Moved to Poverty? A Legacy of the Apartheid Experiment in South Africa*

Bladimir Carrillo Carlos Charris Wilman Iglesias

July 21, 2022

Abstract

During the South African apartheid, Black people were forced to move to homelands during the 1960s and 1970s, resulting in one of history's largest segregation policy experiments. We examine how and why relocation to the homelands affected human capital attainment. Exploiting the staggered timing of homeland establishment in a cross-cohort identification strategy, we find that moving to the homelands during childhood significantly reduces educational attainment, labor earnings and employment rates in adulthood. The data suggest an important role for place effects. Moving to the homelands in childhood implies greater exposure to poorer neighborhoods and it disproportionally reduces human capital attainment.

JEL codes: E24, J15, O15, N37

Keywords: Apartheid; Homelands; Segregation; Migration; Human capital; Africa

*Contact information: Carrillo: Department of Economics, Universidade Federal de Pernambuco, AV. Prof. Moraes Rego, 1235 - Cidade Universitaria, Recife - PE, Brazil, 50670-420 (e-mail: bladimir.carrillo@ufpe.br). Charris: Department of Economics, Catholic University of Brasilia, Brasilia, Brasil, 71966-700 (e-mail: ccharris1988@gmail.com). Iglesias: Department of Agricultural Economics and Agribusiness, University of Arkansas, Fayetteville, AR 72701 (e-mail: wi001@uark.edu). C. Kirabo Jackson was coeditor for this article. We thank Daniel Araujo, Sascha Becker, Leah Boustan, Raj Chetty, Eric Chyn, Francisco Costa, Lilyan Fulginiti, Pauline Grosjean, Yifan Gong, Dyeggo Guedes, Felipe Iachan, Philip Oreopoulos, Elias Papaioannou, Pedro Queiroz, Imran Rasul, Breno Sampaio, Jesse Shapiro, Daniel Tannenbaum, Brenden Timpe, Andre Trindade, and participants at various conferences and seminars for helpful comments and suggestions. Carlos Charris acknowledges financial support by the National Council for Scientific and Technological Development (CNPq). We are solely responsible for this paper's contents. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

1 Introduction

The creation of South Africa's system of racial segregation known as *apartheid* was one of the most remarkable events in modern world history. A critical aspect of apartheid was the delimitation of "homeland" areas within the country where several million African Blacks were forced to move and live separately from the White minority (Simkins, 1983). Black families were suddenly uprooted from their traditional way of life and sent to distant reserves characterized by poor sanitary and physical infrastructure conditions. Many historians and other social scientists have argued that this apartheid experiment is a prime reason why South Africa is today one of the most unequal societies in the world (Clark, 2014), but research quantifying its economic and social legacies remains surprisingly limited. The consequences of this segregation system are of interest not only for historical reasons but also because they relate to contemporary debates on the role of places in shaping children's economic prospects and inequality.

Although many aspects of apartheid are unique to its particular history, governments that practice analogous discriminatory policies have been common throughout history and remain a pervasive feature of many contemporary societies. Salient examples include Jews in Medieval and Modern Europe, Chinese in the Philippines and Indonesia, Muslims in India, and non-Muslims in Pakistan and Saudi Arabia.[1] Some scholars have even compared the South African apartheid to the historical treatment of African Americans in the United States, particularly during the Jim Crow era (Massey and Denton, 1993). This study seeks to contribute to the literature on the legacies of apartheid and other instances of governmental ethnic discrimination.

We bring together a rich collection of microdata to provide systematic evidence on whether, how, and to what extent relocation to the homelands affected children's human capital formation. We take advantage of a policy change introduced by the apartheid government in the late 1950s which mandated the formalization of existing Native Reserve areas into homelands for the Black population. The primary goal of the government was to remove Blacks residing in "White" areas and make certain areas racially homogeneous. With this policy change, Blacks would be resident only of their assigned homeland, and migration out of or between homelands was severely controlled through a complex passport system. Since the homelands were formalized at different moments in time, there were differences in the time when individuals moved to the homelands.

Our identification strategy exploits variation in the date of homeland establishment in an

[1] In Medieval and Modern Europe, Jewish people were forcibly confined to certain areas and faced restrictions on spatial mobility (see Becker and Pascali (2019)). In the Philippines and Indonesia, the Spanish and the Dutch colonizers introduced anti-Chinese laws, resulting in many Chinese families being expelled and their properties looted and burned. In India, discrimination against Muslims has been legitimized by policies and legislation, mainly anti-cow slaughter laws, that exacerbated systematic violent attacks and state-sanctioned expulsions of India's Muslim communities. In Pakistan and Saudi Arabia, non-Muslim subjects have faced severe official discrimination in employment and education and were forbidden from openly practicing their religion.

intent-to-treat framework. Specifically, we predict for each migrant of White-place origin the age at which he or she would have arrived in the homelands based on *when* his or her homeland was established and then assign a measure of childhood exposure. Since population mobility was under strict control and thus most "White place migrants" resided in the homelands where they moved to, our assignment of childhood exposure is likely to be highly accurate in the apartheid censuses.[2] In line with historical narratives, we show that homeland establishment is a powerful determinant of the timing of arrival in the homelands. Of course, one might be concerned that the timing of homeland establishment is correlated with pre-existing conditions and other coincident changes in the homelands. We thus extend our comparison group by including homeland natives in our analysis, a group that certainly was not subject to relocation.[3] The inclusion of this extended comparison group enables us to gain statistical power and control for a detailed set of homeland $\times$ birth-year fixed effects, which absorb any pre-existing trend or contemporary economic circumstances that affect all individuals in a homeland similarly.

By implementing our cross-cohort identification strategy with census data, we find that moving to the homelands during childhood significantly reduces educational attainment. The estimates are very precise and economically meaningful. Our preferred estimates imply that a predicted move at age 9 would reduce educational attainment by about 0.55 years, or equivalently 11 percent of the sample mean. Extending our analysis to labor market outcomes, we find that the relocation policy led to a 20-percent decline in the likelihood of working and a 5.6-percent reduction in adult income. These effects tend to be more pronounced for individuals with earlier predicted ages at arrival, an age profile consistent with the importance of circumstances and events during critical periods of child development.

A causal interpretation of these patterns requires the assumption that the outcomes of migrants and homeland natives would have followed similar cohort trends in the absence of the relocation policy. We provide a variety of evidence supporting this identifying assumption. Most importantly, we use an event-study-like specification and provide visually clear evidence that the policy had no effect on the outcomes of older cohorts who had already completed their schooling decisions before the policy change, exactly where one would expect to find zero impacts if the identification condition holds. Moreover, we show that the effects are completely absent for individuals who were not directly induced to move to the homelands. In general, the patterns we document are robust to a variety of alternative specifications, different datasets, and inference methods.

[2] We use the term "White place migrants" to refer to those individuals who moved to their designated homelands in response to the policy, either directly or indirectly, rather than those who moved for other reasons.

[3] The term "locals" refers to individuals born and reside in a homeland at census time. We use the terms locals, homeland natives, and permanent homeland residents interchangeably throughout the paper. Similar migrant and native comparison designs have been used in other settings. For a recent example, see Hornbeck (2020).

In the final section of the paper, we take a closer look at the data to better understand why the relocation policy had adverse effects. The possible causes of these effects can be divided into two broad categories: changes in individual-specific circumstances and changes in childhood environment. The former may include, for example, loss of immobile assets, discrimination, and changes in preferences. Changes in childhood environment potentially include employment opportunities for parents, school quality, peers, access to healthcare, and other important places-specific inputs. We shed light on these forces by directly investigating the causal effects of places. Changes in childhood environment are a natural explanation for our results given that the homelands were disadvantaged places relative to White areas and may have affected migrants by offering an environment unfavorable for human capital acquisition.

To investigate the importance of places, we use an empirical strategy similar in spirit to that of Chetty and Hendren (2018). Specifically, we compare migrants who moved earlier and later to higher and lower schooling areas relative to their origin places. If places have causal effects, then children's schooling should tend to converge toward the levels of schooling where they move to. When the schooling gap between origin and destination is high, the decline in migrant schooling would be larger in magnitude than if the schooling gap is small. This is exactly what we find. Moving during childhood to lower education areas has a larger negative effect on educational attainment. These effects are gender-specific. Moving to the homelands in childhood has a more pronounced effect if one's gender schooling in the destination is particularly low, but origin-destination differences in the other-gender schooling have no significant effects. This pattern is consistent with a causal interpretation of our results. A back-of-envelope calculation suggests that changes in childhood environment account for approximately 70 percent of the overall decline in migrant schooling.

Summarizing, relocation to the homelands had adverse consequences on human capital, and this occurred primarily through a change in an individual's early childhood environment. We interpret the entirety of these findings as evidence on one of the channels by which the South African government's project of apartheid had long-lasting consequences.

These findings represent one of the first rigorous, empirical analyses of a significant historical event that has not received much attention in the economics literature.[4,5] We view this itself as an important contribution. After the mass murder of million Jews and mass population movements in the aftermath of World War II, the South African apartheid is among the most infamous chapters of the twentieth century. The United Nations declared apartheid a "crime

[4]There is a large literature in other social sciences on the cultural, social, and psychological impacts of the apartheid regime, with a qualitative focus. Within economics, Abel (2019) studies the homelands by making a cross-section comparison of areas closer to and farther away from the relocation camps several decades after the relocation policy. His focus is on inter-ethnic trust, and he does not provide any evidence regarding human capital formation or labor market outcomes.

[5]While evidence on the legacies of apartheid is limited, a rapidly growing literature examines the impacts of post-apartheid policies. See, for example, Duflo (2000), Tanaka (2014), and Ito and Tanaka (2018).

against humanity" once the discriminatory policies implemented under this regime became widely known. Our findings are likely to be relevant for other historical and contemporary episodes of governmental ethnic discrimination.

This paper also speaks to a large body of work on forced migration, whose consequences are not necessarily the same as those of voluntary migration.[6,7]This literature studies the effects of forced migration on long-run aggregate outcomes of receiving areas (e.g., Dippel, 2014), as well as on individual outcomes disentangling locals (e.g., Borjas and Monras, 2017; Morales, 2018) and migrants (e.g., Chyn, 2018; Becker et al., 2020). Our findings contribute to the literature by providing new causal evidence that heterogeneous conditions between origin and destination places play a key role in shaping the outcomes of displaced people. As noted by Becker and Ferrara (2019), with a few exceptions, the literature on forced migration typically treats displaced individuals as moving between large regions or countries without explicitly considering that the path of migrants' outcomes could depend on where they come from and on where they settle at. This may help understand why estimates from previous studies differ in magnitude and in some cases, the direction of the effects.

Our findings are also related to the abundant literature on the effects of neighborhoods/places on children. Recent experimental and quasi-experimental studies document that children growing up in better places exhibit higher socioeconomic status in adulthood (Chetty and Hendren, 2018; Deutscher, 2020; Nakamura et al., 2016; Chyn, 2018). A unique feature of our study is the focus on families who were forcefully relocated to lower-opportunity areas, and the fact that residence in such places was strictly enforced. This provides a rare opportunity to evaluate the consequences of relocating people from "good" to "bad" places and shed light on mechanisms. This is important because it is not clear that this type of relocation can have corresponding causal effects[8] and because several developing countries are implementing large-scale housing programs that in many cases imply relocation to lower-opportunity areas.[9]

The rest of the paper is organized as follows. Section 2 provides detailed background information. Section 3 introduces the data. Section 4 describes the empirical approach and main results. Section 5 explores mechanisms. Section 6 concludes.

[6]See Becker and Ferrara (2019) for an excellent review of this vast literature.

[7]Prominent studies on voluntary migration include Abramitzky et al. (2012), Bazzi et al. (2016) and, more recently, Abramitzky et al. (2020) and Sequeira et al. (2020).

[8]For example, if there are critical factors (such as adequate information) in good places that permanently define the technology of skill formation, then relocating families to worse places could have only limited consequences.

[9]Several developing countries have been implementing housing programs to reduce urban slums by relocating people to city peripheries (UN-Habitat, 2004). These countries include for example Brazil (Dasgupta and Lall, 2009), China (Day and Cervero, 2010), India (Barnhardt et al., 2017), Indonesia (Some et al., 2009), and Thailand (Viratkapan and Perera, 2006). These programs have some important benefits including improved housing, but also imply losing the opportunities urbanicity offers, such as access to schools and other public services as well as proximity to employment opportunities (Barnhardt et al., 2017).

2 Institutional Background

2.1 The Homeland System and Relocation Policy

The apartheid regime was a system of racial segregation implemented in South Africa between 1948 and 1991. This regime began with the National Party's ascension in the 1948 general elections, which implemented several discriminatory reforms against African Blacks.[10] The key goal of this regime was to sustain the political and socioeconomic domination of the nation's White minority over the Black majority by creating a segregated-based government system. It would imply the complete separation between a White state and several Black ethnic-based states within South Africa, with racial classification determining the place of residence. Blacks would not be residents of South Africa but only of their assigned homeland.

Staggered timing of homeland establishment. The Bantu Authorities Act (1951) and Bantu Resettlement Act (1954) allowed the creation of ten homeland states, also known as Bantustans, based on existing Native Reserve areas (see Figure 1). The establishment of the homelands was staggered, as listed in Appendix Table A.1. The first homelands to be formally established were Transkei, Bophuthatswana, Ciskei, and Venda between 1959 and 1962. Over the next fifteen years, the rest of the reserves were officially created.

The formalization date of the homelands was not random across reserves. The status of the reserves was formalized after the apartheid government had completely established bureaucratic structures in these rural areas. This included the creation of special authorities to enforce mobility restrictions and ensure continued control of the homelands over the long run (O'Malley, 2007). Furthermore, there were important political issues involved in creating the homelands, such as the length of the time required to forge relationships between the central government and local chiefs which varied across reserves. It was necessary to ensure that local chiefs would be compliant with executing massive removals and monitoring homeland administration (Evans, 1997). After introducing our identification strategy, we discuss the implications of non-random establishment dates.

Relocation and homeland assignment. Once formally established, African Blacks were assigned a homeland according to ethnic status and those residing in White areas were forced to relocate to the reserves. The government used the establishment of the homelands to politically

[10] The apartheid government introduced a number of acts that explicitly discriminated against Blacks in several dimensions, including interethnic marriage, the labor market, education, and internal mobility within South Africa. The regime received strong international opposition and many of its laws were viewed as a violation of basic human rights. Since 1986, the White government gradually began to repeal several restrictions through a series of negotiations between the National Party and the African National Congress. The regime officially ended in 1994 with the first free democratic election by which Nelson Mandela became the first Black president of South Africa.

justify the removals of people on a large scale, typically leading to the relocation of entire communities and villages. Removals without homeland formalization and coordination with local chiefs could have jeopardized the long-term control of the homelands. This notion is consistent with several historical observations. Skelcher (2003) notes that *"[i]n 1970, the South African government established the Zululand Territorial Authority... Not coincidental, these changes came at a time when there was an upsurge in removals."* Luwaya (2018) succinctly highlights that "*[t]he establishment and consolidation of the homelands resulted in . . . a process of vicious forced removals."*

Between 1960 and 1980, approximately 3.5 million Black Africans were forcefully removed to the homelands (Simkins, 1983). Communities received notification eviction and, in some cases, public services like bus and water supply were cut off. Communities that refused to move were threatened with imprisonment. Bulldozers and government garage trucks were sent to destroy houses and transport people to the homelands. At the homelands, relocated people received small plots of land with a hut and latrine. Many of the displaced individuals had resided in White areas for more than one generation. As described by Hanoman (2017), "*many Blacks who had been born in the urban areas [in White South Africa] were suddenly forced to live in a homeland that was created for them by the White government.*"

The key rule for homeland assignment was ethnicity. With the political discourse of promoting tribal identity, the government targeted specific communities and relocated them according to ethnicity. In many cases, this rule was infeasible or too costly to implement because not all communities in White areas were perfectly ethnically organized. Using information from the 1980 census, we find that about 65 percent of migrants from White areas are residing in the homeland of their ethnic group.[11]

Strict migration control. Control of internal mobility was crucial for the success of the homeland system. The apartheid government developed a complex bureaucracy to enforce strict control over mobility through the Pass Laws Act. Blacks were required to carry a passbook known as *dompas* everywhere and at all times. It contained information about ethnicity, residential address, employer, employment history, among others. Police could ask Black people to show them this document at any time. Blacks could not move out of or between homelands, and illegal mobility was penalized with imprisonment. It was not easy to escape these high barriers to mobility. As Lemon (1984) notes, *"Probably no avowedly capitalist country controls its labor market to the same degree as South Africa."* This strict control of internal migration

[11]This imperfect homeland assignment along ethnicities is not an issue for our research design we describe below, as our assignment of childhood exposure is not based on ethnicity but on the homeland where migrants are observed at census time. As we discuss below, the strict control on internal migration during the apartheid era means that the most migrants were likely to reside in the homelands where they were relocated during the apartheid regime.

means that at least during the apartheid era, the vast majority of migrants was likely to reside in the homelands where they were relocated.

It is important to note that, although free migration between and out of homelands was prohibited, homeland residents were allowed to look for work in White areas under special permissions granted by White authorities (Greenberg and Giliomee, 1983). These migrant workers were engaged in short-term contracts, after which they had to return to their homelands for some time.[12] According to the 1980 census, approximately 13 percent of employed individuals in the homelands reported that their place of work was in a White area.[13]

2.2 Education in South Africa

The formal education system in South Africa involves seven years of primary education (which is subdivided into four grades in a foundation phase, three grades in an intermediate phase, and a grade in the senior phase), and five years of secondary school (which include two grades in a senior phase and three grades in the Further Education and Training phase). Although primary school was compulsory for all racial groups by law, it was rarely enforced in practice. The central government offers grants to partly fund school operational costs, including salaries, books, and other educational materials. School governments must supplement these grants with other sources of income, including that from local governments and school fees. This is especially true in the homelands where the grants received from the government were less generous.

There were important discriminatory reforms that affected all African Blacks. A major reform was the introduction of the Bantu Education Act in 1953. It implied the implementation of separated educational facilities for Blacks, and the content of education was aimed to prepare them for the unskilled labor market. Although this national reform certainly had profound consequences on education inequality between Whites and Blacks, it does not threaten the validity of our research design because we exploit a relative narrow source of variation within Black birth cohorts.

3 Data Sources, Samples and Definitions

This section provides an overview of the data sources and estimation samples used in our analysis. Additional details on the data and descriptive statistics are provided in Appendix A.

[12] As described by the Minister of Bantu Administration in a public speech, *"...the Bantu is only allowed to be in White area for the labor that he offers and the moment he no longer meets this condition the grounds for his presence in the White area are no longer valid"* (Mare, 1980).

[13] Labor migration may affect our identification strategy if we are less likely to observe information on migrant workers in the data. Fortunately, census enumerators collected information for all household members, including those who were temporarily absent due to labor migration.

3.1 Census Individual-Level Data

Our analysis focuses on census data conducted during the apartheid era. Census enumerators collected information for all household members, including those who were temporally absent.[14] This avoids selection issues due to temporary labor migration.[15] The provinces under apartheid include Cape, Natal, Transvaal, Oranje Free State, and the 10 homelands. Each homeland is labeled as a different province in the census, and we use both terms interchangeably.

We use individual-level data from the 1980, 1985, and 1991 censuses, which provide basic information on demographic characteristics as well as information on educational attainment, our key outcome of interest. We use the 100 percent census samples obtained from Statistics South Africa (1980, 1985, 1991). The core census files provide information on all individuals in South Africa but exclude a few homelands (listed in Appendix Table A.1). We use information on an individual's place of birth to identify migrants of White-place origin and assign a measure of childhood exposure, which is likely to be accurate given the strict restrictions on internal mobility during the apartheid regime. Information on place of birth is particularly geographically detailed at the Magisterial level in the 1980 census round. This information is more limited in the other census waves: census enumerators asked respondents to indicate the homeland/province of birth, but provinces in White areas were grouped into a single category called "Republic of South Africa."

In supplementary exercises, we use the post-apartheid census conducted in 1996 (Statistics South Africa, 1996). We use the 10 percent randomly drawn sample available from Statistics South Africa. Unlike previous census data, the 1996 census covers all the former homelands. Census enumerators did not collect information on an individual's place of birth in 1996, but we can use the following questions to identify migrants of White-place origin: "Here do you live now? Where did you live before this?." As discussed in Appendix A, the potential for measurement error in these data is unlikely to be severe because the fraction of individuals who moved more than once in South Africa is relatively low —about 5 percent.

3.2 Outmigration after 1986

One important caveat with using the censuses conducted after 1986 is that all legal restrictions on internal migration were removed this year. Consequently, individuals could move outside

[14]Information for temporarily absent residents was provided by other household members or collected after returning.

[15]One might still be worried about measurement error for individuals who were temporarily absent due to labor migration. However, it does not seem plausible that this measurement error was differential for locals and migrants and across different birth cohorts. Consistent with this notion, the results remain virtually unchanged when we exclude Magisterial Districts where a high fraction of the labor force is labor migrant (i.e., more than 20 percent). In this case, the estimated coefficient of interest is -0.10 (standard error =0.024) —extremely similar to our baseline -0.11 (standard error=0.020).

their homelands for living or other purposes. If individuals systematically moved out of the homelands, then it may introduce noise in our intent-to-treat design and thus the estimates may represent attenuated impacts of policy. However, historical accounts suggest that the anticipated mass outmigration did not occur in the immediate years following the abolition of the internal passport system (Reed, 2013). Appendix C documents that only 1 percent of individuals moved out of the homelands between 1986 and 1991, and 6 percent between 1986 and 1996.[16] Hence, outmigration seems unlikely to be a major issue in the 1991 census, but it could introduce an important attenuation bias in the census conducted after democratization.

3.3 Estimation Sample

We construct several estimation samples, one for each census wave. We first limit the data to individuals residing in a homeland at census time, all of whom we match to homeland formalization laws based on their place of residence. We then restrict the sample to individuals born in the 1937-1969 birth cohorts who were age 23 or younger when their homeland of residence was established. In the 1980 census, the number of observations drops sharply when considering individuals who were under age 2 at homeland formalization because several homelands were established after 1970, and we observe their outcomes in 1980. Therefore, we exclude these individuals to increase precision.[17]

Preferred sample. We chose the 1980 census sample as the benchmark for two reasons. First, while outmigration was relatively low between 1986 and 1991, our measure of childhood exposure is likely to be more accurate in 1980 than in subsequent censuses conducted after 1986. Moreover, the 1991 census suffered from severe incomplete coverage within the homelands due to political violence and budget restrictions in 1991 (p. 20 Human Sciences Research Council, 2007). Second, the 1980 census is the only wave with granular information on the place of birth, which allows us to carry out a comprehensive investigation on the role of places in driving the basic findings. While the 1980 census is our preferred sample, we do exploit the relative strengths of the other census data in supplementary exercises.

[16]Data from the South African Internal Migration Survey 1999-2000, which provides detailed migration histories, indicate that these figures are remarkably similar when comparing migrants and homeland natives.

[17]The fraction of individuals who were 2 years old at homeland formalization is about 6 percent. The same figure for those who were only one-year younger is about 3 percent. The inclusion of these individuals in the baseline sample does not appreciably affect the overall estimate of interest. However, the number of observations is insufficient to reliably estimate the impact of the policy separately at each age of exposure. To increase the number of observations of these groups, one would have to include individuals who are too young at census time and thus less likely to have completed their schooling decisions. We overcome this limitation with the census conducted in 1991, when the birth cohorts were 11 years older.

4 The Impact of Moving to the Homelands

In this section, we estimate the long-run impacts of the relocation policy on human capital using a cross-cohort identification strategy. We first provide a detailed description of the research design and present the key identifying assumption. We then report baseline estimates and a number of specification checks to address potential concerns regarding the interpretation and internal validity of the results.

4.1 Research Design

4.1.1 Overview

We employ an *age-at-policy* approach to identify the impacts of moving to a homeland on human capital attainment.[18] We leverage variation in the timing of homeland establishment for identification. Operationally, we predict for each migrant from White areas the age at which she or he would have arrived in the homelands based on *when* his or her homeland of residence was formally established. Since migration out of or between homelands was under strict control, most relocated people were residing in the homelands where they were placed, and therefore our assignment will be highly accurate —particularly before 1986.[19] Using policy-induced variation in the age at arrival in the homelands rather than the actual age at arrival (when available), we can avoid several potential selection issues —e.g., some individuals may have voluntarily moved to their homelands due to other factors such as changes in preferences or specific, individual economic circumstances. Our research design captures variation in mobility shocks that were, to a great extent, unanticipated and over which individuals had no control.

In principle, one would only use the (within-cohorts) differential timing of relocation due to differences in the timing of homeland establishment as a source of variation. However, a major concern with this approach is that the timing of homeland establishment might be correlated with unobserved factors. For example, the government may have found it easier to establish the homelands of poorer ethnic groups who were on a downward trend in socioeconomic con-

[18]The core idea of this approach is that individuals who were younger at the policy change had greater exposure during critical periods of child development than those who were older and completed their schooling decisions. A variety of studies has implemented age-at-policy designs to estimate the long-run impacts of specific shocks on human capital. A salient example is Duflo (2001).

[19]An alternative strategy would be to assign exposure to the homelands based on an individual's ethnicity since the relocation policy affects different ethnicities at different times. Under this approach, one would consider all Blacks born in White areas (stayers and movers) and define exposure to the policy based on the year their ethnic homeland was established but ignoring whether they actually migrated to the ethnic homeland. While we find this ethnicity-oriented approach appealing in principle, it has limitations in its implementation. First, in many cases, it was unfeasible to relocate individuals based on their ethnicity because not all communities in White areas were perfectly ethnically organized. Second, a substantial fraction of Blacks was not ultimately relocated. Together, these factors would imply a weak "first-stage", rendering the ethnicity-oriented approach less appealing.

ditions and educational investments. In addition, formalization may have coincided with other important changes in the homelands that might be correlated with human capital investments. To address these concerns, we use homeland locals as a comparison group in our analysis. This approach allows us to remove any pre-existing differential trends across birth cohorts and homeland-specific shocks that affect all individuals in a homeland similarly.

While the use of locals as a comparison group has important strengths, we acknowledge that it is not a panacea. A major concern with this approach is that it may violate the stable unit treatment value assumption (SUTVA) if the influx of migrants from urban areas affects educational resources and opportunities for the natives. We note that *a prior*, such spillover effects are likely to go in the same direction as the effects for migrants, and in this case, our estimates would represent lower bounds of true effects. We will present evidence suggesting that this is the case. While our focus is on models that use locals as a comparison group, we will also show results from a specification that includes only migrants in the estimation sample as well as from a specification that uses an alternative comparison group. While the estimates differ somewhat, they provide reasonable bounds of the likely effects of moving to the homelands.

4.1.2 Homeland Establishment and Timing of Moves

As discussed in Section 2, African Blacks in White areas were forced to move to their homelands following the formalization of these areas. To motivate our research design, we empirically examine how systematic this link is. While there is no complete information on the extent and timing of removals, we can use retrospective information on migration histories in the 1996 census to examine this question.[20] The census conducted in 1996 is the only one with information on the timing of moves. Based on the information on the year of arrival in the homelands, we define an event in individuals' residential history as a binary variable equal to one for the years after they arrived in the homeland and zero otherwise. This allows us to generate a microdata panel at the individual-year level.

[20]It is perhaps unsurprising that the apartheid government did not systematically collect official data on the exact number or timing of relocated people. An alternative source of information is the Surplus People Project (1983), which collected data on removals in several homelands. However, these data often cover a limited period, and information on the date of removals is missing in the vast majority of the cases.

We then estimate the following event-study model of the timing of in-migration:[21]

$$H_{ijt} = \alpha \; + \underbrace{\sum_{z=0}^{T} (Establishment)_{jtz} \cdot \beta_z}_{\text{post-policy period}} + \underbrace{\sum_{z=-T}^{-2} (Establishment)_{jtz} \cdot \beta_z}_{\text{pre-policy period}} \qquad (1)$$
$$+ \; \lambda_i \; + \; \gamma_t \; + \; \xi_{it}$$

where H_{ijt} is a dummy indicator for the years in which the migrant of White-place origin i is residing in the homeland j. *Establishment* is an indicator for z years between homeland formalization and year t for the individual belonging to the homeland j. The omitted category is -1. We include fixed effects for individual (λ_i) and year (γ_t). We allow the idiosyncratic error term to be correlated across individuals within a homeland-birth decade cohort. To estimate this event study on a fully balanced sample of individuals, we exclude individuals who are not observed during the entire pre- and post-policy periods. Balancing the panel implies a smaller sample but alleviates selection concerns due to different birth cohorts being observed in different moments in time (although the results are quite similar when using the extended, unbalanced sample).[22]Our baseline sample also excludes never movers —those who have always resided in a homeland.

Figure 2 plots β_z coefficients with 95 percent confidence intervals. It shows that there are no statistically significant trends in in-migration patterns in the pre-establishment period. The overall pre-policy trend is clearly flat, and the point estimates are close to zero. By contrast, the estimates increase sharply immediately after homeland establishment and remain positive during the entire post-policy period. The effects are large in magnitude. By the fifth year since homeland formalization, the likelihood of residing in the homelands increased by about 18 percentage points. Relative to the pre-formalization mean, this represents an increase of as much as 360 percent.

Mobility between homelands was severely restricted, except for those Blacks who were misplaced in a homeland other than that of their ethnic group.[23]Therefore, one could expect an effect on mobility between homelands driven by misplaced Blacks if there were something together with formalization that made it more lucrative to reside in the homelands. However, in Appendix Figure D.2, we show that homeland establishment does not predict the timing

[21]In an event-study framework with variable treatment timing where all units are eventually treated, the Difference-in-differences parameter is a weighted average of all 2 × 2 estimators in the data (Goodman-Bacon, 2018). As Goodman-Bacon (2018) highlights, this does not imply that the differences-in-differences design is wrong. It implies that trying to summarize the effects into a single post-treatment dummy could be problematic when treatment effects vary over time. Figure 2 suggests that this appears to be the case in our setting. For this reason, we focus on the event-study figures to summarize our results, and not a single post-treatment coefficient.

[22]For the interested reader, the results using the extended, unbalanced sample of individuals are shown in Appendix Figure D.1.

[23]In our data, more than 90 percent of migrants from other homelands correspond to individuals who were previously residing in a homeland belonging to a different ethnolinguistic group.

of homeland-homeland migration. The point estimates are small and statistically indistinguishable from zero, with the few statistically significant coefficients plausibly due to sampling variance. More importantly, the overall trend is clearly flat both before and after homeland establishment. This pattern is consistent with the apartheid government's focus on relocating Black people from White areas and not those who were already in a homeland.

Average lag in the age at arrival. The fact that the probability to live a homeland does not go from 0 to 1 rapidly could be in part due to measurement error, as we are using data from the Post-Apartheid census conducted in 1996 and it does not record an individual's entire migration history. However, the pattern documented in Figure 2 does suggest that there exists a lag in the time when people actually moved to a homeland. Therefore, we use the data to understand the average lag in the age at arrival. Figure 3 compares the ages at homeland establishment and arrival to a homeland. On average, individuals arrived in the homelands 6.8 years after their formal establishment. We use this result as inputs in what follows to construct a simulated measure of exposure to the homelands.

4.1.3 Simulated Exposure

Motivated by the evidence above, we construct a simulated measure of childhood exposure to the homelands based on the timing of homeland establishment and average lag in the age at arrival. In doing so, we first compute the simulated age at arrival for individual i born in year t and observed in homeland j established in year T is computed as follows:

$$\text{simulated age at arrival}_{ijt} = \underbrace{T^j - t}_{\substack{\text{age at} \\ \text{homeland} \\ \text{establishment}}} + \underbrace{\ell}_{\text{average lag}} \tag{2}$$

where ℓ is the average lag in the age at arrival, which we set to 7 years based on the evidence in Figure 3. We then calculate the predicted years of childhood exposure:

$$\text{simulated exposure}_{ijt} = (\bar{\mathrm{B}} - \text{simulated age at arrival}_{ijt}) \cdot \mathbb{1}\left\{\text{simulated age at arrival}_{ijt} \leq \bar{\mathrm{B}}\right\} \tag{3}$$

In this expression, $\bar{\mathrm{B}}$ is the upper age limit of exposure. In the simplest terms, equation (3) represents the number of critical childhood years predicted to have been spent living in a homeland. A large body of work in economics and medicine indicates that exposure to events and circumstances at earlier ages has greater impacts on individuals' lifetime outcomes (Heckman, 2007). However, these studies do not provide a guide on the exact years of childhood that matter. To guide our choice of the childhood years of exposure, we conduct a series of trend-break tests on the coefficients obtained from a flexible model that measures the effects at different exposure ages (described in detail below). As documented in Appendix E and the event-study figures shown below, these tests suggest that the upper age limit of exposure is 14. Thus, we set $\bar{\mathrm{B}}$ to 14.

In the previous subsection, we show that the timing of homeland establishment is a key determinant of in-migration patterns in the homelands using the 1996 census. In Figure 4, we further document that our predicted-based measure of exposure (3) is a strong predictor of the corresponding actual exposure measure. The actual measure of exposure increases by 0.8 years for each additional simulated year of exposure. Since we do not have information on actual exposure in our main estimation sample, we focus throughout the paper on intent-to-treat estimates based on the simulated exposure measure. However, after presenting the basic findings, we also will use the "first-stage" results documented in Figure 4 to calculate an approximate estimate of the treatment-on-the-treated estimates.

4.1.4 Estimating Equation

To estimate the long-run impacts of the relocation policy on human capital attainment, we employ the following specification:

$$\begin{aligned} S_{ijkt} = \alpha \; + \; \beta \left(\text{simulated exposure}\right)_{jt} \times \left(\textit{migrant}\right)_k \\ + \; \mathbf{X}'_{ijkt}\Omega \; + \; \eta_{jt} + \mu_k \; + \; \xi_{ijkt} \end{aligned} \tag{4}$$

where S is completed years of schooling of individual i in homeland j born in province k and year t. *Simulated exposure* is the predicted number of childhood exposure years, as defined in equation (3). *Migrant* is an indicator for migrants of White-place origin. Our key independent variable of interest is the product of the interaction between these two variables. The vector $\mathbf{X}'_{ijkt}$ denotes the set of additional individual characteristics, such as gender, ethnicity, and Magisterial District of residence. This specification includes a full set of homeland × birth-cohort (η_{jt}) and province-of-birth (μ_k) fixed effects.[24] Finally, ξ_{ijkt} is a random disturbance term.[25]

Identification. A causal interpretation of our estimates requires the counterfactual assump-

[24] We should remind readers that we identify migrants from White areas using information on an individual's province of birth. Hence, the inclusion of province-of-birth fixed effects obviates the need to include an additional indicator for migrants of White-place origin as a separate control.

[25] As in any difference-in-differences design with variation in treatment timing, the single parameter of interest β in equation (4) could be biased if the effects are heterogeneous across cohorts (Goodman-Bacon, 2018). However, note that since our specification includes a large group of "untreated" individuals (i.e., local natives), this issue is largely diminished. Moreover, our interaction term allows for (linear) heterogeneous effects across cohorts exposed early in life, a specification consistent with the pattern in the data. As Goodman-Bacon (2018) discusses, summarizing the findings in this way considerably reduces bias from time-varying effects even if there are no untreated units. This is particularly important when estimating our baseline model using either migrants or local natives only in the sample.

tion that migrants and locals would have followed similar cohort trends in schooling in the absence of the resettlement program. The inclusion of the full set of fixed effects absorbs a number of potential factors, including for example the Bantu Education Act and other discriminatory policies set nationally by the apartheid government. Any differences between locals and migrants that are invariant across birth cohorts will be absorbed by the province-of-birth fixed effects. By conditioning on Magisterial District fixed effects, we compare migrants and homeland natives residing in the same location within the homelands.[26] With the homeland $\times$ year-of-birth fixed effects our coefficient of interest is identified from within homeland variation only and *not* from differences between homelands.

One might still be concerned about differences between homelands that were formalized earlier and later. As mentioned above, the timing of homeland establishment was not assigned at random. The government may have found it easier to establish bureaucratic structures in the reserves of poorer ethnic groups that were less likely to invest in human capital. It is important to emphasize, however, that the variation we exploit is *not* between homelands but *within* homeland-birth cohorts. This is a relatively narrow source of variation: an unobservable factor would need to have different effects both across birth cohorts *and* across local natives and migrants *within* a given homeland. Our rich set of homeland-specific $\times$ cross-cohort fixed effects is able to remove all time-varying homeland-specific factors, including pre-existing trends and coincident homeland-specific shocks (e.g., changes in school or health spending in homelands), that affect a birth cohort belonging to a homeland similarly. Thus, while the dates of homeland formalization were not randomly assigned, our specially demanding specification is likely to capture many differences across homelands, across cohorts, and across migrants and natives, making our identifying assumption plausible. After presenting the basic patterns, we provide a variety of evidence that is consistent with the validity of our research design.

Inference. As in recent studies using cross-cohort identification strategies (e.g., La Ferrara and Milazzo (2017) and Cantoni et al. (2017)), we use standard errors clustered at the homeland-birth decade level to account for correlation in error terms across individuals within the same homeland and birth cohort (24 clusters). We chose this approach, rather than clustering by homeland only, for two reasons. First, the variation in our key independent variable of interest occurs at the homeland $\times$ birth-cohort level. Second, we have a relatively small number of homelands (only seven). In Appendix F, we critically evaluate the robustness of our core results to alternative inference procedures. We find similar results when we cluster by home-

[26]One might be worried that the actual place of residence within the homelands may be endogenous to the treatment exposure measure and thus affect the interpretation of the results. In Appendix G, we exclude the detailed set of Magisterial District fixed effects as a robustness check. While the inclusion of these detailed location fixed effects has little impact on the estimated coefficient, it reduces substantially sampling variance and increases the precision with which our parameter of interest is estimated.

land only and correct for the small number of clusters using wild bootstrap (Cameron et al., 2008). Moreover, we conduct a permutation test that compares the estimated coefficient to the distribution of placebo effects obtained by *randomly* assigning placebo formalization dates to homelands. Again, our statistical inference is not appreciably affected when implementing this permutation test. In light of this evidence, we focus on our computationally simpler baseline approach throughout the paper.

4.2 Baseline Estimates

Table 1 reports results from estimating equation (4). Columns (1) through (4) present results based on the 1980 analysis sample. In column (1), we start with our baseline specification that includes homeland × birth-year fixed effects, province-of-birth fixed effects, and individual controls for gender, ethnicity, and Magisterial District of residence. The estimated coefficient of interest is -0.116 (standard error=0.019) and highly significant at the conventional levels of significance. It implies that migrants who potentially moved to a homeland at an earlier age attain fewer years of schooling.

Column (2) adds other individual-level characteristics, including religion, marital status, and first home language. Adding these controls increases the precision with which the parameter of interest is estimated, but the magnitude of the coefficient hardly changes. It is now -0.1102, with a standard error of 0.016. We do not include these additional individual controls in our preferred specification because they could be endogenous. For example, educational attainment could also determine religiosity and marital status. The inclusion of these controls is nonetheless a useful and simple check that our results are unlikely to capture differential trends in potential determinants of schooling.

Columns (3) and (4) repeat the baseline specification separately for migrants and homeland natives to assess how the use of the latter as a comparison group affects the results. Note that in this specification, we cannot use the rich set of homeland × birth-year fixed effects because the key underlying source of variation is at the homeland and cohort level. And we should remind readers that the possibility of homeland-year unobserved shocks is the motivation to use locals as a comparison group in the first place. To mitigate these possible confounders, we include homeland-specific linear pre-trends.[27]

The estimated coefficient for the subsample of migrants is -0.132 with a standard error of 0.057 (column 3). This estimate is somewhat more negative than the baseline. Column (4)

[27] An alternative approach would be to include homeland-specific linear trends. However, these controls are inadequate in our context because the policy change likely affected the outcomes in levels and trends. As we will show below , younger cohorts with an earlier predicted move experienced larger negative impacts. Consequently, the inclusion of homeland-specific linear trends could introduce an important bias in our estimates, an issue recognized in several studies using difference-in-differences analyses (e.g., Lee and Solon, 2011; Goodman-Bacon, 2016). The use of homeland-specific linear *pre*-trends addresses this issue to a great extent.

shows the results for locals only. One could observe significant effects for locals if there were homeland-specific shocks differentially affecting the native group that are not captured by the homeland-specific linear cohort trends. One could also observe significant effects in the presence of spillovers. If there are negative spillover effects from migrants to natives (e.g., migrants crowding out natives' enrollment, increases in pupil/teacher ratios), then the coefficient for homeland natives will go in the same direction as the main estimates. Alternatively, if there are positive spillovers (e.g., increased competition making local kids work harder), then the coefficient for locals will be positive and thus would bias our main results toward zero.

We find that the coefficient for locals is also negative but statistically insignificant. Most remarkably, the point estimate is about 80 times smaller in magnitude than that for migrants. These results suggest either that there are no unobservable, specific-cohort shocks differentially affecting locals or that there are no meaningful spillover effects (or if anything, they are negative and attenuate our main coefficient). In any case, these results are reassuring and leave us more confident about our identification strategy.

In columns (5)-(6), we repeat our baseline specification but using the 1985 and 1991 apartheid censuses. In these samples, all cohorts are older and thus we can extend the baseline sample by including younger cohorts at formalization. The 1985 sample includes fewer homelands than the 1980 sample (6 versus 7). The 1991 census includes an additional homeland that is not available in the 1980 census, yet the former sample is approximately 6 percent smaller. When comparing the same homelands that are observed both in the 1980 and 1991 censuses, we find that the resulting sample derived from the latter is about 20 percent smaller. As mentioned above, this difference in sample sizes reflects the fact that many poor and remote areas within homelands were not enumerated in 1991 due to political violence and budget limitations at the time (p. 20 Human Sciences Research Council, 2007). Therefore, caution is warranted when interpreting the results from this census wave.

In line with the baseline, we find a negative and statistically significant difference-in-differences coefficient in the 1985 and 1991 samples. Quantitatively, the point estimate derived from the 1985 sample is extremely similar to the baseline (-0.10 versus -0.11). On the other hand, the estimated coefficient obtained from the 1991 sample is significantly attenuated (-0.043 versus -0.11). This difference in the point estimates is most likely the result of the lower coverage of the 1991 census, confirming the strengths of our main approach. Other reasons for these differences might be i) outmigration after 1986 and ii) some cohorts being able to complete school at a later age. The former cannot explain much of these differences given the relatively low rate of outmigration between 1986 and 1991 (see Appendix C). Similarly, returning to school does not seem to be an important explanation because the results are extremely similar in the 1980 and 1985 samples, and even the youngest members of the 1985 sample (individuals aged 16) had certainly completed their schooling decisions.[28] In any case, the results from the 1991 apartheid census replicate the general picture presented so far with precision, suggesting that these data are useful to examine the impacts of the relocation policy.

A possible concern with the baseline estimates is that all censuses carried out during the totalitarian apartheid regime might be of questionable validity. Appendix Table G.1 investigates whether our findings continue to appear in the post-apartheid census conducted in 1996. As discussed above, it is important to keep in mind that our measure of childhood exposure is likely to be somewhat less accurate in this census because of increased outmigration after 1994 and because we use previous place of residence (rather than place of birth) to identify migrants due to data constraints. Using these data, we find an estimated coefficient that is smaller in magnitude than the baseline (-0.11 versus -0.072), but statistically significant and meaningful. This suggests that the decline in schooling we observe is unlikely to be an artifact of the data collected during the apartheid regime.

Magnitude. Overall, the results of this section suggest that the relocation policy reduces educational attainment of individuals exposed at earlier ages. Estimates from our preferred specification in column (1) suggest that a predicted move at age 9 would reduce educational attainment by about 0.55 years (=-0.11 × 5 years of childhood exposure). Relative to the sample mean, this estimate represents a decline of approximately 11 percent.

Treatment-on-the-treated. Our estimates represent an intent-to-treat (ITT) effect because our measure of exposure is simulated based on an individual's age at homeland establishment. We can calculate an approximate estimate of the treatment-on-the-treated (TOT) effect by dividing the ITT estimate by the actual-versus-simulated exposure gradient reported in Figure 3.[29]As shown in the figure, each additional year of simulated childhood exposure to the homelands is associated with an increase of 0.8 years of actual childhood exposure to the homelands. This implies a TOT coefficient of -0.13 ($\approx -0.11/0.8$). Thus, moving at age 9 to a homeland would reduce educational attainment by 0.65 years, an effect extremely similar to that implied by our ITT estimate.

[28]Moreover, when we limit the 1985 sample to individuals over 18 years of age, the estimates remain virtually unchanged. This is unsurprising because educational attainment in South Africa is low and thus most individuals have dropped out of school by age 14.

[29]This is a "back-of-the-envelope" calculation of the TOT effect. This calculation approximates the IV coefficient:

$$\beta_{IV} = \frac{\beta_{ITT}}{\beta_{\text{first stage}}} = \frac{\text{cov}(y,z)}{\text{cov}(x,z)}$$

where y is the outcome of interest, z the simulated exposure measure, and x the actual exposure variable.

4.3 Validation of the Research Design

In the previous section, we have argued that the differential age-at-exposure patterns reflect causal impacts of the relocation policy rather than pre-existing differential trends or other coincident shocks. In this section, we present evidence in support of this claim.

4.3.1 Effects at Different Ages of Exposure

The validity of our statistical approach relies on the common trend assumption discussed above. To assess the plausibility of this identification condition, we estimate a flexible version of equation (4) where we estimate the effects at different predicted ages of exposure:

$$\begin{aligned} S_{ijkt} = \alpha \ + \sum_{\tau=0} \beta_\tau \ \mathbb{1}\big\{\text{simulated age at arrival} = \tau\big\} \times \big(\mathit{migrant}\big)_k \\ + \ \mathbf{X}'_{ijkt}\Omega + \eta_{jt} \ + \ \mu_k \ + \ \xi_{ijkt} \end{aligned} \tag{5}$$

where $\mathbb{1}\{\cdot\}$'s are indicators for ages $\tau = 7, 8, 9, \ldots, 30$ of exposure. β_τ is normalized so that it is equal to zero for individuals who were 25 years of age when their homeland was formally established. The rest of variables and parameters are the same as in equation (4). The coefficients $\beta'_\tau s$ provide a detailed depiction of the policy impacts at different ages of exposure in a transparent way. If the baseline estimate in Table 1 reflects causal effects of moving to the homelands, then one should not observe significant effects on cohorts who were exposed during non-school-age years (i.e., after age 16), when most of individuals have likely completed their schooling decisions. Meaningful estimates and any clear tendency toward improving or deteriorating schooling among cohorts exposed during non-school-age years would indicate that our main estimates are driven by preexisting differential trends across birth cohorts.

Figure 5 plots estimates of β_τ and 95 percent confidence intervals based on the non-parametric model (5). There are clear significant and negative effects when exposure occurs at earlier ages, particularly below 15 years of age. By contrast, there are no significant effects on educational attainment of individuals with predicted ages of exposure above 16. In addition, the placebo estimates are very small in magnitude and without any tendency toward improving or deteriorating schooling. Appendix Figure G.1 repeats the same exercise separately for migrants and locals. For the subsample of migrants, the pattern is essentially the same as that observed in Figure 5: migrant schooling declines when the move is predicted to occur below 15 years of age, without any tendency above age 16. By contrast, for locals, the placebo age at arrival does not predict schooling. The path of the placebo coefficients is clearly flat, either below or above age 16.

Overall, the results of this section provide evidence that our estimates are very unlikely to be driven by pre-existing differential trends.

4.3.2 Falsification

To further evaluate the plausibility of the identification condition, we exploit a specific feature of the relocation policy to implement a falsification exercise: its focus on removing Blacks residing in White areas rather than those individuals who were already residing in a homeland (even if the homeland was not that of their ethnic group). Although misplaced Blacks in a homeland of other ethnic groups were allowed to move if they wanted so, they were not directly induced to move. Consistent with this notion, all of the removals documented in Surplus People Project (1983) involved individuals who were residing in White areas. And the estimates in Appendix Figure D.2 show no corresponding significant effects of homeland establishment on the timing of homeland-homeland migration. If our results are causal and do not reflect other factors correlated with homeland establishment, then we should not observe corresponding significant effects on education for homeland-homeland migrants since our exposure measure does not predict the timing of these moves.

Note that this is not a test of voluntary homeland-homeland migration versus forced White place-to-homeland migration. It is possible that homeland-homeland migration has causal effects, but this exercise will not (and is not intended to) evaluate this point. The purpose of this exercise is to evaluate if our design would yield similar results for a group whose timing of migration is uncorrelated with the date of homeland establishment. Meaningful effects on education would suggest that our main results are capturing unobserved factors correlated with the timing of homeland establishment.

Following this reasoning, we perform this falsification test. We replace our group of migrants of White-place origin with the group of homeland-homeland migrants and rerun our baseline specification. With homeland-homeland migrants, we mean individuals born in a homeland but moved subsequently to another homeland. These results are displayed in Table 2 and in Appendix Figure G.2. Column (1) repeats the baseline for ease of comparison. Column (2) shows no systematic association between educational attainment and predicted childhood exposure. Reassuringly, the estimated coefficient is not only statistically insignificant but very small in magnitude. It is estimated at -0.01 with a standard error of 0.009. Our baseline estimates allow us to rule out an effect five times as large as the placebo coefficient.

4.3.3 Selection of Migrants

An objection to the validity of our results is migrant selection. If individuals with poorer economic prospects were differentially induced to move, our estimates might be driven, not by the relocation policy but by selection.[30] However, note that any differences due to selective

[30] Selection into migration may occur in our setting if for example some individuals moved to the homelands because they believed that homeland establishment would make it more lucrative (e.g., more employment

targeting of migrants will be a threat only if they are correlated with an individual's age at formalization. The absence of differential trends for exposure at older ages (as shown in Figure 5) already suggests that this is unlikely to be the case. Perhaps most importantly, the nature of the removals, which were typically executed en masse and within a short timeframe, makes such stories difficult to construct.

To explore this issue, we compare the levels of human capital attainment of pre-policy cohorts of movers and stayers (i.e., those born before 1950). If different individuals were induced to move, one would observe significant differences when comparing both groups. Figure 6, Panel A provides the mean educational attainment separately for movers and stayers, with 95 percent confidence intervals. As can be seen, educational attainment between both groups is strikingly similar, with a difference of only 0.05 years that is statistically insignificant. This suggests that low- and high-education individuals were not differentially induced to move, which is consistent with the nature of the removals were implemented. As an additional check to assuage concerns that individuals with a declining trend in human capital were more likely to move, Panel B compares cohort trends in schooling attainment between both groups. There are no differences in cohort trends in schooling.

Overall, movers and stayers who had completed their schooling decisions before the policy change have similar educational attainment, both in levels and trends. These pictures suggest that selection is not responsible for our results.

4.3.4 Selection into the Sample

i) Mortality. One may be worried about the possibility that moving to the homelands at earlier ages increased child mortality, which may change the composition of the sample. However, child mortality is of less concern for our identification strategy as it most likely would affect lower quality individuals at the margin of survival. This kind of differential mortality selection would make it more difficult to detect meaningful impacts of relocation, and if so, our results showing significant estimates become even more telling.

Of greater concern is attrition due to mortality in the older migrant cohorts. It would be problematic if migration to the homelands had larger negative impacts on life expectancy for the least educated in the older migrant cohorts. This mortality impact would change the pool of older migrants and skew the observed education levels upward, biasing our difference-in-differences coefficient toward finding a negative effect on schooling. Several pieces of evidence suggest that this is not a major issue. First, the oldest cohort in our sample was 42 years of

opportunities, more school investments) to reside in such areas. The results in Appendix Figure D.2 showing that homeland residents did not move to their native homeland areas in response to formalization indicates this was not the case.

age in 1980, and life expectancy in South Africa during the 1960-1980 period was about 53.[31] While life expectancy was low during the study period, the oldest individuals in our sample are relatively young for this mortality impact to materialize and systematically change the composition of the sample.

Second, if the least educated individuals in the older migrant cohorts are systematically less likely to be observed in the 1980 census due to mortality and this mortality impact is large enough to introduce an important bias in our estimates, then we should observe a clear tendency toward improving schooling when exposure occurs at older ages. But we do not. Estimates from the flexible specification (5) reveal that the pattern of age-specific effects is clearly flat for exposure after school-age years (see Figure 5). Appendix Figure G.3 repeats the flexible specification but extends the baseline sample by including older individuals in the 1980 sample. Again, we do not observe any clear tendency toward improving schooling even when including individuals whose predicted age at arrival is above 30.[32] We observe significant effects only for migrant cohorts potentially exposed during the school-age period.

Third, if older migrant cohorts at the predicted move were less likely to survive in 1980, then one would observe a differential reduction in cohort sizes for these individuals relative to younger individuals. Appendix Figure G.4 shows no evidence of significant differential changes in cohort sizes for older migrant cohorts. This is true even when we extend our baseline sample and include individuals who were as old as 30 years of age at homeland establishment. More importantly, Appendix Figure G.5 documents that these patterns are not significantly different for less- and more-educated cohorts who most likely had completed schooling decisions at homeland establishment. This is consistent with the flexible estimates above showing no differential trends in schooling for older migrant cohorts. In sum, there is no indication that attrition due to mortality is a significant problem.

ii) Fertility. A related concern is that parents might have changed fertility decisions if they anticipated the date of homeland establishment. It is possible that homeland establishment implied increased uncertainty about the future and, as a consequence, families reduced the demand for children before the expected move, potentially changing the composition of birth cohorts in our estimation sample. However, it seems implausible that families would have anticipated the exact date of homeland establishment given the political and other arbitrary factors involved in the formalization process of the homelands, such as the length of the time that it took to establish bureaucratic and administrative structures. It is more plausible that families changed fertility decisions *after* the move, but this does not affect the composition of

[31] See https://www.theglobaleconomy.com/South-Africa/Life_expectancy/.

[32] Not surprisingly, the difference-in-differences coefficient obtained from our baseline specification (4) remains virtually unchanged when using this extended sample (-0.116 versus -0.114).

our estimation sample because we focus on cohorts born before homeland establishment.[33]

If families systematically changed fertility decisions before homeland establishment, then one should observe meaningful changes in cohort sizes. However, Appendix Table G.2 and Figure G.4 show no indication that our measure of childhood exposure significantly predicts cohort sizes.[34] We can further explore the possibility of changes in fertility decisions by looking at the October Household Survey (1994-1995) from Statistics South Africa (1994, 1995), which provides detailed retrospective information on birth histories for all women aged 12 to 54, including the timing of all births.[35] As described in Appendix G.5, we construct a pseudo panel of women-by-year and explore trends in the probability of giving birth before homeland establishment. Consistent with the birth cohort results, there are no differential trends in the probability of giving birth in the 9 years prior to homeland establishment. Moreover, when we look separately at adult women who are more and less educated, the patterns are very similar. Summarizing, changes in the composition of the birth cohorts due to fertility are unlikely to be responsible for the decline in schooling we document.

4.3.5 Selective Outmigration from South Africa

Another source of selection is that many individuals in White areas may have left South Africa in response to the relocation policy. This would affect our estimates only if younger individuals at the homeland formalization were differentially more or less likely to move out of South Africa. We begin by noting that the pre-policy cohorts of stayers and movers are comparable, as documented in Section 4.3.3. These results suggest that if there was increased migration out of South Africa, it does not appear to have been selective. As a further check, we reconstruct our sample by excluding all Magisterial Districts in White areas close to the international border with neighboring countries, where the costs of outmigration are arguably lower than in more distant areas. As shown in Appendix Table G.3 and Figure G.8, this restriction does not materially affect our baseline results. Selective outmigration is thus unlikely to explain the patterns in Table 1.

[33] Nevertheless, changes in fertility and thus family size after the move may be a mechanism behind our results, and we discuss this possibility in Section 5.

[34] Another possible consideration is that the "shock" to fertility might have come when the first round of homelands was created. Perhaps this information shock affected the fertility of women who knew they would move to a homeland at some point (even if the move came several years later). Conceptually, this would imply that the timing of establishment of the actual homeland is not correlated with such changes in fertility, and if so, it would not necessarily introduce a selection bias in our analysis. Moreover, the information shock would affect most women at the same time and in a similar way. Only differential changes in fertility correlated with homeland establishment would be an issue for our identification strategy. In any case, if such possible information "shock" is important for our findings, we should observe appreciable changes in cohort sizes. But the data do not suggest that this is the case.

[35] Alternatively, one would directly examine changes in birth rates through official vital registries. However, the government did not maintain official vital registries of birth for African Blacks during the apartheid era (Moultrie and Timæus, 2003).

4.3.6 Movers-versus-Stayers Design

Our baseline analysis uses locals as a comparison group. As a robustness check, we use stayers as the control group. This exercise represents an additional test of the research design. If this comparison shows corresponding negative effects, then it would be consistent with the interpretation that moving to the homelands has causal effects on human capital. But if this alternative design reveals no negative effects, that is stayers and movers are similarly affected, then it would suggest that our basic results are the product of an unobservable factor in origin places correlated with the timing of homeland establishment. For example, the discrimination against Blacks may have intensified in where they had lived in the wake of homeland establishment. This possibility is important given the lag in the time when people actually moved to the homelands.

To explore this issue, we estimate an analogous version of model (5) but replacing locals for stayers in the estimation sample. Since the model includes province-of-birth fixed effects, the parameter of interest is identified by comparing movers and stayers from the same origin place. As shown in Appendix Figure G.9, the results closely resemble our baseline estimates. Importantly, the timing of the effects is essentially the same: movers experience a differential decline in their educational attainment when they potentially move at ages below 15, without any tendency toward improving or deteriorating schooling above age 16. In sum, these results suggest that our estimates are not the product of unobserved shocks in origin places.

4.4 Further Robustness Checks

We have conducted a variety of additional checks to evaluate the robustness of our baseline estimates. To save space, the results from these additional analyses are shown in the Online Appendix G. In particular, we investigate the robustness of our main estimates to: alternative specifications (removing Magisterial District of residence fixed effects, controlling for Magisterial District of birth, exploring alternative measures of exposure), and alternative sample restrictions (excluding homeland-to-homeland movers, excluding big homelands, Magisterial Districts, or provinces). The results from these robustness exercises are broadly consistent with the overall pattern in Table 1.

4.5 Labor Market

A key question of independent interest is whether the homeland policy affected individual success in the labor market. Job opportunities in the homelands are limited, and many of the existing ones are dominated by the local government. To access these positions, particularly those that pay higher salaries, it is necessary to have some degree of formal education. While this suggests that the net gains from schooling in the homelands may be positive, it is still possible they are negligible so it is not worth investing in human capital acquisition. This might explain why we observe negative impacts on schooling in the first place. If this reasoning is correct, then income wages would be similar for all individuals irrespectively of when they arrived in the homelands, and thus childhood exposure should not have impacts on labor market outcomes. In this section, we provide evidence that is inconsistent with this claim: exposed cohorts are earning lower income wages as adults.

To explore this question systematically, we use the 1991 census where all cohorts are old enough to observe their labor market outcomes consistently.[36] The results are shown in Table 3 and Figure 7. We begin by examining the extensive margin. We create an indicator equal to one for individuals employed at census time and zero otherwise. Figure 7, panel A shows no evidence of pre-trends, with effects that are insignificant and close to zero when the move is predicted to occur at age 16 or older. By contrast, exposure at earlier ages leads to a reduction in the likelihood of working in adulthood. Column (1) shows that individuals with a predicted move at age 7 are 6.3 percentage points less likely to work in adulthood —or a 21-percent reduction relative to the sample mean. This effect is large but plausible given the particularly low employment rates in the homelands relative to White areas.

Column (2) repeats the baseline specification using the log income as the dependent variable, which is conditional on working and thus excludes the zeros. Here too the picture is very clear. We find that potential exposure to the homelands at age 7 is associated with a decline in income of about 5.6 percent and Figure 7 panel B suggests that this result is not driven by pre-existing differential trends.

One complication with interpreting these results is that moving to the homelands has negative effects on employment and this could change the sample of individuals with positive income in our data. To address this issue, in columns (3) and (4), we calculate lower and upper bounds by imputing information under best- and worst-case scenarios for individuals missing in the income sample due to the relocation policy.[37] The lower bound assumes that the "extra" unemployed individuals as a consequence of the policy change would be located at the top of the income distribution. The upper bound assumes that "extra" unemployed individuals would be located at the bottom of the income distribution. Appendix G.9 provides details on this procedure. Both lower and upper bounds imply negative effects on income. In practice, the most reasonable scenario is that lower-income individuals are less to be employed. In this case, the

[36] While individuals in the 1985 census are old enough to consistently observe their labor market outcomes, information on economic outcomes such as income was not collected.

[37] An alternative possibility is to use the inverse hyperbolic sine of income as the dependent variable, including the zeros for those who report zero income. However, this approach arbitrarily aggregates relatively large and small effects on the extensive and intensive margins. When we use this method, we estimate an effect on income that seems implausibly large (about 50 percent), particularly when compared to the upper bound estimated in Appendix G.9. This confirms the arbitrariness of aggregating zero and positive incomes.

upper bound estimate implies that our results could be underestimated by about 60 percent.

Overall, moving to the homelands in childhood led to poorer labor market outcomes in adulthood. The magnitude of the results is relatively large, and Appendix G.6 shows that schooling cannot account for all of the effects on income. This suggests that any mechanisms that affect education are likely to have also independent effects on subsequent labor market outcomes. These mechanisms could include for example changes in school quality, parental investments in general, or health deficiencies. This interpretation is consistent with recent work documenting that improvements in early-life conditions by, for example, reducing disease burdens (Bleakley, 2010) or increasing access to health care (Goodman-Bacon, 2016) or preventing nutritional deficiencies (Adhvaryu et al., 2020) has substantial effects on labor market outcomes but smaller effects on educational attainment.

We should emphasize at this point that these estimates are likely to be conservative relative to the lifetime effect due to life-cycle bias (Haider and Solon, 2006). Additionally, the relocation policy might have negatively affected the labor market outcomes of individuals who were exposed in adulthood, introducing a downward bias in our difference-in-differences estimates. Consequently, we take these results as lower bounds of the effect of the relocation policy on labor market outcomes for individuals who were younger at the policy change.

5 Homelands and Place Effects

Thus far, the results have demonstrated that moving to the homelands has negative effects on human capital. The possible causes behind these effects can be divided into two broad categories: changes in individual-specific circumstances and changes in childhood environment. The former could include for example loss of immobile assets such as housing, increased uncertainty, discrimination in the homelands, school interruption, and changes in preferences. Changes in childhood environment potentially include employment opportunities for parents, school quality, peers, access to healthcare, and other important place-specific inputs.[38] In this section, we shed light on these forces by directly investigating the causal effects of childhood environment.

5.1 Empirical Strategy

Overview. We follow an approach similar in spirit to Chetty and Hendren (2018) to investigate the overall effects of childhood environmental conditions. Intuitively, if places have causal effects, then children's schooling should tend to converge toward the levels of schooling in where they move to. When the schooling gap between origin and destination is high, the impact on

[38]As mentioned in the Introduction of this paper, living conditions in the homelands were poorer than in White areas. Appendix B provides a more detailed discussion of living conditions in the homelands.

schooling would be larger in magnitude than if the schooling gap is small. We therefore expect to see larger reductions in schooling among migrants from areas where the expected level of educational attainment (or simply, predicted schooling) is high relative to that of destination areas.

Following this reasoning, we proxy schooling predictions with Magisterial District-level average years of schooling, calculated from pre-policy cohorts born before 1950.[39] We use an individual's Magisterial District of birth rather than that of actual residence to compute average years of schooling, based on the sample that only includes Blacks. Migration before the homeland policy was relatively low, and thus most pre-policy cohorts attended school in their place of birth.[40] We then calculate the predicted schooling gap for each individual, defined as origin minus destination predicted schooling.[41] On average, migrants moved to areas where schooling predictions are 0.75 years lower —a difference of about 25 percent.

Figure 8 illustrates the intuition of the empirical strategy. It shows how origin-destination differences in predicted schooling are related to migrant schooling. Specifically, we calculate the mean difference in schooling between migrants aged 9-12 and 28-30 at the predicted move from both the same origin and destination places. We then break down the predicted schooling gap into twenty equal-sized bins and collapse the data into these cells. As can be seen from the non-parametric binned scatter plot in Figure 8, the decline in migrant schooling is larger in magnitude when the origin-destination difference in predicted schooling is larger. The slope of this relationship is negative and highly statistically significant.

Estimating equation. We can generalize these comparisons by using the full sample of migrants and estimating the following equation:

$$\begin{aligned} S_{idojt} = \alpha \; + \; \beta \left(\text{simulated exposure}\right)_{jt} \times \Delta_{od} \\ + \; \mathbf{X}'_{idojt}\Omega \; + \; \mu_o + \; \mu_d + \; \eta_{jt} + \; \xi_{idojt} \end{aligned} \tag{6}$$

where S_{idojt} is completed years of schooling. $\Delta_{od} = \overline{S}_o - \overline{S}_d$ is the difference in predicted schooling between origin o and destination d. We include a set of Magisterial District of origin and destination fixed effects, captured by μ_o and μ_d respectively. This expression is analogous

[39]The results are almost unchanged if we instead use older cohorts such as those born prior to 1945 or 1940.

[40]An alternative approach would be to calculate educational attainment predictions based on permanent residents (i.e., individuals residing in their place of birth at census time). However, one might be concerned that permanent residents are a particular, selected group that does not necessarily represent the true expected level of schooling in a given area. Despite this potential issue, we find results that are similar to the baseline when constructing predicted schooling gaps using permanent residents.

[41]One might be concerned with the possibility of measurement error due to individuals moving between districts within homelands after the initial relocation, as there were no barriers to mobility within the homelands. But if this occurs systematically, then one should observe a high fraction of individuals that moved more than once across districts in their entire life cycle. But data from the South African Internal Migration Survey 1999-2000 indicate that this figure is only 5 percent, suggesting that measurement error due to multiple movers is unlikely to be a major issue.

to equation (4) but replaces the migrant dummy for Δ_{od} and includes Magisterial District of origin fixed effects. The rest of the terms are the same as in equation (4). The parameter of interest is β, which measures the effects of earlier exposure to destination areas with a lower and higher predicted schooling relative to origin places. The estimation sample is based only on migrants from White areas, but we also present results from an extended specification that includes homeland locals as a comparison group.

Identification. The parameter of interest in our estimation of equation (6) is identified by comparing migrants from the same origin and destination places who differ in the predicted age at the move. With the inclusion of homeland $\times$ birth-year fixed effects (η_{jt}), we limit these comparisons within homelands only. Note that even if families with less propensity to invest in children's education were more likely to move to worse areas, this does not threaten the validity of the identifying assumption. Identification requires only that the timing of such moves is uncorrelated with differential trends in unobservable factors across cohorts. This is reasonable given that we exploit policy-induced variation in the timing of the moves, something over which individuals had no control.[42]

Identification would be threatened if the apartheid government was more likely to selectively relocate individuals who were on a cross-cohort tendency toward deteriorating schooling from higher- to lower-education areas. However, the large-scale nature of the relocation policy makes this kind of selection stories implausible, and the evidence in Section 4.3.3 indicates that relocation appears to be largely unrelated to individuals' education. Moreover, we will also present results from a non-parametric version of equation (6) and show that there is no clear tendency toward improving or deteriorating schooling among cohorts exposed during non-school-age years.

5.2 Results

The results are presented in Table 4. For ease of interpretation, we normalize the predicted schooling gap term by dividing it by the standard deviation (or equivalently 90 percent of the interquartile range). We present results from our main specification and some alternative versions. Column (1) presents results from a specification that adjusts only for origin place fixed effects and birth-year fixed effects. Columns (2) and (3) sequentially add destination place fixed effects and homeland $\times$ birth-year fixed effects. In column (4), we use homeland locals as a comparison group, so the key variable of interest is now interacted with an indicator for migrant status. In addition to the control variables included in column (3), this specification includes migrant $\times$ birth-year fixed effects to account for overall heterogeneity between migrants and

[42] This rules out potential biases due to more "ambitious" and higher-education parents moving earlier from worse to better areas.

locals. In all specifications, we include indicators for gender and ethnicity.

The results confirm the basic picture in Figure 8. Moving from higher-to lower-education areas leads to a larger decline in migrant schooling. The point estimate is extremely similar across alternative specifications and highly significant at the conventional levels of significance. The coefficient from our main specification in column (3) is -0.08 with a standard error of 0.013. It implies that migrants who potentially moved at age 9 to an area where predicted schooling is a one-standard-deviation higher relative to origin completed 0.4 fewer years of schooling.

We also present results from a non-parametric version of equation (6) that replaces *simulated exposure* with indicators for simulated ages at arrival. Figure 9 plots the age-specific coefficients and 95 percent confidence intervals. One can see that the path of these effects closely mirrors the overall pattern in Figure 5. The estimated coefficients are small and statistically insignificant for non-school ages. By contrast, there is a clear statistically significant decline in schooling for individuals who were exposed at earlier ages, particularly before age 15. The lack of pre-cohort trends yields strong support for the identification assumption that migrants in worse and better destination places would have experienced similar cohort trends in the absence of the mobility shock.

Gender-specific effects. As in recent work in the place effects literature (Chetty and Hendren, 2018), we also examine the gender-specificity of the heterogeneous effects documented above. Areas that are good for boys also tend to be good for girls. But boys could make better than girls and vice versa in certain areas and these differences can be substantial.[43] These differences could be the result of a number of factors, including for example gender differences in the opportunity cost of schooling, the returns to school investments, and the marriage market. Irrespective of the exact mechanisms, these differences suggest a simple placebo test: conditional on one's own gender schooling gap, moving to areas with particularly low schooling predictions for the other gender should have no corresponding significant effects on one's educational attainment.

To implement this test, we compute own-gender and other-gender predicted differences in schooling and interact them with our baseline measure of exposure. Column (5) of Table 4 replaces our key independent variable for these new gender-specific exposure measures. In this specification, we also include full interactions between an indicator for gender and the full set of fixed effects. The last line of Table 4 presents the p-value for the test of equality between the own- and other-gender exposure coefficients. Consistent with the place effects channel, we find that only the own-gender exposure measure has significant effects on migrant schooling. The own-gender exposure coefficient is -0.060 (standard error =0.019), which is very similar to the baseline and three times as large as the other-gender exposure coefficient. The coefficient

[43]For example, male schooling is twice as large as female schooling in some areas, whereas in other areas this pattern is completely reversed.

on the other-gender prediction is close to zero and statistically insignificant.

The gender-specificity of these results is striking and consistent with previous estimates by Chetty and Hendren (2018) in the American context. Since other unobservable factors and alternative mechanisms such as discrimination or school interruption are unlikely to generate childhood exposure effects that vary in proportion to the predicted schooling gap *in a gender-specific manner*, we conclude that the set of results in Table 4 must reflect the causal effects of places.

5.3 Discussion

We can combine our baseline results in Table 1 with the estimated effects of places above to evaluate the extent to which changes in childhood environment explain the overall decline in migrant schooling caused by the relocation policy. Our preferred estimate in column (1) of Table 1 suggests an overall decline of 0.55 (-0.11×5) years of schooling for individuals potentially exposed since age 9. Combined with our main coefficient in column (3) of Table 4, we estimate that changes in childhood environment can account for approximately 72 percent ($\frac{-0.08\times 5}{-0.11\times 5}$) of the overall decline in migrant schooling.

This magnitude is substantial but leaves space for other explanations. A prominent possibility is a direct effect of eviction. Forced relocation implied the loss of immobile assets such as housing, and as a consequence, parents could have changed investments in children's human capital even if they moved to destination places that do not differ from origin ones. As education is a normal good, wealth loss may have reduced educational investments. However, recent well-identified work on the direct effects of eviction in other settings suggests that this mechanism could work in the opposite direction. Perhaps the best-identified work in the literature is that of Becker et al. (2020). Studying a setting where they can arguably isolate the direct effect of eviction from other mechanisms such as place effects, they find an increase in migrant schooling. They argue and provide compelling evidence that this effect is driven by a shift in preferences from immobile assets to education, where education would serve as informal insurance in case of subsequent displacement/migration. Other studies find similar results in other settings (Nakamura et al., 2016; Sarvimäki et al., 2019).[44] If this mechanism is salient enough in our context, then it could have mitigated other adverse effects of relocation, including poorer conditions in destination areas.[45]

Another possible explanation is discrimination. While migrants were not discriminated

[44] Nakamura et al. (2016) find that children who were forced to move due to an environmental shock that destroyed their houses have increased educational attainment. Sarvimäki et al. (2019) document a similar pattern by studying a resettlement program that forcefully relocated Finish individuals during World War II. In these settings, families moved to destination places that do not appear to significantly differ from origin areas.

[45] However, unlike other settings in previous literature, the education insurance mechanism may have been less salient in our context because families moved to areas where they arguably knew they would be stuck.

against in their access to education and employment in the homelands by law, the de facto situation may have been different. The homelands were ancestral lands of natives who may not have liked that urban dwellers arrived. The local tribal structure had already been formed before the relocation policy and resettled people may have been excluded. If discrimination were widespread, it would likely have manifested in the labor market where discriminatory practices are especially difficult to regulate. One may expect, for example, the returns to schooling to be lower for migrants than locals if discrimination was widespread. However, Appendix Table G.7 shows no evidence that the labor market returns to education differ between migrants and locals. Of course, this evidence alone cannot rule out other more subtle forms of discrimination in the labor market and other dimensions, but it does suggest that discrimination is not the entirety of the story.

Another plausible possibility is that fertility and family structure may have changed after families arrived in the homelands. To explain part of the overall decline in schooling we find, parents should have increased fertility given the evidence that increased family size has negative effects on children (e.g., Black et al., 2010).[46] However, as documented in Appendix G.5, there is no evidence that the probability of a migrant woman giving birth increased in the years following the predicted move.

Finally, a possible reason behind the decline in schooling is school interruption. It is possible that when dozens of families arrived in a homeland, there would not be enough schools or teachers to attend migrant kids at the time of arrival. While the central government and local chiefs attempted to mitigate these issues by allowing newly arriving children to attend school as soon as possible, these measures were imperfect and not always enforced. Moreover, it is possible that forced moves are disruptive to a family. Families would need time to adapt to an unknown place and find schools, which may have caused children to delay school and many may have not caught up subsequently in terms of grade attainment. While plausible, this possibility is inconsistent with the pattern of exposure effects we find. To the extent this hypothesis is important it would imply larger impacts for individuals who moved later during their school-going ages as the scope for catching up is more limited at older than younger ages. However, we observe the exact opposite pattern: exposure at younger ages has larger effects on schooling.

6 Conclusion

We study the consequences of the homeland system implemented under the South African apartheid, one of the largest segregation policy experiments in history. Under this policy

[46] It bears noting that not all studies find evidence that family size significantly affects children. For example, Black et al. (2005) find that once they account for birth order, the relationship between family size and children's outcomes becomes zero. Similarly, Angrist et al. (2010) find limited evidence of a child-quantity/child-quality tradeoff using quasi-experimental variation due to twin-births and sibling-sex composition.

experiment, many Blacks were forcefully removed from their homes in "White" areas to the homelands. We leverage variation in the timing of homeland establishment to identify the impacts of this policy on human capital formation. Our empirical strategy is a difference-in-differences setup that compares migrants and homeland natives who were younger and older when their homeland was established. We find that the relocation policy reduced educational attainment of White-place migrants by about 0.55 years or 11 percent of the sample mean. Our results also indicate that childhood exposure to the policy led to poorer labor market outcomes in adulthood, with a decline of 20 percent in the likelihood of working, and a 5.6-percent decline in the intensive margin of income.

Our exploration of possible mechanisms suggests that an important part of these results reflects a change in childhood environmental conditions. Individuals who moved to the homelands at earlier ages had greater exposure to lower-opportunity places during critical periods of child development, and this negatively affected human capital attainment. These findings have implications for policy. Relocation to lower opportunity areas is increasingly common. Several developing countries across the globe have been undertaking housing projects to reduce urban slums (UN-Habitat, 2004), which offer some benefits but in many cases can imply relocation to lower-opportunity areas (Barnhardt et al., 2017). This is especially true in lower-income countries, given more complicated trade-offs and lower state capacities. Therefore, the implications of our results may not be limited to the specific experience of South Africa.

Bibliography

Abel, Martin, "Long-run effects of forced resettlement: evidence from Apartheid South Africa," *The Journal of Economic History*, 2019, *79* (4), 915–953.

Abramitzky, Ran, Leah Boustan, and Katherine Eriksson, "Do Immigrants Assimilate More Slowly today than in the past?," *American Economic Review: Insights*, 2020, *2* (1), 125–41.

_, Leah Platt Boustan, and Katherine Eriksson, "Europe's tired, poor, huddled masses: Self-selection and economic outcomes in the age of mass migration," *American Economic Review*, 2012, *102* (5), 1832–56.

Adhvaryu, Achyuta, Steven Bednar, Teresa Molina, Quynh Nguyen, and Anant Nyshadham, "When It Rains It Pours: The Long-Run Economic Impacts of Salt Iodization in the United States," *Review of Economics and Statistics*, 2020, *102* (2), 395–407.

Angrist, Joshua, Victor Lavy, and Analia Schlosser, "Multiple experiments for the causal link between the quantity and quality of children," *Journal of Labor Economics*, 2010, *28* (4), 773–824.

Barnhardt, Sharon, Erica Field, and Rohini Pande, "Moving to opportunity or isolation? network effects of a randomized housing lottery in urban india," *American Economic Journal: Applied Economics*, 2017, *9* (1), 1–32.

Bazzi, Samuel, Arya Gaduh, Alexander D Rothenberg, and Maisy Wong, "Skill transferability, migration, and development: Evidence from population resettlement in Indonesia," *American Economic Review*, 2016, *106* (9), 2658–98.

Becker, Sascha O and Andreas Ferrara, "Consequences of forced migration: A survey of recent findings," *Labour Economics*, 2019, *59*, 1–16.

_ and Luigi Pascali, "Religion, Division of Labor, and Conflict: Anti-Semitism in Germany over 600 Years," *American Economic Review*, 2019, *109* (5), 1764–1804.

Becker, Sascha O., Irena Grosfeld, Pauline Grosjean, Nico Voigtlander, and Ekaterina Zhuravskaya, "Forced Migration and Human Capital: Evidence from Post-WWII Population Transfers," *American Economic Review*, May 2020, *110* (5), 1430–63.

Black, Sandra E, Paul J Devereux, and Kjell G Salvanes, "The more the merrier? The effect of family size and birth order on children's education," *The Quarterly Journal of Economics*, 2005, *120* (2), 669–700.

_ , _ , and _ , "Small family, smart family? Family size and the IQ scores of young men," *Journal of Human Resources*, 2010, *45* (1), 33–58.

Bleakley, Hoyt, "Malaria eradication in the Americas: A retrospective analysis of childhood exposure," *American Economic Journal: Applied Economics*, 2010, *2* (2), 1–45.

Borjas, George J and Joan Monras, "The labour market consequences of refugee supply shocks," *Economic Policy*, 2017, *32* (91), 361–413.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller, "Bootstrap-based improvements for inference with clustered errors," *The Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Cantoni, Davide, Yuyu Chen, David Y Yang, Noam Yuchtman, and Y Jane Zhang, "Curriculum and ideology," *Journal of Political Economy*, 2017, *125* (2), 338–392.

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

Chyn, Eric, "Moved to opportunity: The long-run effects of public housing demolition on children," *American Economic Review*, 2018, *108* (10), 3028–56.

Clark, Nancy L, "Structured inequality: Historical realities of the post-apartheid economy," *Ufahamu: A Journal of African Studies*, 2014, *38* (1).

Dasgupta, Basab and Somik V Lall, "Assessing benefits of slum upgrading programs in second-best settings," in "Urban Land Markets," Springer, 2009, pp. 225–251.

Day, Jennifer and Robert Cervero, "Effects of residential relocation on household and commuting expenditures in Shanghai, China," *International journal of urban and regional research*, 2010, *34* (4), 762–788.

Deutscher, Nathan, "Place, Peers, and the Teenage Years: Long-Run Neighborhood Effects in Australia," *American Economic Journal: Applied Economics*, 2020, *12* (2), 220–49.

Dinkelman, Taryn, "Long-run Health Repercussions of Drought Shocks: Evidence from South African Homelands," *The Economic Journal*, 2017, *127* (604), 1906–1939.

Dippel, Christian, "Forced coexistence and economic development: evidence from Native American Reservations," *Econometrica*, 2014, *82* (6), 2131–2165.

Duflo, Esther, "Child health and household resources in South Africa: evidence from the old age pension program," *American Economic Review*, 2000, *90* (2), 393–398.

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

Evans, Ivan, *Bureaucracy and race: Native administration in South Africa*, Vol. 53, Univ of California Press, 1997.

Ferrara, Eliana La and Annamaria Milazzo, "Customary norms, inheritance, and human capital: evidence from a reform of the matrilineal system in Ghana," *American Economic Journal: Applied Economics*, 2017, *9* (4), 166–85.

Goodman-Bacon, Andrew, "The long-run effects of childhood insurance coverage: Medicaid implementation, adult health, and labor market outcomes," Technical Report, National Bureau of Economic Research 2016.

— , "Difference-in-differences with variation in treatment timing," Technical Report, National Bureau of Economic Research 2018.

Greenberg, S and H Giliomee, "Labour Bureaux and the African Reserves," *South African Labour Bulletin*, 1983, *8* (4), 37–50.

Haider, Steven and Gary Solon, "Life-cycle variation in the association between current and lifetime earnings," *American Economic Review*, 2006, *96* (4), 1308–1320.

Hanoman, Jacqueline, *Hunger and Poverty in South Africa: The Hidden Faces of Food Insecurity*, Routledge, 2017.

Heckman, James J, "The economics, technology, and neuroscience of human capability formation," *Proceedings of the national Academy of Sciences*, 2007, *104* (33), 13250–13255.

Hornbeck, Richard, "Dust Bowl Migrants: Identifying an Archetype," *Unpublished*, 2020.

Human Sciences Research Council, "Using the 2001 Census: Approaches to analysing data," Technical Report, Human Sciences Research Council and Statistics South Africa 2007.

Ito, Takahiro and Shinsuke Tanaka, "Abolishing user fees, fertility choice, and educational attainment," *Journal of Development Economics*, 2018, *130*, 33 – 44.

Lee, Jin Young and Gary Solon, "The Fragility of Estimated Effects of Unilateral Divorce Laws on Divorce Rates," Technical Report, National Bureau of Economic Research 2011.

Lemon, Anthony, "State control over the labor market in South Africa," *International Political Science Review*, 1984, *5* (2), 189–208.

Luwaya, Nolundi, "Land, status and security–a burden borne by women," *Agenda*, 2018, *32* (4), 103–110.

Mare, Gerry, *African population relocation in South Africa*, SA Institute of race relations, 1980.

Massey, Douglas and Nancy Denton, *American apartheid: Segregation and the making of the underclass*, Harvard University Press, 1993.

Morales, Juan S, "The impact of internal displacement on destination communities: Evidence from the Colombian conflict," *Journal of Development Economics*, 2018, *131*, 132–150.

Moultrie, Tom A and Ian M Timæus, "The South African fertility decline: Evidence from two censuses and a Demographic and Health Survey," *Population studies*, 2003, *57* (3), 265–283.

Nakamura, Emi, Jósef Sigurdsson, and Jón Steinsson, "The gift of moving: Intergenerational consequences of a mobility shock," Technical Report, National Bureau of Economic Research 2016.

O'Malley, P, "The Heart of Hope–South Africa's Transition from Apartheid to Democracy," *Nelson Mandela Foundation's Centre of Memory and Dialogue*, 2007.

Reed, Holly E, "Moving across boundaries: migration in South Africa, 1950–2000," *Demography*, 2013, *50* (1), 71–95.

Sarvimäki, Matti, Roope Uusitalo, and Markus Jäntti, "Habit formation and the misallocation of labor: evidence from forced migrations," *Available at SSRN 3361356*, 2019.

Sequeira, Sandra, Nathan Nunn, and Nancy Qian, "Immigrants and the Making of America," *The Review of Economic Studies*, 2020, *87* (1), 382–419.

Simkins, Charles Edward Wickens, *Four essays on the past, present & possible future of the distribution of the black population of South Africa*, Southern Africa Labour and Development Research Unit, University of Cape Town, 1983.

Skelcher, Bradley, "Apartheid and the removal of black spots from Lake Bhangazi in KwaZulu-Natal, South Africa," *Journal of Black Studies*, 2003, *33* (6), 761–783.

Some, Wawan, Wardah Hafidz, and Gabriela Sauter, "Renovation not relocation: the work of Paguyuban Warga Strenkali (PWS) in Indonesia," *Environment and urbanization*, 2009, *21* (2), 463–475.

Statistics South Africa, "South African Census 1980," Technical Report, Cape Town: DataFirst, https://www.datafirst.uct.ac.za/dataportal/index.php/catalog/252 1980.

—, "South African Census 1985," Technical Report, Cape Town: DataFirst, https://www.datafirst.uct.ac.za/dataportal/index.php/catalog/146 1985.

—, "South African Census 1991," Technical Report, Cape Town: DataFirst, https://www.datafirst.uct.ac.za/dataportal/index.php/catalog/253 1991.

—, "October Household Survey 1994," Technical Report, Cape Town: DataFirst, https://www.datafirst.uct.ac.za/dataportal/index.php/catalog/407 1994.

—, "October Household Survey 1995," Technical Report, Cape Town: DataFirst, https://www.datafirst.uct.ac.za/dataportal/index.php/catalog/410 1995.

—, "South African Census 1996," Technical Report, Cape Town: DataFirst, https://www.datafirst.uct.ac.za/dataportal/index.php/catalog/255 1996.

Surplus People Project, "Forced removals in South Africa : Volume 1-5 of the Surplus People Project report," Technical Report, Cape Town: Surplus People Project 1983.

Tanaka, Shinsuke, "Does abolishing user fees lead to improved health status? Evidence from post-Apartheid South Africa," *American economic Journal: economic policy*, 2014, *6* (3), 282–312.

UN-Habitat, "The challenge of slums: global report on human settlements 2003," *Management of Environmental Quality: An International Journal*, 2004, *15* (3), 337–338.

Viratkapan, Vichai and Ranjith Perera, "Slum relocation projects in Bangkok: what has contributed to their success or failure?," *Habitat international*, 2006, *30* (1), 157–174.

Figure 1: South African Homelands

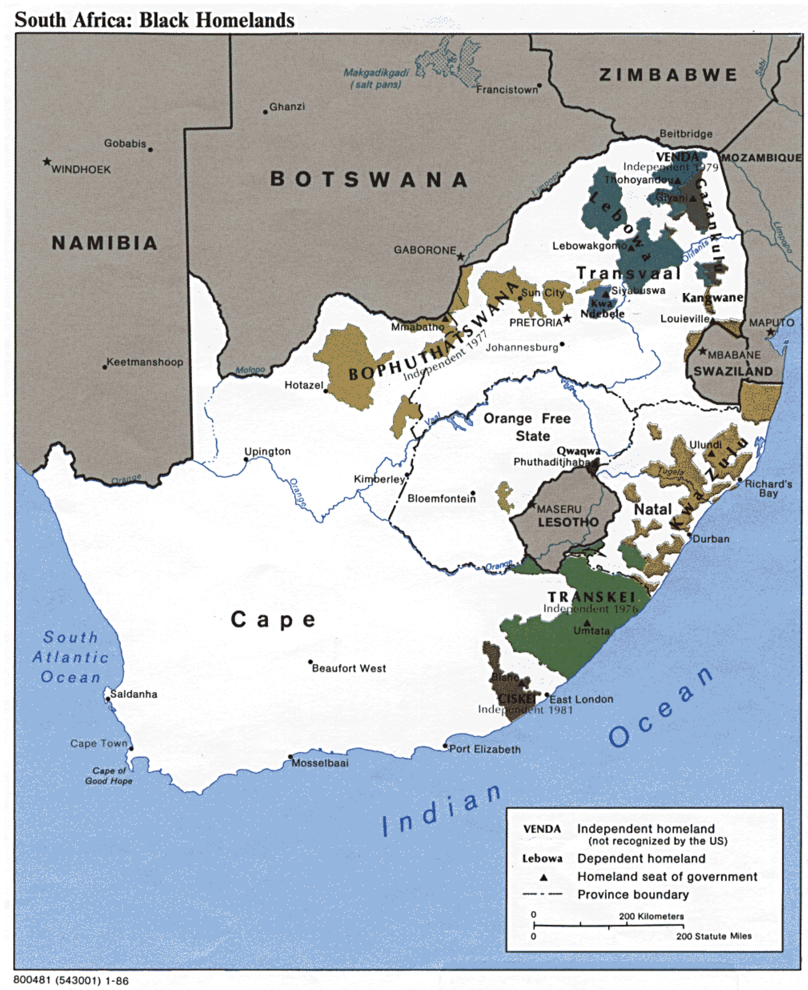

Sources: Political map of South Africa 1986, Perry-Castañeda Library Map Collection, http://www.lib.utexas.edu/maps/south_africa.html, accessed July 2020.

Figure 2: Estimated Effects on Probability of Living in a Homeland

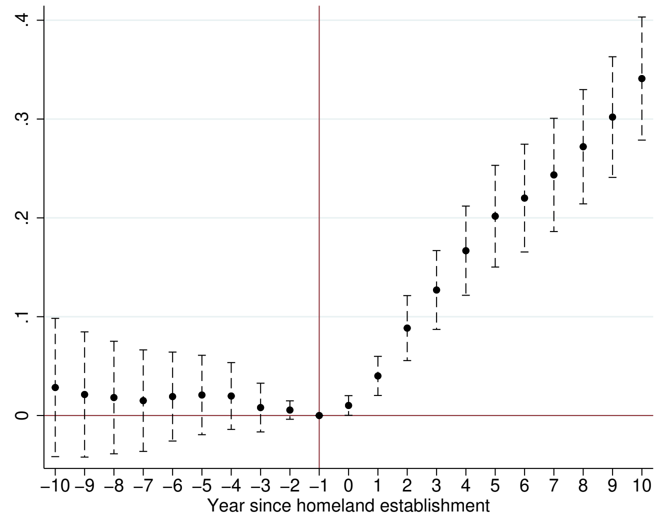

Notes: This figure plots β_z (and 95 confidence intervals) from estimating equation (1). Results based on the 1996 census. Using information on year of arrival, we define an event in individuals' residential history as a binary variable equal to one for the years in current homeland of residence. This generates a microdata panel at the individual-year level. The sample is limited to individuals whose previous place of residence is in a White area. To estimate this event study on a fully balanced sample of individuals, the sample is limited to individuals who are observed during both the entire pre- and post-policy periods. Controls include fixed effects for year and individual. Robust standard errors are clustered at the homeland × birth decade level.

Figure 3: Timing of Moves relative to Homeland Establishment

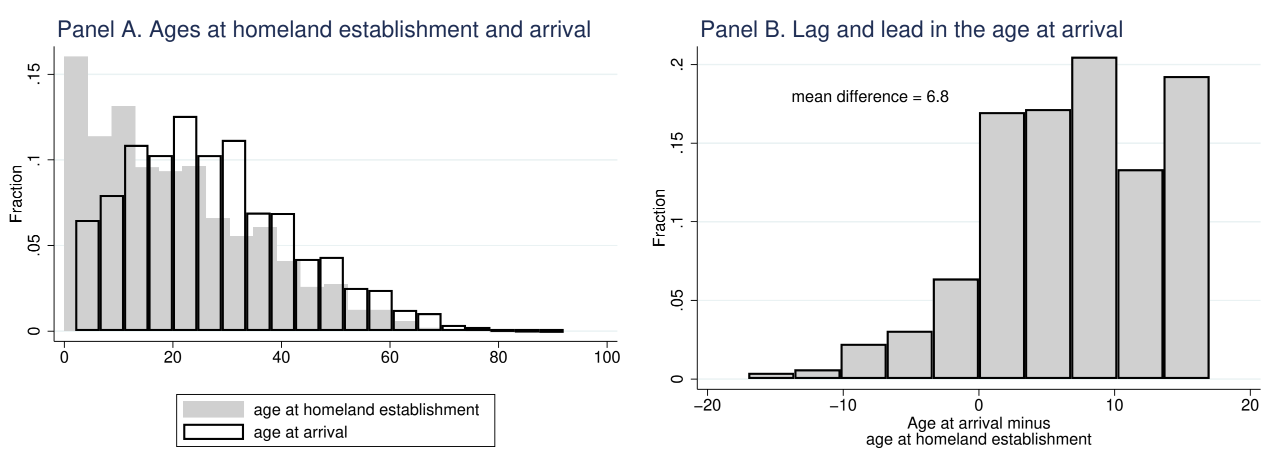

Notes: This figure compares an individual's age at arrival and homeland establishment using data from the 1996 Census.

Figure 4: Actual versus Simulated Exposure to Homelands

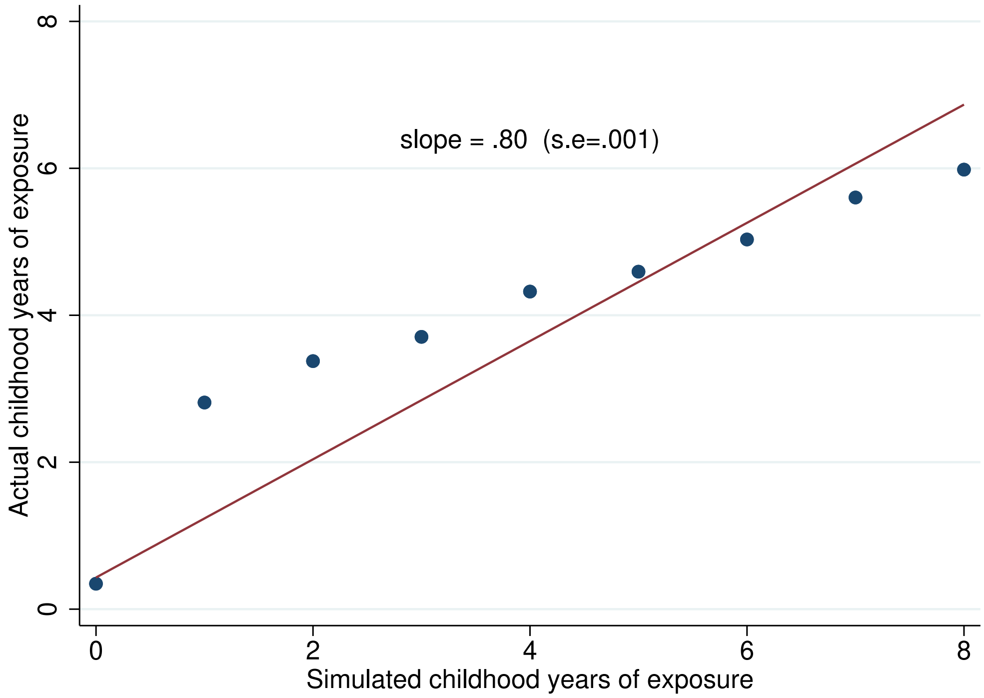

Notes: This figure presents a binned scatter plot of the relationship between individuals' actual and simulated exposure to the homelands using the 1996 Census. The simulated exposure measure is constructed as described in Section 4.1.3, considering the average lag in the age at arrival.

Figure 5: Estimated Effects on Years of Education

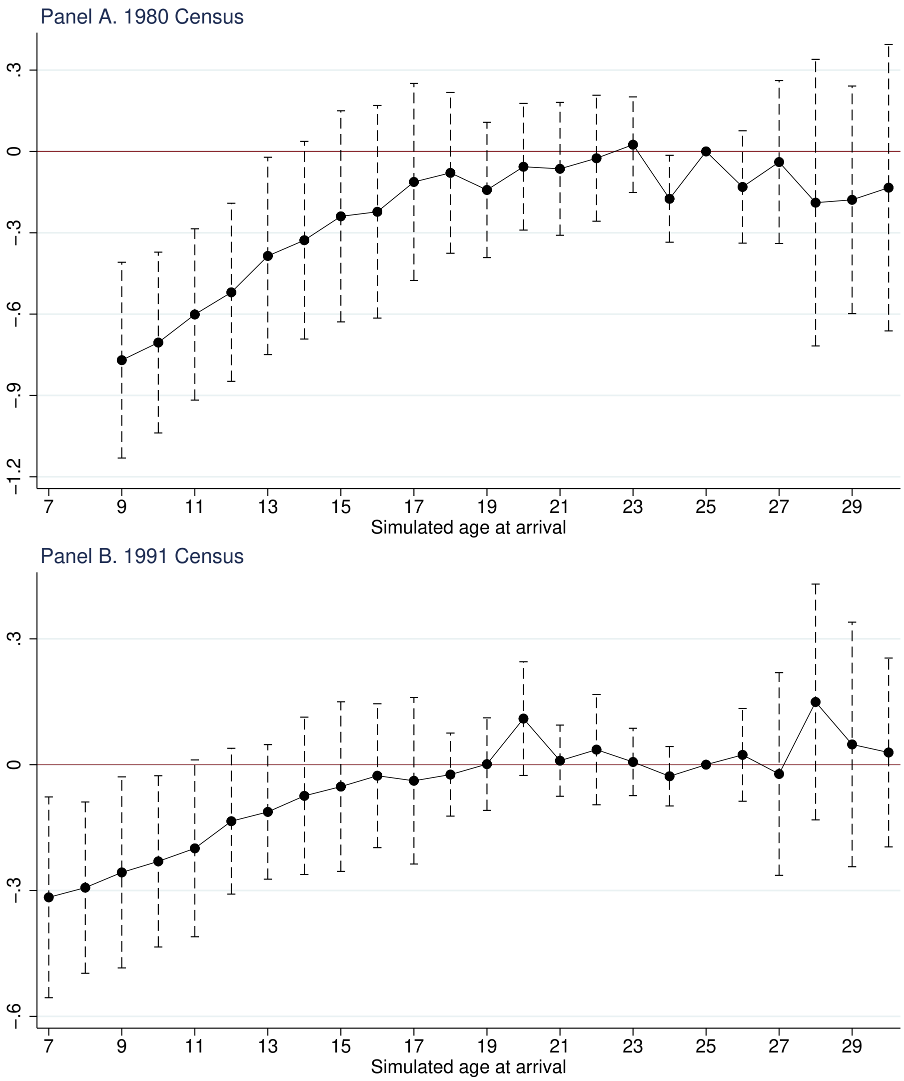

Notes: This figure plots β_τ (and 95 percent confidence intervals) from estimating the equation (5). Robust standard errors are clustered at the homeland-birth decade level.

Figure 6: Educational Attainment: Stayers versus Movers
(pre-policy cohorts)

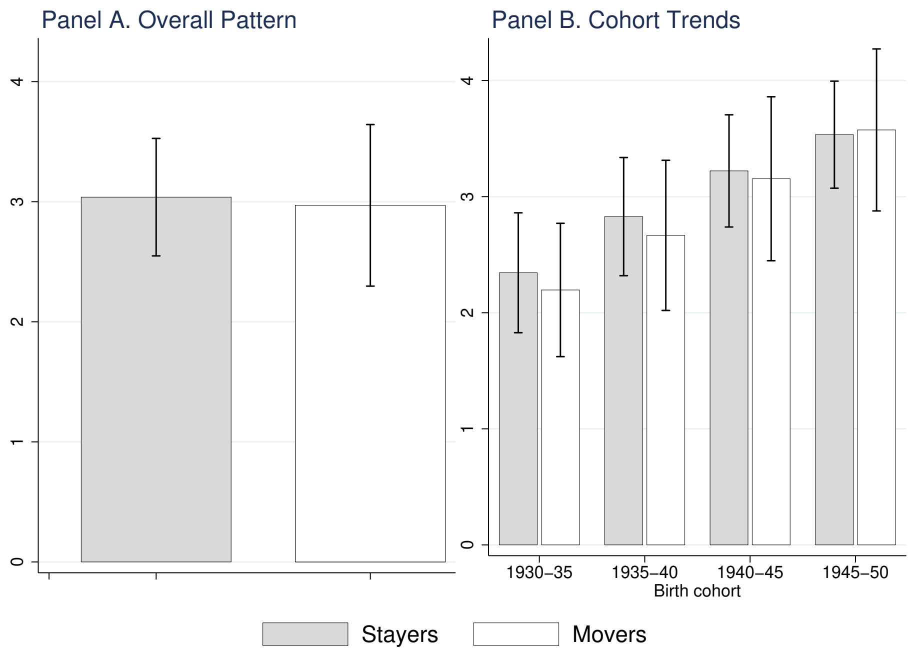

Notes: This figure shows the patterns of education for cohorts born before 1950, separately for individuals in White areas who did not move to homelands (stayers) and those whose who migrated to the homelands (movers). Sample based on the 1980 census.

Figure 7: Estimated Effects on Employment and Income

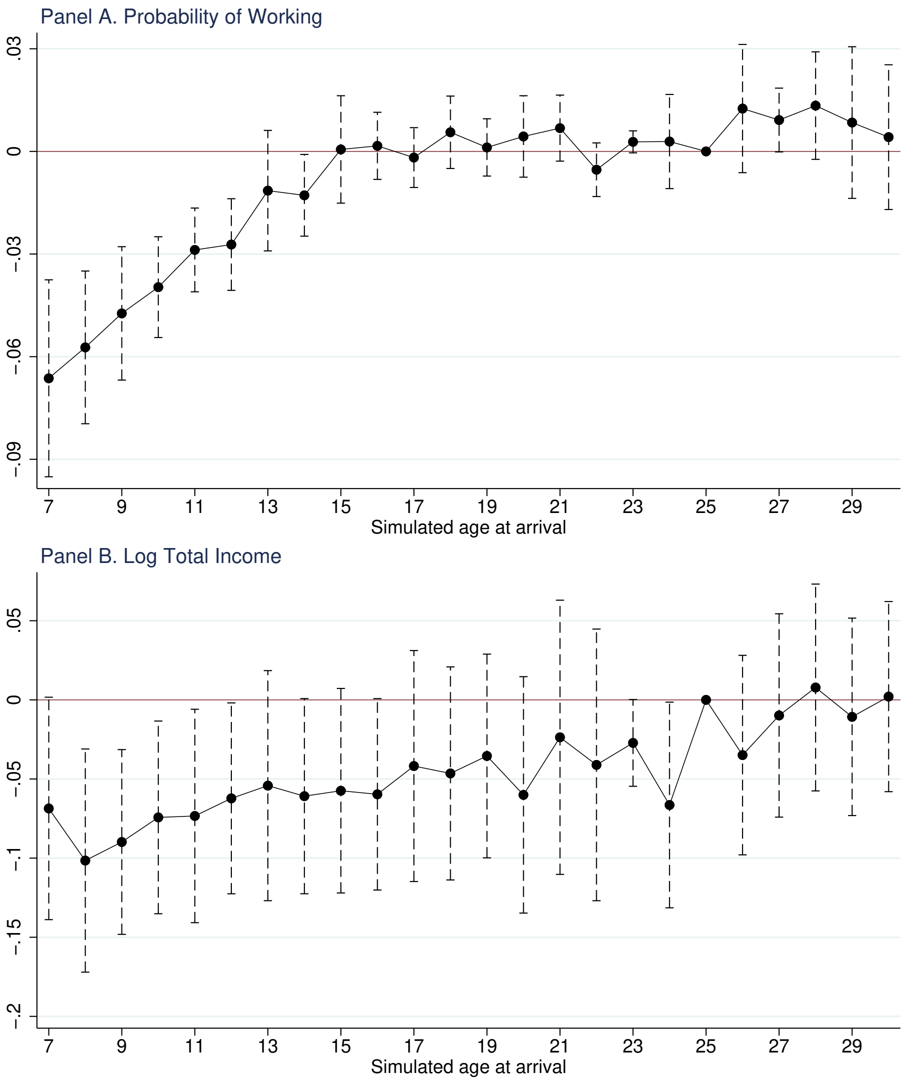

Notes: This figure plots β_τ (and 95 percent confidence intervals) from estimating the equation (5). Sample is based on the 1991 census. Robust standard errors are clustered at the homeland-birth decade level.

Figure 8: Cross-cohort Difference in Migrant Schooling by predicted Differences in schooling

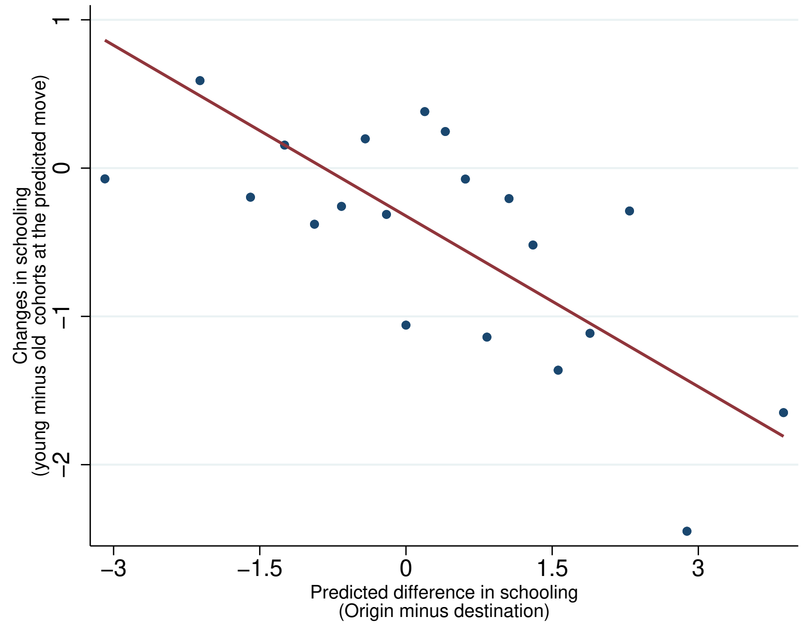

Notes: To construct this figure, we calculate the mean difference in schooling between migrants aged 9-12 and 28-30 at the predicted move within both the same Magisterial District of birth and residence. We then break down the predicted schooling gap into twenty equal-size bins and collapse the data to these cells. The scaled slope by one-standard deviation of the schooling gap is -0.21 (standard error=0.019)

Figure 9: Migrant Schooling and Differences in Origin-Destination Predicted Schooling

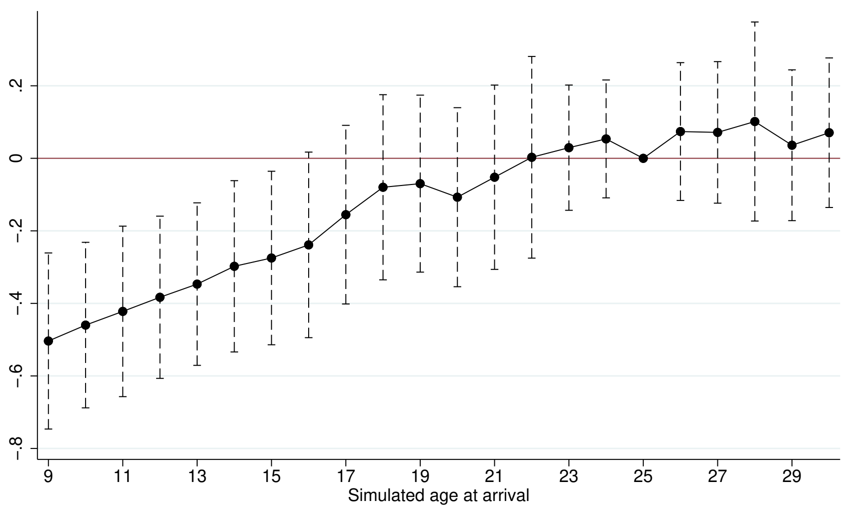

Notes: This figure plots coeficients and 95 percent confidence intervals from estimating a non-parametric version of equation (6) which includes interactions between origin-destination differences in predicted schooling and indicators for simulated age at arrival. Robust standard errors are clustered at the homeland-birth decade level.

Table 1: Estimated Effects on Schooling

	Dependent variable: years of education					
	(1)	(2)	(3)	(4)	(5)	(6)
	1980 Census				1985 Census	1991 Census
	Baseline sample	Baseline sample	Only immigrants from White areas	Only homeland natives	Baseline sample	Baseline sample
Simulated Exposure × Migrant	-0.116	-0.1102			-0.10	-0.043
	[0.0198]	[0.0159]			[0.0168]	[0.012]
Simulated Exposure			-0.1326	-0.0016		
			[0.0573]	[0.0384]		
Homeland × birth-year FE	✓	✓			✓	✓
Province-of-birth FE	✓	✓	✓	✓	✓	✓
Magisterial District FE	✓	✓	✓	✓	✓	✓
Demographic controls	✓	✓	✓	✓	✓	✓
Other individual controls		✓				
Homeland-specific linear pre-trends			✓	✓		
Observations	2,489,596	2,489,596	425,818	2,063,778	2,346,208	2,139,840

Notes: Demographic controls include gender and ethnicity. Ethnicity is not available in the 1991 census. Other individual controls include marital status, religion, and first home language. Homeland-specific linear pre-trends represent interactions between a linear cohort trend and homeland dummies, which are set to zero for simulated ages at arrival below 18. Robust standard errors (reported in brackets) are clustered at the homeland-birth decade level.

Table 2: Falsification

	Dependent variable: years of education	
	Baseline sample Treatment: migrants from white areas	Placebo sample Treatment: migrant from other homelands
	(1)	(2)
Simulated Exposure × Migrant	-0.116	-0.010
	[0.0198]	[0.0099]
Homeland × birth-year FE	✓	✓
Province-of-birth FE	✓	✓
Magisterial district FE	✓	✓
Observations	2,489,596	2,063,778

Notes: Estimates based on the 1980 census. Homeland Migrant is an indicator for homeland-homeland migrants. The sample in column (2) excludes migrants from White areas. All regressions include controls for gender and ethnicity. Robust standard errors (reported in brackets) are clustered at the homeland-birth decade level.

Table 3: Estimated Effects on Employment and Income

	Dependent variable:			
	(1)	(2)	(3)	(4)
		Log total income		
			Accounting for sample selection	
	Probability of working	baseline estimate	Lower bound	Upper bound
Simulated Exposure × Migrant	-0.009	-0.008	-0.003	-0.012
	[0.001]	[0.002]	[0.002]	[0.002]
Homeland × birth-year FE	✓	✓	✓	✓
Province-of-birth FE	✓	✓	✓	✓
Magisterial district FE	✓	✓	✓	✓
Observations	2,179,334	631,070	632,742	632,742

Notes: Estimates based on the 1991 census. All regressions include an indicator for gender. Lower and upper bounds in columns (3) and (4) are computed by imputing information for observations who are missing as result of the homeland policy under best and worst-scenarios. See Appendix G.9 for details. Robust standard errors (reported in brackets) are clustered at the homeland-birth decade level.

Table 4: Migrant Schooling by Origin-Destination Differences in Predicted Schooling

	Basic				Gender-specific effects
	(1)	(2)	(3)	(4) (Locals as control)	(5)
	Origin FE, Birth-year FE	Destination FE, Origin FE, Birth-year FE	Homeland × Birth-year FE, Destination FE, Origin FE	Homeland × Birth-year FE, Destination FE, Origin FE	Homeland × Birth-year × Gender FE, Destination × Gender FE, Origin × Gender FE
Simulated Exposure × Predicted Diff. in Schooling	-0.0818	-0.0976	-0.0809	-0.0723	
	[0.0169]	[0.0173]	[0.0133]	[0.0143]	
Simulated Exposure × Predicted Diff. in Schooling (own gender exposure: β_{own})					-0.0604
					[0.0195]
Simulated Exposure × Predicted Diff. in Schooling (placebo: other-gender exposure: β_{other})					-0.0195
					[0.0218]
Observations	425,818	425,818	425,818	2,451,695	425,818

Notes: Estimates based on the 1980 census and alternative versions of specification (6). The predicted difference in schooling is defined as the origin minus destination expected level of schooling. We normalize the predicted schooling gap term by dividing it by the standard deviation (or equivalently 90 percent of the interquartile range). All regressions include controls for gender and ethnicity. Column (4) in addition includes migrant × birth-year fixed effects. Robust standard errors (reported in brackets) are clustered at the homeland-birth decade level.