Slums: Poverty Traps or Stepping-Stones?*

Diogo Britto † Lucas Conwell ‡ Alexandre Fonseca §
Breno Sampaio ¶ Gabriel Ulyssea || Lucas Warwar **

First Version: January 2025
This Version: March 2025

Abstract

Are slums stepping-stones to better lives or poverty traps? To address this crucial question, we leverage unique Brazilian administrative data sources to link children's long run outcomes with those of their parents and examine intergenerational mobility at the neighborhood level within cities in most of urban Brazil. Slums, on average, offer no worse intergenerational mobility than formal neighborhoods due to offsetting effects of crime rates versus school quality and formal employment opportunities. We develop a spatial model with overlapping generations which captures the trade-offs parents face when choosing where to live and raise their children. The model integrates intergenerational human capital investment within an urban structure, emphasizing housing costs, access to educational opportunities, exposure to crime, and proximity to employment as key determinants of location decisions. We use the estimated model to quantify the relative importance of the different determinants of intergenerational mobility in slums versus non-slums in general equilibrium and thus shed light on the traits which make slums "traps" or "stepping-stones."

*We thank Varsha Vejalla for outstanding research assistance, and Clement Imbert for useful comments. All errors are ours. Ulyssea gratefully acknowledges the funding received from the European Research Council (ERC) Starting Grant under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 852744).

†University of Milan
‡University College London. l.conwell@ucl.ac.uk
§UFPE
¶UFPE
||University College London. g.ulyssea@ucl.ac.uk
**Stanford University

I. Introduction

By 2024, around 1.1 billion people or 25% of the world's urban population resided in slums or informal settlements (UN 2024). The bulk of this contingent is concentrated in developing countries, where rapid urbanization frequently outpaces infrastructure development, leading a substantial proportion of residents to live in substandard conditions (J Vernon Henderson et al. 2016; Michaels et al. 2021a). With urban populations in these regions projected to expand 64.7% by 2050 (UNCTAD 2021), these challenges are expected to intensify significantly in the coming decades, with 2 billion more people expected to live in slums or slum-like conditions in cities in the next 30 years (UN 2024).

Despite the sheer magnitude of this phenomenon, we still have limited evidence on the characteristics of slum dwellers and their livelihoods. Crucially, a key open question refers to slums' long-term effects on the economic outcomes of their residents. On the one hand, slums may function as stepping-stones, where affordable housing enables residents to establish an urban foothold, invest in their children's education, and eventually access the full economic benefits of urban centers (Glaeser 2011). On the other hand, the adverse living conditions in slums—which often include inadequate infrastructure, high crime, and environmental hazards—may significantly impair residents' economic prospects and their children's educational achievements, perpetuating cycles of poverty (Marx et al. 2013). The existing literature has yet to conclusively distinguish between these competing views.

This paper addresses this critical question: are slums stepping-stones to better lives or poverty traps? To do so, we leverage unique administrative data sources from Brazil covering most of the urban population in the country. We construct a large, representative sample of 7.3 million individuals born between 1983 and 1990 for whom we can link their long-run economic outcomes to those of their parents. We estimate different measures of intergenerational mobility (IGM) across 77,000 census blocks (a more granular level than U.S. census block groups), and document IGM in slums *versus* non-slums within and across the 19 largest metropolitan areas in Brazil. We establish new facts about IGM in slums *versus* city proper, and how they are shaped by factors such as access to schools, formal employment opportunities, housing quality, and crime rates. Based on these facts, we develop a spatial model with overlapping generations which captures the trade-offs parents face when choosing where to live and raise their children. The model integrates intergenerational human capital investment within an urban structure, emphasizing housing costs, access to educational opportunities, exposure to crime, and proximity to employment as key determinants of location decisions. We use the estimated

model to quantify the relative importance of these different determinants for slum dwellers' welfare and social mobility, as well as to quantify the traits that make slums "traps" or "stepping-stones" in general equilibrium.

In the first part of the paper, we build on Britto et al. (2022) and leverage different administrative data sources that include individual-level information from population-wide administrative registries and large-scale surveys, enriched by unique tax and residential address data provided by the Brazilian tax authority. We link parents and children using the tax registry, welfare registry, and school census. We then measure formal labor income from individual tax returns as well as RAIS matched employer-employee data. We observe informal income for a subset of individuals in large-scale surveys conducted by the Brazilian Institute of Geography and Statistics and impute informal income for all individuals using random forest and a rich set of characteristics available in administrative data. Finally, we geocode tax registry addresses to link children with the census tract where they grew up. As our key measure of IGM, we compute the expected adult income rank of children with parents at the 25th income percentile who grew up in a given census block. In more densely populated areas, block and neighborhood boundaries tend to coincide, and we use both terms interchangeably throughout the paper.

We start by characterizing slum dwellers and basic slum characteristics, contrasting them with those observed in formal neighborhoods with average income per capita below the city's median (the "Poor Formal"), and those with average income per capita above the city's median ("Rich Formal"). For that, we focus on the 77,000 census blocks for which we can compute IGM measures, which corresponds to 86% of all census blocks in Brazil. Unsurprisingly, we find that residents of poor neighborhoods (formal and informal) have substantially worse access to basic infrastructure compared to households "rich formal" neighborhoods. However, if anything, slum households have slightly better coverage than those in poor but formal neighborhoods. This difference is not entirely driven by variation across cities, as half of it remain even after controlling for metropolitan area fixed effects. Similarly, we show that residents of slums and poor formal neighborhoods have equally precarious floor and wall materials. However, slums do differ along key dimensions relative to poor formal neighborhoods: on average, they offer greater access to formal jobs and to higher quality education, are less distance-segregated (Harari 2024), but show higher crime (measured as victimization rates). Overall, slum residents have on average 14% lower monthly income per capita compared to households in poor formal neighborhoods, as well as lower average schooling (by around half of a standard deviation), with these gaps being slightly larger within cities.

Turning to the main question of the paper – whether slums are poverty traps or stepping-stones – we approach it by establishing four key facts. First, we show that slums are not poverty traps *on average*: we find a very small difference on average IGM between slums and formal neighborhoods, and no difference between slums and poor but formal neighborhoods. These results are confirmed by different specifications at the neighborhood-level, but also using our individual-level data to estimate rank-rank regressions, and a "mover design" specification.

Second, the absence of a "slum effect" on IGM hides important offsetting effects of the differences in characteristics between slums and formal neighborhoods. We do a weighted Oaxaca-Blinder decomposition of the zero "slum effect", and find that slums' higher victimization rates (assault hospitalizations) and lower distance segregation decrease their intergenerational mobility relative to poor formal neighborhoods. These negative effects are offset by the fact that slums have lower informal employment, better access to higher school quality, and better commuter market access. Thus, better labor market access and educational opportunities offset higher crime in Brazilian slums, such that they on average have no worse mobility than formal neighborhoods.

Third, there is substantial heterogeneity in mobility across slums. Some slums are characterized by very low to zero upward social mobility, while some slums have high upward mobility across generations in terms of income rank of the children when adults, with a 90-10 ratio across slums of 1.9. Thus, some slums are poverty traps, while others are stepping-stones. Furthermore, we show that the slums that have the lowest IGM are also those associated with higher probability of their children living in slums as adults. This reinforces the evidence that these slums indeed constitute intergenerational poverty traps that tend to reproduce poverty cycles across generations.

Fourth, most of the mobility gap across slums can be explained by differences in crime, access to schools, and informality levels. We show that the strongest predictors of which slums are poverty traps – those that lie in the bottom tercile of the IGM distribution – are these three factors. We then conduct another Oaxaca-Blinder decomposition to identify the main determinants of the 19-percentile difference in mean mobility between the "stepping stones" and "poverty traps". We decompose this IGM gap among slums using the same set of covariates as in the "slum effect", and find that school quality, informal employment, and assault hospitalizations account for 60% of the gap between the low- and high-mobility slums.

Based on these facts, we develop a novel model of internal city structure with overlapping

generations. At the heart of the model, parents choose a home neighborhood, draw a formal or informal job offer and then choose to invest in their children's human capital. Investments vary across neighborhoods based on access to formal labor markets, fixed neighborhood characteristics, e.g. crime, and housing rents, which determine the net income available for investment. The risk of losing an informal job further constrains investment. In contrast to standard spatial models, our model generates rich within-neighborhood heterogeneity in human capital (investments) and highlights feedback between location choice, the housing market, and human capital.

We calibrate the model using our microdata and then use it to compare resident welfare and social mobility in slums versus formal neighborhoods. While slums, on average, offer no worse mobility for local children, the model reveals that slum-dwellers suffer from substantially lower amenities and housing supply. Residents thus appear to choose to live in slums due to idiosyncratic tastes, such as social ties or heterogeneity in preferences for child schooling. Finally, we decompose the poverty-trap versus stepping-stone difference in IGM into the relative contributions of commuter market access, frictions to finding formal jobs, crime risk, and education quality. In particular, we equalize one factor at a time between slums and formal neighborhoods and compute the resulting differential. We find that, even taking into account equilibrium changes in parental investment, rents, and location choices, equalization of school quality offers the greatest potential to turn poverty traps into stepping stones.

Related Literature

The existing literature investigating the relationship between slums and mobility has been severely limited by data constraints, providing mixed evidence coming from small sample surveys. On the one hand, Banerjee et al. 2012 and Wong (2019) provide suggestive evidence of educational gains across generations, and strong presence of migrants and substantial churn in slums in Jakarta and Delhi, respectively. On the other hand, Krishna (2013) and Marx et al. (2013) provide evidence more consistent with the "poverty trap view" of slums. More recently, Rojas-Ampuero and Carrera (2024) study the effects of slum clearance programs in Chile on the long-run outcomes of displaced children. In our context of Brazil, Barza et al. (2024) find larger urban wage premium in the south, potentially due to lower skill segregation in the labor market. In contrast to these papers, we directly measure slum residence and quantify inter-generational income mobility at a granular level and for the near-universe of urban Brazil.

A large literature studies the determinants of slum formation (J. Vernon Henderson et

al. 2021; Gonzalez-Navarro and Undurraga 2023), re-development (Gechter and Tsivanidis 2023; Harari and Wong 2024), and public goods provision (Michaels et al. 2021b; Harari 2024). In the context of Brazil, Cavalcanti et al. (2019) develop a macro model of housing choice and slum formation, while Ferreira et al. (2024) develop a model of structural change and slum formation to investigate whether slums are intergenerational stepping stones. Our quantitative spatial model (e.g. Ahlfeldt et al. 2015) highlights how micro-geographic variation in education quality, job access, crime, and housing supply translate into families' neighborhood choices and into parental human capital investment, and therefore social mobility.

II. Data

We combine multiple comprehensive data sources to conduct our analysis: population-wide administrative registries, granular Decennial Population Census data, national household surveys, and proprietary tax and residential address records from the Brazilian tax authority. This unique data combination enables detailed characterization of slum residents, their housing conditions, and direct comparisons with formal city neighborhoods. Critically, these integrated sources allow us to link children's long-term outcomes with their parents', facilitating fine-grained analysis of intergenerational mobility (IGM) at the census block level. As mentioned before, we use census block and neighborhood interchangeably throughout the paper.

In this Section, we first provide a brief description of the three main steps in constructing our individual-level dataset: (i) recovering family links, (ii) measuring income, and (iii) georeferencing households and detecting those living in slums. We then describe our sources of neighborhood-level and expenditure data. Steps (i) and (ii) constitute the data constructed by Britto et al. (2022). We start from this data and extend it with rich residential data in step (iii).

II.1 Family links

We link children to their mothers and fathers separately and then take the intersection of both linked datasets as our main sample. We start with an initial linking step and then proceed by rounds, additionally matching children to parents using other sources.

For mothers, the first step uses the Brazilian Person Registry (CPF), which is maintained by the Brazilian Tax Authority (RFB) and covers the entire population. All individuals are

identified by their tax ID, full name, and mother's full name. If the mother can be uniquely identified by her name, we link the child's identifier to her mother's based on the mother's name.[1]

For fathers, we first use the Welfare Registry (*CadÚnico*) from the Ministry of Social Development (MDS), which covers families eligible for social programs, around two-thirds of the Brazilian population. We implement the same procedure as before, linking children to their fathers conditional on the father having a unique name in the country. We complement this with information from the Brazilian Tax Authority on the dependent claims in individual tax returns. Parents report children aged 0-24 for the purpose of tax deductions, in which case we can directly link them to each other through the unique person codes.

Around 40% of the observations in our sample are obtained in this initial step. We then perform additional links by using information on (i) names and residential addresses from the Person Registry, (ii) names, state of residence, and household composition from the Welfare Registry, and (iii) names from the 2014 School Census. Appendix XX provides additional details and reports matching rates across steps. Our main sample contains 7.3 million individuals born between 1983 and 1990 linked to their parents. Table A.1 shows that the linking procedure yields a fairly representative sample of the population in terms of several individual characteristics.

II.2 Income

Our main outcome for intergenerational mobility is individuals' total income, defined as the sum of formal and informal income before tax and transfers. We define children's income at the individual level and parental income as the sum of the father's and mother's total incomes. We compute each component of individuals' income separately and then aggregate over sources (formal and informal) and individuals (in the case of parental income). In our main analysis, we measure the average income of children when they are 25-31 years old and the average income of their parents at the time when children were 3-18 years old.

FORMAL INCOME. There are two types of formal income: labor and non-labor. Our preferred source of individuals' formal income is tax returns – which include both labor

[1]The share of individuals with a unique name in the country is large because Brazilians typically carry one or more surnames from both their parents. These individuals are easily identified in the person registry, which includes names and individual identifiers. **britto2021** show that individuals with unique names do not strongly differ from the overall population along several characteristics.

TABLE 1. DESCRIPTIVE STATISTICS OF MAIN SAMPLE

Variable		Raw		Weighted	
	Population (1)	Sample (2)	Std. Diff. (3)	Sample (4)	Std. Diff. (5)
Non-white	0.58	0.59	0.03	0.58	0.00
Female	0.49	0.50	0.00	0.49	0.00
Formal Employment	0.79	0.81	0.02	0.78	0.00
Social Registry	0.54	0.78	0.36	0.74	0.18
Unique Name	0.57	0.60	0.06	0.57	0.01
State capital	0.26	0.24	0.09	0.26	0.01
Less than Primary	0.19	0.22	0.13	0.19	0.02
Primary	0.18	0.21	0.15	0.18	0.02
High School	0.46	0.45	0.02	0.46	0.00
College	0.09	0.09	0.02	0.09	0.00
North	0.09	0.09	0.04	0.09	0.01
Northeast	0.28	0.31	0.12	0.28	0.02
Southeast	0.42	0.38	0.09	0.42	0.01
South	0.14	0.14	0.03	0.14	0.00
Center-West	0.08	0.08	0.01	0.08	0.00

Notes: The table compares the average characteristics of our main sample of children born in 1983-1990 (omitting missing values) with the average characteristics of the same cohorts in the general population. The means for each variable are presented (columns 1-2), along with the standardized difference (column 3), the mean in the main sample after re-weighting observations to match the first and second moments of population characteristics (**hainmueller_2012**) (column 4), and the standardized difference between the samples in columns 1 and 4 (column 5). All variables are recorded as dummy indicators.

and non-labor components. However, tax data are not always available for two reasons: only a third of Brazilians file taxes each year, and tax data are available from 2006 onwards, limiting our ability to measure parental income.

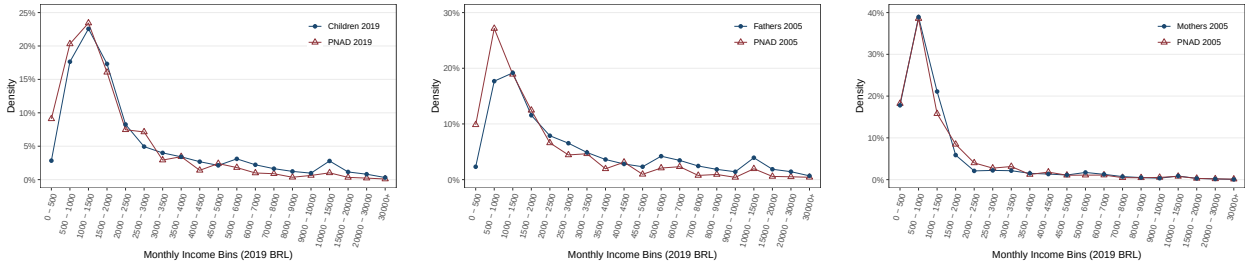
If tax records are unavailable for a given individual in a given year, we track the labor and non-labor components separately. Formal labor income is directly available from administrative employment data covering the population of formal jobs for the 1985-2019 period (RAIS dataset from the Ministry of Labor). Formal non-labor income – e.g., dividends, rents, interests, and capital gains – is not available in administrative registries other than tax data. We thus follow an imputation procedure using survey data. The procedure is the same one used to input informal income, which we describe next.

INFORMAL INCOME. We measure informal income using individual-level data from large-scale surveys implemented by the Brazilian Institute of Geography and Statistics (IBGE) covering 1991–2019.[2] We impute informal income based on a rich array of individual

[2]Specifically, we use PNAD, a cross-sectional household survey covering about 400,000 individuals per

FIGURE 1. INCOME DISTRIBUTION: DATA *vs.* PNAD

(A) CHILDREN (B) FATHERS (C) MOTHERS



Notes: The graphs compare the income distribution in our main sample (blue dots) and in the PNAD survey (red triangles), separately for children, fathers, and mothers. The PNAD sample includes "children" born in 1988-1990 that were interviewed in the wave 2019, and "parents" of children born in 1988-1990 that were interviewed in the wave 2005. PNAD sample weights are used to compute the income distribution in PNAD data.

characteristics available in both administrative registries and survey data fed to a Random Forest (RF) algorithm (Athey et al. 2019).

Using the survey data, we train a separate RF to predict informal income in each year from 1991 to 2019. The vector of predictors includes state of residence (27), a dummy for the metropolitan region, gender, age, race (white vs. non-white), education dummies (4), and occupation category (formal worker, formally self-employed, and firm owner, with informal workers as a residual category). After training the model, we predict informal income for all individuals in our sample, including formal workers and owners who may earn part of their total income in the informal sector. We repeat the same process for estimating formal non-labor income, which is necessary when tax data are not available.

Britto et al. (2022) estimate that the RF model predicts income ranks based on a single year with a fairly high R-squared of .57, which helps mitigate measurement error issues. As a comparison, using a fully saturated OLS leads to a significantly smaller R-squared of .29, likely due to overfitting issues.[3] They also provide evidence that averaging out income over multiple years further increases the precision of our predictions. Figure 1 shows that the distribution of total income in our sample matches the distribution of total income in the PNAD survey, apart from some (small) differences at the bottom of the father's income distribution.

year for the 1992-1999 and 2001-2019 periods, and Population Census surveys covering 10% of the population in 1991, 2000, and 2010.

[3]R-squared statistics are based on out-of-sample predictions using a random subsample of the survey data – not used for training the prediction models.

II.3 Geocoding

For every individual (parents and children) in our sample, we obtain their full residential address history from the Brazilian Tax Authority. As in Britto et al. 2022, we define children's location in adulthood as their address in 2019 and their location when growing up as their parents' address in 2000.[4] After a cleaning and standardization process of the addresses, we geolocalize each unique location to a latitude-longitude point using commercial mapping services.

We then use IBGE's shapefiles to map each location to a census tract, which categories allow us to identify formal *vs.* slum census tracts.[5]

II.4 Other Data Sources

IBGE POPULATION CENSUS. use information on households' access to piped water, bathrooms, sewer, and garbage collection available at the block level. We also obtain the informal share of employment, in which we include informal wage workers as well as the self-employed, and average rents. Finally, we calculate the Harari (2024) measure of distance segregation, namely the average distance from a neighborhood to wealthy neighborhoods in the top quartile of the metro area by income per capita. All variables are from the 2000 Census.

SUS HOSPITALIZATIONS We use individual-level data of hospitalizations from the public universal health care system (SUS) matched with the residential addresses data to obtain block-level measures of assault hospitalizations as a proxy for violent crime in the year 2000.[6]

RAIS FORMAL EMPLOYMENT We use geocoded establishment-level data on formal employment to obtain block level job counts, which we then use to compute commuter market access in 2002. Specifically, for a given tract i we take the inverse distance-weighted sum of job counts across blocks j within the same metro area, with weights given by $1/\exp\left(k \cdot \text{distance}_{ij}\right)$. We set k equal to the product of the workplace choice elasticity and disutility of travel time from Tsivanidis (2023), namely, in his notation, $\theta\kappa = .0374$, and use straight-line distance.

IDEB EDUCATION INDEX We use school-level data from the Basic Education Development

[4] We first consider the residential address linked to the father, and if it is unavailable, that of the mother. We set the year 2000 as a reference point due to the age of our cohorts and data quality limitations.

[5] Slums are classified as "aglomerados subnormais" by IBGE and have a specific census tract category.

[6] Details on matching, IDC codes for assault, etc.

Index (IDEB), a measure computed by the Ministry of Education using information on approval rates and performance in standardized tests. We combine this with schools' geolocalization data from the Institute of Applied Economic Research (IPEA), and create block level measures of school quality in 2007.

IBGE HOUSEHOLD BUDGET SURVEY (POF) We use microdata on households' income sources and expenditures from the 2017 edition of POF, a detailed survey conducted by IBGE with the purpose, among other goals, of tracking households' consumption basket. In particular, we use POF. to obtain expenditure measures on education and income derived from government transfers.

III. SLUMS: A CHARACTERIZATION

We now leverage the breadth of our combined datasets to paint a picture of slum dwellers and their livelihoods in urban Brazil.

Residents We start by contrasting basic characteristics of slum dwellers with those living in formal neighborhoods with average income per capita below the city's median – the "Poor Formal" – and with those of residents of formal and more affluent neighborhoods with average income per capita above the city's median – the "Rich Formal".

As expected, Table 2 shows that residents of poor neighborhoods have substantially worse access to basic infrastructure – which we proxy by coverage by sewer system – compared to households in more affluent areas across the country. However, if anything slum households have slightly better coverage than those in poor but formal neighborhoods. This difference is not entirely driven by variation across cities, as half of it remain even after controlling for metropolitan area fixed effects (results no reported). Regarding the materials used in their houses, we rely on CadUnico data to measure those, which means we are comparing low SES individuals across areas. With this caveat in mind, Table 2 shows that residents of slums and poor formal neighborhoods have equally precarious floor and wall materials, while low SES households in rich formal neighborhoods fair substantially better.

However, slum residents have on average 14% lower monthly income per capita compared to households in poor formal neighborhoods, as well as lower average schooling (by around half of a standard deviation). Interestingly, this gap is magnified when comparing slums and poor formal neighborhoods within the same city. We find no meaningful differences in the likelihood of being informally employed.

TABLE 2. CHARACTERISTICS OF RESIDENTS

	Slum		Poor Formal		Rich Formal	
	Mean	SD	Mean	SD	Mean	SD
Share w/access to sewer	0.475	0.408	0.409	0.399	0.777	0.339
Monthly income p.c.	333.7	142.6	381.8	115.6	1509.7	1246.0
Average years of schooling	4.6	0.9	5.2	1.1	8.3	2.1
Share informally employed	0.457	0.225	0.448	0.158	0.427	0.150
House materials						
Floor: cement or dirt	0.421	0.246	0.425	0.233	0.261	0.225
Wall: uncoated brickwork	0.200	0.158	0.189	0.143	0.118	0.142
	7,054		30,907		39,006	

Notes: Share with access to sewer, monthly income per capita, avergae years of schooling, and share informally employed are computed from the Demographic Census. House materials are computed from CadUnico data. "Poor Formal" corresponds to formal neighborhoods below the city's median income, while "Rich Formal" are those above the median. Informal share of employment includes both informal wage workers and the self-employed.

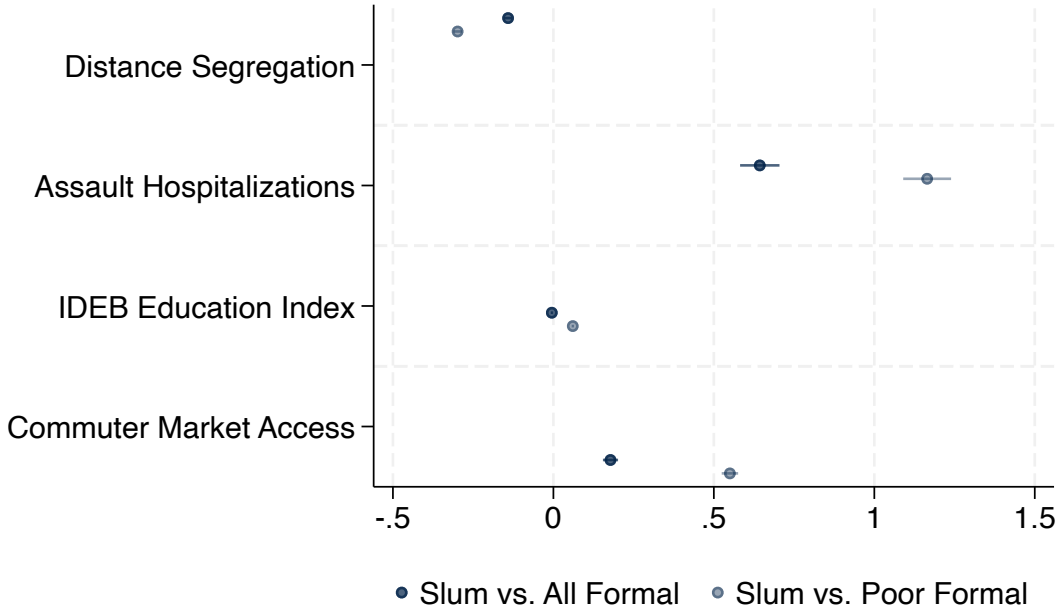
Neighborhoods Turning to neighborhood characteristics, we leverage the richness and granularity of our data to construct four key variables. The first is the *distance segregation measure* proposed by Harari (2024), which we construct at the block level. The greater the value of this measure the more segregated a given block is within the city.

The second is a *block-level measure of victimization*. We use the hospitalizations microdata from DataSus (see Section II) to compute the number of hospitalizations by assault per 100 thousand residents, which is arguably an extreme measure of victimization (there is no granular crime data for the whole of Brazil). We locate a given hospitalization entry by the block of residence of the victim. For a given block, we construct a weighted average that includes the victimization rate across neighboring blocks, where the weights are given by the inverse of distance between the blocks' centroids.

The third is a *block-level measure of access to education*, which is the IDEB index discussed in Section II. As in the victimization measure, we construct a weighted average using neighboring blocks indexes weighted by the inverse distance between centroids. The fourth is the block-level *Commuter Market Access*, which measures access to formal jobs as discussed in Section II.

Figure 2 shows striking patterns. First, slums are substantially less distance segregated than poor formal neighborhoods, consistent with previous results in the literature (e.g.).

FIGURE 2. NEIGHBORHOOD CHARACTERISTICS (PERCENTAGE DIFFERENCES)

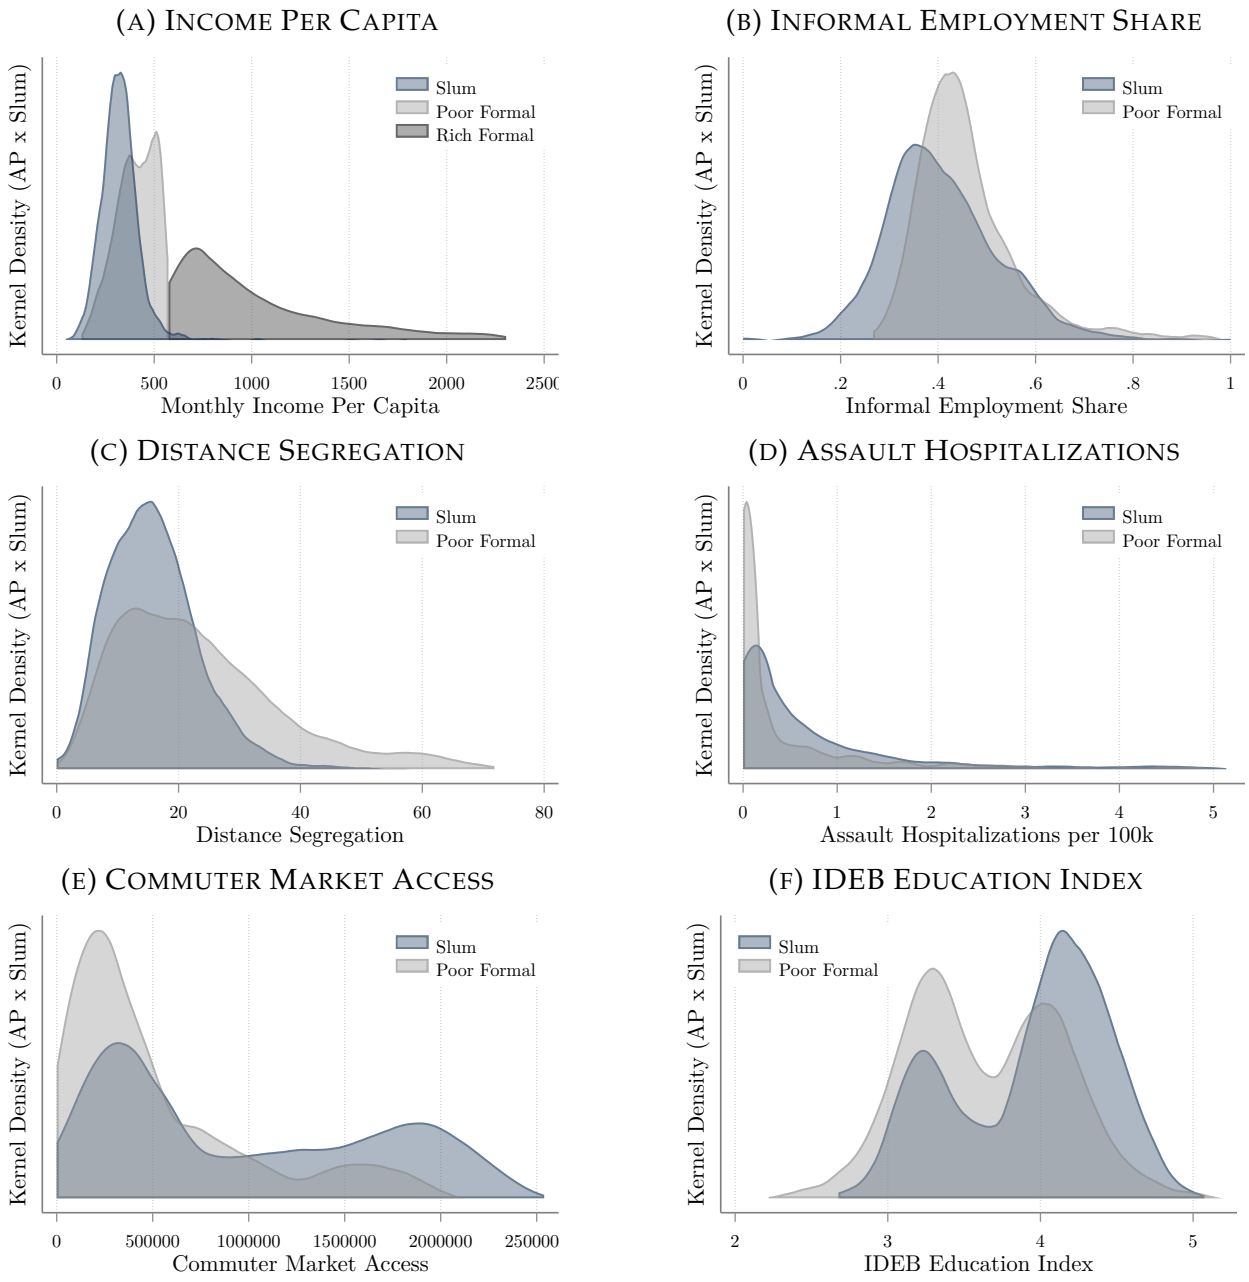


Notes: Figure displays coefficients with 95% confidence intervals from a regression of the logarithm of each neighborhood characteristic on a slum dummy. The dark blue dots are based on our full sample of slum and formal neighborhoods, while the light blue dots reflect only the sample of slums and "poor" below-median-income formal neighborhoods, where we define the median income across the population of our sample of metro areas

Second, and perhaps consistent with annecdotal knowledge, slum dwellers are subject to substantially higher victimization rates, more than twice as large compared to poor formal neighborhoods. Third, slums offer greater access to quality education compared to poor formal neighborhoods, while there is no difference relative to rich formal ones. Finally, commuter market access to formal jobs is 50% higher in slums than in poor formal neighborhoods, consisten with them being more distance integrated to rich neighborhoods.

These difference in averages are quite stark and already suggest important trade-offs faced by slum dwellers, which we will fully explore in the next sections. Examining the distributions of key variables also reveals extreme inequality across slums and formal neighborhoods (both poor and rich). Figure 3 displays the density plots for the three types of neighborhoods across the main variables discussed so far. In addition to the large inequality across slums, the plots reveal a salient bimodality in the distributions of both commuter market access and access to education.

FIGURE 3. SLUMS VS. FORMAL NEIGHBORHOODS: DISTRIBUTIONS OF KEY CHARACTERISTICS



Notes: Figures display kernel density across neighborhoods for monthly income per capita (Panel A), informal employment shares (Panel B), distance segregation (Panel C), assault hospitalizations per 100k residents (Panel D), commuter market access (Panel E), and the IDEB Education Index measure of school quality (Panel F), separately for slums and below-median-income formal neighborhoods, where we take all percentiles across the population of our sample of metro areas. For readability, we winsorize income at the 95th, distance segregation at the 99th and assault hospitalizations at the 90th percentiles.

IV. Slums and Social Mobility

We now take advantage of our rich data on intergenerational mobility across urban Brazil to answer our central question: are slums poverty traps or stepping stones to opportunity? For that, we focus on our representative sample of individuals born between 1983 and 1990 for whom we can link their long-run economic outcomes to those of their parents.

To reduce noise coming from census blocks with few observations, we use the National Bureau of Statistics (IBGE) crosswalks to map each census tract to a one-level higher statistical unit ("àrea de ponderação", AP), which roughly corresponds to the US Census tracts.[7] Our final unit of analysis is an AP $\times$ slum/formal cell. Throughout this section, we use the term "neighborhood" to refer to our AP $\times$ slum units, and we calculate income percentiles across all residents of our sample of metropolitan areas.

We now discuss 5 key facts about slums and intergenerational mobility.

Fact 1: Slums, on average, are not poverty traps.

We measure social mobility as the expected income rank, as adults, of children born to parents at the 25th percentile of the all-metropolitan-area income distribution. We compute this measure for all neighborhoods, and then run simple regressions of this summary statistic for mobility on a dummy for slums. Table 3 shows the results. In Column 1, we include all neighborhoods in our metropolitan areas and find a modest "slum penalty": children who grow up in slums end up around 1.5 percentiles lower in the income distribution as adults, relative to their peers who grow up in formal neighborhoods. This relatively mild effect disappears in Column 2 when we exclude formal neighborhoods with incomes per capita above the (population-weighted) median in our entire sample. A modest slum disadvantage in expected child rank reappears when we control for metro area fixed effects in Column 3, but overall, the average slum offers no worse – or better – intergenerational mobility than the average formal neighborhood.

We next exploit our rich individual-level data to confirm this striking lack of a slum penalty to mobility. We run the standard rank-rank regressions of a child's income rank on that of their parents, interacted with a dummy for growing up in a slum. Table 4 presents the results. Once we restrict the sample to slums and below-median-income formal neighborhoods in Column 2, we detect a modest slum penalty for the poorest

[7] Each census tract in the Brazilian census is designed to account for around 400 households. In major cities, like São Paulo and Rio de Janeiro, this can be as little as a single apartment building (or less).

TABLE 3. SLUM EFFECT ON SOCIAL MOBILITY

Sample:	*Full*	*Below-Median-Income*	
	(1) Expected income rank, children of p25 parents	(2) Expected income rank, children of p25 parents	(3) Expected income rank, children of p25 parents
Slum	-1.570***	-0.348	-0.864***
	(0.428)	(0.463)	(0.219)
Constant	38.91***	37.69***	37.79***
	(0.139)	(0.225)	(0.0805)
Metro Area FE			✓
Observations	3,419	1,966	1,966
R-squared	0.005	0.000	0.779

Notes: Robust standard errors in parentheses; *** indicates p<0.01, ** p<0.05, and * p<0.1. Table presents regressions at the neighborhood level (census weighting area by slum status) of the expected income rank of children growing up in the neighborhood with parents at the 25th income percentile on a dummy for slum status and a constant, with metro area fixed effects included as noted. We weight all regressions by neighborhood population. The "full" sample in Column 1 includes all neighborhoods in our metropolitan area sample; the "below-median-income" sample in Columns 2-3 includes all slums plus below-median-income formal neighborhoods. We measure all percentiles according to the aggregate income distribution in our metro area sample.

children, who lose up to 2 income ranks if they grow up in a slum, but this penalty disappears rapidly up the income distribution. Slums continue to show little effect on expected mobility even if we control, in Columns 3-5, for characteristics of parents and of the census tract where children grew up.

In Table 5, we re-estimate the rank-rank regressions including fixed effects at different levels of geographic aggregation for the adulthood location of children. These specifications, which purge the effect of slum residence on children's later location choice, also fail to uncover any economically significant slum effect on intergenerational mobility.

To better tease apart the effects of place from spatial sorting, we now focus on the subset of children whose parents move exactly once during childhood between a slum and a formal neighborhood or vice versa and compare children who move earlier versus later during childhood. We estimate "mover design" specifications by regressing a child's income rank as an adult on the number of years he or she spent in the slum as a child, either before or after the move, also interacted with parent income rank. Table 6 shows the results. Focusing on movers from slums to formal neighborhoods (Columns 3-4), we find a negative, significant exposure effect of additional years in a slum. For children at the very bottom of the income distribution, this level effect amounts to 6 income percentiles for

TABLE 4. RANK-RANK REGRESSIONS BY SLUM STATUS

Dependent Variable:	Child income rank				
	(1)	(2)	(3)	(4)	(5)
Parental income rank	0.539***	0.423***	0.326***	0.374***	0.321***
	(0.0005)	(0.001)	(0.001)	(0.001)	(0.001)
Slum	0.279**	-2.24***	-1.17***	-1.01***	-0.976***
	(0.135)	(0.139)	(0.137)	(0.141)	(0.140)
Parental income rank × Slum	-0.097***	0.019***	0.013***	0.014***	0.015***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Sample	Full	Below	Below	Below	Below
Parental Controls	No	No	Yes	No	Yes
Tract controls	No	No	No	Yes	Yes
Observations	2,460,825	1,024,719	1,024,103	1,024,718	1,024,102
R^2	0.29066	0.15748	0.20157	0.18849	0.20354
Adjusted R^2	0.29066	0.15748	0.20155	0.18848	0.20351

Notes: IID standard errors in parentheses; *** indicates p<0.01, ** p<0.05, and * p<0.1. Table presents individual-level regressions of child on parental income rank, fully interacted with a dummy for growing up in a slum. We variously control for, as noted, characteristics of the parents, including age, race, education, labor market status, region, and number of children, as well as of the census tract where the child grew up, including the informality share, the IDEB index, hospitalizations by assault, the infrastructure index, distance segregation, and CMA. The "Full" sample in Column 1 includes all neighborhoods in our metropolitan area sample; the "Below" sample in Columns 2-5 includes all slums plus below-median-income formal neighborhoods. We measure all percentiles according to the aggregate income distribution in our metro area sample.

children who spend a full 18 years in a slum. However, the positive interaction between slum exposure and parental rank renders the slum exposure effect insignificant for children of higher-income parents. Thus, in an echo of our individual results, additional years of slum exposure seem to hurt only the poorest out-movers.

The lack of an average effect of slum residence on social mobility might mask substantial heterogeneity among demographic groups. Table 7 thus re-estimates Column 2 of Table 4 by regions (northern vs. southern Brazil), gender, and race. Children who grow up in northern Brazil and non-white children suffer a marginally worse slum penalty, but on balance, the data provide little support for the idea that slums trap only certain, perhaps already-disdvantaged demographic groups in poverty.

TABLE 5. RANK-RANK REGRESSIONS, CONTROLLING FOR ADULTHOOD LOCATION

Dependent Variable:	Child income rank			
	(1)	(2)	(3)	(4)
Parental income rank	0.344***	0.400***	0.351***	0.332***
	(0.005)	(0.014)	(0.003)	(0.006)
Slum	-2.25	-2.17	-1.31	-0.802**
	(1.54)	(1.58)	(0.921)	(0.346)
Parental income rank × Slum	0.020	0.014	0.013	4.03×10^{-5}
	(0.016)	(0.016)	(0.017)	(0.010)
Fixed Effects				
North/South	Yes			
Slum-Below Median Income		Yes		
Region			Yes	
Immediate Region				Yes
Observations	1,024,325	932,872	1,024,325	1,024,325
R^2	0.20901	0.17219	0.21231	0.23317
Within R^2	0.10734	0.14004	0.10962	0.09435

Notes: xxx standard errors in parentheses; *** indicates p<0.01, ** p<0.05, and * p<0.1. Table presents individual-level regressions of child on parental income rank, fully interacted with a dummy for growing up in a slum. We include, as noted, fixed effects for the adulthood location of children, variously defined by northern vs. southern Brazil, slum vs. below-median vs. above-median-income formal neighborhoods, 5 regions, and 510 immediate regions (local labor markets). All specifications include only children growing up in slums or below-median-income formal neighborhoods, and we measure all percentiles according to the aggregate income distribution in our metro area sample.

Fact 2: The zero "slum effect" reflects offsetting forces

In Section III, we show that slums have very distinct characteristics than formal neighborhoods. To understand to what extent these different characteristics explain the zero estimated "slum effect" on IGM, we follow Firpo et al. (2018) and conduct a weighted Oaxaca-Blinder decomposition of the estimate in Column 2 of Table 3. Intuitively, the decomposition runs a regression of our expected child rank outcome on covariates among formal neighborhoods only; we then use the estimated coefficients to predict the change in average slum mobility were slums to have the same average value of a given characteristic as formal neighborhoods.

Following the labor literature, we term these estimated "counterfactual" changes in the slum effect on intergenerational mobility resulting from equalization of a given covariate *composition effects* and plot them in light blue in Figure 4. This statistical decomposition reveals that slums' higher victimization rates (assault hospitalizations) and lower dis-

TABLE 6. CHILDHOOD MOVER DESIGN: SLUM EXPOSURE EFFECT

Dependent Variable:	Child's Income Rank			
Sample:	*Formal to Slum*		*Slum to Formal*	
	(1)	(2)	(3)	(4)
Years in slum	-0.048	0.276**	0.121**	-0.350***
	(0.049)	(0.135)	(0.060)	(0.129)
Years in slum × Parental Rank		-0.006***		0.009***
		(0.002)		(0.002)
Observations	33,669	33,669	27,364	27,364
R^2	0.70878	0.70898	0.75902	0.75943

Notes: Robust standard errors in parenthesis, clustered at the origin by destination neighborhood (census weighting area by slum status) by parent income decile level. *** indicates p<0.01, ** p<0.05, and * p<0.1. Table presents individual-level regressions of child income rank on number of years spent in a slum neighborhood before age 18, interacted, as noted, with parental income rank. The sample includes only children whose parents moved exactly once before the child reached age 18 from a formal to a slum neighborhood (Columns 1-2) or vice versa (Columns 3-4). All specifications include fixed effects for origin by destination neighborhood by parent income decile, and we measure all percentiles according to the aggregate income distribution in our metro area sample.

TABLE 7. RANK-RANK REGRESSIONS BY DEMOGRAPHIC GROUP

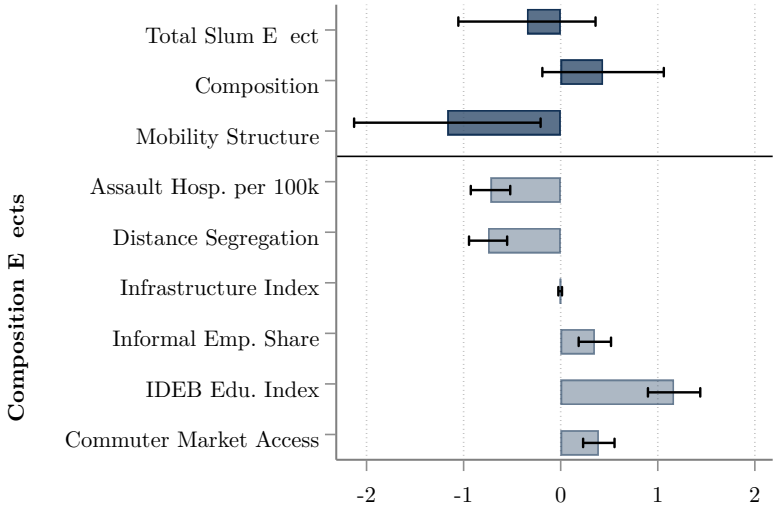
Dependent Variable:	Child income rank					
	(1)	(2)	(3)	(4)	(5)	(6)
Demographic Group:	South	North	Male	Female	White	Non-White
Parental income rank	0.377***	0.352***	0.394***	0.459***	0.424***	0.385***
	(0.002)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
Slum	-1.11***	-2.90***	-2.31***	-1.66***	-1.54***	-2.15***
	(0.234)	(0.172)	(0.178)	(0.184)	(0.284)	(0.162)
Parental income rank × Slum	0.004	0.023***	0.020***	0.011**	0.019***	0.012***
	(0.005)	(0.004)	(0.004)	(0.004)	(0.006)	(0.004)
Observations	514,752	509,963	486,573	538,146	316,167	694,312
R^2	0.11656	0.11269	0.17566	0.19085	0.16291	0.12365
Adjusted R^2	0.11656	0.11268	0.17565	0.19085	0.16290	0.12365

Notes: IID standard errors in parentheses; *** indicates p<0.01, ** p<0.05, and * p<0.1. Table presents individual-level regressions of child on parental income rank, fully interacted with a dummy for growing up in a slum, separately for children growing up in northern vs. southern Brazil, by gender, and by race. All specifications include only children growing up in slums or below-median-income formal neighborhoods, and we measure all percentiles according to the aggregate income distribution in our metro area sample.

tance segregation decrease intergenerational mobility in slums, relative to poor formal neighborhoods. However, these negative effects are offset by the fact that slums have lower informal employment, higher school quality (i.e. IDEB Education Index values),

and better commuter market access. Thus, in a statistical sense, better labor market access and educational opportunities offset higher crime in Brazilian slums, such that they on average, as we find, have no worse mobility than formal neighborhoods.

FIGURE 4. OAXACA DECOMPOSITION OF SLUM MOBILITY EFFECT



Notes: Figure displays the results of a re-weighted Oaxaca-Blinder decomposition (Firpo et al. 2018) of mean difference in mobility between slums and below-median-income formal neighborhoods, i.e. the "slum effect" in the neighborhood-level regression of the expected income rank of children growing up in the neighborhood with parents at the 25th income percentile on a dummy for slum status (Column 2 of Table 3). The dark blue bars in the upper panel display the coefficient on the slum dummy, the total composition effect, and the total "mobility structure" effect, analogous to the usual wage structure effect. The light blue bars in the lower panel display the contribution of each covariate to the composition effect. Distance segregation, the IDEB education index, and commuter market access enter the model in logs. Note that the "counterfactual" in the decomposition uses formal coefficients and slum covariates, reweighted using a logit model which includes all covariates. We weight all regressions by neighborhood population; the sample includes all slums plus below-median-income formal neighborhoods. We measure all percentiles according to the aggregate income distribution in our metro area sample, and error bars reflect robust standard errors. Note that informal share of employment includes both informal wage workers and the self-employed.

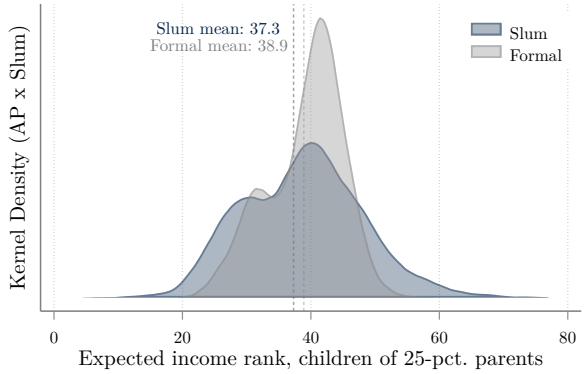
Fact 3: Some slums are poverty traps, some are stepping-stones

The fact that offsetting effects leave the average slum-dwelling child with no worse intergenerational opportunities masks substantial differences in social mobility across different slums. Indeed, in Figure 5, we plot the distribution across neighborhoods of expected income ranks for children with parents at the 25th percentile. Panel A compares slums to all formal neighborhoods, while Panel B compares slums to our preferred below-median-income formal neighborhood sample. Both neighborhood types exhibit substantial heterogeneity in mobility outcomes, but slums display a particularly striking range of intergenerational mobility, with a 90-10 ratio across slums of 1.9. Figure 6 uncovers a huge

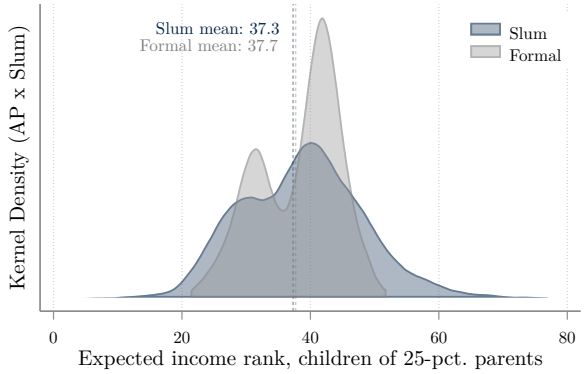
range of opportunity, or lack thereof, in terms of expected child ranks even across slums within the same metro area.

FIGURE 5. DISTRIBUTION OF MOBILITY IN SLUMS, VERSUS...

(A) ALL FORMAL



(B) BELOW-MEDIAN-INCOME FORMAL



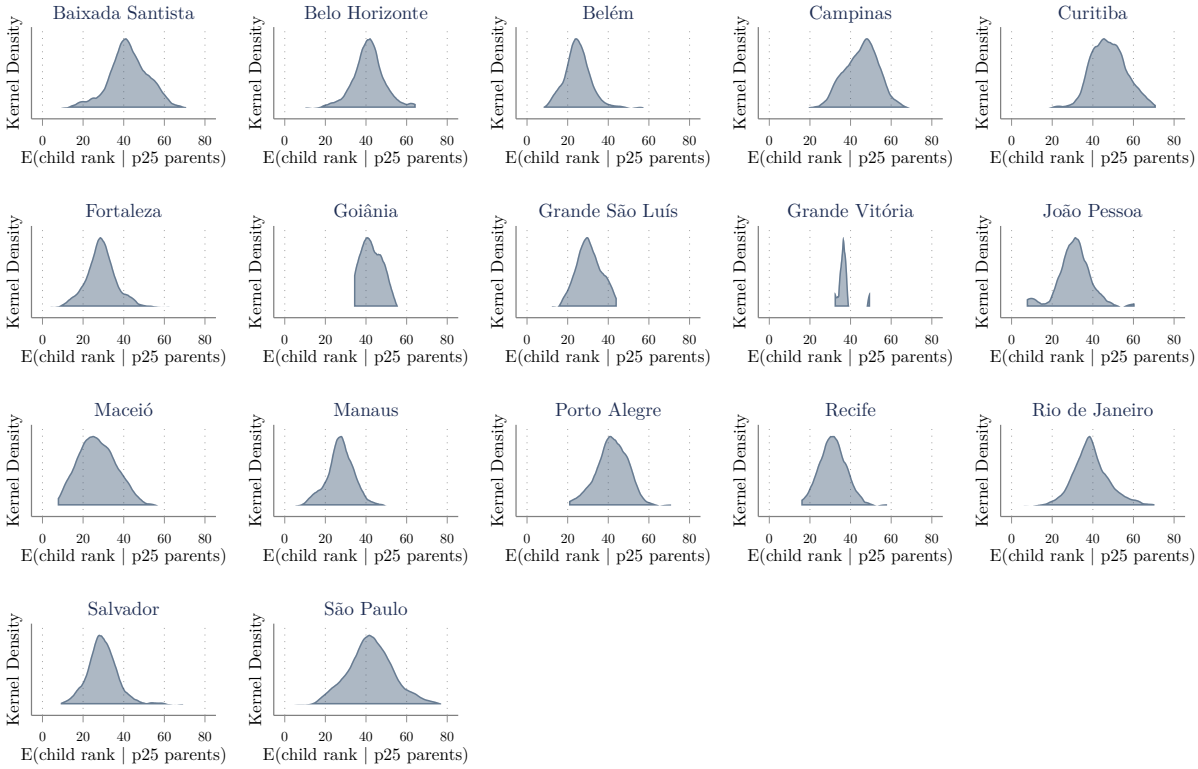
Notes: Figures display kernel density across neighborhoods of the expected income rank of children born in a neighborhood with parents at the 25th income percentile, in Panel (A) for slums and all formal neighborhoods separately, and in Panel (B) for slums and all below-median-income formal neighborhoods, where we take all percentiles across the population of our sample of metro areas.

Thus, this evidence indicates that a subset of slums are characterize by very low to zero upward social mobility, with children of poor parents remaining poor and therefore reproducing poverty cycles across generations. In contrast, some slums have high upward mobility across generations in terms of income rank of the children when adults. For the remainder of the paper, we define as *poverty traps* the slums whose IGM measure (expected child rank with 25th-percentile parents) falls in the bottom tertile across slums; symmetrically, we define as *stepping stones* those slums that fall in the top tercile of the distribution.

In Figure 7, we investigate whether slums that have lower IGM are also associated with higher probability of their children living in slums as adults. As the figure shows, the slums with the lowest expected income rank for children of 25th-percentile parents also have much higher probability to live in slums as adults.[8] This reinforces the fact that these slums indeed constitute intergenerational poverty traps.

[8]Even children growing up in high-opportunity slums have a higher probability of ending up in slums as adults, relative to their peers raised in formal neighborhoods, as Appendix Figure A.2 demonstrates.

FIGURE 6. DISTRIBUTION OF MOBILITY IN SLUMS BY METRO AREA



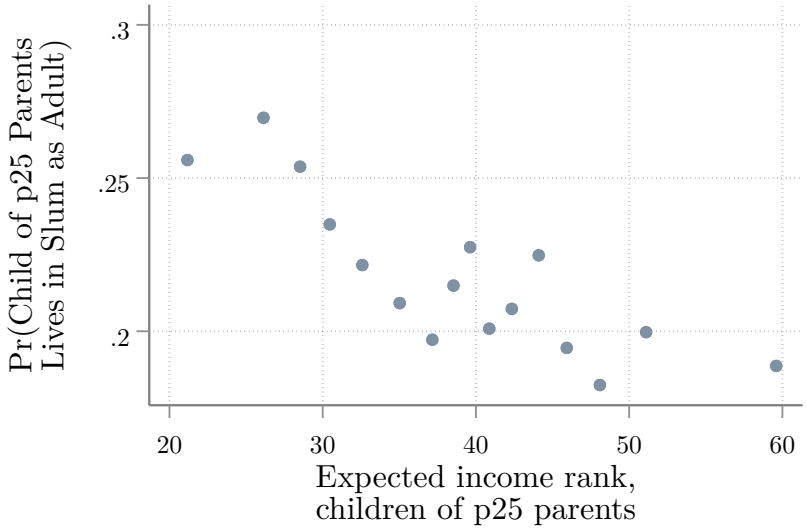
Notes: Figures display kernel densities across slum census tracts of the expected income rank of children born in a neighborhood with parents at the 25th income percentile, separately for each metro are in our sample. We exclude 2 metro areas (Natal and Grande Teresina) with very small numbers of slum tracts, and we take all percentiles across the population of our sample of metro areas.

Fact 4: Crime, schools, and informality explain the mobility gap across slums

Given the huge heterogeneity in intergenerational mobility across slums, a vital question for policy involves which characteristics predict whether a slum traps its residents in poverty. In Table 8, we display a regression, among our sample of slums, of an indicator for poverty trap status on our set of relevant neighborhood characteristics. We find, in Column 1, that higher crime, a higher informal share of employment, and poor quality schools tend to predict which slums are traps (i.e. fall in the bottom-tertile intergenerational mobility).

We can take one step further and conduct the same weighted Oaxaca-Blinder decomposition as in Figure 4 to understand the main determinants of the 19-percentile difference in mean mobility between the "stepping stones" and "poverty traps". We decompose this IGM gap among slums using the same set of covariates as before. As the light-blue composition

FIGURE 7. PROBABILITY OF LIVING IN SLUM AS ADULT VERSUS EXPECTED INCOME RANK OF CHILDREN OF 25TH-PERCENTILE PARENTS, ACROSS SLUMS



Notes: Figure displays binned scatterplot, across slum neighborhoods, of the probability that a child growing up in a given slum neighborhood lives in a slum as an adult, versus the expected income rank of children born in the neighborhood, both measured given parents at the 25th income percentile. We take all percentiles across the population of our sample of metro areas.

effect bars in Figure 8 make clear, school quality in and of itself already accounts for 6 percentiles of the mobility gap; informal employment accounts for another two, as do assault hospitalizations. Thus, most of the gap in mobility between the low- and high-mobility slums is explained by these three factors

V. A SPATIAL MODEL OF INTERGENERATIONAL HUMAN CAPITAL INVESTMENT

To rationalize the empirical facts discussed so far, this section develops a novel model of internal city structure with an overlapping generations structure that incorporates the main determinants of neighborhood-level IGM discussed in Section IV.

Environment

We consider a closed city composed of N locations, or neighborhoods, each of which belongs to one of two mutually-exclusive and exhaustive sets: formal neighborhoods, $\mathscr{F}$, and informal slum neighborhoods, $\mathscr{S}$. Each neighborhood's offers a fixed supply $\overline{H}_i$ of housing.

TABLE 8. PREDICTORS OF SLUM POVERTY TRAP STATUS

	(1) Pr(Poverty Trap)	(2) Pr(Poverty Trap)
Assault Hospitalizations per 100k	0.00416***	0.00128
	(0.000640)	(0.00113)
Log Distance Segregation	-0.121***	0.0375
	(0.0298)	(0.0507)
Infrastructure Index	-0.0237	0.000887
	(0.130)	(0.0922)
Informal Employment Share	0.506***	0.102
	(0.162)	(0.143)
Log IDEB Education Index	-1.551***	-0.0927
	(0.149)	(0.207)
Log Commuter Market Access	-0.0203	0.0164
	(0.0133)	(0.0233)
Metro Area FE		✓
Observations	1,268	1,268
R-squared	0.549	0.678

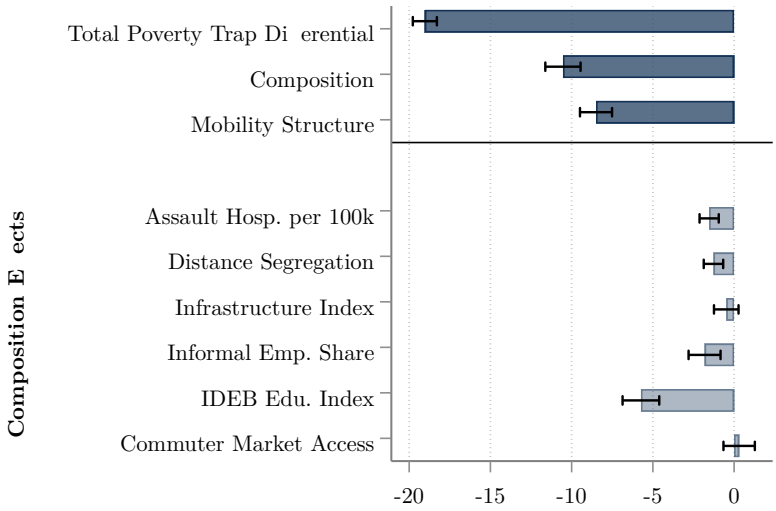
Notes: Robust standard errors in parentheses; *** indicates p<0.01, ** p<0.05, and * p<0.1. Table presents regressions, across slum neighborhoods (census weighting area by slum status), of a dummy for a slum having a child expected income rank, given parents at the 25th income percentile, in the bottom tertile of slums ("poverty trap") on neighborhood covariates. We weight all regressions by neighborhood population, and all specifications include a constant and, if noted, metro area fixed effects. Note that informal share of employment includes both informal wage workers and the self-employed.

The workers who populate our city belong to one of an infinite series of non-overlapping generations of fixed size $\overline{L}$. Workers in each generation enter the labor force with heterogeneous levels of human capital h and draw idiosyncratic preference shocks over home locations. They then choose a home location, indexed by i and with fixed amenity B_i.

Workers receive a formal job offer with a home-location-specific probability λ_i, which reflects search frictions and potential discrimination against residents of certain neighborhoods. Workers who receive a formal job offer draw idiosyncratic preferences over and then choose a (formal) work location. Each formal work location, indexed by j, pays a fixed wage w_j per unit of human capital. Iceberg commute costs τ_{ij} reduce formal work income proportionately. Workers without formal job offers work informally in their home neighborhoods for an informal wage $\tilde{w}_i$, itself a neighborhood-invariant proportion χ of the local formal wage.

Subsequently, each generation of worker-parents chooses a human capital level h' for their children, which requires some amount m of direct parental investment. After they choose

FIGURE 8. OAXACA DECOMPOSITION OF POVERTY TRAP MOBILITY EFFECT



Notes: Figure displays the results of a two-fold Oaxaca-Blinder decomposition of the mean difference in mobility between poverty-trap and stepping-stone slums, those in the bottom and top tertiles of 25th-percentile child mobility across slums. The dark blue bars in the upper panel display this poverty trap effect, the total composition effect, and the total "mobility structure" effect, analogous to the usual wage structure effect. The light blue bars in the lower panel display the contribution of each covariate to the composition effect. Distance segregation, the IDEB education index, and commuter market access enter the model in logs. Note that the composition effects use poverty-trap covariates and coefficients from a regression of the expected income rank of children at the 25th percentile on the aforementioned covariates within the entire sample of slums. We weight all regressions by neighborhood population. Error bars reflect robust standard errors. Note that informal share of employment includes both informal wage workers and the self-employed.

their children's human capital, *informal* workers face income risk, namely a probability δ of job loss. Workers who lose their jobs lose all work income and receive only a government benefit $\underline{w}$. Finally, workers choose consumption of a homogeneous, freely-traded good as well as housing. We solve the model in steady-state and thus omit time subscripts from all variables.

Preferences

Workers derive utility from consumption during their own lifetime as well as, in contrast to standard spatial models, from the expected lifetime utility of their one child. We begin with the former. A worker ω resident in neighborhood i has Stone-Geary preferences over his own consumption of a freely-traded final numeraire good, c, and housing, l,

$$u_i(c,l) \equiv B_i c^{1-\gamma}\left(l-\bar{l}\right)^{\gamma}. \tag{1}$$

This consumption utility depends directly on the home neighborhood amenity B_i, which may differ systematically across formal and slum locations. The parameter γ denotes the housing share, and the minimum housing requirement $\bar{l} > 0$ makes housing a subsistence good.

We then define a worker *lifetime* utility which aggregates the consumption utility laid out in the preceding paragraph with the expected utility of her child. In particular, the lifetime utility $U_{ilj\omega}$ of a worker resident in i, employed in sector s, and who works in location j takes a Cobb-Douglas form,

$$U_{ijs\omega}(c, l, h') \equiv u_i \left(c, l\right)^{1-\beta} \overline{U}\left(h'\right)^{\beta} \varepsilon_{i\omega} \zeta_{j\omega}^{\mathbb{1}\{s=F\}}. \tag{2}$$

The share β controls intergenerational altruism, i.e. the weight the worker places on the expected utility, $\overline{U}\left(h'\right)$, of her child, given the latter's human capital h'. Lifetime utility also depends on an idiosyncratic home location preference shock $\varepsilon_{i\omega}$ and, for the formally-employed, a work location preference shock $\zeta_{j\omega}$.

Human Capital Production

We first lay out the neighborhood-specific human capital production technology at the heart of our model. We assume that children derive human capital from direct parental investments m, whose productivity A_i varies across neighborhoods. In particular, the child of a worker ω in neighborhood i receives human capital

$$h'_{\omega} \equiv A_i m_{\omega}^{\alpha}. \tag{3}$$

The elasticity α controls the returns to direct parental investments, and we allow productivity to depend on a vector of neighborhood characteristics x_{ki} indexed by k, in a simple Cobb-Douglas fashion, $A_i \equiv \iota_0 \Pi_k x_{ki}^{\iota_k}$. A positive elasticity ι_k for a given covariate thus induces a natural complementarity with parental investment.

Choices

We now solve the sequential choices which make up the worker problem backwards and begin with the choices of goods and housing consumption.

Consumption

Workers choose consumption only after their choices of home, work, and child human capital and, if applicable, job loss. Consider a worker ω resident in i with income I_ω who has invested m_ω in his child's human capital. Given local housing rent r_i, workers' consumption must satisfy a budget constraint, $c + r_i l + m_\omega = I_\omega$. Thus, indirect consumption utility takes the form, $v_{i\omega} \equiv \overline{\gamma}\frac{B_i}{r_i^{\overline{\gamma}}}\left(I_\omega - m_\omega - r_i\bar{l}\right)$, where $\overline{\gamma}$ denotes the usual Cobb-Douglas constant.

Before workers choose consumption, informal workers may suffer a probability-δ job loss shock. To derive the resulting expected utility, we substitute indirect consumption utility into lifetime utility in (2) and take expectations over the job loss shock. An informally-employed worker ω with human capital h who lives in i and has chosen child human capital h' thus enjoys expected lifetime utility, net of idiosyncratic preferences, equal to

$$EU_{iI\omega} \equiv \left(\overline{\gamma}\frac{B_i}{r_i^{\overline{\gamma}}}\right)^{1-\beta}\left[(1-\delta)\left(\chi w_i h_\omega - m_\omega - r_i\bar{l}\right)^{1-\beta} + \delta\left(\underline{w} - m_\omega - r_i\bar{l}\right)^{1-\beta}\right]\overline{U}\left(h'_\omega\right)^{\beta}. \quad (4)$$

Note that informal-sector workers all work locally and so pay no commute costs out of their wage $\tilde{w}_i = \chi w_i$ if they remain employed. Informal workers who lose their jobs must instead finance consumption and human capital investment out of the neighborhood-invariant government benefit $\underline{w}$. A formal-sector worker, in turn, who commutes from i to j has (degenerate) expected utility

$$EU_{ijF\omega} \equiv \left[\overline{\gamma}\frac{B_i}{r_i^{\overline{\gamma}}}\left(\frac{w_j h_\omega}{\tau_{ij}} - m_\omega - r_i\bar{l}\right)\right]^{1-\beta}\overline{U}\left(h'_\omega\right)^{\beta}. \quad (5)$$

Formal workers thus earn a certain income but must lose some of their wages w_j to the commute costs τ_{ij}.

Child Human Capital

Conditional on home and work location as well as sector of work, workers choose a human capital level h' for their child. For each h', the production function in (3) implies a necessary level of direct parental investment $m_\omega = \left(\frac{h'_\omega}{A_i}\right)^{1/\alpha}$. Neighborhoods whose characteristics, e.g. school quality, endow them with higher educational productivity A_i thus require less

investment to attain a given human capital level. We can rewrite the expected utilities (4) and (5) as functions $EU_{iI}(h, h')$ and $EU_{ijF}(h, h')$ only of workers' own human capital and that which they choose for their child. Workers' choice of child human capital then maximizes these ex-ante expected utilities and, in doing so, trades off this direct cost against the resulting marginal increase in child utility $\overline{U}(\cdot)$. Thus, our model generates rich within-neighborhood heterogeneity in parents' human capital h, the amounts they invest, and therefore the human capital h' their children attain.

Workplace

After they choose a home location, workers draw formal job offers and, if they receive one, choose a formal work location. Consider first workers who receive a probability-λ_i formal employment offer, i.e. the opportunity to choose a formal workplace.[9] Workers draw multiplicative idiosyncratic workplace preference shocks $\zeta_{j\omega}$ from a Frechet distribution with shape parameter θ. They then choose a workplace based on the product of the utility $EU_{ijF}(h, h')$ available at the optimal child human capital h' and the idiosyncratic preference. The usual Frechet results imply an expected value of formal employment in i of

$$\Omega_{iF}(h) \equiv \Gamma\left(\frac{\theta-1}{\theta}\right)\left\{\sum_{j'}\left[\max_{h'} EU_{ij'F}(h, h')\right]^{\theta}\right\}^{1/\theta}, \tag{6}$$

and probability

$$\pi_{j|iF}(h) \equiv \left[\max_{h'} EU_{ijF}(h, h')\right]^{\theta} / \sum_{j'}\left[\max_{h'} EU_{ij'F}(h, h')\right]^{\theta}. \tag{7}$$

that workers in i who receive a formal offer choose to work in j. Workers who fail to receive the formal employment offer receive the expected utility of local informal employment, given optimally-chosen child human capital h', $\Omega_{iI}(h) \equiv \max_{h'} EU_{iI}(h, h')$.

Home

In the model's outermost layer, workers first choose home locations based on the expected utility across formal and informal employment. We assume that they draw Frechet-distributed idiosyncratic preferences with shape η. This distributional assumption again yields the standard home neighborhood choice probabilities, conditional on workers' own

[9] We assume that formal exceed informal wages sufficiently such that all workers who receive formal offers accept them.

human capital h,

$$\pi_i(h) = \frac{\left[\lambda_i \Omega_{iF}(h) + (1-\lambda_i)\,\Omega_{iI}(h)\right]^{\eta}}{\sum_{i'}\left[\lambda_{i'} \Omega_{i'F}(h) + (1-\lambda_{i'})\,\Omega_{i'I}(h)\right]^{\eta}}. \tag{8}$$

The expected location value in the numerator reflects not only the amenities, rents, and wages relevant to workers' own utility but, in our intergenerational framework, also the local productivity of human capital investment. Note that a worker with human capital h enjoys lifetime expected utility

$$\overline{U}(h) \equiv \Gamma\left(\frac{\eta-1}{\eta}\right)\left\{\sum_{i'}\left[\lambda_{i'} \Omega_{i'F}(h) + (1-\lambda_{i'})\,\Omega_{i'I}(h)\right]^{\eta}\right\}^{1/\eta}. \tag{9}$$

Equilibrium

We now define our steady-state equilibrium. Denote the pdf of the invariant, steady-state distribution of human capital by $g(h)$. Given the Stone-Geary assumption on housing demand and the aforementioned home-location choice probabilities, housing market clearing requires

$$\begin{aligned}\int \overline{L} g(h)\pi_i(h)\Bigg\{\gamma\lambda_i \sum_j \pi_{j|iF}(h)\left[\frac{w_j h}{\tau_{ij}} - \left(h'^{*}_{ijF}(h)/A_i\right)^{1/\alpha} - r_i\bar{l}\right] + \\ \gamma(1-\lambda_i)\Bigg[(1-\delta)\left(\chi w_i h - \left(h'^{*}_{ijF}(h)/A_i\right)^{1/\alpha} - r_i\bar{l}\right) + \\ +\delta\left(\underline{w} - \left(h'^{*}_{ijF}(h)/A_i\right)^{1/\alpha} - r_i\bar{l}\right)\Bigg] + r_i\bar{l}\Bigg\}dh = r_i\overline{H}_i. \end{aligned} \tag{10}$$

Note that we denote the optimal child human capital choice of parents with human capital h, living in i and who work formally in j as $h'^{*}_{ijF}(h)$ and the corresponding choice if informally employed as $h'^{*}_{iI}(h)$. Equilibrium in our model depends on one set of universal parameters and another model geography of parameters specific to home locations, work locations, or both. These parameters determine six sets of endogenous variables or functions, as follows.

Equilibrium. *Given parameters* $\left\{\overline{L}, \eta, \theta, \underline{w}, \gamma, \bar{l}, \beta, \boldsymbol{\iota}, \alpha, \chi, \delta\right\}$ *and model geography* $\{\mathscr{F}, \mathscr{S}, \boldsymbol{B}, \boldsymbol{w}, \boldsymbol{\tau}, \boldsymbol{x}, \boldsymbol{\lambda}, \overline{\boldsymbol{H}}\}$*, an equilibrium is a vector of formal-sector expected utility functions* $EU_{ijF}(h, h')$ *and work-location choice probability functions* $\pi_{j|iF}(h)$ *for each home-work pair, a vector of informal-sector expected utility functions* $EU_{iI}(h, h')$ *and (unconditional) residential choice probability functions* $\pi_i(h)$ *for each home location, an expected lifetime utility function* $\overline{U}(h)$*, and a vector*

of neighborhood rents r_i satisfying (i) formal-sector expected utility (5) and the work location choice probability equation (7); (ii) informal-sector expected utility (4) and the home location choice probability equation (8); (iii) the dynamic program in (9); and (iv) housing market clearing (10).

VI. Quantification of Intergenerational Mobility

We adapt our baseline model from Section V for quantification and then calibrate our model. To reduce the computational burden, in this preliminary exercise we restrict our estimation and counterfactual exercises to the metropolitan area of Rio de Janeiro. Rio is a good case study, as it has a large number of neighborhoods, and its statistics are comparable to the overall sample of neighborhoods and slums discussed in Section IV.

We first externally calibrate a subset of parameters based on aggregate and neighborhood-level data as well as established estimates in the literature. Second, we internally calibrate key human capital production and preference parameters to match moments in our combined datasets and invert the model to obtain home-location-specific amenities and work-location-specific wages. Table 9 summarizes our calibration strategy.

Model Modifications

We undertake minor modifications to the model from the preceding section to facilitate numerical solutions and calibration. To avoid negative consumption values, we assume that the government pays all workers a housing benefit ς equal to the maximum cost of subsistence housing consumption across neighborhoods such that $\varsigma \equiv \max_i r_i \bar{l}$. In a similar vein, we bound the range of possible human capital values below by some $h_{\min}$ and then modify the production function in (3) such that workers can endow their children with this minimum level at zero investment cost, i.e. $h'_\omega \equiv h_{\min} + A_i m_\omega^\alpha$. Finally, we parameterize formal commute costs as a function of straight-line distance Δ_{ij} between home and work locations, such that $\tau_{ij} \equiv \exp\left(\kappa \Delta_{ij}\right)$, where κ captures the disutility of commute distance.

Externally-Calibrated Parameters

First, we externally calibrate labour market, commuting, and miscellaneous parameters based on the literature and data. We take the home and workplace Frechet shapes $\eta = 2.16$ and $\theta = 3.398$, respectively, from Tsivanidis (2023), where the latter equals the average of his values for low- and high-skill workers. The informal job loss probability $\delta = 0.084$ equals the yearly transition rate from informal employment to unemployment

TABLE 9. CALIBRATION STRATEGY

Parameter	Description	Method	Data/Source
Labour Market			
η	Home Frechet Shape	External	Tsivanidis (2023)
θ	Formal Work Frechet Shape	External	Tsivanidis (2023)
κ	Commute Distance Disutility	Internal: commute gravity	Chaves et al. (2023)
χ	Informal Wage Factor	Internal: ratio of mean informal to formal income	Census (2000)
δ	Informal Pr(Job Loss)	External	Ulyssea (2010)
Human Capital Production			
$\boldsymbol{\iota}$	Covariate Coefficients	Internal: neighborhood-level rank-rank regression with covariates	Individual-level + IDEB + SUS
α	Parental Investment Elasticity	Internal: neighborhood-level rank-rank regression with covariates	Individual-level + IDEB + SUS
Preferences			
γ	Housing Share	Internal: regression of housing expenditure share on income	POF
$\bar{l}$	Housing Subsistence Requirement	Internal: regression of housing expenditure share on income	POF
β	Intergenerational Altruism	Internal: mean educational expenditure share of income	POF
Geography			
$\boldsymbol{B}$	Amenities	Model Inversion: population	Census (2000)
$\boldsymbol{w}$	Formal Wages	Model Inversion: formal employment by workplace	RAIS
$\boldsymbol{\Delta}$	Commute Distances	External	geobr
$\boldsymbol{\lambda}$	Pr.(Formal Job Offer)	External	Census (2000)
$\overline{\boldsymbol{H}}$	Housing Supply	External	Census (2000)
$\boldsymbol{x}$	Neighborhood Characteristics	External	IDEB + SUS
Miscellaneous			
$\underline{w}$	Government benefit	External	POF
$\overline{L}$	City Population	External	Census (2000)

Notes: Table lists model parameters along with calibration method and data or source used.

from Ulyssea (2010). We set the government benefit $\underline{w} = 5.3$ such that the ratio between between $\underline{w}$ and the mean formal wage across work locations in our model equals the ratio the average government benefit, conditional on receipt of nonzero government benefits, and average household income observed in the POF expenditure data. Finally, we set $\overline{L} = 10,948,463$ equal to the Rio de Janeiro metro area population observed in the 2000 Census.

Second, we externally calibrate a subset of the model geography. Our set of home locations consists of the 624 AP $\times$ slum units in the Rio de Janeiro metropolitan area. For tractability,

we define more aggregated spatial units for formal workplaces: following Chaves et al. (2023), our workplaces consist of 51 units: IBGE subdistricts within the municipality of Rio de Janeiro and municipalities elsewhere within the metro area.[10] We calculate straight-line distances Δ_{ij} between home and work location centroids using shapefiles provided by the geobr package. The formal offer probabilities λ_i equal the formal share among wage workers in a given AP $\times$ slum in the 2000 Census. We approximate housing supply $\overline{H}_i$ using the number of households in non-collective, permanent accommodations in the 2000 Census. Finally, the neighborhood characteristics in $\boldsymbol{x}$ capture school quality and crime prevalence. To measure the former, we employ the IDEB Education Index, and to approximate crime rates, we measure the number of hospitalizations due to assault per 100,000 residents in the SUS data. In each case, as in Section **??**, we take an inverse-distance-weighted average for each AP $\times$ slum.[11]

Internally-Calibrated Parameters

We solve the full model to match five auxiliary models and thereby calibrate the remaining human capital production, labour market, and preference parameters. Conditional on these parameters, we follow in the tradition of Ahlfeldt et al. (2015) and invert the model to obtain home-location amenities and work-location formal wages.

Our first two auxiliary models involve labour market data. For our first auxiliary moment, we employ mobile phone data published by Chaves et al. (2023) to estimate a commute gravity regression. For an origin-destination geography which corresponds to subdistricts of Rio municipality plus other municipalities in the metropolitan area, we approximate commute flows using their aggregations of mobile phone pings. Specifically, we measure the probability $\pi_{j|i}$ that a resident of location i appears most frequently during working hours on workdays in a given neighborhood j. Our auxiliary model then consists of a simple regression of these workplace choice probabilities on straight-line origin-destination distance, $\log \pi_{j|i} = \beta_0^1 + \beta_\Delta^1 \Delta_{ij} + \delta_i + \delta_j + \varepsilon_{ij}$, with origin-destination fixed effects and where ε_{ij} denotes the idiosyncratic error. Our second auxiliary moment equals the ratio of mean informal to mean formal wages as measured, nationwide, in the 2000 Census.

The third auxiliary model runs a rank-rank regression across neighborhoods i in our sample

[10] More specifically, we use aggregations of our AP x slum units which very closely correspond to subdistricts within Rio municipality and municipalities elsewhere. In almost all cases, AP x slum units can be matched many-to-one with a subdistrict or municipality. In the few cases where AP x slum units cross subdistrict boundaries, we group the entire AP x slum with the subdistrict in which it has greater formal employment.

[11] Given straight-line distance Δ_{ij}, the weights equal $1/\exp\left(0.5 \times \Delta_{ij}\right)$.

of metropolitan areas. We measure the expected income rank r_i^c of children growing up in AP $\times$ slum i as well as the average income rank r_i^p of their parents and then run a population-weighted regression,

$$r_i^c = \beta_0^3 + \beta_{IGE}^3 r_i^p + \mathbf{X}_i \beta_x^3 + \vartheta_i. \tag{11}$$

We additionally include the same aforementioned assault hospitalizations and IDEB education index covariates, in logs, in the vector X_i, and ϑ_i denotes the usual idiosyncratic error.

Our fourth and fifth auxiliary models employ POF expenditure data. For a household l, we first calculate educational expenditure m_l and income I_l as well as expenditure on housing rents, r_l. For our fourth model, we then regress, among the sample of households with nonzero rental expenditure, the share of income net of educational expenditure spent on rent, on household income, or, augmented with an idiosyncratic error,

$$\frac{r_l}{I_l - m_l} = \beta_0^4 + \beta_I^4 \log I_l + \epsilon_l. \tag{12}$$

The fifth and final model is the mean, across households, of educational expenses as a share of income, $E\left(m_l / I_l\right)$.

Conditional on all preceding parameters, we then invert the model to recover the rest of the geography. Namely, we match population by home location from the 2000 Census to recover home location amenities B_i. To obtain work location formal wages w_j, we exactly match formal employment by workplace, as observed in RAIS.

Identification

We now provide the usual heuristic argument for identification of the eight parameters estimated via internal calibration; Table 10 lists parameters and corresponding moments in both the data and model.

First, the commute gravity regression approximates speaks to the disutility κ of commute distance. Though this "usual" commute gravity regression does not hold exactly in our model, the sensitivity β_Δ^1 of location choices to commute distance nonetheless increases with κ. Second, the ratio of mean formal to informal wages, though also a function of home and work location choices, varies strongly with the within-location informal wage factor of proportionality, χ.

Third, the across-neighborhood rank-rank regression identifies the human capital produc-

TABLE 10. INTERNAL CALIBRATION

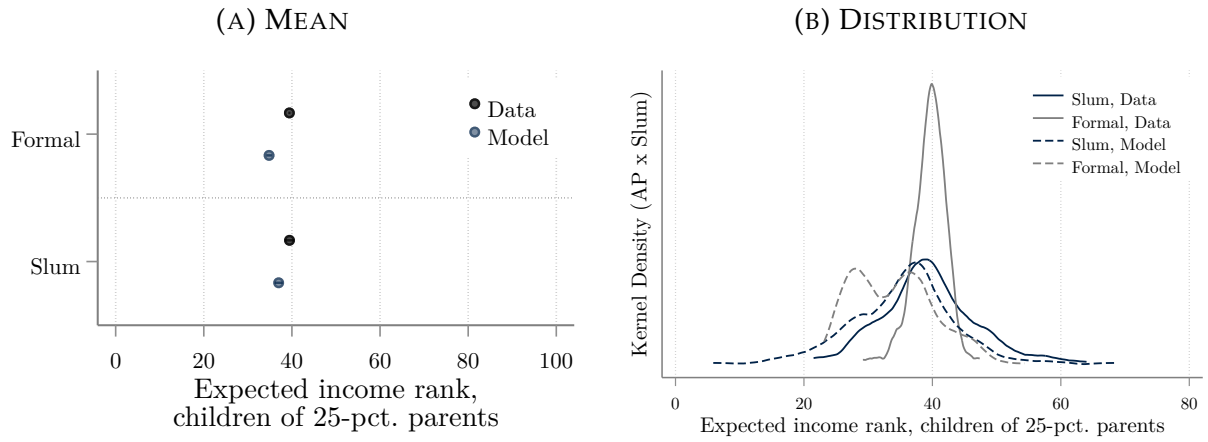
Moment			Parameter		
Description	*Data*	*Model*	*Description*		*Value*
Commute Gravity: Distance Coefficient	-0.058	-0.051	κ	Commute Distance Disutility	0.035
Mean Informal/Mean Formal Wage	0.86	0.86	χ	Informal Wage Factor	0.75
Rank-Rank: Parental Rank Coefficient	0.61	0.63	α	Parental Investment Elasticity	0.65
Rank-Rank: IDEB Index Coefficient	0.18	0.36	ι_{school}	School Quality Coefficient	0.8
Rank-Rank: Hospitalizations Coefficient	-0.003	-0.004	ι_{crime}	Crime Coefficient	-0.011
Housing Expenditure: Intercept	0.97	0.88	γ	Housing Share	0.27
Housing Expenditure: Income Slope	-0.08	-0.14	$\bar{l}$	Housing Subsistence Requirement	0.05
Educational Expenditure Share	0.015	0.03	β	Intergenerational Altruism	0.047

Notes: This table displays the moments used in internal calibration, listing the values in the data and in the model, using which we simulate 100,000 workers and run equivalent specifications. I also list the model parameter heuristically corresponding to each moment, along with its internally-calibrated value.

tion function parameters. In our model, parents invest directly in their children's education in approximate proportion to their income, so this regression closely approximates the human capital production function (3). The across-neighborhood intergenerational income elasticity β^3_{IGE} thus pins down the parental investment elasticity α, and the coefficients β^3_x on the IDEB education index and assault hospitalizations pin down their equivalents $\boldsymbol{\iota}$ in the human capital production function.

Fourth and fifth, the expenditure models quantify the housing share γ, the housing subsistence requirement $\bar{l}$, and the intergenerational altruism share β. The intercept β^4_0 in the housing expenditure regression varies with the average proportion of income spent on housing and thus the Cobb-Douglas share γ. In our model, due to the presence of non-homothetic educational expenditure, the share of overall income spent on housing may not decrease with income even under a positive subsistence requirement. However, the share of income, net of educational spending, devoted to housing does decrease with income in the usual fashion. Indeed, $\beta^4_l < 0$ increases in magnitude, the larger the housing subsistence requirement $\bar{l}$. In turn, parents' altruism share β determines their level of spending on their children's human capital and thus varies the mean educational expenditure share.

FIGURE 9. MOBILITY IN SLUMS: DATA VS. MODEL



Notes: Figure displays, in Panel (A), the population-weighted mean, and in Panel (B), the kernel density across neighborhoods, of the expected income rank of children born in a neighborhood with parents at the 25th income percentile, for slums and all formal neighborhoods separately, in the data and as predicted by our calibrated model.

VII. ACCOUNTING FOR SOCIAL MOBILITY

To validate our model, we return to the key patterns in urban social mobility discussed in Section IV and compare the corresponding model quantities with their empirical counterparts for the Rio de Janeiro metro area. Most crucially, our model must match the distributions, across slums and formal neighborhoods, of the expected adult income rank of children born to parents at the 25th income percentile if we hope to use it to then interpret these patterns. Recall that slums have essentially equal mobility to formal neighborhoods, on average. Despite the fact that none of our model parameters vary based explicitly on slum vs. formal status, the model matches this lack of a slum effect, as demonstrated in Figure 9a. Furthermore, the model replicates the distribution, in particular the wide spread, of mobility across slums in Rio plotted in Figure 9b.

Furthermore, the relationships between slum characteristics and poverty trap status turns out strikingly similar in our model equilibrium. Table 11, Column 1 replicates the poverty-trap-prediction regression, estimated in the data across all metro areas, from Section IV; Column 2 then displays the results of the same regression in our model equilibrium across slum neighborhoods in Rio. The model replicates both the negative relationship between distance segregation and the likelihood that a slum offers poor mobility – even though distance segregation does not directly enter our model. As in the data, slums with superior schools, as measured by the IDEB Education Index, are much less likely to be poverty

TABLE 11. POVERTY TRAP PREDICTORS: DATA VS. MODEL

	Data	*Model*
	(1) Pr(Poverty Trap)	(2) Pr(Poverty Trap)
Assault Hospitalizations per 100k	0.00416***	-0.142***
	(0.000640)	(0.0455)
Log Distance Segregation	-0.121***	-0.280**
	(0.0298)	(0.119)
Infrastructure Index	-0.0237	
	(0.130)	
Informal Employment Share	0.506***	0.262
	(0.162)	(0.356)
Log IDEB Education Index	-1.551***	-2.967***
	(0.149)	(0.240)
Log Commuter Market Access	-0.0203	0.00307
	(0.0133)	(0.0647)
Observations	1,268	211
R-squared	0.549	0.483

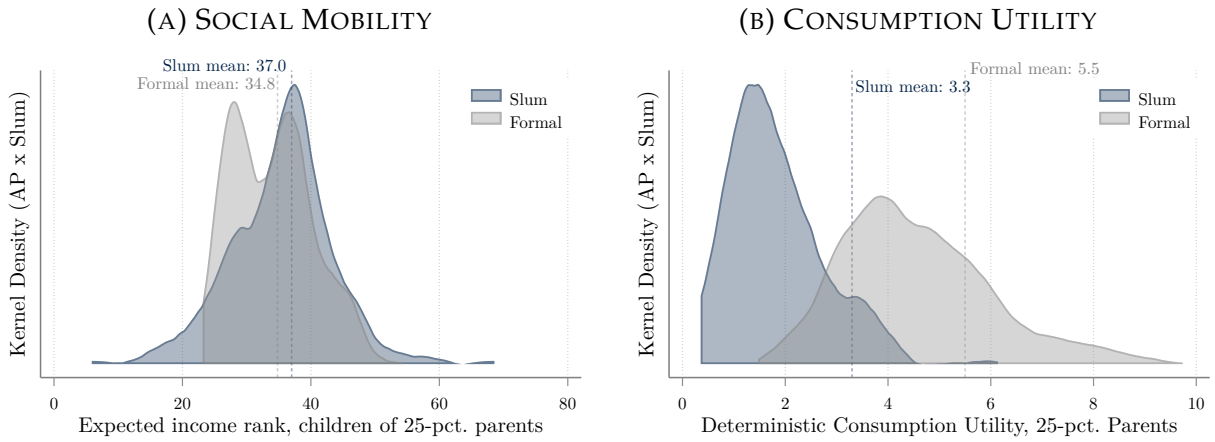
Notes: Robust standard errors in parentheses; *** indicates p<0.01, ** p<0.05, and * p<0.1. Table presents regressions, across slum neighborhoods (census weighting area by slum status), of a dummy for a slum having a child expected income rank, given parents at the 25th income percentile, in the bottom tertile of slums ("poverty trap") on neighborhood covariates. Column (1) runs the regression in the data across all metro areas, and Column (2) runs the regression in the model for Rio de Janeiro only, omitting the infrastructure index, which does not enter the model. We weight all regressions by neighborhood population, and all specifications include a constant.

traps.

VIII. MOBILITY, WELFARE, AND POVERTY TRAPS IN GENERAL EQUILIBRIUM

In this final section, we move beyond empirical measurement of social mobility and leverage our estimated model for two purposes. First, we compare lean on the structure of the model to compare not only income mobility but also welfare across slums and formal neighborhoods. Second, we decompose the determinants of our "poverty trap" slums, taking into account equilibrium adjustments in human capital investment, rents, and location choices.

FIGURE 10. MOBILITY VERSUS WELFARE IN SLUMS



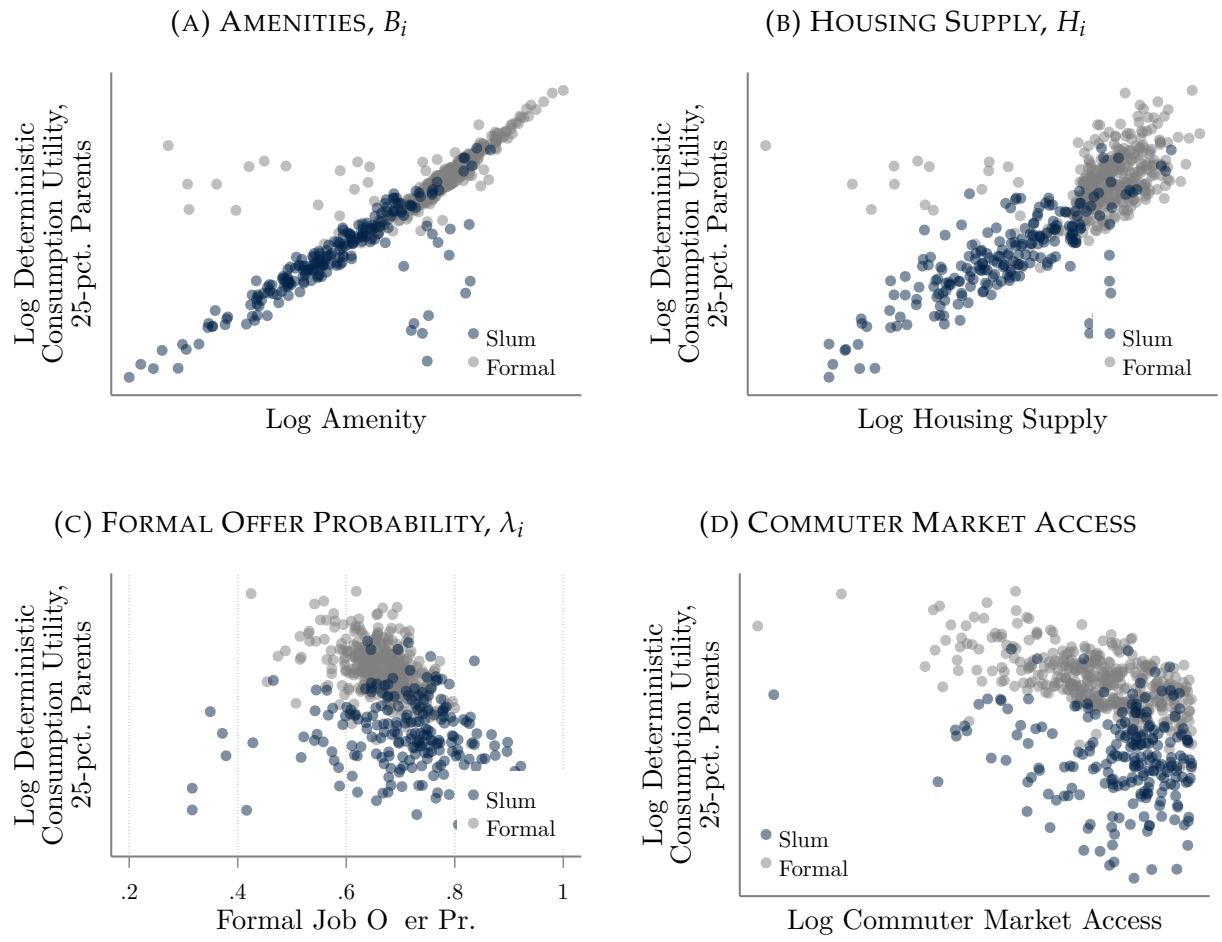
Notes: Figure displays, in Panel (A), the kernel density across neighborhoods, of the expected income rank of children born in a neighborhood with parents at the 25th income percentile, for slums and all formal neighborhoods separately, as predicted by our calibrated model. Panel (B) instead displays the distribution of deterministic consumption utility of parents at the 25th percentile, where we define consumption utility as $u_i(c,l)^{1-\beta}$, i.e. parents' utility net of idiosyncratic shocks and the utility they derive from their child's expected welfare.

Mobility Versus Welfare

Slums may offer their children no worse opportunities, but our model, which crucially allows us to calibrate local amenities, reveals that slum residents do incur a substantial penalty to the common component of their utility. In particular, in Figure 10a, we again plot the distribution of expected 25th-percentile child income rank across slum and formal neighborhoods, as predicted by our model. In Figure 10b, we instead plot the distribution of deterministic consumption utility $u_i(c,l)^{1-\beta}$, i.e. 25th-percentile parents' utility net of idiosyncratic shocks and the utility they derive from their child's expected welfare. Not only do slum residents enjoy only 40% of the deterministic utility of formal residents, on average, but the slum distribution lies so far to the left that the lion's share of slums offer lower utility than even the worst formal neighborhood.

Why do slums offer such substantially worse utility to the "parents" living there? In Figures 11a-11d, we display scatterplots of (log) deterministic consumption utility of 25th-percentile parents against home-neighborhood characteristics and distinguish between slums and formal neighborhoods. Consider first amenities: amenities covary strongly and positively with utility, and slums evidently offer lower amenities. Housing supply similarly increases consumption utility via lower rents, yet slums have, on average, smaller numbers of housing units. On the other hand, formal job offer probabilities and commuter

FIGURE 11. DETERMINANTS OF LOWER SLUM WELFARE



Notes: Figure displays scatterplots of the deterministic consumption utility of parents at the 25th percentile as implied by our model versus (home) neighborhood characteristics, as noted. We define consumption utility as $u_i(c,l)^{1-\beta}$, i.e. parents' utility net of idiosyncratic shocks and the utility they derive from their child's expected welfare.

market access display no clear relationship to consumption utility across neighborhoods and do not differ as substantially between slums and formal neighborhoods. Thus, our model points towards an intuitive conclusion: both lower amenities and a more limited housing supply contribute to substantially inferior utility in slums.

Why, then, does anyone stay in Rio's low-amenity, high-rent slums? To rationalize these location choices, we must consider the idiosyncratic home location preferences, which we have assumed to follow a Frechet distribution. The Frechet distribution famously generates a knife-edge case, whereby, gross of the utility they derive from their children's welfare and gross of all idiosyncratic shocks, residents with a given human capital level are equally well off in each neighborhood. Slum-dwellers, then, willingly accept substantially worse

common utility because they possess such strong idiosyncratic tastes for their homes. The latter represent not only social ties and networks but also heterogeneity in preferences, e.g. intergenerational altruism. Thus, slums offer their residents a distinct bundle: much lower amenities combined with valuable social ties and no-worse opportunities for their children.

What Makes a Poverty Trap?

Section IV already provided a first indication of the factors which create poverty trap slums; we now make use of our model to study these determinants in equilibrium. In contrast to the Oaxaca decomposition in Section IV, we now take into account the feedback between changes in neighborhood characteristics, housing rents, and human capital investment. In particular, we focus on four neighborhood characteristics: (i) crime rates, i.e. assault hospitalizations in $\boldsymbol{x}_i$; (ii) school quality, or the IDEB Education Index also in $\boldsymbol{x}_i$; (iii) formal job offer probabilities λ_{ij}; and (iv) commute costs to formal jobs, i.e. commute distance Δ_{ij}. For one characteristic at a time, we set the corresponding model object for a poverty-trap slum equal to its average for stepping-stone slums, for poverty traps where such a change represents an improvement.[12] We then recompute the resulting equilibrium and difference, between poverty traps and stepping stones in the expected income rank of children with parents at the 25th percentile.

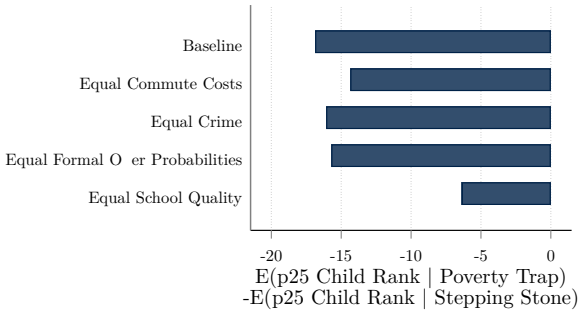
Figure 12a displays the resulting poverty trap-stepping stonne mobility differentials, starting with the difference in expected child rank in our baseline model as estimated. In an echo of our empirical decomposition, we find that, even taking into account equilibrium effects, equalization of school quality would substantially close the gap in mobility between poverty traps and stepping stones. Indeed, these school quality improvements would translate the entire distribution of poverty trap mobility to the right, as shown in Figure 12b. Elimination of differences in commute costs, i.e. distances, formal job offer probabilities, or crime between poverty traps and stepping stones, on the other hand, would each do little to improve mobility.

Important general equilibrium adjustments underlie this closure of the poverty-trap-gap in mobility. In Figure 13a, we plot the change in neighborhood rents from the baseline to the equalized-school-quality model; rents directly capitalize these educational improvements and rise most acutely in the poverty traps which previously had the lowest school quality. Nonetheless, parents in these same poverty traps, whose school quality previously fell

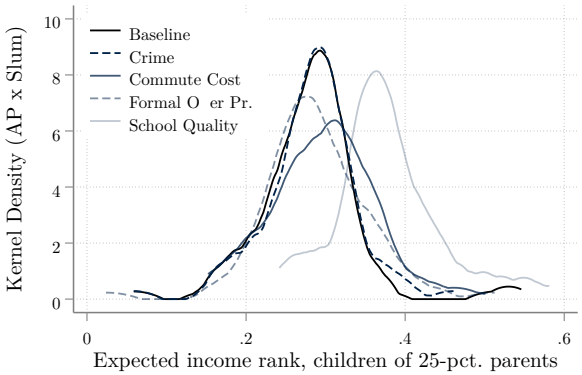
[12]When individual poverty trap slums offer superior values of a characteristic, we leave these values unchanged.

FIGURE 12. ACCOUNTING FOR POVERTY TRAPS IN GENERAL EQUILIBRIUM

(A) POVERTY TRAPS VS. STEPPING STONES



(B) MOBILITY ACROSS POVERTY TRAPS

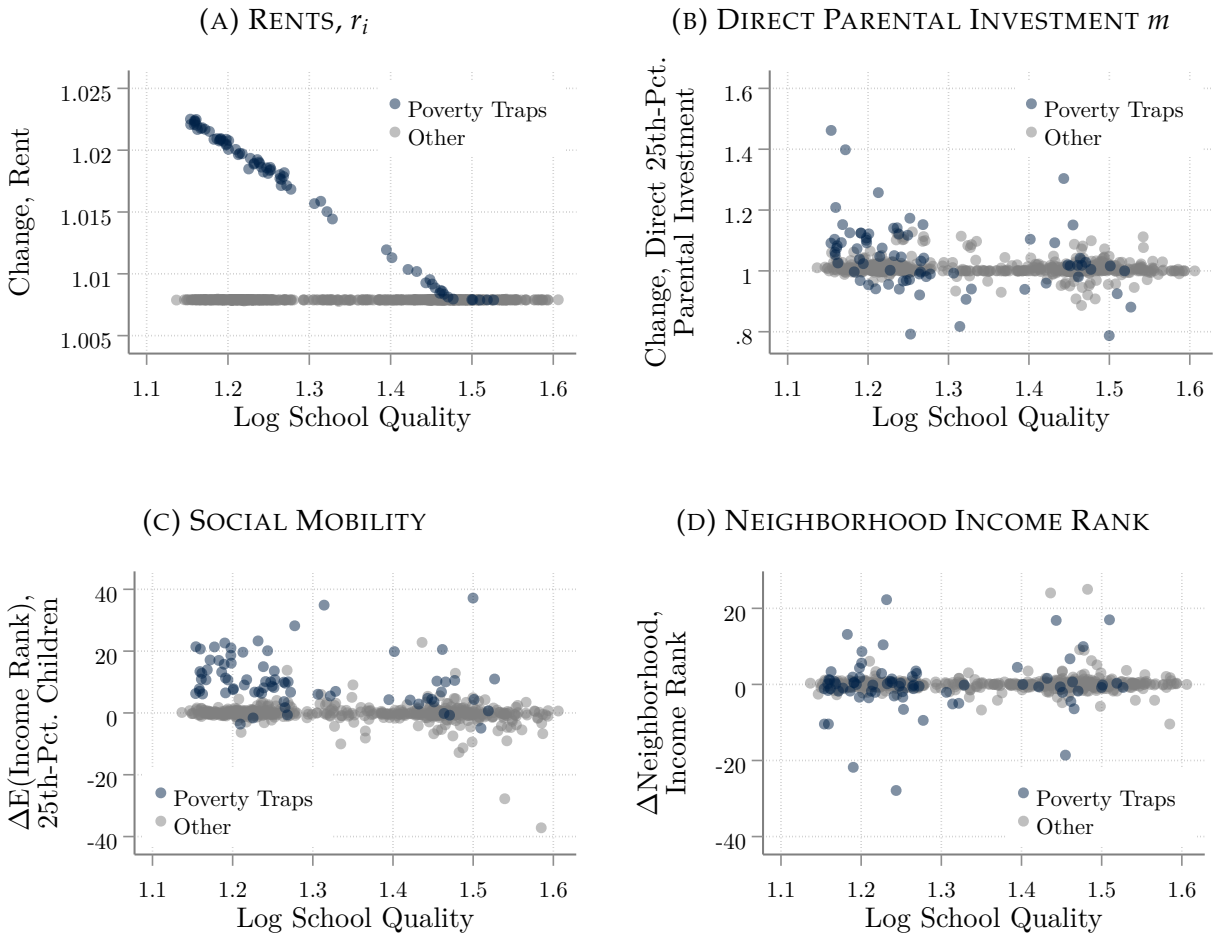


Notes: Figure displays, in Panel (A), the difference between poverty-trap and stepping-stone slum neighborhoods in the expected income rank of children with parents at the 25th percentile of the citywide income distribution, in our baseline model and under different scenarios where, for all poverty trap slums, we set a given neighborhood characteristic equal to the maximum of its current value or the average value among stepping-stone slums. Panel (B) displays the resulting distributions of expected 25th-percentile child mobility across poverty-trap slums in each scenario.

at the bottom of the spectrum, typically increase their direct investments to complement the better schools (plotted for parents at the 25th income percentile in Figure 13b). The combined boosts to school quality and parental investment then augment the expected income rank of children born at the 25th percentile in virtually every poverty trap, particularly in those which previously had the poorest schools, as Figure 13c demonstrates. However, perhaps due to offsetting effects of the compensating differentials in rents and the schooling improvements, the income ranks of poverty traps as compared to other neighborhoods do not systematically change, as plotted in Figure 13d. Thus, these school quality improvements generate little in the way of "gentrification."

In sum, the empirical Oaxaca decomposition of the poverty trap-stepping stone differential and this model-based exercise generate qualitatively similar results, namely that school quality makes poverty traps. However, the model-based decomposition additionally accounts for the parental investment response, which serves to magnify the effect of school quality improvements, as well as changes in housing rents, which would tend to depress investment and might push poorer households out of these newly-improved neighborhoods. In Rio de Janeiro, then, the investment complementarities outweigh the rent increases, and, as a result, upgrading the schools in poverty-trap slums would go a sizeable way towards establishing them as stepping stones to opportunity.

FIGURE 13. SCHOOL QUALITY EQUALIZATION: MARGINS OF ADJUSTMENT



Notes: Figure displays scatterplots across home neighborhoods in the model of the changes in equilibrium outcomes from the baseline model to the model where poverty trap slums' school quality is equalized to the mean among stepping-stone slums, if the latter is higher. In each plot, we distinguish between poverty trap slums and all other slum or formal neighborhoods.

IX. CONCLUSION

To date, a lack of data has precluded a clear understanding of whether slums in developing countries serve as poverty traps or stepping stones to opportunity. We merge a range of administrative datasets covering formal income, family links, and residential locations in urban Brazil and combine them with a random forest methodology to predict informal income. As a result, we provide novel estimates of intergenerational mobility at the level of individual slums. On average, slums turn out to offer no worse intergenerational mobility than formal neighborhoods due to offsetting effects of crime and distance segregation, versus school quality and formal job access. These averages mask substantial across-slum

heterogeneity; the "poverty trap" slums tend to have higher crime as well as informal employment and lower-quality schools. We then introduce a model of intergenerational human capital investment to separate welfare from mobility and study the determinants of poverty traps in equilibrium. Our model, which generates rich within-neighborhood heterogeneity in human capital investment, reveals that slums, despite the similar opportunities they offer children, have substantially lower amenities. In turn, school quality proves a central determinant of whether slums function as poverty traps or stepping stones to opportunity.

REFERENCES

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf. 2015. "The Economics of Density: Evidence from the Berlin Wall." *Econometrica* 83 (6): 2127–2189.

Athey, Susan, Julie Tibshirani, and Stefan Wager. 2019. "Generalized random forests." *The Annals of Statistics* 47 (2): 1148–1178. https://doi.org/10.1214/18-AOS1709. https://doi.org/10.1214/18-AOS1709.

Banerjee, Abhijit, Rohini Pande, Yashas Vaidya, Michael Walton, and Jeff Weaver. 2012. "Delhi's slum-dwellers: Deprivation, preferences and political engagement among the urban poor." *London School of Economics, London: International Growth Centre Working Paper*.

Barza, Radu, Edward L. Glaeser, César A. Hidalgo, and Martina Viarengo. 2024. "Cities as Engines of Opportunities: Evidence from Brazil." NBER Working Paper 32426.

Britto, Diogo, Alexandre Fonseca, Paolo Pinotti, Breno Sampaio, and Lucas Warwar. 2022. "Intergenerational Mobility in the Land of Inequality." *Revise & Resubmit REStat*.

Cavalcanti, Tiago, Daniel Da Mata, and Marcelo Santos. 2019. "On the determinants of slum formation." *The Economic Journal* 129 (621): 1971–1991.

Chaves, Júlio César, Moacyr A.H.B. da Silva, Ricardo de Souza Alencar, Alexandre G. Evsukoff, and Vinícius da Fonseca Vieira. 2023. "Human mobility and socioeconomic datasets of the Rio de Janeiro metropolitan area." *Data in Brief* 51:109695.

Ferreira, Pedro Cavalcanti, Alexander Monge-Naranjo, and Luciene Torres de Mello Pereira. 2024. "Of Cities and Slums." *Forthcoming, Journal of Political Economy*.

Firpo, S. P., N. M. Fortin, and T. Lemieux. 2018. "Decomposing wage distributions using recentered influence function regressions." *Econometrics* 6 (2): 28.

Gechter, Michael, and Nick Tsivanidis. 2023. "Spatial Spillovers from High-Rise Developments: Evidence from the Mumbai Mills." *Revise & Resubmit accepted, Econometrica*.

Glaeser, Edward. 2011. *Triumph of the city: How our greatest invention makes us richer, smarter, greener, healthier, and happier*. Penguin.

Gonzalez-Navarro, Marco, and Raimundo Undurraga. 2023. "Immigration and Slums."

Harari, Mariaflavia. 2024. "Residential Patterns and Local Public Goods in Urban Brazil."

Harari, Mariaflavia, and Maisy Wong. 2024. "Slum Upgrading and Long-run Urban Development: Evidence from Indonesia." *Accepted, Review of Economic Studies*.

Henderson, J Vernon, Anthony J Venables, Tanner Regan, and Ilia Samsonov. 2016. "Building functional cities." *Science* 352 (6288): 946–947.

Henderson, J. Vernon, Tanner Regan, and Anthony J. Venables. 2021. "Building the City: From Slums to a Modern Metropolis." *Review of Economic Studies* 88:1157–1192.

Krishna, Anirudh. 2013. "Stuck in place: Investigating social mobility in 14 Bangalore slums." *The Journal of Development Studies* 49 (7): 1010–1028.

Marx, Benjamin, Thomas Stoker, and Tavneet Suri. 2013. "The Economics of Slums in the Developing World." *Journal of Economic Perspectives* 27 (4): 187–210.

Michaels, Guy, Dzhamilya Nigmatulina, Ferdinand Rauch, Tanner Regan, Neeraj Baruah, and Amanda Dahlstrand. 2021a. "Planning ahead for better neighborhoods: Long-run evidence from Tanzania." *Journal of Political Economy* 129 (7): 2112–2156.

———. 2021b. "Planning Ahead for Better Neighborhoods: Long-Run Evidence from Tanzania." *Journal of Political Economy* 129 (7): 2112–2156.

Rojas-Ampuero, Fernanda, and Felipe Carrera. 2024. "Sent Away: The Long-Term Effects of Slum Clearance on Children."

Tsivanidis, Nick. 2023. "Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio." *Conditionally accepted, American Economic Review*.

Ulyssea, Gabriel. 2010. "Regulation of entry, labor market institutions and the informal sector." *Journal of Development Economics* 91:87–99.

UN. 2024. *The Sustainable Development Goals Report 2024.* Technical report.

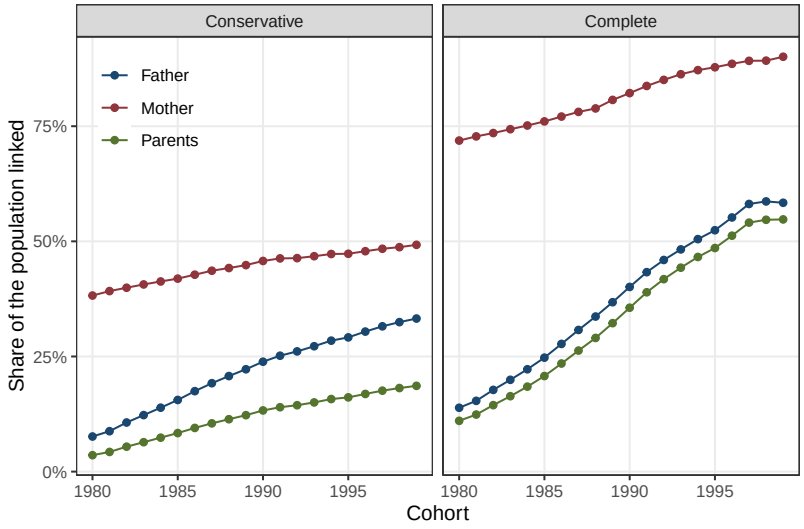
UNCTAD. 2021. *Handbook of Statistics.* Technical report. United Nations Conference on Trade and Development.

Wong, Maisy. 2019. "Intergenerational mobility in slums: Evidence from a field survey in Jakarta." *Asian Development Review* 36 (1): 1–19.

A. Empirical Appendix

A.1 Data Appendix

Figure A.1. Number of parent-child links relative to the population by cohort



Notes: The figure plots the share of the population that can be linked to their parents by cohort following our baseline, more conservative method (left graph) and the alternative, less conservative method (right graph). The first method links children to parents using unique person codes in dependent claims tax data and using names for uniquely named parents in population and welfare registries. The second method allows for additional links using individual names and addresses.

A.2 Facts Appendix

TABLE A.1. DESCRIPTIVE STATISTICS OF MAIN SAMPLE

Metro Area	Slum % (Data, 2000)	Slum % (Census, 2010)
Manaus	6.9%	15.0%
Belem	17.7%	53.9%
Sao Luis	5.7%	24.5%
Teresina	3.6%	13.4%
Fortaleza	7.1%	11.9%
Natal	0.5%	6.0%
Joao Pessoa	3.6%	8.5%
Recife	4.8%	23.2%
Maceio	1.7%	10.6%
Aracaju	2.8%	9.8%
Salvador	4.7%	26.1%
Belo Horizonte	5.5%	9.1%
Vitoria	0.9%	10.6%
Rio de Janeiro	4.0%	14.4%
Sao Paulo	3.6%	11.0%
Curitiba	3.0%	5.7%
Porto Alegre	2.9%	6.2%
Brasilia	0.7%	3.7%

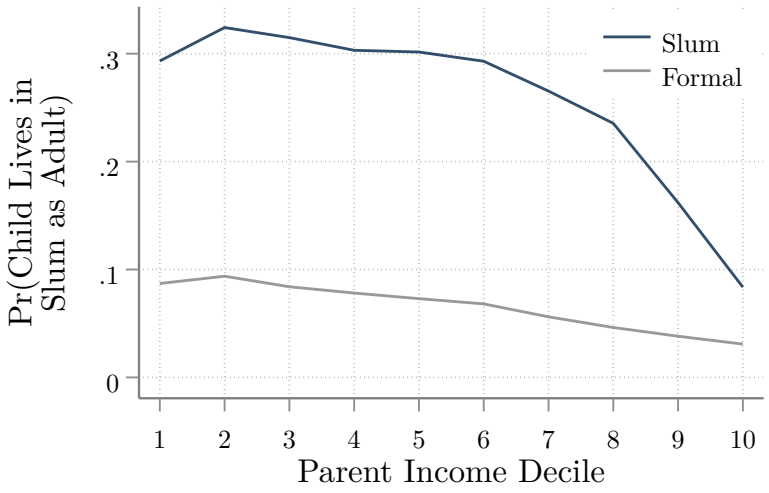
Notes: The table compares the share of individuals living in slums in our sample with the slum population share in the 2010 Census, by metropolitan area. There are several caveats for this comparison. Most importantly, we are targeting individual's location in the year 2000, a decade before the Census. Also, the Census definition of the boundaries for each metropolitan area is slightly different than the one we use throughout the paper.

TABLE A.2. SLUM EFFECT ON OTHER OUTCOMES (CHILDREN OF 25TH-PERCENTILE PARENTS)

	(1) Income Rank	(2) Live in Slum	(3) Graduate High School	(4) Informal Employment
Slum Dummy	-0.348	0.180***	-0.0344***	0.00722
	(0.463)	(0.00827)	(0.00544)	(0.00633)
Constant	37.69***	0.0796***	0.600***	0.487***
	(0.225)	(0.00305)	(0.00279)	(0.00250)
Observations	1,966	1,966	1,965	1,966
R-squared	0.000	0.318	0.021	0.001

Notes: Robust (?) standard errors in parentheses; *** indicates $p<0.01$, ** $p<0.05$, and * $p<0.1$. Regressions at AP-by-slum level.

FIGURE A.2. PROBABILITY OF CHILDREN LIVING IN SLUM AS ADULT, BY PARENT INCOME AND CHILDHOOD SLUM RESIDENCE



Notes: Figure displays the share of children living in a slum as adults, by parent income decile and childhood neighborhood type, either slum or below-median-income formal neighborhood. We take deciles across the population of our sample of metro areas.