of employment, as shown in appendix figure 2.13.

Instead, the increase in earnings appears mainly driven by an increase in hourly wages as shown in the left panel of figure 2.2. Moving to an environment one log-point denser before the age of 6 increases hourly wages by around 2% compared to making the same move after the age of 12. This increase in hourly wages is very similar to the increase in earnings shown in figure 2.1. To provide a better sense of the magnitude and statistical significance of the various urban exposure effects, appendix 2.C.1 presents the urban exposure estimates under two different parameterizations. The urban exposure effects are similar to those inferred from the figures, with the childhood urban exposure effects for incomes, earnings and wages significant at the 0.1% level under both parameterizations.

Figure 2.2. Estimates of b_m for log hourly wage (left) and hours worked (right).

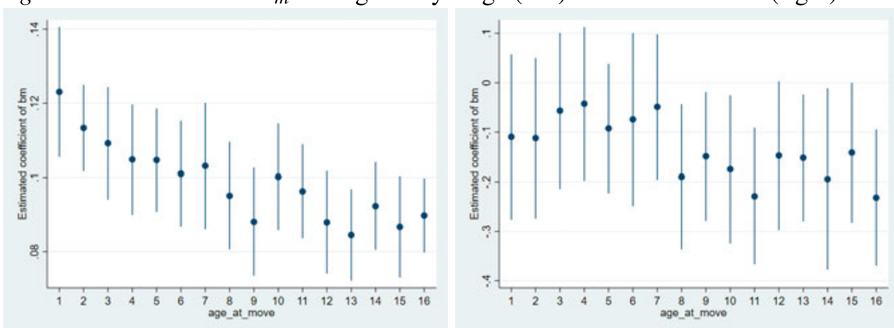

Note: Figure displays the estimates of b_m obtained by estimating equation 2.1 on the log of hourly wages (left) and hours worked (right) at ages 25 - 29 for the sample of 267.751 movers described in section 2.3.2. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

2.4.2 Effect on human capital accumulation

There is a variety of channels through which childhood urban residency could affect earnings, including through health outcomes, network effects and educational attainment. Although information on networks and health outcomes is unfortunately not available from the census, information on educational attainment is readily available. Figure 2.3 below displays the urban exposure estimates for primary school completion (left panel) and years of schooling (right panel). There is again clear evidence of positive selection of households into cities, with children who migrate to a denser environment at any age being more likely to have completed primary school and to have obtained more years of education.

At the same time, the effect on educational attainment strongly depends on the length of exposure to a more urban environment. Moving to an environment one log-point denser before the age of 6 rather than after the age of 12 raises primary school completion by around 1.5 percentage points (mean 51%) and the years of schooling by around 0.2 years (mean 6.7 years). The urban exposure effects are largely concentrated between the age of 6 and 12, with little evidence that an additional year of exposure before age 6 or after age 12 affects educational attainment. Appendix 2.C.2 shows the urban exposure estimates of figure 2.3 under two different parametarizations and finds that the estimates are similar to those derived from the figure and statistically signif-

Figure 2.3. Estimates of b_m for primary school completion (left) and years of schooling (right).

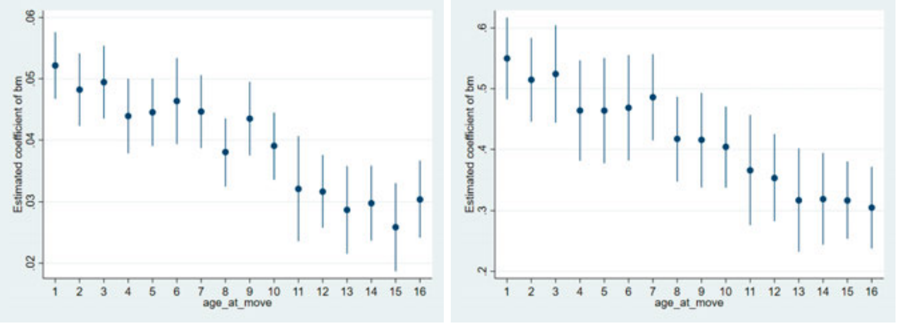

Note: Figure displays the estimates of b_m obtained by estimating equation 2.1 on years of schooling at ages 25 - 29 for the sample of 267.751 movers described in section 2.3.2. The years of schooling was not part of the 2010 census and the estimates in the right figure are thus based on the 1991 and 2000 census. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

icant at the 0.1% level in both cases. Furthermore, the effects on secondary school completion are virtually identical to the results for primary school completion, as shown in figure 2.14 and table 2.7 in the appendix.

Using a standard 10% return to years of education in line with estimates for Brazil (Tannen, 1991; Brotherhood, Ferreira and Santos, 2019) implies that the increase in educational attainment can explain a large part of the increase in wages observed in figure 2.2. This is indeed the case. Appendix 2.C.3 shows the childhood urban exposure effects on income, earnings and wages when controlling for the years of education. The urban exposure effects estimates remain statistically significant in all three cases when controlling educational attainment, but the point estimate decreases by around half. This in fact might be an underestimate of the importance of educational attainment in explaining the urban exposure effects insofar as the quality of education differs between urban and rural environments. Hence, educational attainment appears to be an important channel in explaining the effect of childhood urban exposure on earnings.

2.4.3 Robustness

The results above suggest that childhood urban exposure increases earnings, wages and educational attainment. This section discusses the robustness of the main results to four potential concerns: (1) the dependence on the urban definition, (2) selective migration during adulthood and (3) different selection into cities by the age of children and (4) mismeasurement in the region of origin.

Urban definition

One concern is that the results depend on the selected urban definition.[8] An alternative to the baseline urbanization measure is to use population density, in which case $\Delta urbanization$ in equation 2.1 is defined as the difference in population density between the current and previous region of residence.[9] Figure 2.4 shows the estimates of b_m for log earnings (left panel) and primary school completion (right panel) when using population density with very similar estimates to the baseline results.

A second possibility is to follow the urban status of the households provided in the census. Whether a municipality is urban or rural is determined by the municipalities themselves and hence potentially endogenous, but might nonetheless be an accurate representation. The median municipality claims to be around 68% urban, which rises to 85% when the municipalities are weighted by population. Figure 2.5 below shows the estimates when the urbanization measure is based on the share of households living in an urban area according to the census. The urban exposure effects are again clearly visible for both the earnings and educational outcomes.

Figure 2.4. Estimates of b_m for log earnings (left) and primary school completion (right) when using log population density as urbanization measure.

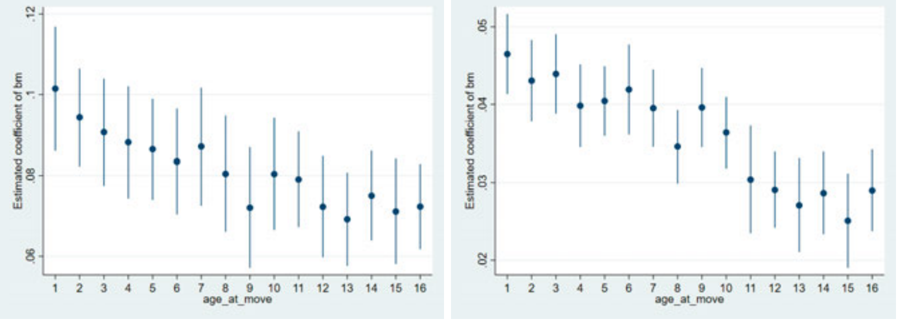

Note: Figure displays the estimates of b_m obtained by estimating equation 2.1 on the log of hourly wages (left) and primary school completion (right) at ages 25 - 29 for the sample of 267.751 movers described in section 2.3.2. The urbanization measure is based on the log of the population density. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

Finally, a significant number of the migration events are to and from Brazil's three largest cities: Brasília, Rio de Janeiro and São Paulo. Figure 2.6 shows the results when analyzing childhood exposure to the federal states containing the cities of São Paulo, Rio de Janeiro and Brasília, using moves to and from these states to estimate the urban exposure effects.[10] The exposure effects are

[8]For a recent overview of density measures, see Henderson, Nigmatulina and Kriticos (2021).

[9]Or rather, the log of the change in population density as density is far from normally distributed.

[10]In all three cases, the metropolitan area of the city contains the vast majority of the state its population. The municipalities of Rio de Janeiro and São Paulo only contain a small percentage the metropolitan area population.

Figure 2.5. Estimates of b_m for log earnings (left) and primary school completion (right) based on the census urban definition.

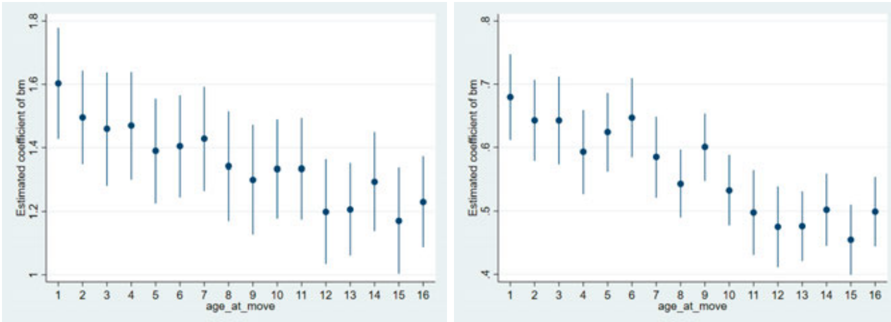

Note: Figure displays the estimates of b_m obtained by estimating equation 2.1 on the log of hourly wages (left) and primary school completion (right) at ages 25 - 29 for the sample of 267.751 movers described in section 2.3.2. The urbanization measure is based the share of households in a municipality that live in an urban region. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

again clearly visible with childhood exposure to the cities of Brasília, Rio de Janeiro and São Paulo significantly increasing adulthood earnings and primary school completion. Appendix 2.C.4 shows the urban exposure estimates for the three alternative urbanization measures under two different parameterizations, with the estimates significant at the 0.1% level in all cases. Thus, the childhood urban exposure effects are not dependent on the selected measure of urbanization, but visible across a variety of characterizations.

Figure 2.6. Estimates of b_m for log earnings (left) and primary school completion (right) using exposure to São Paulo, Rio de Janeiro and Brasília.

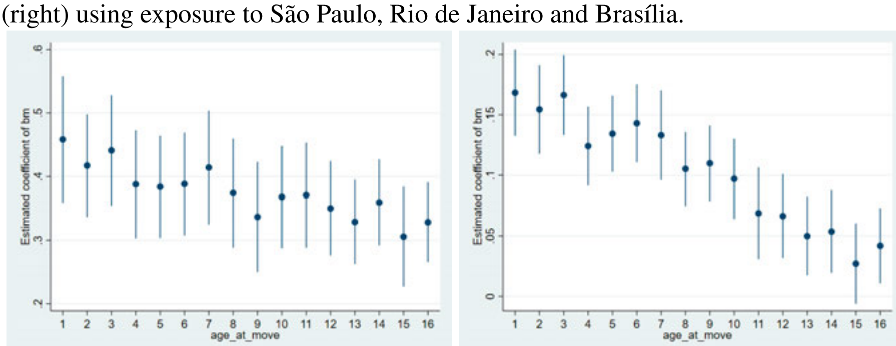

Note: Figure displays the estimates of b_m obtained by estimating equation 2.1 on the log of hourly wages (left) and primary school completion (right) at ages 25 - 29 for the sample of 267.751 movers described in section 2.3.2. The exposure effects are solely based on exposure to the states of São Paulo, Rio de Janeiro and Distrito Federal, with $\Delta urbanization_{od}$ in equation 2.1 defined as (I(Current state of residence = São Paulo/Rio de Janeiro/Distrito Federal) − I(Previous state of residence = São Paulo/Rio de Janeiro/Distrito Federal)). Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

Selection into adult migration

The second concern is that the censuses only record information for the most recent migration event across state and municipal boundaries. Hence, the age-at-move design can only be estimated for individuals who moved to their current region of residence during childhood and did not move again afterwards. One concern is that migration decisions during early adulthood may depend on the initial destination during the childhood, which would result in a selective sample of observed childhood migrants. This would for instance be the case if the most ambitious individuals who moved to a rural region during childhood are more likely to migrate again during early adulthood, in which case the observed rural childhood migrants would be negatively selected.

One way to assess the potential impact of selective adulthood migration is to analyze whether the childhood urban exposure estimates depend on the age at which outcomes are observed. If selective adulthood migration drives the results, then we should expect the childhood urban exposure effects to only gradually appear throughout adulthood. To investigate this, I estimate the model separately for each age group between ages 17 and 29. To facilitate comparison between the age groups, I impose a linear functional form on the urban exposure estimates as shown in equation 2.3. The first and second terms of equation 2.3 are the same as equation 2.1 and capture the quality of the region of origin and the disruption effects of moving. The third term functions as an intercept, with θ capturing the effect of moving to a place one log point more urban between ages 13 to 16. Finally, the γ coefficient in the fourth term is the coefficient of interest and captures the slope of the urban exposure effects between ages 1 - 12.[11]

$$\text{Individual outcome}_i = \alpha_{ocs} + \sum_{m=1}^{16} I(m_i = m)\phi_{ocm}$$
$$+ \theta * \Delta urbanization_{od} + \gamma \sum_{m=1}^{12} I(m_i = m)\frac{12 - m_i}{12}\Delta urbanization_{od} + \varepsilon_i \quad (2.3)$$

Figure 2.7 displays the γ estimates obtained by estimating equation 2.3 on log hourly wage and primary school completion by age group. The childhood urban exposure effects for both wages and primary school completion are present and statistically significant at all ages from age 17 onward. Given that the urban exposure effects are visible already during early adulthood, it appears unlikely that selective migration during adulthood is driving the results.

Selection into cities

A different concern is that the selection of households into cities may change depending on the age of children at the time of the move, thus violating the

[11] The results are virtually identical when imposing a piecewise linear functional form as shown in appendix 2.E.1.

Figure 2.7. Estimates of γ for log hourly wage (left) and primary school completion (right) by age for the ages 17 - 29.

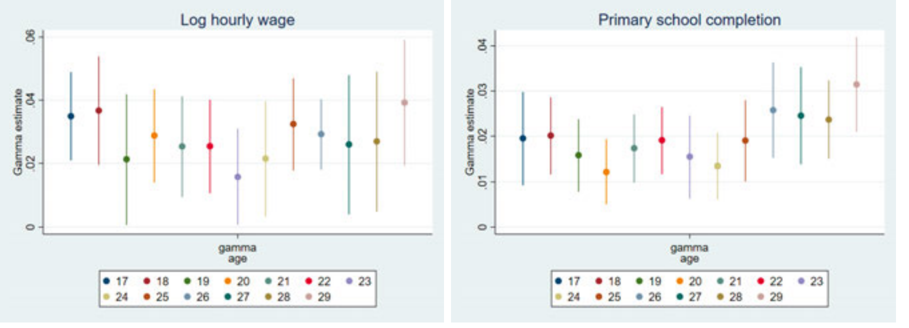

Note: The figure shows the estimates of γ obtained by estimating equation 2.3 by age for those aged 17 - 29 at the time of the census. The sample is restricted to those who live in their current municipality since childhood, have been born in brazil and whose last move was across state boundaries between age 1 and 16. The left panel displays the γ estimates for log hourly wage, whereas the right panel has primary school completion as outcome variable. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

identification assumption. It is difficult to test directly for such selection among the baseline sample of 25 - 29 year olds, since we observe very few characteristics for this group which are not potentially themselves affected by childhood urban exposure. Therefore, I focus instead on the group aged 17 - 20 to investigate the differences in selection by the age of moving. The majority of this group (65%) still lives with their parents, which allows me to control for various individual and household characteristics. At the same time, around half of this group (51%) is employed and thus has observed wages, with similar urban exposure estimates as the other age groups as shown in the previous section.
Table 2.1 below shows the effect of controlling for various individual and household characteristics on the urban exposure estimates for log wages and primary school completion. The first and fourth column show the coefficients obtained by estimating the baseline model of equation 2.3 for the group of 17 - 20 year olds. The second and fifth column incorporate a variety of observed individual level controls, such as ethnicity, gender and religion, into the fixed effect term of equation 2.3. The third and sixth column also include the years of education and literacy of the household head in the fixed effect term. Controlling for individual characteristics and the household head educational attainment does not significantly change the urban exposure estimates for log wages and primary school completion, with the estimates statistically significant at the 0.1% level in all cases.[12][13] This suggests that different selection

[12]The results are very similar when instead analyzing earnings and years of schooling.

[13]The γ estimate for primary school completion is somewhat smaller than those shown in figure 2.7. The reason is that the sample in table 2.3 is self-selective in that the group still lives with their parents, which is far more common when children still attend school and thus have completed primary school, which limits the variation within this group.

based on household characteristics by the age of migration is unlikely to drive the results.

An different concern is that there might be selection based on unobserved household characteristics. To test whether selection on unobserved characteristics might explain the urban exposure estimates, I use the adjustments of Oster (2019) and use the selection based on observed characteristics as a guide for the selection based on unobserved characteristics. The question is what value for the max R^2 to use for the Oster (2019) adjustment. Following the suggestion of Oster (2019), I select the R^2 of 0.62 obtained from sibling regressions in the Brazilian data, which should provide an upper-bound for the variance that can be explained by unobserved household characteristics.

Using the Oster (2019) correction for the primary school estimates changes little, since the inclusion of the control variables actually increases the urban exposure estimates in columns (5) and (6). However, the inclusion of control variables does slightly reduce the urban exposure estimates for log wages as shown in column (3). Using the assumption of equal selection based on observed and unobserved characteristics together with the maximum R^2 of 0.62 results in an adjusted estimate for log wages of 1.62, which remains statistically significant at the 0.1% level.[14][15]

Mismeasurement in region of origin

Finally, while the information on the current place of residence is very detailed as shown in appendix figure 2.10, the previous place of residence is based on the state of birth. This means that individuals may originate from very different parts on the state of origin depending on their age, and introduces some mismeasurement in case children move repeatedly during childhood. Section 2.E.2 in the appendix uses the fact that the censuses record more detailed information on the region of origin for those who moved in the nine years prior to the census. Neither issue significantly affects the estimates.

2.4.4 Heterogeneity

The heterogeneity analyses reveal some additional findings of interest. Table 2.2 shows the urban exposure estimates interacted with gender and ethnicity. The childhood urban exposure effects are similar by gender with no significant differences, suggesting that boys and girls benefit equally much from growing up in an urban environment. With regards to ethnicity, the urban exposure effects are slightly smaller for non-white children, suggesting that they benefit less from childhood urban exposure. This may also partially reflect that non-white migrant households tend to end up in worse neighborhoods within urban

[14]The adjusted estimate is similar when using the simple adjustment provided at the end of section 3.2 in Oster (2019) or when using the unrestricted estimator discussed in section 3.3 of Oster (2019)

[15]An alternative is to use the recent adjustments of Finkelstein, Gentzkow and Williams (2021) which specifically adjusts for spatial selection based on unobserved characteristics. The adjusted coefficient in this case is slightly larger (1.70).

Table 2.1. *Robustness of γ estimates for log wages and primary school completion to the inclusion for individual and household controls.*

	Log wages			Primary school completion		
	(1)	(2)	(3)	(4)	(5)	(6)
γ	0.0251***	0.0252***	0.0224***	0.00998**	0.0109***	0.0111***
	(0.00558)	(0.00540)	(0.00516)	(0.00311)	(0.00308)	(0.00273)
delta	0.0771***	0.0757***	0.0682***	0.0294***	0.0243***	0.0124***
	(0.00509)	(0.00492)	(0.00574)	(0.00271)	(0.00257)	(0.00152)
N	86853	86853	86853	186234	186234	186234
R^2	0.141	0.194	0.327	0.195	0.253	0.399
Indv ctrl.	No	Yes	Yes	No	Yes	Yes
HH. ctrl.	No	No	Yes	No	No	YSes

Note: The table displays the estimates of γ obtained by estimating equation 2.3 on log wages and primary school completion for the sample of 17 - 20 year olds who live with their parents and last moved during childhood across state boundaries. Columns (1) and (4) display the baseline coefficient, columns (2) and (5) include individual controls (gender, ethnicity and religion) in the fixed effect terms of equation 2.3. Columns (3) and (6) in addition include the educational attainment of the household head and literacy of the household head in the fixed effects term. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

communities, with non-white children in urban areas being less likely to live in a home that is connected to electricity, piped water or the sewage system. However, the childhood urban exposure estimates do remain statistically significant and meaningful in economic terms for both groups.

On the other hand, the results change quite substantially when the analyses are repeated for the individual census waves as shown in table 2.3. The urban exposure effects for earnings and wages are statistically significant in all three censuses, but the effect declines by around $2/3^{th}$ between 1991 and 2010. A similar decline is visible in the urban exposure estimates for primary school completion and years of schooling as shown in table 2.4.[16] Hence, while the urban exposure effects remain statistically significant for all census years, the effects are significantly smaller in 2010 compared to 1991.

An interesting question is whether the decline in the childhood exposure effects over time is specific to Brazil or whether this reflects a general trend. Brazil implemented extensive social reforms between 1991 and 2010, which appear to have somewhat favored rural communities (Higgins, 2012). At the same time, the decline of the urban exposure effects is also consistent with cross-sectional results, with richer countries typically having far smaller urban - rural gaps in income and educational outcomes compared to developing countries. Unfortunately, is impossible to distinguish between the two expla-

[16] And finally, a similar effect is also visible for secondary school completion as shown in table 2.12 in the appendix.

nations within a single case study. Investigating whether the decline in the urban exposure effects is due to specific investments in rural communities during this period in Brazil or whether it reflects a more general pattern as countries develop has important implications for policy and thus forms an interesting avenue for future research.

Table 2.2. *Heterogeneity by gender and ethnicity*

	Log hourly wage		Primary school completion	
	(1)	(2)	(3)	(4)
γ	0.0217**	0.0337***	0.0286***	0.0287***
	(0.00662)	(0.00904)	(0.00442)	(0.00448)
$\gamma * female$	0.0125		-0.00638	
	(0.00767)		(0.00387)	
$\gamma * non-white$		-0.00896		-0.00773*
		(0.00769)		(0.00340)
N	162565	161650	239635	238644
R^2	0.121	0.146	0.152	0.171

The table displays the estimates of γ obtained by estimating equation 2.3 on log wages and primary school completion for the baseline sample of 25 - 29 year olds. All terms in equation 2.3 have been interacted with gender for the odd columns and with ethnicity (white vs non-white) for the even columns. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

Table 2.3. *Estimates of γ for log earnings and log hourly wage by census year.*

	Log earnings			Log hourly wage		
Year	1991	2000	2010	1991	2000	2010
γ	0.0684***	0.0382***	0.0219**	0.0656***	0.0369***	0.0203*
	(0.0114)	(0.00683)	(0.00730)	(0.0106)	(0.00632)	(0.00869)
delta	0.105***	0.0851***	0.0577***	0.112***	0.0869***	0.0593***
	(0.0107)	(0.00820)	(0.00563)	(0.00966)	(0.00679)	(0.00596)
N	50656	60343	51661	50656	60343	51661
R^2	0.132	0.104	0.089	0.133	0.106	0.071

The table displays the estimates of γ obtained by estimating equation 2.3 on log earnings (columns 1-3) and log hourly wage (columns 4-6) for the baseline sample of 25 - 29 year olds by census year. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

Table 2.4. *Estimates of γ for primary school completion and years of schooling by census year.*

	Primary school completion			Years of schooling		
Year	1991	2000	2010	1991	2000	2010
γ	0.0413***	0.0261***	0.0132*	0.313***	0.180***	0.108*
	(0.00560)	(0.00509)	(0.00576)	(0.0557)	(0.0319)	(0.0438)
delta	0.0306***	0.0319***	0.0280***	0.323***	0.297***	0.210***
	(0.00727)	(0.00495)	(0.00537)	(0.0634)	(0.0426)	(0.0506)
N	76385	91636	71672	76467	92403	67595
R^2	0.091	0.085	0.060	0.125	0.120	0.087

The table displays the estimates of γ obtained by estimating equation 2.3 on primary school completion and years of schooling for the individual census years. The years of schooling is not available for the 2010 census and I instead infered based on the highest completed education. To improve comparability between census years, I also replaced the 1991 and 2000 years of schooling with the estimated years of schooling based on educational attainment. The effect of this adjustment on the coefficients for the 1991 and 2000 census are minor. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

2.5 Discussion and conclusion

This paper investigates the effects of childhood urban residency on labor market outcomes in Brazil. Using an age-at-move design combined with Brazilian census data, I find that spending the childhood in an environment one log-point denser increases earnings at ages 25 - 29 by 2 - 3 percent. The earnings increase is primarily driven by an increase in hourly wages rather than a change in hours worked or labor market participation. The childhood urban exposure effects are also visible for educational outcomes, with the increase in educational attainment accounting for more than half of the earnings increase. Spending the childhood in an environment one log-point denser raises primary school completion by 1.5 percentage points (mean 51%) and the years of schooling by around 0.2 years (mean 6.7 years). The urban exposure estimates do not dependent on the measure of urbanization and are not driven by selective adulthood migration or different selection into urban migration by the age of the move.

The findings of this paper complement the previous literature on rural to urban migration in developing countries by estimating the impact of childhood urban exposure. Quantitatively, the gains of spending the childhood in an urban environment are similar, and in addition to, the wage effects of rural to urban migration during adulthood, which are estimated to be around 5% in the case of Brazil (Alvarez, 2020). Furthermore, the elasticity of wages with respect to childhood density of 0.02 - 0.03 is only somewhat smaller than the elasticity of wages with respect to adulthood density, which is estimated to be 0.04 by

Ahlfeldt and Pietrostefani (2019) generally and 0.054 for Brazil specifically by Chauvin et al. (2017).

These previous estimates have exclusively used adulthood migrants for identification, and thus have been unable to estimate the effects of childhood urban residency on economics outcomes. The findings of this paper suggest that the previous literature has thus missed an important part of the explanation of why urban residents are wealthier. Furthermore, this paper is the first to extend the literature on the effects of childhood urban residency on educational attainment beyond the African and Indian context and to show that similar educational gains from urban residency are visible in the context of a middle income country.

Finally, a key challenge for future research is to decompose the effect of childhood urban exposure on labor market outcomes into the individual mechanisms. This paper shows that educational attainment appears to be an important factor to explain the positive effects of childhood urban residency on wages, but it remains unclear which mechanisms drive this increase. Urban and rural areas differ on a wide bundle of factors, such as the returns to education, opportunity costs of education, household income, quality of peers and access to education. Finding exogenous sources of variation within each of these factors to causally identify their impact is beyond the scope of this paper. However, understanding which factors are the main drivers of the childhood urban exposure effect will further enhance our understanding of the welfare effects of rural to urban migration and urbanization in the developing world.

Acknowledgements

The author wishes to acknowledge the statistical offices that provided the underlying data making this research possible: Institute of Geography and Statistics, Brazil.

References

Adukia, Anjali, Sam Asher, and Paul Novosad. 2020. "Educational investment responses to economic opportunity: evidence from Indian road construction." *American Economic Journal: Applied Economics*, 12(1): 348–376.

Ahlfeldt, Gabriel M, and Elisabetta Pietrostefani. 2019. "The economic effects of density: A synthesis." *Journal of Urban Economics*, 111: 93–107.

Alesina, Alberto, Sebastian Hohmann, Stelios Michalopoulos, and Elias Papaioannou. 2021. "Intergenerational mobility in Africa." *Econometrica*, 89(1): 1–35.

Alvarez, Jorge A. 2020. "The agricultural wage gap: Evidence from brazilian micro-data." *American Economic Journal: Macroeconomics*, 12(1): 153–173.

Brotherhood, Luiz Mário, Pedro Cavalcanti Ferreira, and Cezar Santos. 2019. "Education quality and returns to schooling: evidence from migrants in Brazil." *Economic Development and Cultural Change*, 67(3): 439–459.

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. 2014. "Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh." *Econometrica*, 82(5): 1671–1748.

Chauvin, Juan Pablo, Edward Glaeser, Yueran Ma, and Kristina Tobio. 2017. "What is different about urbanization in rich and poor countries? Cities in Brazil, China, India and the United States." *Journal of Urban Economics*, 98: 17–49.

Chetty, Raj, and Nathaniel Hendren. 2018. "The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects." *The Quarterly Journal of Economics*, 133(3): 1107–1162.

Chetty, Raj, Nathaniel Hendren, and Lawrence F Katz. 2016. "The effects of exposure to better neighborhoods on children: New evidence from the Moving to Opportunity experiment." *American Economic Review*, 106(4): 855–902.

Chiovelli, Giorgio, Stelios Michalopoulos, Elias Papaioannou, and Sandra Sequeira. 2021. "Forced Displacement and Human Capital: Evidence from Separated Siblings." National Bureau of Economic Research.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon. 2008. "Spatial wage disparities: Sorting matters!" *Journal of urban economics*, 63(2): 723–742.

D'Costa, Sabine, and Henry G Overman. 2014. "The urban wage growth premium: Sorting or learning?" *Regional Science and Urban Economics*, 48: 168–179.

De La Roca, Jorge, and Diego Puga. 2017. "Learning by working in big cities." *The Review of Economic Studies*, 84(1): 106–142.

Deutscher, Nathan. 2020. "Place, Peers, and the Teenage Years: Long-Run Neighborhood Effects in Australia." *American Economic Journal: Applied Economics*.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams. 2021. "Place-based drivers of mortality: Evidence from migration." *American Economic Review*, 111(8): 2697–2735.

Giorgio, Chiovelli, Michalopoulos Stelios, Papaioannou Elias, and Sequiera Sandra. 2021. "Forced Displacement and Human Capital Evidence from Separated Siblings."

Glaeser, Edward L. 1999. "Learning in cities." *Journal of urban Economics*, 46(2): 254–277.

Glaeser, Edward L, Jed Kolko, and Albert Saiz. 2001. "Consumer city." *Journal of economic geography*, 1(1): 27–50.

Hamory, Joan, Marieke Kleemans, Nicholas Y Li, and Edward Miguel. 2020. "Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata." *Journal of the European Economic Association*.

Harris, John R, and Michael P Todaro. 1970. "Migration, unemployment and development: a two-sector analysis." *The American economic review*, 60(1): 126–142.

Henderson, J Vernon, Dzhamilya Nigmatulina, and Sebastian Kriticos. 2021. "Measuring urban economic density." *Journal of Urban Economics*, 125: 103188.

Higgins, Sean. 2012. "The Impact of Bolsa Família on Poverty: Does Brazil's Conditional Cash Transfer Program Have a Rural Bias?"

Lagakos, David, Samuel Marshall, Ahmed Mushfiq Mobarak, Corey Vernot, and Michael E Waugh. 2020. "Migration costs and observational returns to migration in the developing world." *Journal of Monetary Economics*, 113: 138–154.

Minnesota Population Center. 2020. "Integrated Public Use Microdata Series, International: Version 7.3 [dataset]. Minneapolis, MN: IPUMS, https://doi.org/10.18128/D020.V7.2."

Nakamura, Emi, Jósef Sigurdsson, and Jón Steinsson. 2021. "The Gift of Moving: Intergenerational Consequences of a Mobility Shock." *The Review of Economic Studies*, rdab062.

Oster, Emily. 2019. "Unobservable selection and coefficient stability: Theory and evidence." *Journal of Business & Economic Statistics*, 37(2): 187–204.

Tannen, Michael B. 1991. "New estimates of the returns to schooling in Brazil." *Economics of Education Review*, 10(2): 123–135.

van Maarseveen, Raoul. 2022. "The Effect of Urban Migration on Educational Attainment: Evidence from Africa." *Available at SSRN 3836097*.

Young, Alwyn. 2013. "Inequality, the urban-rural gap, and migration." *The Quarterly Journal of Economics*, 128(4): 1727–1785.

Appendices

2.A Sample selection

Figure 2.8. Labor force participation by age

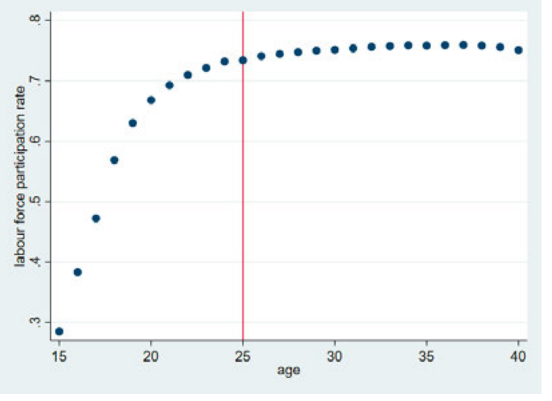

Note: Figure shows the labor force participation by age for the group aged 15-40. No labor market participation is recorded in the census for children below the age of 15. Individuals are part of the labor force if they are either engaged in work or have been looking for work during the last 2 months.

Figure 2.9. Density distribution of all 25 - 29 year olds (left panel) and the mover sample (right panel).

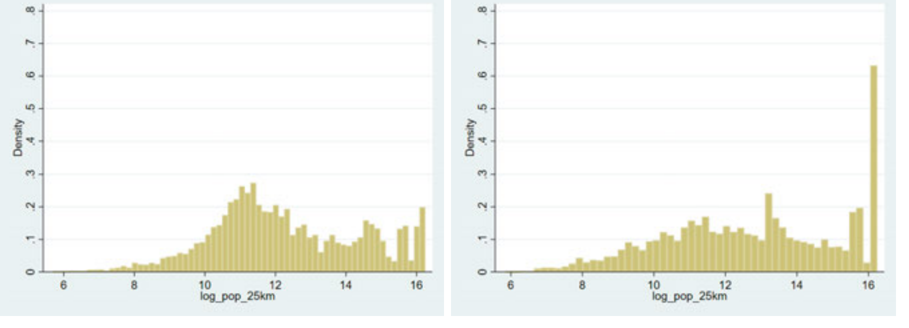

Note: Figure displays the density distribution, measured as the log of the number of individuals living within a 25km radius, for all 25 - 29 year olds (left) and the mover sample described in section 2.3.2 (right). The spike at 16 is due to the municipalities in and around the cities of São Paulo, Rio de Janeiro and Brasília.

2.B Urbanization measure

There are many different ways to characterize the degree of urbanization or density of a given region, as recently discussed by Henderson, Nigmatulina and Kriticos (2021). In the baseline urbanization measure, I define the urbanization measure as the log of the number of individuals within a given distance following De La Roca and Puga (2017). For each of the 2,422 municipalities, I draw a 25km circle around the centroid of the municipality and calculate the number of individuals within this 25km radius, assuming a constant population density within municipalities. The average municipality has 6.5 municipalities within the 25km radius, although for the very rural municipalities in the West, the only municipality included in the 25 km radius is often the municipality itself.[17]

The resulting urbanization map is displayed in figure 2.11. The highest density areas are on the East coast around the cities of Rio de Janeiro, São Paulo and Porto Alegre and in the center of Brazil around Brasília. As would be expected, the lowest density areas are in the Western parts of the country in the Amazon basin. Despite the very large size of some of the municipalities in particular in the State of Amazonas, the total number of inhabitants is usually quite low.

In contrast to the high geographic detail of the current place of residence, the previous place of residence is only observed at the level of the state of birth. Within each state, I take the average of the urbanization measure of the municipalities contained within the same census (weighted by population) and assign this value as the urbanization measure of the state of birth. The lack of detail on the previous place of residence raises some concerns regarding repeat migration and selection within the state of origin. Appendix 2.E.2 discusses these concerns and shows that they do not significantly affect the baseline estimates.

[17]This is more generally the case in 5% of the municipalities across Brazil.

Figure 2.10. Map of municipalities of Brazil in 2010

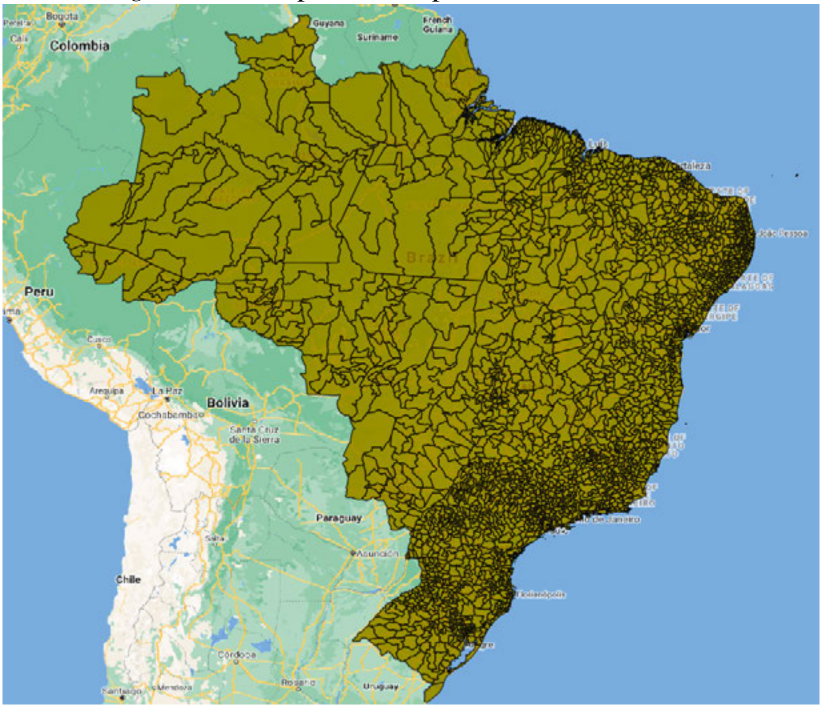

Figure 2.11. Urbanization measure by municipality for Brazil in 2010

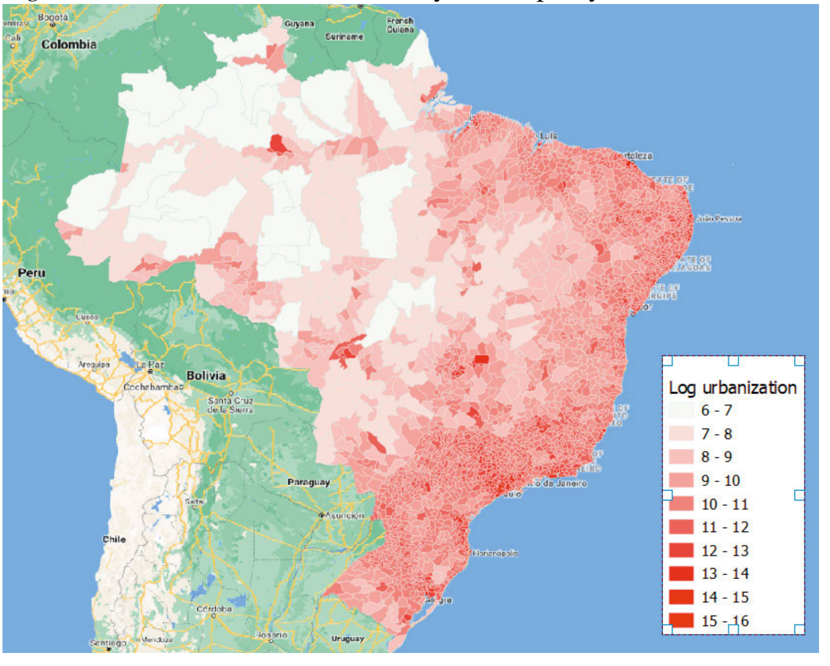

Note: Figure displays the urbanization measure for the 2422 municipalities based on the 2010 census population and municipality boundaries. For the purpose of the figure, the municipalities are divided into ten bins of 1 log-point each.

2.C Parameterization of main effects

This section puts some functional form assumption on the results discussed in section 2.4 to test for statistical significance of the urban exposure effects and to provide more detailed information on the size of the urban exposure effects. In order to do so, I impose two different functional forms. First, I impose a linear functional form on the b_m estimates, which is introduced and discussed in the main text as equation 2.3 in section 2.4.3. The linear functional form assumption appears to match well with the b_m estimates shown section 2.4.1. Second, to show that the results do not depend on the linear functional form assumption, I also impose a more flexible piecewise linear functional form assumption as shown in equation 2.4. The δ term in equation 2.4 functions as the intercept and captures the effect of migrating to an environment one log-point denser between ages 13 - 16. The two γ terms capture whether there are any additional effects from moving earlier, between ages 1 - 6 and ages 7 - 12.

$$\text{Individual outcome}_i = \alpha_{ocs} + \sum_{m=1}^{16} I(m_i = m)\phi_{ocm} + \delta \Delta urbanization_{od} \\ + \gamma_{1-6} I(m_i \leq 6)\Delta urbanization_{od} + \gamma_{7-12} I(7 \leq m_i \leq 12)\Delta urbanization_{od} + \varepsilon_i \tag{2.4}$$

2.C.1 Labor market outcomes

Table 2.5 below shows the estimates obtained by estimating the linear functional form of equation 2.3 on the various labor market outcomes. As can be seen from table 2.5, the estimates of γ for income, earnings and wages are similar in size and statistically significant at the 0.1% level in all three cases. The effect on employment is insignificant and the effect on hours worked very small (0.13 hours) which is in line with the discussions in the main text on the basis of the figures.

The estimates are similar in terms of effect size and statistical significance when instead imposing the piecewise linear function of equation 2.4 as shown in table 2.6. Income, earnings and wages are around 2 percent higher for those who move to an area one log-point denser before the age of 6 rather than after the age of 12, which is closely in line with the conclusions based on the figures in the main text. Furthermore, the effects on hours worked (0.10 hours) and employment are again fairly small and insignificant in case of employment.

Table 2.5. *Estimates for labor market outcomes using a linear functional form.*

	Log income	Log earnings	Log wage	Employed	Hours worked
γ	0.0324***	0.0327***	0.0302***	-0.000596	0.132*
	(0.00452)	(0.00451)	(0.00469)	(0.00183)	(0.0618)
delta	0.0925***	0.0850***	0.0884***	0.00913***	-0.180***
	(0.00466)	(0.00556)	(0.00534)	(0.00108)	(0.0456)
N	172111	162660	162660	240825	167900
R^2	0.897	0.909	0.898	0.019	0.035

Note: The table displays the estimates of γ obtained by estimating equation 2.3 on the baseline sample of 25 - 29 years olds for the various labour market outcomes. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

Table 2.6. *Estimates for labor market outcomes using a piecewise linear functional form.*

	Log income	Log earnings	Log wage	Employed	Hours worked
γ_{1-6}	0.0226***	0.0232***	0.0210***	-0.00139	0.0994*
	(0.00345)	(0.00342)	(0.00361)	(0.00138)	(0.0481)
γ_{7-12}	0.00664*	0.00822**	0.00662*	-0.00285**	0.0275
	(0.00301)	(0.00285)	(0.00284)	(0.00107)	(0.0402)
delta	0.0927***	0.0845***	0.0885***	0.0104***	-0.182***
	(0.00459)	(0.00545)	(0.00519)	(0.00118)	(0.0510)
N	172111	162660	162660	240825	167900
R^2	0.897	0.909	0.898	0.019	0.035

Note: The table displays the estimates of γ obtained by estimating equation 2.4 on the baseline sample of 25 - 29 years olds for the various labour market outcomes. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

2.C.2 Educational outcomes

Tables 2.7 and 2.8 repeat the functional form estimations for the various educational outcomes. The effects of childhood urban exposure on primary school completion, secondary school completion and years of schooling are statistically significant at the 0.1% level across all three outcomes under both parameterizations. Given the pattern of the estimates presented in figure 2.3, the piecewise linear functional form assumption would appear most appropriate. The regression results obtained by estimating the piecewise linear function form are again very similar to the coefficients derived from the figures in the main text.

Table 2.7. *Estimates for educational outcomes using a linear functional form.*

	Primary school degree	Secondary school degree	Years of schooling
γ	0.0249***	0.248***	0.0244***
	(0.00293)	(0.0256)	(0.00191)
delta	0.0306***	0.331***	0.0191***
	(0.00301)	(0.0350)	(0.00298)
N	239693	167780	239693
R^2	0.142	0.135	0.137

Note: The table displays the estimates of γ obtained by estimating equation 2.3 on the baseline sample of 25 - 29 years olds for the various labour market outcomes. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

Table 2.8. *Estimates for educational outcomes using a piecewise linear functional form.*

	Primary school degree	Secondary school degree	Years of schooling
γ_{1-6}	0.0188***	0.183***	0.0184***
	(0.00200)	(0.0181)	(0.00131)
γ_{7-12}	0.00958***	0.0929***	0.00912***
	(0.00132)	(0.0115)	(0.00120)
delta	0.0287***	0.314***	0.0173***
	(0.00306)	(0.0343)	(0.00292)
N	239693	167780	239693
R^2	0.142	0.134	0.137

Note: The table displays the estimates of γ obtained by estimating equation 2.4 on the baseline sample of 25 - 29 years olds for the various labour market outcomes. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

2.C.3 Urban exposure estimates when controlling for education

Finally, one question is to what degree the increase in educational attainment can explains the urban exposure effect for wages. To investigate this, I estimate the linear model of equation 2.3 while controlling for the years of education.

To control as flexibly as possible for the years of education, I include them in the first two terms of equation 2.3. Since the years of education are not included in 2010 census, I limit the analyses to the 1991 and 2000 census years. Table 2.9 below shows the γ estimates for log income, earnings and wages both with and without controls for educational attainment. In all three cases, the coefficient declines by more than 50% when educational attainment is controlled for, suggesting that education plays an important role in explaining why childhood exposure to cities raises educational attainment. The childhood urban exposure estimates do remain statistically significant when controlling for educational attainment. This implies that other factors, such as health and network effects, may also play a role. At the same time, the quality of education might also be higher in urban areas, which means that the analyses may still somewhat underestimate the true importance of educational attainment in explaining the childhood urban exposure effects.

Table 2.9. *Effect of controlling for education on urban exposure estimates.*

	Log income		Log earnings		Log wages	
γ	0.0464***	0.0203***	0.0477***	0.0217***	0.0453***	0.0170***
	(0.00537)	(0.00482)	(0.00525)	(0.00471)	(0.00524)	(0.00448)
delta	0.0959***	0.0709***	0.0955***	0.0708***	0.0996***	0.0723***
	(0.00594)	(0.00847)	(0.00611)	(0.00858)	(0.00579)	(0.00829)
N	114721	113988	111022	110329	111022	110329
R^2	0.922	0.943	0.924	0.945	0.921	0.944
Contr. for educ.?	No	Yes	No	Yes	No	Yes

Note: The table displays the estimates of γ obtained by estimating equation 2.3 on the baseline sample of 25 - 29 years olds for log income, log earnings and log wages. The uneven columns are the baseline estimates, whereas the even columns display the estimates when including the years of education in the first two terms of equation 2.3. Since years of education was not part of the 2010 census, the observations for the year 2010 have been removed from all columns to improve comparability. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

2.C.4 Urban exposure estimates for alternative urbanization measures.

Finally, one might be concerned that the urban exposure effects depend on the selected measure of urbanization. Table 2.10 shows the results when imposing the functional form of equation 2.3 on the estimates displayed in figures 2.4 - 2.6. The urban exposure estimates in columns (1) and (4) are based on the difference in log population density between the current and previous region of residence, corresponding to figure 2.4 in the main text. The estimates in columns (2) and (5) are based on the change in urban share between the current and previous place of residence according to the census, corresponding to figure 2.5 in the main text. Finally, the estimates in columns (3) and (6) are based on the exposure to the three largest Brazilian cities, corresponding to figure 2.6 in the main text. The urban exposure estimates are statistically significant at the 0.1% level across both outcome measures for all thee urbanization measures. The results are very similar and again statistically significant at the 0.1% level when instead imposing the piecewise -linear functional form of equation 2.4 on the various urbanization measures (not shown here).

Table 2.10. *Parameteric estimates for three alternative urbanization measures.*

	Log earnings			Primary school completion		
	(1)	(2)	(3)	(4)	(5)	(6)
γ	0.0271***	0.350***	0.116***	0.0201***	0.212***	0.138***
	(0.00425)	(0.0576)	(0.0268)	(0.00284)	(0.0306)	(0.0124)
delta	0.0722***	1.233***	0.335***	0.0290***	0.490***	0.0541***
	(0.00535)	(0.0644)	(0.0329)	(0.00249)	(0.0226)	(0.0150)
N	162660	162660	162660	239693	239693	239693
R^2	0.908	0.910	0.907	0.143	0.150	0.124

Note: The table displays the estimates of γ obtained by estimating equation 2.3 on log earnings and primary school completion for the baseline sample of 25 - 29 years olds. The urbanization measure in columns (1) and (4) is based on the change in log density between the current and previous region of residence corresponding to figure 2.4 in the main text. The urbanization measure in columns (2) and (5) is based on the change in the urban share between the current and previous region of residence corresponding to figure 2.5 in the main text. Finally, columns (3) and (6) are based on the moves to and from the three largest cities corresponding to figure 2.6 in the main text. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

2.D Urban exposure effects for additional outcomes

Figure 2.12. Estimates of b_m for log income

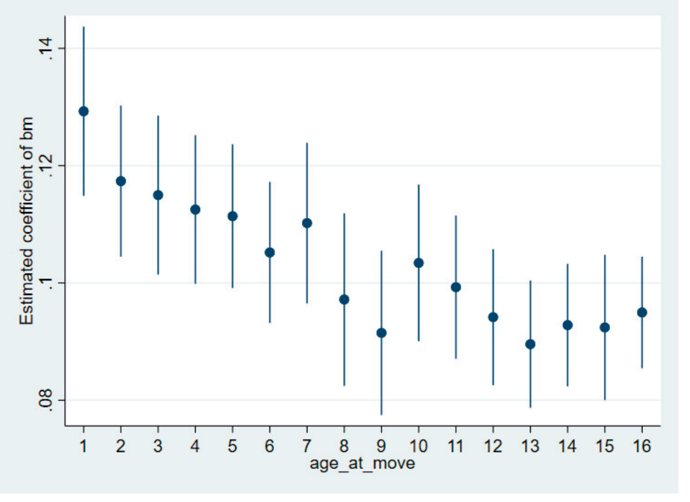

Note: Figure displays the estimates of b_m obtained by estimating equation 2.1 on the log of income for the sample of 267.751 moves described in section 2.3.2. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

Figure 2.13. Estimates of b_m for employment probability

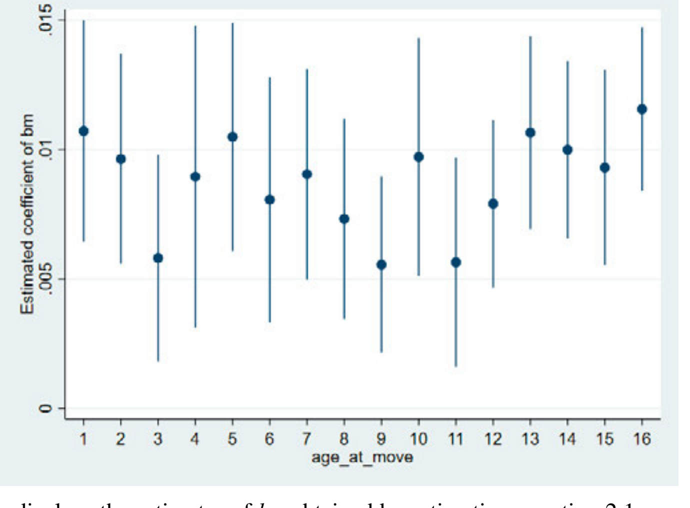

Note: Figure displays the estimates of b_m obtained by estimating equation 2.1 on employment for the sample of 267.751 moves described in section 2.3.2. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

Figure 2.14. Estimates of b_m for secondary school completion.

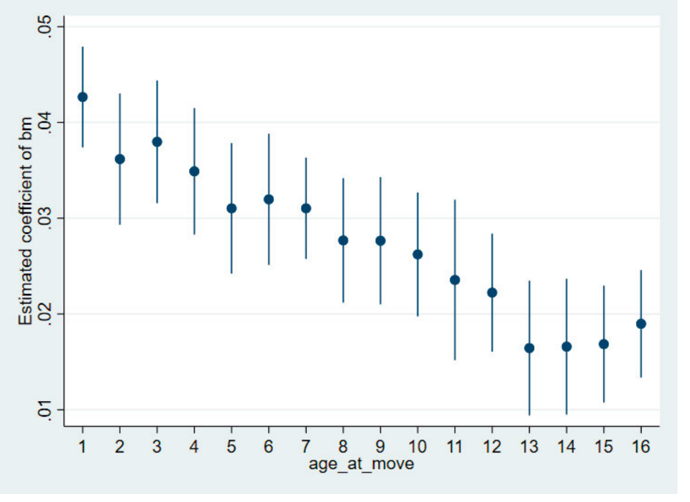

Note: Figure displays the estimates of b_m obtained by estimating equation 2.1 on primary school completion (left) and secondary school completion (right) at ages 25 - 29 for the sample of 267.751 movers described in section 2.3.2. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

2.E Further robustness tests

2.E.1 Urban exposure effects by age under piecewise linear functional form

Figure 2.15. Effects of urban exposure for primary school completion (left) and log hourly wage (right) by age under a piecewise-linear functional form assumption.

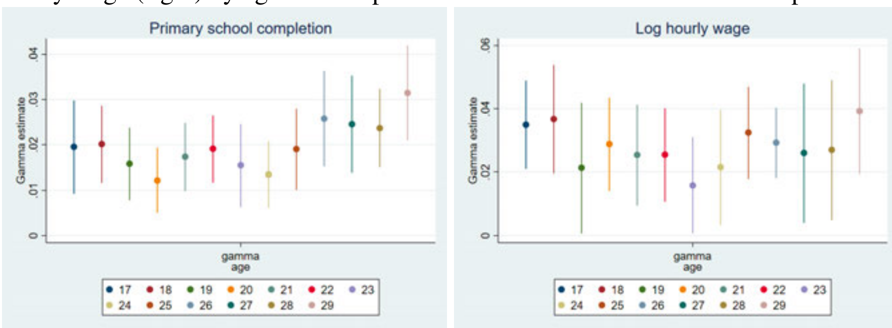

Note: The figure shows the estimates by age when imposing a piecewise linear functional form on the estimates. The estimating equation is similar to equation 2.3, but the fourth term is replaced by $\gamma_{1-6}I(\text{age at move} \leq 6) * \Delta urbanization_{od} + \gamma_{7-12}I(7 \leq \text{age at move} \leq 12) * \Delta urbanization_{od}$. The sample is restricted to those who live in their current municipality since childhood, have been born in brazil and whose last move was across state boundaries between the ages 1-16. The left panel displays the γ_{1-6} for the log of hourly wages, whereas the right panel has primary school completion as outcome variable. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

2.E.2 Mismeasurement in region of origin

One concern is that there might be mismeasurement in the region of origin. Whereas the information on the current region of residence is quite detailed as shown in figure 2.10 in the appendix, far less information is available on the previous region of residence, which is based on the state of birth. The large geographical size of some Brazilian states allows for the possibility that children could originate from different regions within the state of birth depending on the age at migration. In addition, the region of origin might be misclassified for those who moved repeatedly during childhood.

To assess whether such mismeasurement in the region of origin affects the results, I use more detailed information on the last migration event available for individuals who moved across state boundaries in the nine years prior to the census. Since this nine year window does not overlap with childhood migration for the baseline sample of 25 - 29 year olds, I focus on individuals between the ages of 15 and 17, for whom we thus have more detailed information on migration events between ages 8 - 14. Given that only few individuals at this age participate in the labor market (see appendix figure 2.8), I restrict attention to educational outcomes.

The first two columns in table 2.3 test for the mismeasurement of the region of origin due to repeat migration, using the fact that information is available for both the previous region of residence as well as the region of birth for the selected group of 15 - 17 year olds who migrated between ages 8 - 14. The first column of table 2.3 shows the estimates of equation 2.3 for primary school completion with the previous place of residence based on the place of birth, whereas column (2) restricts the sample to those for whom the state of previous residence is the same as the state of birth (86% of the sample). The estimated urban exposure effects are very similar in both cases, suggesting that mismeasurement in the region of origin due to repeated childhood migration is unlikely to drive the results.

Assessing potential selection within the region of origin is more difficult to test due to the lack of more detailed data on the region of origin. One exception is the 1991 census which includes a question for those who migrated in the last 9 years on whether the previous place of residence was urban (64%) or rural (36%). Column (3) repeats the analysis exclusively for the 1991 census, with column (4) also incorporating the urban/rural status of the previous place of residence into the fixed effect term of equation 2.3. Adding the urban/rural status of the previous place of residence does not significantly change the estimates of gamma, suggesting that selection within the region of origin is unlikely to drive the the results.

Table 2.11. *Robustness of γ estimates for primary school completion to the mismeasurement of the region of origin.*

	All census years		1991 census	
	(1)	(2)	(3)	(4)
γ	0.0216**	0.0226**	0.0386***	0.0353***
	(0.00647)	(0.00718)	(0.00747)	(0.00783)
delta	0.0161***	0.0145***	0.00957**	0.00690*
	(0.00256)	(0.00239)	(0.00349)	(0.00306)
N	92720	78781	29216	29216
R^2	0.205	0.212	0.110	0.138
Previous region same as region of birth?	No	Yes	No	Yes

The table displays the estimates of γ from estimating equation 2.3 on primary school completion for the sample of 15 - 17 year olds who moved between ages 8 - 14. Column (1) displays the baseline coefficient of equation 2.3 with the region of origin based on birth region, whereas column (2) restricts the sample to those for whom the previous region of residence is the same as the region of birth. Columns (3) and (4) restrict the sample to 1991, with column (4) incorporating the urban/rural status of the previous place of residence into the fixed effect term of equation 2.3. Note that the larger estimates for 1991 are a general pattern for the childhood urban exposure effects explored further in section 2.4.4. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

2.F Schooling estimates by census year

Table 2.12. *Estimates of γ for secondary school completion and tertiary education enrollment/completion by census year.*

	Secondary school completion			Tertiary education enrollment		
Year	1991	2000	2010	1991	2000	2010
γ	0.0281***	0.0241***	0.0185**	0.0130**	0.00982***	0.0165***
	(0.00498)	(0.00320)	(0.00518)	(0.00412)	(0.00265)	(0.00401)
delta	0.0156*	0.0176**	0.0258***	0.00936**	0.00907**	0.0124*
	(0.00646)	(0.00496)	(0.00596)	(0.00327)	(0.00310)	(0.00554)
N	76385	91636	71672	76467	92403	71955
R^2	0.075	0.077	0.065	0.055	0.059	0.074

The table displays the estimates of γ obtained by estimating equation 2.3 on secondary school completion (columns 1-3) and tertiary education enrollment/completion for the individual census years. Standard errors are twoway clustered at the region of residence and region of birth $*$ cohort level.

3. The urban-rural education gap: do cities indeed make us smarter?

This article has been published as: van Maarseveen R., The urban-rural education gap: do cities indeed make us smarter?, Journal of Economic Geography, 2021, 21(5), p 683-714, by permission of Oxford University Press.

Journal of Economic Geography (2021) **21**, 683–714 doi: 10.1093/jeg/lbaa033
Advance Access Publication 23 November 2020

The urban–rural education gap: do cities indeed make us smarter?*

Raoul van Maarseveen†

Department of Economics, Uppsala University, 752 36 Uppsala, Sweden
†Correspondence to: *raoul.vanmaarseveen@nek.uu.se*

Abstract

Despite the large urban–rural education gap observed in most countries, little attention has been paid to whether cities actually enjoy a comparative advantage in the production of human capital. Using Dutch administrative data, this paper finds that children growing in urban regions consistently attain higher levels of human capital compared with children in rural regions, conditional on observed cognitive ability and various family characteristics. The elasticity of university attendance with respect to population density is 0.07, which is robust across a variety of specifications. Hence, the paper offers a different explanation to explain the recent success of cities.

Keywords: urban–rural education gap, agglomeration economies, human capital
JEL classifications: I20, J24, R10
Date submitted: 17 June 2020 **Editorial decision:** 26 October 2020 **Date Accepted:** 4 November 2020

1. Introduction

The resurgence of the city has received widespread attention from both academics and policymakers in recent decades. Explanations typically highlight the success of cities at attracting well educated and skilled individuals from elsewhere, both by offering higher wages (see Combes et al., 2008) and superior consumption amenities (see Glaeser et al., 2001). Furthermore, cities are thought to have a comparative advantage in employing human capital in the production of goods and services, due to the sharing, matching and learning mechanisms described in Duranton and Puga (2004) and Rosenthal and Strange (2004).

A small literature has analyzed the role of cities in the production of human capital itself. Both Glaeser and Mare (2001) and De La Roca and Puga (2017) find evidence that workers in larger cities accumulate human capital at a faster rate. However, little attention has been paid to the role of cities in human capital accumulation in the period prior to labor market entry. This is surprising, since agglomeration economies are likely to

* The author would like to express his gratitude to the editor and two anonymous referees, as well as Niklas Bengtsson, Cristina Bratu, Matz Dahlberg, Peter Fredriksson, Edward Glaeser, Matthew Lindquist, Björn Öckert, Albert Saiz, Anna Tompsett, Santiago Truffa and Dinand Webbink, and the participants to the 2019 Urban Economics Association conference and VU Amsterdam/Uppsala University seminar series for their excellent comments and suggestions. Furthermore, he would like to thank CPB Netherlands Bureau for Economic Policy Analysis for making the data available for this research, as well as their excellent suggestions. All results in this paper are based on calculations by the author using non-public microdata from Statistics Netherlands. The code used to generate the results is available upon request.

influence educational investment decisions through various channels. In this paper, I investigate whether exposure to urban environments during childhood affects the educational attainment of children and young adults.

Cross-sectional evidence reveals a sizeable rural–urban gap in educational outcomes across a wide range of countries.[1] It is unclear whether these differences are fully explained by the spatial sorting of households or whether they partially reflect an advantage of cities in educating their population. To answer this question, I make use of the particular institutional setting in one country, the Netherlands. The Dutch educational system requires students to make conscious decisions about their desired level of human capital investment both at the end of primary school and at the end of secondary school. Just before making the decisions, students participate in national tests of academic ability. These test scores are highly predictive of future academic outcomes and thus provide an excellent measure of the students' cognitive and academic ability precisely at the moment when they make educational investment decisions. Furthermore, I observe a wide range of household characteristics, including detailed measures of parental education and household income.

The analyses reveal substantial differences in the human capital choices between children growing up in urban areas and rural areas. Conditional on family background and observed cognitive ability, a one log-point increase in population density is associated with a 1.68 percentage point increase in the likelihood that a child enrolls in an academic secondary school, from a base of 23%. Similarly, among high school students who have obtained all prerequisites to enroll for a university degree, a one log-point increase in population density is associated with a 0.8 percentage point increase in the likelihood that a child attends university, from a base of 84%. Children who grow up in more urban environments are thus significantly more likely to select into the schooling tracks that provide higher levels of human capital accumulation. Taken together, a one log-point increase in population density is associated with a 1.4 percentage point increase in the probability that a child attends university, which implies an elasticity of university attendance with respect to population density of 0.07.[2] Expressed differently, growing up in the center of Amsterdam rather than a place at the 25th percentile of the density distribution would increase the probability of attending university by 3 percentage points. Similar to Combes et al. (2008) in the case of the urban wage premium, I find an elasticity about twice as large when individual and family characteristics are not controlled for.

To assess whether differences in unobserved characteristics between urban and rural students might be responsible for differences in human capital investment decisions, I employ the methodology of Altonji et al. (2005) and Oster (2019), which uses the selection on observed variables as guide to the selection on unobserved variables. Under the assumption of Altonji et al. (2005) and Oster (2019) that selection on unobserved characteristics is weakly less than selection on observed characteristics, omitted variable bias cannot account for the results obtained in this paper. In addition, the coefficients are similar across

1 See Appendix A for the urban–rural education gap based on a wide variety of countries.

2 A total of 1.23 pp of the increase is due to the higher probability that a child enrolls in an upper secondary school, taking into account high school dropout rates and that not all students with an academic secondary school degree pursue university studies. A further 0.16 pp is due to the 0.8% increase in university attendance among the children who complete upper secondary school, which contains about 20% of the population. Given that a one-log point increase in density raises university enrollment by 1.4 pp and the mean university enrollment rate of 0.2, an elasticity of 0.07 is obtained.

subgroups, not driven by functional form assumptions and robust to the inclusion of regional fixed effects. Using the historical densities of 1840 as instrument for modern densities to account for the potential endogeneity of urban settlements also leaves the results unaffected.

The findings of this paper contribute to three strands of literature. First, the paper contributes to the urban economics literature by showing that density affects human capital decisions of children. Earlier literature in this field has typically focused on the effect of density at a single decision moment (Frenette, 2006; Newbold and Brown, 2015) or on a single outcome variable (Gibbons and Silva, 2008), whereas this paper estimates the effect of population density on educational outcomes throughout the entire schooling career. In addition, the institutional setting and rich administrative data allow me to condition on a wide range of family characteristics and detailed measures of cognitive ability, which has typically remained unobserved in earlier studies. The closest related work is by Gibbons and Silva (2008), who show that density affects high school test scores in the UK. This paper in comparison takes test scores as given, and focuses on students' educational choices conditional on the primary and secondary school test scores. The resulting elasticity is around three times larger than the elasticity reported by Gibbons and Silva (2008), reflecting that a substantial part of the effect of density on educational attainment operates through differences in educational choices. The finding that children in cities accumulate more human capital also well complements the work by Glaeser and Mare (2001) and De La Roca and Puga (2017), who find that workers in cities accumulate human capital at a faster rate, and the recent work by Eckert et al. (2020), who use a natural experiment to show that refugees who are placed in cities are more likely to start working in skill-intensive occupations.

Second, the results have implications for the long-term economic growth of regions and the persistence of regional differences. It is well known that the stocks of human capital are typically lower in rural regions. The findings of this study suggest that rural regions are in addition at a disadvantage in the expansion of their human capital stock, since children in rural areas invest less in human capital accumulation throughout their childhood, even when they have the same cognitive ability and family background as children in urban areas. While the lower human capital investment in rural regions might be efficient from the perspective of the individual, it does have significant implications for the long-term inequality between urban and rural regions. This is particularly important due to the fact that human capital is the most important predictor of economic growth in regions (Gennaioli et al., 2013),[3] and hence provides a mechanism that could explain the slow observed convergence in incomes between regions within countries as noted by Gennaioli et al (2014) and others.

Finally, the paper contributes to a rapidly growing literature on the spatial (in)-equality of opportunity, following Chetty et al. (2014a) and Chetty and Hendren (2018) in the USA, Alesina et al. (2019) in Africa and Deutscher (2020) in Australia. In this paper, I focus on one specific aspect of places, namely population density, and examine the effects on educational attainment. The results indicate that children who grow up in rural communities do not appear to enjoy or take advantage of the same educational opportunities as children in urban communities, even when they appear to have the same cognitive

3 Although Gennaioli et al. (2013) may overestimate the returns to entrepreneurial education relative to worker education due to selection into entrepreneurial activities, as pointed out by Behrens et al. (2014).

ability and family background as their urban peers. Hence, the findings of this paper highlight that the urban–rural differences in educational outcomes may not just be a result of spatial sorting of households, but may also reflect a lack of opportunities or awareness of opportunities for children residing in rural regions.

Finally, a limitation of this study is that it cannot attribute the overall effect of density on educational attainment to individual mechanisms. Based on the existing literature, there are three main channels through which density is likely to affect educational decisions. First, the returns to education are higher in cities as agglomeration forces mainly complement the productivity of high-skilled workers (Baum-Snow et al., 2018; Autor, 2019), which raises the incentive to invest in human capital in urban areas. Second, the costs of obtaining education are lower in cities, both due to the smaller distance to educational institutes (Frenette, 2006), as well as the possibility that the more diverse school choice in cities improves the match between students their needs and interests and the educational institutes. Third, it might be costlier for youth to acquire information about future educational possibilities in rural areas, due to the absence of university outreach programs as well as the absence of strong network linkages to higher educational institutes in rural communities (Hoxby and Avery, 2012). However, the high correlations between the measures of the individual mechanisms make it difficult to attribute the overall effect of density to the individual mechanisms within the scope of this paper, and a detailed investigation along the lines of De La Roca and Puga (2017) and Dauth et al. (2018) in the case of the urban wage premium is thus left for future research.

The remainder of the paper proceeds as follows. Section 2 reviews the existing theoretical and empirical literature linking density to educational outcomes. Section 3 provides an overview of the context and data. Section 4 discusses the methodology and identification strategy. Section 5 presents the results and the various robustness analyses. Section 6 discusses the implications and concludes.

2. Related literature

There are good reasons why children with similar capabilities might make different educational choices depending on whether they live in urban or rural communities. This section provides an overview of the theoretical and empirical support for three potential mechanisms, as well as review the empirical literature which links educational outcomes to urbanization.

2.1. Higher returns to education

A key determinant of educational choice in any model of educational investment is the expected returns to schooling. Most models of educational investment, such as the seminal Ben-Porath model, assume that students invest in education until the point where the discounted increase in future wages is equal to the opportunity costs and direct costs of obtaining education, thus providing a tight link between perceived returns to education and human capital investments. Theoretically one would expect highly skilled and educated workers to benefit most from agglomeration economies (see Berry and Glaeser, 2005; Behrens et al., 2014; Davis and Dingel, 2019), which matches the empirical findings that the returns to education are higher in urban environments (Gould, 2007; Combes and Gobillon, 2015; Baum-Snow et al., 2018). This difference has become more pronounced over the last twenty years, during which the urban wage premium for non-college educated workers has essentially disappeared in the US (Autor, 2019). Hence, in the absence of perfect mobility, one would expect the higher returns to education in cities to lead to higher educational investment by children growing up in urban places. However, a key question is whether children and young adults correctly infer the returns to education from their environment and if they respond accordingly. The way in which people form beliefs with regards to the returns to education is still largely unknown, but the experiment by Jensen (2010) suggests that these beliefs are not necessarily deeply held, and that beliefs are updated when new information is presented. Recent work by Adukia et al. (2020) provides evidence for this in the context of India. Rural villages which were connected to cities by roads experienced significant increases in educational attainment, with the effect increasing with the returns to education in newly connected cities.

2.2. Costs of schooling

A second channel through which population density can affect educational outcomes is through a reduction in the perceived costs of attending education. At the level of primary and secondary school, the higher density of schools in urban areas means that children on average will spend less time and financial resources on commuting to a school of a given level and quality. In the case of tertiary education, commuting to college or university may become simply unfeasible for children in rural areas, hence necessitating a costly move to attend further education. The empirical literature finds some supporting evidence for the importance of commuting costs on educational decisions, particularly in the case of tertiary education. Card (1993) uses distance to college as an instrument for educational attainment and finds a significant effect in the first stage. Frenette (2006) finds in Canada that children whose families live more than 80 km from the nearest university have a 40% lower probability to enroll at university compared with individuals who grow up within 40 km of a university. However, both studies include only a very limited set of control variables, and as such, it is not clear to what degree the results are driven by the spatial selection of households.

Furthermore, the higher density of educational institutions in cities may also allow for better matching between schools and students, as long as students and schools are heterogeneous in some dimensions. Burgess et al. (2019) show that the majority of students do not list the nearest secondary school as their first preference in the UK, indicating that students do not perceive schools as a homogenous good. Such school-child match specific component can be based on academic preferences of the child, such as the level of instruction (see for instance Bau, 2019) and the focus on certain specialization tracks, or can be based on more personal preferences of the child, such as the religious orientation of a school (Cohen-Zada and Sander, 2008). Models of school choice typically assume that an outside option is available with zero utility, which depending on the context might consist of staying home, working or choosing a different level of schooling. The smaller the number of schools available, the more likely that unfavorable draws of the school-student match quality will result in the student choosing the outside option, which may reduce human capital investment in rural areas. Gibbons and Silva (2008) argue that greater school choice and competition in cites may explain the better performance of urban schools in the UK.

2.3. Information and network effects

Third, density may affect human capital decisions through the availability of information on future schooling prospects or due to network effects. Hoxby and Avery (2012) find in the USA that high-performing students who apply to non-selective colleges are disproportionally located in rural areas, as the lower density of high-performing students makes it less profitable for colleges to engage in outreach campaigns. In addition, Hoxby and Avery (2012) suggest that the lower Social-Economic Status (SES) composition of households in rural areas limits the exposure of children to alumni of various educational institutes, as well as reduce the expertise of study counselors on how to advise high-performing students. Hence, the limited information about future educational possibilities in rural areas provides a third potential channel through which density may affect educational investment.

2.4. Empirical evidence

Despite theoretical reasons to expect children and young adults in urban areas to invest more in their human capital, the empirical evidence on this subject is limited. Knight and Shi (1996) show large differences in years of educational attainment in China between urban and rural regions, although it remains unclear to what degree this is driven by household heterogeneity. Katz and Goldin (2008, 222) instead report a negative relationship between town size and high school graduation rates in the early 20th century in the USA, which they suggest might be the result of the higher opportunity costs of education in cities. More recently, Newbold and Brown (2015) show that the likelihood that Canadian youth attend university is increasing in city size.

A key challenge in this literature has been dealing with the spatial sorting by households and unobserved heterogeneity between children. Previous work has established that households sort into cities based on parental education (De la Roca, 2017 ;Ahlin et al., 2018; Bosquet and Overman, 2019) and cognitive skills (Bacolod et al., 2010; Bütikofer and Peri, 2020). Separating spatial selection from area effects is particularly difficult in the case of education, since panel data on educational outcomes is typically unavailable on the individual level. As such, the methods developed to estimate the urban wage premium, which rely heavily on the inclusion of individual fixed effects, cannot easily be applied in the case of educational outcomes.

As a way to overcome the spatial sorting of households, some studies have relied on quasi-experimental variation to quantify the effect of exposure to neighborhood on educational outcomes. Chetty et al. (2016) use the Moving-to-Opportunity program to study neighborhood effects, which provided a randomly selected group of families the opportunity to move to a better neighborhood with less poverty and higher incomes. They find that children who moved to better neighborhoods before the age of 13 years had significantly higher schooling outcomes compared with those who did not receive the opportunity to move. However, such experiments are typically difficult to scale beyond the individual city. In a more general approach, Chetty and Hendren (2018) use the difference in the age of children when they move between regions to identify neighborhood effects. They find a negative relationship between population density and upward mobility as measured by the income at age 26 years, although it remains unclear to what extent differences in human capital accumulation play a role.

3. Context and data

In this paper, I utilize the Dutch context to investigate the relationship between population density and educational outcomes, which is particularly well suited for this question.

Tracking in the Dutch educational system starts at the age of 11–12 years, which means that meaningful decisions on human capital investment are taken from an early age onwards. Furthermore, students participate in high-stakes tests of academic ability shortly before making their key educational decisions. Hence, I observe detailed measures of cognitive ability precisely when students make educational decisions, which remains typically unobserved in most other settings. One concern with the Netherlands is that it is relatively urbanized by international standards. However, earlier studies in the Netherlands on the urban wage premium found results in line with other countries (Groot et al., 2014; Verstraten et al., 2019).[4]

3.1. Context

Figure 1 provides an overview of the Dutch school system, including the flows between school types.[5] Compulsory education starts at the age 6 years, when all students enroll in primary school. The primary school education lasts for 6 years, at the end of which students participate in a national test measuring their ability in reading comprehension, mathematics and vocabulary, as well as their reasoning and studying ability.[6] Tracking begins at the end of primary school when students can enroll in three different levels of secondary school: upper, middle and lower secondary school. The three levels differ both in length of study, the difficulty of the material and the access to tertiary education which a degree grants. Students are free to apply to any of the three levels of secondary school after finishing primary school, but secondary schools can choose whether to accept students. During the admission decision, secondary schools primarily rely on the scores on the standardized test at the end of the primary school, as well as the recommendation of the primary school teacher. In the years 2003–2015, the score on the national standardized test was considered the leading admission criteria for secondary schools.[7]

Figure 2 shows the distribution of the end of primary school test score for the cohort born in 1996 (panel a) as well as the probability that students with a given test-score enroll in an upper secondary school (panel b). The test score is highly predictive of whether a student enrolls in upper secondary school. Students who score less than 535 (which is about the median test score) are highly unlikely to be enrolled in an upper secondary school 4 years later, whereas among the group of students who obtain the maximum score of 550,[8] 95% enrolls at an upper secondary school. Panel c in Figure 2 shows how these

4 Verstraten et al. (2019) find an elasticity of wages with respect to employment density of 0.02, whereas Groot et al. (2014) find an elasticity of 0.048 without including individual fixed effects. Both are very similar to the estimates cited by Combes and Gobillon (2015) for models with and without person fixed effects in other countries.

5 All data on school enrollment en educational attainment is collected by the Ministry of Education and is made available for research via Statistics Netherlands. The school enrollment data for all levels is taken from the "Onderwijsdeelnemerstab" register, the end of primary school test from the "Citotab" registry and the secondary school final examination grades from the "Examvotab" registry.

6 This test consists of 200 multiple choice questions (100 on reading comprehension and vocabulary, 60 on Mathematics and 40 on studying and reasoning ability). The tests are centrally graded by the Citigroup agency, which also develops the annual test.

7 At the end of 2015, the system was reformed to make the primary school teacher's recommendation binding and secondary schools were no longer allowed to use the end-of-primary school test score. Politicians and teachers felt that the test had become the only selection criteria on which secondary schools evaluated children, which they argued put undue pressure on the children to perform at one specific moment in time.

8 The maximum score of 550 is obtained by 5.7% of the students in the 1996 cohort.

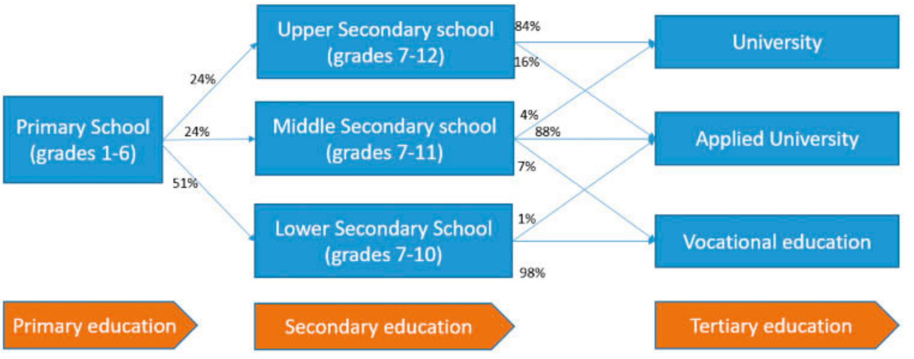

Figure 1. Tracking through the Dutch education system for the 1996 cohort. The figure is based on the cohort born in 1996. The secondary school enrollment percentages reflects enrollment 4 years after completing primary school. The tertiary education percentages reflect enrollment within 3 years after finishing the highest secondary school degree. There are some streams within the secondary schooling system (for instance, students dropping from upper secondary school to middle secondary school), which are not displayed here. In addition, some students continue with applied university after obtaining a vocational education degree or with university after obtaining an applied university degree. These streams are not displayed here either, as this usually happens outside of the 3-year window.

probabilities differ between urban and rural regions, providing a first indication of the differences in educational decisions between urban and rural regions.

Once students enroll in a secondary school, they continue for another 4–6 years depending on the selected type of secondary school. At the end of secondary school, all students take a national examination which determines if they are granted a degree. This end of secondary school test is specific for their level of secondary school.[9] Once students obtain a secondary school degree, they continue to tertiary education.[10] The tracking system is designed such that certain levels of secondary school feed into a specific level of tertiary education. As can be seen from Figure 1, 84% of the students who obtain a degree from an upper secondary school are registered at a university 3 years after graduation. Similarly, 88% of the students who obtain a degree from a middle secondary school and 98% of the students completing the lower secondary school are, respectively, registered at an applied university[11] and a vocational education 3 years

9 All students have a basic set of subjects (such as Dutch, English and Math A), as well as a set of subjects related to their chosen subject specialization (health, humanities, natural sciences or social sciences) in which they complete the national examination. The material covered in the examination as well as the difficulty varies depending on the type of secondary school.

10 Upon obtaining a middle or lower secondary school degree, students with excellent grades have the opportunity to take the final two years of the next secondary school level. When discussing degree completion or conditioning on degree, these students are always classified according to the highest degree they obtained.

11 The applied university is similar to the Fachhochschule in the German system. Applied university is more focused on practical skills compared with universities and part of the degree requirement typically involves various long-term internships at companies. For illustrational purposes, individuals with a university/applied university/vocational education on average had an annual income of respectively 50.00/36.000/25.000 euro between 2007 and 2009 (Statistics Netherlands, 2011) .

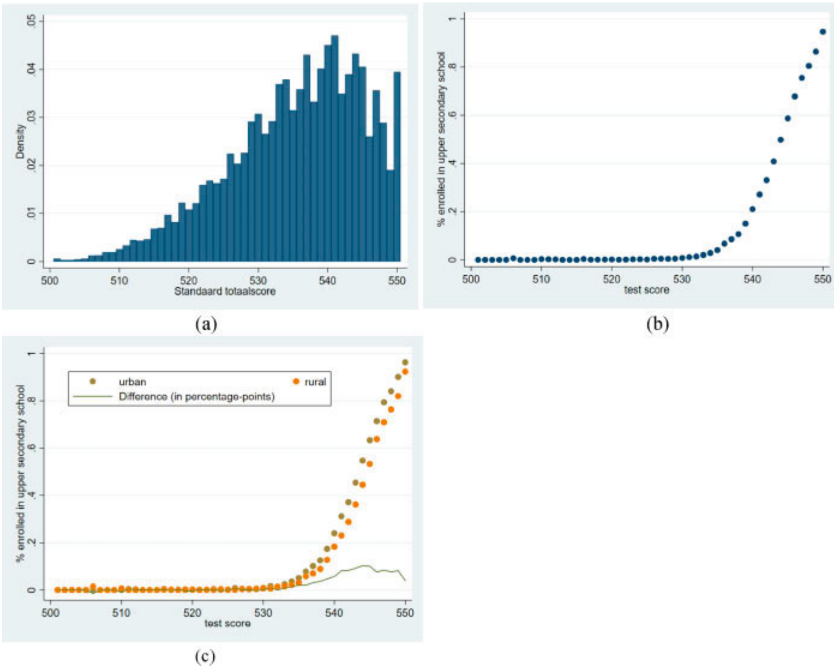

Figure 2. Frequency and probability of enrolling in upper secondary school by end-of-primary school test score. (**a**) Displays the frequency for each of the 50 potential scores (501–550) that students could receive at the end-of-primary school test for the cohort born in 1996. (**b**) Displays the percentage of students that are enrolled at an upper secondary school 4 years after taking the test by the end-of-primary school test score. (**c**) Splits the sample between those who grew up in a place with less than the median density (rural) and those who grew up in a place with above-median density (urban). The green line shows the difference in upper secondary school enrollment between urban and rural places in percentage points.

after graduation.[12] Hence, the choice that students make at the end of primary school matters greatly for their future educational prospects.

Tertiary education in the Netherlands differs from many other European countries in the sense that few university degrees have a binding constraint on the number of students they accept.[13] Students who have obtained the relevant qualifications can typically register for their study of choice at the university of their choice, without additional entry conditions. Tuition fees are around 1800 euros per year for universities and applied universities and around 1000 for vocational education, with governmental student loans available if needed. Primary schools and secondary schools are free of tuition. Educational institutes

12 These figures only count students who remained in the Netherlands after completing their secondary school. Around 2% of Dutch students register for a bachelor- or master degree abroad at some point during their education (Department of Education, 2016). As robustness check in the main analysis, I exclude all students from the analysis who are not registered at any educational institution at least 2 out of the 3 years between the ages 19 and 21 years.

13 The most important exceptions are medicine, veterinary medicine and dentistry, as these studies typically are comparatively expensive to offer.

at all levels are financed directly by the national government, with the financing based on the number and type of students enrolled.

3.2. Data

The data used in the analysis are primarily for the cohorts born between 1994 and 1998.[14] I restrict the sample to individuals born in the Netherlands for whom both parents can be identified. For these individuals, I observe the place of residence between the ages 1 and 20 years, the enrollment status in all types of secondary and tertiary education and the results on the national tests at the end of primary school[15] and secondary school. In addition, a large number of family characteristics are available, including parental income, education, country of birth and year of birth. The summary statistics for all variables are provided in Table 1.

Parental income is defined similarly as in Chetty and Hendren (2018) and consists of the sum of income for both parents when the child is between ages 14 and 18 years, divided by the number of parent-years with non-missing income. Some parents have a negative income, top-coded income or have missing income for more than 5 parent-years. For the parents with negative income and parents with fewer than 5 parent-year income observations, income is set to zero and a dummy is included for both groups in each regression. A dummy is also included for top-coded incomes. The robustness checks reveal that the results are not sensitive to exclusion of these (small) groups, which together contain about 1.5% of the observations. As the coverage of the income tax data may have improved slightly over the years, I allow the coefficient of parental income as well as the coefficients of the three dummy variables to vary across cohorts.

For each parent, I observe 13 possible levels of education.[16] I do not observe education for all parents, particularly for older parents who completed schooling before some of the national education registers started. Discussions with Statistics Netherlands reveal that parents with missing education are more likely to be low educated, since highly educated parents are more likely to be captured by the various educational registers. In the baseline result, missing education is included as a 14th education type. However, the results are robust to the exclusion of this group. To opt for an as flexible approach as possible when controlling for parental education, I include dummies for each of the 196 possible parental education combinations. Similarly, parental country of birth is grouped into 'The Netherlands', 'Europe' and 'Non-European' for each parent, and nine dummies are included to account for each possible parental combination. Finally, the age of the oldest parent at the time of the birth is added as an additional control variable.

14 Some analyses can be carried out using a larger sample. For consistency, results in the main text are based on the cohorts 1994–1998 unless specified otherwise. The main tables also report results for a larger sample as robustness test whenever possible.

15 A small number of students is exempted from making the test due to disabilities and not all schools have agreed to make the test scores available to Statistics Netherlands. Nonetheless, I observe the test score for the large majority of the population (70%). The coefficients on the urbanization measure of columns (1) and (2) of Tables 2 and 3 are virtually unchanged (within 10% in all cases) when estimating it either on the full sample of individuals born between 1994 and 1998 or the 70% sample for whom I observe the test score. Hence, different selection between rural and urban regions into which schools agreed to make the test scores available to Statistics Netherlands is not driving the results.

16 These groups are kindergarten, primary school, some secondary education (low, middle or high), secondary education degree (low, middle or high), some university or applied university, applied university bachelor, university bachelor, university master or applied university master, and doctoral degree. Vocational education is not listed separately, but included in the three levels of secondary education.

Table 1. Summary statistics for individuals born in the Netherlands between the years 1994 and 1998

	N	Mean	Sd. Dev.	p1	p99
Child characteristics					
Urbanization measured at age 11 years	631,815	12.07	0.87	9.99	13.64
Year of birth	631,815	1995.99	1.42	1994	1998
Parental characteristics					
Log parental income	631,890	10.36	0.83	8.82	11.68
Insufficient income data (dummy)	631,890	0.001	0.03	0.00	0.00
Negative household income (dummy)	631,890	0.003	0.05	0.00	0.00
Top coded incomes (dummy)	631,890	0.011	0.10	0.00	1.00
Age of oldest parent at the time of birth	631,890	33.46	4.97	23.00	48.00
Country of birth mother (categorical)	631,890	1.22	0.58	1.00	3.00
Country of birth father (categorical)	631,890	1.23	0.59	1.00	3.00
Education level mother (categorical)	361,625	9.49	3.45	2.00	14.00
Education level father (categorical)	340,446	9.10	3.55	2.00	14.00
Secondary school enrollment and graduation variables					
End of primary school test score	631,890	535.43	9.66	510.00	550.00
Upper secondary school enrollment	631,890	0.23	0.42	0.00	1.00
Middle secondary school enrollment	631,890	0.24	0.43	0.00	1.00
Lower secondary school enrollment	631,890	0.50	0.50	0.00	1.00
Upper secondary school graduation	631,890	0.19	0.39	0.00	1.00
Middle secondary school graduation	631,890	0.25	0.43	0.00	1.00
Lower secondary school graduation	631,890	0.45	0.50	0.00	1.00
GPA upper secondary school degree	121,106	6.63	0.69	5.50	8.50
Specialization track upper secondary school degree (categorical variable)	121,106	4.47	2.65	1.00	10.00
Tertiary education enrollment (within 3 years of high school graduation, conditional on graduating)					
University enrollment	400,305	0.20	0.40	0.00	1.00
Applied University enrollment	400,305	0.29	0.46	0.00	1.00
Vocational education enrollment	400,305	0.50	0.50	0.00	1.00

Note: The table shows the summary statistics for the individuals born between 1994 and 1998 in the Netherlands. Due to the privacy-sensitive nature of the microdata, it is not possible to report minima or maxima, hence the 1st and 99th percentiles values are displayed instead. As explained in the main text, education is not available for all parents. Secondary school enrollment is measured 4 years after completing primary school. The GPA and specialization track of the upper secondary school degree are only observed for students who graduated from upper secondary school. Tertiary education enrollment is based on cohorts born in 1994–1996, as for the cohorts born in 1997/1998 not enough time has passed for all students to measure. Tertiary education enrollment is defined as the highest enrollment within 3 years of high-school graduation. The graduation rate of middle secondary school is higher than enrollment due to the fact that about 10% of the children enrolled in an upper secondary school drop down to a middle secondary school during their final three years in secondary school.

The next step is to construct an index of urbanization. For each individual, the zip code in which he or she resided between the years 1995 and 2018 is observed. The size of the zip codes is relatively small at 8 km^2. I define urbanization as the log of the number of people living within 10 km of the centroid of the zip code in which an individual resides (see Appendix B for the details of the procedure).[17] The threshold of 10 km is selected

17 To avoid potential reverse causality problems, I calculate the density based on the spatial distribution of the population in 1995. As the robustness analysis reveals, the results remain virtually unchanged when I use the spatial distribution of 1840 as instrumental variable.

based on the fact that the majority of Dutch children travel to school by bicycle, where 10 km seems a reasonable upper bound for their reach. Furthermore, the choice of a 10-km radius is in line with recent studies of agglomeration economies on wages (De La Roca and Puga, 2017; Verstraten et al., 2019). Nonetheless, the number of individuals that live within, respectively, 5, 10 and 20 km is highly correlated, and as such, the exact distance cut-off has little influence on the results.[18] Densities are the highest in the urbanized Western part of the Netherlands and are relatively low in the North and East, as shown in Figure 3.

4. Empirical approach

The Dutch educational system provides two key decision moments that can be exploited to analyze the impact of urbanization on educational outcomes.[19] The first decision moment is at the end of primary school, when children select one of the three levels of secondary school. At this point, I observe the cognitive ability of the child as measured by the score on the national end of primary school test, as well as a wide range of family characteristics. This allows me to study whether, conditional on observed academic ability and parental background, children who grow up in urban areas make different educational choices compared with children who grow up in rural environments. The estimating equation is displayed in Equation (1), where the family characteristics contain the variables discussed in Section 3.2 and dummies are included for all possible primary school test scores. The educational outcome measure is whether a child attends an upper secondary school four years after graduation from primary school. The log-linear relationship between density and educational outcomes is well supported by the data, as will be shown in Section 5.2,

$$\begin{aligned} P(\text{attends upper secondary school})_i = {} & \beta_1 * \text{Child characteristics}_i \\ & + \beta_2 * \text{Family characterstics}_i + \beta_3 * f(\text{test score}_i) \\ & + \beta_4 * \text{Urbanization}_i + \gamma_t + \varepsilon_{1i}. \end{aligned} \tag{1}$$

The second decision moment is at the end of secondary school. Here, I focus on the group of students who obtained a degree from an upper secondary school which provides access to the large majority of university degrees without further conditions. Equation (2) estimates the effect of density on the probability of attending an academic university within 3 years after graduating from high school,[20] conditional having obtained a degree from

18 For instance, the correlation between density within 10 km and the density within 5 km and the density within 20 km is 0.89 in both cases. Alternatively one could also define density based on the number of schools within a given distance as in Gibbons and Silva (2008) or only use the population density of the zip code itself. The results are similar and the conclusions of this paper do not depend on the selected measure of density.

19 In the empirical framework, I treat the decision for secondary school and tertiary school as two separate decisions. It would also be possible to estimate the two decisions in a joint framework, in which the decision for secondary school and tertiary education are made jointly. While a joint framework might be a better representation for the decision process of certain students, it would significantly complicate the exposition and the interpretation of the coefficients. Furthermore, the most important effect of a joint framework would be a change in the timing when a student does or not does decide to attend university, which should leave the joint effect of density on university attendance across the two decision moments unaffected.

20 The reason for allowing a 3-year lag is that some students opt for a gap year between completing secondary school and starting tertiary education.

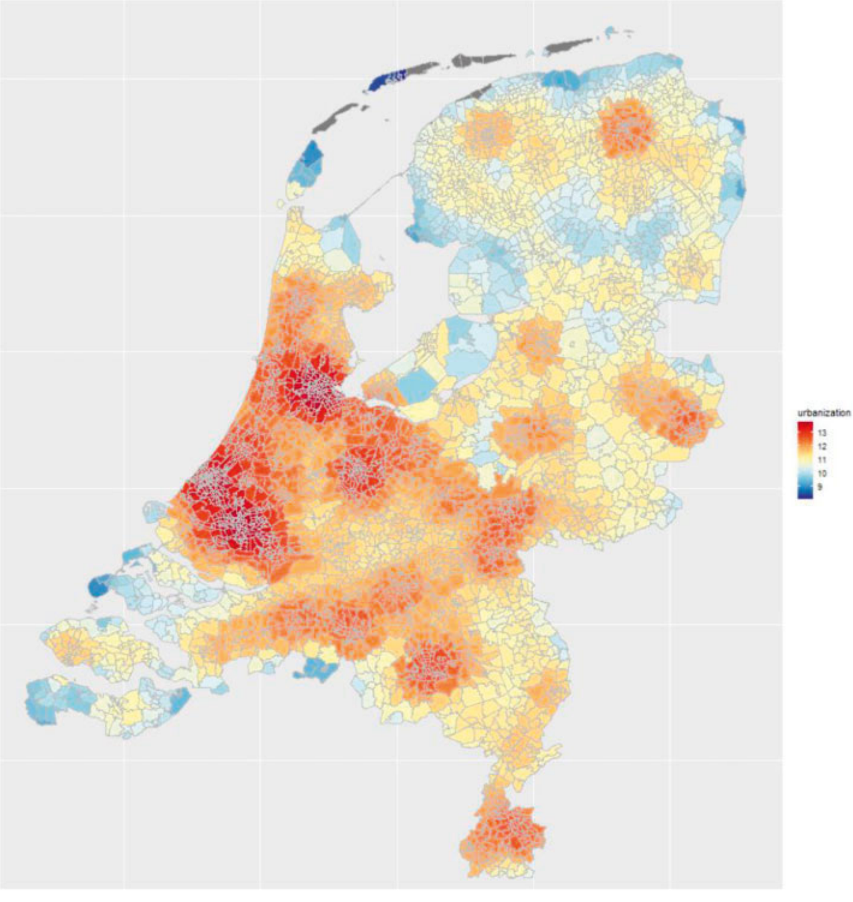

Figure 3. Urbanization measure per zip code. The map displays the log of number of individuals within 10 km, based on the 1995 population distribution. For graphical purposes, areas with an urbanization grade less than 9 are gray (containing 230 out of the 631.890 observations in the baseline sample). Appendix B provides details on the construction of the urbanization measure.

an upper secondary school, high school Grade Point Average (GPA) and family characteristics.

$$\begin{aligned} &P(\text{attends university}|\text{upper secondary school degree})_i \\ &\quad = \alpha_1 * \text{Child characteristics}_i + \alpha_2 * \text{Family characteristics}_i + \alpha_3 * f(\text{GPA high school}_i) \\ &\qquad + \alpha_4 * \text{Urbanization}_i + \gamma_t + \varepsilon_{2i}. \end{aligned} \tag{2}$$

Under the assumption that the covariance between ε_{it} and the urbanization measure is zero, Equations (1) and (2) will correctly identify the effect of growing up in an urban environment on educational outcomes. One concern is that parents move to places best fitted to realize the potential outcomes of their children, in which case the estimates of α_4 and

β_4 would be biased. However, I find no evidence of this in the Dutch setting. For the cohorts born between 1994 and 1998, only 5.4% of the families with children between the ages 6 and 17 years make a substantial move (more than 20 km), thus limiting the degree to which families respond to the realized potential of their children. In addition, within the group of children who move, there is no correlation between the change in density and the observed academic ability of the children.[21] Nonetheless, I show that the results of Equations (1) and (2) are robust to estimating it on either the full sample or on the group of children who do not move across municipalities between the ages 6 and 17 years.[22]

A second concern is that parental characteristics may vary between urban and rural places in ways that are not fully captured by the control variables. For instance, it might be that parents in urban areas are more ambitious than parents in rural areas and that this is not fully captured by the differences in wages. Separating spatial sorting from area effects has been a key challenge in the urban literature and solutions to this have typically relied heavily on individual-fixed effects (Combes et al., 2008; De La Roca and Puga, 2017). However, such approaches are typically not possible in the case of educational outcomes due to the lack of consistent national panel data on the educational outcomes of children.[23] Section 5.2 explores this potential threat to identification, using the methodology of Altonji et al. (2005) and Oster (2019) to assess the potential importance of sorting on unobserved characteristics.

Finally, conditioning on the test scores in Equation (1) and (2) might risk overcontrolling for the effect of density on educational attainment, as test scores themselves may also be affected by density. Gibbons and Silva (2008) find that students in urban schools in the UK show more rapid improvements in test scores compared with rural students. Hence, to the degree that urban–rural differences in test scores reflect better learning outcomes in urban areas rather than heterogeneity in cognitive ability between students, the reported coefficients may somewhat underestimate the full effect of population density on educational outcomes.

5. Results

5.1. Baseline results

Table 2 presents the results of Equation (1), analyzing how the decision of children to enroll in an upper secondary school depends on population density. Column (1) shows the regression when only the urbanization measure is included, whereas column (2) adds child and family characteristics and column (3) adds dummies for the end of primary school test scores. The coefficient of urbanization is statistically significant in all three specifications

21 Result not separately included in the paper but available from the author upon request.

22 A different approach would be to use the Chetty and Hendren (2018) identification strategy and use children who move between regions at different ages to generate urban exposure effects. However, within the group of movers there is a substantial decline in the average end-of-primary school test score by the age of moving, even after the age of 12 years when the test score is essentially predetermined. This differential selection into migration by the age of children hence makes it difficult to apply the Chetty and Hendren (2018) identification strategy within this context.

23 Such approach is difficult to implement for educational outcomes, as most countries lack repeated outcomes for children on a national level, which would allow fixed-effect estimations similar to Combes et al. (2008) with grades instead of wages. Individual-fixed effect approaches have been successfully applied within schools in the economics of education literature to identify effects of school level variables, such as teacher quality (see Chetty et al., 2014b).

Table 2. Effect of urbanization on probability of enrolling in upper secondary school

	Baseline estimates			Sensitivity analysis			Area fixed effects		IV
	Ind. level controls (1)	Ind. + Fam. controls (2)	Ind. + Fam controls + test-score (3)	Movers excluded (4)	Different cohorts (1999–2002) (5)	Full parental education (6)	Municipality-fixed effects (7)	Provincial-fixed effects (8)	IV estimates (9)
Urbanization at age 11 years	0.0342***	0.0260***	0.0167***	0.0180***	0.0167***	0.0141***	0.0123**	0.0196***	0.0178***
	(0.0067)	(0.0037)	(0.0017)	(0.0019)	(0.0020)	(0.0015)	(0.0045)	(0.0020)	(0.0027)
Individual controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Family controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Test score dummies	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.005	0.13	0.50	0.50	0.52	0.53	0.51	0.50	0.50
Number of observation	631.731	631.731	631.731	554.826	503.200	242.156	631.731	631.731	628.396

Note: All results apart from column (5) are based on individuals born between 1994 and 1998. Dependent variable is whether a child is enrolled in an upper secondary school 4 years after completing primary school. Individual controls include a gender dummy and cohort dummies. Family controls include dummies for parental education combinations (168 dummies), parental nationality combinations (9 dummies), year-of-birth of oldest parent (40 dummies), family income (interacted with cohort-fixed effects), low-income dummy (interacted with cohort-fixed effects), top-coded income dummy (interacted with cohort-fixed effects) and dummy for insufficient income data (interacted with cohort-fixed effects). For a detailed explanation of the construction of the urbanization measure or the family controls, see Section 3. Column 4 excludes all children who moved municipalities during school-going age. Column 5 instead estimates the model on the cohorts born between the years 1999 and 2002. Column 6 removes all parents for whom uncertainty exists about the education level of one or both parents. Column 7 adds municipality-fixed effects for the municipality in which children live at age 11 years (430 dummies), whereas column (8) adds province-fixed effect for the province in which children live at age 11 years (12 dummies). Column (9) instruments for modern densities (based on 1995 population distribution, see Appendix B) with historical densities. All standard errors are clustered on the municipality level. Columns (1–8) are estimated by OLS, column (9) by 2SLS. First stage results of column (9) are reported in Appendix C3 (first stage F-statistic: 2349.94).

at the 1% level. In economic terms, the baseline result of column (3) indicates that an increase in density by one log-point (which coincidentally is close to one standard deviation, see Table 1) is associated with a 1.68 percentage point increase in the probability that a child attends upper secondary school. Given that the mean percentage of children attending the upper secondary school is about 23%, this increase is substantial. When comparing the individual columns, the coefficient drops somewhat when family characteristics are added, which is largely driven by the differences in parental education between urban and rural regions. The coefficient declines further when the test scores are added, which might reflect a lower initial academic ability of children in rural regions, or alternatively might be the effect of overcontrolling to the degree that growing up in an urban environment may also affect end-of-primary school test scores.[24]

Columns 4–6 test the robustness of the results by excluding various subgroups. Column (4) excludes children who moved between municipalities during their school-going age, as their families may have sorted themselves into places based on the potential outcomes of their children. Column (5) changes the sample and instead estimates the effect on the cohorts born between 1999 and 2002, for whom information is also available. Finally, column (6) reduces the sample to the individuals for whom the exact education level of both parents is known. The coefficient is very stable across the various subgroups, suggesting that these groups are not driving the results.

One concern might be that children in urban environments enroll in classes that are too difficult for their level of ability, resulting in higher drop-out rates in urban areas.[25] Around 21% of the students who are enrolled in an upper secondary school four years after finishing primary school eventually do not obtain an upper secondary school degree, which is far from negligible.[26] To control for the possibility that differences in dropout rates between urban and rural areas are driving the results, Table 3 instead analyzes the probability that a student obtains a degree from an upper secondary school. The results are very similar to Table 2, indicating that a misallocation of students at the end of primary school in urban communities is unlikely to drive the results. Furthermore, Table C1 shows the results when directly analyzing drop-out rates, by estimating the probability that a student obtains a degree from an upper secondary school, conditional on being enrolled in an upper secondary school 4 years after finishing primary school. The results indicate that children in urban areas are actually slightly less likely to drop out from an upper

24 Children in urban environments indeed have higher average test scores on the end-of-primary school tests, even when conditioning on household and child characteristics. A one-log point increase in population density raises the average end-of-primary school test score by 0.2 points and increases the probability that a child has at least a score of 545 by about 1 percentage point (mean 15 pp).

25 Alternatively, one may also worry that children in rural areas obtain a degree from a middle secondary school, then continue with some years of applied university before switching to university. In this case, children from rural areas would simply follow a different track to end up at university. I do find some evidence for this, as children who obtain a middle secondary school degree are slightly more likely to enroll in university in the seven years after graduating from middle secondary school if they grow up in a rural area rather than an urban area. However, the increased usage of this alternative route in rural areas is far too small to compensate for the negative effects described in Tables 2 and 4. Overall, a one-log point increase in density reduces the probability that a student enrolls for university by 0.6% among the group of middle secondary school graduates (25% of the sample, see Table 1). Hence, taking this path to university into account reduces the overall effect of density on university enrollment by 0.25 * 0.6 = 0.15%. Hence, this alternative route can compensate for about 10% of the total effect of density on university enrollment of 1.4%.

26 The majority (62%) of these students instead obtains a degree from a middle secondary school. The drop-out figure of 21% is in line with the statistics reported for this period by Statistics Netherlands (2019).

Table 3. Effect of urbanization on probability of obtaining a degree in upper secondary school

	Baseline estimates			Sensitivity analysis		Area-fixed effects		IV
	Ind. level controls	Ind. + Fam. controls	Ind. + Fam controls + test score	Movers excluded	Full parental education	Municipa-lity-fixed effects	Provincial-fixed effects	IV estimates
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Urbanization at age 11 years	0.0296***	0.0226***	0.0146***	0.0156***	0.0130***	0.0113**	0.0153***	0.0158***
	(0.0060)	(0.0030)	(0.0013)	(0.0014)	(0.0013)	(0.0039)	(0.0016)	(0.0016)
Individual controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Family controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Test score dummies	No	No	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.004	0.12	0.44	0.45	0.47	0.45	0.44	0.44
Number of observations	631.731	631.731	631.731	554.826	242.156	631.731	631.731	628.396

Note: All results are based on individuals born between 1994 and 1998. Dependent variable is whether a child graduates from an upper secondary school. Individual controls include a gender-dummy and cohort-dummies. Family controls include dummies for parental education combinations (168 dummies), parental nationality combinations (nine dummies), year-of-birth of oldest parent (40 dummies), family income (interacted with cohort fixed effects), low-income dummy (interacted with cohort-fixed effects), top-coded income dummy (interacted with cohort-fixed effects) and dummy for insufficient income data (interacted with cohort-fixed effects). For a detailed explanation of the construction of the urbanization measure or the family controls, see Section 3. Column (4) removes all parents for whom uncertainty exists about the education level of one or both parents. Column (5) adds municipality-fixed effects for the municipality in which children live at age 11 years (430 dummies), whereas column (6) adds province-fixed effect for the province in which children live at age 11 years (12 dummies). Column (7) instruments for modern densities (based on 1995 population distribution, see Appendix B) with historical densities. All standard errors are clustered on the municipality level. Columns (1–7) are estimated by OLS, column (8) by 2SLS. First stage results of column (8) are reported in Appendix C3 (first stage F-statistic: 2349.94).

Table 4. Effect of urbanization on probability of enrolling at university/applied university, conditional on having an upper secondary school degree

	Baseline estimates			Sensitivity analysis			Area-fixed effects		IV
	Ind. level controls	Ind. + Fam. controls	Ind. + Fam controls + GPA	Movers excluded	Cohorts of Table 2 (1994–1998)	Full parental education	Municipality-fixed effects	Provincial-fixed effects	IV estimates
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Urbanization at age 11 years	0.0232***	0.0083***	0.0080***	0.0084***	0.0087***	0.0072***	0.0098*	0.0135***	0.0071**
	(0.0022)	(0.0016)	(0.0017)	(0.0018)	(0.0019)	(0.0018)	(0.0047)	(0.0022)	(0.0025)
Individual controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Family controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
GPA and specialization track	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.003	0.03	0.06	0.07	0.06	0.06	0.07	0.07	0.06
Number of observation	289.109	289.109	289.109	253.430	139.559	110.921	289.109	289.109	287.763

Note: All results apart from column (5) are based on individuals born between 1989 and 1998. Dependent variable is whether a child attends university within 3 years of graduating from upper secondary school. Individual controls include a gender dummy and cohort dummies. Family controls include dummies for parental education combinations (168 dummies), parental nationality combinations (nine dummies), year-of-birth of oldest parent (40 dummies), family income (interacted with cohort-fixed effects), low-income dummy (interacted with cohort-fixed effects), top-coded income dummy (interacted with cohort-fixed effects) and dummy for insufficient income data (interacted with cohort-fixed effects). For a detailed explanation of the construction of the urbanization measure or the family controls, see Section 3. Column (4) excludes all children who moved municipalities during school-going age. Column (5) limits the sample to the cohorts born between 1994 and 1998, in line with the baseline sample of Table II. Column (6) removes all parents for whom uncertainty exists about the education level of one or both parents. Column (7) adds municipality-fixed effects for the municipality in which children live at age 11 years (430 dummies), whereas column (8) adds province-fixed effect for the province in which children live at age 11 years (12 dummies). Column (9) instruments for modern densities (based on 1995 population distribution, see Appendix B) with historical densities. All standard errors are clustered on the municipality level. Columns (1–8) are estimated by OLS and column (9) by 2SLS. First stage results of column (9) are reported in Appendix C3 (first stage F-statistic: 1028.11).

secondary school, again showing that differential dropout rates between urban and rural areas are not driving the results of Table 2.

The second key educational decision moment is after children obtain an upper secondary school degree, when they have to decide on the level of tertiary education. Table 4 shows the estimates for Equation (2), analyzing the effect of urbanization on the probability that a student enrolls at a university within 3 years of obtaining an upper secondary school degree, which provides access to the large majority of university–subject combinations. Column (1) only includes the urbanization measure, whereas column (2) adds student and family characteristics and column (3) adds controls for the specialization track and the Grade Point Average (GPA) on the national examination at the end of upper secondary school.[27] Since I condition on the national end of high school examination scores (instead of end of primary school test scores), it is possible to use a larger sample and to also include the cohorts born between 1989 and 1993.[28]

The results in Table 4 indicate that children who grow up in urban areas are more likely to enroll at a university, conditional on their academic ability and family characteristics. The preferred specification in column (3) indicates that children who grow up in an area where log density is one-point higher (about one standard deviation) are 0.8 percentage points more likely to enroll in university, from a base of 84%. The coefficient remains very similar when excluding children who moved between municipalities during their childhood (column 4), limiting the sample to the cohorts born between 1994 and 1998 (column 5) or limiting the sample to the children for whom no uncertainty exists over parental education (column 6).

The next question that arises is which alternatives these children select when they decide not to attend university. One alternative is to enroll at the lower rated applied universities. As Figure 1 shows, about 16% of the students who obtained an upper secondary school degree decide to enroll at an applied university, despite having the necessary qualifications to start at a university. Table C2 shows the effect of density on the probability that a student enrolls at an applied university, conditional having obtained a degree from an upper secondary school. The coefficient is significant and negative, indicating that students growing up in rural communities are more likely to instead enroll at an applied university. The coefficients are nearly identical to the coefficients in Table 4, suggesting that diversion of rural students into the lower rated applied universities fully explains the lower enrollment rates at universities found in Table 4.

5.2. Robustness

The results in Section 5.1 indicate that population density may play a role in the educational decisions of children. Naturally, some concerns arise over the interpretation and the robustness of this result. This section will discuss four potential threats to identification: the log-linear functional form assumption, the influence of general regional differences, the endogeneity of population density and spatial sorting on unobserved characteristics.

27 The GPA is based on the national examinations at the end of high school. In the final 2 years of secondary school, students can enroll in one of the possible four tracks (humanities, social sciences, biology or natural sciences). The track determines the subjects in which the students take their final examinations.

28 For the students born in 1994–1998, it is possible to include both the secondary school test score as well as end of primary school test score. However, adding the end of primary test score results offers little explanatory power over the high school GPA and leaves the coefficient virtually unchanged. Hence, the loss of this variable due is more than outweigh by the higher precision obtained due to the larger sample.

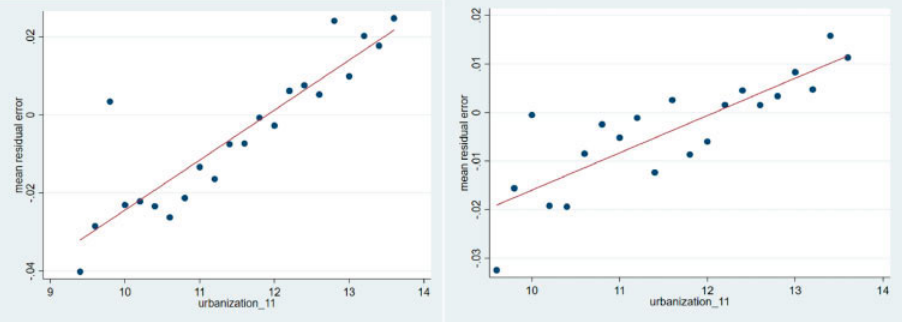

Figure 4. Residuals of baseline estimates plotted against urbanization measure. Average residuals of columns (3) of Tables 2 (left) and 4 (right) without the urbanization measure plotted against the urbanization measure. Each dot corresponds to the average residual in a 0.2 log-point bin. Bins containing fewer than 500 observations are not depicted.

5.2.1. Functional form assumption

All estimations have been based on the assumption of a log-linear relationship between population density and educational outcomes, which may or may not accurately represent the true functional form. If the functional form is misspecified, it might bias the coefficients. To test this possibility, Figure 4 plots the residuals of the baseline estimations (columns 3 in Tables 2 and 4) without the urbanization measure against the urbanization measure itself. The residuals are averaged over 0.2 log point intervals. As Figure 4 shows, the log-linear functional form assumption seems to fit the data quite well and hence is unlikely to bias the results. In addition, the residuals reveal that the results do not depend on any specific part of the density distribution: the log-linear relationship seems to describe the actual relationship well throughout the observed density distribution.

5.2.2. Influence of broad regional differences

Second, the map of the urbanization measure displayed in Figure 3 shows that urbanization is highest in the West and relatively low in the North and East of the country. As a result, the urbanization measure may at least partially reflect broader economic or cultural regional differences within the country, rather than urbanization *per se*. To test this possibility, columns (7) and (8) in Tables 2 and 4 add provincial and municipality fixed effects to the model. The coefficient of the urbanization measure in this case is identified on the variation in population density within provinces and municipalities. However, the coefficients change relatively little and remain statistically significant in all cases, even though the standard errors increase substantially as most of the variation in the urbanization measure is discarded. Hence, the results are not driven by broader differences between the various regions.[29]

29 Kelly (2019, 2020) shows that spatial autocorrelation in the dependent and independent variables means that the standard errors might be too small in many spatial regressions. Using the standard error correction proposed by Kelly (2020) increases the standard errors of column 3 in Table 2 by a factor of 1.8, which thus remains significant at the 0.1% level. The correction for Table 4 standard errors is somewhat more complicated due to the low number of observations in many zip codes, which results in very noisy estimates of the zip code fixed effect. If we assume the same spatial autocorrelation in Table 2 as in Table 4 then the estimates in Table 4 remain significant at the 5% level with a *t*-statistic of 2.6.

5.2.3. Endogeneity of Urbanization

Third, one might be concerned that population density itself is endogenous. Factors that attract population to certain areas and hence contribute to city formation may also directly affect educational outcomes. Furthermore, reverse causality can play a role if individuals migrate to areas with good schools or favorable schooling policies. Combes et al. (2008) highlight the potential of such contemporary factors to bias estimations in the case of the urban wage premium and use historical densities as an instrumental variable (IV). To alleviate concerns that the results in this paper are driven by the endogeneity of the urbanization measure, I follow Combes et al. (2008) and instrument current population densities with the population density based on the 1840 Dutch census.[30,31] The details of this procedure and a map of the population densities of 1840 are contained in Appendix C3. Even though it is not a perfect historical instrument, as 4 of the 13 Dutch universities predate 1840, it should substantially reduce any bias arising from the endogeneity of the urbanization measure. The first stage is significant as shown in Appendix C3, which indicates that the instrument is relevant. The results of the 2SLS estimations are provided in column (9) in Tables 2 and 4. In all three cases, the coefficients hardly change when using the 2SLS estimator. Hence, the endogeneity of urbanization is not driving the results.

5.2.4. Sorting on unobserved variables

Finally, a key concern for many urban and regional studies is the spatial sorting of households on unobserved characteristics. Despite observing some of the key factors of importance for educational decisions, such as parental education and cognitive ability of the child, some other factors such as ambition and non-cognitive skills remain unobserved. One method to assess the importance of such sorting on unobserved variables is provided by Altonji et al. (2005) and Oster (2019). Their procedure relies on the assumption that the sorting on observed variables is informative of sorting on unobserved variables, and in particular that the sorting on unobserved characteristics is weakly less than the sorting on observed characteristics. As such, the degree to which the coefficient of interest and the R^2 change when control variables are added can provide an indication of the potential importance of omitted variable bias.

Altonji et al. (2005) and Oster (2019) argue that sorting on unobserved variables in most cases is less severe than sorting on observable characteristics for two reasons. First of all, they argue that in the most extreme case, researchers observe a random set of variables since they typically have no direct influence over the data collection in surveys and administrative data. In such cases, the observed variables are a random subset, and as such, sorting on observed and unobserved variables should be equally important in generating bias. However, Altonji et al. (2005) argue that researchers typically direct their effort toward obtaining and including the control variables which are seen as being most likely to generate omitted variable bias. Hence, the included control variables in all likelihood contribute more to the omitted variable bias compared with the unobserved variables. For this reason, both Altonji et al. (2005) and Oster (2019) argue that a reasonable upper

30 I would like to express my gratitude to Paul Verstraten (CPB Netherlands Bureau for Economics Policy Analysis) for providing the shape files containing the boundaries of the 1840 municipalities as well as the digitalized 1840 census.

31 A small number of zip codes cannot be matched to the densities of 1840, as they are located on land that has been reclaimed from the sea after 1840. Nonetheless, the instrument is available for 99.6% of the observations.

bound for the degree of selection on unobserved variables is the selection on observed variables. Furthermore, Altonji et al. (2005) argue that when there is a lag between observing the explanatory variables and the outcome measures, any idiosyncratic shocks that occur between observing the explanatory variables and the outcome measure cannot bias the effect of the predetermined explanatory variable. Hence, to the extent that (unobserved) idiosyncratic shocks are present, they provide an additional reason why sorting on unobserved variables may be less important than sorting on observed variables in generating omitted variable bias.

A key decision when applying the methodology of Altonji et al. (2005) and Oster (2019) is the choice for the upper bound on $R^2_{\max}$, that is, how much of the remaining variance the model variables which are unobserved by the researcher would explain. Oster (2019) argues that an upper bound of 1 on the R^2 is too restrictive, as measurement error and the true idiosyncratic error term contained in most models would prevent the researcher from reaching an R^2 of 1, even in cases where all relevant variables are observed. Instead, based on a simulation exercise, Oster (2019) suggests using an R^2 1.3 times larger than the R^2 of the regression with full controls, or to find a reasonable upper bound based on earlier research. For instance, in the case of child outcomes, Oster (2019) argues that sibling correlations may provide a reasonable upper bound for the explanatory power of environmental and family characteristics. Given the decision with respect to the $R^2_{\max}$, it is then possible to rescale the coefficient to account for the sorting on unobserved variables, using the simplified estimator provided in Oster (2019)[32]:

$$\beta^* = \tilde{\beta} - \left[\dot{\beta} - \tilde{\beta}\right] \frac{R_{\max} - \tilde{R}}{\tilde{R} - \dot{R}}, \tag{3}$$

where $\dot{R}$ and $\dot{\beta}$ are obtained from the short regression of density on educational outcomes reported in columns (1) of Tables 2 and 4, and $\tilde{R}$ and $\tilde{\beta}$ are obtained from the full regression model specified in columns (3) of Table 2 and 4. Table 5 shows the adjusted coefficients of Table 2 and 4 when applying the correction for unobserved variables of Oster (2019), under the assumption that selection on observed variables is equal to selection on unobserved variables. The upper half of Table 5 displays the original coefficients and R^2s taken from Tables 2 and 4, which form the inputs for Equation (3). The second half of Table 5 provides the adjusted coefficients obtained from Equation (3) under two different sets of $R^2_{\max}$. In the case of the choice for secondary school level (Table 2), the coefficient remains statistically significant and economically relevant when adjusting for the omitted variable bias, both when using the R^2 based on sibling correlations in the data as well as Oster's suggested value of $1.3 * \tilde{R}$. The second column of Table 5 adjusts the coefficients based on university enrollment (Table 4). In this case, the sibling correlations are hard to interpret as they are only available for the small subset of siblings where both children graduate from an upper secondary school. When the coefficient is adjusted for omitted variable bias using Oster's suggestion of $1.3 * \tilde{R}$, the coefficient diminishes substantially, but remains positive. Taken together, the results in Table 5 indicate that under the assumption that selection on unobserved variables is weakly less than selection on observed

32 The reported coefficients in Table 5 are based on the general estimator as reported in Oster (2019, Section 3.2), which allows for a more general covariance structure of the unobserved variable with the observed explanatory and outcomes variables. However, the adjusted coefficients are virtually identical when using the simplified and general estimator in this study.

Table 5. Importance of selection on unobserved variables

	Table 2	Table 4
Original coefficients taken from Tables 2/4		
R^2 of short regression (column (1) in Tables 2/4): ($\dot{R}$)	0.005	0.003
Coefficient of short regression (column (1) in Tables 2/4): ($\dot{\beta}$)	0.0342***	0.0232***
R^2 of long regression (column (3) in Tables 2/4): ($\tilde{R}$)	0.502	0.064
Coefficient of long regression (column (3) in Table 2/4): ($\tilde{\beta}$)	0.0167***	0.0080***
Coefficient adjusted for unobserved variables following Oster (2019)	Table 2	Table 4
$R_{\max}$ obtained from sibling regressions	0.54	–
Adjusted coefficient (β^*) using $R_{\max}$ based on sibling regressions	0.0153***	–
$R_{\max}$ as $1.3 * \tilde{R}$	0.65	0.083
Adjusted coefficient (β^*) using $R_{\max}$ of $1.3 * \tilde{R}$	0.0106***	0.0028

Note: Adjustment for unobserved variables based on Oster (2019). The R^2 based on sibling regression comes from a regression of the model in Table 2, column (3) on the subset of siblings, with the secondary educational choice of the sibling added as control variable. To provide a conservative inference, the standard errors of Table 2/4 have been used to calculate significance.

variables, omitted variable bias cannot explain the majority of the effect of population density on educational outcomes.

6. Discussion and conclusion

The results in this paper show evidence that population density affects the educational investment decisions of children in the context of the Netherlands. Conditional on observed ability and parental background, children who grow up in urban areas consistently choose to invest more in their education compared with children who grow up in more rural environments. This result is robust across various specifications, subgroups and spatial scope, and cannot be accounted for by the endogeneity of the urbanization measure or sorting on unobserved variables under the assumptions of Oster (2019). Taken together, the results imply that conditional on family characteristics and academic ability, a one log-point increase in density is associated with a 1.4 percentage point increase in the probability that a child attends university. Given the mean university attendance rate of 20%, this implies an elasticity of university attendance with respect to density of 0.07.[33] Expressed differently, moving from a place at the 25th percentile of the density distribution to Amsterdam increases the probability that a child attends university by 3 percentage points.

It is surprising that the potential for agglomeration economies to affect human capital decisions of children and young adults has not received more attention in the literature, particularly given the implications for long-term regional and national growth. Children in rural communities do not seem to enjoy or take the same educational opportunities as children who grow up in urban communities, even in a country such as the Netherlands where the rural areas are relatively accessible from an international perspective. Hence, the findings suggest that differences in educational attainment between urban and rural

33 1.23% of this increase is due to the increase 1.46% in probability that a child graduates from upper secondary school as seen in Table 3, taking into account that on average 84% continues with university. A further 0.15% of the increase is due to the 0.8% increase in university attendance among the children who complete upper secondary school, which contains about 20% of the population (see Table 1).

communities observed in a wide range of countries may reflect more than just the spatial sorting of households.

Finally, one question that remains is which mechanisms drive the increased human capital accumulation of children in urban environments. Sections 2.1 and 2.2 indicated that based on existing literature, one would expect the agglomeration mechanisms to play a role in the early-life human capital decisions, either through an increase in the returns to education in cities or by increasing the availability of schooling and by reducing the commuting or moving costs to attend further education. Furthermore, there might also be a role for network effects as highlighted in Section 2.3. However, a detailed exposition of specific mechanisms in the spirit of De La Roca and Puga (2017) or Dauth et al. (2018) in the case of the urban wage premium is beyond the scope of this paper and will be an interesting avenue for future research.

Acknowledgments

The author wishes to acknowledge the statistical offices that provided the underlying data making this research possible: National Institute of Statistics and Censuses, Argentina; National Statistical Service, Armenia; Bangladesh Bureau of Statistics, Bangladesh; Ministry of Statistics and Analysis, Belarus; National Institute for Statistics and Economic Analysis, Benin; National Institute of Statistics, Bolivia; Central Statistics Office, Botswana; Institute of Geography and Statistics, Brazil; National Institute of Statistics and Demography, Burkina Faso; National Institute of Statistics, Cambodia; Central Bureau of Census and Population Studies, Cameroon; Statistics Canada, Canada; National Bureau of Statistics, China; National Administrative Department of Statistics, Colombia; National Institute of Statistics and Censuses, Costa Rica; National Statistics Office, Dominican Republic; National Institute of Statistics and Censuses, Ecuador; Central Agency for Public Mobilization and Statistics, Egypt; Bureau of Statistics, Fiji; National Institute of Statistics and Economic Studies, France; Federal Statistical Office, Germany; National Institute of Statistics, Guatemala; Institute of Statistics and Informatics, Haiti; Central Statistical Office, Hungary; BPS Statistics Indonesia, Indonesia; Central Statistics Office, Ireland; Central Bureau of Statistics, Israel; National Institute of Statistics, Italy; Statistical Institute, Jamaica; National Bureau of Statistics, Kenya; National Statistical Committee, Kyrgyz Republic; Statistics Bureau, Laos; Bureau of Statistics, Lesotho; Institute of Statistics and Geo-Information Systems, Liberia; National Statistical Office, Malawi; Department of Statistics, Malaysia; National Directorate of Statistics and Informatics, Mali; National Institute of Statistics, Geography, and Informatics, Mexico; National Institute of Statistics, Mozambique; Central Bureau of Statistics, Nepal; National Institute of Information Development, Nicaragua; Statistics Division, Pakistan; Central Bureau of Statistics, Palestine; Census and Statistics Directorate, Panama; National Statistical Office, Papua New Guinea; General Directorate of Statistics, Surveys, and Censuses, Paraguay; National Institute of Statistics and Informatics, Peru; Central Statistics Office, Poland; National Institute of Statistics, Portugal; National Institute of Statistics, Romania; Federal State Statistics Service, Russia; National Institute of Statistics, Rwanda; Government Statistics Department, Saint Lucia; Statistics Sierra Leone, Sierra Leone; Statistical Office of the Republic of Slovenia, Slovenia; Statistics South Africa, South Africa; National Bureau of Statistics, South Sudan; Central Bureau of Statistics, Sudan; Bureau of Statistics, Tanzania; National Statistical Office, Thailand; National Institute of Statistics, Togo; Bureau of Statistics, Uganda; State Committee of Statistics, Ukraine; Bureau of the Census, United States; National Institute of Statistics, Venezuela; General Statistics Office, Vietnam; Central Statistics Office, Zambia and National Statistics Agency, Zimbabwe.

References

Adukia, A., Asher, S., Novosad, P. (2020) Educational investment responses to economic opportunity: evidence from Indian road construction. *American Economic Journal: Applied Economics*, 12: 348–376.

Ahlin, L., Andersson, M., Thulin, P. (2018) Human capital sorting: the "when" and "who" of the sorting of educated workers to urban regions. *Journal of Regional Science*, 58: 581–610.

Alesina, A., Hohmann, S., Michalopoulos, S., Papaioannou, E. (2019) Intergenerational Mobility in Africa (No. w25534). National Bureau of Economic Research.

Altonji, J. G., Elder, T. E., Taber, C. R. (2005) Selection on observed and unobserved variables: assessing the effectiveness of Catholic schools. *Journal of Political Economy*, 113: 151–184.

Autor, D. (2019) Work of the past. Work of the future. American Economic Association: Papers and Proceeding.

Bacolod, M., Blum, B. S., Strange, W. C. (2010) Elements of skill: traits, intelligence, education, and agglomeration. *Journal of Regional Science*, 50: 245–280.

Bau, N. (2019) *Estimating an Equilibrium Model of Horizontal Competition in Education*. London: Centre for Economic Policy Research.

Baum-Snow, N., Freedman, M., Pavan, R. (2018) Why has urban inequality increased? *American Economic Journal: Applied Economics*, 10: 1–42.

Behrens, K., Duranton, G., Robert-Nicoud, F. (2014) Productive cities: sorting, selection, and agglomeration. *Journal of Political Economy* 122: 507–53.

Berry, C. R., Glaeser, E. L. (2005) The divergence of human capital levels across cities. *Papers in Regional Science*, 84: 407–444.

Bosquet, C., Overman, H. G. (2019) Why does birthplace matter so much? *Journal of Urban Economics*, 110: 26–34.

Burgess, S., Greaves, E., Vignoles, A. (2019) School choice in England: evidence from national administrative data. *Oxford Review of Education*, 1–21.

Bütikofer A Peri G. How cognitive ability and personality traits affect geographic mobility. *Journal of Labor Economics, forthcoming*.

Card, D. (1993) Using geographic variation in college proximity to estimate the return to schooling (No. w4483). National Bureau of Economic Research.

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

Chetty, R., Friedman, J. N., Rockoff, J. E. (2014b) Measuring the impacts of teachers I: evaluating bias in teacher value-added estimates. *American Economic Review*, 104: 2593–2632.

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

Chetty, R., Hendren, N., Katz, L. F. (2016) The effects of exposure to better neighborhoods on children: new evidence from the Moving to Opportunity experiment. *American Economic Review*, 106: 855–902.

Cohen-Zada, D., Sander, W. (2008) Religion, religiosity and private school choice: implications for estimating the effectiveness of private schools. *Journal of Urban Economics*, 64: 85–100.

Combes, P. P., Duranton, G., Gobillon, L. (2008) Spatial wage disparities: sorting matters! *Journal of Urban Economics*, 63: 723–742.

Combes, P. P., Gobillon, L. (2015) The empirics of agglomeration economies. In G. Duranton, J. V. Henderson, W. C. Strange (eds) *Handbook of Regional and Urban Economics*, Vol. 5, pp. 247–348. Amsterdam: Elsevier.

Dauth, W., Findeisen, S., Moretti, E., Suedekum, J. (2018). Matching in cities (No. w25227). National Bureau of Economic Research.

Davis, D. R., Dingel, J. I. (2019). A spatial knowledge economy. *American Economic Review*, 109: 153–70.

Deutscher, N. (2020) Place, peers, and the teenage years: long-run neighborhood effects in Australia. *American Economic Journal: Applied Economics*, 12: 220–249.

Duranton, G., Puga, D. (2004) Micro-foundations of urban agglomeration economies. In G. Duranton, J. V. Henderson, W. C. Strange (eds) *Handbook of Regional and Urban Economics*, Vol. 4, pp. 2063–2117. Amsterdam: Elsevier.

Eckert, F., Hejlesen, M., Walsh, C. (2020) *The Return to Big City Experience: Evidence from Refugees in Denmark*. NBER Working Paper.

Frenette, M. (2006) Too far to go on? Distance to school and university participation. *Education Economics*, 14: 31–58.
Gennaioli, N., La Porta, R., Lopez-de-Silanes, F., Shleifer, A. (2013) Human capital and regional development. *The Quarterly Journal of Economics*, 128: 105–164.
Gennaioli, N., La Porta, R., De Silanes, F. L., Shleifer, A. (2014) Growth in regions. *Journal of Economic Growth*, 19: 259–309.
Gibbons, S., Silva, O. (2008) Urban density and pupil attainment. *Journal of Urban Economics*, 63: 631–650.
Glaeser, E. L., Kolko, J., Saiz, A. (2001) Consumer city. *Journal of Economic Geography*, 1, 27–50.
Glaeser, E. L., Mare, D. C. (2001). Cities and skills. *Journal of Labor Economics*, 19: 316–342.
Gould, E. D. (2007) Cities, workers, and wages: a structural analysis of the urban wage premium. *The Review of Economic Studies*, 74: 477–506.
Groot, S. P., de Groot, H. L., Smit, M. J. (2014) Regional wage differences in the Netherlands: micro evidence on agglomeration externalities. *Journal of Regional Science*, 54: 503–523.
Hoxby, C. M., Avery, C. (2012) The missing "one-offs": the hidden supply of high-achieving, low income students (No. w18586). National Bureau of Economic Research.
Jensen, R. (2010). The (perceived) returns to education and the demand for schooling. *The Quarterly Journal of Economics*, 125: 515–548.
Katz, L. F., Goldin, C. D. (2008) *The Race between Education and Technology*. Cambridge: Harvard University Press.
Kelly, M. (2019) *The Standard Errors of Persistence*. London: CEPR
Kelly, M. (2020) *Understanding Persistence*. London: CEPR
Knight, J., Shi, L. (1996) Educational attainment and the rural–urban divide in China. *Oxford Bulletin of Economics and Statistics*, 58: 83–117.
Newbold, K. B., Brown, W. M. (2015) The urban–rural gap in university attendance: determinants of university participation among Canadian youth. *Journal of Regional Science*, 55: 585–608.
Minnesota Population Center. (2019). *Integrated Public Use Microdata Series*, International: Version 7.2 [dataset]. Minneapolis (MN): IPUMS (https://doi.org/10.18128/D020.V7.2).
Oster, E. (2019) Unobservable selection and coefficient stability: theory and evidence. *Journal of Business & Economic Statistics*, 37: 187–204.
De La Roca, J., Puga, D. (2017) Learning by working in big cities. *The Review of Economic Studies*, 84: 106–142.
De la Roca, J. (2017). Selection in initial and return migration: evidence from moves across Spanish cities. *Journal of Urban Economics*, 100, 33–53.
Rosenthal, S. S., Strange, W. C. (2004) Evidence on the nature and sources of agglomeration economies. In G. Duranton, J. V. Henderson, W. C. Strange (eds) *Handbook of Regional and Urban Economics*, Vol. 4, pp. 2119–2171. Amsterdam: Elsevier.
Statistics Netherlands. (2011) Studeren loont. Inkomens van afgestudeerden in het mbo, hbo en wetenschappelijk onderwijs [Online publication].
Statistics Netherlands. (2019). Statline. https://opendata.cbs.nl/statline/#/CBS/en/
Verstraten, P., Verweij, G., Zwaneveld, P. J. (2019) Complexities in the spatial scope of agglomeration economies. *Journal of Regional Science*, 59: 29–55.

Appendix

A. The Urban–Rural Educational Gap

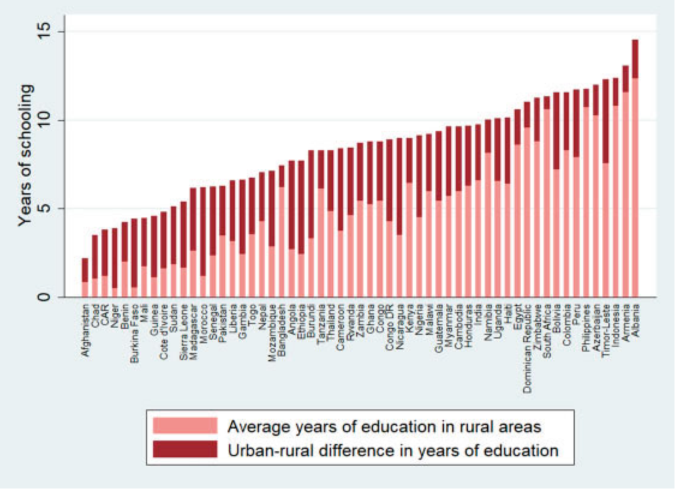

Figure A1. Urban–rural education gap based on 25-year-old males in DHS data. The figure is based on 57 countries included in the Demographic and Health Survey (DHS). Reported are the average years of schooling in urban and rural areas for 25-year-old males. The urban–rural definition follows the definition of DHS, which is based on the urban–rural definition of the country in question. Countries with fewer than 100 18-year old males for either the rural or urban region of the country have been excluded from the original sample of 75 countries.

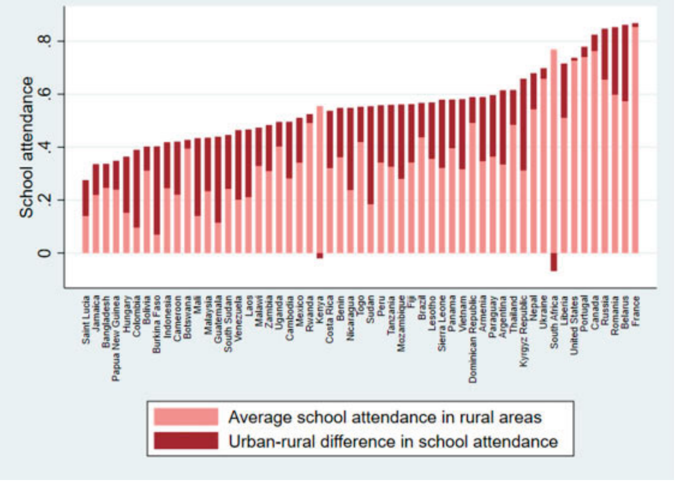

Figure A2. Urban–rural school attendance gap based on 18-year-old males in the IPUMS International database (Minnesota Population Center, 2019). The figure is based on the latest census for 55 countries included in the IPUMS International data which include an urban/rural definition in their census (which excludes most developed countries). The *y*-axis shows the percentage of children that indicate that they were attending school at the time of the census in urban and rural areas for 18-year-old males. The urban–rural definition is based on the definition of the stat-

B. Construction of Urbanization Measure

The administrative data include information on place of residence from January 1995 onward. Hence, I start by calculating the number of individuals registered per zip code on the 5th January 1995. As the next step, I determine the centroid of each zip code using GIS software. The map of the Netherlands with all the centroids is displayed in Figure A1. For each centroid, I determine which other zip codes lie within a 10-km radius of the centroids and add up the population of these zip codes. For zip codes that lie partially in the 10-km radius, I multiply the share of the zip code area covered by the 10-km radius with the population of that particular zip code.[34,35] The average zip code has 43 other zip-codes within a 10-km radius and even the first percentile of zip codes has 8 other zip-codes within a 10-km radius. The correlation between the number of individuals living within 5, 10 and 20 km is fairly high at 0.89. Finally, I take the log of the number of individuals within the 10-km radius as urbanization measure. Figure 3 in the main text shows the resulting map of the urbanization measure. In addition, Figure B2 shows a histogram of the urbanization measure.

Figure B1. Map of Dutch zip codes and centroids.

34 This means that I implicitly assume that the population is spread equally across zip codes. However, as zip codes are fairly small (average of 8 km^2) compared to the 10-km radius (314 km^2) radius, this assumption has little effect on the relative differences in urbanization between regions.

35 For some very isolated zip codes (for instance, on the islands in the North), there are no other zip codes within 10 km, and hence the population with 10 km is simply the zip codes own population.

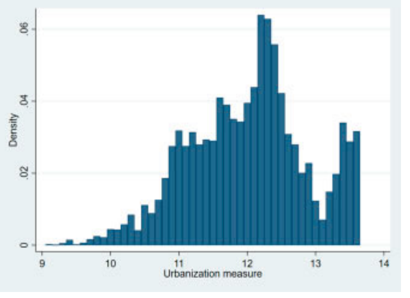

Figure B2. Histogram of urbanization measure at the age 11 years. The histogram of urbanization measure for the cohorts born between 1994 and 1998. Areas with an urbanization score below 9 (containing 230 out of the 631.815 observations in the baseline sample) not displayed here.

C. Additional Robustness Analyses

C.1. Probability of upper secondary school graduation, conditional on being registered at an upper secondary school

Table C1. Effect of urbanization on probability of obtaining a degree in upper secondary school, conditional on enrollment in upper secondary school

	Baseline estimates			Sensitivity analysis		Area-fixed effects		IV
	Ind. level controls	Ind. + Fam. controls	Ind. + Fam controls + test score	Movers excluded	Full parental education	Municipality-fixed effects	Provincial-fixed effects	IV estimates
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Urbanization at age 11 years	0.0075	0.0060***	0.0068**	0.0069**	0.0062*	0.0071	−0.0005	0.0071*
	(0.0041)	(0.0022)	(0.0021)	(0.0022)	(0.0025)	(0.0059)	(0.0028)	(0.0034)
Individual controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Family controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Test score dummies	No	No	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.000	0.03	0.08	0.08	0.09	0.08	0.08	0.08
Number of observation	146.879	146.879	146.879	130.634	60.945	146.879	146.879	146.248

All results are based on individuals born between 1994 and 1998. Dependent variable is whether a child graduates from an upper secondary school, conditional on being enrolled at the third grade of upper secondary school. Individual controls include a gender dummy and cohort dummies. Family controls include dummies for parental education combinations (168 dummies), parental nationality combinations (nine dummies), year-of-birth of oldest parent (40 dummies), family income (interacted with cohort-fixed effects), low-income dummy (interacted with cohort-fixed effects), top-coded income dummy (interacted with cohort-fixed effects) and dummy for insufficient income data (interacted with cohort-fixed effects). For a detailed explanation of the construction of the urbanization measure or the family controls, see Section 3. Column (4) excludes all children who moved municipalities during school-going age. Column (5) removes all parents for whom uncertainty exists about the education level of one or both parents. Column (6) adds municipality-fixed effects for the municipality in which children live at age 11 years (430 dummies), whereas column (7) adds province-fixed effect for the province in which children live at age 11 years (12 dummies). Column (8) instruments for modern densities (based on 1995 population distribution, see Appendix B) with historical densities. All standard errors are clustered on the municipality level. Columns (1–7) are estimated by OLS and column 8 by 2SLS. First stage results of column (8) are reported in Appendix C3 (first stage F-statistic: 596.93).

C.2. Applied university enrollment

Table C2. Effect of urbanization on applied university attendance

	Baseline estimates			Sensitivity analysis			Area-fixed effects		IV
	Ind. level controls	Ind. + Fam. controls	Ind. + Fam controls + GPA	Movers excluded	Cohorts of Table 2 (1994–1998)	Full parental education	Municipality-fixed effects	Provincial-fixed effects	IV estimates
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Urbanization at age 11 years	−0.0233***	−0.0094***	−0.0089***	−0.0095***	−0.0101***	−0.0081***	−0.0090*	−0.0129***	−0.0089***
	(0.0023)	(0.0016)	(0.0016)	(0.0017)	(0.0017)	(0.0017)	(0.0041)	(0.0018)	(0.0025)
Individual controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Family controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
GPA and specialization track	No	No	Yes	Yes	Yes	Yes	Yes	No	Yes
R^2	0.00	0.03	0.07	0.07	0.06	0.07	0.07	0.07	0.07
Number of observation	289.109	289.109	289.109	253.430	139.559	110.921	289.109	289.109	287.763

Note: All results apart from column (5) are based on individuals born between 1989 and 1998. Dependent variable is whether a child attends applied university as highest education within 3 years of graduating from upper secondary school. Individual controls include a gender dummy and cohort dummies. Family controls include dummies for parental education combinations (168 dummies), parental nationality combinations (nine dummies), year-of-birth of oldest parent (40 dummies), family income (interacted with cohort-fixed effects), low-income dummy (interacted with cohort-fixed effects), top-coded income dummy (interacted with cohort-fixed effects) and dummy for insufficient income data (interacted with cohort-fixed effects). For a detailed explanation of the construction of the urbanization measure or the family controls, see Section 3. Column (4) excludes all children who moved municipalities during school-going age. Column (5) limits the sample to the cohorts born between 1994 and 1998, in line with the baseline sample of Table 2. Column (6) removes all parents for whom uncertainty exists about the education level of one or both parents. Column (7) adds municipality-fixed effects for the municipality in which children live at age 11 years (430 dummies), whereas column (8) adds province-fixed effect for the province in which children live at age 11 years (12 dummies). Column (9) instruments for modern densities (based on 1995 population distribution, see Appendix B) with historical densities. All standard errors are clustered on the municipality level. Columns (1–8) are estimated by OLS and column (9) by 2SLS. First stage results of column (9) are reported in Appendix C3 (first stage F-statistic: 1028.11).

C.3. IV first stage

Figure C3.1 displays the population densities of Dutch municipalities in 1840. The number of municipalities (1340) is relatively large compared the modern day number of municipalities (380). The map below is used as input to construct the IV measure, namely the number of individuals in 1840 living within 10 km of all current zip codes, using the procedure outlined in Appendix B. Notice that I make the implicit assumption that the population in 1840 was spread homogenously within municipalities, as the 1840 population statistics are only available at the municipality level. Figure C3.2 shows the resulting IV-density measure for the zip codes.

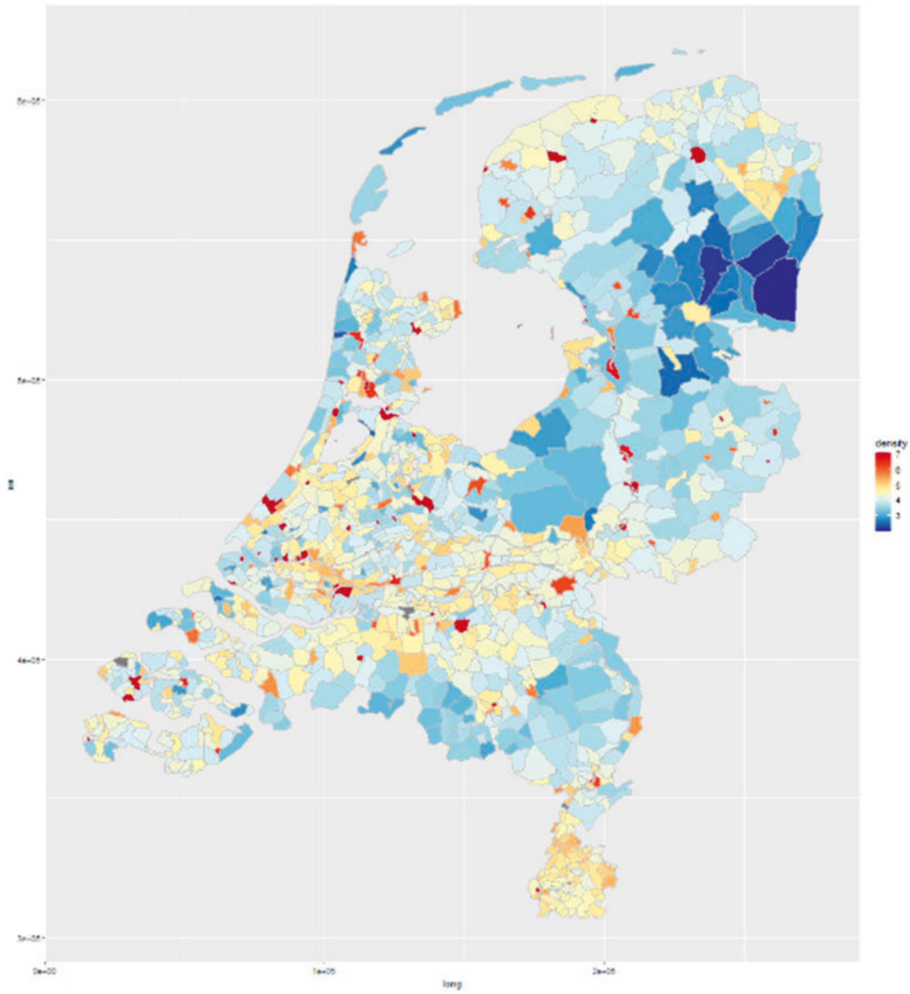

Figure C3.1. Population density per municipality in 1840. When calculating the densities based on the 1840 population map, I implicitly assumed that population is spread homogenously within municipalities. Notice that the Netherlands contained a substantially larger amount of inland water compared with the contemporary Netherlands, due to the land reclamation programs in the 19th and 20th centuries.

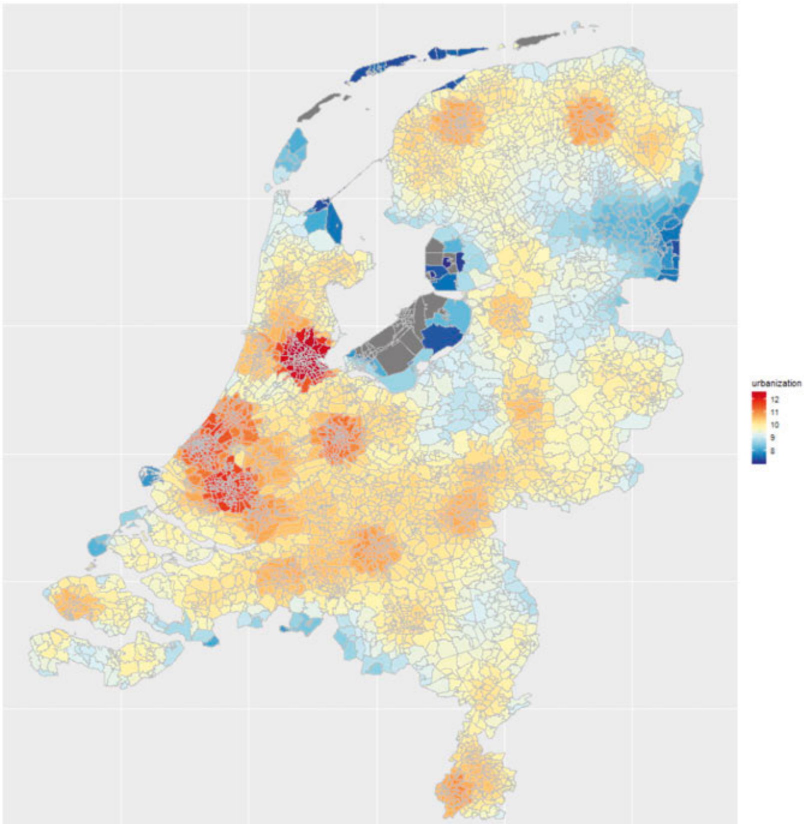

Figure C3.2. Urbanization measure based on 1840 population distribution. The map displays the urbanization measure for each zip code (log of number of people living within 10 km), based the 1840 population distribution. See Appendix B for the construction of the measure and Figure C3.1 for the 1840 population distribution. The gray areas have an urbanization measure below 8. The gray mass in the center of the Netherlands consists of land that has been reclaimed since 1840, and hence had very few to no individuals living with 10 km in 1840, as can be also seen in Figure C3.1.

Table C3. First stage of regressing urbanization based on the 1840 densities on contemporary urbanization measure (based on 1995 densities)

	First stage
Urbanization based on 1840 densities	0.635***
	(0.0045)
F-statistic	2349.94
R^2	0.58
Number of observation	628.396

Note: First stage of 2SLS regression of baseline urbanization measure (based on the 1995 population distribution) on the historical urbanization measure (based on the 1840 population distribution). The first stage of Tables 4 and C2 are virtually identical and not separately reported here, with a slightly smaller *F*-statistics (1028.11) due to the smaller number of observations.

4. The social cost of industrialization: Evidence from the 19th century Swedish iron industry

with Niklas Bengtsson and Adrian Poignant

We would like to express our gratitude to Tessa Bold, Matz Dahlberg, Johannes Haushofer, Mounir Karadja, Andreas Madestam, Jakob Molinder, Eva Mörk, Anders Ögren, Ola Ohlsson, Göran Rydén, Mara Squicciarini, Anna Tompsett, and the participants of the Urban Economics Association conference, Economic history Society Conference, Uppsala Urban Lab seminar and UCBH seminar for their helpful comments and suggestions on earlier versions of this paper.

4.1 Introduction

The transition from artisanal to industrial production is a ubiquitous feature of economic development. Whether in the context of 19th century Europe, China in the 1970s, or developing countries today, the first step toward industrialization often involves the concentration of capital and workers in emergent industrial towns and the parallel decline of artisanal production. This transition can benefit workers since industrialization generally leads to higher productivity and increased wages. However, it could also have adverse consequences for some workers, since it spatially reallocates production and changes the demand for various skills. The net effect of any such industrialization episode on individual workers thus remains unclear. How many workers are permanently displaced by industrialization? In what sectors are the displaced workers reinstated? Do workers migrate to find new jobs? Does industrialization lift all boats in the long run or create a permanent class of "economic losers"?

In this paper, we study these questions in the context of the industrial transformation of the Swedish iron industry in the late 19th century. The iron industry was one of the first to adopt modern large-scale production techniques and is therefore often seen as the harbinger of the second industrial revolution (Mokyr, 1992). During this period, pressures exerted by technological advances and trade liberalization caused hundreds of furnaces and forges to close down, while production moved to increasingly large-scale industrial producers. As a result, the output of the Swedish iron industry increased threefold between 1860 and 1890, whereas the number of production sites declined by more than half.

To study the development of the iron industry, we digitized the annual reports covering the universe of production sites in Sweden between 1860 and 1894. The production data is combined with the population censuses from the North Atlantic Population Project, covering the full Swedish population in 1880, 1890, and 1900. The linked census data allow us to track individual workers and their descendants across time, including those who moved. The ironworkers are identified in 1880 from their listed occupations in the census data and linked to the iron production sites at the parish level. We exploit variation in the timing and location of site closures to estimate the effects of the industrial transformation on individual ironworkers.

We find that the industrial transformation of the iron industry significantly affected the ironworkers. Ironworkers displaced by the closure of ironworks were 23 - 25 percentage points less likely to be employed in the iron industry in 1890, compared to non-displaced ironworkers. At the same time, displaced ironworkers were 25 percentage points more likely to migrate to a different parish, which is more than double the migration rate among the non-displaced groups. Although migration may have alleviated some of the negative effects of displacement, the displaced ironworkers nonetheless tended to end up in lower-paid professions. On average, displaced ironworkers lost roughly 105

- 123 SEK in annual earnings (from a mean of 1012 SEK) compared to their non-displaced counterparts, which implies a relative earnings loss of around 10 percent. Interestingly, we find no evidence that the displacement effects diminish over time. The effects on occupational change, migration, and income are similar in 1890 and in 1900. However, in the long run, the adverse effects are increasingly concentrated among the group of workers who chose not to migrate. Around half of the displaced ironworkers migrated to a different parish between 1880 - 1900. Despite suffering an income loss at the time of the 1890 census, we find that their incomes are no longer significantly different from non-displaced workers by 1900. Thus, our estimates suggest that regional migration dampened the displacement shock in the long run.

Next, we analyze the effects of displacement on the outcomes of children of displaced workers to study the intergenerational spillovers. Although the children of displaced ironworkers are significantly less likely to be employed in the iron industry and more likely to migrate compared to the children of non-displaced workers, their average annual earnings are not statistically different from those of children of non-displaced workers. We thus conclude that this episode of industrialization created a relatively transitory group of "economic losers" with long-lasting effects for the individual workers, but without significantly affecting the subsequent generation.

We interpret these estimates as the causal effects of industrialization. The causal interpretation relies on the assumption that the outcomes of ironworkers observed in 1880 would have evolved similarly if not for the industrial transformation that caused the spatial reallocation of iron production. One concern is that the results may reflect broader changes in economic fortunes, and that factory closings and workers' responses are both symptomatic of a region in decline. We include all non-ironworkers in the estimations to control for such a possibility. We generally find no differences in the occupational changes, migration patterns, and earnings trajectory among non-ironworkers between the regions where the iron industry closes and remains open. The same applies to manufacturing workers not employed in the metal industry. Hence, the results are specific to workers in the iron industry and appear not to be driven by broader regional trends.

A second concern is that workers may engage in anticipatory sorting, which could influence the characteristics of workers employed at a site prior to closure. To address this possibility, we first restrict the sample to workers living in their parish of birth and thus are less likely to have sorted based on short-term prospects. Second, we restrict the analysis to sites that closed after 1885, which makes it less likely that workers were able to anticipate the closure when first observed in 1880. In both cases we find very similar results. Finally, we look at site-level employee counts from the production books and find that they start to decline only around two years before closure, thus further limiting our concerns over sorting.

The findings of the paper contribute to several strands of literature. First of all, the paper contributes to the literature on the effect of technological change on workers' outcomes. Innovations in agriculture have displaced workers globally since the early 19th century and continue to do so in the developing world today (Olmstead and Rhode, 2001). The shift from artisanal spinning and weaving to industrial textile production led to widespread displacement of workers in England (Mantoux 1928) and triggered the so-called Luddite movement; a technology-skepticism that survives to this day (Mokyr 1992). While productivity gains and the reinstatement of labor have historically counterbalanced the displacement effects in the long run (Acemoglu and Restrepo, 2019), the effects on the displaced workers in the short run might be substantial. This paper advances the literature by studying one of history's largest displacement shocks, the transition from artisinal production to industrial production, and by using worker-level data to estimate the effects on occupational changes, income and migration.

Surprisingly, we find earning losses suffered by workers that are similar to, or even lower than, the earnings losses found in the modern plant closure literature (Jacobson, LaLonde and Sullivan, 1993; Eliason and Storrie, 2006; Couch and Placzek, 2010; Seim, 2019). The moderate effects on earnings is noteworthy given the absence of unemployment programs in the late 19th century, which have been shown to reduce the impact of contemporary displacement shocks (Edin et al., 2019). Key to our understanding of this result is the role of local migration. Workers that migrated to a new location experienced no significant decrease in their earnings in the long run. Except for Huttunen, Møen and Salvanes (2018), who highlight a similarly important role for internal migration in coping with worker displacement in Norway, the role of worker migration in adjusting to local shocks has received little attention in work using either historical or contemporary data.

Furthermore, our study contributes to contemporary research on the social and intergenerational effects of worker displacement. Several papers find negative effects on worker health (Browning and Heinesen, 2012; Eliason and Storrie, 2009), spousal health (Marcus, 2013; Bubonya, Cobb-Clark and Wooden, 2017), and parental health and family separations (Mörk, Sjögren and Svaleryd, 2020), which suggest that worker displacement can permanently harm human capital. When it comes to intergenerational effects, the evidence is mixed. Some papers find that children of displaced fathers have lower annual earnings (Oreopoulos, Page and Stevens, 2008), worse health and nutritional outcomes (Schaller and Zerpa, 2019; Pieters and Rawlings, 2020) and lower educational performance (Rege, Telle and Votruba, 2011; Stevens and Schaller, 2011), while other papers find small or insignificant effects (Bratberg, Nilsen and Vaage, 2008; Hilger, 2016; Mörk, Sjögren and Svaleryd, 2020). Our paper contributes to this literature by finding that the displacement of workers during the industrialization of the Swedish iron industry in the 19th century did not appear to have significant scarring effects on the next generation – at least in terms of occupational earnings.

Finally, our paper appears to be one of the first to estimate the effects of worker displacement during the early stages of economic development. Canonical work in development economics focuses on the shift from agriculture to artisinal production (Lewis, 1954; Harris and Todaro, 1970). We are unaware of empirical studies addressing the consequences of job displacement during the next step in development: the shift from artisanal production to large-scale production. By highlighting the importance of spatial worker mobility in transition economies, our findings are relevant for policies that aim to facilitate worker mobility. policies that can effectively facilitate migration in response to adverse shocks, such as the provision of transportation for rural households (Bryan, Chowdhury and Mobarak, 2014), might be effective in reducing the costs of industrial transformation. Conversely, our results suggest that policies that limit internal migration, such as China's Hukou system, might increase the economic and social costs of the industrialization process.
The remainder of the paper continues as follows. Section 4.2 provides an overview of the historical context, the data and the empirical approach. Section 4.3 presents the main findings whereas section 4.4 discusses the effectiveness of the migration response. Finally, section 4.5 concludes.

4.2 Historical context, data and methodology

4.2.1 Background

The Swedish iron industry in 1850 consisted of hundreds of small blast furnaces and finery forges scattered across the countryside. Skilled artisans and craftsmen produced iron with centuries-old techniques, with most production sites operating since at least the 17th century. By 1900, this traditional industry had almost entirely been replaced by large ironworks organized as modern factories and the real output of Swedish iron more than quadrupled (BISOS, 2019). The transformation of the iron industry marks the beginning of a period of rapid industrialization in Sweden's economic history and is typical of the economic changes associated with the second industrial revolution.
The transformation of the Swedish iron industry was a consequence of technological developments beginning in Great Britain (Eriksson, 1957). In 1784 the British inventor Henry Cort patented an improved version of the so-called "puddling process" for converting crude pig iron into malleable bar iron, which revolutionized British iron production in terms of quality and cost (Attman, 1958). Swedish iron producers recognized the benefits of this process at an early stage (see Rydén, 2017), but could not immediately adopt it due to a lack of domestic coal deposits in Sweden, which was required to reach the higher temperatures demanded by the new process. This problem raised concern in the Swedish iron industry, which exported around a quarter of its output to the British market. Figure 4.1 displays the rise of British iron production in the first half of the 19th century. While British outpit of pig iron was compara-

ble to Sweden's in 1800, it exceeded Swedish output by more than an order of magnitude by 1850. Simultaneously, the declining production costs caused bar and pig iron prices in the British markets to fall.

The challenges arising from British puddle iron generated dedicated efforts to

Figure 4.1. Swedish and British iron Industry 1800-1850

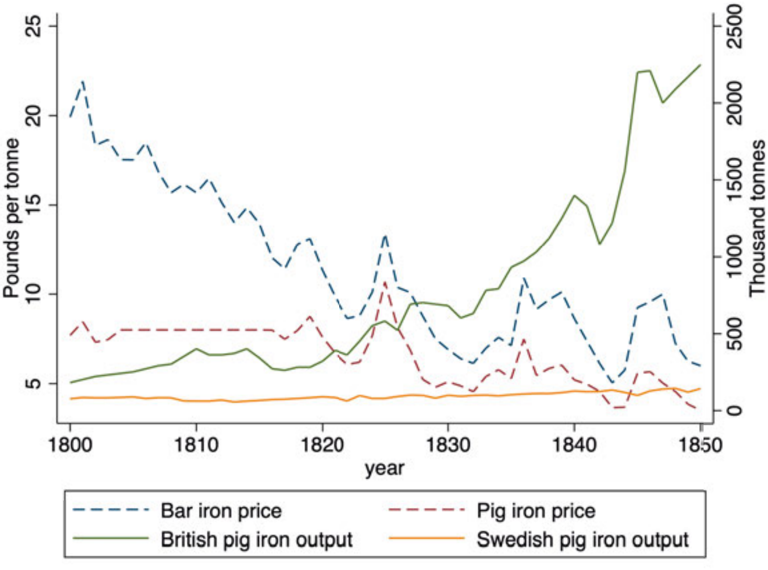

Note: Prices are from Gayer, Schwarts and Rostow (1953) and refer to British markets. British output is from Riden (1977) and Swedish output is from Schön (1988) and Olsson (2007). Pig iron is the main intermediary good in all iron products at this time (such as bar iron, cast iron and steel), which makes it a useful proxy for the industry overall.

close the technological gap in the first half of the 19th century. However, the Swedish iron industry struggled to keep up with the mounting foreign competition (Attman, 1986). The first breakthrough was made in 1845 when the Swedish metallurgist Gustaf Ekman devised a new charcoal-based method for producing bar iron inspired by British techniques. The so-called "Swedish Lancashire process" improved productivity in ironmaking to such an extent that Swedish producers could again begin to compete with their British counterparts. Heckscher (1954) notes the significance of this achievement as the foremost development in "one of the most glorious pages in Swedish economic history". A second breakthrough came with the so-called "ingot steel processes", which include the Bessemer process, invented in 1856, and the Siemens-Martin process, invented in 1865. Parallel to the transformation of the iron sector discussed in this paper, these breakthroughs laid the groundwork for the rise of the contemporary Swedish steel industry.[1]

[1]Here we have recounted the traditional view represented by Heckscher (1954) and Attman (1958, 1986), which emphasises technological breakthroughs in the second half of the 19th

In addition to the technological improvements, the advance of liberal economic policies around the mid-19th century brought essential policy changes that began to erode some pre-industrial institutions that had restricted the growth of the Swedish iron industry. Most importantly, all ironworks had been subject to output restrictions due to a monopolistic policy dating back to the 1700s when Swedish iron dominated the European iron trade (Heckscher, 1949). The so-called "smidesregleringen" was gradually lifted through the legalization of pig iron exports in 1835, increased production quotas for bar iron in 1838, and was finally abolished in 1846 (Isacson, 1997; Olsson, 2007). Other archaic policies sought to protect the traditional iron industry by rationing forest resources for charcoal production and regulating work in detail (Boëthius, 1958). The so-called "hammarsmedsordningarna" and "masmästarordningen" dating back to 1637 and 1649 were both abolished in 1858 (Isacson, 1997), which finally removed a significant part of the regulations which, according to Heckscher (1954), had imposed "insuperable obstacles to industrial reorganization" in the Swedish iron industry.

The liberalizing reforms facilitated the emergence of new organizational structures that accompanied the technological transition. By the end of the 19th century, hierarchical firm structures, which separated the roles of capital owners, industrial engineers, and factory workers, had replaced the once-dominant cooperative associations of craftsmen. These institutional changes, in combination with technological developments, contributed to the Swedish iron industry entering an exceptionally expansive phase around the mid-19th century. Figure 4.2 shows that Swedish pig iron and bar iron output increased four-fold between 1850 and 1900.[2]

The new technologies generally required more extensive facilities and continuous production to operate efficiently, which meant that iron producers needed to scale up their operations to stay competitive. In addition, the expansion of railway networks made larger-scale production possible by allowing the charcoal to be transported over greater distances, the cost of which had exerted a centrifugal force on the geographical distribution of iron production in previous decades. The transformation of the Swedish iron industry in the latter half of the 19th century thus involved the emergence of large-scale ironworks, which displaced hundreds of producers who could not make the necessary ad-

century as the starting point of a new era in Swedish iron making: marking a clear departure from old artisanal traditions in favor of scientifc metallurgy and industrial-scale production. However, Rydén (1998) challenges this narrative as somewhat over-simplified. In his study of the history of iron production in the district of Gästrikland, Rydén argues that the traditional view tends to exaggerate the disruptive effect of specific technologies while ignoring continuous productivity improvements going back to the 18th century.

[2]Between 1850 and 1900 the total output of pig iron increased from 141 to 526 thousand tonnes, while the output of bar iron increased from 97 to 489 thousand tonnes. This translates to an annual growth rate of around 2.6% for pig iron and 3.2% for bar iron. In the preceding period, 1800 to 1850, real output growth was more limited, amounting to annual growth rates of around 1.2% for pig iron and 1.3% for bar iron

Figure 4.2. Production in the Swedish iron industry: 1800-1900

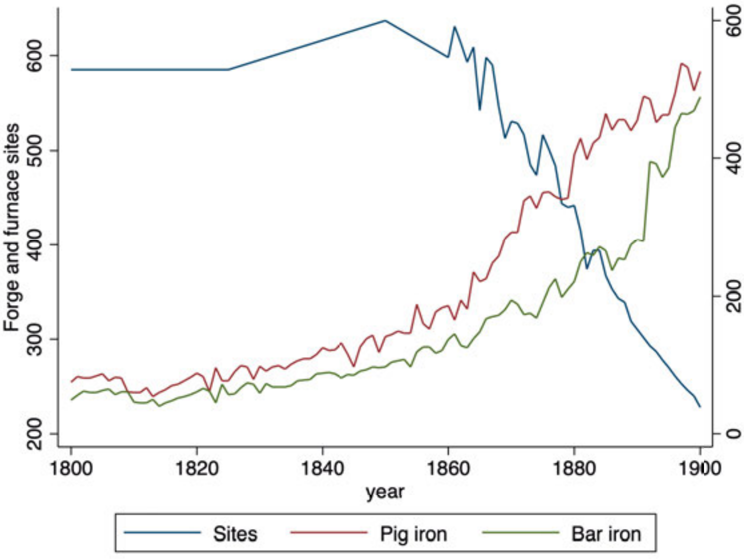

Note: Swedish bar and pig iron output for 1800-1900 are from Olsson (2007). The total number of operational furnaces and finery forges (production sites) is from Statistics Sweden (BiSOS C) for the period 1860-1900, and from Åberg (1999) for the years 1800, 1825 and 1850.

justments. This decline of the old Swedish iron industry is commonly referred to as "the death of the forges" ("den stora bruksdöden"), and is among the more well-documented events of 19th-century Swedish economic history. Figure 4.2 illustrates this development by plotting the number of production sites over time. The number of sites declined by about two-thirds between 1860 and 1900, from roughly 600 sites to around 200 sites. Most of the production sites had been operating for centuries[3], which sets this episode apart from the high turnover observed in many other industries during the industrial revolution.[4] At the same time, the average number of workers per site increased sixfold as the expanding iron industry increasingly concentrated in fewer and larger sites that could accommodate the scale required by new production methods. The simultaneous expansion and concentration of the iron sector meant that surviving sites often turned into small industrial towns as they attracted workers from the surrounding countryside. Boëthius (1958)

[3] A historical register called "Harmens register" tracks the earliest historical mention of the iron production sites. The median iron production site covered in the mining reports first entered the historical records in 1673, with the 10th and 90th percentile first mentioned in 1628 and 1775 respectively. These records highlight that the iron production sites typically had been operating for well over a century prior to the closures in the latter half of the 19th century.

[4] For instance, Juhász, Squicciarini and Voigtländer (2020) find a 93% turnover rate in the cotton industry in France between 1800 - 1840.

notes that by 1913, only 37 ironworks represented over 90 percent of the Swedish iron and steel output.

4.2.2 Data

In order to study the transformation of the Swedish iron industry, we digitalized data from a series of annual production reports covering the period between 1860 and 1894 ("bergmästarrelationer"). These production reports were compiled by the administrators of each mining district and submitted to the National Board of Trade. The production reports contain detailed annual accounts of each district's mining and metal processing industries. They include annual site-specific information on the activities of individual mines and ironworks (including all blast furnaces and finery forges). Unlike an annual census or survey, which would only observe sites operating at a given moment in the year, the production reports include every production site in Sweden during this period regardless of how long it operated. We digitized the original handwritten reports stored in the Swedish National Archives in Stockholm and constructed a panel of year-site observations covering the universe of the Swedish iron and steel production sites between 1860 and 1894.[5]

The digitized data contains information on the output of each production site for every year of operation, as well as employee counts for specific years. Figure 4.3 illustrates a page from the production reports, detailing the output and employee count of all finery forges operating in the Northern mining district in 1861. For our purposes, the production reports offer crucial data on the periods of activity for ironworks that operated during the latter half of the 19th century.

To observe worker outcomes, we use data from the Swedish population censuses of 1880, 1890, and 1900 digitized by the North Atlantic Population Project (NAPP). The censuses contain individual-level data on various characteristics, including occupation, family structure, current place of residence, and place of birth. The NAPP has matched individual census records across censuses based on personal information provided in the census and successfully linked 64% (55%) of males aged 25 - 45 between the 1880 and 1890 (1880 and 1900) censuses.[6] The census data thus allow us to track the outcomes of individual workers and their children between 1880 and 1900.

To identify ironworkers in the census data, we rely on the occupational descriptions from the census (originally recorded by the parish priest) that have been coded into five-digit HISCO codes by the NAPP. The occupational codes used to classify ironworkers are listed in table 4.1 below. Unfortunately, due to incomplete information in the occupational descriptions in the census, it is

[5]The final reporting year varies slightly depending on the district

[6]The linkages are based on name, age, and parish of birth. The Swedish parishes are, on average, relatively small (around 2000 inhabitants), which significantly aids the linking procedures. For more detailed information, we refer the reader to the documentation on the NAPP website. There are no differences in the matching rate between the parishes where the iron industry remains active and the parishes where the iron industry closes, either for the total population or the ironworkers specifically.

Figure 4.3. Extract from production reports

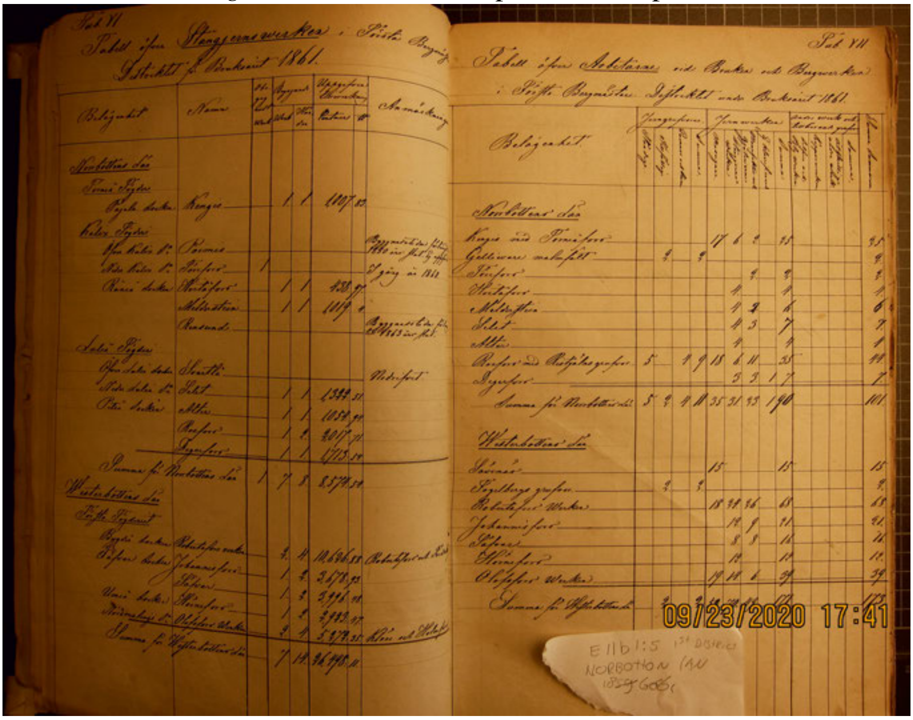

not always possible to distinguish ironworkers from workers in related industries. For instance, some of the occupations we use to identify ironworkers, such as "metal smelters", could refer to workers in other metal industries. However, the dominance of iron within the Swedish metal processing sector in this period means that this problem is not likely to be severe.[7] In other cases, potential ironworkers may be listed in the census as "factory workers" or simply "workers", which could lead us to erroneously exclude workers actually employed in the iron industry. In these cases, we take a conservative approach and only include workers whose listed occupation refers to smelting and forging metals.[8]

The census data does not include explicit information about individual income or social status, but we can approximate the average annual earnings of individual workers using the listed occupations. We combine the Swedish census data with data on annual earnings by occupation provided by Abramitzky, Boustan and Eriksson (2012) for Norway in 1900. The Norwegian earnings data offer the advantage of a much larger number of occupations than the Swedish earnings data for this period. We have verified that the annual earnings reported in the Norwegian data are comparable to Sweden, at least for the

[7]In 1890 iron accounted for roughly 85% of employment in the mining and metal processing sector according to the National Statistics. See BISOS C, Bergshantering 1890.

[8]We test the sensitivity of our results to different occupational specifications in appendix 4.E

Table 4.1. *Occupations classified as iron industry workers*

Occupation	Observations
forgemen	263
hammersmiths	530
metal annealers, temperers and hardeners	5
metal melters and reheaters, and workers in metal melting mills	111
metal rollers	1013
metal smelter and furnacemen	2140
other (specialised) metal workers	174
other blacksmiths, hammermen and forgeing-press operators	2970
other metal workers nec	317
Total	7523

Note:: Number of observations based on the occupational codings in 1880, for individuals who are linked across the 1880 and 1890 census.

occupations for which there is Swedish data (see appendix 4.A). The process used to link the earnings data to occupations in the Swedish census data is described in appendix 4.B.

Since the earnings data from Abramitzky, Boustan and Eriksson (2012) represents a cross-sectional sample of occupations, it cannot be used to capture changes in wages within occupations over time. However, it provides a unified measure of workers' position in the occupational hierarchy, as indicated by average annual earnings around 1900. Therefore, as long as the relative earnings of occupations remained more or less stable between 1880 and 1900, the data allows us to quantify the relative impact of occupational changes over time for individuals and groups of workers.[9]

4.2.3 Methodology

Since the Swedish census data starts in 1880 and the mining books extend to 1894, the focus of this paper will be on the period between 1880 and 1890. In the baseline analysis, we compare the worker and their offspring in the parishes where the ironworks closed down to their counterparts in parishes where at least some ironworks continued to operate. We cross-reference the closure dates reported in the mining books between 1880 and 1890 using external sources and exclude closures that represent a transition from iron production into other closely related types of manufacturing not covered in the

[9]The assumption of constant relative earnings over time is one of the strongest ones we make in this paper. Bagge, Lundberg and Svennilson (1933) suggest that average earnings between 1880 and 1890 grew by somewhat different rates in the categories of metal industries (19.5%), textile industries (15.1%), forestry (18.5%) and agriculture (16.5%). However, despite some differential changes in occupational earnings, we expect the data from Abramitzky, Boustan and Eriksson (2012) to provide a reasonably stable approximation of the relative social status of occupations over this period.

mining book data. Figure 4.5 in the appendix shows a map of parishes where iron production continued uninterrupted and where iron production ceased. Throughout the analysis, we focus on three worker outcomes: (1) whether the worker continues to be employed in the iron industry, (2) the migration responses, and (3) the changes in the average annual income as imputed from worker occupations. To estimate the effect of displacement, we estimate equation 4.1, where i refers to the individual and j refers to the parish (församling) of residence in the 1880 census. The terms $\text{occupation}_{i,1880}$ and $\text{region}_{i,1880}$ represent occupation and region (län) fixed effects as measured in 1880. The coefficient δ captures the general differences in worker outcomes between "treatment parishes" (where the iron production closes down between 1880 and 1890) and "control parishes" (where iron production did not close down between 1880 and 1890).[10] Finally, the coefficient γ captures the corresponding difference specifically for ironworkers.

$$\begin{aligned}\text{Worker\_outcome}_{i,1890/1900} &= \text{occupation}_{i,1880} + \text{region}_{i,1880} + \\ \delta(\text{plant\_death}_{j,1880-1890}) &+ \gamma(\text{iron\_worker}_{1880} \times \text{plant\_death}_{j,1880-1890}) + \varepsilon_i \end{aligned} \tag{4.1}$$

The coefficient γ can be interpreted as the causal effect of the site closures on ironworker outcomes under two essential assumptions. First of all, we assume that the individual characteristics of workers employed in the iron industry in 1880 are similar across treated and control parishes, so that initial sorting (conditional on the occupation and region fixed effects) does not explain differences in outcomes in 1890 or 1900. Second, we assume that there are no unobserved factors that (1) affected iron workers differently than non-iron workers and (2) affect the probability of site closures between 1880 and 1890. These two assumptions are discussed further in the robustness section.
Since the empirical framework compares worker outcomes across parishes it is crucial to select the most relevant control group. Agruably, the closest counterfactual for workers in parishes where iron production closed down between 1880 and 1890 would be workers in parishes where iron production continued between 1880 and 1890. However, since many surviving production sites rapidly expanded, it is unclear whether such a comparison would reflect losses among displaced workers or gains among workers in the control group. Therefore, we also evaluate the outcomes of displaced ironworkers against a broader control group that includes parishes where no iron production was taking place. We use three alternative control groups: (1) parishes where iron production continued between 1880 and 1890, (2) parishes where iron pro-

[10] Note that since we define treatment at the parish-level, we leave out a subset of displacement events beteen 1880 and 1890 that took place in parishes where some production sites survive. In appendix 4.E we present alternative approaches using more detailed geographical information and discuss their potential trade-offs in terms of causal inference.

Figure 4.4. Change in number of iron workers between 1880 - 1890 according to the census by plant closure status

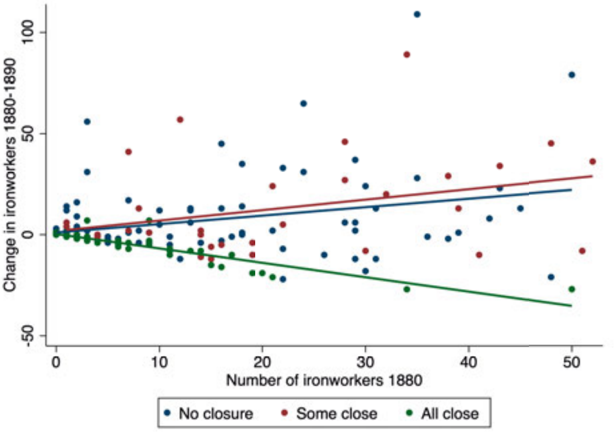

duction took place at any time between 1860 and 1894, and (3) all Swedish parishes.[11] The results are similar across the three groups, which suggests that the choice of control does not matter much for our interpretation.

Finally, we verify that the census data accurately identifies ironworkers and that site closures identified in the production report data are visible as displacement events in the census data. Figure 4.4 plots the change in parish level employment in the iron industry according to the census between 1880 and 1890 for three types of parishes: (1) those where all iron production sites close (green), (2) those where some close and some remain open (red) and (3) those where all remain open (blue). In parishes where the production reports indicate that iron production closed down, the census data also show a clear decline in the iron industry employment. This suggests that the site closures observed in the mining reports align well with the aggregate employment changed observed in the census data.

4.3 Results

4.3.1 Effect on worker outcomes

We start by analyzing the effect of the plant closures on the probability that a worker remains employed in the iron industry after the displacement shock. Table 4.2 shows the estimates obtained by estimating equation 4.1 on a dummy variable indicating whether a worker is employed in the iron industry when ob-

[11] Not all of Sweden has employees in the iron industry. A few parishes have employees in the same occupations, but in other types of metal manufacturing (listed in table 4.1 – f.i. Zinc and Lead) that can also serve as a control group. Furthermore, the different control groups each help to identify the δ estimate in equation 4.1 based on different groups.

served in 1890. Consistent with expectations, the plant closures significantly reduce the likelihood of working in the iron industry in 1890. In parishes where iron production continues, around 62% of the ironworkers observed in 1880 remain employed in the iron industry in 1890. By contrast, this probability is 23 - 25 percentage points lower for workers in parishes where iron production ceased. The sharp and significant drop in the likelihood of continued employment in the iron industry indicates that workers switched to different occupations following the initial displacement shock.

Table 4.2. *Occupation 1880 - 1890*

Outcome:	Working in Iron Industry			
Sample:	male workers observed in 1880 and 1890			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	-0.00637***	-0.00339***	-0.00692***	-0.0105***
	(0.00143)	(0.00123)	(0.00201)	(0.00265)
plant death × ironworker 1880		-0.225***	-0.244***	-0.250***
		(0.0593)	(0.0590)	(0.0593)
R-squared	0.310	0.312	0.363	0.366
Observations	529387	529387	110846	85754

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18-59 at the time of the 1880 census. The lists of occupations considered as iron workers are displayed in table 4.1.

The fact that many individuals remain employed in the iron industry, even though the production sites in their parish of residence closed down, suggests that migration might play an important role. Internal migration was substantial during the period under study, with 25% of adult males moving to a different parish between 1880 and 1890.[12] Table 4.3 shows the effect of site closures on a dummy variable indicating whether the worker moved to a different parish between 1880 and 1890. The closure of the iron industry leads to a 25 percentage point increase in migration rates among ironworkers, representing a doubling of the baseline migration rate. No such differences are visible among workers outside of the iron industry. The out-migration for non-ironworkers is slightly lower in parishes where the iron industry disappeared, suggesting that the changes in migration rates are specific to the iron industry workers.

[12]These figures exclude international migration which was also significant during this period as discussed in Karadja and Prawitz (2019).

Table 4.3. *Migration between 1880 - 1890*

Outcome:	Migration			
Sample:	male workers observed in 1880 and 1890			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	-0.0196**	-0.0228***	-0.00864	0.000544
	(0.00855)	(0.00845)	(0.00998)	(0.00851)
plant death × ironworker 1880		0.248***	0.253***	0.257***
		(0.0491)	(0.0490)	(0.0490)
R-squared	0.143	0.143	0.115	0.111
Observations	529387	529387	110846	85754

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18-59 at the time of the 1880 census. Migration is defined as a change of parish between the 1880 and 1890 census.

Iron industry workers thus responded to the upheaval in their industry by migrating to other parishes and moving out of the iron industry. The question is how effective these responses were in adapting to the displacement shock. Table 4.4 reports the effect of iron site closures on log income at the time of the 1890 census. On average, displaced ironworkers earn around 9 - 10% less in 1890 than their non-displaced counterparts. The difference is stable across the three comparison groups. Furthermore, table 4.4 shows that non-ironworkers in the parishes where the iron industry closed between 1880 and 1890 do not appear to suffer similar earnings losses. The results are similar when using income levels as shown in table 4.11 in the appendix.
Given that the income measure does not capture changes in wages within occupations, the income losses would be more significant if the iron production site closures also decreased the general wage levels in the affected parishes.[13] At the same time, it should be noted that the income of displaced ironworkers may still have increased slightly between 1880 and 1890 in absolute terms, despite their shift into lower-paid occupations. Assuming that earnings in the various occupations grew at the same rate as the average annual real wage in this period (1.6% according to Prado, 2010), our estimates indicate that the earnings of displaced workers would still have increased by about 6.5% over the ten years between 1880 and 1890 (compared to 17.2% for other workers).

Finally, we utilize the linked 1880 and 1900 censuses to analyze whether the effects on occupational changes, migration, and income persist in the long run. The displacement indicator is likely to be somewhat less accurate in this

[13]However, it seems unlikely that declines in the general wage level would have been very substantial, given the high mobility of workers in this period and the absence of out-migration among non-iron workers in the parishes where the iron industry closed.

Table 4.4. *Log income in 1890*

Outcome:	Log income			
Sample:	male workers observed in 1880 and 1890			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	0.00583	0.00715	0.0116	0.0230**
	(0.00844)	(0.00852)	(0.00897)	(0.00945)
plant death × ironworker 1880		-0.0987***	-0.0899***	-0.0984***
		(0.0253)	(0.0278)	(0.0271)
R-squared	0.420	0.420	0.463	0.463
Observations	512003	512003	107095	83238

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18-59 at the time of the 1880 census. The income data is described in Appendix 4.B.

instance, as we do not have information for all the iron site closures between 1890 and 1900. Therefore, some of the control group's ironworkers may also have become displaced when observed in 1900. Nonetheless, the effect of worker displacement on occupational changes, migration and income remain clearly visible in the 1900 census as shown in tables 4.19 - 4.22 in the appendix. The estimated effects of displacement on migration and occupational change between 1880 - 1900 are very similar to the estimates on migration and occupational change between 1880 - 1890, suggesting that most of the migration response and occupational changes took place shortly after the plant closures. At the same time, the similar estimates for earnings suggests that workers seem unable to adapt to the displacement shock even when extending the time frame by ten years.

4.3.2 Robustness

One concern is that differences in outcomes between the displaced and non-displaced ironworkers might reflect broader trends in these regions. In general, the δ coefficient in equation 4.1 should capture such differences. However, this coefficient is estimated based on all non-ironworkers, which might not form a reasonable counterfactual for ironworkers. To investigate whether broader trends within the manufacturing sector may explain our results, we repeat the analyses for manufacturing workers working with non-metallic inputs, such as carpenters and tailors. The results are displayed in tables 4.12 -

4.14 in the appendix. There are no significant effects of the iron plant closures on the migration behavior or incomes of manufacturing workers outside the iron industry. Hence, the results for the ironworkers do not appear to reflect broader trends in the manufacturing industry between the regions but seem specific to the iron industry.
It is possible to take this analysis one step further by estimating separate versions of equation 4.1 for each occupation listed in the census. Tables 4.15 - 4.18 in the appendix shows the result of this exercise for the three outcomes. The migration response, occupational change, and income losses are substantially stronger for those employed in the iron industry in 1880 compared to the other occupations. This auxiliary analysis indicates that the changes within the affected parishes are specific to the iron industry and not driven by changes in the region's broader economic development.
A second concern is that the estimates may be driven by the selection of workers. Selection could be an issue if the most talented workers leave the closing sites a few years before their actual closing date, thus leaving us with a negatively selected sample of employees. One way to investigate this is to compare the estimates for site closures close to when we first observed workers in 1880 with the estimates based on sites which closed in the late 1880's, which would arguably be more difficult to predict in 1880. Appendix table 4.23 shows the estimates for the split sample. There are no significant differences in the displacement effect estimates for the plants that closed in 1880 - 1884 and those that closed in 1885 - 1889, suggesting that worker selection based on plant prospects is unlikely to drive the results. Furthermore, to exclude the possibility that migration of highly skilled workers into the expanding sites is driving the results, we restrict the sample in appendix table 4.24 to workers still living in their parish of birth at the time of the 1880 census. The results are again similar. Finally, figure 4.6 in the appendix shows that the employee level counts in the mining books also do not appear to decline in the years leading up to plant closure, further alleviating our concerns about anticipatory sorting.

4.3.3 Effect on the next generation

One further question is whether the displacement shock also affected the outcomes of subsequent generations. To investigate this, we study a sample of workers' sons observed in 1880.[14] To ensure a sufficiently large sample of children old enough to have observable labor market outcomes, we focus on the outcomes measured in the 1900 census. We restrict the sample to children between 0 and 16 at the time of the 1880 census and under 18 at the time of the plant closure. We estimate the model in equation 4.2 which is similar to

[14]While it would be interesting to also estimate the models on the labor market outcomes of daughters, the labor force participation of women during this period is around 5%, meaning that occupations are observed for only a very small sub-sample.

equation 4.1, but with the occupation and ironworker status on the right-hand side based on the occupation of the household head in 1880. Equation 4.1 additionally includes a control for the age of the child in 1880.

$$\begin{aligned}\text{Childhood\_outcome}_{i,1900} = {} & \alpha_i * \text{occupation father}_{i,1880} + \beta_i * \text{region}_{i,1880} \\ & + \beta_i * \text{age}_{i,1880} + \delta * \text{plant\_death}_{j,1880-1890} \\ & + \gamma * \text{father iron\_worker}_{1880} \times \text{plant\_death}_{j,1880-1890} + \varepsilon_i \end{aligned} \tag{4.2}$$

Tables 4.5 to 4.7 show the effect of worker displacement on the outcomes of their children. The site closures have a negative effect on the probability that the child is employed as an ironworker. Whereas 25% of the sons of non-displaced ironworkers worked in the iron industry by 1900, only 15% of the sons of displaced ironworkers did. At the same time, we also find a significant 22 - 25% increase in migration among the children of displaced workers, which is similar to the effects on their parents. Given the relatively high baseline mobility of the children of ironworkers during this period (44% lives in a different parish in 1900 compared to 1880), the overwhelming majority of children of the displaced workers thus no longer live in their parish of birth at the time of the 1900 census. Hence, the migration response seems to be an essential and lasting response to the plant closures. At the same time, no such migration response is visible for the children of non-iron workers, again suggesting that broader regional trends in migration patterns do not drive the effects.

Finally, table 4.7 presents the estimates of the worker displacement on their children's income in 1900. The coefficient on log income is comparatively small and statistically insignificant across all specifications, suggesting that the worker displacement did not strongly affect the occupational outcomes of their children. Even though the estimates are somewhat imprecise due to the relatively small sample of 130 male children of displaced workers (out of which 101 report an occupation)[15], the estimates are nonetheless significantly different from the parental earning losses estimates in table 4.21.

4.4 Effectiveness of the migration response

The results in the previous section highlight that migration appears to be an important part of the adjustment process following the iron plant closures. One question is how effective migration was as a strategy to mitigate the impact of the displacement shock. While the group that chose to migrate is selective, the results may indicate whether migration can insulate workers against displacement shocks. Around 43% of the displaced ironworkers migrated to a different parish by 1890, which rises to 54% by 1900. Table 4.8 below shows

[15] The percentage not reporting an occupation is very similar among the children of the displaced workers (22%) and the children of the non displaced workers (20%), suggesting that differences in labor force participation are unlikely to drive the results.

Table 4.5. *Employment of children in the iron industry in 1900*

Sample:	male children observed in 1880 and 1900			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	-0.00332	-0.00169	-0.00924***	-0.0181***
	(0.00213)	(0.00207)	(0.00302)	(0.00423)
plant death × ironworker 1880		-0.0829**	-0.0996***	-0.105***
		(0.0324)	(0.0331)	(0.0338)
R-squared	0.0562	0.0565	0.0884	0.0924
Observations	291950	291950	65886	50573

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 0 - 16 at the time of the 1880 census and below age 18 at the time of the plant closure. The lists of occupations considered as iron workers are displayed in table 4.1.

Table 4.6. *Migration of children 1880 - 1900*

Sample:	male children observed in 1880 and 1900			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	-0.0355***	-0.0408***	-0.00552	-0.00252
	(0.0135)	(0.0134)	(0.0136)	(0.0155)
plant death × ironworker 1880		0.270***	0.239***	0.240***
		(0.0525)	(0.0513)	(0.0521)
R-squared	0.119	0.119	0.0977	0.0900
Observations	291950	291950	65886	50573

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 0 - 16 at the time of the 1880 census and below age 18 at the time of the plant closure. Migration is defined as a change of parish between the 1880 and 1890 census.

Table 4.7. *Log income of children in 1900*

Sample:	male children observed in 1880 and 1900			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	-0.0332***	-0.0327***	-0.0261**	-0.0307***
	(0.00924)	(0.00937)	(0.0116)	(0.0111)
plant death × ironworker 1880		-0.0190	-0.0404	-0.0474
		(0.0283)	(0.0294)	(0.0293)
R-squared	0.147	0.147	0.149	0.149
Observations	201243	201243	44531	35011

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 0 - 16 at the time of the 1880 census and below age 18 at the time of the plant closure. The income data is described in Appendix 4.B.

the estimated displacement effect by migration status in 1890. As expected, the effect of displacement on the likelihood of working in the iron industry is smaller among those who migrated, suggesting that some displaced workers found new jobs in the iron industry in their destination parish. The effects on income are nonetheless negative for both groups and not significantly different from each other.

This conclusion changes when we analyze worker outcomes in 1900. The effect of displacement on income is now smaller and statistically insignificant for the ironworkers who migrated to a different parish. Furthermore, the differences between migrating and remaining workers are now statistically significant. These findings suggest that, at least in the long run, migration appears to have been a viable way to mitigate the impact of the original displacement shock. However, since the adverse effects of displacement on income are visible for both migrants and non-migrants in 1890, this adjustment may take some time.

Table 4.8. *Effects of worker displacement on outcomes in 1890 by migration status*

Sample:	non-migrants			migrants		
Outcome 1890:	Iron employment	Log Income	Income	Iron employment	Log Income	Income
plant death	-0.0105***	0.0213**	20.11**	-0.0108***	0.0213**	20.08**
	(0.00268)	(0.00962)	(9.213)	(0.00268)	(0.00962)	(9.204)
plant death × ironworker 1880	-0.306***	-0.117***	-123.7***	-0.178**	-0.0734**	-93.79***
	(0.0760)	(0.0329)	(30.73)	(0.0772)	(0.0345)	(33.67)
R-squared	0.372	0.463	0.505	0.373	0.463	0.505
Observations	83854	83854	83854	83824	83824	83824

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18 - 59 at the time of the 1880 census. The income data is described in Appendix 4.B. The sample is restricted to parishes with iron production in 1880 to save space. The results are virtually identical when using different control groups.

Table 4.9. *Effects of worker displacement on outcomes in 1900 by migration status*

Sample:	non-migrants			migrants		
Outcome 1900:	Iron employment	Log Income	Income	Iron employment	Log Income	Income
plant death	-0.0239***	-0.0151	-10.17	-0.0237***	-0.0150	-10.11
	(0.00615)	(0.0124)	(8.957)	(0.00614)	(0.0125)	(8.966)
plant death × ironworker 1880	-0.408***	-0.166***	-143.9***	-0.147	-0.0501	-50.53
	(0.0490)	(0.0375)	(25.77)	(0.0895)	(0.0356)	(36.25)
R-squared	0.271	0.326	0.378	0.273	0.327	0.378
Observations	56322	56322	56322	56330	56330	56330

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18 - 49 at the time of the 1880 census. The income data is described in Appendix 4.B. The sample is restricted to parishes with iron production in 1880 to save space. The results are virtually identical when using different control groups.

4.5 Conclusion

The transition from artisanal to industrial production is an important milestone in economic development, but the consequences for individual workers remain poorly understood. We study the industrial transformation of the Swedish iron industry in the late 19th century and find that many ironworkers were negatively affected by the closure of the old iron production sites. Ironworkers displaced by site closures between 1880 and 1890 often left the iron industry and ended up in lower-paid professions. We estimate that displaced workers, on average, ended up with 10 percent lower earnings than non-displaced workers. While these losses do not wholly offset the general improvements in wages in this period, the downward movement within the income distribution implies that displaced ironworkers missed out on much of the real wage growth between 1880 and 1890.

Many ironworkers responded to displacement through migration. We find that displaced ironworkers were 25 percentage points more likely to migrate to a different parish than non-displaced workers to cope with the local labor market shock. This response appears to have been effective in the long run. We find that the initial earnings losses associated with displacement are similar for migrants and non-migrants. However, one decade later, the earnings of the displaced ironworkers who chose to migrate are similar to those of non-displaced ironworkers. Furthermore, we find no adverse effects on the earnings of children of displaced workers. Therefore, while industrialization negatively affected the displaced workers, it appears that it did not generate long-term scarring effects for the subsequent generations.

It should be noted that migration is most likely a costly response to the displacement shock in itself, as it typically involves leaving family and social networks behind. In addition, displaced iron-workers are more likely to permanently leave their parish of birth, which could carry particularly large social costs. These social and economic costs of the migration response are difficult to measure. However, they may constitute an essential part of the overall welfare effects of displacement, which calls for more research.

The occupational earning losses we find are comparable to the earnings losses found in the contemporary plant closure literature, which is remarkable given the lack of reeducation and training programs to mitigate the costs of job loss in this period. One explanation may be the importance of regional migration as an adaptation mechanism. Furthermore, as noted by Söderberg (2000), the social consequences of the destruction of the old Swedish iron industry must be interpreted in light of the general industrial expansion happening simultaneously. This sets the event studied in this paper apart from many similar examples of mass displacement, such as the steel crises that affected the English midlands, the Rhur area in Germany, the Rust Belt in the US, and the Swedish Bergslagen in the 1970s and 1980s, where the plant closures coincided with a wider contraction of the manufacturing industry.

Since this is the first paper to quantitatively study the consequences of worker displacement during the industrial revolution, we cannot compare our estimates to those of other countries. Nevertheless, the findings of this paper provide a rational motivation for the early social mobilization against technological change, such as the Luddite movement in Great Britain, as we find that displaced workers suffer permanent income losses. At the same time, the unions and political parties that grew out of Scandinavia's industrial revolution were less technology-averse than the United Kingdom. Historical records indicate that opposition to technology was seldom explicit in Swedish union-orchestrated demonstrations in the early 20th century (Molinder, Karlsson and Enflo, 2021). As such, the direct negative consequences of industrialization for Swedish workers may still have been relatively mild compared to the experiences in other European countries.

References

Åberg, Kjell G. 1999. *Tusen svenska hamrar och järnbruk, anlagda o. 1540-1870, samt hundra bergsmanshamrar.* . 4 [rättade ed., KG Data, Mellerud.

Abramitzky, Ran, Leah Platt Boustan, and Katherine Eriksson. 2012. "Europe's tired, poor, huddled masses: Self-selection and economic outcomes in the age of mass migration." *American Economic Review*, 102(5): 1832–1856.

Acemoglu, Daron, and Pascual Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives*, 33(2): 3–30.

Attman, Arthur. 1986. "Svenskt Jarn och Stal 1800-1914." *Jernkontorets bergshistoriska skriftserie*, 21: 67.

Attman, Artur. 1958. *Fagerstabrukens Historia: Adertonhundratalet.* Almqvist & Wiksell.

Bagge, Gosta, Erik Lundberg, and Ingvar Svennilson. 1933. *Wages, cost of living and national income in Sweden, 1860-1930. Vol. II: Wages in Sweden 1860-1930.* PS King and Son Ltd, London.

Boëthius, Bertil. 1958. "Swedish iron and steel, 1600–1955." *Scandinavian Economic History Review*, 6(2): 144–175.

Bratberg, Espen, Øivind Anti Nilsen, and Kjell Vaage. 2008. "Job losses and child outcomes." *Labour Economics*, 15(4): 591–603.

Browning, Martin, and Eskil Heinesen. 2012. "Effect of job loss due to plant closure on mortality and hospitalization." *Journal of health economics*, 31(4): 599–616.

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. 2014. "Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh." *Econometrica*, 82(5): 1671–1748.

Bubonya, Melisa, Deborah A Cobb-Clark, and Mark Wooden. 2017. "Job loss and the mental health of spouses and adolescent children." *IZA Journal of Labor Economics*, 6(1): 1–27.

Couch, Kenneth A, and Dana W Placzek. 2010. "Earnings losses of displaced workers revisited." *American Economic Review*, 100(1): 572–589.

Edin, Per-Anders, Tiernan Evans, Georg Graetz, Sofia Hernnäs, and Guy Michaels. 2019. "Individual consequences of occupational decline." Centre for Economic Policy Research (CEPR) Discussion Paper DP13808.

Eliason, Marcus, and Donald Storrie. 2006. "Lasting or latent scars? Swedish evidence on the long-term effects of job displacement." *Journal of Labor Economics*, 24(4): 831–856.

Eliason, Marcus, and Donald Storrie. 2009. "Job loss is bad for your health–Swedish evidence on cause-specific hospitalization following involuntary job loss." *Social science & medicine*, 68(8): 1396–1406.

Eriksson, Gösta A. 1957. "The Decline of the Small Blast-furnaces and Forges in Bergslagen After 1850: With Special Reference to Enterprises in the Valley of Kolbäck River." *Geografiska Annaler*, 39(4): 257–277.

Gayer, Arthur David, Anna Jacobson Schwarts, and Walt Whitman Rostow. 1953. *The growth and fluctuation of the British economy, 1790-1850: an historical, statistical, and theoretical study of Britain's economic development.* Clarendon Press.

Harris, John R, and Michael P Todaro. 1970. "Migration, unemployment and development: a two-sector analysis." *The American economic review*, 60(1): 126–142.

Heckscher, E F. 1949. "Sveriges ekonomiska historia från Gustav Vasa: Det moderna Sveriges Grundläggning (Vol. 2: 1)." *Stockholm: Albert Bonniers Förlag*.

Heckscher, Eli Filip. 1954. *An economic history of Sweden.* Harvard University Press.

Hilger, Nathaniel G. 2016. "Parental job loss and children's long-term outcomes: Evidence from 7 million fathers' layoffs." *American Economic Journal: Applied Economics*, 8(3): 247–283.

Huttunen, Kristiina, Jarle Møen, and Kjell G Salvanes. 2018. "Job loss and regional mobility." *Journal of Labor Economics*, 36(2): 479–509.

Isacson, Maths. 1997. "Bergskollegium och den tidigindustriella järnhanteringen." *Daedalus: Tekniska Museets årsbok 1998, Människa, teknik, industri*, 43–58.

Jacobson, Louis S, Robert J LaLonde, and Daniel G Sullivan. 1993. "Earnings losses of displaced workers." *The American economic review*, 685–709.

Juhász, Réka, Mara P Squicciarini, and Nico Voigtländer. 2020. "Technology Adoption and Productivity Growth: Evidence from Industrialization in France." National Bureau of Economic Research.

Karadja, Mounir, and Erik Prawitz. 2019. "Exit, voice, and political change: Evidence from Swedish mass migration to the United States." *Journal of Political Economy*, 127(4): 1864–1925.

Klovland, Jan Tore. 2004. "Historical exchange rate data 1819–2003." Norges Bank Occasional Papers 1.

Lewis, William Arthur. 1954. "Economic development with unlimited supplies of labour." *The Manchester School*, 22(2): 139–191.

Mantoux, Paul. 1928. *The Industrial Revolution in the Eighteenth Century.* Jonathan Cape, London.

Marcus, Jan. 2013. "The effect of unemployment on the mental health of spouses–Evidence from plant closures in Germany." *Journal of health economics*, 32(3): 546–558.

Mokyr, Joel. 1992. *The lever of riches: Technological creativity and economic progress.* Oxford University Press.

Molinder, Jakob, Tobias Karlsson, and Kerstin Enflo. 2021. "More power to the people: Electricity adoption, technological change, and labor conflict." *The Journal of Economic History*, 81(2): 481–512.

Mörk, Eva, Anna Sjögren, and Helena Svaleryd. 2020. "Consequences of parental job loss on the family environment and on human capital formation-Evidence from workplace closures." *Labour Economics*, 67: 101911.

Olmstead, Alan L, and Paul W Rhode. 2001. "Reshaping the landscape: the

impact and diffusion of the tractor in American agriculture, 1910-1960." *Journal of Economic history*, 663–698.
Olsson, Fredrik. 2007. "Järnhanteringens dynamik: Produktion, lokalisering och agglomerationer i Bergslagen." PhD diss. UmeåUniversity: Department of Economic History.
Oreopoulos, Philip, Marianne Page, and Ann Huff Stevens. 2008. "The intergenerational effects of worker displacement." *Journal of Labor Economics*, 26(3): 455–483.
Pieters, Janneke, and Samantha Rawlings. 2020. "Parental unemployment and child health in China." *Review of Economics of the Household*, 18(1): 207–237.
Prado, Svante. 2010. "Nominal and real wages of manufacturing workers, 1860–2007." *Exchange rates, prices, and wages*, 1277–2008.
Rege, Mari, Kjetil Telle, and Mark Votruba. 2011. "Parental job loss and children's school performance." *The Review of Economic Studies*, 78(4): 1462–1489.
Riden, Philip. 1977. "The output of the British iron industry before 1870." *The Economic History Review*, 30(3): 442–459.
Rydén, Göran. 1998. "Lancashiresmide och omvandling. Teknisk utveckling, strukturomvandling och produktionstillväxt i Gästrikland under 1800-talet." *Med Hammare och Fackla*, 35: 105–151.
Rydén, Göran. 2017. "Responses to coal technology without coal. Swedish iron making in the nineteenth century." In *The Industrial Revolution in Iron*. 111–127. Routledge.
Schaller, Jessamyn, and Mariana Zerpa. 2019. "Short-run effects of parental job loss on child health." *American Journal of Health Economics*, 5(1): 8–41.
Schön, Lennart. 1988. "Historiska Nationalräkenskaper för Sverige: Industri och Hantverk 1800–1980 Vol LIX."
Seim, David. 2019. "On the incidence and effects of job displacement: Evidence from Sweden." *Labour Economics*, 57: 131–145.
Söderberg, Johan. 2000. "Sociala och demografiska verkningar av bruksdöden i Värmland." *Den moderna människans uppkomst och andra uppsatser. Carlssons, Stockholm*, 99–131.
Stevens, Ann Huff, and Jessamyn Schaller. 2011. "Short-run effects of parental job loss on children's academic achievement." *Economics of Education Review*, 30(2): 289–299.

Acknowledgements

The author wishes to acknowledge the statistical offices that provided the underlying data making this research possible:
- Sweden 1880 Census: The Swedish National Archives, Umeå University, and the Minnesota Population Center. National Sample of the 1880 Census of Sweden, Version 1.0. Minneapolis: Minnesota Population Center [IPUMS], 2014.
- Sweden 1890 Census: The Swedish National Archives, Umeå University, and the Minnesota Population Center. National Sample of the 1890 Census of

Sweden, Version 1.0. Minneapolis: Minnesota Population Center [IPUMS], 2010.
- Sweden 1900 Census: The Swedish National Archives, Umeå University, and the Minnesota Population Center. National Sample of the 1900 Census of Sweden, Version 2.0. Minneapolis: Minnesota Population Center [IPUMS], 2010.

Appendices

4.A Comparing Norwegian and Swedish income data

In this paper, we use Norwegian data on average annual earnings by occupation in 1900 to quantify the economic significance of changes in occupational status. The use of Norwegian data in the Swedish context is not ideal since the relative earnings of different occupational groups could differ between the two countries. However, given that data on the annual earnings of Swedish workers is not available at a sufficiently detailed level, we use Norwegian data as the second-best option. That said, the Norwegian data likely offers a good approximation due to the large structural similarities and high level of integration between the Swedish and Norwegian economies at the time (Sweden and Norway were in a political, economic and monetary union until 1905). Table 4.10 reports the average annual earnings of different categories of Norwegian workers in 1900 provided by Abramitzky, Boustan and Eriksson (2012), as well as the average annual earnings of similar categories of Swedish workers in 1900 provided by Bagge, Lundberg and Svennilson (1933). Table 4.10 suggests that, for those worker categories where a comparison is possible, the Norwegian and Swedish earnings data appear to be well aligned.

Table 4.10. *Annual earnings in 1900 by occupation: Swedish and Norwegian data*

	Norwegian data		Swedish data	
Category	Description	Kr	Description	Kr
Metal workers	Metal smelter and furnacemen; Metal casters and workers in metal casting plants; Metal annealers, temperers and hardeners; Metal melters and reheaters, and workers in metal melting mills; Metal rollers	1294	Metal manufacturing and engineering industries	1229
	Hammersmiths; Forgemen; Other blacksmiths, hammermen and forgeing-press operators	1202		
	Other metal workers n.e.c.	1103		
Textile workers	Textile workers, specialisation unknown	701	Workers in textile industries	655
Farmers	general farmers and farmers nfs	1470	Farmers (owner-occupiers)	1457*
	Farm workers, specialisation unknown	655	Male agricultural day-laborers	544
	Husbandmen or cottars	425	Farm servants (statare)	472

Note The Norwegian data is from Abramitzky, Boustan and Eriksson (2012) (converted to Swedish Kronor using the exchange rate for year 1900 provided by Klovland (2004)). The Swedish data is from Bagge, Lundberg and Svennilson (1933) who collected information on wages between 1860 and 1930. The reported income if statare is adjusted for in-kind payments. The annual earnings of male agricultural day-laborers is calculated by multiplying the daily wage rate (markegång rate) by 297 (assuming a six day working week).

*To calculate the annual earnings of owner-occupier farmers we divide the total value of the 1900 harvest (after correcting for input costs following Abramitzky, Boustan and Eriksson (2012)) by the number of farm owners in Sweden in 1900 obtained from Bidrag till Sveriges officiella statistik (BiSOS) 1851-1917.

4.B Imputation of wage data

The linkage of the historical occupation to the wage data comes from Abramitzky, Boustan and Eriksson (2012). This data is coded according to the same historical 5-digit occhisco classification used by IPUMS International for the digitalized Swedish censuses, which means that the historical occupations are standardized according to the same 5-digit occupational classification. For 73% of the observations, we directly observe the income data from Abramitzky, Boustan and Eriksson (2012), out of which a substantial fraction (73%) are farming households. For the remaining 27%, we average the wages from occupations within the same 1-/2-/3- digit historical ISCO classification with available wage data, weighting the occupation by the frequency at which they are observed. Using the 3-digit classification pushes the matching rate up to 95% and with the 2-digit classification to 99%. Finally, we use wages on the 1-digit ISCO classification level to obtain a 100% matching rate on wages. The occupation for which we have to rely on the 1-digit occupational structure are not very obviously selected, with the 5 largest occupations being accountants, charcoal burners, hotel managers, musicians and "other jewelry workers". However, the results are virtually identical if we restrict attention to the occupations for which wages are known at the 3-digit level.

To investigate the applicability of the Norwegian wage structure to the Swedish context, we compared the Norwegian wages with the wage levels which are known for a select number of Swedish occupations in appendix 4.A. Finally, for some individuals the occupation is classified as missing, but we know from the occupational string that these individuals were living of church charity or begging. For these individuals, we assign the income to be the lowest observed income among the occupations, which are domestic servants.

4.C Iron production locations

Figure 4.5. Iron production site closures 1880-1890

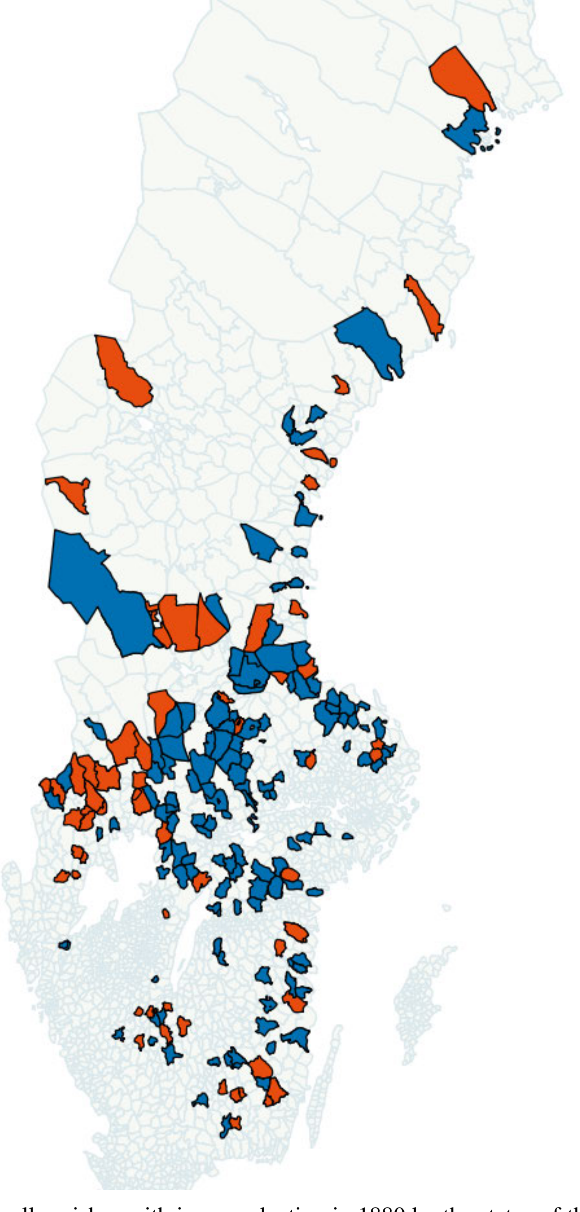

Note: Figure displays all parishes with iron production in 1880 by the status of their iron production in 1890. Blue indicators that all iron production sites survived, whereas red indicated that all production sites closed. Parishes where some sites closed and some remained open are not displayed on the map.

4.D Additional outcomes

4.D.1 Effect of displacement on 1890 income of ironworkers

Table 4.11. *Income in 1890*

Outcome:	Income			
Sample:	male workers observed in 1880 and 1890			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	3.353	5.004	9.774	21.83**
	(7.587)	(7.679)	(8.522)	(8.910)
plant death × iron-worker 1880		-122.6***	-104.5***	-111.0***
		(24.25)	(26.28)	(25.86)
R-squared	0.495	0.495	0.506	0.505
Observations	512003	512003	107095	83238

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ Sample restricted to those between ages 18-60 at the time of the 1880 census. The income data is described in Appendix 4.B.

4.D.2 Effects of plant closures on craftsmen in 1890

Table 4.12. *Migration between 1880-1890*

Outcome:	Migration			
Sample:	male workers observed in 1880 and 1890			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	-0.0196**	-0.0180**	-0.00361	0.00442
	(0.00855)	(0.00854)	(0.00981)	(0.00874)
plant death × craftsman 1880		-0.0354**	-0.0292	-0.0120
		(0.0149)	(0.0207)	(0.0155)
R-squared	0.143	0.143	0.114	0.110
Observations	529387	529387	110846	86495

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Craftsmen are defined as three occupational groups, namely "Boot and shoe makers and repairers", "Tailors and tailoresses", "Carpenters".

Table 4.13. *Income in 1890*

Outcome:	Income			
Sample:	male workers observed in 1880 and 1890			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	3.353	4.091	8.793	18.56**
	(7.587)	(7.832)	(8.650)	(9.288)
plant death × craftsman 1880		-16.67	-12.56	-5.158
		(11.50)	(12.25)	(12.91)
R-squared	0.495	0.495	0.506	0.504
Observations	512003	512003	107095	83947

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Craftsmen are defined as three occupational groups, namely "Boot and shoe makers and repairers", "Tailors and tailoresses", "Carpenters".

Table 4.14. *Log income in 1890*

Outcome:	Log income			
Sample:	male workers observed in 1880 and 1890			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	0.00583	0.00697	0.0110	0.0202**
	(0.00844)	(0.00869)	(0.00908)	(0.00967)
plant death × craftsman 1880		-0.0259**	-0.0176	-0.0103
		(0.0108)	(0.0114)	(0.0120)
R-squared	0.420	0.420	0.463	0.463
Observations	512003	512003	107095	83947

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Craftsmen are defined as three occupational groups, namely "Boot and shoe makers and repairers", "Tailors and tailoresses", "Carpenters".

4.D.3 Effects of plant closure by occupation

Table 4.15. *Estimates of γ for each occupation separately on probability of working in the iron industry.*

	gamma
Metal smelter and furnacemen	-.3212057
Other blacksmiths, hammermen and forgeing-press operators	-.0838042
Factory labourers (unspecified)	-.0409795
Sawyers and other titled wood/sawmill operatives	-.0098577
Day labourers (e.g., journalier)	-.0042322
Labourers nfs	.0004746
Military, rank unknown	.0018223
Worker nfs	.0019844
Seamen	.0024938
General farmers and farmers nfs	.0030728
Foremen and supervisors, nec	.0031168
Tailors and tailoresses	.0034527
Painters, construction	.0037315
Teachers (primary)	.0038176
Husbandmen or cottars	.0039149
Boot and shoe makers and repairers	.004046
Farm workers, specialisation unknown	.0045755
Bookkeepers and bookkeeping clerks	.0046891
Foresters and wood wardens	.0048584
Dealer, merchant etc. (wholesale and retail trade)	.0078011
Carpenters	.0082665
Millers and related workers	.0120077
Blacksmiths	.0502774
Total	-.0145946

Note: estimates obtained by estimating equation 4.1 seperately for all occupation with at least 50 observations in the communities where the iron plants close between 1880 and 1890.

Table 4.16. *Estimates of γ for each occupation separately on the probability of migration.*

	gamma
Other blacksmiths, hammermen and forgeing-press operators	.3071771
Metal smelter and furnacemen	.2855957
Labourers nfs	.1532894
Bookkeepers and bookkeeping clerks	.1301142
Blacksmiths	.1076954
Millers and related workers	.0376488
Husbandmen or cottars	.0229289
Worker nfs	.0185562
Carpenters	.0171508
Teachers (primary)	.0153584
Dealer, merchant etc. (wholesale and retail trade)	.0097882
General farmers and farmers nfs	.0057026
Foresters and wood wardens	.0044715
Day labourers (e.g., journalier)	-.0008961
Foremen and supervisors, nec	-.0095637
Seamen	-.0223469
Tailors and tailoresses	-.0445002
Factory labourers (unspecified)	-.050493
Boot and shoe makers and repairers	-.0545428
Painters, construction	-.0622196
Farm workers, specialisation unknown	-.0897436
Military, rank unknown	-.0928522
Sawyers and other titled wood/sawmill operatives	-.1192087
Total	.0247439

Note: estimates obtained by estimating equation 4.1 seperately for all occupation with at least 50 observations in the communities where the iron plants close between 1880 and 1890.

Table 4.17. *Estimates of γ for each occupation separately on 1890 log income.*

	gamma
Metal smelter and furnacemen	-.1411351
Other blacksmiths, hammermen and forgeing-press operators	-.0876114
Dealer, merchant etc. (wholesale and retail trade)	-.0851558
Seamen	-.0707601
Husbandmen or cottars	-.0654878
Factory labourers (unspecified)	-.0579437
Blacksmiths	-.0490396
Carpenters	-.0379942
Millers and related workers	-.0370944
Boot and shoe makers and repairers	-.0306444
Worker nfs	-.0271859
Tailors and tailoresses	-.0087069
Painters, construction	-.0079238
Sawyers and other titled wood/sawmill operatives	-.0062354
Bookkeepers and bookkeeping clerks	-.005761
Foremen and supervisors, nec	-.0021034
Farm workers, specialisation unknown	.010238
Foresters and wood wardens	.0191371
Military, rank unknown	.0201311
Labourers nfs	.0296351
Teachers (primary)	.0309602
Day labourers (e.g., journalier)	.0437726
General farmers and farmers nfs	.0700754
Total	-.0216015

Note: estimates obtained by estimating equation 4.1 seperately for all occupation with at least 50 observations in the communities where the iron plants close between 1880 and 1890.

Table 4.18. *Estimates of γ for each occupation separately on 1890 income.*

	gamma
Metal smelter and furnacemen	-151.2554
Dealer, merchant etc. (wholesale and retail trade)	-137.8477
Other blacksmiths, hammermen and forgeing-press operators	-119.6495
Factory labourers (unspecified)	-110.9953
Seamen	-76.16467
Husbandmen or cottars	-57.65844
Blacksmiths	-55.22545
Millers and related workers	-51.86971
Boot and shoe makers and repairers	-33.76601
Painters, construction	-24.89331
Worker nfs	-15.88531
Sawyers and other titled wood/sawmill operatives	-12.15184
Carpenters	-12.05427
Bookkeepers and bookkeeping clerks	-8.811954
Tailors and tailoresses	4.34342
Farm workers, specialisation unknown	17.32788
Military, rank unknown	19.45932
Foremen and supervisors, nec	20.65549
Foresters and wood wardens	36.61045
Labourers nfs	44.02872
Teachers (primary)	48.33202
Day labourers (e.g., journalier)	53.82156
General farmers and farmers nfs	63.80629
Total	-24.34103

Note: estimates obtained by estimating equation 4.1 seperately for all occupation with at least 50 observations in the communities where the iron plants close between 1880 and 1890.

4.D.4 Effects of displacement on iron workers in 1900

Table 4.19. *Occupational changes 1880-1900*

Outcome:	Working in Iron Industry			
Sample:	male workers observed in 1880 and 1900			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	-0.00638**	-0.00346	-0.0129***	-0.0232***
	(0.00283)	(0.00282)	(0.00453)	(0.00604)
plant death × iron-worker 1880		-0.209***	-0.251***	-0.261***
		(0.0570)	(0.0581)	(0.0609)
R-squared	0.200	0.201	0.258	0.266
Observations	355165	355165	74406	58283

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18-49 at the time of the 1880 census. The lists of occupations considered as iron workers are displayed in table 4.1.

Table 4.20. *Migration between 1880-1900*

Outcome:	Migration			
Sample:	male workers observed in 1880 and 1900			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	-0.0126	-0.0163	0.00599	0.0114
	(0.0122)	(0.0122)	(0.0107)	(0.0101)
plant death × iron-worker 1880		0.264***	0.255***	0.263***
		(0.0483)	(0.0483)	(0.0485)
R-squared	0.147	0.147	0.117	0.116
Observations	355165	355165	74406	58283

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18-49 at the time of the 1880 census. Migration is defined as a change of parish between the 1880 and 1900 census.

Table 4.21. *Log income in 1900*

Outcome:	Log income			
Sample:	male workers observed in 1880 and 1900			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	-0.0130	-0.0115	-0.0129	-0.0151
	(0.00997)	(0.0100)	(0.0111)	(0.0124)
plant death × iron-worker 1880		-0.108***	-0.0980***	-0.105***
		(0.0226)	(0.0225)	(0.0230)
R-squared	0.303	0.303	0.333	0.326
Observations	341445	341445	71705	56403

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18-49 at the time of the 1880 census. The income data is described in Appendix 4.B.

Table 4.22. *Income in 1900*

Outcome:	Income			
Sample:	male workers observed in 1880 and 1900			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	-11.46	-9.928	-9.339	-10.13
	(7.279)	(7.343)	(8.048)	(8.954)
plant death × iron-worker 1880		-106.7***	-87.42***	-94.63***
		(20.68)	(20.11)	(20.63)
R-squared	0.377	0.377	0.386	0.378
Observations	341445	341445	71705	56403

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18-49 at the time of the 1880 census. The income data is described in Appendix 4.B.

4.D.5 Robustness: Selection of workers into plants

Table 4.23. *Effects of worker displacement on 1890 outcomes by closing data*

Sample:	Plant closed < 1885			Plant closed >= 1885		
Outcome 1900:	Ironworker	Migration	Log Income	Ironworker	Migration	Log Income
plant death	-0.0106***	0.000445	0.0231**	-0.0105***	0.000488	0.0230**
	(0.00265)	(0.00853)	(0.00945)	(0.00262)	(0.00849)	(0.00946)
plant death × ironworker 1880	-0.285***	0.298***	-0.105***	-0.200**	0.203***	-0.0885**
	(0.0764)	(0.0595)	(0.0342)	(0.0907)	(0.0706)	(0.0422)
R-squared	0.366	0.111	0.463	0.367	0.110	0.463
Observations	85659	85659	83146	85628	85628	83114

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18 - 59 at the time of the 1880 census. The income data is described in Appendix 4.B. The sample is restricted to parishes with iron production in 1880 to save space. The results are virtually identical when using different control groups.

Table 4.24. *Effects of worker displacement on outcomes in 1890 for the sample living in their place of birth in 1880*

Sample:	Those living in place of birth in 1880			
Outcome 1890:	Ironworker	Migration	Log Income	Income
plant death	-0.00884***	-0.00357	26.04*	0.0281*
	(0.00221)	(0.00798)	(10.34)	(0.0112)
plant death × ironworker 1880	-0.196*	0.231**	-110.1**	-0.105**
	(0.0915)	(0.0777)	(33.04)	(0.0388)
R-squared	0.316	0.0916	0.448	0.440
Observations	47299	47299	45697	45697

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18 - 59 at the time of the 1880 census. The lists of occupations considered as iron workers are displayed in table 4.1.

Figure 4.6. Worker counts prior to plant closure

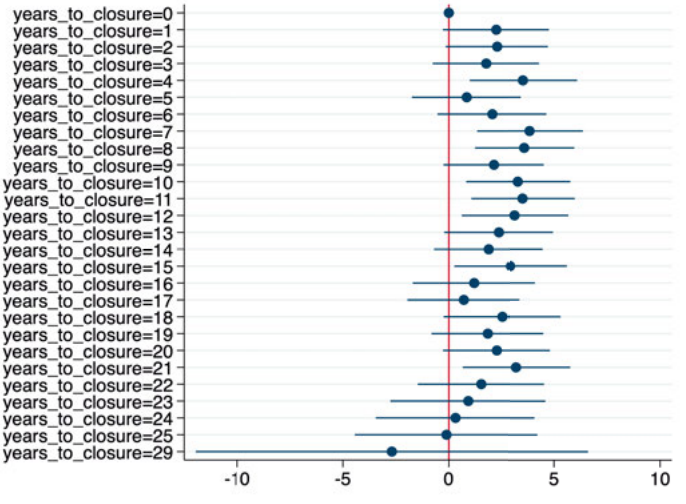

Note:: The figure displays estimates of site-level employment in the years leading up to closure. We run the following regression $emp_{i,t} = \sum_{t=1}^{30} \delta_t + \gamma_i + \varepsilon_{i,t}$, where $emp_{i,t}$ is the number of workers employed at site i in the period t years prior to closure (i.e. the last year of operation), and δ_t captures the the number of workers employed t is years prior to closure relative to the baseline $t = 0$). The sample consists of all sites that closed between 1880 and 1890.

4.E Alternative treatment definitions

One of the main challenges of this paper is to create an accurate mapping between site closures and individual workers. Between 1880 and 1890 almost one hundred production sites closed down, representing rougly one quarter of sites operating in 1880. But, in order to identify the workers who were displaced by these closures, we are forced to draw inferences from the limited information availible in census record. Since we do not have access to perfect information, this process inevitably creates some degree of measurement error in the treatment variable that could bias the estimates. In this section, we present different approaches of identifying displaced workers to address this potential issue.

In all specifications we identify displaced workers using two overlapping criteria: (1) geographical location of residence and (2) listed occupation. Thus, a worker is considered displaced if he resides in a location where iron production is shut down, *and* has an occupation that is connected to the iron industry. In the main specification we use a relatively broad geographical definition (based on the parish of residence), and a narrow occupational definition (see table 4.1). We prefer this specification as it allows us to avoid overly restrictive assumtptions regarding worker (im)mobility, while simultantiously limiting the assignment of treatment to workers who are highly likely to work in the iron sector based on their listed occupations. However, this means that we cannot expoit a handful of site closures between 1880 and 1990 which took place in parishes where some sites continue to operate. Furthermore, the narrow occupational definition implies that we could fail to identify some displaced ironworkers who are listed as general blacksmiths or unspecified factory workers.

In order to evaluate these concerns we experiment using a broader set of occupational descriptions, as well as a narrower geographical definition of treatment. We create a broader occupational definition of ironworkers by also including the occupations listed in table 4.25. On the spatial dimension, we construct a narrower geographical definition using the transcribed place of residence (hemort) listed in the census, rather than the parish. For each parish in the census data containing at least one production site, we list each recorded place name mentioned in the census and (when possible) manually assign it to a specific production site. This allows us to identify individuals residing in the immediate town or village where the site is located, or at the production site itself.

In each case, changing the treatment definition is associated with certain trade-offs. A narrower geographical definition allows us to exploit a larger number of closures, but may also lead to incorrectly discarding displaced workers residing in nearby locations. Similarly, by expanding the occuptional definition we may succeed in including a larger number of displaced ironworkers, but also risk diluting the treatment variable by including workers belonging to other sectors. Tables 4.27 to 4.30 report the key estimates of this

Table 4.25. *Additional occupations classified as iron industry workers*

Occupation	Observations
blacksmiths	15877
factory labourers (unspecified)	9926
metal casters and workers in metal casting plants	1624
metal drawers and workers in metal drawing	64
metal grinders and sharpeners	64
metal turners	453
sheet metal worker, general	2320
Total	30328

Note: Number of observations based on the occupational codings in 1880, for individuals who are linked across the 1880 and 1890 census.

paper when varying the treatment definition. The tables also report the number of displaced workers (treated units) generated by using each definition. Table 4.26 repeats the main specifications reported in the text of this paper, and is included for reference.

Table 4.26. *Main specifications*

Sample:	male workers observed in 1880 and 1900		
Outcome:	Working in Iron Industry	Migration	Log income
Model:	(1)	(2)	(3)
plant death	-0.0105***	0.000400	0.0230**
	(0.00265)	(0.00852)	(0.00945)
plant death × ironworker 1880	-0.250***	0.257***	-0.0986***
	(0.0593)	(0.0490)	(0.0271)
R-squared	0.366	0.111	0.463
Observations	85596	85596	83088
No. displaced workers	221	221	221

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18-49 at the time of the 1880 census. The income data is described in Appendix 4.B.

Table 4.27 uses the transcribed place of residence ("hemort") to define the geographical dimension of the treatment variable. The estimated migration response is significantly smaller compared to the main specification. This is likely due to the inclusion of site closures in parishes where some iron production continued; such workers would be more likely to find similar employment within the parish. Otherwise the results are similar.

Table 4.27. *Restricted spatial definition*

Sample:	male workers observed in 1880 and 1900		
Outcome:	Working in Iron Industry	Migration	Log income
Model:	(1)	(2)	(3)
plant death	-0.00120	0.0430**	-0.00838
	(0.00897)	(0.0183)	(0.0139)
plant death × ironworker 1880	-0.249***	0.118**	-0.0800***
	(0.0584)	(0.0548)	(0.0272)
R-squared	0.366	0.110	0.462
Observations	85596	85596	83088
No. displaced workers	249	249	249

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18-49 at the time of the 1880 census. The income data is described in Appendix 4.B.

In table 4.28 we instead use the transcribed place of residence (hemort) only in parishes where some iron production continued, and otherwise rely on the parish-level definition used in the main specifications. This combined measure allows us to maintain our preferred definition for the most part, without throwing away the variation caused by additional displacement events in parishes where some iron production continued.

Table 4.28. *Combined spatial definition*

Sample:	male workers observed in 1880 and 1900		
Outcome:	Working in Iron Industry	Migration	Log income
Model:	(1)	(2)	(3)
plant death	-0.00964***	0.00170	0.0210**
	(0.00252)	(0.00813)	(0.00888)
plant death × ironworker 1880	-0.269***	0.192***	-0.114***
	(0.0498)	(0.0478)	(0.0226)
R-squared	0.369	0.111	0.463
Observations	85596	85596	83088
No. displaced workers	338	338	338

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18-49 at the time of the 1880 census. The income data is described in Appendix 4.B.

In table 4.29 we limit the geographical dimension of treatment using the transcribed place of residence (hemort), but also expand the occupational definition using the wider set of occupations listed in table 4.25.

Table 4.29. *Restricted spatial and expanded occupational definition*

Sample:	male workers observed in 1880 and 1900		
Outcome:	Working in Iron Industry	Migration	Log income
Model:	(1)	(2)	(3)
plant death	0.00196	0.0290*	0.00354
	(0.00524)	(0.0151)	(0.0154)
plant death × ironworker 1880	-0.142***	0.111**	-0.0836***
	(0.0361)	(0.0429)	(0.0280)
R-squared	0.364	0.110	0.462
Observations	85596	85596	83088
No. displaced workers	464	464	464

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18-49 at the time of the 1880 census. The income data is described in Appendix 4.B.

In table 4.30 we again use the transcribed place of residence (hemort) in combination with the parish-level treatment, but also use the expanded set of occupations.

Table 4.30. *Combined spatial and expanded occupational definition*

Sample:	male workers observed in 1880 and 1900		
Outcome:	Working in Iron Industry	Migration	Log income
Model:	(1)	(2)	(3)
plant death	-0.00976***	0.000161	0.0223**
	(0.00255)	(0.00816)	(0.00900)
plant death × ironworker 1880	-0.164***	0.162***	-0.105***
	(0.0363)	(0.0418)	(0.0230)
R-squared	0.366	0.111	0.463
Observations	85596	85596	83088
No. displaced workers	553	553	553

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18-49 at the time of the 1880 census. The income data is described in Appendix 4.B.

4.F DD and triple differences estimates

The main specification described in equation 4.1 can be shown to be numerically equivalent to a modified triple differences (DDD) estimator that controls for occupation-specific and region-specific differential trends. Therefore, one alternative way to state the key identifying assumption underlying our causal claims is that; conditional on occupation-specific and region-specific trends, outcomes would have evolved similarly over time for ironworkers in parishes where the iron production ceased (treatment) and where it did not stop (control) if it were not for the closure of iron production sites. If this assumption holds, then any observed relative changes in outcomes between workers in control and treatment parishes (after adjusting for occupation and region effects) can be attributed to the treatment.

The central identifying assumption underlying our preferred specification is thus closely related to the standard parallel trends assumption commonly referred to in the context of DD and DDD estimators, but offers the additional flexibility of also allowing for differential region- and occupation-level dynamic effects. However, this added flexibility also means that, in contrast to canonical DD and DDD estimators, our estimated treatment effects do not represent unconditional differences in expectation between the potential and realized outcomes of treated individuals (ATT), as per the Neyman-Rubin potential outcomes framework. Therefore, in this section we reproduce the main results using a standard triple differences strategy, and show that it allows us to reach the same conclusions. We estimate the regression model:

$$\begin{aligned} y_{i,t} = \alpha_i &+ \delta\left(\text{plant\_death}_j \times \mathbf{1}(year = 1890)\right) \\ &+ \gamma\left(\text{plant\_death}_j \times \mathbf{1}(year = 1890) \times \text{ironworker\_1880}_i\right) \\ &+ \theta\left(\mathbf{1}(year = 1890) \times \text{ironworker\_1880}_i\right) + \varepsilon_{i,t} \end{aligned}$$

where $y_{i,t}$ now represents the outcome of individual i in time $t \in \{1880, 1890\}$. The parameter γ specified in this model is the DDD estimator, which can be described as follows:

$$\begin{aligned} \gamma = &\left((\bar{Y}^{T,ironworker1880}_{1890} - \bar{Y}^{T,ironworker1880}_{1880}) - (\bar{Y}^{T,not\ ironworker1880}_{1890} - \bar{Y}^{T,not\ ironworker1880}_{1880}) \right) \\ &- \left((\bar{Y}^{C,ironworker1880}_{1890} - \bar{Y}^{C,ironworker1880}_{1880}) - (\bar{Y}^{C,not\ ironworker1880}_{1890} - \bar{Y}^{C,not\ ironworker1880}_{1880}) \right) \end{aligned}$$

where $\bar{Y}^{T,ironworker1880}_t$ indicates the mean of outcome Y for workers in year t belonging to a treated parish T (as opposed to a control parish C) who were ironworkers in the 1880 (as opposed to having another occupation in 1880). Treated parishes are parishes were iron production closed down between 1880 and 1890. If the difference in changes in outcome $Y_{i,t}$ between ironworkers and non-ironworkers would have been the same across treated and control parishes in the absence of treatment (i.e. if iron production sites did not close), then γ captures the causal effect of the site closures on the outcomes of ironworkers

in the treated parishes.
We also estimate a simple difference-in-differences estimator (corresponding to the estimates presented in column 1 of the main tables), which is given by:

$$Y_{i,t} = \hat{\alpha}_i + \hat{\delta}\left(\text{plant\_death}_j \times \mathbf{1}(year = 1890)\right) + \hat{\varepsilon}_{i,t}$$

where the difference-in-differences estimator $\hat{\delta}$ is given by

$$\hat{\delta} = (\bar{Y}^T_{1890} - \bar{Y}^T_{1880}) - (\bar{Y}^C_{1890} - \bar{Y}^C_{1880})$$

If the changes in outcome Y_t would have been the same across treated and control parishes in the absence of treatment (i.e. if iron production sites did not close), then $\hat{\delta}$ captures the causal effect of the site closures on the outcomes of workers in the treated parishes. Table 4.31 below replicates the main results for ironworker status. Table 4.32 replicates the main results for migration. Table 4.33 replicates the main results for log income.

Table 4.31. *Occupational changes 1880 - 1890*

Outcome: Sample:	Working in Iron Industry panel of male workers observed in 1880 and 1890			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	-0.00664***	-0.00121	-0.0131***	-0.0186***
	(0.00169)	(0.000883)	(0.00191)	(0.00234)
plant death × ironworker 1880		-0.249***	-0.271***	-0.272***
		(0.0602)	(0.0610)	(0.0611)
R-squared	0.000586	0.159	0.167	0.167
Observations	1058774	1058774	221692	172990

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18-60 at the time of the 1880 census. The lists of occupations considered as iron workers are displayed in table 4.1.

Table 4.32. *Migration between 1880 - 1890*

Outcome:	Migration			
Sample:	panel of male workers observed in 1880 and 1890			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	-0.0885***	-0.0921***	-0.0356**	-0.0307**
	(0.0139)	(0.0138)	(0.0161)	(0.0131)
plant death × ironworker 1880		0.263***	0.276***	0.276***
		(0.0497)	(0.0510)	(0.0511)
R-squared	0.245	0.245	0.187	0.182
Observations	1058774	1058774	221692	172990

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18-60 at the time of the 1880 census. Migration is defined as a change of parish between the 1880 and 1890 census.

Table 4.33. *Log income in 1890*

Outcome:	Log income			
Sample:	panel of male workers observed in 1880 and 1890			
Parishes:	all	all	Ever iron produced	Producing Iron 1880
Model:	(1)	(2)	(3)	(4)
plant death	-0.000612	0.00101	-0.0115	-0.0187*
	(0.00888)	(0.00912)	(0.00976)	(0.0101)
plant death × ironworker 1880		-0.0961***	-0.0692**	-0.0649*
		(0.0351)	(0.0351)	(0.0352)
R-squared	0.00154	0.00170	0.00536	0.00761
Observations	1041390	1041390	217941	170442

Note: Standard errors in parentheses are adjusted for clustering at the parish level: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Sample restricted to those between ages 18-60 at the time of the 1880 census. The income data is described in Appendix 4.B.

Economic Studies

1987:1 Haraldson, Marty. To Care and To Cure. A linear programming approach to national health planning in developing countries. 98 pp.

1989:1 Chryssanthou, Nikos. The Portfolio Demand for the ECU. A Transaction Cost Approach. 42 pp.

1989:2 Hansson, Bengt. Construction of Swedish Capital Stocks, 1963-87. An Application of the Hulten-Wykoff Studies. 37 pp.

1989:3 Choe, Byung-Tae. Some Notes on Utility Functions Demand and Aggregation. 39 pp.

1989:4 Skedinger, Per. Studies of Wage and Employment Determination in the Swedish Wood Industry. 89 pp.

1990:1 Gustafson, Claes-Håkan. Inventory Investment in Manufacturing Firms. Theory and Evidence. 98 pp.

1990:2 Bantekas, Apostolos. The Demand for Male and Female Workers in Swedish Manufacturing. 56 pp.

1991:1 Lundholm, Michael. Compulsory Social Insurance. A Critical Review. 109 pp.

1992:1 Sundberg, Gun. The Demand for Health and Medical Care in Sweden. 58 pp.

1992:2 Gustavsson, Thomas. No Arbitrage Pricing and the Term Structure of Interest Rates. 47 pp.

1992:3 Elvander, Nils. Labour Market Relations in Sweden and Great Britain. A Comparative Study of Local Wage Formation in the Private Sector during the 1980s. 43 pp.

12 Dillén, Mats. Studies in Optimal Taxation, Stabilization, and Imperfect Competition. 1993. 143 pp.

13 Banks, Ferdinand E.. A Modern Introduction to International Money, Banking and Finance. 1993. 303 pp.

14 Mellander, Erik. Measuring Productivity and Inefficiency Without Quantitative Output Data. 1993. 140 pp.

15 Ackum Agell. Susanne. Essays on Work and Pay. 1993. 116 pp.

16 Eriksson, Claes. Essays on Growth and Distribution. 1994. 129 pp.

17 Banks, Ferdinand E.. A Modern Introduction to International Money, Banking and Finance. 2nd version, 1994. 313 pp.

18 Apel, Mikael. Essays on Taxation and Economic Behavior. 1994. 144 pp.

19 Dillén, Hans. Asset Prices in Open Monetary Economies. A Contingent Claims Approach. 1994. 100 pp.

20 Jansson, Per. Essays on Empirical Macroeconomics. 1994. 146 pp.

21 Banks, Ferdinand E.. A Modern Introduction to International Money, Banking, and Finance. 3rd version, 1995. 313 pp.

22 Dufwenberg, Martin. On Rationality and Belief Formation in Games. 1995. 93 pp.

23 Lindén, Johan. Job Search and Wage Bargaining. 1995. 127 pp.

24 Shahnazarian, Hovick. Three Essays on Corporate Taxation. 1996. 112 pp.

25 Svensson, Roger. Foreign Activities of Swedish Multinational Corporations. 1996. 166 pp.

26 Sundberg, Gun. Essays on Health Economics. 1996. 174 pp.

27 Sacklén, Hans. Essays on Empirical Models of Labor Supply. 1996. 168 pp.

28 Fredriksson, Peter. Education, Migration and Active Labor Market Policy. 1997. 106 pp.

29 Ekman, Erik. Household and Corporate Behaviour under Uncertainty. 1997. 160 pp.

30 Stoltz, Bo. Essays on Portfolio Behavior and Asset Pricing. 1997. 122 pp.

31 Dahlberg, Matz. Essays on Estimation Methods and Local Public Economics. 1997. 179 pp.

32 Kolm, Ann-Sofie. Taxation, Wage Formation, Unemployment and Welfare. 1997. 162 pp.

33 Boije, Robert. Capitalisation, Efficiency and the Demand for Local Public Services. 1997. 148 pp.

34 Hort, Katinka. On Price Formation and Quantity Adjustment in Swedish Housing Markets. 1997. 185 pp.

35 Lindström, Thomas. Studies in Empirical Macroeconomics. 1998. 113 pp.

36 Hemström, Maria. Salary Determination in Professional Labour Markets. 1998. 127 pp.

37 Forsling, Gunnar. Utilization of Tax Allowances and Corporate Borrowing. 1998. 96 pp.

38 Nydahl, Stefan. Essays on Stock Prices and Exchange Rates. 1998. 133 pp.

39 Bergström, Pål. Essays on Labour Economics and Econometrics. 1998. 163 pp.

40 Heiborn, Marie. Essays on Demographic Factors and Housing Markets. 1998. 138 pp.

41 Åsberg, Per. Four Essays in Housing Economics. 1998. 166 pp.

42 Hokkanen, Jyry. Interpreting Budget Deficits and Productivity Fluctuations. 1998. 146 pp.

43 Lunander, Anders. Bids and Values. 1999. 127 pp.

44 Eklöf, Matias. Studies in Empirical Microeconomics. 1999. 213 pp.

45 Johansson, Eva. Essays on Local Public Finance and Intergovernmental Grants. 1999. 156 pp.

46 Lundin, Douglas. Studies in Empirical Public Economics. 1999. 97 pp.

47 Hansen, Sten. Essays on Finance, Taxation and Corporate Investment. 1999. 140 pp.

48 Widmalm, Frida. Studies in Growth and Household Allocation. 2000. 100 pp.

49 Arslanogullari, Sebastian. Household Adjustment to Unemployment. 2000. 153 pp.

50 Lindberg, Sara. Studies in Credit Constraints and Economic Behavior. 2000. 135 pp.

51 Nordblom, Katarina. Essays on Fiscal Policy, Growth, and the Importance of Family Altruism. 2000. 105 pp.

52 Andersson, Björn. Growth, Saving, and Demography. 2000. 99 pp.

53 Åslund, Olof. Health, Immigration, and Settlement Policies. 2000. 224 pp.

54 Bali Swain, Ranjula. Demand, Segmentation and Rationing in the Rural Credit Markets of Puri. 2001. 160 pp.

55 Löfqvist, Richard. Tax Avoidance, Dividend Signaling and Shareholder Taxation in an Open Economy. 2001. 145 pp.

56 Vejsiu, Altin. Essays on Labor Market Dynamics. 2001. 209 pp.

57 Zetterström, Erik. Residential Mobility and Tenure Choice in the Swedish Housing Market. 2001. 125 pp.

58 Grahn, Sofia. Topics in Cooperative Game Theory. 2001. 106 pp.

59 Laséen, Stefan. Macroeconomic Fluctuations and Microeconomic Adjustments. Wages, Capital, and Labor Market Policy. 2001. 142 pp.

60 Arnek, Magnus. Empirical Essays on Procurement and Regulation. 2002. 155 pp.

61 Jordahl, Henrik. Essays on Voting Behavior, Labor Market Policy, and Taxation. 2002. 172 pp.

62 Lindhe, Tobias. Corporate Tax Integration and the Cost of Capital. 2002. 102 pp.

63 Hallberg, Daniel. Essays on Household Behavior and Time-Use. 2002. 170 pp.

64 Larsson, Laura. Evaluating Social Programs: Active Labor Market Policies and Social Insurance. 2002. 126 pp.

65 Bergvall, Anders. Essays on Exchange Rates and Macroeconomic Stability. 2002. 122 pp.

66 Nordström Skans, Oskar. Labour Market Effects of Working Time Reductions and Demographic Changes. 2002. 118 pp.

67 Jansson, Joakim. Empirical Studies in Corporate Finance, Taxation and Investment. 2002. 132 pp.

68 Carlsson, Mikael. Macroeconomic Fluctuations and Firm Dynamics: Technology, Production and Capital Formation. 2002. 149 pp.

69 Eriksson, Stefan. The Persistence of Unemployment: Does Competition between Employed and Unemployed Job Applicants Matter? 2002. 154 pp.

70 Huitfeldt, Henrik. Labour Market Behaviour in a Transition Economy: The Czech Experience. 2003. 110 pp.

71 Johnsson, Richard. Transport Tax Policy Simulations and Satellite Accounting within a CGE Framework. 2003. 84 pp.

72 Öberg, Ann. Essays on Capital Income Taxation in the Corporate and Housing Sectors. 2003. 183 pp.

73 Andersson, Fredrik. Causes and Labor Market Consequences of Producer Heterogeneity. 2003. 197 pp.

74 Engström, Per. Optimal Taxation in Search Equilibrium. 2003. 127 pp.

75 Lundin, Magnus. The Dynamic Behavior of Prices and Investment: Financial Constraints and Customer Markets. 2003. 125 pp.

76 Ekström, Erika. Essays on Inequality and Education. 2003. 166 pp.

77 Barot, Bharat. Empirical Studies in Consumption, House Prices and the Accuracy of European Growth and Inflation Forecasts. 2003. 137 pp.

78 Österholm, Pär. Time Series and Macroeconomics: Studies in Demography and Monetary Policy. 2004. 116 pp.

79 Bruér, Mattias. Empirical Studies in Demography and Macroeconomics. 2004. 113 pp.

80 Gustavsson, Magnus. Empirical Essays on Earnings Inequality. 2004. 154 pp.

81 Toll, Stefan. Studies in Mortgage Pricing and Finance Theory. 2004. 100 pp.

82 Hesselius, Patrik. Sickness Absence and Labour Market Outcomes. 2004. 109 pp.

83 Häkkinen, Iida. Essays on School Resources, Academic Achievement and Student Employment. 2004. 123 pp.

84 Armelius, Hanna. Distributional Side Effects of Tax Policies: An Analysis of Tax Avoidance and Congestion Tolls. 2004. 96 pp.

85 Ahlin, Åsa. Compulsory Schooling in a Decentralized Setting: Studies of the Swedish Case. 2004. 148 pp.

86 Heldt, Tobias. Sustainable Nature Tourism and the Nature of Tourists' Cooperative Behavior: Recreation Conflicts, Conditional Cooperation and the Public Good Problem. 2005. 148 pp.

87 Holmberg, Pär. Modelling Bidding Behaviour in Electricity Auctions: Supply Function Equilibria with Uncertain Demand and Capacity Constraints. 2005. 43 pp.

88 Welz, Peter. Quantitative new Keynesian macroeconomics and monetary policy 2005. 128 pp.

89 Ågren, Hanna. Essays on Political Representation, Electoral Accountability and Strategic Interactions. 2005. 147 pp.

90 Budh, Erika. Essays on environmental economics. 2005. 115 pp.

91 Chen, Jie. Empirical Essays on Housing Allowances, Housing Wealth and Aggregate Consumption. 2005. 192 pp.

92 Angelov, Nikolay. Essays on Unit-Root Testing and on Discrete-Response Modelling of Firm Mergers. 2006. 127 pp.

93 Savvidou, Eleni. Technology, Human Capital and Labor Demand. 2006. 151 pp.

94 Lindvall, Lars. Public Expenditures and Youth Crime. 2006. 112 pp.

95 Söderström, Martin. Evaluating Institutional Changes in Education and Wage Policy. 2006. 131 pp.

96 Lagerström, Jonas. Discrimination, Sickness Absence, and Labor Market Policy. 2006. 105 pp.

97 Johansson, Kerstin. Empirical essays on labor-force participation, matching, and trade. 2006. 168 pp.

98 Ågren, Martin. Essays on Prospect Theory and the Statistical Modeling of Financial Returns. 2006. 105 pp.

99 Nahum, Ruth-Aïda. Studies on the Determinants and Effects of Health, Inequality and Labour Supply: Micro and Macro Evidence. 2006. 153 pp.

100 Žamac, Jovan. Education, Pensions, and Demography. 2007. 105 pp.

101 Post, Erik. Macroeconomic Uncertainty and Exchange Rate Policy. 2007. 129 pp.

102 Nordberg, Mikael. Allies Yet Rivals: Input Joint Ventures and Their Competitive Effects. 2007. 122 pp.

103 Johansson, Fredrik. Essays on Measurement Error and Nonresponse. 2007. 130 pp.

104 Haraldsson, Mattias. Essays on Transport Economics. 2007. 104 pp.

105 Edmark, Karin. Strategic Interactions among Swedish Local Governments. 2007. 141 pp.

106 Oreland, Carl. Family Control in Swedish Public Companies. Implications for Firm Performance, Dividends and CEO Cash Compensation. 2007. 121 pp.

107 Andersson, Christian. Teachers and Student Outcomes: Evidence using Swedish Data. 2007. 154 pp.

108 Kjellberg, David. Expectations, Uncertainty, and Monetary Policy. 2007. 132 pp.

109 Nykvist, Jenny. Self-employment Entry and Survival - Evidence from Sweden. 2008. 94 pp.

110 Selin, Håkan. Four Empirical Essays on Responses to Income Taxation. 2008. 133 pp.

111 Lindahl, Erica. Empirical studies of public policies within the primary school and the sickness insurance. 2008. 143 pp.

112 Liang, Che-Yuan. Essays in Political Economics and Public Finance. 2008. 125 pp.

113 Elinder, Mikael. Essays on Economic Voting, Cognitive Dissonance, and Trust. 2008. 120 pp.

114 Grönqvist, Hans. Essays in Labor and Demographic Economics. 2009. 120 pp.

115 Bengtsson, Niklas. Essays in Development and Labor Economics. 2009. 93 pp.

116 Vikström, Johan. Incentives and Norms in Social Insurance: Applications, Identification and Inference. 2009. 205 pp.

117 Liu, Qian. Essays on Labor Economics: Education, Employment, and Gender. 2009. 133 pp.

118 Glans, Erik. Pension reforms and retirement behaviour. 2009. 126 pp.

119 Douhan, Robin. Development, Education and Entrepreneurship. 2009.

120 Nilsson, Peter. Essays on Social Interactions and the Long-term Effects of Early-life Conditions. 2009. 180 pp.

121 Johansson, Elly-Ann. Essays on schooling, gender, and parental leave. 2010. 131 pp.

122 Hall, Caroline. Empirical Essays on Education and Social Insurance Policies. 2010. 147 pp.

123 Enström-Öst, Cecilia. Housing policy and family formation. 2010. 98 pp.

124 Winstrand, Jakob. Essays on Valuation of Environmental Attributes. 2010. 96 pp.

125 Söderberg, Johan. Price Setting, Inflation Dynamics, and Monetary Policy. 2010. 102 pp.

126 Rickne, Johanna. Essays in Development, Institutions and Gender. 2011. 138 pp.

127 Hensvik, Lena. The effects of markets, managers and peers on worker outcomes. 2011. 179 pp.

128 Lundqvist, Heléne. Empirical Essays in Political and Public. 2011. 157 pp.

129 Bastani, Spencer. Essays on the Economics of Income Taxation. 2012. 257 pp.

130 Corbo, Vesna. Monetary Policy, Trade Dynamics, and Labor Markets in Open Economies. 2012. 262 pp.

131 Nordin, Mattias. Information, Voting Behavior and Electoral Accountability. 2012. 187 pp.

132 Vikman, Ulrika. Benefits or Work? Social Programs and Labor Supply. 2013. 161 pp.

133 Ek, Susanne. Essays on unemployment insurance design. 2013. 136 pp.

134 Österholm, Göran. Essays on Managerial Compensation. 2013. 143 pp.

135 Adermon, Adrian. Essays on the transmission of human capital and the impact of technological change. 2013. 138 pp.

136 Kolsrud, Jonas. Insuring Against Unemployment 2013. 140 pp.

137 Hanspers, Kajsa. Essays on Welfare Dependency and the Privatization of Welfare Services. 2013. 208 pp.

138 Persson, Anna. Activation Programs, Benefit Take-Up, and Labor Market Attachment. 2013. 164 pp.

139 Engdahl, Mattias. International Mobility and the Labor Market. 2013. 216 pp.

140 Krzysztof Karbownik. Essays in education and family economics. 2013. 182 pp.

141 Oscar Erixson. Economic Decisions and Social Norms in Life and Death Situations. 2013. 183 pp.

142 Pia Fromlet. Essays on Inflation Targeting and Export Price Dynamics. 2013. 145 pp.

143 Daniel Avdic. Microeconometric Analyses of Individual Behavior in Public Welfare Systems. Applications in Health and Education Economics. 2014. 176 pp.

144 Arizo Karimi. Impacts of Policies, Peers and Parenthood on Labor Market Outcomes. 2014. 221 pp.

145 Karolina Stadin. Employment Dynamics. 2014. 134 pp.

146 Haishan Yu. Essays on Environmental and Energy Economics. 132 pp.

147 Martin Nilsson. Essays on Health Shocks and Social Insurance. 139 pp.

148 Tove Eliasson. Empirical Essays on Wage Setting and Immigrant Labor Market Opportunities. 2014. 144 pp.

149 Erik Spector. Financial Frictions and Firm Dynamics. 2014. 129 pp.

150 Michihito Ando. Essays on the Evaluation of Public Policies. 2015. 193 pp.

151 Selva Bahar Baziki. Firms, International Competition, and the Labor Market. 2015. 183 pp.

152 Fredrik Sävje. What would have happened? Four essays investigating causality. 2015. 229 pp.

153 Ina Blind. Essays on Urban Economics. 2015. 197 pp.

154 Jonas Poulsen. Essays on Development and Politics in Sub-Saharan Africa. 2015. 240 pp.

155 Lovisa Persson. Essays on Politics, Fiscal Institutions, and Public Finance. 2015. 137 pp.

156 Gabriella Chirico Willstedt. Demand, Competition and Redistribution in Swedish Dental Care. 2015. 119 pp.

157 Yuwei Zhao de Gosson de Varennes. Benefit Design, Retirement Decisions and Welfare Within and Across Generations in Defined Contribution Pension Schemes. 2016. 148 pp.

158 Johannes Hagen. Essays on Pensions, Retirement and Tax Evasion. 2016. 195 pp.

159 Rachatar Nilavongse. Housing, Banking and the Macro Economy. 2016. 156 pp.

160 Linna Martén. Essays on Politics, Law, and Economics. 2016. 150 pp.

161 Olof Rosenqvist. Essays on Determinants of Individual Performance and Labor Market Outcomes. 2016. 151 pp.

162 Linuz Aggeborn. Essays on Politics and Health Economics. 2016. 203 pp.

163 Glenn Mickelsson. DSGE Model Estimation and Labor Market Dynamics. 2016. 166 pp.

164 Sebastian Axbard. Crime, Corruption and Development. 2016. 150 pp.

165 Mattias Öhman. Essays on Cognitive Development and Medical Care. 2016. 181 pp.

166 Jon Frank. Essays on Corporate Finance and Asset Pricing. 2017. 160 pp.

167 Ylva Moberg. Gender, Incentives, and the Division of Labor. 2017. 220 pp.

168 Sebastian Escobar. Essays on inheritance, small businesses and energy consumption. 2017. 194 pp.

169 Evelina Björkegren. Family, Neighborhoods, and Health. 2017. 226 pp.

170 Jenny Jans. Causes and Consequences of Early-life Conditions. Alcohol, Pollution and Parental Leave Policies. 2017. 209 pp.

171 Josefine Andersson. Insurances against job loss and disability. Private and public interventions and their effects on job search and labor supply. 2017. 175 pp.

172 Jacob Lundberg. Essays on Income Taxation and Wealth Inequality. 2017. 173 pp.

173 Anna Norén. Caring, Sharing, and Childbearing. Essays on Labor Supply, Infant Health, and Family Policies. 2017. 206 pp.

174 Irina Andone. Exchange Rates, Exports, Inflation, and International Monetary Cooperation. 2018. 174 pp.

175 Henrik Andersson. Immigration and the Neighborhood. Essays on the Causes and Consequences of International Migration. 2018. 181 pp.

176 Aino-Maija Aalto. Incentives and Inequalities in Family and Working Life. 2018. 131 pp.

177 Gunnar Brandén. Understanding Intergenerational Mobility. Inequality, Student Aid and Nature-Nurture Interactions. 2018. 125 pp.

178 Mohammad H. Sepahvand. Essays on Risk Attitudes in Sub-Saharan Africa. 2019. 215 pp.

179 Mathias von Buxhoeveden. Partial and General Equilibrium Effects of Unemployment Insurance. Identification, Estimation and Inference. 2019. 89 pp.

180 Stefano Lombardi. Essays on Event History Analysis and the Effects of Social Programs on Individuals and Firms. 2019. 150 pp.

181 Arnaldur Stefansson. Essays in Public Finance and Behavioral Economics. 2019. 191 pp.

182 Cristina Bratu. Immigration: Policies, Mobility and Integration. 2019. 173 pp.

183 Tamás Vasi. Banks, Shocks and Monetary Policy. 2020. 148 pp.

184 Jonas Cederlöf. Job Loss: Consequences and Labor Market Policy. 2020. 213 pp.

185 Dmytro Stoyko. Expectations, Financial Markets and Monetary Policy. 2020. 153 pp.

186 Paula Roth. Essays on Inequality, Insolvency and Innovation. 2020. 191 pp.

187 Fredrik Hansson. Consequences of Poor Housing, Essays on Urban and Health Economics. 2020. 143 pp.

188 Maria Olsson. Essays on Macroeconomics: Wage Rigidity and Aggregate Fluctuations. 2020. 130 pp.

189 Dagmar Müller. Social Networks and the School-to-Work Transition. 2020. 146 pp.

190 Maria Sandström. Essays on Savings and Intangible Capital. 2020. 129 pp.

191. Anna Thoresson. Wages and Their Impact on Individuals, Households and Firms. 2020. 220 pp.

192. Jonas Klarin. Empirical Essays in Public and Political Economics. 2020. 129 pp.

193. André Reslow. Electoral Incentives and Information Content in Macroeconomic Forecasts. 2021. 184 pp.

194. Davide Cipullo. Political Careers, Government Stability, and Electoral Cycles. 2021. 308 pp.

195. Olle Hammar. The Mystery of Inequality: Essays on Culture, Development, and Distributions. 2021. 210 pp.

196. J. Lucas Tilley. Inputs and Incentives in Education. 2021. 184 pp.

197. Sebastian Jävervall. Corruption, Distortions and Development. 2021. 215 pp.

198. Vivika Halapuu. Upper Secondary Education: Access, Choices and Graduation. 2021. 141 pp.

199. Charlotte Paulie. Essays on the Distribution of Production, Prices and Wealth. 2021. 141 pp.

200. Kerstin Westergren. Essays on Inflation Expectations, Monetary Policy and Tax Reform. 2021. 124 pp.

201. Melinda Süveg. Finance, Shocks, Competition and Price Setting. 2021. 137 pp.

202. Adrian Poignant. Gold, Coal and Iron. Essays on Industrialization and Economic Development. 2022. 214 pp.

203. Daniel Bougt. A Sequence of Essays on Sequences of Auctions. 2022. 188 pp.

204. Lillit Ottosson. From Welfare to Work. Financial Incentives, Active Labor Market Policies, and Integration Programs. 2022. 219 pp.

205. Yaroslav Yakymovych. Workers and Occupations in a Changing Labour Market. The Heterogeneous Effects of Mass Layoffs and Social Safety Nets. 2022. 212 pp.

206. Raoul van Maarseveen. Urbanization and Education. The Effect of Childhood Urban Residency on Educational Attainment. 2022.