In-situ Upgrading or Population Relocation? Direct Impacts and Spatial Spillovers of Slum Renewal Policies*

Paul Gertler
University of California Berkeley and NBER

Marco Gonzalez-Navarro
University of California Berkeley

Raimundo Undurraga
Universidad de Chile

Joaquin A. Urrego
Wake Forest University

February 10, 2026

Abstract

Slums are widely viewed as disadvantaged areas characterized by substandard housing and negative neighborhood spillovers that constrain local economic development. Governments typically respond with one of two place-based interventions: (a) *in-situ* upgrading that improves infrastructure and housing conditions within existing settlements, or (b) relocation of slum households to formal housing elsewhere in the city, freeing valuable urban land for alternative uses. We compare the direct effects and local spatial spillovers of these two approaches in Chile. We construct a 20-year national panel of slum polygons that integrates satellite imagery, census microdata, construction permits, administrative records, crime reports, and property-tax data. Using this dataset, we estimate how each policy affects the physical characteristics of slums and adjacent areas, and whether either intervention shifts the socioeconomic composition of residents. Our empirical strategy relies on Synthetic Difference-in-Differences for causal identification. We find that both policies reduce the share of land devoted to housing within the original slum perimeter — *in-situ* upgrading by reallocating land toward neighborhood infrastructure, and population relocation by reducing the number of slum households and housing structures. However, only *in-situ* upgrading generates durable improvements in housing quality and the socioeconomic status of residents. In addition, upgrading produces substantial positive spillovers in nearby neighborhoods, increasing formal housing investment and reducing crime. Moreover, administrative cost data show that *in-situ* upgrading is roughly one-third less expensive per household than population relocation, underscoring its greater cost-effectiveness.

Keywords: Cities, Urbanization, Slums, Slum Policy, Spatial Spillovers
JEL Codes: R14, R31, R38, R50

*Contact: gertler@berkeley.edu; marcog@berkeley.edu; raimundo.undurraga@dii.uchile.cl; urregoj@wfu.edu. We thank Matias Labbé for his superb research assistance and coordination with the Chilean government and TECHO. We thank seminar participants at UC-Berkeley, RIDGE, UEA, Wharton, NBER SI Urban Economics Group, for their comments. We acknowledge funding by MIGRA Millenium Nucleus under grant ANID-MILENIO-NCS2025_69. Undurraga acknowledges funding by Fondecyt Regular ID 15894 Folio 1230510. All errors are our own.

1 Introduction

Slums continue to be a defining feature of urban life in low- and middle-income countries, as one in four urban residents —more than a billion people— live in slum conditions (UN-Habitat, 2020). Slums are generally viewed as spaces of disadvantage, characterized by substandard housing and inadequate access to essential services such as clean water and sanitation, and as sources of negative externalities for nearby neighborhoods, limiting local economic development. Governments often respond by implementing place-based policies aimed at fostering development in slums and reducing housing inequality within cities. In designing such interventions, governments typically consider two options: upgrading slums in place — by improving neighborhood infrastructure and housing conditions — or offering to relocate slum households to formal housing elsewhere in the city, freeing valuable urban land for alternative uses, albeit at the cost of uprooting established slum communities.

The central question for policymakers, then, is whether to improve slums where they stand or to replace them entirely through relocation. *In-situ* upgrading programs provide slum dwellers with basic services such as land titling, safe water, electricity, and sanitation infrastructure, and in many cases also improve the housing stock itself, all without displacing residents. This approach aligns with efforts to protect informal dwellers from clearance, ensuring they benefit from progressive formalization and urban development rather than being displaced by it. As Jane Jacobs asserted "*The best way to plan for downtown is to see how people use it today; to look for its strengths and to exploit and reinforce them*" (Jacobs, 1961, p.238).

By contrast, population relocation strategies first offer to resettle households in formal housing located elsewhere and then repurpose the vacated land for other urban uses such as commercial development, public infrastructure, or green and recreational spaces. While relocation may be necessary when slums are located in hazard-prone areas such as flood zones, in most settings *in-situ* upgrading remains a viable alternative, giving urban planners meaningful discretion between the two approaches.

As rapid urbanization continues to strain housing markets and infrastructure, credible evidence on the cost-effectiveness of these policies becomes increasingly urgent. In this paper, we help fill this gap by examining the direct and spillover effects of slum renewal policies in the Chilean context, where both *in-situ* upgrading and population relocation interventions have been implemented simultaneously since 2011. One of the primary challenges in generating rigorous evidence on slum policy is the absence of high-quality, representative data that track the evolution of slums over time and space (Marx et al., 2013). A central contribution of this study is to assemble a high-frequency panel of the universe of slum areas that spans more

than two decades. The starting point is the set of georeferenced polygons that were ever officially designated as slums at any point in time between 2011 and 2021. Each year, active slums were identified by the Chilean government based on three conditions: (i) an area that is home to at least 8 spatially close households, (ii) households do not have property rights, and (iii) houses are of substandard quality and lack access to at least one of the basic services (water, sanitation, and/or electricity). Since slum boundaries change over time, we define each slum's area as the combined footprint of all its observed polygons across years.

We used historical Google Earth satellite images to track the evolution of land use within these polygons beginning in 2000 and follow them up to 2021. Using mask image-processing CNN algorithms, we identify built structures inside slums and up to 500 meters just outside slums. Following Michaels et al. (2021), we measure physical characteristics such as building area, building density, similarities in orientation angle, and proximity between buildings. We merge building investment information from georeferenced building permits and property tax records for all formal housing in surrounding formal neighborhoods adjacent to slums. We also merge socio-economic information on residents from population censuses, as well as information on criminal activity using geocoded administrative crime records.

The data provide a nuanced picture of the formation and evolution of slums that is consistent with several of the central predictions in the theoretical literature.[1] When slums appear, they form overwhelmingly at city boundaries in areas situated at higher elevations, on steeper slopes, in locations with greater overall ruggedness, closer to rivers and closer to firms that employ low-skill workers. Slums on average have worse access to public amenities such as schools, supermarkets, and public safety services. Initially, slums are populated with low-educated single males. However, over time the socio-economic characteristics of slum residents (e.g., education, gender, and marital status) converge to those of their nearby non-slum counterparts. Slum populations grow faster in cities where quality-adjusted housing rents increase rapidly, but wages for low-skilled workers do not keep up.

In this setting, our main objective is to evaluate and compare the impacts of *in-situ* upgrading and population relocation policies on living conditions in slums and spillovers to adjacent formal neighborhoods. Importantly, both approaches are voluntary —participation by slum residents is not mandated— and both aim to facilitate the transition of informal

[1]The literature has identified multiple structural determinants of slum formation, including weak enforcement of property rights, underdeveloped capital markets, and low levels of human capital (Collier and Venables, 2015; Brueckner and Lall, 2015; Duranton and Venables, 2020). For models of slum formation, see Jimenez (1984), Duranton (2008), Brueckner and Selod (2009), Duranton and Puga (2015), Cavalcanti et al. (2019), Henderson et al. (2021), Duranton and Puga (2023) and Cavalcanti et al. (2025). For studies focusing on how local and market-wide economic conditions (including labor and housing markets) interact with the formation and evolution of slums, see Glaeser (2011), Marx et al. (2013), Wong (2019), Cavalcanti et al. (2019), Henderson et al. (2021), Alves (2021), and Gechter and Tsivanidis (2023).

settlements into formal urban areas.

Both policies began implementation in 2011 by the Chilean Ministry of Housing and Urban Development (MINVU), but were rolled out gradually across slum areas. We use Synthetic Difference-in-Differences to estimate causal effects of each policy (Arkhangelsky et al., 2021; Clarke et al., 2023). By tracking whether and when each slum was treated, the SDiD estimator combines Synthetic Controls methods (Abadie and Gardeazabal, 2003; Abadie, 2021) with a Difference-in-Differences approach that accounts for differential timing of treatment. The control group for each treatment is chosen from the set of slums that have not yet been treated, and SDiD assigns an optimal weight to each potential control, so that the control group looks similar to treated group in terms of the pre-intervention trends. Since both *in-situ* upgrading and relocation projects typically take between two and four years from assignment to implementation, we measure policy impacts six years after treatment assignment to allow sufficient time for effects to materialize.

We begin our analysis by inquiring whether the two policies achieve their respective stated objectives. Recall that population relocation seeks to reduce the number of people living in a slum perimeter to zero by offering formal housing outside the slum perimeter. Our data shows that six years after a slum polygon is slated for relocation, the population living within the slum perimeter is only 16 percent lower than it was originally, a level far from its stated objective and one which would preclude redevelopment of the land towards other uses. Consistent with this finding, there is a concomitant 12 percent average reduction in the amount of land used for residential purposes. This result is likely due to the program's voluntary nature — many residents choose not to leave, often due to social ties or proximity to employment or households may decide to split to keep possession of both properties — and partly due to rapid repopulation of vacated houses when closure measures like barriers are not effectively implemented.

Recall that *in-situ* upgrading's objective is to minimize depopulation while improving living conditions within slums. In this case, our results point to the policy achieving its stated objective of minimizing population displacement. Indeed, we cannot reject that average changes in total polygon population differ from zero six years post treatment designation, consistent with the program's objective of enabling current residents to stay in place.

We then turn to investigating the effects of the two policies on living conditions within slum perimeters. Across a variety of measures, our results point to *in-situ* upgrading having significant positive effects on housing quality and available infrastructure. Among others, the share of paved streets in upgraded slums is double relative to that in synthetic control slums. Upgraded sites are more likely to have new formal housing and improvements in existing structures. Housing units in upgraded slums are also more likely to be larger and have two or

more bedrooms. *In-situ* upgrading also generates more regularly oriented structures that are spaced further apart, suggesting coordinated formalization. Finally, upgraded slums attract a better educated population, an important indicator of a more desirable neighborhood.

In marked contrast, we find no statistically significant effects of population relocation on any of the housing quality or demographic composition measures in intervened slum areas compared to synthetic controls. We interpret this result as indicating that slum upgrading achieves its objective of creating more desirable neighborhoods, while population relocation generates no meaningful difference in conditions for those living within the slum perimeter.

After establishing the effects of slum policies within slum perimeters, we turn to investigating spillovers onto neighboring non-slum areas. This is an important and until now overlooked aspect of slum policies, since surrounding formal neighborhoods often have more political voice than slum-dwellers and can thus influence which policies get implemented.

We find that *in-situ* upgrading has strong positive spillover effects in adjacent neighborhoods across a range of outcomes. Specifically, we find increases in investment in both new and existing housing. New housing starts are 10 percent higher and existing buildings are more likely to have undergone renovation compared to neighborhoods surrounding synthetic control slums. These improvements in housing attract higher-SES residents in terms of education and employment. Finally, we find a significant reduction in both violent and property crimes. On average, there was 11 percent less property crime, 21 percent less violent crime, and the number of homicides dropped precipitously. To our knowledge, this is the first documentation of such positive crime externalities arising from *in-situ* upgrading. This evidence is consistent with the positive direct impacts evidence for *in-situ* upgrading. If the slum perimeter becomes a more desirable neighborhood, it is natural for surrounding non-slum neighborhoods to also be perceived as safer and more desirable.

For population relocation policies, we find no economically nor statistically significant changes in formal neighborhoods adjacent to treated slums compared to synthetic control slums. This result is consistent with the direct impacts finding that being designated a population relocation slum did not generate important changes in the living conditions of people residing in the slum polygons. It is thus unsurprising that spillover effects were unimportant in population relocation neighborhoods.

Taken together, the evidence on direct impacts and spillover effects provides strong support for *in-situ* upgrading as an effective strategy to promote the development of more desirable neighborhoods without displacing slum populations. In contrast, the evidence is quite clear that population relocation programs do not achieve their objectives of depopulating slum areas. As a corollary, because investment in the slum polygon is very limited, living conditions do not differ from those in control slums, and as a consequence, spillovers

are non-existent. These findings are especially striking given that administrative expenditure data showing that average government expenditure per treated slum household is approximately one third lower for *in-situ* upgrading relative to population relocation. We believe these findings are novel and can inform policymakers when making decisions about which policy to implement.

The analysis so far has compared each policy to its own counterfactual. However, these treatment effects may not be directly comparable because the group of untreated slums selected as relevant counterfactuals can differ across the two policy arms. For example, *in-situ* upgrading may be less likely to occur in informal settlements located in areas highly exposed to natural hazards, such as floodplains or landslide-prone zones. To assess whether the two interventions generate different impacts when applied in comparable contexts, we conduct a head-to-head comparison that directly contrasts *in-situ* upgrading with relocation. We implement an SDiD estimator that constructs the synthetic counterfactual not from untreated slums, but from slums that received the alternative policy in the same year and exhibit overlapping predicted assignment probabilities, thereby ensuring that comparisons are made within a region of common support. This subsample represents a large share of all slums, indicating substantial overlap between the groups. The results are consistent with our earlier findings: even when evaluated in comparable contexts, *in-situ* upgrading generates larger improvements in neighborhood quality than population relocation.

This article contributes to the slums literature by evaluating the relative effectiveness of the two main remedial policies commonly implemented worldwide to address slum growth. It provides empirical evidence on both the direct impacts of each policy on treated slum polygons and the spatial spillovers they generate in surrounding areas. A key input for this analysis is a national high-resolution yearly panel of slum areas spanning more than two decades, which requires methodological advances in slum detection and spatial analysis. Our analysis builds on recent advances in remote sensing in economics that have led to powerful analytical tools to map spatial development patterns (Burchfield et al., 2006; Henderson et al., 2012; Michalopoulos and Papaioannou, 2013; Henderson et al., 2021), and particularly innovative methods for identifying and delineating the spatial extent of slums (Harari and Wong, 2025; Gechter and Tsivanidis, 2023; Abascal et al., 2022; Najmi et al., 2022; Fisher et al., 2022).

Previous work has studied the impact of both slum renewal policies, but only individually, such as Dasgupta and Lall (2009); Collins and Shester (2013); Gonzalez-Navarro and Quintana-Domeque (2016); Barnhardt et al. (2017); Galiani et al. (2017); Bazzi et al. (2019); Picarelli (2019); and Franklin (2020).[2] Overall, none of the existing papers provide direct

[2]For a thorough review on relocation policies, see Belsky et al. (2013).

comparison of the relative effectiveness of these policies nor the spillover effects associated to them. Our work complements studies on non-remedial preventive policies aimed at avoiding the emergence of slums through infrastructure planning (Michaels et al., 2021; Henderson et al., 2025) and land titling programs (Field, 2007; Galiani and Schargrodsky, 2010).

Our paper is closely related to Harari and Wong (2025), who combine high-resolution policy maps, administrative data, and an innovative photographic survey of formal and informal housing markets to study the long-term impact of Jakarta's Kampung Improvement Program (KIP), one of the world's largest slum upgrading initiatives. After four decades of exposure, the authors find that areas targeted by KIP became more informal, exhibited lower land values, and contained fewer high-rise buildings, with the evidence being robust across multiple localized comparison methods. In other words, although KIP neighborhoods experienced improvements, many of them persisted as slums while comparable non-KIP areas transitioned to formality, suggesting that the *in-situ* approach may have delayed rather than accelerated formalization.

Importantly, while KIP substantially improved access to basic public services, it did not formalize property rights through ownership titles, a feature that may help explain the delayed formalization observed in KIP neighborhoods. This distinction likely explains why KIP's direct effects contrast with those observed under Chile's *in-situ* upgrading program, which *de facto* formalized property rights by providing beneficiaries with land titles along with formal housing. In fact, under MINVU's administrative procedure, the government re-classifies treated polygons as non-slum after the intervention since inhabitants obtained formal property rights and access to all basic services. This comprehensive formalization likely accounts for the strong direct effects observed in our setting, in addition to spillovers to surrounding areas. This is consistent with Henderson et al. (2021)'s model, which predicts that overcoming institutional frictions in land markets that impede the conversion of slum structures into durable housing can spur redevelopment of formal neighborhoods over time.

Another closely related study is Rojas-Ampuero and Carrera (2023), which analyzes an involuntary slum clearance program implemented in Santiago, Chile during the 1980s. Unlike modern voluntary slum policies, the Chilean dictatorship forcibly evicted slum households and relocated them to formal housing projects elsewhere, while a subset of households received improved dwellings through *in-situ* upgrading. The authors show that children from relocated households experienced significantly lower educational attainment and earnings relative to those who remained in their original locations and benefited from improved conditions.[3]

[3]For evidence on the effects of forced displacement, see the comprehensive review by Becker and Ferrara (2019). The impacts of forced displacement are highly context-specific and do not uniformly translate into adverse outcomes; in many settings, the evidence is mixed. For instance, Bryan et al. (2025) find that slum

Our findings are not directly comparable by design to those from Rojas-Ampuero and Carrera (2023), but they are complementary in several respects. First, whereas the authors study individual-level outcomes — specifically, child human capital — our analysis considers spatial outcomes at the level of slum polygons. The unit of analysis is thus fundamentally different: we examine how policies shape neighborhood conditions, built environment quality, and spillovers into surrounding non-slum areas. Second, their setting contrasts *in-situ* upgrading with *forced* relocation during an authoritarian regime, a context in which displaced households had no choice in the intervention they received. By contrast, our study evaluates slum policies implemented in a *voluntary* framework. Third, their empirical design identifies differences between displaced and non-displaced (on-site upgraded) households, but does not include a separate counterfactual group that would allow estimation of the absolute effects of each policy, nor does it recover spatial spillovers. Our approach addresses both dimensions: it separately identifies the impact of *in-situ* upgrading and relocation policies and quantifies their external effects on neighboring areas.

Finally, our evidence contributes to the longstanding "people versus places" debate over the appropriate target of public policy for vulnerable populations (Glaeser, 2001; Glaeser and Gottlieb, 2008; Neumark and Simpson, 2015; Gaubert et al., 2025; Garin and Rothbaum, 2025). In the standard spatial equilibrium model with perfectly mobile workers and a perfectly inelastic housing supply, assisting individuals directly through transfers is typically preferred to place-based interventions, largely because of the equity concern that the gains from place-based policies accrue to landowners rather than to local residents. This concern is substantially mitigated in our context: *in-situ* slum upgrading in Chile includes the provision of land titles, meaning that the capitalization of public investments remains with slum residents.

Place-based interventions in poor neighborhoods are also often criticized on efficiency grounds if they stimulate activity in relatively unproductive areas at the expense of more productive ones. Yet market failures —such as under-provision of public goods, mobility constraints, or incomplete insurance and credit markets— can make improvements to the locations that low-income households actively choose welfare-enhancing in ways that individual transfers cannot easily replicate. As emphasized by Gaubert et al. (2025) and Kline and Moretti (2014a), place-based policies can unlock agglomeration forces and productivity

dwellers in central Addis Ababa, Ethiopia, who were forced to leave their settlements and relocate to subsidized housing located farther from their original homes (about 14 minutes away), on average, experienced declines in social networks and a higher perception of serious crime relative to stayers but gained higher labor force participation, higher earnings, better housing quality, and greater environmental amenities. Chiovelli et al. (2021) show that, in the context of the Mozambican civil war, internally displaced individuals attained substantially more schooling than their siblings who remained behind, yet they exhibit lower social capital and worse mental health compared with local residents.

spillovers. In our setting, the positive externalities we document for nearby formal neighborhoods —*in-situ* upgrading reduces crime and increases both housing quality and the share of highly educated residents in neighboring areas— represent local multipliers that standard efficiency arguments should not overlook. Moreover, coordination failures can depress measured productivity without implying that the underlying locations are intrinsically low productive; slum households' revealed preferences for proximity to jobs and key urban amenities suggest substantial latent value that becomes tradable once basic frictions such as substandard housing are alleviated.[4] More broadly, our findings highlight that the equity-efficiency trade-off that typically dominates debates on place-based policies is less binding when these interventions target low-income groups such as slum dwellers.

2 Slum Renewal Policy in Chile

2.1 Evolution

Adequate housing for slum dwellers, estimated at 10-20% of the population in the 1960s, has been a long-standing policy issue in Chile (Hidalgo, 2019). The government established the Ministry of Housing and Urban Development (MINVU) in 1965 to address housing conditions.

In the latter half of the 1960s, MINVU launched its first major initiative, Operación Sitio, which provided low-income households with urban plots equipped with basic services —water, electricity, and sanitation— on which beneficiaries could construct their own homes. During President Allende's administration (1970-1973), public housing production expanded substantially, with slum dwellers prioritized for access to new units. Nevertheless, rapid rural-to-urban migration kept housing demand far above available supply, and informal settlements continued to grow.

Under the Pinochet dictatorship (1973-1989) new public housing was often located on the periphery of cities in areas with low land costs but high unemployment and limited access to public services. The average size of public housing units decreased from an average of 60 m^2 (650 sq. ft.) in the 1960s to less than 30 m^2 (320 sq. ft.) in the 1980s. A signature policy of the Pinochet era was the "Program for Urban Marginality" (1979–1985), a large-scale slum

[4] A large literature documents sizable gains for individuals who can move to better neighborhoods, with programs like MTO in the U.S. being a prominent example (Kling et al., 2007; Chetty et al., 2016). These findings suggest that relocating slum residents to high-productivity neighborhoods could, in principle, be more efficient than *in-situ* upgrading. Our findings suggest that even offering outright ownership in formal areas is not enough to entice many households to move away from their slum communities. This low take-up rate of housing subsidies intended to facilitate access to better-located housing has been well documented (Selman, 2025, Barnhardt et al., 2017).

clearance and relocation initiative that affected 340 slums and more than 40,000 residents. The program provided formal housing ownership to slum dwellers, but was implemented through forced relocation. Rojas-Ampuero and Carrera (2023) assess the long-term effects of such displacements. By comparing forcibly relocated households with those who received on-site public housing, they find that displaced children earn 9% less and complete 0.8 fewer years of education on average. They also tend to work in more precarious, informal jobs and are more likely to live in high-poverty neighborhoods as adults.

Following the return to democracy in 1990, the supply of public housing again expanded. For the first time, private firms developed most of the new public housing units financed through demand-side subsidies. The new housing helped contain the expansion of slums, although informal settlements remained a prominent feature of urban landscapes. Between 2010 and 2020, however, slums grew dramatically, partly due to large immigration inflows (Gonzalez-Navarro and Undurraga, 2023). By 2020, there were nearly 1,000 slums, and their population had risen to almost 100,000 —a figure that increased even further after the pandemic.

2.2 Contemporary Policy

In 2011, the Chilean government conducted the first Slum Census with universal geographic coverage, encompassing the entire country (MINVU, 2011). Following the census, MINVU was mandated to address the substandard living conditions of slum dwellers through one of two slum renewal policies: (i) *in-situ* upgrading, which improves infrastructure and housing conditions within slum areas without displacing residents; or (ii) population relocation, which provides slum residents with subsidized ownership of formal housing elsewhere, freeing up the vacated land for alternative urban uses.

According to internal documents from MINVU, the choice between *in-situ* upgrading and population relocation reflects a policy decision shaped by a combination of place-specific factors such as the feasibility of new construction —determined by terrain conditions and exposure to natural hazards— as well as the slum's location within the urban structure, proxied by its access to amenities like public transit, schools, health services, and/or green areas (DIPRES, 2019).[5]

Importantly, in contrast to policies during the Pinochet years, both approaches are voluntary —participation by slum residents is not mandated— and both aim to facilitate the transition of informal settlements into formal urban areas. In practice, MINVU visits each identified slum area to assess local conditions, select an intervention strategy, and confirm

[5]Section 6.1 provides an empirical analysis of the determinants of *in-situ* versus relocation interventions.

that residents agree with the intervention.[6] The process is summarized in Appendix Figure A1.

Given fiscal constraints, program implementation has proceeded incrementally, with the pace of expansion dictated by the annual appropriations authorized by the National Congress. Between 2011 and 2020, MINVU implemented *in-situ* upgrading in 209 slums and population relocation in 440 slums, investing more than US$400 million in total.

2.2.1 *In-situ* Upgrading

In-situ upgrading consists of two components: (i) improving infrastructure and (ii) improving housing conditions, both without displacing existing residents. The first component finances and contracts for the provision of basic infrastructure, including street paving, piped water, sewerage, and electricity. This component is provided at no cost to households.

The second component supports the construction of improved housing in cases where existing dwellings are deemed substandard. This effort is financed through demand-side subsidies, specifically the DS49 voucher program, which is means-tested and targeted to households in the lowest 40% of the national income distribution. Slum residents are provided with technical assistance and guidance throughout the DS49 application process, and all of them are *de-facto* eligible for it. The average subsidy is approximately USD 15,000 per household and beneficiaries are required to contribute an additional USD 800 copay. These funds are used to construct housing units within the slum polygon, often involving the demolition of existing structures and their replacement with a formal housing project. As part of the new construction, land titles are awarded to the beneficiaries.[7] Although the subsidy terms prohibit the sale or rental of housing units for the first five years, intergenerational transfers are permitted. In some cases, MINVU may invite residents from other slums or people in need of housing to join a new housing project.

Detailed cost data for *In-situ* upgrading from MINVU are summarized in Table 1. Throughout the period 2011-2020, a third of intervened slums received *in-situ* upgrading, with an average cost per household of USD $28,636. Over half of slum dwellers received a DS49 housing voucher, i.e., they accepted to stay in the slum and participate in the *in-situ* upgrading. About one-third of the program costs are spent on direct public investment in infrastructure and two-thirds on housing subsidies.

[6] The original list of target slums was established in 2011 and updated in 2018 (DIPRES, 2019).

[7] Implementation tends to be more straightforward when the slum is located on publicly owned land, allowing the government to deal with a single counterpart rather than multiple private owners.

2.2.2 Population Relocation

The population relocation program moves residents out of slums by offering them access to formal housing in non-slum areas. This strategy comprises two primary components. The first involves providing DS49 housing vouchers that recipients can use to acquire private or public formal housing elsewhere in the city. These units must be priced below the eligibility threshold defined by the subsidy program.

A second component of this strategy—though not consistently undertaken—involves direct public investment by MINVU in the original slum area. This can include the demolition of existing structures, debris removal, and, in some cases, the reallocation of the cleared land for alternative public uses such as parks or community facilities. In instances where the slum was located on privately owned land, the land is usually returned to its original owners, who may develop it at a later stage.

An important feature of the population relocation strategy is its voluntary nature: although the intervention is designed and implemented at the slum level, individual households are not mandated to participate. Relocation decisions are made independently by each household, and in practice, roughly one in five choose not to accept MINVU's offer. As a result, some intervened slum areas persist as informal settlements even after relocation efforts, either due to individual noncompliance or subsequent repopulation dynamics. Vacated homes often attract new informal settlers, and in some cases, only part of a household relocates while others remain, thereby perpetuating the slum's existence. Moreover, even fully cleared slums can be rapidly repopulated if closure measures are not effectively planned in advance.

Although population relocation was the most widely implemented intervention throughout the period 2011-2020 –roughly two-thirds of treated slums– it is, on average, one-third more expensive than *in-situ* upgrading. As shown in Table 1, the average per-household fiscal cost of population relocation was $42,532. Implementation relies primarily on subsidies: about four in five beneficiary households receive housing vouchers to relocate elsewhere. The remaining cases are financed via direct investments or have not yet been assigned a subsidy.

3 Data and Description

3.1 Slums Identification

One of the major contributions of this paper is to assemble a long panel of time-varying characteristics of the universe of slums in Chile. We begin by delineating every area ever classified as a slum, which were identified by the Chilean government based on three conditions: (i) an area that is home to at least 8 spatially close households, (ii) households do not

have property rights, and (iii) houses are of substandard quality and lack access to at least one of the basic services (water, sanitation, and/or electricity). Chile applied this definition in national slum censuses conducted between 2011 and 2020 by MINVU and TECHO.[8] Working with municipalities, field teams geocoded candidate locations, verified eligibility on site, and surveyed households. Overall, the picture that emerges is that between 2011 and 2020 the number of slums increased from roughly 650 to nearly 900 (a 40% rise), while the number of households living in slums more than doubled, from about 31,000 to more than 80,000 (see Appendix Figure G1).

A challenge with the raw slum-census data is that slum perimeters expand and contract over time, so geographic boundaries are not consistent across waves. We address this by defining a constant geographic unit for each slum: the spatial union of all observed perimeters in the censuses —its "ever-slum area." This approach fixes the analysis footprint during the study period and ensures that the measured changes are not mechanically driven by shifting boundaries. Figure 1 provides a graphical example.

Using the ever-slum areas as constant geographic units, we construct a 2000–2022 panel of physical characteristics for each slum and its adjacent formal neighborhoods from Google Earth imagery. We downloaded approximately 90,000 images following Huang et al. (2021).[9] Each image covers about 1 km^2, centered on the ever-slum polygon and extending into surrounding non-slum neighborhoods. This setup allows us to measure built structures inside and outside slums, including years prior to the first administrative designation of an area, producing a panel at the slum level spanning more than two decades. The panel covers the universe of slums in Chile, conditional on appearing in any slum census between 2011 and 2020.

3.2 Slum Location

Geocoded data allow us to visualize the location of slums within the urban footprint. For example, Appendix Figure C3 illustrates where slums are located in the city of Valparaiso, along with the 1993 and 2020 urban footprint. The figure shows that slums are found primarily in the city's periphery. However, newer slums, those established after 2010, tend to be in the expansion areas of the city. These patterns hold nationwide and suggest a dynamic process: slums emerge on the periphery of the city and are gradually absorbed into the urban footprint as the city expands. This fits the description in Henderson et al. (2021)'s model, in which the city first develops informally and then formally through redevelopment.

[8] MINVU conducted the 2011 and 2019 censuses, and TECHO carried out the 2016, 2017, 2018, and 2020 waves.

[9] Code to download images from Google Earth is available at this GitHub repository.

The data also show that slums tend to occupy less favorable terrains and are more exposed to hazards than non-slum areas. Table 2 reports pooled block-level regressions in which each terrain or distance-to-amenity measure is regressed on an indicator for "ever-slum" status, combining slum areas and non-slum census blocks in a single sample. The results confirm that the slums are in more rugged terrain, higher elevations, steeper slopes, and closer to rivers, which is consistent with a higher risk of flooding.[10] In addition, slum blocks are farther from essential public and private amenities than non-slum census blocks. In particular, they are further from fire and police stations, schools, supermarkets, and financial institutions. A standardized amenity-distance index shows that slum blocks are, on average, 0.6 standard deviations further away from urban amenities than non-slum blocks.

Finally, we examine the extent to which slums are located close to labor demand. We map the locations of clusters of establishments by skill intensity and measure distances to each slum polygon.[11] As exemplified in Appendix Figure C4 for the case of Coquimbo, slums systematically lie closer to low-skill firm clusters and farther from high-skill clusters, consistent with informal households trading amenity access and housing quality for access to jobs that are likely to match their skills. Similar patterns are observed in other large cities.

3.3 Built Structures and Land Use in and around Slums

We extract building footprints from satellite imagery based on a U-Net convolutional neural network trained on 3-band (RGB) images, which is calibrated to capture the irregular geometry typical of informal housing (Abascal et al., 2022; Fisher et al., 2022). Trained on more than 280,000 images and evaluated on a held-out set of 60,000, the model attains precision 0.94 and recall 0.95. Details on the image processing procedure and model calibration are provided in Online Appendix Section C.[12] Using the predicted footprints, we construct annual measures of total built area, building density, and geometric regularity inside ever-slum polygons and in adjacent formal neighborhoods over time. Appendix Figure G2 shows the

[10] See Online Appendix Section B for variable definitions.

[11] We use data from the 2017 National Firms Census from the National Statistics Institute (INE), which provides exact establishment coordinates and primary industry. "Low-skill" industries are agriculture, mining, construction, retail, transport, and low-tech manufacturing (food, wood, paper, plastic, non-metal minerals, furniture). "High-skill" industries include chemicals, pharmaceuticals, electronics, machinery, automobiles, education, professional, and scientific services.

[12] We pre-process each $4,800 \times 4,800$ Google Earth image into 256 tiles of 300×300 pixels, run inference with test-time augmentations on RGB imagery, relax shape/thresholding to capture irregular informal structures, and then post-process and mosaic tiles back to full scenes. Second, we check the robustness of machine learning algorithms by means of Human Observational Data (HOD). For each slum-year, research assistants examined available imagery and classified the conditions within the slum polygon on the image, reporting the coverage of residential land, vacant land, and paved roads. Each slum was independently coded by the RAs, allowing us to check agreement on the coverage classifications; a supervisor resolved any discrepancies. More details in Online Appendix Section C.

evolution of building density within and outside slums over the 20 year period. Both were largely vacant in the early 2000s; density within slums then climbed rapidly from roughly 2007-2012, paused through about 2016 and accelerated again through 2021. The pattern points to a dynamic slum development that, while in general increasing, can have periods of reduced activity.

In addition to image-based measures, we incorporate two administrative datasets. First, building permit microdata reporting residential and commercial building starts (National Statistics Institute, INE), which is compiled from municipal planning departments covering roughly 90% of municipalities. Second, the Chilean Internal Revenue Service (IRS) cadastre for formal parcels, which records building size, main construction material, year built, and last renovation. Since the cadastre only covers formal properties, comparable details for slums are unavailable. Therefore, for slums we rely on imagery within slum polygons, while the cadastre helps measure spillover effects in nearby formal areas neighborhoods.[13]

3.4 Socio-economic Composition in and around Slums

We complement the imagery panel with block-level population censuses from 2002 and 2017 to characterize sociodemographics in and around slums. We link census blocks to ever-slum polygons using block centroids: a block is labeled as "slum" if its centroid falls within the slum polygon.[14] The median block contains 100 people and 33 households, close to the median slum size —implying that centroid matching provides a relatively precise method of mapping slums. Likewise, we categorize a block as "non-slum" if its centroid lies 200 to 500 meters from the slum boundary.[15]

Appendix Table G1 reports descriptive statistics comparing census blocks inside ever-slum areas with census blocks 200-500 meters from slum borders. The table shows that between 2002 and 2017, population grew more than twice as fast in slum blocks (11.2% vs. 4.9%) than in formal blocks in the vicinity. The changes encompass more than population: The adult male-to-female ratio in slums fell from 1.3 to 0.97, converging to nearby areas. This is consistent with a shift from predominantly male settlers to family households. Human capital and employment indicators show similar convergence, with rising shares reporting secondary or tertiary education and employment. Overall, across key observables, slum and adjacent formal blocks look increasingly similar over time as the slum blocks develop.

[13] Although the cadastre includes assessed value, the valuation uses a mechanical appreciation formula that does not fully capture changes in the real estate market, and hence we do not use it in the analysis.

[14] Because a block's centroid depends on its size and geometry, large or irregular blocks can overlap a slum without their centroid falling inside; to avoid missed matches, we also include the census block that contains the slum centroid.

[15] We exclude immediately adjacent blocks because some can have overlap with the slum area.

3.5 Municipal Level Slum Population Growth

As our final descriptive overview of the data, we examine factors determining growth of slum population at the municipal level. We merge WorldPop's 100-meter gridded population estimates for 2011 and 2017 with the ever-slum polygons, yielding directly comparable slum population counts for the same years in which household surveys and municipal accounts are available. As in the previous section with census blocks, WorldPop grid cells are classified as slums based on their centroid locations: if a centroid falls within a slum boundary, the grid's population is counted as living in a slum. We then estimate long-difference regressions of the log change in municipal slum population (2011–2017) on a set of variables identified in the literature as drivers of slum growth, including housing costs, wages, employment, and municipal expenditures. As a robustness check, we replicate the exercise using census counts (2002–2017) in place of WorldPop's 100-meter grids.

Appendix Table G2 reports the results. In both specifications, higher quality-adjusted rents are strongly associated with faster slum growth: a one–standard deviation increase in the municipality's quality-adjusted rent predicts a 24–29% rise in slum population.[16] This is consistent with tighter housing budgets pushing low-income households toward informal options when formal housing becomes more expensive. Second, higher wages —particularly those for low-skilled workers— are correlated with a slower rate of slum expansion: an increase of one standard deviation in wages is associated with a 21-22% reduction in slum population, again consistent with the affordability channel. Finally, employment growth and municipal social expenditures are not statistically significant predictors. Overall, these patterns indicate that slum populations grow faster in municipalities where quality-adjusted housing rents increase rapidly, but wages for low-skilled workers do not keep up.

4 Methods for Estimating Direct Impacts and Spatial Spillovers

Our objective is to estimate the causal effects of the two policy interventions on slum area characteristics and potential spillovers to adjacent non-slum neighborhoods. The set of outcomes is drawn from our panel of slums and includes variables measured both at the slum level and in adjacent non-slum neighborhoods, including population, residential land coverage, streets paved, number of building permits, number of large buildings, standard

[16] The quality-adjusted rent is the municipality fixed effect from a hedonic regression of rents on housing attributes (e.g., floor area, water access, roof materials, roof condition) using the 2009 and 2017 CASEN surveys. For details on variable definitions, see Online Appendix Section B.

deviation of building main angle, average distance between building footprints, violent and property crime rates in adjacent neighborhoods, and buildings undergoing renovations in adjacent neighborhoods.[17]

Both policies began to be implemented in 2011, but were rolled out gradually across slum areas. Our dataset thus provides repeated observations over a 20 year period with at least 10 years of observations before the slum intervention and up to 10 years after implementation.

Because assignment to either treatment was a function of multiple place-based factors, our econometric strategy addresses potential selection bias using the Synthetic Difference-in-Differences (SDiD) approach developed by Arkhangelsky et al. (2021). The SDiD estimator combines Synthetic Controls methods (Abadie and Gardeazabal, 2003; Abadie, 2021) with a Difference-in-Differences approach that tracks whether and when each slum was affected by one of the two renewal policies. While traditional Synthetic Controls methods focus primarily on matching levels, SDiD emphasizes matching trends in the pre-intervention period, which allows for balance in both levels and trends. Specifically, SDiD constructs a weighted counterfactual for the treated units by assigning both unit and time weights to untreated slums, ensuring that the synthetic control group closely replicates the pre-intervention trends of the treated group. In doing so, SDiD reduces reliance on the parallel trends assumption as in traditional Diff-in-Diff. Another key advantage of SDiD is its ability to accommodate staggered treatment timing. In our setting, slums were treated between 2011 and 2017. SDiD constructs a separate synthetic control for each treated cohorts using weighted combinations of untreated slums observed over the same period, accounting for differential timing of treatment.

More formally, consider the following standard equation for the outcomes of interest, y_{it}, which are measured annually:

$$y_{it} = \alpha_0 + \alpha_i + \alpha_t + \beta \mathbb{D}_{it} + \gamma X_{it} + \varepsilon_{it} \tag{1}$$

where α_0, α_i and α_t represent the constant term, and the slum- and time-fixed effects, respectively, and X_{it} are time-variant slum characteristics. $\mathbb{D}_{it}$ is a treatment indicator equal to one in the year slum i is first treated and thereafter. We estimate two separate models: one for slums under *in-situ* upgrading and one for the population relocation treatment.

For N slums and T time periods in equation 1, we can express the objective optimization function for the SDiD model as follows:

$$(\hat{\beta^{sdid}}, \hat{\alpha}, \hat{\gamma}) = \arg\min_{\beta,\alpha,\gamma} \left\{ \sum_{i=1} \sum_{t=1} (Y_{it} - \alpha_0 - \alpha_i - \alpha_t - \beta \mathbb{D}_{it} - \gamma X_{it})^2 \hat{\omega}_i^{sdid} \hat{\lambda}_t^{sdid} \right\} \tag{2}$$

[17] For details on variable definitions, see Online Appendix Section B.

with $\hat{\omega}_i^{sdid}$ and $\hat{\lambda}_t^{sdid}$ representing the individual slum and time weights, respectively. The slum weights $\hat{\omega}_i^{sdid}$ are chosen to minimize the average squared difference in pre-treatment trends between slums exposed to a given treatment strategy and non-treated slums. This is conceptually similar to the Synthetic Control method, although SC focuses on matching levels of the outcome rather than trends. As for the time weights, $\hat{\lambda}_t^{sdid}$, they are selected to minimize the sum of squared differences between the time-weighted pre-treatment outcomes of the control slums and the simple average of their post-treatment outcomes. In other words, SDiD assigns more weight to pre-treatment periods that are more predictive of post-treatment outcomes for the control group.[18]

Both sets of weights, $\hat{\omega}_i^{sdid}$ and $\hat{\lambda}_t^{sdid}$, are optimally chosen to construct a counterfactual slum for each cohort-treatment combination. Although this weighting is implemented directly through the SDiD model, researchers can refine the pool of control slums to improve comparability. Two relevant variables predicting the probability of receiving treatment are inclusion in the 2011 slum census and the number of slum households.[19] Thus, we restrict the sample to slums created before 2017, ensuring never-treated comparison slums were either listed in the 2011 census or formed a few years after the 2011 census. Second, SDiD allows for the inclusion of time-varying confounding controls to further strengthen the robustness of the estimates, hence we include one-year lagged population data, which helps account for differences in outcome trends related to slum size and the likelihood of receiving any treatment. The annual LandScan population control we use covers a 1 km^2 area centered on the slum, capturing broader population trends while preserving local geographic variation.

For implementation, we follow Clarke et al. (2023) and compare treated and synthetic control slums over time relative to their time-weighted pre-treatment outcomes. Confidence intervals are constructed using a nonparametric bootstrap. SDiD is estimated separately for each intervention type —*in-situ* upgrading and population relocation— and within each, for each treated cohort corresponding to the years 2011 to 2017. This procedure produces cohort-specific synthetic controls tailored to the timing and nature of the intervention. After estimating the optimal weights, we compute post-treatment effects by directly comparing treated slums with their synthetic counterparts. While this yields point estimates, we assess statistical significance by repeating the entire estimation procedure 250 times to generate standard errors. The bootstrapped distribution supports hypothesis testing across intervention types.

It should be noted that all potential beneficiary slums are drawn from the 2011 census.

[18] Notice that, depending on the values assigned to the weights, the SDiD estimator nests other well-known estimators: standard DiD (when $\hat{\omega}_i^{sdid} = 1$ and $\hat{\lambda}_t^{sdid} = 1$) and SC (when only $\hat{\lambda}_t^{sdid} = 1$).

[19] For a detailed analysis of the predictors for a slum to be intervened through *in-situ* or population relocation, see Section 6.1.

This census remained unchanged until 2018, meaning newer slums that appeared after 2018 are excluded. Second, there is typically a two- to four-year lag between the selection of an intervention strategy and its implementation. Although we observe the timing and magnitude of government investments, as well as the number of subsidies allocated to each slum, we lack precise information on the completion dates of infrastructure works and the timing of subsidy disbursement and uptake by beneficiary households. To allow sufficient time for program effects to unfold, we focus on SDiD estimates six years after treatment assignment, which are interpreted as Intention-to-Treat (ITT) effects.

Finally, spillover effects are estimated using independent SDiD models at the buffer level. For each slum, we construct concentric rings at 0–200 m and 200–500 m from the slum boundary; these ring-by-year panels constitute the units of analysis. For each policy arm and buffer, treated units are rings surrounding slums assigned to that intervention, while the control pool consists of analogous rings around never-treated slums. SDiD then selects buffer and time weights to construct synthetic controls whose pre-treatment trends closely parallel those of the treated rings.

5 Results

This section presents the main empirical results on the impacts of Chile's two slum-renewal strategies. Section 5.1 examines direct effects within slum perimeters, estimating how *in-situ* upgrading and population relocation alter land use, housing quality, and the socioeconomic composition of residents six years after treatment assignment. Section 5.2 then assesses local spatial spillovers, evaluating whether interventions generate changes in formal neighborhoods adjacent to treated slums, such as changes in formal housing investment, renovations, demographic composition, and crime. Together, these sections characterize both the internal transformation of intervened slums and the broader neighborhood-level consequences of each policy, providing a comprehensive assessment of their effectiveness.

5.1 Direct Effects

We present six-years intention-to-treat effects of each renewal policy across three domains: Slum Characteristics, Housing Investment, and Housing Quality. Since multiple indicators are included within each family of outcomes, we also present the estimated effect on a composite index constructed following the methodology of Anderson (2008). Specifically, each outcome variable is standardized using the control group's mean and standard deviation and combined using inverse covariance weights. Variable signs are adjusted so that higher

index values correspond to more desirable outcomes.[20]

Direct Effects on Slum Characteristics. Table 3 presents our main direct effects results. For each outcome variable, column 1 shows the mean of the control group, while columns 2 and 3 present intent to treat coefficients in year six for each of the slum policies, from independent regression models. Column 4 reports the p-value of a test of equality of coefficients. Our first result is that *in-situ* upgrading does not lead to significant changes in population, consistent with the program's objective of improving conditions without displacing residents. In contrast, column 3 shows that the population is 16 log points lower in areas assigned to population relocation. As is shown by Appendix Figure D1, there is a gradual decline in population over time, reaching statistical significance only by year six.[21] While the direction of the effect aligns with the program's stated objective of depopulating slum areas, the magnitude falls substantially short of complete population relocation. This attenuated impact may reflect incomplete program take-up, partial household movements, and re-occupation of vacated slum areas by new residents.

Second, we find the share of the slum polygon devoted to residential use —including built structures and private-use spaces such as driveways and patios— declines by 12 percentage points following relocation, with no detectable change under *in-situ* upgrading. These directional effects are in line with the relocation program's objective of reducing residential occupancy in targeted areas, although the estimated magnitudes fall short of what might be expected given the program's stated goals. Finally, in the fourth row, we observe that the share of paved streets increases by 11 percentage points under *in-situ* upgrading, whereas no significant effect is found for population relocation. This result is consistent with the design of the upgrading intervention, which includes infrastructure improvements such as street paving, whereas relocation programs do not involve such investments. In all cases, the null hypothesis of equal effects across treatment arms is rejected, as shown in column 4.

Effects on Housing Investment. As is shown in Table 3, the probability of having any formal housing building permits issued within the slum polygon increases by 5 percentage points, a 5 fold increase from the mean in the control group. This points to *In-situ* upgrading initiating a neighborhood transformation that attracts formal housing investment. We do not detect significant changes for the population relocation program. Qualitatively similar differences are found when we use the number of building permits. Overall, an Anderson

[20] We collapse a family of related outcomes into a single Anderson (2008) index. By replacing K correlated tests with one composite outcome, we avoid inflating the family-wise Type I error that arises from multiple hypothesis testing. This reduces concerns that significant effects may emerge by chance, even if there are no treatment effects.

[21] The pre-treatment coefficients serve as a diagnostic for the credibility of the counterfactual constructed by SDiD. Indeed, the absence of pre-trends and the presence of narrow confidence intervals indicate that the method achieves a close match between treated and synthetic control units prior to intervention.

index of housing investment indicates that *in-situ* upgrading generates levels of investment that are 0.56 standard deviations higher than those observed in the control group. In contrast, we do not find statistically significant effects for the population relocation strategy.

Effects on Housing Quality. The results on housing quality are reported in the bottom panel of Table 3. We first consider the average distance between a dwelling and its nearest eight neighbors, a proxy for residential spaciousness (the opposite of crowding) within slum areas. Slum dwellings are typically built very close to each other; thus, increases in the distance between a building and its neighbors can be interpreted as reflecting a higher neighborhood quality. The estimated effects reveal that *In-situ* upgrading leads to a statistically significant increase of 8.7 meters in the average inter-household distance, equivalent to a 14 percent rise, which points to improvements in urban form and housing quality. By contrast, the effect of population relocation on this measure is small in magnitude and statistically indistinguishable from zero.

Next, we examine the similarity in building orientation, measured using the angle of the principal axis of the minimum bounding rectangle that surrounds each slum structure, following the methodology of Michaels et al. (2021). We compute the standard deviation of the angles between an observation and its eight nearest neighboring structures and then average this measure across the whole polygon. This serves as an indicator of building layout regularity. Higher values of this measure correspond to greater heterogeneity in building orientation and are interpreted as indicative of more disorderly or unplanned construction patterns. In more orderly neighborhoods, most houses tend to face the same direction, resulting in a standard deviation of orientation angles close to zero. We find that *in-situ* upgrading is associated with significantly more orderly urban development. Specifically, the standard deviation in building orientation decreases by 15% following upgrading, indicating increased alignment of structures, whereas the corresponding estimate for the population relocation intervention is small and not statistically distinguishable from zero.

Finally, we examine the prevalence of large structures, defined as buildings with a footprint exceeding 64 m^2. Such structures are unlikely to correspond to single-family dwellings and more plausibly reflect formal constructions intended for commercial use or multi-unit residential occupancy, including multi-story buildings. Our estimates show that *in-situ* upgrading led to a 16 percent increase in large buildings within the slum perimeter ($p < 0.10$) and no statistically detectable change for the population relocation.

Overall, the Anderson index of housing quality points to a 0.09 standard deviation increase in housing quality within the slum perimeter for *in-situ* upgrading ($p- =< 0.01$), and no statistically detectable changes in housing quality under the population relocation strategy.

Effects on Sociodemographics. We next examine the effects of slum interventions on the sociodemographic characteristics of residents using data from the 2002 and 2017 population censuses.[22] Given these outcomes are measured at just two points in time, before and after the interventions, we rely on a canonical Two-Way Fixed Effects (TWFE) framework instead of SDiD.[23,24] A caveat of the TWFE strategy is that given that the interventions were initiated only after 2011 and typically require two to four years for full implementation, the 2017 census may capture outcomes that reflect only partial treatment exposure. As a result, the estimated effects may be attenuated toward zero due to the lag between treatment assignment and actual program completion.

With these caveats in mind, results are reported in Table 4. First, we find that the share of adults with high school diploma or more is approximately 3 percentage points higher in slums that received *in-situ* upgrading, relative to the control group[25]. No statistically significant differences are observed for slums exposed to the population relocation program. Estimates are robust to restricting the sample to adult individuals aged 30, 35, or 40 years, with point estimates remaining stable across age thresholds (not shown). Second, we did not detect significant differences in unemployment rates under either intervention. However, employment rates among adults are already high; around 83 percent in the control group .

The Censuses provide information on aspects of housing quality that are not visible in satellite imagery, such as the number of bedrooms and the quality of wall materials, which are used to test for robustness. Consistent with improved housing conditions, *in-situ* upgrading is associated with a statistically significant increase of approximately 4.5 percentage points

[22]While censuses attempt to cover the universe of the territory, targeting informal areas without formal addresses is challenging, and thus not all of them are included. As a result, our analysis sample here does not correspond to the universe of slum areas, but only those that were identified in the population censuses.

[23]We allow treatment effects to vary by intervention type by estimating a TWFE pooled model of the form $y_{it} = \alpha_0 + \beta_1 D^1_{it} + \beta_2 D^2_{it} + \alpha_i + \alpha_t + \gamma X_{it} + \varepsilon_{it}$, where y_{it} is the outcome of interest measured using 2002 and 2017 population census data, D^1_{it} is a dummy variable equal to one when slum i received the *in-situ* upgrading treatment at time t and remains one thereafter, D^2_{it} is constructed similar to D^1_{it} but for those receiving the population relocation strategy, α_i and α_t are slum and year fixed effects, and X_{it} is a set of baseline covariates. As in the SDiD specification, we include lagged LandScan population as a control to account for differences in slum size and potential confounding trends.

[24]The credibility of the TWFE estimates hinges on how comparable never-treated slums are to those that received interventions before 2017. To strengthen identification, we restrict the sample in two ways. First, we include only slums that existed by 2017, ensuring all units had a similar baseline probability of being treated. Second, we focus on treated slums that had been exposed to their assigned intervention for at least three years prior to the 2017 census, allowing time for effects to materialize. These restrictions help mitigate concerns about differential selection and ensure a more balanced comparison. While this strategy cannot exploit the rich temporal variation used in the SDiD analysis, it offers a feasible alternative for estimating the effects of the interventions on socioeconomic outcomes that are only captured in census years.

[25]Similar results are found for a Sites and Services program in Tanzania, where the authors find that intervened neighborhoods attract more educated residents, who can afford to pay for the higher quality on offer (Michaels et al., 2021)

in the share of dwellings with two or more bedrooms. By contrast, we do not find evidence that population relocation leads to improvements in housing quality. In contrast, the share of dwellings constructed with high-quality wall materials exhibits a modest decline in areas subject to the relocation program.

In sum, the evidence presented in this subsection demonstrates that *in-situ* upgrading is consistently more effective than population relocation in improving conditions within slum polygons, in multiple outcomes and empirical strategies. *In-situ* upgrading is associated with significant enhancements in housing and neighborhood quality, including a higher prevalence of paved streets, the construction of larger and more uniformly oriented structures, and increased spacing between dwellings. Moreover, it appears to catalyze neighborhood transformation by attracting formal housing investment, as indicated by an uptick in building permits. The intervention also contributes to shifts in the sociodemographic profile of slum areas, with evidence of increased residence by individuals with higher socioeconomic status (SES). In contrast, the population relocation program fails to produce measurable improvements in housing and neighborhood quality and achieves only limited success in reducing population within slum perimeters, its primary stated objective.

5.2 Spatial Spillovers

We exploit the same SDiD identification strategy to estimate spatial externality effects of both *in-situ* upgrading and population relocation on surrounding formal neighborhoods. The analysis focuses on two spatial buffers surrounding slum boundaries: 0-200 meters and 200-500 meters, meaning we capture spillover effects within the immediately adjacent areas as well as in neighboring areas located farther from the treated zones. We begin with crime—available only for non-slum areas—and then present results for housing investment, housing quality, and sociodemographic outcomes in the same sequence used in the direct-effects section.

Crime. We use geocoded data on police incident reports from the Under-Secretary for Crime Prevention (SPD). The dataset covers the entire country and includes georeferenced events from 2013 to 2021, which allows us to recover the spatial crime externalities of each policy up to five years of assignment to treatment.[26]

Table 5 reports the results. The top row reports results for an aggregate crime index, standardized such that estimated coefficients can be interpreted in units of standard deviations. After 5 years of assignment, we find that *in-situ* upgrading is associated with

[26]Because geocoding relies on standardized addresses, crimes occurring within slums, where formal addresses are uncommon, are systematically underreported or completely excluded. This data limitation prevents us from estimating direct crime impacts within slums.

a statistically significant reduction in crime in the 0-200 meter buffer surrounding treated slums of -0.12 standard deviations ($p < 0.01$) in reported crime incidents relative to synthetic control areas. In contrast, the estimated effect of population relocation on crime is positive and statistically indistinguishable from zero, and we reject the null hypothesis of equality between the two policy coefficients ($p < 0.01$), as reported in column 4.[27]

Disaggregating the crime index into its three constituent components (property crimes, robberies, and homicides per square kilometer) we find that *in-situ* upgrading consistently leads to statistically significant reductions across all categories. It is important to exercise caution when interpreting crime reports as proxies for actual crime levels, due to well-known issues with under-reporting. With this in mind, among all crime categories, homicide is generally considered the least susceptible to such reporting biases. It is therefore reassuring that our findings also hold for homicides. In contrast, population relocation is associated with statistically insignificant changes in each category. Accordingly, we can reject the null hypothesis of coefficient equality in policies in all three domains.

Next, columns (5) through (8) present estimates for the 200-500 meter buffer surrounding slum perimeters. While most point estimates remain directionally consistent with those observed in the adjacent buffer, neither policy exhibits statistically significant effects at conventional levels.

The observed reduction in crime in areas adjacent to slums that undergo *in-situ* upgrading is a plausible consequence of the infrastructure improvements in slum conditions, as new constructions have better security standards, and specific investments such as street pavement and connection to basic services increase state presence and reduce the cost for police to actively patrol these areas. Likewise, the arrival of individuals with higher socioeconomic status, who may be argued are less likely to engage in criminal behavior, could also have contributed. Importantly, this mechanism does not depend on the assumption that offenders disproportionately reside within slum areas. If that were the case, the population relocation strategy would similarly be expected to reduce criminal activity, due to the stronger population decline with the relocation strategy.

Housing Investment. Next, we turn to housing investment outcomes, with results presented in Table 6. Within 200m of a slum, we find a 0.22SD ($p < 0.01$) increase in the housing investment index from *in-situ* upgrading, which is significantly larger than the 0.14SD increase observed for population relocation. When we consider individual components of the index, a similar pattern holds. Specifically, there is a 4 percentage point increase in new buildings and a 32% increase in renovated buildings. For areas near population relocation slums, we find a smaller effect on new buildings and no significant effect on building

[27] Appendix Figure D3 illustrates the dynamic of spillover effects on crime.

renovations. Second, the probability that at least one building permit is approved in a formal neighborhood within 200 meters of the slum increases by 10% and there is a 13% increase in the number of building permits relative to the counterfactual buffers. For formal neighborhoods near population relocation slums, we still find positive but smaller point estimates. Overall, this suggests that communities near upgraded slums end up seeing a larger share of new housing starts and renovated properties, which we interpret as positive spatial spillovers in the extensive margin (new buildings). Finally, the spatial spillovers are somewhat smaller but persist for neighboring areas located farther from the treated zones (200-500m), and these are always statistically larger for *in-situ* upgrading than for population relocation.

Housing Quality. The bottom panel of Table 6 reports estimates of spatial spillovers on intensive margin outcomes for all housing stock at different distance bands. Overall, impacts are muted, as the overall index is not significantly different from zero for either policy. Yet, when considering the different subcomponents, we find a significant reduction in the average building age in areas within 200 meters of *in-situ* upgraded slums, as well as an increase in the building size. In contrast, no significant effects are found on building orientation angle and distance between buildings, although this result is not surprising in more established neighborhoods where changes in building regularity and distribution can take much longer to appear.

Sociodemographics. In Table 7, we use the two population census rounds and the TWFE specification to examine whether the interventions generated changes in the socioeconomic characteristics of individuals who live close to slums.[28] Subject to the caveats discussed previously about this analysis relying on just two waves of data, the results indicate that adults living near *in-situ* upgraded areas are statistically more likely to have high school diploma than adults living near never-treated slums. They are also more likely to be employed. In contrast, adults living close to population relocation slums are not statistically different from those living near never-treated slums. This is true whether we consider the 0-200m or the 200-500m distance band.

Second, cumulative population growth in areas near slums assigned to population relocation is significantly lower than in areas near never-treated slums. Specifically, between 2002 and 2017, population growth is 21 percent lower within 200 meters of the slum perimeter and 18 percent lower in the 200–500 meter band. This finding is consistent with the view that slums subject to population relocation interventions are, in equilibrium, less attractive in terms of amenities than untreated slums. This surprising result is also consistent with peripheral formal areas experiencing decline when population relocation strategies are imple-

[28]The sample size reflects the universe of formal areas surrounding slums, which are fully identified in the censuses. This is different for slum areas, which are not comprehensively identified in the censuses.

mented in nearby slums, which may be due to density losses in nearby slums. If there is less population in the vicinity, amenities such as stores and buses may be harder to sustain even in formal neighborhoods. Finally, the lower panel revisits housing indicators, confirming the findings of the SDiD specification discussed previously, namely, *in-situ* upgrading results in improved housing in surrounding areas when considering a housing index consisting of wall quality and homes with 2 or more bedrooms.

Overall, our analysis demonstrates that *in-situ* upgrading generates substantial positive spillover effects in adjacent formal neighborhoods, whereas population relocation strategies do not yield comparable benefits. Notably, we document significant reductions in crime rates, pointing to improvements in local safety. Together with greater public safety, areas within 200 meters of *in-situ* upgraded slums exhibit increased building activity and housing investment, as indicated by higher probabilities of approved building permits and a marked rise in building renovations. These neighborhoods also undergo favorable sociodemographic changes, which attract a greater share of high socioeconomic status (SES) residents. Taken together, these findings underscore the broader effectiveness of *in-situ* upgrading not only in improving conditions within targeted slums but also in catalyzing positive transformations in the surrounding neighborhood.

6 *In-situ* Upgrading versus Population Relocation

Up to this point, we have provided a parallel analysis of both slum policies. The pattern of results presented thus far strongly suggests that if the objective is to generate higher quality neighborhoods within slum polygons and in surrounding non-slum areas, the *in-situ* upgrading strategy achieves superior results at a lower cost than population relocation.

However, as pointed out in section 4, the SDiD models are independently estimated for each policy and there is no guarantee that the synthetic controls are the same or similar to each other across the two policies. In this section, we provide a head-to-head comparison of the two policies when the counterfactual is the same. We proceed in two steps. First, we empirically examine the determinants of each renewal policy and assess the extent of overlap in the probability distribution of being assigned to each policy in the control slums. Second, we directly compare *in-situ* upgrading and relocation slums using an SDiD estimator subject to the subsample of slums whose predicted assignment probabilities to either program overlap. These estimates will directly compare both policies to each other, creating the synthetic counterfactual not from the untreated set of slums, but from the set of slums that received the other policy in the same year, with the null hypothesis being that effects are equal for the two policies.

6.1 The Spatial Allocation of Slum Renewal Policies

According to internal MINVU documents, the choice between *in-situ* upgrading and population relocation is a policy decision based on a cost-benefit analysis that often depends on a combination of local factors, including consideration of property rights, feasibility for new construction (terrain characteristics and exposure to natural hazards) and location within the city structure, typically measured in terms of access to amenities like public transit, green areas and other public services (DIPRES, 2019).

To empirically examine this process, we compile, for each ever-slum area, a set of potential determinants of treatment assignment grouped into three categories: slum, terrain, and location characteristics. We then estimate two separate logistic regressions: (i) a model predicting the likelihood of receiving any treatment, and (ii) a model predicting the likelihood of receiving *in-situ* upgrading, conditional on having been treated.

The results are presented in Table 8. In columns (1) and (2) the dependent variable is an indicator for whether a slum receives any type of intervention, with column (2) including region fixed effects. As expected from discussions with program managers, we observe a strong positive association between inclusion in the 2011 slum census and the likelihood of treatment, highlighting the program's political mandate to rely on this registry to allocate resources. We also find that treatment is less likely for larger slums, measured by the number of households, possibly reflecting higher coordination costs. Slums situated on terrain unsuitable for construction are significantly less likely to receive any intervention, again suggesting that the higher fiscal burden associated with large-scale improvements deters government action. The probability of intervention increases with distance from the city center —consistent with lower land values, and therefore lower intervention costs, in peripheral areas— and also rises when surrounding neighborhoods exhibit higher building density or when municipalities spend more on social programs. Overall, these results indicate that slums identified in the 2011 census, with fewer households, and located on flatter terrain in the urban periphery, are more likely to receive an intervention.

Next, in columns (4) and (5), we analyze the determinants of receiving *in-situ* upgrading, conditional on a slum being selected for any intervention. The results indicate that slums with larger geographic footprints and greater numbers of households are more likely to undergo *in-situ* upgrading. Second, slums located in risky areas (i.e., at risk of *tsunami*, volcano, wildfire, or flooding) or close to rivers are significantly more likely to be assigned to population relocation rather than *in-situ* upgrading, consistent with MINVU's objective of mitigating risks. Similarly, slums located on terrain deemed unsuitable for construction are also more likely to be targeted for relocation. In contrast, rugged terrain is positively associated with the likelihood of receiving *in-situ* upgrading conditional on risk, suggesting

that adaptation of existing structures may be preferred over new construction even in such contexts. Finally, slums located near or beyond the urban periphery exhibit a higher probability of receiving *in-situ* interventions. Overall, the evidence is consistent with intervention decisions being driven by cost considerations: *in-situ* upgrading is favored when feasible due to its lower fiscal burden, whereas relocation is favored in high-risk areas.

6.2 A Direct Comparison of the Two Policies

The previous subsection showed that, on average, the two slum policies tend to be applied in different contexts. In this subsection, we aim to conduct a head-to-head comparison of the two interventions. As a first step, we assess whether the distributions of treatment probabilities overlap across the two policies. To do so, we re-estimate the specification in column (2) of Table 8, restricting the sample to *in-situ* upgraded slums and never-treated slums. Panel (a) of Appendix Figure E1 plots the predicted probabilities of treatment for the in-sample *in-situ* upgraded slums, along with the corresponding out-of-sample predictions for relocation slums. The distributions exhibit full overlap. Panel (b) performs the analogous exercise for relocation slums by restricting the sample to relocation and never-treated slums, and then plotting predicted probabilities for both the in-sample relocation slums and the out-of-sample *in-situ* upgraded slums. Again, the distributions fully overlap. Taken together, these results indicate that although the two strategies are, on average, implemented in different contexts, both policies are in fact used across the full range of slum characteristics observed in the data.[29]

Having established substantial overlap in treatment probabilities, we next implement an SDiD analysis analogous to that in Section 5, except that one of the two policies is classified as the "treated" group and synthetic controls are drawn from the pool of slums assigned to the other policy. This procedure constructs a synthetic control from the pool of slums treated with the alternative policy. Importantly, this is conditional on having assigned treatment in the same year (same cohort), such that the synthetic controls choice of pretreatment periods are similar across arms. Then, by matching pre-treatment trends, we can assess whether the two policies produce differential effects when applied to comparable slums. We conduct this exercise in both directions, and the corresponding results are reported in Table 9.

Conducting this head-to-head comparison reproduces the patterns documented in Section 5 — when SDiD effects were estimated using never-treated slums as synthetic controls —

[29] An alternative exercise is to plot the counterfactual synthetic controls that correspond to each treatment arm (i.e., the SDiD estimated series for the never treated slums associated to each policy) and test whether they are statistically balanced. As shown by Appendix Figure E2, the distributions of the estimated counterfactual synthetic controls of each policy are never statistically different across a wide range of outcomes.

reinforcing the evidence that *in-situ* upgrading generates higher neighborhood quality than population relocation. For example, when *in-situ*–upgraded slums are classified as treated, both the probability and number of building permits increase relative to slums assigned to relocation. When the labels are reversed, the magnitudes are similar in absolute value but the signs flip, implying *in-situ* upgrading areas attract substantially more investment than relocation sites. This mirrored pattern also appears consistently across other outcomes, providing further support for the conclusion that *in-situ* upgrading yields greater improvements than population relocation.

Note, however, that one exception is the average distance to neighbors, which falls in the direct comparison. At first glance, this contrasts with results shown in Table 3, which suggests that the treatment effects on distance to neighbors are larger in magnitude for *in-situ* upgrading than for relocation. One possible explanation is that the SDiD models in Table 3 are estimated separately for each policy, and thus the synthetic controls used for *in-situ* upgrading and relocation slums may differ for this particular outcome. Instead, enforcing a head-to-head comparison under a common counterfactual —as we do in the direct-comparison exercise we show in this section— precisely addresses that concern.

Overall, the direct-comparison results reinforce our earlier findings: housing investment, the issuance of building permits, the presence of larger buildings, and the regularity of building structures are all more prevalent within slum perimeters undergoing *in-situ* upgrading.

7 Discussion

We empirically evaluate the two most widely used slum-intervention policies to assess whether they achieve their stated objectives. In the Chilean context, we find that population relocation, as currently implemented, does not meet its goal of clearing a slum from residential use. Six years after intervention, relocation slums retain roughly 85% of their original population, and housing quality shows no improvement relative to comparable never-treated slums. Consistent with these patterns, we also find no evidence of spatial spillovers from relocation —whether in crime reduction, new investment, or broader neighborhood upgrading.

The second policy we evaluated, *in-situ* slum upgrading, was considerably more successful in achieving its stated objective of improving housing and neighborhood conditions without displacing residents. Consistent with the intervention making the area more desirable, we observe increased investment and an inflow of higher–socioeconomic-status residents. These improvements in neighborhood quality also extend to adjacent formal neighborhoods, where we additionally find evidence of reductions in crime.

We interpret these results within the broader debate on place-based versus person-based

policy interventions (Bartik, 1991, 2020; Glaeser, 2001; Glaeser and Gottlieb, 2008; Kline and Moretti, 2014b; Neumark and Simpson, 2015; Gaubert et al., 2025; Garin and Rothbaum, 2025). The population-relocation program aligns conceptually with a person-based approach. However, our findings indicate that it does not achieve its core objective of moving individuals out of harm's way, suggesting that its design and implementation should be reconsidered. One avenue for improvement would be to condition access to formal housing on fully vacating the slum premises. Re-occupation could be deterred by demolishing existing structures and undertaking defensive investments —such as installing barriers or converting the land into parks— to prevent rapid repopulation. More generally, redevelopment of the vacated slum area for non-residential uses is more likely to succeed when the program operates as a collective, take-it-or-leave-it offer rather than relying on individual participation decisions.

The MTO literature shows that children's outcomes improve the longer they reside in higher-quality neighborhoods (Kling et al., 2007; Chetty et al., 2016). This evidence suggests that relocating slum residents to well-located formal neighborhoods could, in principle, be more effective than *in-situ* upgrading, particularly for low-income households whose budget constraints and limited access to credit prevent them from moving to high-rent, high-productivity areas.[30] However, our findings underscore the practical difficulty of inducing individuals to leave the places they have chosen as home and where they have invested resources in constructing their dwellings. The limited existing evidence similarly documents strong reluctance to relocate. For instance, Barnhardt et al. (2017) show that in Ahmedabad, India, one-third of eligible slum households declined subsidized formal housing on the urban periphery, and among those who initially accepted, another third eventually returned to the slum rather than remain in formal housing.

One could argue that this reluctance reflects the undesirable location of the offered formal housing, and that providing homes in highly desirable neighborhoods would increase take-up. However, this view clashes with the reality that acquiring land or building formal housing in high-demand, high-productivity neighborhoods is prohibitively costly. This practical constraint sharply limits the feasibility of large-scale relocation policies and reinforces the relevance of *in-situ* upgrading as a viable policy instrument.

Moreover, homeownership is a central consideration for households, and temporary rental

[30]This may vary across developed and developing countries. For instance, in the United States, Garin et al. (2025) show that more than one-third of people living in high-poverty neighborhoods move to a less poor neighborhood within eight years, and that higher earnings are a significant predictor of migration to better neighborhoods. In contrast, upward residential mobility among slum dwellers appears to be quite limited in India, as movement out of slums —particularly into non-slum neighborhoods— is rare (Rains and Krishna, 2020). Likewise, in Tanzania, rural-to-urban migrants often experience an initial improvement in housing conditions, but then show limited upward residential mobility within the city, with many remaining in low-standard neighborhoods despite multiple moves (Patel et al., 2020).

subsidies offer substantially weaker incentives to move than owning even a substandard dwelling within a slum. Consistent with this, there is substantial evidence that rental-subsidy interventions would face very low take-up (Selman, 2025), in stark contrast to the MTO context, where the eligible population consisted entirely of renters.

A related tension arises when relocation programs aim to keep costs manageable by offering housing in more affordable, peripheral areas. Convincing slum dwellers to move far from opportunity zones requires strong improvements in accessibility —amenities typically found in precisely those neighborhoods where land values are highest. Evidence from multiple developing-country cities shows that shocks to labor market access and transit connectivity reshape welfare and spatial sorting (Balboni et al., 2025; Bryan and Morten, 2025). This suggests that the welfare consequences of relocation critically hinge on whether destination neighborhoods provide sufficiently high connectivity to compensate for their distance.

In contrast to population relocation, the *in-situ* upgrading program can be interpreted as combining elements of both place-based and person-based policies. On the one hand, the provision of higher-quality housing accompanied by secure property rights can be justified on equity grounds — particularly given that residing in a slum is itself a credible signal of economic distress, and that the provision of land titles guarantees the place-based intervention gains accrue primarily to slum dwellers.[31] On the other hand, the program also exhibits the features of a place-based intervention, as it supplies missing infrastructure with clear public-good characteristics, including paved streets, reliable water supplies, sewerage, and electrification. These improvements are consistent with the broader rationale for public provision of infrastructure. Local governments generally have an interest in ensuring that all residents meet a minimum standard of housing quality, a principle reflected in building codes and inspection regimes in high-income countries. In this sense, slum upgrading can be viewed as a second-best policy response in contexts where the private sector is unable or unwilling to supply adequate-quality housing.

This perspective connects to the U.S. literature on place-based policies, which has focused primarily on interventions such as wage subsidies for firms in distressed areas (e.g., Busso et al., 2013), where concerns about distortions arising from firm relocation or selective worker inflows are central. In contrast, these mobility-related concerns are largely absent in the context of *in-situ* slum upgrading: the intervention is tightly targeted to poor households and does not distort broader economic location decisions. Importantly, the positive localized spillovers we document — particularly the sizable reductions in crime in surrounding

[31] In a standard spatial equilibrium model with perfectly inelastic housing supply, the gains from place-based interventions accrue primarily to landowners rather than to renters, implying equity concerns. Yet this is less of a concern in the context of an *in-situ* upgrading program that includes the provision of land titles to slum dwellers, as in the MINVU program.

neighborhoods — provide an additional rationale for government involvement, as upgrading helps internalize these external benefits.[32] More generally, the equity-efficiency trade-off that typically dominates debates on place-based policies is less binding when these interventions target low-income groups such as slum dwellers.

Note, however, that the program design can play a central role in shaping long-term outcomes. In Jakarta, Harari and Wong (2025) show that upgrading slums without providing land titles to slum dwellers improves basic amenities yet slows the emergence of high-rise formal structures in the long-run, illustrating how design choices can influence the broader trajectory of urban development. In Tanzania, descriptive evidence from Michaels et al. (2021) indicates that upgrading pre-existing settlements often yields housing quality that is no better than that of non-intervened slums. They further document that upgraded roads and water mains deteriorate rapidly when no maintenance systems are established, underscoring how vulnerable infrastructure becomes without sustained institutional support. By contrast, in Chile, the emphasis on tenure regularization and local infrastructure investment appears to catalyze complementary private improvements and neighborhood safety in upgraded slums. This suggests that securing property rights and planning for long-term maintenance are critical elements for unlocking household investment and preventing the deterioration that undermines many upgrading efforts.

Finally, a key concern with these kinds of interventions is the potential for dynamic moral hazard, whereby slum formation increases in anticipation of future government benefits. Although this concern applies to both strategies, our evidence from the Chilean context (see Section 3.5) does not suggest it is empirically salient. For instance, municipal social expenditures are not statistically significant predictors of slum growth. Still, it remains possible that some specific interventions could induce slum formation through these forward-looking incentives, making this an important avenue for future research that lies beyond the scope of this paper.

8 Conclusion

As rapid urbanization continues to strain housing markets and infrastructure, credible evidence on the cost-effectiveness of different policies in informal settlements becomes increasingly urgent. By assembling a 2000-2021 high-frequency panel of satellite imagery, building

[32] As emphasized by Gaubert et al. (2025) and Kline and Moretti (2014a), place-based policies can be more cost-effective because they unlock agglomeration forces and productivity spillovers, none of which are obvious outcomes of individual transfers. In our setting, not only the capitalization of public investments remains with slum residents, but also the positive externalities we document for nearby formal neighborhoods represent local multipliers that standard efficiency arguments should not overlook.

permits, population censuses and crime reports for the universe of slum areas in Chile, this paper employs a SDiD approach to evaluate the direct effects of the two most important slum renewal policies —*in-situ* upgrading and population relocation— as well as their spillover effects on surrounding formal neighborhoods.

Our findings indicate that *in-situ* upgrading is more effective than population relocation at improving conditions within slum polygons and at generating positive spillover effects in surrounding formal neighborhoods. *In-situ* upgrading yields substantial improvements in housing quality, including a higher incidence of paved streets, larger and more regularly oriented buildings, and increased spacing between dwellings. It also catalyzes neighborhood transformation by attracting formal housing investment. In adjacent formal areas, we observe favorable changes in sociodemographic composition as well as significant reductions in property and violent crime. In contrast, population relocation generates only limited improvements in housing quality and sociodemographic outcomes, producing merely a modest reduction in the population residing within the slum polygon.

Taken together, the evidence on direct impacts and spillover effects provides strong support for *in-situ* upgrading as a more effective strategy than population relocation to promote the development of more desirable neighborhoods. The average fiscal expenditure per slum household is approximately one third lower for *in-situ* upgrading relative to population relocation, implying that upgrading is more cost-effective, further strengthening its policy appeal. Yet the fact that *in-situ* upgrading often outperforms relocation does not imply that urban planners should always select it. In many cases, *in-situ* upgrading is simply not viable: numerous slums occupy steep slopes, floodplains, or other high-hazard terrain where upgrading would be unsafe or infeasible. In such settings, relocation may be the only appropriate instrument, but we propose changes to make it more effective: Make take up of formal housing conditional on vacating the slum, and investing more in defensive investments such as installing barriers or converting land to parks to prevent rapid repopulation.

Despite our extensive data collection and analysis, the study has limitations that point to areas needing further research. Notably, we lack detailed information on the post-relocation outcomes of individuals moved from population relocation areas. Understanding whether these individuals resettle in better neighborhoods and improve their well-being, or if the disruption of their community ties adversely affects them, is an important consideration (Rojas-Ampuero and Carrera, 2023). Future research should also focus on the relationship between local labor markets and slums, particularly how the availability of low-skilled jobs influences slum formation and growth. Gaining insights into these dynamics could inform more effective policies that address both housing and employment challenges faced by low-income populations.

References

Abadie, A. (2021). Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*, 59(2):391–425.

Abadie, A. and Gardeazabal, J. (2003). The economic costs of conflict: A case study of the basque country. *American Economic Review*, 93(1):113–132.

Abascal, A., Rodríguez-Carreño, I., Vanhuysse, S., Georganos, S., Sliuzas, R., Wolff, E., and Kuffer, M. (2022). Identifying degrees of deprivation from space using deep learning and morphological spatial analysis of deprived urban areas. *Computers, Environment and Urban Systems*, 95.

Alves, G. (2021). Slum growth in Brazilian cities. *Journal of Urban Economics*, 122:103327.

Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry preschool, and early training projects. *Journal of the American Statistical Association*, 103(484):1481–1495.

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

Balboni, C., Bryan, G., Morten, M., O'Connor, C., and Siddiqi, B. (2025). Impacts of infrastructure: People versus place.

Barnhardt, S., Field, E., and Pande, R. (2017). Moving to opportunity or isolation? network effects of a randomized housing lottery in urban India. *American Economic Journal: Applied Economics*, 9(1):1–32.

Bartik, T. (1991). Who benefits from state and local economic development policies? *W.E.Upjohn In., MI.*

Bartik, T. (2020). Using place-based jobs policies to help distressed communities. *Journal of Economic Perspectives*, 34:99–127.

Bazzi, S., Gaduh, A., Rothenberg, A., and Wong, M. (2019). Unity in diversity? how intergroup contact can foster nation building. *American Economic Review*, 109(11):3978–4025.

Becker, S. and Ferrara, A. (2019). Consequences of forced migration: A survey of recent findings. *Labour Economics*, 59:1–16.

Belsky, E., DuBroff, N., McCue, D., Harris, C., McCartney, S., and Molinsky, J. (2013). Advancing inclusive and sustainable urban development: Correcting planning failures and connecting communities to capital. *Joint Center for Housing Studies of Harvard University*.

Brueckner, J. and Lall, S. (2015). Cities in developing countries: Fueled by rural-urban migration, lacking in tenure security, and short of affordable housing. *in Handbook of Regional and Urban Economics, ed. by G. Duranton, J. V. Henderson, and W. C. Strange, Elsevier Science*, 5(21).

Brueckner, J. K. and Selod, H. (2009). A theory of urban squatting and land-tenure formalization in developing countries. *American Economic Journal: Economic Policy*, 1(1):28–51.

Bryan, G., Franklin, S., Getahun, T., and Winton, S. (2025). No place like home? the causal effect of forced relocation from central addis ababa. *Mimeo*.

Bryan, G. and Morten, M. (2025). Spatial economics for low- and middle-income countries. In Donaldson, D. and Redding, S. J., editors, *Handbook of Regional and Urban Economics*, volume 6, chapter 5. Elsevier.

Burchfield, M., Overman, H., Puga, D., and Turner, M. (2006). Causes of sprawl: A portrait from space. *The Quarterly Journal of Economics*, 121(2):587–633.

Busso, M., Gregory, J., and Kline, P. (2013). Assessing the incidence and efficiency of a prominent place based policy. *American Economic Review*, 103(2):897–947.

Cavalcanti, P., Monge-Naranjo, A., and Torres de Melo, L. (2025). Of cities and slums. *Journal of Political Economy*, 133(9):2693–2734.

Cavalcanti, T., Da Mata, D., and Santos, M. (2019). On the Determinants of Slum Formation. *The Economic Journal*, 129(621):1971–1991.

Chetty, R., Hendren, N., and Katz, L. (2016). The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. *American Economic Review*, 106(4):855–902.

Chiovelli, G., Michalopoulos, S., Papaioannou, E., and Sequeira, S. (2021). Forced displacement and human capital: Evidence from separated siblings. *NBER Working Paper No. 29589*.

Clarke, D., Pailañir, D., Athey, S., and Imbens, G. (2023). Synthetic Difference-in-Differences estimation. IZA Working Paper.

Collier, P. and Venables, A. J. (2015). Housing and urbanization in Africa: Unleashing a formal market process. In *The Urban Imperative Towards Competitive Cities, Glaeser, E. and Joshi-Ghani, A., pp.1-19.*

Collins, W. and Shester, K. (2013). Slum clearance and urban renewal in the United States. *American Economic Journal: Applied Economics*, 5(1):239–273.

Czakon, J., Kaczmarek, K. A., Pyskir, A., and TarasiewiczAuthor, P. (2021). Open solution mapping challenge. https://github.com/neptune-ai/open-solution-mapping-challenge. Accessed: 2023-07.

Dasgupta, B. and Lall, S. (2009). Assessing benefits of slum upgrading programs in second-best settings. *In Urban Land markets: Improving land management for successful urbanization, Springer Netherlands*, 9:225–251.
DIPRES (2019). Informe final de Evaluacion Programas Gubernamentales (EPG): Programa campamentos MINVU.
Duranton, G. (2008). Viewpoint: From cities to productivity and growth in developing countries. *Canadian Journal of Economics*, 43(4):689–736.
Duranton, G. and Puga, D. (2015). Chapter 8 - urban land use. In Duranton, G., Henderson, J. V., and Strange, W. C., editors, *Handbook of Regional and Urban Economics*, volume 5 of *Handbook of Regional and Urban Economics*, pages 467–560. Elsevier.
Duranton, G. and Puga, D. (2023). Urban growth and its aggregate implications. *Econometrica*, 91:2219–59.
Duranton, G. and Venables, A. (2020). Place-based policies for development. *NBER WP 24562*.
Field, E. (2007). Entitled to work: Urban property rights and labor supply in Peru. *QJE*, 122(4):1561–1602.
Fisher, T., Gibson, H., Liu, Y., Abdar, M., Posa, M., Salimi, G., Hassaine, A., Cai, Y., Rahimi, K., and Mamouei, M. (2022). Uncertainty-aware interpretable dl for slum mapping and monitoring. *RS*, 14(13).
Franklin, S. (2020). Enabled to work: The impact of government housing on slum dwellers in South Africa. *Journal of Urban Economics*, 118:103265.
Galiani, S., Gertler, P. J., Undurraga, R., Cooper, R., Martínez, S., and Ross, A. (2017). Shelter from the storm: Upgrading housing infrastructure in Latin American slums. *JUE*, 98:187–213.
Galiani, S. and Schargrodsky, E. (2010). Property rights for the poor: Effects of land titling. *Journal of Public Economics*, 94(9):700–729.
Garin, A., Jenkins, E., Mast, E., and Stuart, B. (2025). Dynamic individuals, static neighborhoods: Migration and earnings changes in poor neighborhoods. *mimeo*.
Garin, A. and Rothbaum, J. (2025). The long-run impacts of public industrial investment on local development and economic mobility: Evidence from World War II. *QJE*, 140(1):459–520.
Gaubert, C., Kline, P., Vergara, D., and Yagan, D. (2025). Place-based redistribution. *AER*, 115:3415–50.
Gechter, M. and Tsivanidis, N. (2023). Spatial spillovers from high-rise developments: Evidence from the Mumbai Mills redevelopment. *Working Paper*.
Glaeser, E. (2011). *Triumph of the City: How our greatest invention makes us richer, smarter, greener, healthier, and happier*. Penguin Press.
Glaeser, E. and Gottlieb, J. (2008). The economics of place-making policies. *BPEA*, 39:155–253.
Glaeser, E. L. (2001). The economics of location-based tax incentives. *DP 1932, Harvard, Cambridge, MA*.
Gonzalez-Navarro, M. and Quintana-Domeque, C. (2016). Paving streets for the poor: Experimental analysis of infrastructure effects. *Review of Economics and Statistics*, 98(2):254–267.
Gonzalez-Navarro, M. and Undurraga, R. (2023). Immigration and slums.
Harari, M. and Wong, M. (2025). Slum upgrading and long-run urban development: Evidence from Indonesia. *The Review of Economic Studies*, forthcoming.
Henderson, J., Libano-Monteiro, F., Manara, M., Michaels, G., and Regan, T. (2025). Economics of greenfield urban planning. *CESifo Working Paper No. 11860*.
Henderson, J., Storeygard, A., and Weil, D. (2012). Measuring economic growth from outer space. *American Economic Review*, 102(2):994–1028.
Henderson, J. V., Regan, T., and Venables, A. J. (2021). Building the city: From slums to a modern metropolis. *The Review of Economic Studies*, 88(3):1157–1192.
Hidalgo, R. (2019). La vivienda social en Chile y la construcción del espacio urbano en stgo del S.XX. *RIL*.
Huang, L. Y., Hsiang, S. M., and Gonzalez-Navarro, M. (2021). Using satellite imagery and deep learning to evaluate the impact of anti-poverty programs.
Iglovikov, V. I., Seferbekov, S., Buslaev, A. V., and Shvets, A. (2018). TernausNetV2: Fully convolutional network for instance segmentation.
Jacobs, J. (1961). The death and life of great american cities. *New York: Random House*.
Jimenez, E. (1984). Tenure security and urban squatting. *RESTAT*, pages 556–567.
Kline, P. and Moretti, E. (2014a). Local economic development, agglomeration economies, and the big push: 100 years of evidence from the tennessee valley authority. *Quarterly Journal of Economics*, (1):275–331.
Kline, P. and Moretti, E. (2014b). People, places, and public policy: Some simple welfare economics of local economic development programs. *Annual Review of Economics*, 6:629–62.
Kling, J., Liebman, J., and Katz, L. (2007). Experimental analysis of neighborhood effects. *ECTA*, 75:83–119.
Klotz, M., Wurm, M., Zhu, X., and Taubenböck, H. (2017). Digital deserts on the ground and from space. In *2017 Joint Urban Remote Sensing Event (JURSE)*, pages 1–4.
Kuffer, M., Pfeffer, K., and Sliuzas, R. (2016). Slums from space: 15 years of slum mapping using remote sensing. *Remote Sensing*, 8(6).
Marx, B., Stoker, T., and Suri, T. (2013). The economics of slums in the Developing World. *JEP*, 27:187–210.

Michaels, G., Nigmatulina, D., Rauch, F., Regan, T., Baruah, N., and Dahlstrand, A. (2021). Planning ahead for better neighborhoods: Long-run evidence from Tanzania. *Journal of Political Economy*, 129:2112–56.
Michalopoulos, S. and Papaioannou, E. (2013). Pre-colonial ethnic institutions and contemporary african development. *Econometrica*, 81(1):113–152.
MINVU (2011). Catastro nacional de campamentos. *Ministerio de Vivienda y Urbanismo*.
MINVU (2021). Áreas funcionales de las ciudades chilenas. *Centro de estudios de ciudad y territorio. Ministerio de Vivienda y Urbanismo, available at this link*.
Mohanty, S. P., Czakon, J., Kaczmarek, K. A., Pyskir, A., Tarasiewicz, P., Kunwar, S., Rohrbach, J., Luo, D., Prasad, M., Fleer, S., et al. (2020). Deep learning for understanding SI: Experimental survey. *FAI*, 3.
Müller, I., Taubenböck, H., Kuffer, M., and Wurm, M. (2020). Misperceptions of predominant slum locations? spatial analysis of slum locations in terms of topography based on earth observation data. *RS*, 12(15).
Najmi, A., Gevaert, C. M., Kohli, D., Kuffer, M., and Pratomo, J. (2022). Integrating remote sensing and street view imagery for mapping slums. *ISPRS Int. Journal of Geo-Info.*, 11(12):631.
Neumark, D. and Simpson, H. (2015). Place-based policies. *in Handbook of Regional and Urban Economics, vol. 5, Duranton G., Henderson J.V, and Strange W., eds.*, page 1197–1287.
Patel, A., Joseph, G., Killemsetty, N., and Eng, S. (2020). Effects of residential mobility and migration on standards of living in dar es salaam, tanzania: A life-course approach. *PLoS ONE*, 15(1):e0239735.
Picarelli, N. (2019). There is no free house. *Journal of Urban Economics*, 111:35–52.
Rains, E. and Krishna, A. (2020). Precarious gains: Social mobility and volatility in urban slums. *World Development*, 132:105001.
Rojas-Ampuero, F. and Carrera, F. (2023). Sent away: The long-term effects of slum clearance on children.
Ronneberger, O., Fischer, P., and Brox, T. (2015). U-Net: Convolutional networks for biomedical IS.
Selman, J. (2025). Information frictions in rental voucher programs: Experimental evidence from Chile.
Taubenböck, H., Staab, J., Zhu, X. X., Geiß, C., Dech, S., and Wurm, M. (2018). Are the poor digitally left behind? indications of urban divides based on remote sensing and Twitter data. *ISPRS IJGI*, 7(8).
UN-Habitat (2020). World Cities Report 2020: The value of sustainable urbanization. Tech. Rep.
Wong, M. (2019). Intergenerational mobility in slums: Evidence from a field survey in Jakarta. *Asian Development Review*, 36(1):1–19.
Wurm, M. and Taubenböck, H. (2018). Detecting social groups from space: Assessment of remote sensing-based mapped morphological slums using income data. *Remote Sensing Letters*, 9:41–50.
Wurm, M., Tubbesing, R., Stark, T., Kühnl, M., Sapena, M., Sulzer, W., and Taubenböck, H. (2023). Revealing landslide exposure of informal settlements in Medellín using deep learning. In *JURSE, pp.1-4*.

Tables and Figures

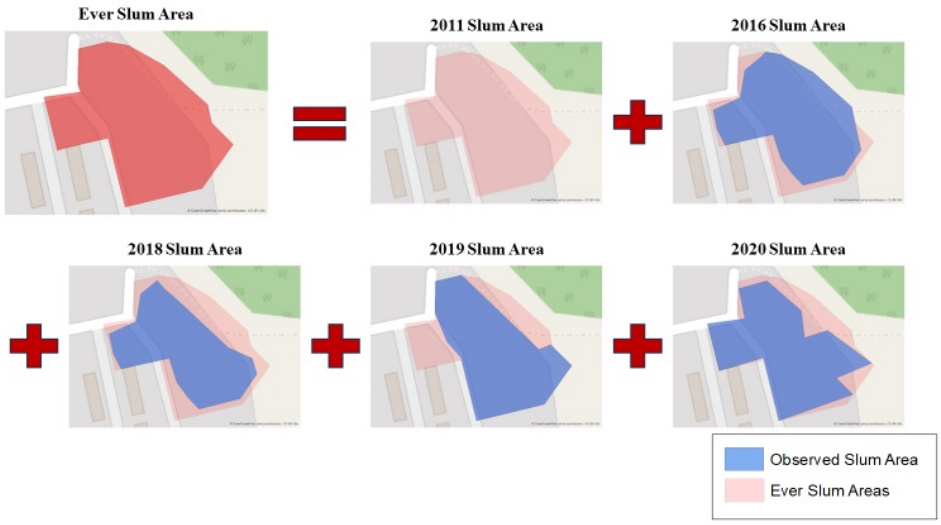
Table 1: *In-situ* Upgrading and Population Relocation Program Expenditures

In-situ Upgrading (N Slums = 209)	Mean	SD	Median
# Households (HH) per Slum	56.22	63.57	38
Share HH Receiving Housing Voucher	0.54	0.48	0.43
Voucher Value per Beneficiary HH	$36,672	$16,534	$34,487
Voucher Value per Slum HH	$19,580	$18,126	$15,917
Infrastructure Investment per Slum HH	$9,057	$14,389	$2,599
Total Expenditure per Beneficiary HH	$45,407	$24,080	$37,469
Total Expenditure per Slum HH	$28,636	$27,307	$23,021
Population Relocation (N Slums = 440)	**Mean**	**SD**	**Median**
# Households (HH) per Slum	34.43	46.92	23
Share HH Receiving Housing Voucher	0.82	0.63	0.79
Voucher Value per Beneficiary HH	$46,710	$2,867	$45,415
Voucher Value per Slum HH	$38,317	$29,087	$36,605
Infrastructure Investment per Slum HH	$4,215	$9,112	$773
Total Expenditure per Beneficiary HH	$50,881	$10,032	$47,209
Total Expenditure per Slum HH	$42,532	$32,620	$40,138

Data source: MINVU. Data covers the period 2011-2020. Monetary values are in 2017 US dollars. The original list of target slums was established in 2011. In 2018, the census was updated, leading to a corresponding update of the list of slums to be intervened. Households (HH) residing in slums correspond to the first slum census in which the slum appears. On population relocation, costs associated to infrastructure investment include slum closure efforts. For details on variable definitions, see Online Appendix Section B.

Figure 1: Unit of Analysis - Ever-Slum Areas

(a) Geographical Union of Observed Slum Areas over Time



(b) Satellite Images of the Fixed Slum Area



Notes: Representation of the ever-slum area of La Iglesia slum in the municipality of Tome. Panel (a) shows the slum boundaries provided by MINVU (2011 and 2019) and TECHO (2016, 2018, 2020). The slum did not exist in 2011 and was included for the first time in the 2016 slum census. The red shaded area is the ever slum area calculated as the union of all observed boundaries (blue polygons). Panel (b) shows satellite images from Google Earth with the ever slum area delineated in red.

Table 2: Distance to Amenities and Differences in Terrain, Slum vs. Non-Slum Census Areas

Dependent Variable:	Mean Dep. Var. Control (1)	Slum Indicator (2)	Romano-Wolf (p-value) (3)
Terrain Characteristics			
Slope (%)	2.68	1.19***	0.00
		(0.14)	
Elevation (m)	329	14.10***	0.00
		(1.98)	
Terrain Ruggedness Index (m)	74	13.55***	0.00
		(2.68)	
Distance Nearest River (m)	2,006	-339.37***	0.00
		(34.05)	
Distance to Amenities			
Amenities Distance index (sd)	0.00	0.06***	-
		(0.02)	
Distance Nearest Fire Station (m)	1,994	213.63**	0.03
		(89.67)	
Distance Nearest Police Station (m)	2,047	210.08***	0.01
		(61.86)	
Distance Nearest School (m)	637	200.99***	0.00
		(32.54)	
Distance Nearest Health Center (m)	1,438	50.42	0.27
		(42.88)	
Distance Nearest Bus Stop (m)	810	64.77**	0.14
		(28.75)	
Distance Nearest Supermarket (m)	2,900	225.41**	0.05
		(99.05)	
Distance Nearest Finance Institution (m)	2,777	524.13***	0.00
		(109.06)	

Notes: Unit of analysis are slums and non-slum census-blocks. Sample size is 156,654 census-blocks. Each row corresponds to a different OLS regression of the distance to the nearest amenity as a function of a slum indicator, the number of households, and Municipality fixed effects. Standard errors in parentheses are clustered at the Municipality level. Column 1 is the mean of the dependent variable for the non-slum areas. Column 2 shows the regression coefficient of the slum indicator. Column 3 presents the Romano-Wolf p-values correcting for multiple hypothesis testing. The Amenities Distance index follows Anderson (2008) and combines all reported individual amenities in the table. Amenities' locations are for 2022. School information includes build year. Distance to the nearest school reflects the proximity at the time the slum was established. For details on variable definitions, see Online Appendix Section B. * p<0.1, ** p<0.05, *** p<0.01.

Table 3: Direct Effect of Slum Renewal Policies —SDiD

Dependent variable:	Mean Control (1)	*In-situ* Upgrading β_1 (2)	Population Relocation β_2 (3)	p-value $\beta_1 = \beta_2$ (4)
Slum Characteristics				
Log Population	2,063	-0.04	-0.16***	0.00
		(0.10)	(0.06)	
Share Residential Land	0.26	-0.03	-0.12**	0.00
		(0.06)	(0.05)	
Share Streets paved	0.09	0.11**	0.04	0.00
		(0.05)	(0.03)	
Housing Investment				
Index Housing Investment (sd)	0.00	0.56**	0.14	0.00
		(0.25)	(0.11)	
Pr(Building Permits > 0)	0.01	0.05**	0.01	0.00
		(0.02)	(0.01)	
Log # Building Permits	0.01	0.02**	0.01	0.00
		(0.01)	(0.01)	
Housing Quality				
Index Housing Quality (sd)	0.00	0.09***	0.05	0.00
		(0.04)	(0.04)	
Avg. Distance to Neighbors (m)	60.05	8.67**	4.74	0.00
		(4.26)	(5.74)	
SD Bldg. Main Angle Across Neighbors	5.27	-0.82**	-0.09	0.00
		(0.41)	(0.38)	
Log # Buildings, size > 64m^2	4.60	0.15*	0.03	0.00
		(0.08)	(0.08)	

Notes: Unit of analysis is the slum. Number of observations are 432 slums in the control group, 209 slums in the *in-situ* upgrading group, and 383 slums in the population relocation group. Each coefficient is from a different SDiD regression (Arkhangelsky et al., 2021) for the dependent variable six years after the indicated treatment (columns 2 and 3). Column 1 shows the control group's dependent variable mean. SDiD identifies a synthetic control group from never-treated slums to match pre-treatment trends, accounting for slum fixed effects, time fixed effects, and lagged slum population data from LandScan (1 km^2 estimates). Inference is built using bootstrapping with replacement, creating a distribution for each estimated coefficient. Column (4) tests the equality of the effects for both policies using the bootstrapped distribution of each estimated effect. *Slum Characteristics* and *Housing Quality* variables come from satellite images. *Housing Investment* variables are from building permits. Indexes for each grouping follow Anderson (2008). Specifically, each outcome is standardized using the control group's mean and standard deviation and combined using inverse covariance weights. Variable signs are adjusted so that higher index values correspond to more desirable outcomes. Baseline mean corresponds to 2009, except for building permits, which is 2010. * p<0.1, ** p<0.05, *** p<0.01.

Table 4: Direct Effect of Slum Renewal Policies —TWFE

Dependent variable:	Mean Control (1)	*In-situ* Upgrading β_1 (2)	Population Relocation β_2 (3)	p-value $\beta_1 = \beta_2$ (4)
Sociodemographics				
Index Sociodemographics (sd)	0.00	0.14	-0.02	0.09
		(0.09)	(0.08)	
Log Population	298	0.14	-0.15	0.07
		(0.18)	(0.15)	
% Adults with $\geq$ High School	55.87	3.05*	1.24	0.32
		(1.62)	(1.58)	
% Employed	83.32	1.28	0.43	0.57
		(1.24)	(1.03)	
Housing				
Index Housing (sd)	0.00	0.15	-0.09	0.06
		(0.13)	(0.08)	
% Houses $\geq$ 2 Bedrooms	70.52	4.49*	0.99	0.15
		(2.60)	(1.84)	
% Houses with High Quality Walls	81.72	0.14	-4.50*	0.18
		(3.31)	(2.68)	

Notes: Unit of analysis is the slum. Number of observations are 196 slums in the control group, 97 slums in the *in-situ* upgrading group, and 168 slums in the population relocation group. While censuses attempt to cover the universe of the territory, targeting informal areas without formal addresses is challenging, and thus not all of them are included. As a result, our analysis sample does not correspond to the universe of slum areas, but only those that were identified in the censuses. Outcome variables are from the 2002 and 2017 population censuses. Column 1 shows the control group's dependent variable mean. Each row corresponds to a separate pooled TWFE regression of the dependent variable indicated against two dummy treatment variables, one for in-situ upgrading (column 2) and another for slums receiving population relocation (column 3), controlling for slum FE, year FE, and lagged LandScan population to account for differences in slum size and potential confounding trends. For the treated sample, only slums treated between 2011 and 2015 are included, which allows time for effects to materialize. Column (4) reports p-values from Wald tests to determine if the treatment coefficients are statistically similar. Indexes for each grouping follow Anderson (2008). Specifically, each outcome is standardized using the control group's mean and standard deviation and combined using inverse covariance weights. Variable signs are adjusted so that higher index values correspond to more desirable outcomes. Clustered standard errors at the Municipality level are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01

Table 5: Spillover Effects of Renewal Policies on Crime in Adjacent Formal Areas —SDiD

	0 - 200 m				200 - 500 m			
	Mean Control	*In-situ* Upgra. β_1	Pop. Reloc. β_2	p-value $\beta_1 = \beta_2$	Mean Control	*In-situ* Upgra. γ_1	Pop. Reloc. γ_2	p-value $\gamma_1 = \gamma_2$
Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Aggregate Crime Index (sd)	0.00	-0.12***	0.11	0.00	0.00	-0.05	0.11	0.00
		(0.03)	(0.18)			(0.05)	(0.17)	
Individual Crimes								
# Property Crime per km^2	44.93	-4.83*	15.79	0.00	48.35	0.17	29.86	0.00
		(2.61)	(13.06)			(2.74)	(25.74)	
# Robberies per km^2	11.55	-2.42**	-1.13	0.00	10.62	-0.52	-0.61	0.22
		(1.07)	(1.69)			(0.80)	(1.04)	
# Homicide per km^2	0.08	-0.08***	0.05	0.00	0.07	-0.03	0.03	0.00
		(0.03)	(0.13)			(0.03)	(0.08)	

Notes: Unit of analysis is the nearby neighborhoods of the slums. Sample sizes correspond to neighborhoods near 433 slums in the control group, 75 slums in the *in-situ* upgrading group, and 121 slums in the population relocation group. Crime outcomes are built based on geocoded data on police incident reports from the Under-Secretary for Crime Prevention (SPD), for the period 2013 to 2021. Since crime data is available from 2013, only slums treated between 2015 and 2018 are considered. Each coefficient is from a different SDiD regression (Arkhangelsky et al., 2021) for the dependent variable five years after the indicated treatment. Columns (2) and (3) show estimates for 0-200 meters spatial buffers surrounding slum boundaries, while columns (6) and (7) for 200-500 meters spatial buffers (areas located farther from the treated zones). Columns (1) and (5) show the control group's dependent variable mean. For each cohort of adjacent areas of treated slums, SDiD finds a synthetic control group using adjacent areas of never-treated slums to match the pre-treatment trends of the treated locations. In addition to matching the pre-treatment trends in the outcome of interest, SDiD also controls for individual FE, time FE, and the lag of slum population measured using LandScan (gridded population estimates for 1 km^2) based on the grid containing the slum polygon. Inference is constructed using bootstrapping with replacement, creating a distribution for each estimated coefficient. Columns (4) and (8) test for the null hypothesis of equality of the effects across the two policies using the bootstrapped distribution of each estimated effect. Property crimes include burglary, theft, auto theft, theft from non-residential areas, and other economically motivated crimes. Crime index follows Anderson (2008), meaning each variable in the individual crime list is standardized using the control group's mean and standard deviation and then combined using inverse covariance weights. * p<0.1, ** p<0.05, *** p<0.01.

Table 6: Spillover Effects of Slum Renewal Policies on Housing in Adj. Formal Areas —SDiD

	0 - 200 m				200 - 500 m			
Dependent variable:	Mean Control (1)	*In-situ* Upgra. β_1 (2)	Pop. Reloc. β_2 (3)	p-value $\beta_1 = \beta_2$ (4)	Mean Control (Level) (5)	*In-situ* Upgra. γ_1 (6)	Pop. Reloc. γ_2 (7)	p-value $\gamma_1 = \gamma_2$ (8)
Housing Investment								
Housing Investment Index (sd)	0.01	0.22***	0.14**	0.00	0.00	0.13***	0.06	0.00
		(0.05)	(0.06)			(0.03)	(0.04)	
% New Buildings (age<5)	12.33	4.04***	2.65**	0.00	12.84	2.50***	2.19**	0.00
		(1.33)	(1.18)			(0.92)	(1.10)	
% Renovated Buildings	0.95	0.32**	0.13	0.00	1.06	0.22***	0.17*	0.00
		(0.14)	(0.14)			(0.08)	(0.10)	
Pr(Building Permits > 0)	0.21	0.10**	0.07	0.00	0.42	0.03	0.00	0.00
		(0.04)	(0.05)			(0.03)	(0.03)	
Log # Building Permits	0.48	0.13**	0.12**	0.98	1.98	0.13**	0.04	0.00
		(0.05)	(0.05)			(0.05)	(0.04)	
Housing Quality								
Housing Quality Index (sd)	0.00	0.03	0.02	0.01	0.00	-0.02	0.016	0.00
		(0.04)	(0.02)			(0.05)	(0.03)	
Log Building Age	18.59	-0.11**	-0.07**	0.00	19.21	-0.02	-0.05*	0.00
		(0.05)	(0.03)			(0.03)	(0.03)	
Log Building Size	68.93	0.11*	0.06	0.00	89.91	0.08	0.07	0.03
		(0.06)	(0.05)			(0.09)	(0.06)	
Log #Buildings, size > 64m^2	14,108	-0.09	-0.01	0.00	22,971	-0.12	-0.04	0.00
		(0.08)	(0.08)			(0.10)	(0.07)	
SD Bldg. Main Angle Neighbors	9.30	0.35	-0.15	0.00	9.21	-0.10	0.09	0.00
		(0.24)	(0.19)			(0.30)	(0.20)	
Avg. Distance to Neighbors (m)	28.18	0.02	-0.18	0.49	29.33	-0.10	-0.20	0.03
		(0.87)	(0.77)			(1.10)	(0.80)	

Notes: Unit of analysis is the nearby neighborhoods of the slums. Sample sizes correspond to neighborhoods near 432 slums in the control group, 209 slums in the *in-situ* upgrading group, and 383 slums in the population relocation group. Each coefficient is from a different SDiD regression (Arkhangelsky et al., 2021) for the dependent variable five years after the indicated treatment. Columns (2) and (3) show estimates for 0-200 meters spatial buffers surrounding slum boundaries, while columns (6) and (7) for 200-500 meters spatial buffers (areas located farther from the treated zones). Columns (1) and (5) show the control group's dependent variable mean. Baseline mean corresponds to 2009, except for building permits, which is 2010. For each cohort of adjacent areas of treated slums, SDiD finds a synthetic control group using adjacent areas of never-treated slums to match the pre-treatment trends of the treated locations. In addition to matching the pre-treatment trends in the outcome of interest, SDiD also controls for individual FE, time FE, and the lag of slum population measured using LandScan (gridded population estimates for 1 km^2) based on the grid containing the slum polygon. Inference is constructed using bootstrapping with replacement, creating a distribution for each estimated coefficient. Columns (4) and (8) test for the null hypothesis of equality of the effects across the two policies using the bootstrapped distribution of each estimated effect. Building age, building size, and renovations come from the land registry dataset. Number large buildings, building main angle, and distance between buildings come from building footprints. Outcomes related to building permits are from housing starts. Indexes for each grouping follow Anderson (2008), meaning each variable within each family of crime outcomes is standardized using the control group's mean and standard deviation and then combined using inverse covariance weights. * p<0.1, ** p<0.05, *** p<0.01.

Table 7: Spillover Effects of Slum Renewal Policies on Adjacent Formal Areas —TWFE

	0 - 200 m				200 - 500 m			
Dependent variable:	Mean Control (1)	*In-situ* Upgra. β_1 (2)	Pop. Reloc. β_2 (3)	p-value $\beta_1 = \beta_2$ (4)	Mean Control (Level) (5)	*In-situ* Upgra. γ_1 (6)	Pop. Reloc. γ_2 (7)	p-value $\gamma_1 = \gamma_2$ (8)
Sociodemographics								
Sociodemographic Index (sd)	0.00	0.07	-0.11***	0.00	0.00	0.05	-0.11***	0.00
		(0.04)	(0.04)			(0.03)	(0.03)	
% Adults with ≥ High School	57.68	1.47*	0.42	0.30	59.97	1.43*	-0.24	0.04
		(0.87)	(0.92)			(0.76)	(0.63)	
% Employed	84.76	1.18*	-0.67	0.01	85.20	0.86*	-0.68*	0.00
		(0.62)	(0.55)			(0.45)	(0.41)	
Log Population	1,117	-0.04	-0.21***	0.01	5,624	-0.07	-0.18***	0.05
		(0.07)	(0.06)			(0.06)	(0.05)	
Housing								
Housing Index (sd)	0.00	0.14**	-0.01	0.05	0.00	0.15***	0.01	0.02
		(0.06)	(0.04)			(0.05)	(0.04)	
% Houses ≥ 2 Bedrooms	76.17	0.51	0.38	0.89	77.38	1.49**	-0.71	0.01
		(0.86)	(0.61)			(0.70)	(0.50)	
% Houses with High Quality Walls	84.27	3.62**	-0.79	0.03	84.54	2.15*	1.22	0.61
		(1.47)	(1.32)			(1.24)	(1.12)	

Notes: Unit of analysis is the nearby neighborhoods of the slums. Sample sizes correspond to neighborhoods near 509 slums in the control group, 137 slums in the *in-situ* upgrading group, and 276 slums in the population relocation group, and they correspond to the universe of formal areas surrounding the slums, as all of them are identified in the censuses. Outcome variables are from the 2002 and 2017 population censuses. Each row corresponds to a separate pooled TWFE regression of the dependent variable indicated against two dummy treatment variables, one for in-situ upgrading and another for slums receiving population relocation, controlling for slum FE, year FE, and lagged LandScan population to account for differences in slum size and potential confounding trends. Columns (2) and (3) show estimates for 0-200 meters spatial buffers surrounding slum boundaries, while columns (6) and (7) for 200-500 meters spatial buffers (areas located farther from the treated zones). For the treated sample, only slums treated between 2011 and 2015 are included, which allows time for effects to materialize. Columns (1) and (5) shows the control group's dependent variable mean. Columns (4) and (8) test for the null hypothesis of equality of the effects across the two policies using the bootstrapped distribution of each estimated effect. Indexes for each grouping follow Anderson (2008). Specifically, each outcome is standardized using the control group's mean and standard deviation and combined using inverse covariance weights. Variable signs are adjusted so that higher index values correspond to more desirable outcomes. Clustered standard errors at the Municipality level are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table 8: Determinants of Treatment Assignment

	Pr(Any Treatment)			Pr(In-situ Upgrading | Any Treatment)		
	Marginal Effect (at the mean)		Covariate Mean	Marginal Effect (at the mean)		Covariate Mean
	(1)	(2)	(3)	(4)	(5)	(6)
Slum Characteristics						
Identified as Slum in 2011	0.88***	0.97***	0.50	0.05	0.08	0.84
	(0.06)	(0.05)		(0.10)	(0.11)	
Log Slum Area	0.00	0.04	9.23	0.08***	0.07***	9.08
	(0.02)	(0.03)		(0.02)	(0.02)	
Log Slum Households	-0.11***	-0.15***	3.45	0.05*	0.07*	3.37
	(0.03)	(0.04)		(0.03)	(0.04)	
Terrain Characteristics						
Rugged (TRI > 497)	-0.00	0.13	0.02	-0.31	-0.09	0.02
	(0.14)	(0.18)		(0.22)	(0.18)	
Steep (Slope > 20)	-0.33***	-0.42***	0.02	0.07	0.11	0.01
	(0.06)	(0.14)		(0.23)	(0.23)	
In Risk Location	0.04	-0.02	0.60	-0.13**	-0.11***	0.63
	(0.06)	(0.06)		(0.06)	(0.04)	
Distance Nearest River (km)	0.01	0.01*	1.12	-0.00	0.04**	1.21
	(0.01)	(0.01)		(0.01)	(0.01)	
Terrain Ruggedness Index - TRI (km)	0.01	-0.23	0.12	0.64***	0.50**	0.11
	(0.25)	(0.24)		(0.20)	(0.21)	
Location & Muni Characteristics						
Located at City Border or Beyond	-0.10**	-0.06	0.78	0.11*	0.11***	0.73
	(0.05)	(0.05)		(0.06)	(0.04)	
Distance to City Center (km)	0.002**	0.003**	18.51	-0.002**	-0.001	20.25
	(0.00)	(0.00)		(0.00)	(0.00)	
Amenities Distance Index	-0.00	-0.00	0.00	-0.01	-0.04	0.03
	(0.02)	(0.03)		(0.03)	(0.04)	
Building Density within 200m	0.62*	0.58*	0.10	-0.08	-0.33	0.11
	(0.35)	(0.35)		(0.24)	(0.28)	
Log Municipality Expenditures	-0.01	0.10***	16.57	-0.07***	-0.02	16.54
	(0.01)	(0.03)		(0.01)	(0.03)	
Obs.	1,102	1,102	1,102	546	546	546
Region FE	No	Yes		No	Yes	
Mean dependent var	0.50	0.50		0.33	0.33	

Notes: Unit of analysis are slums. Columns (1) displays the marginal effects from a logit model on the probability of receiving any treatment, where the outcome is a dummy that equals 1 if the slum was treated by either policy and zero otherwise. Column (2) additionally control for region fixed effects. Each row is a different covariate. Column (3) shows the mean of each covariate, which helps to interpret the order of magnitude of the marginal effects. Columns (4) and (5) replicate the analysis but using only treated slums, where the outcome is a dummy that equals 1 if the slum was treated through *in-situ* upgrading, and zero if treated through population relocation. For details on variable definitions, see Online Appendix Section B. Standard errors clustered at the region level are shown in parenthesis. * p<0.1, ** p<0.05, *** p<0.01.

Table 9: A Direct Comparison of the Two Policies

Dependent variable:	Treated = In-situ Upg. Control = Pop. Relocation (1)	Treated = Pop. Relocation Control = *In-situ* Upg. (2)
Slum Characteristics		
Index Slum Characteristics (sd)	-0.10*	0.09*
	(0.05)	(0.05)
Log Population	0.14**	-0.10
	(0.07)	(0.07)
Share Residential Land	0.13***	-0.13***
	(0.03)	(0.03)
Share Streets paved	0.06***	-0.05**
	(0.02)	(0.02)
Housing Investment		
Index Housing Investment (sd)	0.38**	-0.28***
	(0.19)	(0.10)
Pr(Building Permits > 0)	0.04	-0.04*
	(0.03)	(0.023)
Log # Building Permits	0.03***	-0.04**
	(0.01)	(0.01)
Housing Quality		
Index Housing Quality (sd)	0.03	-0.02
	(0.04)	(0.05)
Log # Buildings, size > 64m^2	0.23***	-0.19**
	(0.08)	(0.08)
SD Bldg. Main Angle Across Neighbors	-0.90*	0.95*
	(0.52)	(0.48)
Avg. Distance to Neighbors (m)	-8.88*	8.14
	(4.56)	(5.11)

Notes: Unit of analysis is the slum. Only treated slums are considered. Number of observations are 197 slums in the *in-situ* upgrading group, and 407 slums in the population relocation group. Each coefficient is from a different SDiD regression (Arkhangelsky et al., 2021) for the dependent variable six years after the indicated treatment. Column 1 presents the results using in-situ upgrading slums as treated units and relocation slums as synthetic controls; Column 2 reverses the roles. SDiD identifies a synthetic control group from slums treated by the alternative policy to match pre-treatment trends, accounting for slum fixed effects, time fixed effects, and lagged slum population data from LandScan (1 km^2 estimates). Inference is built using bootstrapping with replacement, creating a distribution for each estimated coefficient. *Slum Characteristics* and *Housing Quality* variables come from satellite images. *Housing Investment* variables are from building permits. Indexes for each grouping follow Anderson (2008). Specifically, each outcome is standardized using the control group's mean and standard deviation and combined using inverse covariance weights. Variable signs are adjusted so that higher index values correspond to more desirable outcomes. * p<0.1, ** p<0.05, *** p<0.01.

Online Appendix

Table of Contents

A Flowchart of Government Interventions **47**

B List of Variables and Datasets **48**

C Identify Building Footprints **52**

C.1 Pre-processing images 53

C.2 ML Algorithm: Deep Neural Networks - U-Net Architecture 53

C.3 Model calibration 54

C.4 Computing measures using building footprints 55

C.5 Example building footprints 55

D Dynamics of SDiD Effects **60**

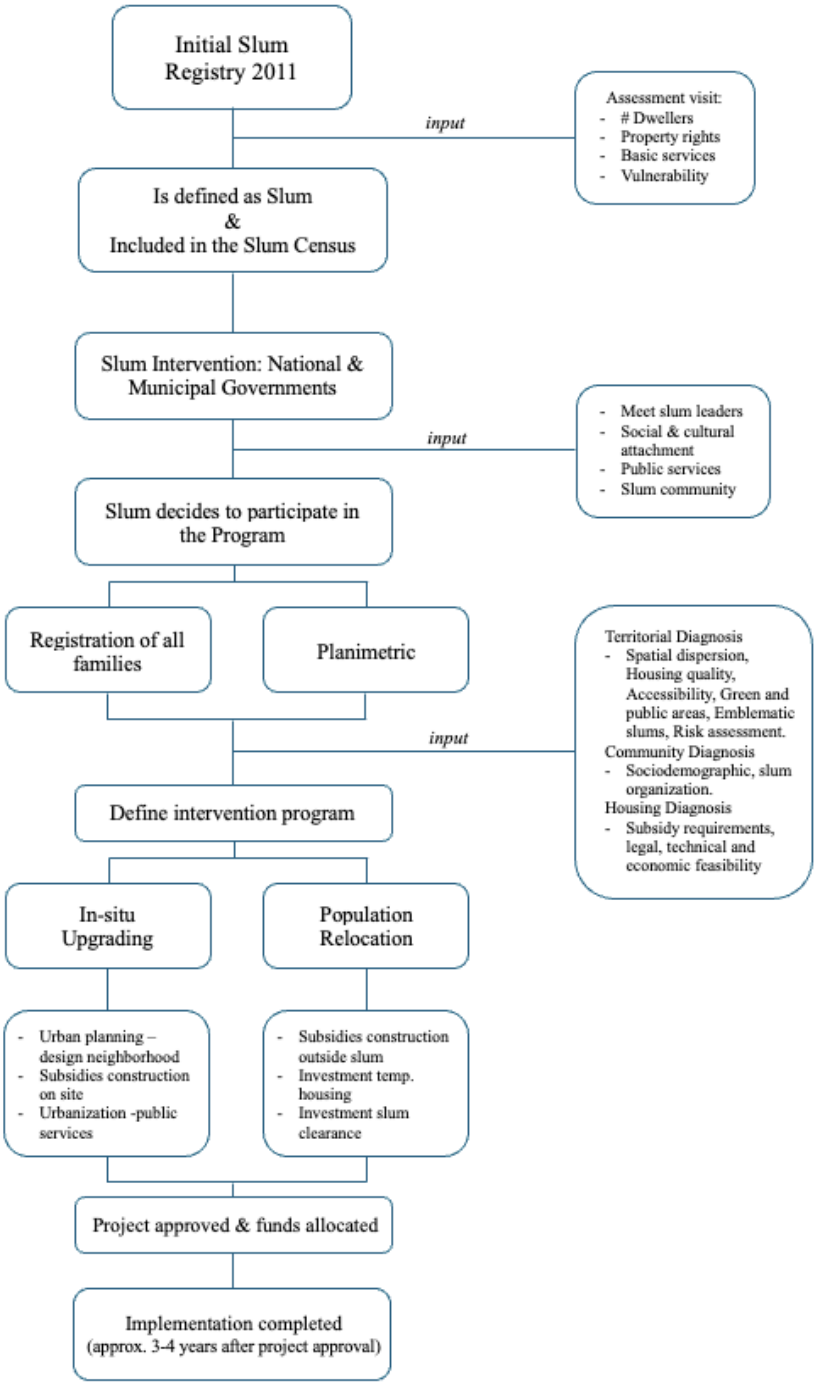
E Counterfactuals **67**

F Location Characteristics: Amenities & Terrain **69**

G Additional Appendix Figures and Tables **72**

A Flowchart of Government Interventions

Figure A1: Flowchart Government Intervention



Source: Elaborated by the authors based on the different law regulations and the DIPRES (2019) report.

B List of Variables and Datasets

Table B1: List of Variables and Data Sources

Variable	Definition	Units	Source	Years	Tables
# Building Permits	Number of building permits registered inside the slum polygon	Count	Instituto Nacional de Estadistica (INE)	2010-2021	5, 8, 11, E2
# Buildings size > 64m2	Number of building footprints larger than 64 square meters	Count	Building footprints - Satellite images	2000-2021	5, 8, 11, E2
# employees	Count of employees in the municipality	Count	Encuesta Nacional de Empleo - INE	2010-2020	4
# employees - high school or more	Count of employees with high school education or more in the municipality	Count	Encuesta Nacional de Empleo - INE	2010-2020	4
# Assaults per km2	Assaults divided by the area in square kilometers	# per square kilometer	Subsecretaria de Prevencion del Delito	2013-2022	7, E1
# Homicide crime per km2	Homicides divided by the area in square kilometers	# per square kilometer	Subsecretaria de Prevencion del Delito	2013-2022	7, E1
# Property crime per km2	Property crimes (burglary, theft, auto theft, theft from non-residential areas, and other economically motivated crimes) divided by the area in square kilometers	# per square kilometer	Subsecretaria de Prevencion del Delito	2013-2022	7, E1
% adults high-school or more educ	Percentage of adults over 25 years with secondary education completed or more	Percentage	Population Census	2002, 2017	6, 9
% Employed	Percentage of the population of working age that is employed	Percentage	Population Census	2002, 2017	3, 6, 9
% houses with high quality walls	Percentage of houses that have concrete, stone, brick, or steel as predominant wall materials	Percentage	Population Census	2002, 2017	6, 9
% houses with two or more bedrooms	Percentage of houses with two or more bedrooms	Percentage	Population Census	2002, 2017	6, 9
% new buildings	Percentage of buildings constructed in the last 5 years	Percentage	Servicio de Impuestos Internos (SII) Cadastro	2009-2024	8, E2
% population with secondary education	Percentage of adults over 25 years with secondary education completed (high school graduates)	Percentage	Population Census	2002, 2017	3
% population with tertiary education	Percentage of adults over 25 years with some tertiary education (vocational training, 2yrs training, colleges)	Percentage	Population Census	2002, 2017	3
% renovated buildings	Percentage of buildings that undergo renovations during a given year (reported alterations to building characteristics)	Percentage	Servicio de Impuestos Internos (SII) Cadastro	2009-2024	8, E2
Aggregate crime index	Anderson (2008) index combining number of property crimes, assaults and homicides per square kilometer	Standard deviations	Subsecretaria de Prevencion del Delito	2013-2022	7, E1
Amenities Distance Index	Anderson (2008) index combining distance to the nearest bus stop, supermarket, library, fire station, police station, school, health center, and finance institution.	Standard deviations	See each component data source	circa 2022	2, 10
Average distance to neighbors	Average distance between a building footprint and its nearest eight neighbors	Meters	Building footprints - Satellite images	2000-2021	5, 8, 11, E2
Building age	Average age of buildings	Years	Servicio de Impuestos Internos (SII) Cadastro	2009-2024	8, E2

Table B2: List of Variables and Data Sources

Variable	Definition	Units	Source	Years	Tables
Building Density within 200 m	Total building footprints area in formal neighborhoods within 200 m of the slum over the area	Share	Building footprints - Satellite images	2000-2021	10
Building Permits > 0	Dummy variable equals one if there are at least one building permit registered inside the slum polygon, zero otherwise	Indicator	Instituto Nacional de Estadistica (INE)	2010-2021	5, 8, 11, E2
Building size	Average reported building size	Square meters	Servicio de Impuestos Internos (SII) Cadastro	2009-2024	8, E2
Distance to City Center	Euclidean distance to the nearest city center	Kilometers	Urban Functional Areas project	2021	10
Distance to Nearest Bus Stop	Euclidean distance to the nearest bus stop	Kilometers	Geoportal IDE	2022	2
Distance to Nearest Finance Institution	Euclidean distance to the nearest finance institution (bank, atms, money transfer, or post office)	Kilometers	OpenStreetMap	2023	2
Distance to Nearest Fire Station	Euclidean distance to the nearest fire station	Kilometers	Geoportal IDE MINVU	2022	2
Distance to Nearest Health Center	Euclidean distance to the nearest health center	Kilometers	Geoportal IDE MINVU	2022	2
Distance to Nearest Police Station	Euclidean distance to the nearest police station	Kilometers	Geoportal IDE MINVU	2022	2
Distance to Nearest River	Euclidean distance to the nearest water body (river, ravine or estuary)	Kilometers	Biblioteca del Congreso Nacional de Chile	2018	2, 10
Distance to Nearest School	Euclidean distance to the nearest school by the time the slum was founded	Kilometers	Centro de Estudios MinEduc	2000-2023	2
Distance to Nearest Supermarket	Euclidean distance to the nearest supermarket	Kilometers	Geoportal IDE	2022	2
Elevation	Distance above sea level	Kilometers	View Finder Panoramas & DIVA-GIS	2022	2
Elevation Slope, percentage rise	Percent of slope is determined by dividing the amount of elevation change by the amount of horizontal distance covered	Percentage	View Finder Panoramas & DIVA-GIS	2022	2, 10
Housing subsidies in slums	Count of the number of housing subsidies allocated to a slum	Count	MINVU	2011-2020	1
Identified as Slum in 2011	Dummy variable equals one if the slum is included in MINVU 2011 census of slums, zero otherwise	Indicator	MINVU Slum Census	2011	10
In Risk Location	Dummy equals 1 if the slum is in an area at risk of tsunami, volcano, wildfire, or flooding	Indicator	Visor Chile Preparado & ARClim	2024 & 1980-2014	10
Investment in slums	Amount of direct investments allocated to a slum	Currency	MINVU	2011-2020	1
Labor salary	Payments received as compensation of work for those employed (not including self-employment)	Currency	Encuesta Suplementaria de Ingresos - INE	2010-2020	4
Labor salary - High school or more	Payments received as compensation of work for those employed (not including self-employment) with high school education or more	Currency	Encuesta Suplementaria de Ingresos - INE	2010-2020	4
Located at City Border or Beyond	Dummy variable equals one if the normalized distance to the city center is greater or equal 0.9	Indicator	Urban Functional Areas project	1993, 2011, 2017, 2019, 2020	10

Table B3: List of Variables and Data Sources

Variable	Definition	Units	Source	Years	Tables
Municipality expenditures	Total expenditures by the local municipality	Currency	SINIM	2011-2020	4, 10
Municipality population	Population in the municipality	Count	Population Census	2002, 2017	4
Population - Census	Count population using census blocks. Census blocks are assigned to each area based on their centroid.	Count	Population Census	2002, 2017	3, 4, 6, 9
Population - LandScan	Estimated count of population in slum areas using 1km grid population estimates based on the grid where the slum is located	Count	Landscan	2000-2021	5, 11
Population - WorldPop	Estimated count of population in slum areas using 100m grid population estimates. Grids are assigned to each area based on their centroid	Count	WorldPop	2000-2021	4
Quality adjusted rent	Municipal FE from a hedonic regression of rent including house size, bedrooms, bathrooms, basic services, and roof materials. FE are adjusted to match mean rent	Currency	CASEN	2009-2017	4
Ratio Adult Males/Females	Number of males between 25-64 years old over females between 25-65 years old	Ratio	Population Census	2002, 2017	3
SD bldg main angle across neighbors	Standard deviation between a building footprint's main angle and its eight nearest neighbors. Building main angle is the direction angle of the minimum bounding rectangle containing the building footprint	Standard deviations	Building footprints - Satellite images	2000-2021	5, 8, 11, E2
Share residential land	Percentage of the slum polygon covered by residential areas. This comes from the question:	Share	HOD	2000-2021	5, 11
Share streets paved	Percentage of the streets inside the slum polygon that are paved. This comes from the question:	Share	HOD	2000-2021	5, 11
Slum Area	Size of the area ever occupied by a given slum (constant fixed boundary)	Square meters	MINVU & Techo		10
Slum Households	Number of households living in the slum	Number	MINVU & Techo	2011-2020	1, 10
Terrain Ruggedness Index (TRI)	TRI is used to express the amount of elevation difference between adjacent cells of a DEM.	Kilometers	Nunn and Puga	circa 2000	2, 10
Value of housing voucher	Estimated expenditures in housing vouchers:	Currency	MINVU	2011-2020	1

C Identify Building Footprints

Identifying building structures has become an active area of research due to the increasing availability of high-resolution satellite images. Over the past decade, computer scientists have made significant advancements in structural identification techniques. Kuffer et al. (2016) provide a comprehensive literature review on slum mapping using remote sensing methods from 2000 to 2015, highlighting that machine learning (ML) methods are particularly effective for identifying the physical attributes of slums. Abascal et al. (2022) used deep neural networks (DNNs) to extract building footprints from different areas in Nairobi, achieving a precision level of 92% —meaning that 92% of all positive predictions were indeed true positives. The high level of accuracy demonstrates the advantages of ML algorithms in advancing urban research. Comparisons between satellite images footprints and official data sources further support the accuracy of these methods (Wurm and Taubenböck, 2018).

Another sign of the improvements in the identification of building footprints is that some new research in computer science has started focusing on other attributes of slums beyond those observed through satellite images. Taubenböck et al. (2018) focus on digital inclusion using Twitter data, and they find that slums' households have less access to the digital world than formal areas. In the case of Klotz et al. (2017), they use location-based social networks data in conjunction with satellite images to identify urban neighborhoods and study their socioeconomic characteristics. Closer to the insights from this paper, Wurm et al. (2023) use Convolutional Neural Network (CNN) for the detection of informal buildings to estimate the number of the exposed population to landslides.

One key factor driving advancements in ML techniques is the emergence of programming challenges hosted by online communities such as Kaggle and AI Crowd. These platforms incentivize innovation by encouraging researchers to develop and share cutting-edge algorithms. The baseline code for our ML algorithm originates from the winners of the first Mapping Challenge hosted by AI Crowd (Mohanty et al., 2020). Their algorithm, which achieved an average precision of 94%, employed strategies similar to those documented by Abascal et al. (2022), further illustrating the collaborative progress in this research area.[33] We explain below the four-stage process to identify polygons representing building footprints.

[33]The winning team members are Jakub Czakon, Kamil A. Kaczmarek, Andrzej Pyskir, and Piotr Tarasiewicz. Complete details about the Mapping Challenge can be found here., and replication materials for the winner team are available here.

C.1 Pre-processing images

Our baseline ML algorithm was trained on 300×300 pixel images containing various building structures similar to those depicted in Figure C1. However, Google Earth (GE) images have dimensions of 4800×4800 pixels, covering approximately $1, \text{km}^2$. This is the highest possible resolution in GE and represents approximately 0.21 meters per pixel. To process these images, we fragmented each original GE image into 256 tiles, each measuring 300×300 pixels. We then applied the model to all tiles and, after obtaining the building footprint predictions, reconstructed the full-size image.

We also experimented with fragmenting the original image into tiles of 100, 600, 900, 1,200, and 2,400 pixels, but the model performance was consistently better with 300-pixel tiles. A similar approach is used by Abascal et al. (2022). As an additional robustness check, we adjusted the original zoom level of the image, varying the meters per pixel between 0.16 and 0.3 meters, but the best performance was still achieved using the original dimensions.

Applying the model to all 90,000 images is extremely time-consuming and requires massive computational power. To reduce the time and resource demands, we selected one image per year for each slum using two approaches:

1. *Human Observational Data (HOD):* Analysts selected the first image of a given year with sufficient quality to clearly observe construction details within a slum, reducing the number of images per slum to 20 or fewer.

2. *File Size Selection:* We chose the image with the largest file size within the first three months of each year, under the assumption that a larger file size correlates with higher image quality.

As expected, the results from these two approaches differed only in a few instances. In cases of discrepancy, we retained the older image as long as its file size exceeded a certain threshold (usually 8,000 KB).

By implementing these strategies, we efficiently managed computational resources while maintaining the quality and consistency of our dataset. This approach allowed us to effectively apply the ML algorithm to a manageable number of images, ensuring accurate building footprint predictions across different slum areas over time.

C.2 ML Algorithm: Deep Neural Networks - U-Net Architecture

The U-Net architecture was initially proposed by Ronneberger et al. (2015) and has proven effective in identifying building footprints (Iglovikov et al., 2018; Abascal et al., 2022). This

architecture consists of a network with a downsampling-upsampling structure. The first part of the method starts from a given image to which different convolutional functions are applied and is downsampled using max-pooling functions. Each time an image is processed by a set of convolutional and max-pooling functions, the dimensions of the data (or original image size) are reduced. The number of blocks in a U-Net network determines the number of times this process is performed. The second part of the process uses the output from the first part and up-samples it until it reaches the original image resolution. By increasing the dimensions of the output, the network allows propagation of contextual information to higher-dimensional layers. Propagation of contextual information is particularly important for identifying high-resolution features.

Additionally, our baseline model implements the strategies proposed by Iglovikov et al. (2018) that enable identification using only three-band images (such as RGB). Previous satellite image analyses have been restricted to the use of multispectral images. The model also incorporates Test Time Augmentation (TTA), meaning that predictions were made on the original image and several of its transformations, such as rotations (90°, 180°, 270°) and flips (vertical, horizontal). U-Net was implemented in Python using the deep learning libraries TensorFlow and PyTorch. All models were run on a virtual machine with access to an NVIDIA Tesla V100 GPU.

C.3 Model calibration

The training images used for the ML algorithm include building structures in different scenarios and contexts. However, this does not guarantee that the model will perform exceptionally well in the Chilean context, especially considering our focus on informal settlements. To address this issue, we calibrated the baseline algorithm for our research purposes. Two main changes were implemented. First, we allowed for more geometric irregularity. Informal constructions are not regular in shape, and restricting the identification of building footprints to regular geometric shapes omits some of the characteristics we aim to observe. Second, we adjusted the masking cutoff, which refers to the threshold at which an observed feature is labeled as a building footprint. We relaxed this restriction because not all of our images have comparable quality to the training sample.

There is a concern that when lowering the threshold, we will get more false positive building footprints. To minimize this effect, we pass every image twice through our ML algorithm. For the first pass, we divide an image into 256 tiles as shown in Figure C1; for the second pass, we shift the 300-pixel grid layout by 100 pixels. The final prediction is the combination of the two procedures.

Unlike other ML exercises, our purpose extends beyond image analysis. Consequently, we allocated substantial effort to retaining the geographical reference of each image. Maintaining a spatial reference is crucial for building a complete dataset that includes the demographic and socioeconomic variables of the slums. We export each masked image as a shapefile using the 1984 World Geodetic System projection.

C.4 Computing measures using building footprints

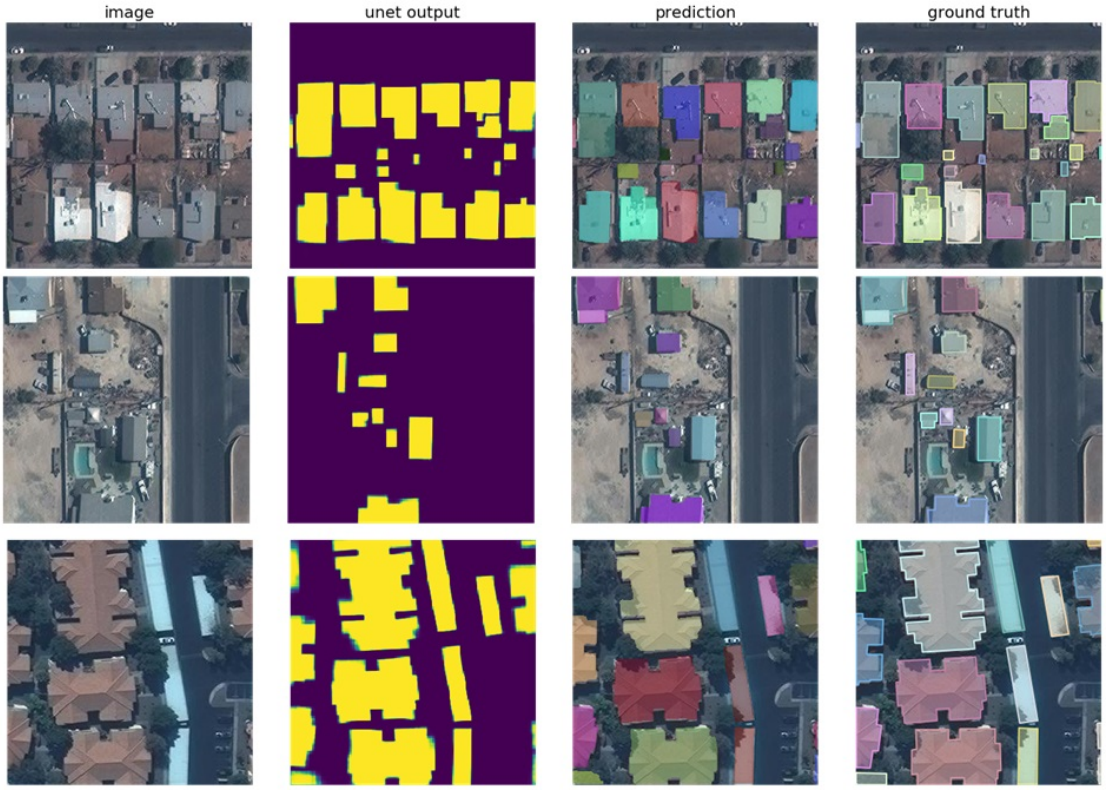
We computed different metrics to measure building activity and regularity inside and outside the slums. We began by calculating the total building area and building density within the slums. We also calculated these measures for areas within a certain distance of the slum border. Additional statistics helped us characterize the built environment inside the slum. The distance to the nearest building footprint provides a proxy for how compact and spatially close the structures are within the slum. We also obtained the orientation angle of each building using the main angle of the minimum bounding rectangle for a given footprint. This measure serves as a proxy for building regularity. In formal neighborhoods, buildings typically face the same direction, which means that the standard deviation of building main angles in that location is close to zero.

C.5 Example building footprints

Figure C2 provides an example of the building footprint prediction for a slum in the Lampa municipality. When processing images using our ML algorithm, we did not differentiate between slum and non-slum locations. By doing so, we ensure that any measurement error due to image quality and other specific image characteristics is homogeneous across slums and nearby neighborhoods. Although Figure C2 demonstrates how accurate the ML algorithm can be, identifying every single building with its exact boundaries is not feasible. There are instances where an existing building is not perfectly delineated or is not identified by the algorithm. Similarly, there are cases where areas are labeled as building footprints when there is actually no building structure in that location. However, these instances do not represent the majority of the identified areas. In fact, the algorithm has a precision level of 94% and a recall rate of 95%.[34] There are about 1,500 slums for which we have geographic references that allow us to obtain satellite images. Processing images of these slums, in addition to calibrating the model, took us over one year.

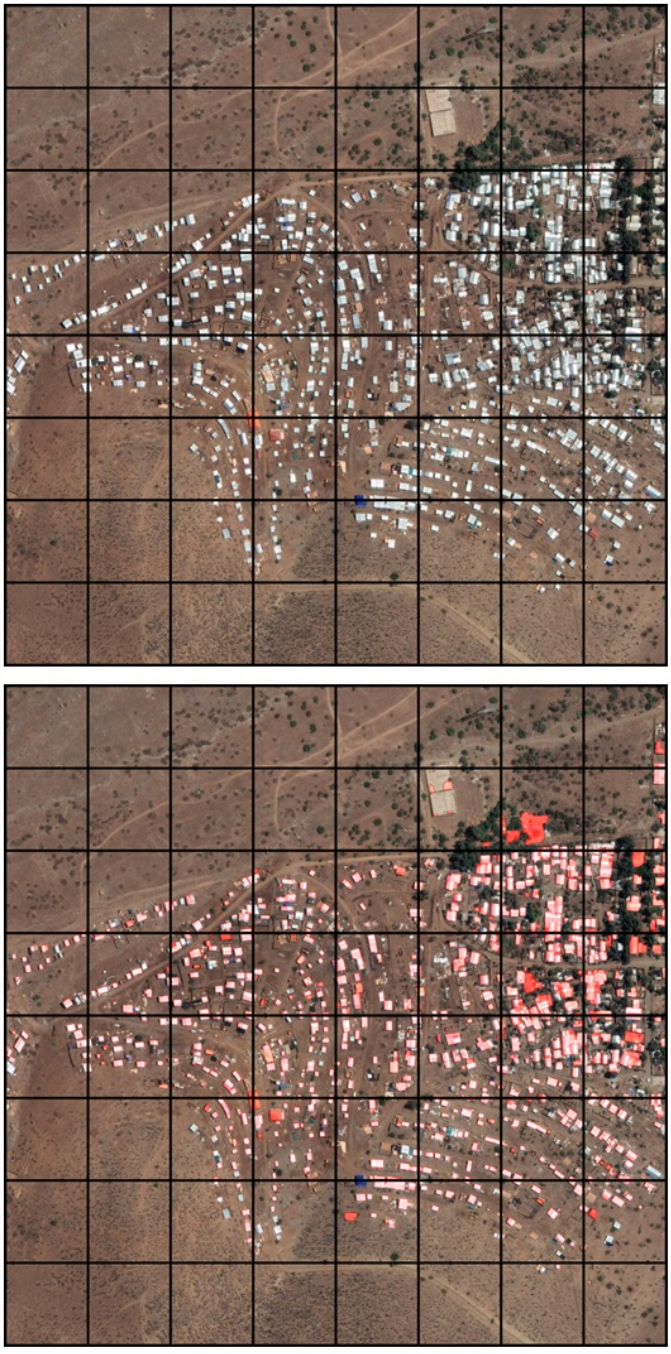
[34]Precision is calculated as the number of true predicted building footprints over the total number of predicted footprints. Recall has the same numerator but uses the actual number of building footprints as the denominator.

Figure C1: Performance ML Algorithm to Identify Building Footprints



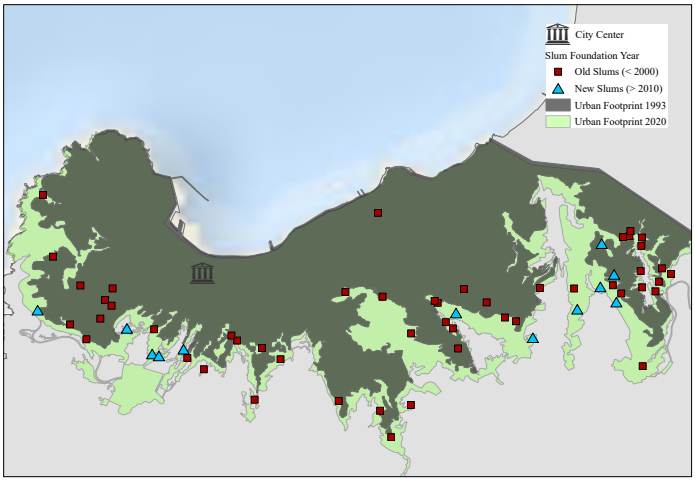
Source: Czakon et al. (2021).

Figure C2: Example Predicted Building Footprints - Slum in Santiago Metropolitan Area



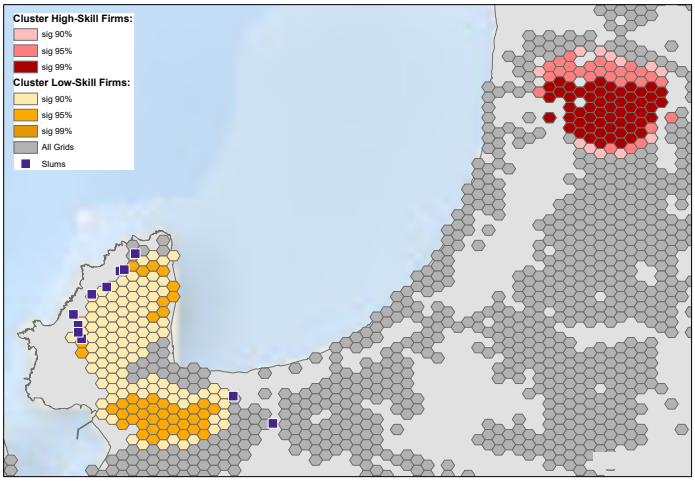
Notes: These images belong to the Slum Bosque Hermoso in the Lampa municipality in the Metropolitan Area of Santiago. The picture of from February 28, 2021. The original image in the top is 4,800 x 4,800 pixels and to improve prediction we fragmented the picture into tiles of 300 x 300 pixels. We applied the ML algorithm to each tile and reconstructed the original picture as it is shown in the bottom image. For exposition purpose we only considered 64 different tiles.

Figure C3: Location of Slums in Valparaiso, Urban Footprints in 1993 vs. 2020



Notes: City boundaries for 1993 and 2020 refer to the city continued urban construction. City center is obtained from the Urban Functional Areas project (MINVU, 2021). Slums that existed before 2000 are marked in red squares, while those that emerged after 2010 are marked in light blue triangles.

Figure C4: Proximity of Slums to Centers of Low- and High-Skill Employment, Coquimbo

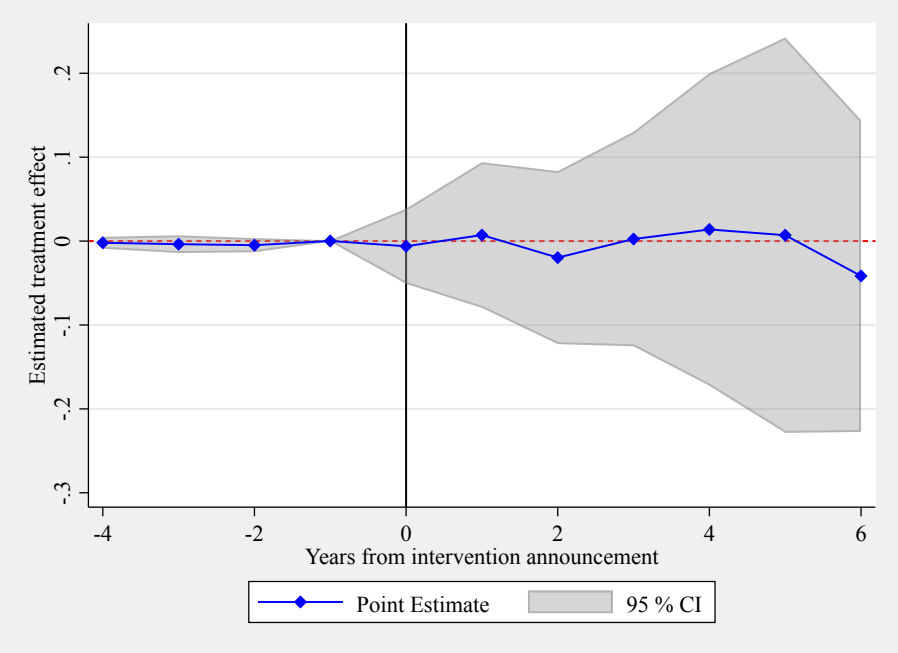


Notes: This figure illustrates spatial clusters of high and low-skills firms and the location of slums in the city of Coquimbo. Similar patterns are observed in other large cities of Chile (not shown here, available upon request). Firms' location comes from the 2017 National Firms Census. The territory is divided into hexagons with sides of approximately 125 m and count the number of firms within each hexagon. Clusters are identified through hot-spot analysis, which detects statistically significant spatial clusters of high and low concentrations of firms at the 90%, 95%, and 99% CIs (see ArcGIS description). Low-skill industries are agriculture, mining, manufacturing (food, beverages, leather, wood, paper, plastic, minerals, and furniture), construction, retail, and transportation. High-skill industries are other manufacturing (chemicals, pharmaceuticals, electronics, machinery, automobiles), services, teaching, professional and scientific activities.

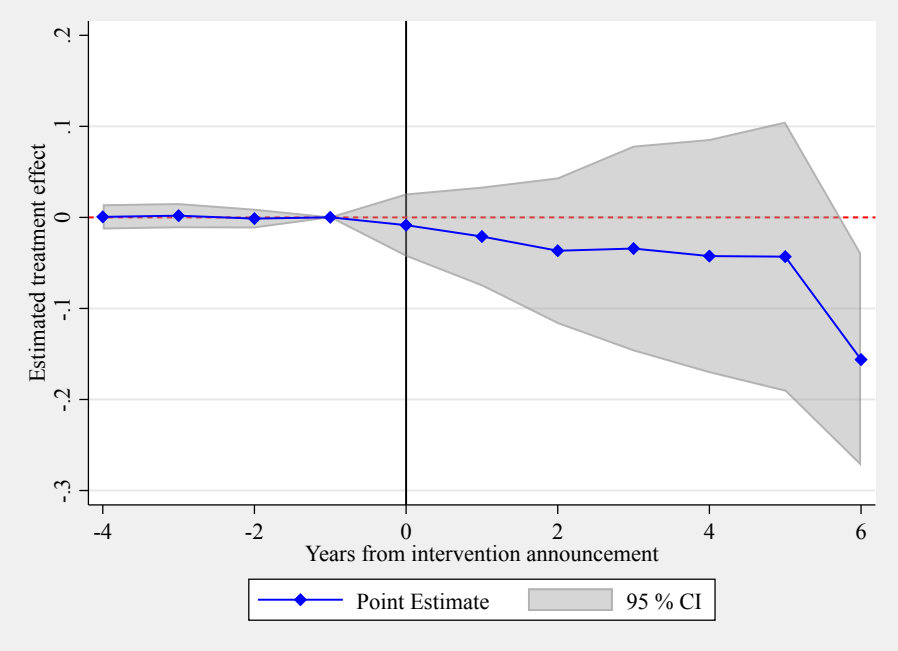
D Dynamics of SDiD Effects

Figure D1: Direct Effects of Slum Renewal Policies on Log Slum Population (SDiD)

(a) Effects of *In-situ* Upgrading on Log Slum Population



(b) Effects of Population Relocation on Log Slum Population

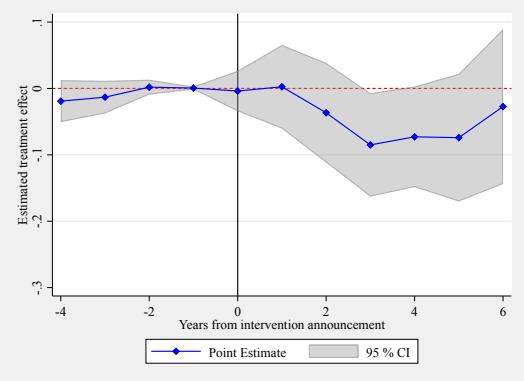


Notes: These are event-study graphs showing the direct effects of *in-situ* upgrading (Panel a) and population relocation (Panel b) on the log of slum population using the SDiD estimator (Arkhangelsky et al., 2021). Slum population is obtained through LandScan (gridded population estimates for 1 km^2) based on the grid containing the slum polygon. For each cohort of treated slums, we find a synthetic control group using never-treated slums to match the pre-treatment trends of treated slums. Once the synthetic control is chosen the standard DiD interpretation applies. Inference is constructed using bootstrapping with replacement. Year zero is the moment the intervention is determined. Only slums treated between 2011 and 2017 are considered.

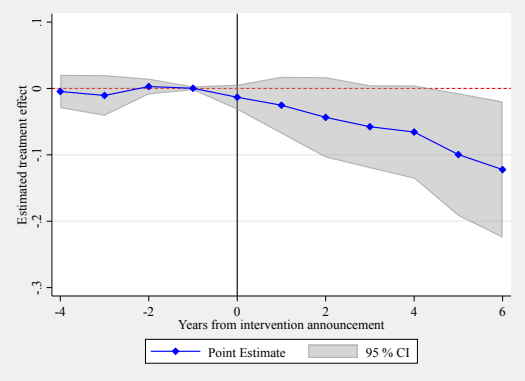
Figure D2: SDiD Direct Effects of Slum Policies

(a) Share of Land Devoted to Residential Housing

In-situ Upgrading

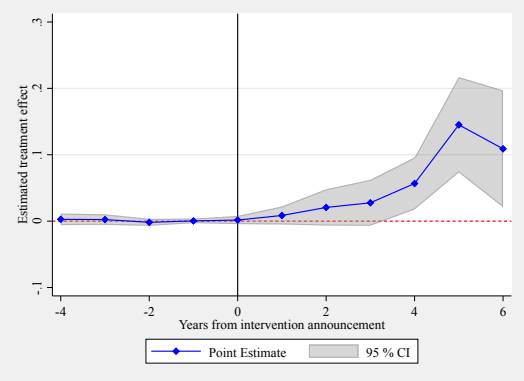


Population Relocation

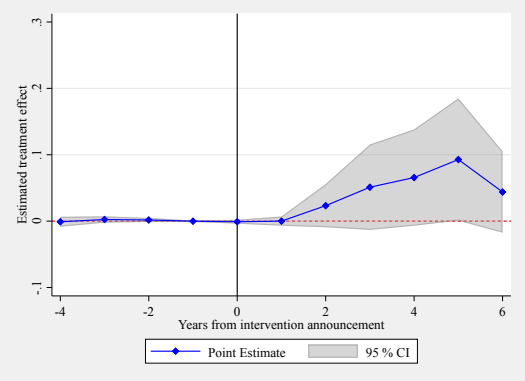


(b) Share of Streets Paved

In-situ Upgrading

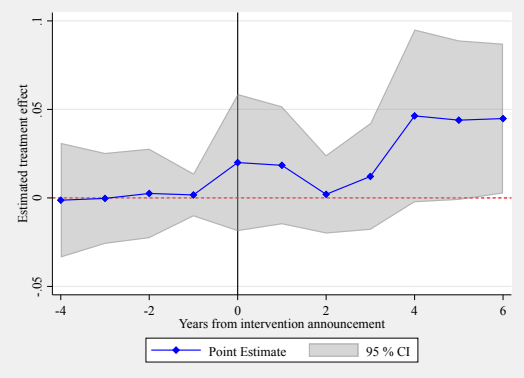


Population Relocation



(c) Pr(Building Permits > 0)

In-situ Upgrading



Population Relocation

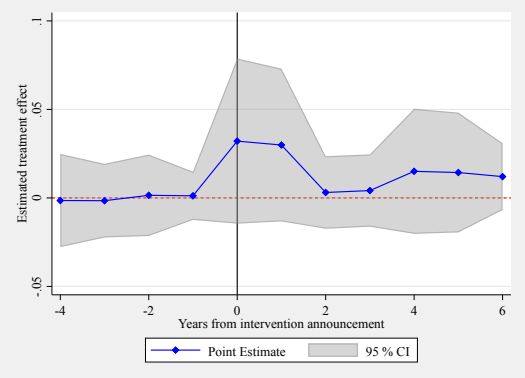
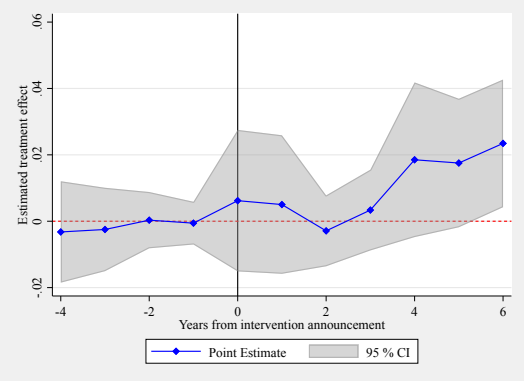


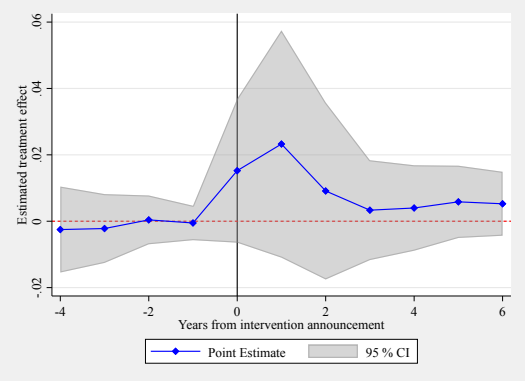
Figure D2: SDiD Direct Effects of Slum Policies (*continued*)

(d) Log # of Building Permits

In-situ Upgrading

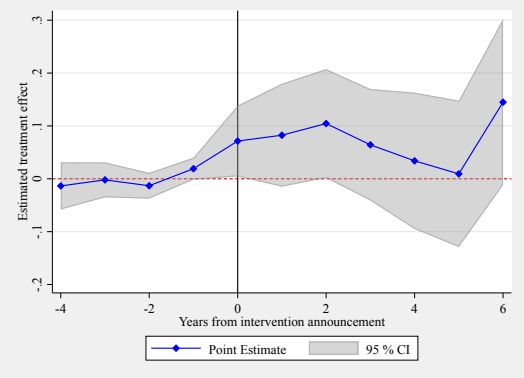


Population Relocation

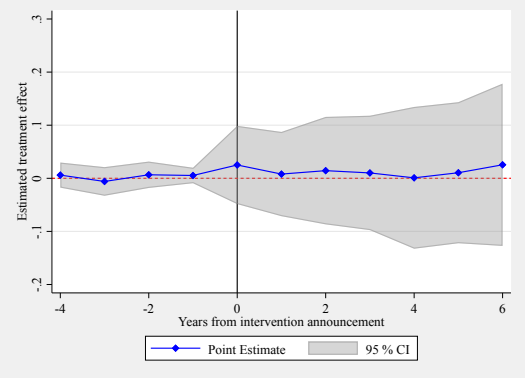


(e) Log # of Large Buildings (size $> 64\text{m}^2$)

In-situ Upgrading

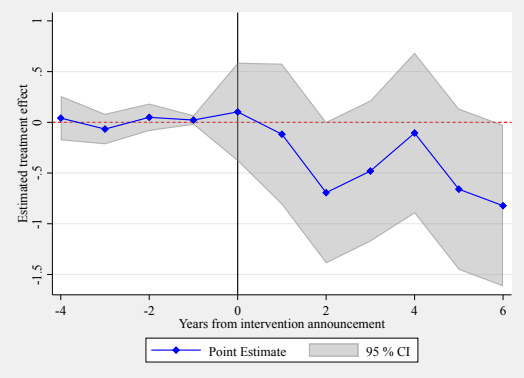


Population Relocation

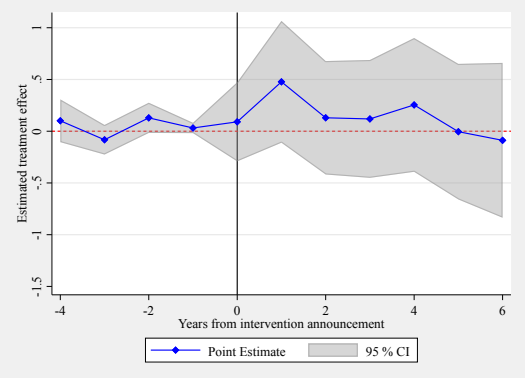


(f) SD Building Main Angle Across Neighbors

In-situ Upgrading



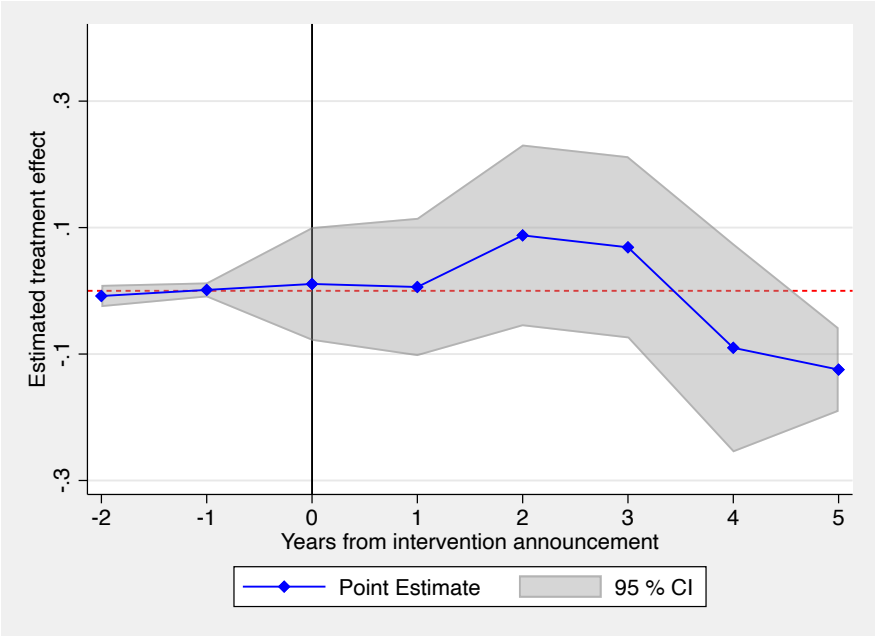
Population Relocation



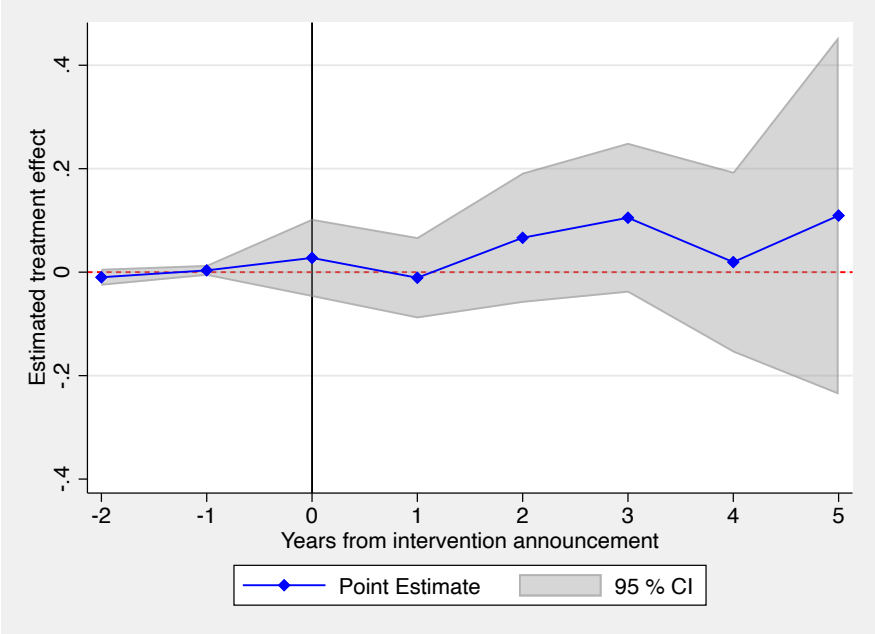
Notes: SDiD - Synthetic Difference-in-Difference methodology developed by Arkhangelsky et al. (2021). Confidence intervals are obtained using bootstrapping. We focus on slums treated between 2011 and 2017.

Figure D3: Spillover Effects of Slum Renewal Policies on Crime Index (sd) in Adjacent Formal Areas

(a) Spillover Effects of *In-situ* Upgrading on Crime



(b) Spillover Effects of Population Relocation on Crime

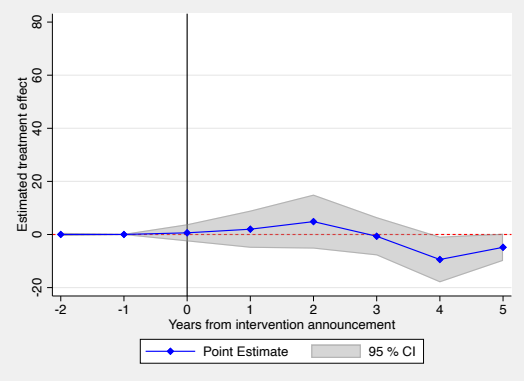


Notes: These are event-study graphs showing the effects of *in-situ* upgrading (Panel a) and population relocation (Panel b) on the crime index in neighborhoods within 200 m of the slum using the SDiD estimator (Arkhangelsky et al., 2021). The crime index follows Anderson (2008) and combines property and violent crime per square kilometer. For each cohort of treated slums, we find a synthetic control group using never-treated slums to match the pre-treatment trends of treated slums. Once the synthetic control is chosen the standard DiD interpretation applies. Inference is constructed using bootstrapping with replacement. Year zero is the moment the intervention is determined. Only treated slums between 2015 and 2018 are considered.

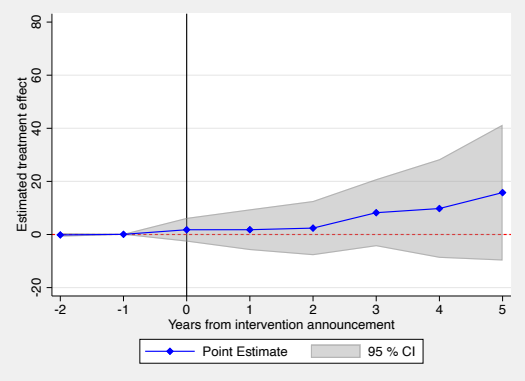
Figure D4: SDiD Spillovers of Slum Policies on Crime -Adjacent Neighborhoods

(a) Property Crime per Km^2 within 200m of a Slum

In-situ Upgrading

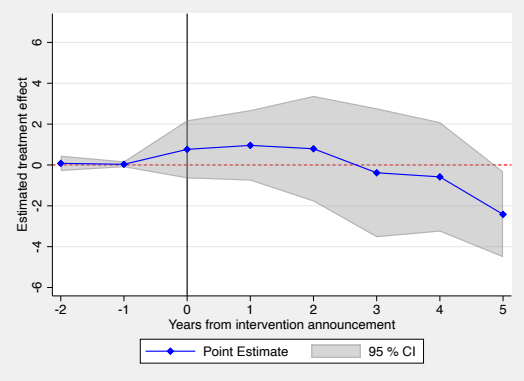


Population Relocation

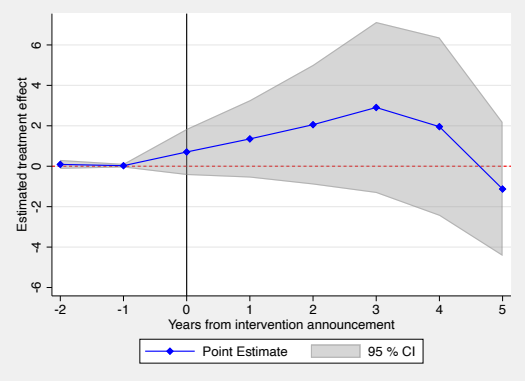


(b) Assaults per Km^2 within 200m of a Slum

In-situ Upgrading

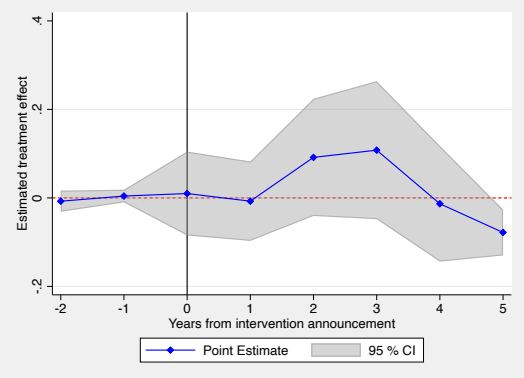


Population Relocation

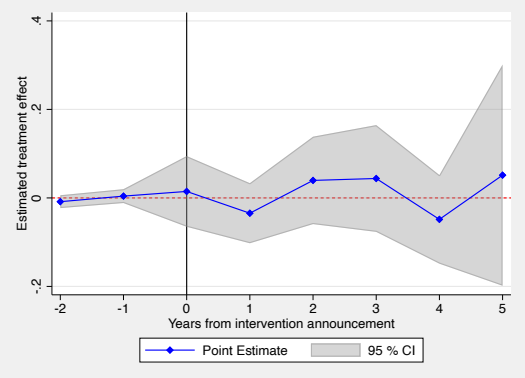


(c) Homicides per Km^2 within 200m of a Slum

In-situ Upgrading



Population Relocation

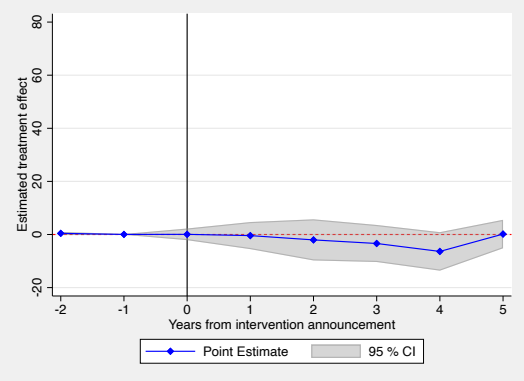


Notes: SDiD - Synthetic Difference-in-Difference (Arkhangelsky et al., 2021). Confidence intervals are obtained using bootstrapping. Crime data come from the Subsecretaria de Prevencion del Delito, and data covers the sample of slums treated between 2015-2018.

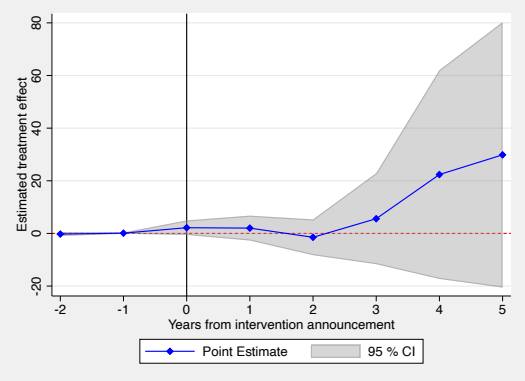
Figure D5: SDiD Spillovers of Slum Policies on Crime -Neighborhoods >200m

(a) Property Crime per Km^2 within 200-500m of a Slum

In-situ Upgrading

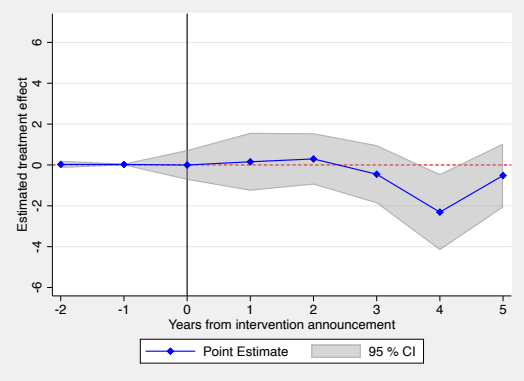


Population Relocation

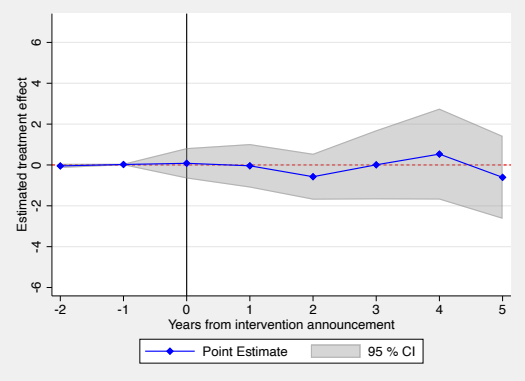


(b) Assaults per Km^2 within 200-500m of a Slum

In-situ Upgrading

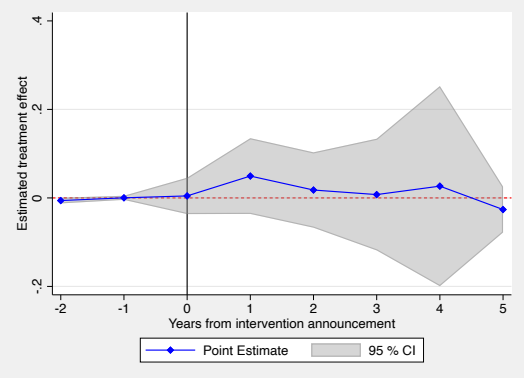


Population Relocation

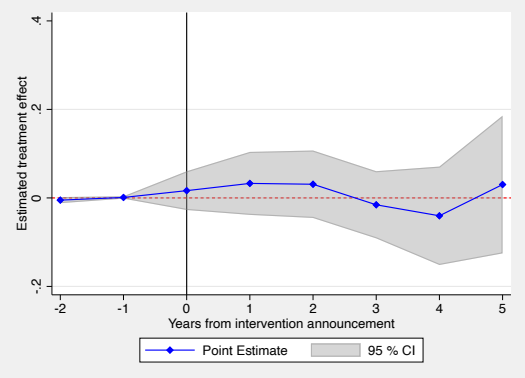


(c) Homicides per Km^2 within 200-500m of a Slum

In-situ Upgrading



Population Relocation

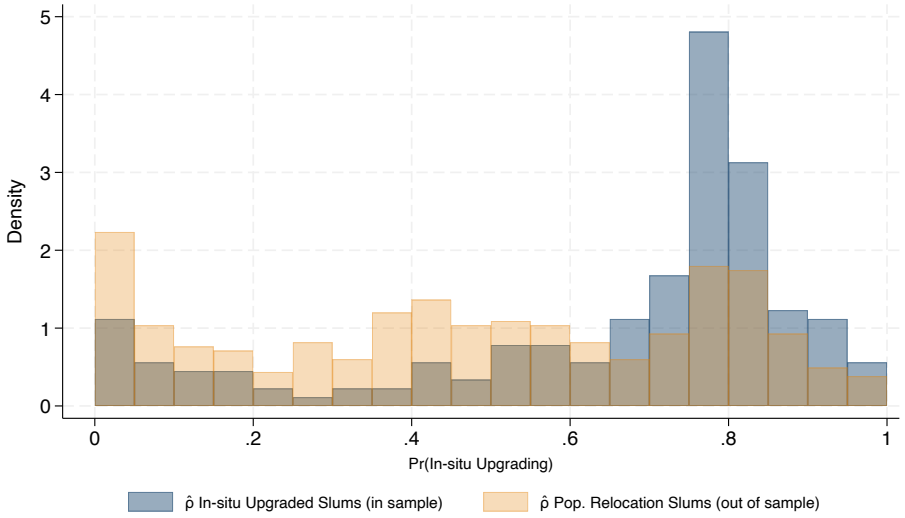


Notes: SDiD - Synthetic Difference-in-Difference (Arkhangelsky et al., 2021). Confidence intervals are obtained using bootstrapping. Crime data come from the Subsecretaria de Prevencion del Delito, and data covers the sample of slums treated between 2015-2018.

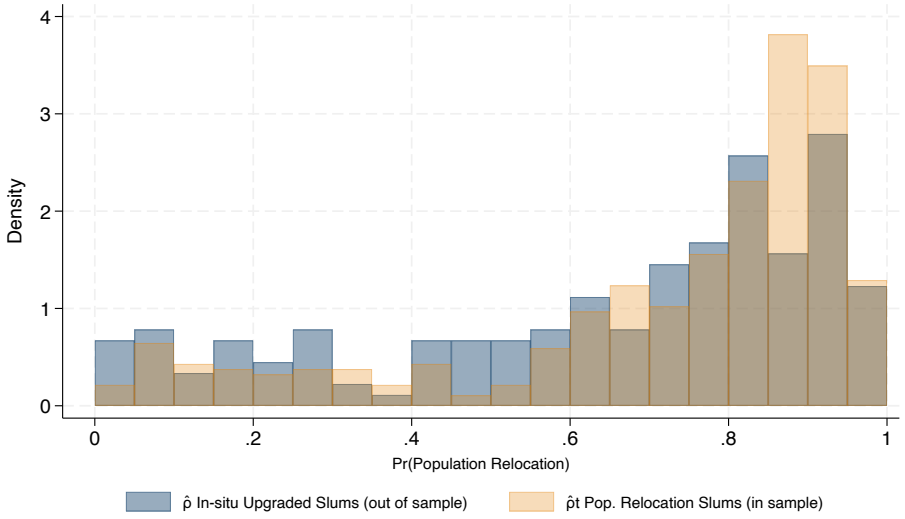
E Counterfactuals

Figure E1: Common Support Treated Slums - Probability Distribution of Each Treatment

(a) Predicted Probabilities from *In-situ* Upgrading Logit Model



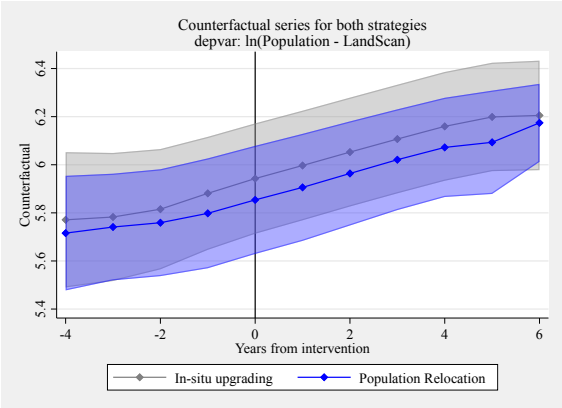
(b) Predicted Probabilities from Population Relocation Logit Model



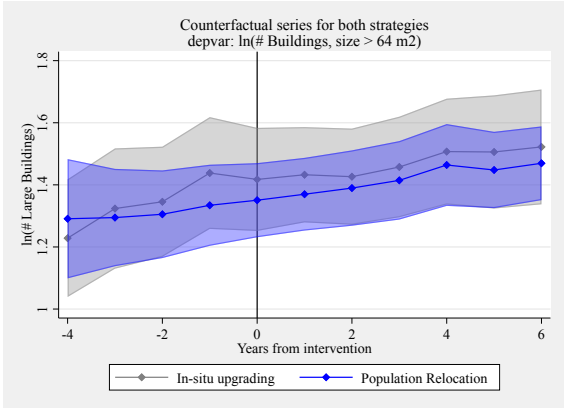
Notes: Each panel shows the distribution of predicted probabilities from logit models that include the full set of covariates listed in Table 8 and region fixed effects. Panel (a) shows predictions from a logit model where the dependent variable equals 1 if the slum was assigned to *in-situ* upgrading and 0 otherwise. The model is estimated using only *in-situ* upgraded slums and never-treated slums. The blue bars display in-sample predicted probabilities for *in-situ* slums (we call them "in-sample" since *in-situ* slums are part of the estimation sample). The tan bars display out-of-sample predictions for slums that actually received population relocation and are therefore not used in the estimation sample; we apply the coefficients from the *in-situ* model to these relocation slums to predict their counterfactual probability of receiving *in-situ* upgrading. Panel (b) mirrors this exercise, i.e., predicted probabilities come from a logit model where the dependent variable equals 1 if the slum was assigned to population relocation and 0 otherwise, estimated using only relocation slums and never-treated slums. Then, blue bars display out-of-sample predicted probabilities for *in-situ* slums, while tan bars display in-sample predictions for relocation slums.

Figure E2: Counterfactual Synthetic Controls for the Two Policies

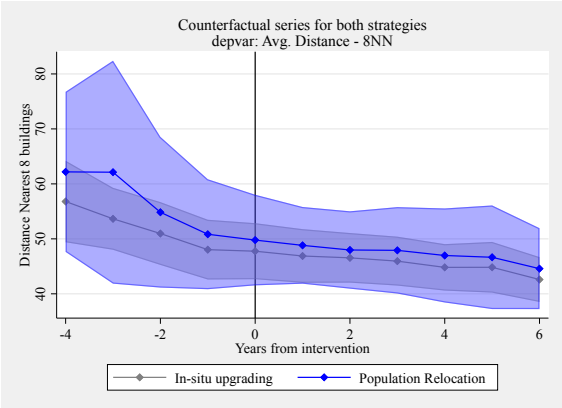
(a) ln(Population - LandScan)



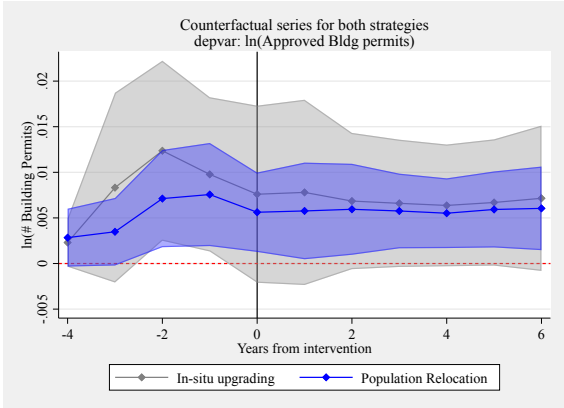
(b) ln(Large Buildings)



(c) Avg. Distance - 8NN



(d) ln(Approved Bldg permits)



Notes: Counterfactual Synthetic Controls correspond to the SDiD estimated series for the control group, i.e., never treated slums. In blue, we show the estimated counterfactual of population relocation slums. In grey, we show the estimated counterfactual of *in-situ* upgrading slums. Each estimation is done separately for each slum policy group, meaning that the objective function for the optimal weights does not take into account the pre-treatment trends of the other treatment arm.

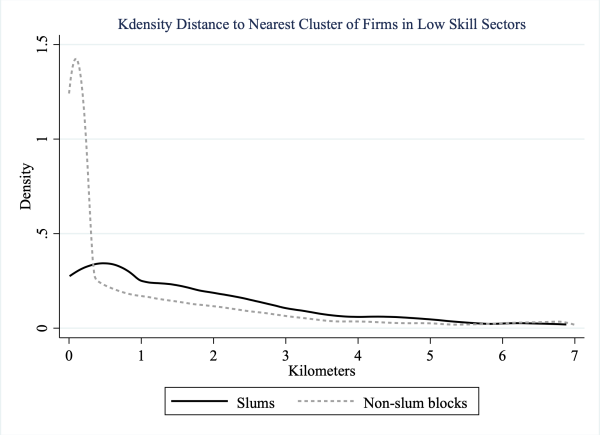
F Location Characteristics: Amenities & Terrain

We have shown that slums develop on the periphery of cities and are eventually absorbed as the city expands. However, slums are not randomly located along the city borders. Instead, they tend to be close to local labor markets, particularly those specializing in low-skill jobs. Table 2 compares a random sample of non-slum census blocks with both old and new slums. Slums are situated farther from police stations, schools, and bus stops than a random set of non-slum blocks. Figure F1 illustrates the distribution of distances to various amenities for slum and non-slum areas. These kernel density plots provide a two-dimensional perspective on the accessibility of local labor markets, public transit, schools, supermarkets, and financial institutions. Most non-slum areas are within 500 meters of low-skill firm clusters and within 1 kilometer of high-skill firm clusters. While slums have limited access to certain amenities compared to non-slum areas, there appears to be a preference for attributes such as public transit and schools. Easy access to public transportation is fundamental for low-income households to commute and access other parts of the city.

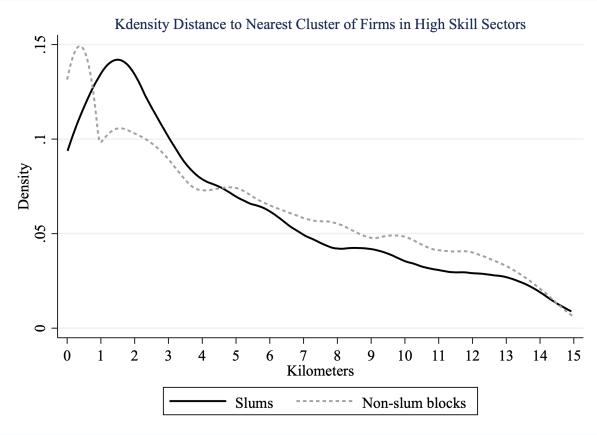
We also assess whether slums are located in riskier geographical areas. While there is a general perception that slums are situated on steep hills—true for some—we compare the geographic conditions between non-slum and slum areas. Figure F2 illustrates the distributions of terrain characteristics for these locations by city size, focusing on three geographical features: the Terrain Roughness Index (TRI), elevation, and slope. In small cities (those with fewer than 300,000 inhabitants), non-slum blocks and slums have similar distributions across these characteristics. However, significant differences emerge between formal and informal areas in larger cities. Slums in big cities are found in areas with lower elevation but higher TRI and steeper slopes. In terms of TRI, values below 80 meters represent level terrain, while values between 240 and 497 meters indicate moderately rugged surfaces (TRI classification by ESRI is available here.) Formal areas are mostly located on level terrain, whereas most slums are on rugged surfaces. Regarding elevation, there is a concentration of non-slum blocks around 500 to 700 meters, corresponding to the city of Santiago. Slums, in contrast, are mostly located below these elevations, where they might be at higher risk of natural hazards such as flooding. In terms of slope, our findings follow those of Müller et al. (2020) in which just few slums are located on hills steeper than 10 degrees.

Figure F1: Slums Proximity to Different City Amenities

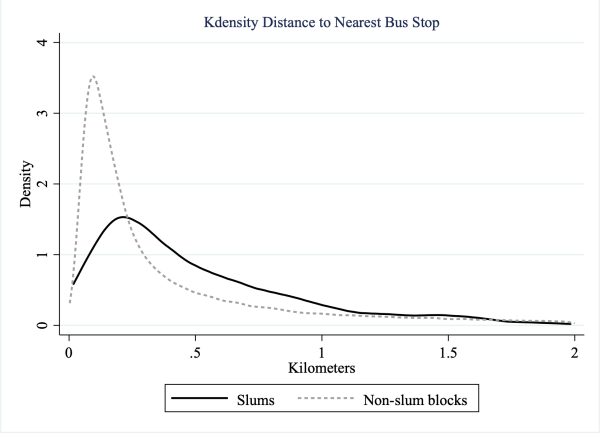
(a) Low-Skill Firms' Cluster



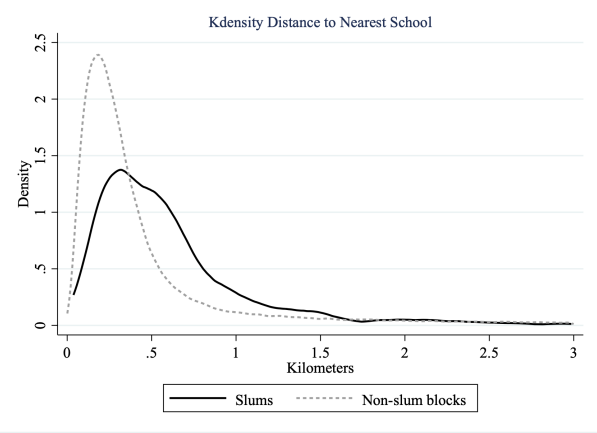
(b) High-Skill Firms' Cluster



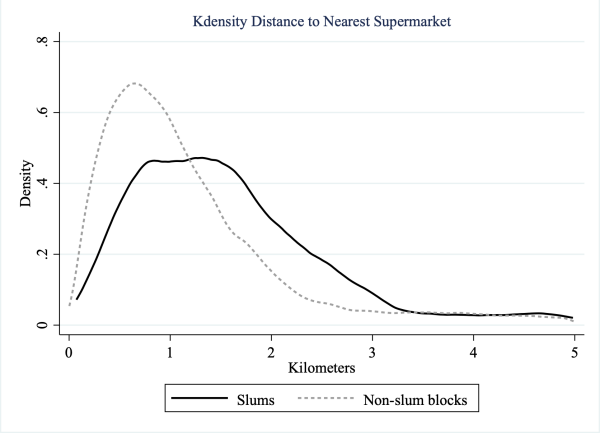
(c) Bus Stop



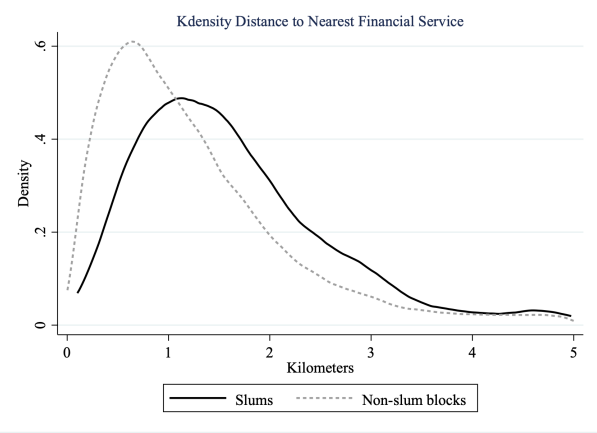
(d) School



(e) Supermarket



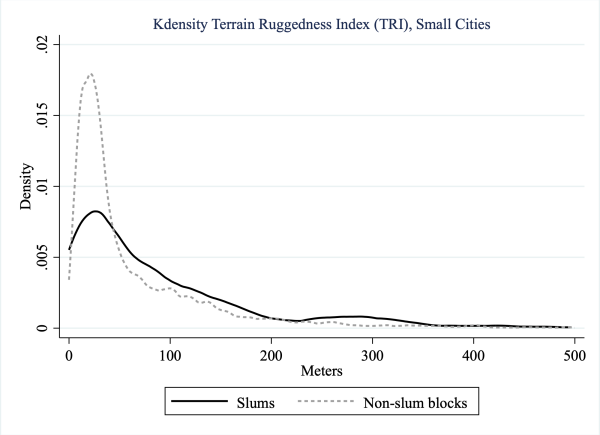
(f) Finance Institution



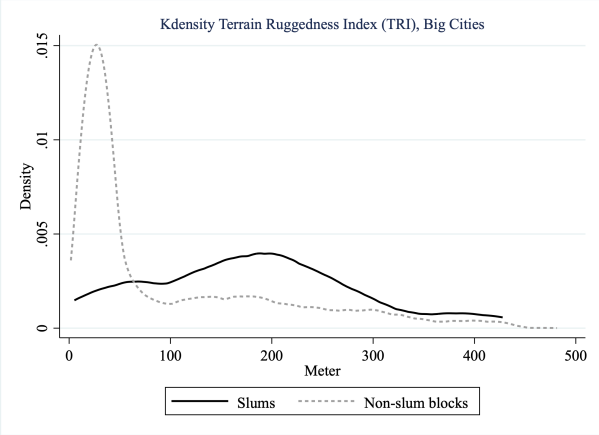
Notes: Each figure compares the distribution of distances to the nearest amenity for slum areas and non-slum areas. Non-slum areas refer to municipality census blocks.

Figure F2: Slums Terrain Characteristics by City Size

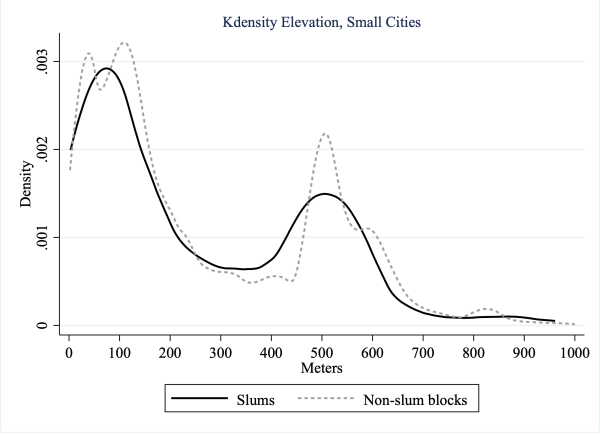
(a) TRI, small cities



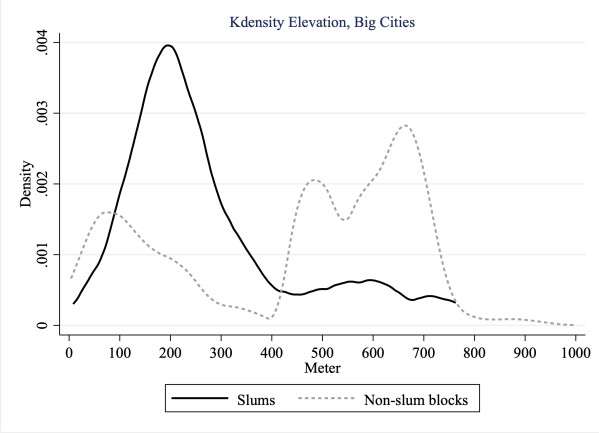
(b) TRI, big cities



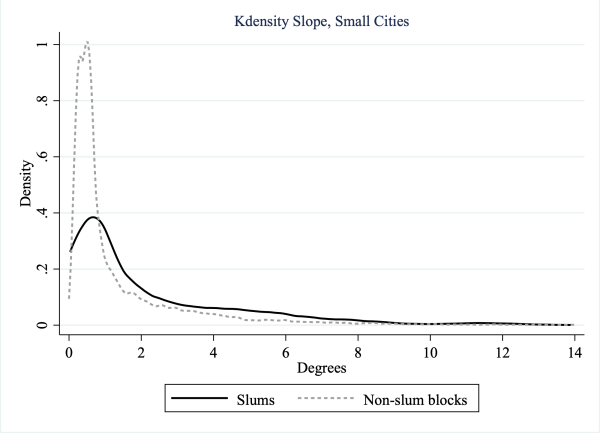
(c) Elevation, small cities



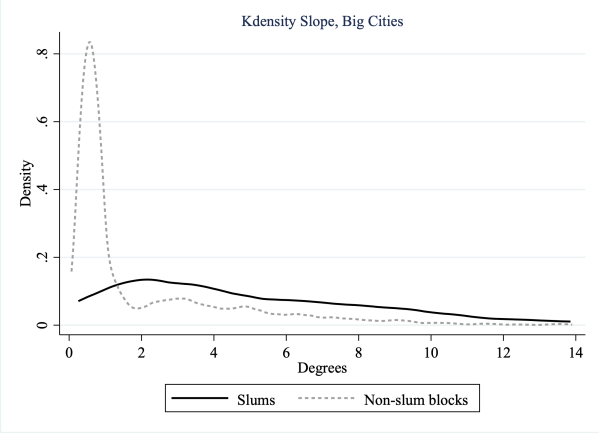
(d) Elevation, big cities



(e) Slope, small cities



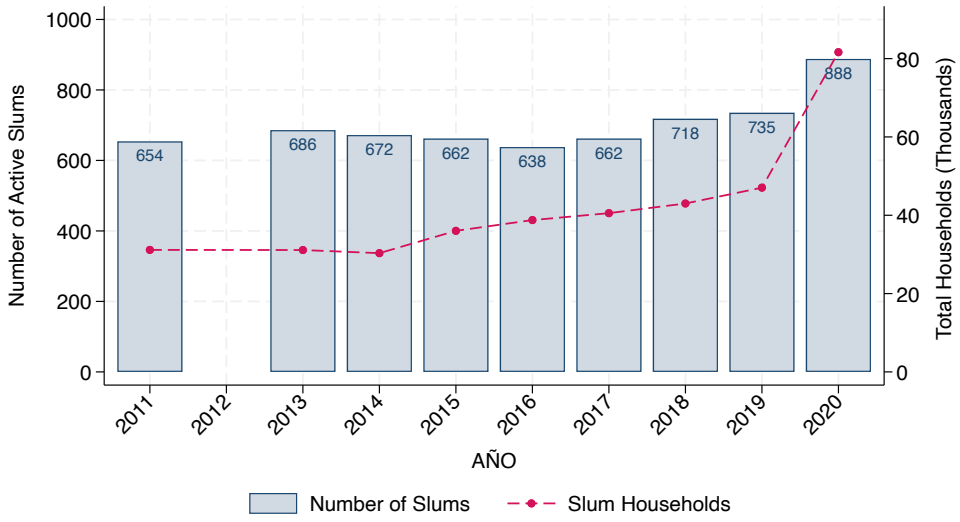
(f) Slope, big cities



Notes: Each figure compares the distribution of three terrain characteristics: Terrain Roughness Index (TRI), elevation, and slope, for slum areas and non-slum areas. Non-slum areas refer to municipality census blocks.

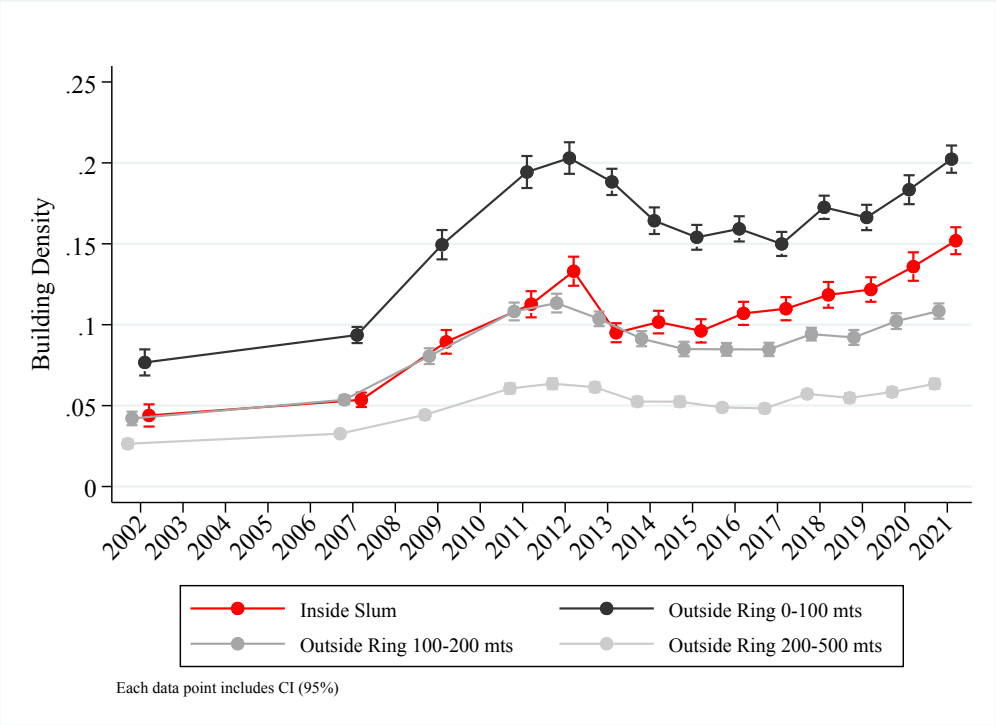
G Additional Appendix Figures and Tables

Figure G1: Number of Active Slums and Slum Households in Administrative Records



Notes: Active slums are informal settlements included in administrative records. An area is considered a slum if it (i) has eight or more households spatially close, (ii) lacks property rights, and (iii) houses are of substandard quality and lack access to at least one of the basic services (water, sanitation, and/or electricity). Source: MINVU (2011,2019) and TECHO (2013-2018, 2020)

Figure G2: Evolution of Building Density Inside Slums and Adjacent Neighborhoods



Notes: Building density is calculated using the building footprints obtained from Google Earth images. See Appendix C for details. Building density inside slums is the proportion of slum areas occupied by building footprints. The outer rings are defined by distance from the slum boundary: the first ring (0-100 meters) measures density within 100 meters, the second ring (100-200 meters) covers the next 100 meters, and the third ring (200-500 meters) extends to 500 meters.

Table G1: Sociodemographic Changes - Population Censuses 2002 and 2017

	Census Year		Δ_t
	2002	2017	(2017-2002)
	(1)	(2)	(3)
Population			
Slums	277	308	31
Adjacent Non-Slum Blocks	6,870	7,206	336
Δ_i (Slums - Non-Slums)	-6,593	-6,898	-305
Ratio Adult Males/Females (25-64 yo)			
Slums	1.30	0.97	-0.33
Adjacent Non-Slum Blocks	0.96	0.95	-0.01
Δ_i (Slums - Non-Slums)	0.34	0.02	-0.32
% Pop. with Secondary Education			
Slums	39.58%	51.03%	11.4%
Adjacent Non-Slum Blocks	43.21%	50.65%	7.4%
Δ_i (Slums - Non-Slums)	-3.62%	0.38%	4.0%
% Pop. with some Tertiary Education			
Slums	10.20%	18.45%	8.3%
Adjacent Non-Slum Blocks	14.45%	20.24%	5.8%
Δ_i (Slums - Non-Slums)	-4.25%	-1.79%	2.5%
% Employed			
Slums	80.63%	90.19%	9.6%
Adjacent Non-Slum Blocks	82.41%	90.31%	7.9%
Δ_i (Slums - Non-Slums)	-1.79%	-0.11%	1.7%

Notes: Unit of analysis are census blocks from 2002 and 2017 censuses. A block is tagged as "slum" if its centroid falls inside the polygon ($N = 555$). Likewise, a block is tagged as adjacent "non-slum" if its centroid lies within 200 to 500 meters of the slum border ($N = 713$). All outcome variables come from the population census. For details on variable definitions, see Online Appendix Section B. Columns (1) and (2) represent the average in the 2002 and 2017 censuses, respectively, while column (3) represents the average change over time (Δ_t). This is computed separately for "Slums" and "Adjacent Non-Slum Blocks". The last row of each grouping presents the across-group difference for the given census year Δ_i.

Table G2: Correlates of Municipal Slum Population Growth

Dependent Variable: Δ Log Municipal Slum Pop.	WorldPop (2011-2017) (1)	Pop. Census (2002-2017) (2)
$\underset{2009-2017}{\Delta}$ Log Quality Adj. Rent	0.24*	0.29**
	(0.13)	(0.14)
$\underset{2010-2017}{\Delta}$ Log Wage	-0.21**	-0.22*
	(0.09)	(0.11)
$\underset{2010-2017}{\Delta}$ Log Wage - High School or more	0.20***	0.17***
	(0.02)	(0.03)
$\underset{2010-2017}{\Delta}$ Log Employment	-0.12	0.25
	(0.14)	(0.16)
$\underset{2010-2017}{\Delta}$ Log # Employment - High School or more	0.14	-0.09
	(0.11)	(0.12)
$\underset{2011-2017}{\Delta}$ Log Municipal Social Expenditures	-0.21	-0.06
	(0.20)	(0.19)
Obs.	261	261
R2	0.17	0.11
Mean Dependent variable	0.27	0.11

Notes: OLS estimates of regressing the log change in the municipality slum population against the log change in rents, employment, salaries, and public expenditure, with $\underset{a-b}{\Delta} X$ representing the log change of variable X between years a and b. For details on variable definitions, see Online Appendix Section B. Each column uses a different data source for calculating slum population growth. Column (1) uses WorldPop's 100-meter resolution estimates for the years 2011 and 2017. Column (2) relies on census data from 2002 and 2017. Both WorldPop grids and Census blocks are classified as slums based on their centroid locations, i.e., if a centroid is within slum boundaries, the population is counted as living in slums. WorldPop 100m cells with fewer than four people are set to zero, likely due to the smoothing algorithm. All regressions control for the log of baseline municipality population (2011 in Column 1 and 2002 in Column 2) and the municipality population growth during the same period. Robust standard errors in parentheses. * p<0.1, ** p<0.05, *** p<0.01.