Long-Term Effects of Slum Upgrading on Climate Resilience and Urban Growth

Juliana Pinillos*

Department of Economics, University of Southern California

Click here for the latest version

October 24, 2026

Abstract

This paper studies the long-term effects of slum upgrading on exposure to natural disasters and urban growth. I study PRIMED, a large-scale program implemented in Medellin, Colombia, between 1993 and 2000 that combined neighborhood improvements with risk-mitigation infrastructure. To identify its effects, I exploit the program's incomplete rollout by comparing slums upgraded during phase I with slums from phase II that were scheduled but never treated. I find that upgrading reduces landslides during extreme rainfall events relative to untreated slums. The program also increased neighborhood population, primarily through an increase in low income residents. Housing and population density also rose, with higher growth on lower-risk land. I embed these empirical estimates in a quantitative spatial model to measure citywide welfare effects. The results indicate that slum upgrading generates welfare gains for the city, with larger gains among low skilled residents. Two mechanisms explain this: improved local fundamental amenities and reduced disaster exposure. However, counterfactual scenarios suggest that more frequent disasters due to climate change lead to substantial welfare losses even when the protection provided by upgrading is maintained. These findings highlight how climate change can limit the long term effectiveness of infrastructure investments within cities.

Keywords: climate change; resilience; slums; urban development.
JEL codes: O18, Q54, R13, R14.

*I am deeply grateful to Professors Paulina Oliva, Andrii Parkhomenko and Thomas Chaney for their guidance, constant support, and encouragement throughout this project. I also thank Professors Jorge De la Roca and Simon Quach for their insightful research advice. I am especially grateful to Professor Matthew Kahn for his thoughtful comments and for his support, together with Professor Oliva, through the USC Presidential Sustainability Initiative. I thank the participants of the USC Applied Microeconomics Reading Group for their valuable feedback and stimulating discussions. I thank participants at the CEGA Research Retreat, UEA, the International conference of AREUA and seminar participants at the universities EAFIT and ICESI. I also appreciate the comments and suggestions of my Third-Year Paper Committee. This paper benefited from research funds from the USC Wrigley Institute through the Graduate Fellowship for Environmental Sustainability Research, the USC Lusk Center for Real Estate through its Research Awards Program, the USC Department of Economics through the 2026 Summer Research Funds, and the USC Dissertation Completion Fellowship. I am very thankful for all this support. I am also grateful to the Colombian National Administrative Department of Statistics (DANE), especially the Dirección de Censos y Demografía, for their assistance with data requests. All remaining errors are my own.
† Department of Economics. Email: pinillos@usc.edu

1 Introduction

Slums are home to nearly one quarter of the world's urban population, roughly one billion people (UN-Habitat, 2021). In many developing cities, slums combine inadequate housing, weak local infrastructure, and high exposure to environmental risk. This is increasingly important as climate change is expected to intensify extreme rainfall, and recent evidence shows that slum residents face disproportionate exposure to floods and landslides (IPCC, 2023; Li et al., 2025; Ozturk et al., 2022). In response to slum growth, governments often face difficult trade-offs. Policy typically centers on two alternatives: *in-situ* upgrading or relocating households to safer land[1]. Relocation can reduce exposure to risk, but it is often costly, politically difficult, and disruptive to jobs and social networks. Upgrading, in turn, can improve services and provide protection against natural hazards, making existing informal settlements more attractive places to live. These improvements may also encourage additional settlement in areas where some disaster risk remains. This paper asks whether slum upgrading reduces disaster vulnerability and leads to positive welfare impacts among vulnerable populations.

To answer this question, it is important to follow both disaster incidence and neighborhood development over a long period. This allows for a comparison between upgraded and non-upgraded neighborhoods in terms of their response to heavy rain events before and after the intervention, while capturing population and housing responses that can take decades to be fully realized after the upgrading has ended. In addition, it is important to identify a credible control group, as upgraded slums are systematically poorer and more vulnerable than the rest of the city, and comparing them to any non-upgraded areas could confound the effects of upgrading with pre-existing differences. Finally, place-based effects alone do not establish the citywide benefits of the policy because people may move in response to the intervention.

I study these questions using a large upgrading program implemented in Medellin, Colombia, between 1993 and 2000: the *Programa de Mejoramiento Integral de Barrios Subnormales de Medellin* (PRIMED). This program provided traditional local infrastructure investments while also building protective infrastructure such as slope stabilization, retaining walls, and drainage systems. PRIMED was one of the first integrated slum upgrading programs in Latin America and received international recognition through its selection in the UN-Habitat/Dubai Best Practices Database in 1998 and 2000 (Habitat, 1998, 2000). The program was designed to be implemented in two phases. During the first phase, 11 neighborhoods were upgraded across three zones of the city with a beneficiary population of 57,900 inhabitants, about 31% of the total population living in slums at the time. The second phase was projected to upgrade an additional 26 neighborhoods, with a total population of about 125,000. Despite being planned, the

[1]Some governments also opt for clearance, but this approach is generally feasible only for smaller slums and raises additional welfare and political concerns.

second phase was not implemented due to administrative closure of the implementing agency in 2001 (Betancur, 2007).

To evaluate the long-term impacts of this policy, I exploit the incomplete rollout of PRIMED by comparing upgraded Phase I with Phase II non-treated slums. As both groups were chosen for upgrading, this makes Phase II a comparison group drawn from the same program. A concern, however, would be that Phase I slums were targeted based on their exposure to natural disasters. I assess their comparability using baseline outcomes and pre-intervention disaster trends. I further examine the effects of upgrading by restricting the control group to neighborhoods at the same level of consolidation as Phase I, using the program's neighborhood historical classification.

To estimate the effects on exposure to natural disasters, I combine disaster records covering 1981-2023 with daily gridded precipitation. I first use a dynamic difference-in-differences model to examine whether upgraded areas had a higher incidence of these events prior to the investment. My results suggest that both groups faced similar exposure before 1993, and that after upgrading, there is a decline in natural disasters primarily driven by a reduction in landslides. I next examine if the risk infrastructure is indeed protective when extreme rainfall materializes. To do this, I estimate an event-study around a heavy-rain episode. I find that slum upgrading substantially reduces the probability of experiencing any natural disaster in response to a heavy rain episode. Event-study estimates indicate that, relative to the pre-intervention period ($\text{year} < 1993$), there is a reduction of roughly 31% in the number of disasters per year during extreme rainfall events, with an estimated reduction of 35% in landslide events.

Then I analyze whether the reduction in these events is likely to result in a reduction of economic damages, as not all events have the same consequences. Although the exact counting of damage is limited over time, I can identify which disasters partially damaged or destroyed at least one home. My findings suggest that there is a reduction in both disasters with and without housing damage during extreme rainy periods, with slightly larger reductions in events with registered damage. I interpret these results as evidence of upgrading as being able to provide protection to homes and people when it is needed the most. These results are economically meaningful and indicate that slum upgrading can enhance resilience to extreme rainfall by reducing both the likelihood and the severity of climate-related shocks.

Having established that upgrading reduced the incidence of natural disasters, I analyze what happens to upgraded slums in terms of urban growth as safer conditions may affect where households choose to live and build. To do this, I use block-level census data and grid-cell-level remote-sensing measures and analyze the trajectory of population, urban density and impervious surface over the last 25 years. For these outcomes, I estimate a difference-in-differences model. I find that blocks in upgraded slums experience an increase of about 49 additional residents relative to the control group. Among adults aged 25+, the increase is concentrated in low-skill residents.

These effects represent about 28% and 21% of the control mean as of baseline (1993). I find no statistically detectable change in the high-skill share. I interpret this result as evidence of no changes in the population composition. This implies that the program did not lead to gentrification or re-sorting effects in the long run. In terms of urban development, my results suggest that housing density increases by 6,384 dwellings per km^2 and population density by 21,075 residents per km^2, corresponding to 41% and 31% of the respective 1993 control means. I do not find effects on dwelling construction characteristics.

However, local population and housing gains do not establish the citywide welfare effects of upgrading because households can relocate and rents and wages can adjust. This is important for three reasons: first, if changes in population and housing reflect simply relocation from somewhere else in the city, then the local gains may be offset by losses in other areas. Second, increases in population in treated areas can lead to higher rents, which can reduce the benefits of the program for the residents. The inflow of population can also create higher local labor supply, which can also reduce wages. Thus, accounting for general equilibrium effects using a quantitative spatial equilibrium model (QSM) is critical to understand the overall welfare implications of PRIMED. The model also helps us understand the distributional consequences of the program by comparing welfare between low and high income groups.

I embed the reduced form estimates into a QSM close in spirit to Ahlfeldt, Redding, Sturm, and Wolf (2015) and Tsivanidis (2026). The model departs from these frameworks by incorporating the risk of natural disasters as a key determinant of residents' utility as higher exposure generates disamenities. I model PRIMED's neighborhood infrastructure improvements as an increase in local fundamental amenities and its risk mitigation investments as a reduction in disaster exposure. I use my event-study estimates around heavy rainfall episodes to specify the reduction in disaster exposure attributable to PRIMED. Conditional on this reduction, I calibrate the fundamental amenity improvement so that the combined population response to both interventions matches the population difference-in-differences identified moment (Nakamura & Steinsson, 2018). The calibrated model can then determine the welfare effects and the contribution of each channel.

My results suggest important welfare gains from the slum upgrading program. Under the closed city results, Medellin experienced a citywide gain of 1.96%, with the majority of the benefits coming from low-income residents (2.1%), whereas high-income residents experienced a smaller gain (1.2%). The results under the open city suggest similar effects, with an overall gain of 0.66%. The lower welfare gain is explained by the increase in rents and the decrease in low skill income, which is due to the increase in population in the city. When decomposing the effects by channel, the reduction in exposure to natural disasters accounts for roughly half of the total welfare effect. The improvement in local fundamental amenities, explains about 28.6%. I also simulate the

welfare gains of completing the program, that is, if the city had implemented Phases I and II. My results indicate larger welfare gains at the aggregate level, which under the open city are 2.1%.

Then I use the model to ask what the welfare effects of this policy would be under different climate change scenarios, as defined by the IPCC. To do this, I use scientific precipitation projections under each scenario together with empirical Bayes methods to project disasters across the city by 2081. Using these projections, I compare future welfare with and without Phase I upgrading relative to today's observed equilibrium. My results suggest that without this infrastructure, the city welfare is 0.96% lower in the low precipitation scenario (SSP1-2.6) and 2.8% lower in the high precipitation one (SSP5-8.5). These city welfare losses are reduced by about 68% in the first scenario, but with high precipitation the program reduces these welfare losses by 23%. This comparison assumes that the infrastructure retains its protective capacity more than eight decades after construction. When I allow protection to deteriorate through depreciation or insufficient maintenance, welfare losses increase and the advantage of upgrading narrows. These results suggest that even if upgraded areas are coping well with current disaster risk, future additional exposure can create unintended consequences as the observed increase in urban density may leave more residents exposed to future higher climate hazards.

This paper contributes to three strands of literature. First, it adds to research on slum policies (Marx, Stoker, & Suri, 2013). Existing studies document improvements in housing, health, and neighborhood conditions (Galiani et al., 2017; Gertler, Gonzalez-Navarro, Undurraga, & Urrego, 2024), alongside lower land values and less vertical development in some settings (Harari & Wong, 2024). I extend this literature by examining whether upgrading protects settlements against extreme rainfall. Second, the paper contributes to work on infrastructure investment under environmental change (Balboni, 2025; Hsiao, 2025; Kahn & Ouazad, 2025). Whereas Balboni (2025) studies road allocation under sea-level rise and Hsiao (2025) examines how anticipated public protection delays coastal retreat, I evaluate the welfare implications of neighborhood investments that jointly improve amenities and reduce disaster risk, and examine whether these benefits persist as climate risk increases. Finally, the paper relates to the literature on the local economic consequences of natural disasters (Aron-Dine, 2025; Boustan, Kahn, Rhode, & Yanguas, 2020; Deryugina, 2017) by studying long-run residential adjustment following public investment in disaster-prone neighborhoods.

This paper is organized as follows: section 2 explains housing informality in Medellín and the institutional details about PRIMED. Section 3 describes the data sources and presents baseline balance checks. Sections 4 and 5 present the empirical analyses of natural disasters and urban outcomes. Section 6 introduces the quantitative spatial model and section 7 presents its quantification and results. Finally, section 8 concludes.

2 Housing Informality and The Slum Upgrading Program

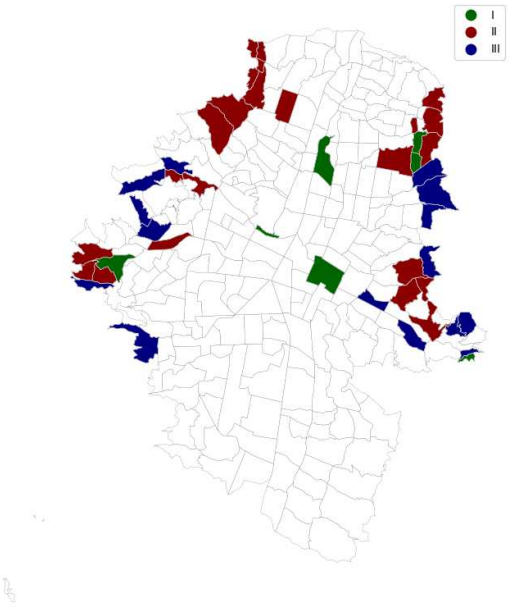
2.1 The Informal Settlements in Medellin

Medellín is the second largest city in Colombia. It has about 2.3 million people, about 5.5% of the Colombian population. The urban area constitutes 101.25 KM2. The city has 249 neighborhoods distributed across 16 districts (*comunas*).

Historically, Medellin was a city marked by strong industrial development. During the decades of 70s and 80s Medellin had an exponential urbanization process, with urban growth rates around 7%. Such rapid sprawl led to the conformation of many informal settlements, especially in the peripheral areas of the city, where land was available. The growth of slums was specially concentrated in the high hills, and sometimes near water streams (locally known as *quebradas*) in the mountains.

In 1992, informal settlements in Medellín contained approximately 41,000 dwellings and 182,000 residents. Due to its informality, slums had poor infrastructure conditions, with unpaved roads, informal paths, and lack of drainage. Slums were classified into three levels according to their degree of urban consolidation. Level I neighborhoods were more consolidated and were excluded from PRIMED. Level II settlements were partially consolidated, while Level III settlements were at an earlier stage of development. The program targeted selected settlements in these latter two categories.

Figure 1: Slum Neighborhoods Classification in Medellin in 1992



Source: Own based on historical archives

2.2 The Slum Upgrading Program (PRIMED)

The *Programa de Mejoramiento Integral de Barrios Subnormales de Medellin* was implemented between 1993-2000 in the city of Medellin, Colombia. The program was executed by the City, with the technical assistance of the German Governmental Bank for Reconstruction and Development (KFW) and the United Nations Development Program (UNDP). The program was the first large holistic approach to informal settlements in the city and one of the first ones in Latin America. It obtained two

PRIMED targeted slums classified as Level II and III. The program was planned to be executed in two phases. Phase I was scheduled to be implemented between 1993-1996, while Phase II was planned to be executed between 1997-2002. In its initial phase, the program started by upgrading Level II slums located in three zones of the city: Central-west, North-Western and the Central-Eastern. The areas upgraded by Phase I represented 30% of housing informality in Medellin. Figure 2 illustrates the location of targeted slums across phases. In practice, phase I was extended to 2000 due to unexpected delays. The next year PRIMED was terminated because the implementing agency (CORVIDE) was dismantled with the start of the new administration in 2001 (Betancur, 2007). Thus, slums scheduled for phase II were not upgraded by this program.

PRIMED had four components: neighborhood upgrading, home improvement, risk mitigation infrastructure and legalization of tenure. Infrastructure investments represented the bulk of the program, while land formalization represented a small share of PRIMED (3.9%)[2]. The neighborhood infrastructure investment provided street improvement in 80% of the area; it also provided 6,000 meters in water pipes and 5,000 mts in drainage, with a respective projected coverage of 95% and 90 % of the households. Similarly, the coverage of pedestrian infrastructure was increased up to 60%; and around 2,800 meters were provided for parks and open spaces. Likewise, it built schools and local health centers. The the mitigation of geological risks, five Hectares (Ha) were recovered and 8.5 Ha. were stabilized, equivalent to around 70% of the areas classified as high risk. Risk mitigation included slope stablizaition, retaining walls, drainage works, and vegetation restoration to control erosion. Figure 3 shows examples of these investments.

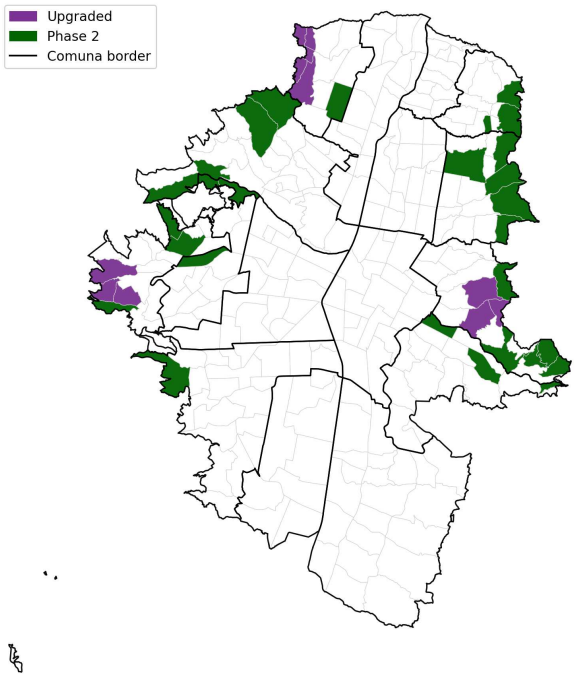
In total, PRIMED invested $63.4 millions as of 2025 ($26.6 Millions in 1991). It was funded through several sources: 33.8% from the city, 35.9% from the National Government, 30% from KFW, and 0.3% from UNDP[3].

[2]The program intended to provide property rights to 5,180 households but only about 2,100 families received legalization. In terms of budget, land tenure represented only 3.9% of the total cost. Tenure was one the most challenging aspects of the program due to the low program investment and complexities related to land ownership, and household ability and willingness to participate (Betancur, 2007).

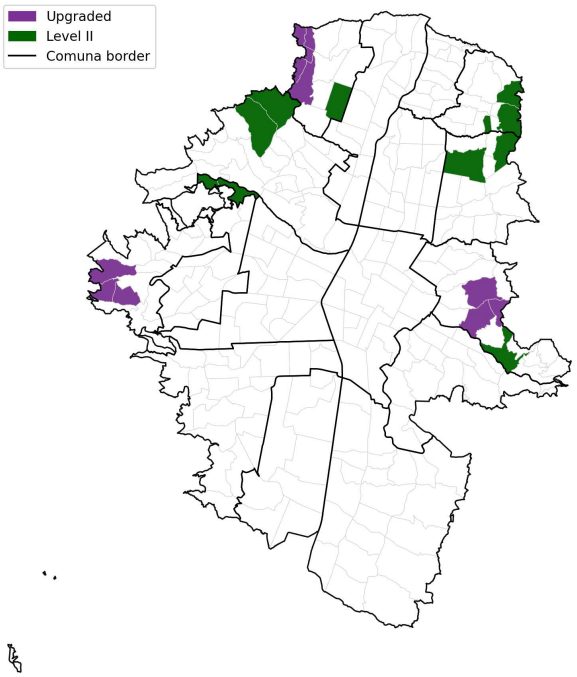
[3]In principle, the community also contributed with resources in order to benefit from certain components of the program such as home subsidy and land tenure. The amount of these contributions was small ($3.7 millions, in 1991). I exclude these from the total cost consolidation due to the lack or records

Figure 2: Targeted Slums by Program Phases

A. All Slums from Phase 2



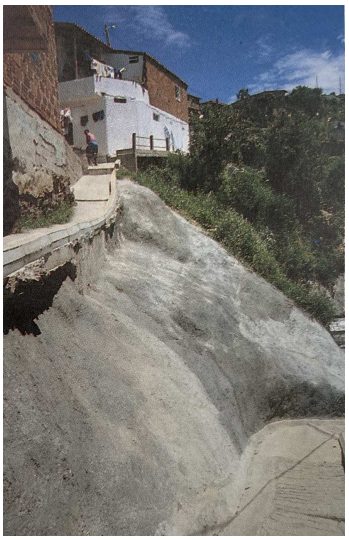
B. Level II Slums from Phase 2



Note: Own elaboration based on historical archives. Panel A displays the current neighborhood representation of Phase I (upgraded) and Phase II. The map contains 11 Phase I and 28 Phase II neighborhood polygons. These are mapped administrative units; the historical planning count and effective estimation samples need not coincide. Panel B restricts the comparison to the mapped Level II neighborhoods.

Figure 3: Examples of PRIMED infrastructure investments

Panel A. Slope protection



Panel B. Stairways and walls



Source: PRIMED archives.

3 Data

3.1 Data Sources

Natural disasters. I compile reports of landslides and floods in Medellín from DesInventar, GeoHazard, PRIMED's historical documentation, and the Unidad Nacional para la Gestión del Riesgo de Desastres (UNGRD). The inventory covers 1981–2023. I include reported occurrences whether or not deaths, injuries, or physical damage were documented, and record damage separately. I reconcile reports across sources by date, neighborhood, and hazard type.

Terrain Features and Precipitation. To compute terrain slope, I use elevation data from the ASTER Global Digital Elevation Model. ASTER data is derived from satellite observations by NASA and METI (NASA & METI, 2019). For the main analysis, I use the average slope at the neighborhood level; for robustness checks, I also compute slope at the block level. Precipitation data comes from the Climate Hazards group Infrared Precipitation with Stations (CHIRPS), Version 2.0. CHIRPS combines "smart" interpolation techniques, high-resolution satellite data (0.05°), and long-term precipitation estimates based on infrared Cold Cloud Duration (CCD) observations (Funk et al., 2015).

Impervious Surface. I measure changes in impervious land cover between 1985 and 2020 using data from the Global Land Cover with Fine Classification (GLC_FCS30D) developed by Zhang et al. (2024). This data is available at a fine resolution of 30 meters and is derived from Landsat imagery. The authors apply machine learning techniques to detect continuous change using a refined classification with 35 land categories. I define an indicator equal to one when a 30 m $\times$ 30 m cell is classified as "impervious surface" and zero otherwise. This category captures artificial land covers such as building rooftops, roads, and industrial facilities, while excluding urban green spaces and water bodies. The measure describes horizontal land coverage and contains no information on building height or floor space. This machine-learning classification provides a significant improvement over index-based measures such as the Normalized Difference Built-up Index (NDBI). NDBI frequently misclassifies bare earth, sand, and fallow land as built-up infrastructure due to their similar reflectance (He, Shi, Xie, & Zhao, 2010). Furthermore, in cities with informal settlements, such as Medellín, which often have fragmented structures, NDBI can fail to detect fine urbanization patterns due to mixed-pixel effects and topographic shadows. The GLC_FCS30D's dynamic approach mitigates these errors.

Urban Outcomes. Urban outcomes such as population, density and demographic composition are computed using granular census data at the block level taken from the Colombian Census Bureau (DANE). Census data covers the universe of residents

confirming the existence of such funds.

in Colombia, particularly those in the city of Medellin. Due to the cost of population data collection, censuses are collected once per decade. Data is available for 1993, 2005 and 2018.

Other data. I use data on workplace and employment location from the Chamber of Commerce from Medellin. Income comes from the *Gran Encuesta Integrada de Hogares* and the *Encuesta de Calidad de Vida*, from where I also obtain data on housing expenditure. Rents are obtained from the cadastre. For commuting flows, I digitized the origin destination survey from 1993 and combine it with subsequent origin destination surveys from 2012 and 2017. Additional data details can be found in A.

3.2 Descriptive Statistics and Balance

Table 1 presents balance tests between Upgraded (Phase 1) and Phase 2 slums. It also compares balance between upgraded areas and the subset of slums that were classified as Level II from Phase 2. The comparison is a pair-wise univariate test from a regression at the block level in baseline (1993), weighted by the log of block population in that year. The results suggest that across most characteristics blocks from upgraded neighborhoods were similar to blocks in both Phase 2 and Level 2 slums. Average household size and the number of low-skill adults show marginal differences at the 10% level when compared with Phase 2 slums. None of the reported differences is significant at the 5% level. These marginal differences are absent with Level II controls, which share the historical consolidation category of upgraded slums.

4 The effects of upgrading on natural disasters

The first thing to examine is whether placing risk mitigation infrastructure really reduced the incidence of natural disasters. One concern is that Phase I slums may have been targeted because of their exposure to natural disasters. If this is the case, a subsequent decline could reflect pre-existing differences rather than the pure impact of PRIMED. To address this concern, I compare the trends between upgraded Phase I slums and Phase II neighborhoods[4], those scheduled for intervention but not upgraded. I estimate a dynamic difference-in-differences model using annual counts of natural disasters. Because disasters are infrequent and some years have very few[5], I group years into bins of roughly three years and estimate one coefficient per bin. In particular, I estimate Equation 1, where y_{st} is the number of disasters recorded in slum s in year t, $Upgraded_s$ indicates whether slum s was upgraded by PRIMED Phase I, and $\mathbf{1}[t \in p]$ indicates that year t belongs to period p. The omitted period is 1990-92, the last one be-

[4]I refer to slums and neighborhoods interchangeably; they correspond to the same boundary.

[5]El Niño episodes can make certain years quite dry, so there are few or no precipitation-triggered events.

Table 1: Balance at baseline between upgraded and control group slums in baseline (1993)

	Upgraded	Phase 2 control			Level II control		
Variable	Mean	Mean	Diff.	p-value	Mean	Diff.	p-value
Physical characteristics							
Block area (m^2)	2,960.354	3,689.243	-728.889	0.169	3,544.337	-583.983	0.323
Mean elevation (m)	1,714.633	1,696.660	17.973	0.593	1,734.503	-19.870	0.598
Distance to CBD (km)	4.695	4.568	0.128	0.864	4.750	-0.054	0.943
Urban density							
Log population density per km^2	11.015	10.929	0.085	0.448	10.992	0.023	0.865
Log housing density per km^2	9.486	9.465	0.021	0.836	9.506	-0.020	0.862
Housing characteristics							
Water access	0.976	0.965	0.012	0.561	0.965	0.011	0.684
Sewerage access	0.922	0.902	0.020	0.555	0.899	0.023	0.633
Electricity access	0.996	0.996	0.000	0.814	0.996	0.000	0.872
Average rooms	3.060	3.426	-0.366	0.136	3.281	-0.221	0.385
Brick walls	0.815	0.821	-0.006	0.845	0.816	-0.001	0.969
Population and household characteristics							
All residents	157.614	200.811	-43.196	0.198	216.531	-58.916	0.222
Low-skill adults	59.222	85.378	-26.156*	0.091	90.912	-31.690	0.163
High-skill adults	8.022	10.636	-2.615	0.314	9.670	-1.648	0.584
Average household size	4.536	4.387	0.148*	0.074	4.437	0.099	0.224
Households per dwelling	1.163	1.153	0.010	0.652	1.154	0.009	0.754
People per room	1.988	1.672	0.316	0.110	1.777	0.211	0.310
Literacy rate	0.877	0.901	-0.024	0.144	0.889	-0.012	0.506
Born in Medellin	0.509	0.545	-0.036	0.122	0.525	-0.016	0.521
Lived in Medellin 5 years earlier	0.825	0.847	-0.022	0.128	0.834	-0.010	0.579
N (baseline blocks)	366	648			397		

Notes: Means are weighted by log 1993 block population. Differences are treated minus control. Standard errors are clustered by neighborhood. Skill counts cover adults aged 25+; high skill means college or higher. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

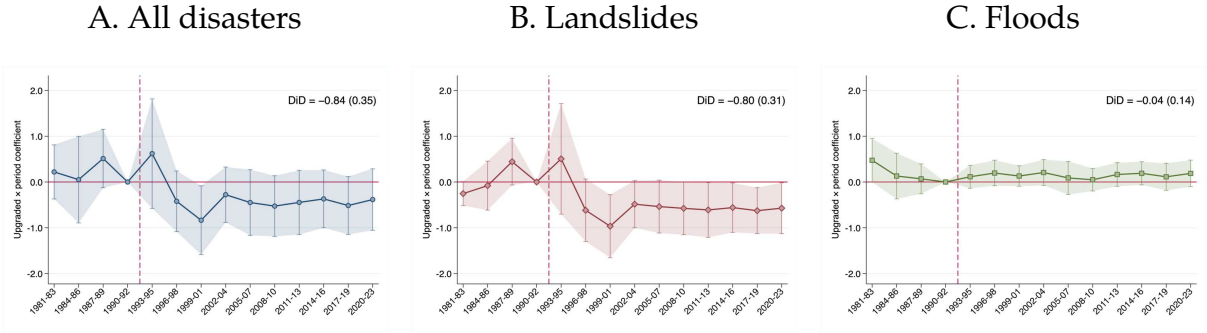
fore the intervention, so each coefficient β_p measures the difference between upgraded and control slums in period p relative to that difference in 1990–92. α_s denotes slum fixed effects, which capture time-invariant characteristics of slums, and λ_t denotes year fixed effects, which capture common shocks. Standard errors are clustered at the slum level, the level of treatment.

$$y_{st} = \alpha_s + \lambda_t + \sum_{p \neq 1990\text{–}92} \beta_p \left[Upgraded_s \times \mathbf{1}[t \in p] \right] + \varepsilon_{st}, \tag{1}$$

Figure 4 presents the results of estimating Equation 1. Panel A shows the effects on the total number of disasters, which includes landslides and floods. Panels B and C disaggregate the result for each type of event. This figure allows to highlight three findings. First, upgraded and control slums followed similar trends in natural disasters before the intervention: none of the pre-period coefficients for total disasters is statistically different from zero. Second, there is a decline in the number of disasters events after the policy is implemented. The decline becomes visible a few years after the intervention began, is largest in 1999-2001, and the estimates remain negative in all

subsequent periods. Third, the large majority of this reduction is due to the decline in landslides, as I do not find a significant reduction in the number of floods. To summarize the magnitude, I replace the period indicators with a single indicator for the years after 1993. The estimated average effect on upgraded slums is a reduction of 0.84 disasters per slum per year relative to 1981-93 (s.e. 0.35, $p = 0.02$). For landslides the reduction is 0.80 per year (s.e. 0.31, $p = 0.01$), while the effect on floods is small and not statistically significant (-0.04, $p = 0.77$). The fact that I observe a reduction only in landslides but not in floods is consistent with landslides being the most important risk that slums face, as they are located on the hillsides of the city, often in very steep areas. For example, Medellín sits at about 1,500 meters above sea level, and a slum can be located 600 meters above this height.

Figure 4: Trends in natural disasters between upgraded and control group slums



Note: Phase II controls. Annual neighborhood reports in three-year bins, except for the final 2020–23 bin; 1990–92 reference period. 95% confidence intervals use standard errors clustered at the slum/neighborhood level.

4.1 Dynamics around rainy events

The previous results suggest that natural disasters declined in upgraded slums relative to the control group. Because heavy rainfall can trigger these events, I next ask whether the decline is also present around heavy-rainfall episodes, when risk mitigation infrastructure is most likely to matter. The key intuition is that upgrading provides protection only when the hazard, atypical precipitation, appears. To do this, I examine how disasters change around heavy-rainfall episodes. I estimate this dynamic response between upgraded and non-upgraded slums by estimating the event-study specification in equation 2.

$$y_{st} = \alpha_s + \lambda_y + \delta D_{sy} + \sum_{k \in \mathcal{K},\, k \neq -1} [\theta_k + \beta_k D_{sy}]\, \mathbf{1}\{\tau_t = k\} + \varepsilon_{st}. \tag{2}$$

Where $D_{sy} = Post_y \times Upgraded_s$, where $Post_y = 1$ for years $y > 1993$ and $Upgraded_s$ indicates whether slum s was treated by PRIMED. The variable τ_t denotes event time

relative to a heavy-rainfall episode. The coefficients β_k trace whether disaster risk changes differently in upgraded slums before and after heavy rainfall. Slum fixed effects absorb permanent differences across neighborhoods, while year fixed effects absorb common shocks. Standard errors are clustered at the slum level. A heavy-rainfall episode occurs when cumulative precipitation over seven days exceeds the 80th percentile of the 1981–2023 distribution. As robustness checks, I analyze dynamics around alternative precipitation thresholds and over alternative rainfall windows.

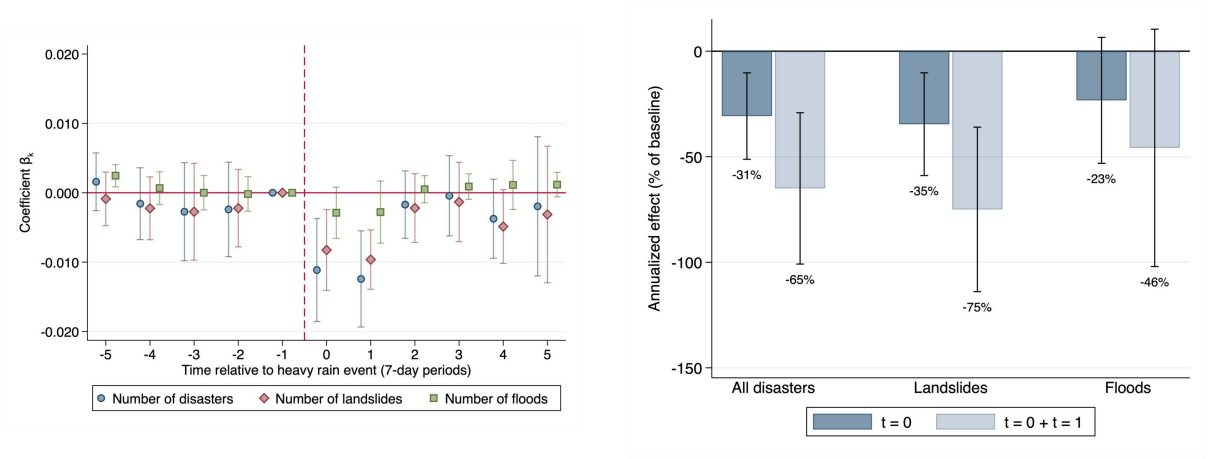
Figure 5 shows the results from equation (2). Before heavy rainfall, the estimates for total disasters and landslides are generally small and statistically indistinguishable from zero. During extreme rainfall events (at $t = 0$), the number of disasters declines significantly in upgraded slums. These results again suggest that this effect is largely due to the decline of landslides, as I do not find effects on the incidence of floods. To interpret the magnitude of these estimates, I annualize the event study point estimates. This allows me to interpret the impacts relative to the number of disasters during heavy-rain events per year. Table 2 presents these effects. My estimates indicate that in total slum upgrading reduced about 0.22 disasters per year. During rainy weeks, relative to the pre-intervention period (year ≤ 1993), there is a reduction in disasters and landslides equivalent to 31% and 35%, respectively. The effects are even larger when we consider that the infrastructure in place still reduces these type of events one week after the heavy rain has occurred ($t = 0 + t = 1$): compared to pre-upgrading, treated slums see a 75% reduction relative to non-upgraded Phase II slums.

These reductions in disasters are consistent with a protective role for the program's infrastructure investments. Descriptive drainage data offer suggestive support for this interpretation: upgraded neighborhoods have a larger share of pipe length in pipes at least 900 mm in diameter and higher average estimated pipe capacity, measured in liters per second (Appendix Table 11 and Figure 16). Drainage was one component of the program's protective works, which also included slope stabilization and retaining walls.

4.2 Disasters with damage

Having examined whether PRIMED Phase I reduced the incidence of natural disasters, I next ask whether it also reduced events that damage housing. This distinction matters because not all climate events have the same consequences for households. Data on the exact counting of destruction is limited over time. However, I can identify whether an event partially damaged or destroyed at least one home. I group these events as disasters with housing damage. I present the results in Figure 6. Panel A compares disasters with VS without housing damage, while Panel B separates the former into landslides and floods. I find that during extreme rainy weeks there is a reduction of both type of events, with slightly larger reductions in events with damage. These results suggest

Figure 5: Dynamic effects of slum upgrading on natural disasters



Note: Coefficients are estimated from equation (2). The outcome is the number of recorded disaster events in a neighborhood-window. Standard errors are clustered at the neighborhood level.

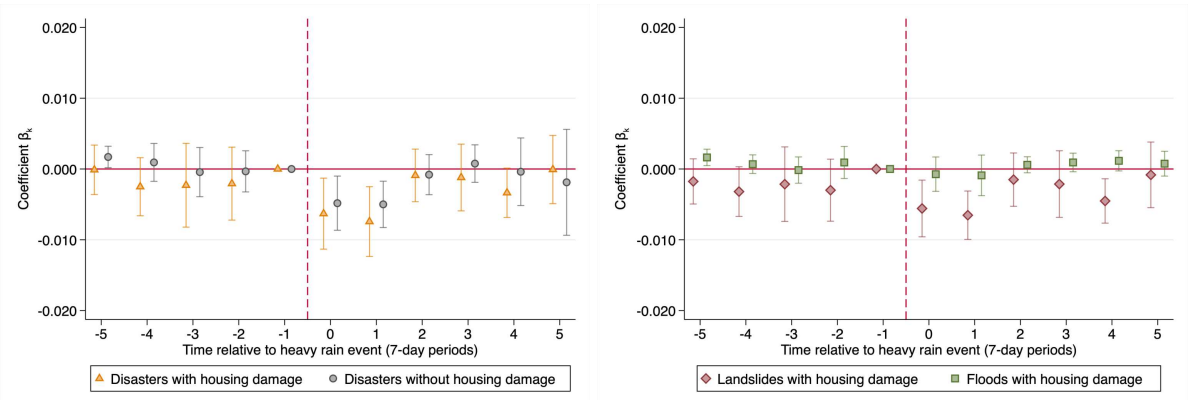
Table 2: Effects of slum upgrading on natural disasters

Panel A. Event-Study Point Effects					
Outcome			Point effect		
			$t=0$	$t=1$	$t=0+t=1$
Disasters			-0.0111***	-0.0124***	-0.0236***
Landslides			-0.0083***	-0.0096***	-0.0179***
Floods			-0.0029	-0.0028	-0.0057
Panel B. Annualized Effects					
Outcome	Baseline rate disasters/yr in heavy-rain windows	Heavy-rain windows/yr	Annualized effect (% of baseline rate)		
			$t=0$	$t=1$	$t=0+t=1$
Disasters	0.347	9.6	-0.107*** (-31%)	-0.119*** (-34%)	-0.226*** (-65%)
Landslides	0.229	9.6	-0.079*** (-35%)	-0.092*** (-40%)	-0.171*** (-75%)
Floods	0.119	9.6	-0.028 (-23%)	-0.027 (-23%)	-0.054 (-46%)

Notes: Coefficients are estimated from equation (2). The outcome is the number of recorded disaster events in a neighborhood-window. I annualized the point effects by multiplying the estimated coefficients by the average number of heavy-rain windows/year. Standard errors are clustered at the neighborhood level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

that upgrading is able to protect homes (and therefore people) when is needed the most.

Figure 6: Effects of slum upgrading on disasters with and without housing damage



Notes: Coefficients β_k from equation (2), estimated separately for each outcome on the full sample of slum-windows. Housing damage indicates that the report records affected, damaged or destroyed houses. Heavy rain is 7-day rainfall above the 80th percentile. 95% confidence intervals; standard errors clustered at the slum level.

4.3 Robustness

The heavy-rainfall results are robust to several changes in the analysis. Appendix Section B varies the rainfall threshold and window length, excludes Phase II controls adjacent to upgraded slums, and restricts the comparison group to Level II slums. The estimated effects on disasters and landslides remain negative across these specifications (Table 13). The appendix also examines possible spillovers to neighboring slums and reports confidence intervals based on Conley standard errors, which allow for spatial correlation.

5 Effects of upgrading on urban growth

The previous section showed that upgrading reduced the incidence of natural disasters, particularly landslides. I now ask whether upgraded slums also followed a different trajectory of urban growth. Evaluating the long-term effects of these infrastructure investments requires examining how population and housing respond, since PRIMED may have affected households' decision about where to live and build. Upgrading may have made treated neighborhoods more attractive, through both the general neighborhood improvement but also due to the risk mitigation infrastructure. If these investments encourage additional settlement, then more residents may live in areas where some risk remains, even if these neighborhoods have become safer on average. These residential

responses help determine the distribution of the program's benefits and motivate the subsequent analysis of citywide welfare.

I examine this residential response in three steps. First, I analyze whether upgrading increased the number of residents and changed the population composition of treated slums, relative to the control group. Second, I examine changes in urban density, paying particular attention to whether housing expanded alongside population. Third, I analyze whether the program affected housing characteristics, as upgrading could have affected dwelling quality.

To examine these changes, I use block-level census data and grid-cell-level remote-sensing measures. The outcomes are resident counts and skill composition, housing and population density, impervious surface, and housing characteristics, including connectivity to water and sanitation, construction materials and crowding. For each outcome, I estimate a difference-in-differences model of the form:

$$Y_{isy} = \beta_1 Upgraded_s + \beta_2 Post_y + \beta_3 Post_y \times Upgraded_s + \gamma_s + \psi_{c(s),y} + \epsilon_{isy}. \quad (3)$$

Where Y_{isy} is an outcome for block or grid cell i in slum s and year y. The variable $Upgraded_s$ indicates whether slum s was treated by PRIMED, and $Post_y$ equals one for census years after 1993 (2005 and 2018). The coefficient of interest is β_3, which captures the differential change in outcomes for upgraded slums after PRIMED. The terms γ_s and $\psi_{c(s),y}$ denote slum fixed effects and comuna-by-year fixed effects, respectively. Standard errors are clustered at the slum level. Because PRIMED combined risk mitigation with other neighborhood infrastructure improvements, these estimates capture the effect of the program as a whole.

5.1 Population Growth and Skill Composition

My results suggest that upgrading increased population, with adult population concentrated among low-skill residents. Table 3 reports these findings. On average, an upgraded block experienced a significant increase of 49.1 residents per block, relative to Phase II control group. This effect represents about 28% of the control mean in 1993. Among residents aged 25+, the estimated increase is 18 people, with the great majority being low-skill residents (17.5). I find no significant effects of increases in the number of high skill residents nor changes in the high-skill share. I interpret the later result as evidence of no changes in the population composition. This implies that the program did not lead to gentrification or re-sorting effects on the long-run.

I also estimate these effects at the neighborhood level by aggregating block-level population in treated and control areas. Appendix Table 14 suggest sizable population increases. On average, an upgraded neighborhood experienced an increase of 2,563

Table 3: Effects of Slum Upgrading on Population Size and Skill Composition

Control group:	Phase 2				
	(1) All residents	(2) Adults 25+	(3) Low-skill adults	(4) High-skill adults	(5) High-skill share
Post × Upgraded	49.075**	18.011**	17.484***	0.527	-0.005
	(21.304)	(7.550)	(6.245)	(2.973)	(0.018)
Control mean (1993)	175.412	83.461	74.065	9.397	0.125
R^2	0.228	0.285	0.305	0.150	0.354
Observations	3,042	3,042	3,042	3,042	3,042
Neighborhood FE	Yes	Yes	Yes	Yes	Yes
Comuna × year FE	Yes	Yes	Yes	Yes	Yes

Notes: Regressions use census-block-year observations from 1993, 2005, and 2018 and are weighted by log 1993 block population. Columns (2)–(4) count residents aged 25+ with classified education; high skill denotes college education or higher. Column (5) is the high-skill share among classified adults aged 25+. Standard errors are clustered at the neighborhood level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

residents of all ages and an increase of 1,357 adults aged 25+, which is corresponds to roughly 55% of the baseline control mean. This estimate in residential adults is going to be informative for the welfare analysis later introduced in section 6.

Disaggregating population by age and sex shows increases in the numbers of children (aged 0-14), working-age residents (15-64), and female residents, alongside a decline in residents aged 65 and above (Appendix Table 25). I also examine literacy and migration background. I find no statistically significant changes in the literacy rate or in the shares of residents born in Medellin or living in the city five years before the census (Appendix Table 26). The census records municipality of birth and municipality of residence five years earlier, but does not identify previous neighborhoods of residence. Thus, the latter indicator distinguishes current residents who already lived in Medellin five years earlier from those who lived outside the city. It cannot distinguish residents who remained in the same neighborhood from those who moved within the city.

5.2 Density and Impervious Surface

Table 4 describes the urban development accompanying the population increase. Housing density rises by 6,384 dwellings per km^2, and population density rises by 21,075 residents per km^2, equivalent to 41% and 31% of their respective 1993 control means. The log housing-density coefficient of 0.256 implies an increase of approximately 29%. The log population-density estimate is positive but not statistically significant. The estimated effect on impervious surface is 0.9 percentage points and is also statistically insignificant. Restricting the control group to Level II slums, which shares the histori-

cal consolidation category of upgraded slums, yields statistically significant estimates for both log population density and impervious surface (Appendix Table 17). With this comparison group, the log population-density estimate implies an increase in population density of 47.1%, and the estimated impervious-surface effect is 6.4 percentage points. Figure 7 presents the dynamics estimates, using 1993 as the reference year for census density outcomes and 1990 for impervious surface.

Table 4: Effects of Slum Upgrading on Density and Impervious Surface

Control group:	Phase 2				
	Density levels		Density logs		Land cover
	(1) Housing per km^2	(2) Population per km^2	(3) Log housing per km^2	(4) Log population per km^2	(5) Impervious surface = 1
Post × Upgraded	6,384***	21,075**	0.256**	0.194	0.009
	(2,087)	(8,548)	(0.105)	(0.124)	(0.030)
Control mean (baseline year)	15,547	67,337	9.396	10.843	0.735
Neighborhood FE	Yes	Yes	Yes	Yes	Yes
Comuna-by-year FE	Yes	Yes	Yes	Yes	Yes
Unit of observation			Block-year		30m cell-year
Sample years			1993, 2005, 2018		1985–2020
Observations	3,042	3,042	3,042	3,042	98,160
R-squared	0.170	0.119	0.326	0.218	0.160

Notes: Columns (1)–(4) use census-block-year observations from 1993, 2005, and 2018, weighted by log 1993 block population. Column (5) uses 30 m × 30 m cell-year observations at five-year intervals over 1985-2020, weighted by log 1993 neighborhood population; its baseline control mean is for 1990. Column (5) equals one if the cell is classified as impervious surface and zero otherwise; it measures land coverage, not building height. Standard errors are clustered at the neighborhood level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

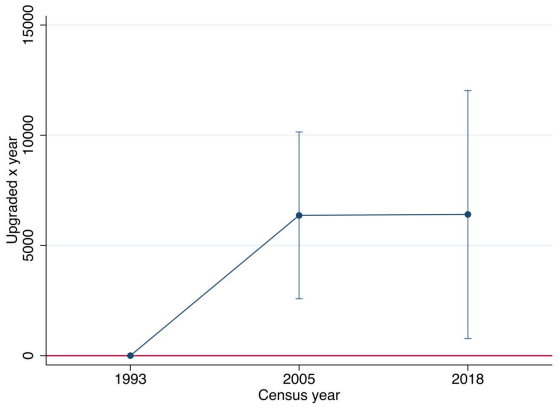
5.3 Housing Characteristics

Table 5 reports the estimated effects on housing characteristics. I find no statistically significant effects on water access, sewerage access, the share of dwellings with brick walls, or rooms per dwelling. These estimates provide limited evidence of changes in the measured housing attributes, but do not rule out improvements in infrastructure quality or housing resilience. In particular, household connection indicators do not capture drainage capacity. Descriptive evidence in Appendix A.6 shows higher estimated pipe capacity in upgraded neighborhoods than in Phase II neighborhoods, although this comparison does not establish a causal effect of PRIMED.

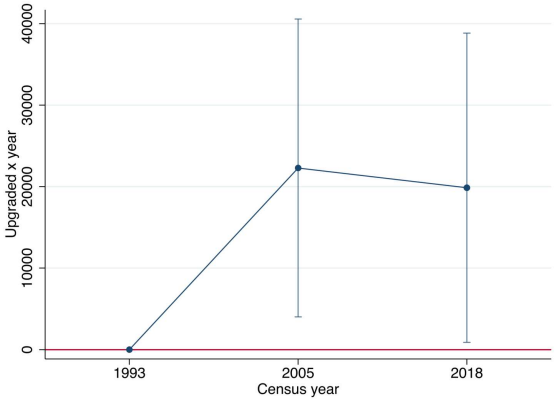
The estimates also suggest a reduction in housing crowding. Between 1993 and 2018, people per room declined by an estimated 0.293 relative to Phase II neighborhoods. This estimate is statistically significant at the 10% level and provides suggestive evidence that housing accommodated residents with less crowding, despite the population growth documented above.

Figure 7: Urban Growth Dynamics

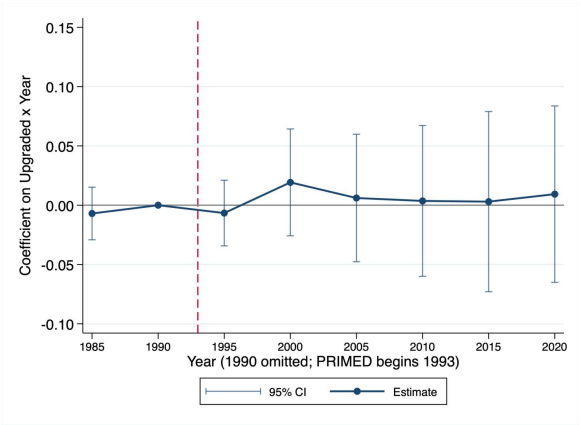
Panel A. Housing Density



Panel B. Population Density



Panel C. Impervious Surface



Note: Panels A and B use census-block-year observations. Panel C uses 30 m × 30 m cell-year observations. Census regressions are weighted by log 1993 block population; the impervious-surface regression is weighted by log 1993 neighborhood population. Standard errors are clustered at the neighborhood level. 95% confidence intervals shown.

5.4 Additional analysis and robustness

I also repeat the urban growth analysis using Level II neighborhoods scheduled for Phase II as controls. These neighborhoods share the historical consolidation category of upgraded slums. Using this more similar group yields even larger estimated effects on population, density, and impervious surface, reinforcing the evidence that PRIMED generated economically substantial urban growth (Appendix Tables 15, 17, and 23). Appendix Figure 22 presents the corresponding dynamic estimates with 95% confidence intervals.

I also examine sensitivity to block and cell fixed effects (Appendix Table 18) and to spatial correlation in inference (Appendix Tables 16, 19, 20, and 24). The positive effects on population growth and urban density remain statistically significant under these alternative specifications and inference procedures.

One concern is that proximity to upgraded neighborhoods may also attract resi-

Table 5: Effects of Slum Upgrading on Housing Characteristics

Control Group:	Phase 2				
	(1)	(2)	(3)	(4)	(5)
	Water access	Sewerage access	Brick walls	Avg. rooms	People per room
Post × Upgraded	-0.017	-0.028	0.023	0.196	-0.293*
	(0.014)	(0.032)	(0.045)	(0.187)	(0.170)
Control mean (1993)	0.956	0.887	0.817	3.389	1.690
Neighborhood FE	Yes	Yes	Yes	Yes	Yes
Comuna-by-year FE	Yes	Yes	Yes	Yes	Yes
N	3042	3042	2028	2028	2028
R-squared	0.275	0.421	0.314	0.422	0.628

Notes: Regressions are weighted by log 1993 block population, with standard errors clustered at the neighborhood level. Columns (1) and (2) use census-block-year observations from 1993, 2005, and 2018; columns (3)–(5) use only 1993 and 2018, when rooms and wall materials were observed. Water, sewerage, and brick-wall outcomes are shares. People per room is the block average of dwelling-level ratios of residents to rooms. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

dents to nearby control areas. Positive population spillovers to these areas would attenuate the estimated treated-control difference. I examine sensitivity to this possibility by excluding adjacent Phase II neighborhoods (four slums) or, separately, by excluding control blocks within distances ranging from 100 to 700m of upgraded areas (Figure 23). Appendix Section C.2.4 shows that population estimates remain positive and are slightly larger after excluding adjacent neighborhoods or control blocks within 500m, although it declines at wider exlusions. Housing-density estimates also remain positive, although they are not consistently larger.

Appendix Section C.2.5 examines potential spillovers more directly by comparing changes in population and housing density between untreated blocks near upgraded areas (0-300 m) and those in more distant reference bands. The comparison includes both Phase II blocks and other untreated blocks. Neither differential is statistically significant.

6 A Model of Slum Upgrading under Climate Risk

The previous results show that PRIMED reduced disaster exposure and increased population and housing density in upgraded neighborhoods, with population gains concentrated among low skill residents. These local changes do not directly measure residents' welfare gains because households can relocate, housing supply can expand, and rents and wages can adjust throughout the city. I therefore develop a quantitative spatial model to translate the estimated effects of upgrading into welfare changes for low and high skill workers. The framework builds on Ahlfeldt et al. (2015); Tsivanidis (2026) but departs from it by incorporating exposure to climate related natural

hazards through amenities. The model distinguishes improvements in neighborhood infrastructure from reductions in disaster exposure, allowing me to quantify their contributions to welfare and examine how the benefits of upgrading change under future climate risk.

6.1 Workers

Workers differ by skill group $g \in \{L, H\}$, where L denotes low skill and H denotes high skill. They choose where to reside in neighborhood $i \in \mathcal{R}$ and to work in location $j \in \mathcal{M}$, where $\mathcal{R} \subseteq \mathcal{M}$. Workers derive utility from consumption of the final good c_{ijg} and residential floor space h_{ijg}. They experience an idiosyncratic residence preference shock z_i and a worker-firm match productivity shock v_j. z_i is drawn from an independent Fréchet distribution with unit scale and shape parameter η_g, $F(z_i) = e^{-z_i^{-\eta_g}}$, where $\eta_g > 1$ governs the dispersion of residence choice for skill group g. v_j is drawn from an independent Fréchet distribution where $\widetilde{T}_g$ is the scale and ε_g is the shape parameter $F(v_j) = e^{-\widetilde{T}_g v_j^{-\varepsilon_g}}$, where $\varepsilon_g > 1$ governs the dispersion of workplace match productivities for skill group g.

The resident's utility is therefore

$$U_{ijg} = \left(\frac{B_{ig} z_i}{d_{ij}}\right) \left(\frac{c_{ijg}}{\beta_g}\right)^{\beta_g} \left(\frac{h_{ijg}}{1-\beta_g}\right)^{1-\beta_g}, \qquad 0 < \beta_g < 1, \tag{4}$$

subject to the budget constraint $c_{ijg} + r_{iR} h_{ijg} = w_{jg} v_j$, where w_{jg} is the wage per effective unit of labor of type g in workplace j, r_{iR} is the residential rent in location i, and $d_{ij} = e^{\kappa t_{ij}}$ is commuting disutility, with t_{ij} denoting travel time and $\kappa > 0$ governing the sensitivity to commuting.

Endogenous amenities. Residential amenities combine neighborhood fundamentals, disaster exposure, and local skill composition:

$$B_{ig} = b_{ig} e^{\gamma D_i} s_{iH}^{\chi}, \tag{5}$$

where b_{ig} captures the fundamental attractiveness of neighborhood i to skill group g. This component includes local infrastructure such as roads, sidewalks, paths, and stairs, as well as other neighborhood characteristics valued by residents. $D_i \geq 0$ denotes the average frequency of natural disasters in neighborhood i, whereas the parameter γ is the semi-elasticity of amenities with respect to exposure: a negative value implies that greater exposure makes a neighborhood less attractive, holding its other characteristics fixed.

Following Tsivanidis (2026) and Diamond (2016), amenities also depend on the

local share of high-skill residents,

$$s_{iH} = \frac{N_{RiH}}{N_{RiL} + N_{RiH}},$$

with χ governing the strength of this residential externality.

PRIMED enters the model through improvements in fundamental amenities, b_{ig}, and reductions in disaster exposure, D_i.

Optimal demands.

$$h^*_{ijg} \;=\; \frac{(1-\beta_g)\, w_{jg} v_j}{r_{iR}}, \qquad c^*_{ijg} \;=\; \beta_g\, w_{jg} v_j. \tag{6}$$

Indirect utility. Substituting (6) into (4) gives

$$U_{ijg} = B_{ig} z_i r_{iR}^{\beta_g - 1} \frac{w_{jg} v_j}{d_{ij}}. \tag{7}$$

Workplace choice. Conditional on residence i, the probability of working in j is

$$\pi_{ij|ig} = \frac{(w_{jg}/d_{ij})^{\varepsilon_g}}{\sum_{k\in\mathcal{M}} \left(\frac{w_{kg}}{d_{ik}}\right)^{\varepsilon_g}} \equiv \frac{(w_{jg}/d_{ij})^{\varepsilon_g}}{\Phi_{Rig}} \tag{8}$$

where Φ_{Rig} is residential commuter market access (RCMA), which summarizes access from residence i to all workplaces.

Given i, the expected value of the best workplace option in (7), $\max_j w_{jg} v_j / d_{ij}$, is $T_g \Phi_{Rig}^{1/\varepsilon_g}$, with $T_g = \Gamma\left(1 - \frac{1}{\varepsilon_g}\right) \widetilde{T}_g^{1/\varepsilon_g}$ and $\Gamma(\cdot)$ the Gamma function.

Mean labor income of residents in location i and skill group g is

$$\bar{w}_{ig} = T_g \sum_{j\in\mathcal{M}} w_{jg} \pi_{ij|ig}^{1-1/\varepsilon_g} \tag{9}$$

Residence choice. Before drawing workplace match productivities, the probability of choosing residence i is

$$\pi_{Rig} = \frac{\left(B_{ig} r_{iR}^{\beta_g - 1} T_g \Phi_{Rig}^{1/\varepsilon_g}\right)^{\eta_g}}{\sum_{\ell\in\mathcal{R}} \left(B_{\ell g} r_{\ell R}^{\beta_g - 1} T_g \Phi_{R\ell g}^{1/\varepsilon_g}\right)^{\eta_g}}. \tag{10}$$

Given city population N_g for group g, the number of residents in location i is

$$N_{Rig} = N_g \pi_{Rig} \tag{11}$$

The joint residence-workplace probability is

$$\pi_{ijg} = \pi_{Rig}\pi_{ij|ig}$$

Labor Market Clearing The supply of workers type g employed in location j, denoted N_{Mjg}, is equal to the number of residents of that group who choose to commute to j. Using the conditional commuting probabilities $\pi_{ij|ig}$, this implies

$$N_{Mjg} = \sum_{i \in \mathcal{R}} \pi_{ij|ig} N_{Rig} = w_{jg}^{\varepsilon_g} \Phi_{Fjg} \tag{12}$$

where $\Phi_{Fjg} = \sum_{i \in \mathcal{R}} d_{ij}^{-\varepsilon_g} \frac{N_{Rig}}{\Phi_{Rig}}$ is firm Commuter Market access (FCMA).

Production uses effective labor, which weights commuter headcounts by selected match productivity:

$$\widetilde{N}_{Mjg} = T_g \sum_{i \in \mathcal{R}} N_{Rig} \pi_{ij|ig}^{1-1/\varepsilon_g}. \tag{13}$$

Welfare. Expected welfare for skill group g is

$$\bar{U}_g = \Gamma\left(1 - \frac{1}{\eta_g}\right)\left[\sum_{i \in \mathcal{R}} \left(B_{ig} r_{iR}^{\beta_g - 1} T_g \Phi_{Rig}^{1/\varepsilon_g}\right)^{\eta_g}\right]^{1/\eta_g}. \tag{14}$$

Disaster exposure and welfare. The model implies that with $\gamma < 0$, greater disaster exposure reduces residential amenities. Holding everything else constant, the direct welfare effect is

$$\frac{\partial \log \bar{U}_g}{\partial D_i} = \gamma \pi_{Rig} < 0, \tag{15}$$

where π_{Rig} is the residential choice probability. A reduction in exposure, $dD_i < 0$, therefore raises expected welfare directly:

$$d \log \bar{U}_g = \gamma \pi_{Rig}\, dD_i > 0.$$

The quantitative model additionally accounts for the resulting equilibrium adjustments in residential choices, rents, and wages.

Migration. Workers choose between Medellín and the rest of Colombia based on expected utility and their preferences for each destination. National population and utility outside Medellín remain fixed. The parameter ζ_g governs migration responsiveness to relative utility. Let N_g denote Medellín's population of skill group g, and π_g^m its base-

line share of the national population of that group. Changes in city population satisfy

$$\widehat{N}_g = \frac{\widehat{\bar{U}}_g^{\zeta_g}}{\pi_g^m \widehat{\bar{U}}_g^{\zeta_g} + 1 - \pi_g^m}. \tag{16}$$

Hats denote ratios to baseline, with $\widehat{\bar{U}}_g$ capturing the change in expected utility from living in Medellín. Conditional on living in the city, workers make the residence and workplace choices described above. Expected welfare, incorporating preferences over the city and the outside option, changes according to

$$\widehat{W}_g = \left[\pi_g^m \widehat{\bar{U}}_g^{\zeta_g} + 1 - \pi_g^m\right]^{1/\zeta_g}.$$

6.2 Production

Firms in location j produce a final good y_j using a composite labor input N_j and commercial floor space H_{jM} under a Cobb-Douglas technology and perfect competition. As in Tsivanidis (2026), the composite labor input is a CES aggregate of low and high skill effective labor with elasticity of substitution σ:

$$y_j = A_j N_j^{\alpha} H_{jM}^{1-\alpha}, \qquad 0 < \alpha < 1. \tag{17}$$

$$N_j = \left(\sum_{g \in \{L,H\}} \alpha_{gj} \widetilde{N}_{Mjg}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \qquad \sigma > 0, \quad \sigma \neq 1, \quad \alpha_{gj} > 0. \tag{18}$$

Let N_{Mjs} denote employment of industry s in location j, $N_{Mj} = \sum_s N_{Mjs}$ total employment there, and α_{sg} the intensity of skill group g in industry s. The skill weights of location j are the employment-weighted average of the industry weights,

$$\alpha_{gj} = \sum_s \omega_{js}\, \alpha_{sg}, \qquad \omega_{js} = \frac{N_{Mjs}}{N_{Mj}},$$

where ω_{js} is the employment share of industry s in location j. These weights are held fixed across equilibria.

Given wages $\{w_{jL}, w_{jH}\}$ and commercial floor-space rents r_{jM}, firms maximize profits. Rearranging the first-order conditions yields the factor price equations

$$w_{jg} = \alpha \alpha_{gj} \frac{y_j}{N_j} \left(\frac{N_j}{\widetilde{N}_{Mjg}} \right)^{1/\sigma}, \qquad g \in \{L, H\}, \tag{19}$$

$$r_{jM} = (1-\alpha) \frac{y_j}{H_{jM}}. \tag{20}$$

Endogenous productivity. Local productivity depends on fundamental productivity and the surrounding density of employment:

$$A_j = a_j \left[\sum_{k \in \mathcal{M}} e^{-\delta t_{jk}} \frac{N_{MkL} + N_{MkH}}{L_k} \right]^{\lambda}. \tag{21}$$

Where L_k is land area, λ governs the strength of productivity spillovers, and δ their spatial decay. Individual firms take A_j as given. Fundamental productivity a_j remains fixed in the policy counterfactuals, while employment relocation changes productivity through the externality.

6.3 Floor Space

Floorspace H_i is produced using land (L_i) and capital (K_i), under a Cobb-Douglas technology, following Ahlfeldt et al. (2015) and Combes, Duranton, and Gobillon (2021), with local construction conditions captured by the supply shifter $\phi_i > 0$. This shifter captures differences in construction conditions that may reflect topography and other local development constraints. A higher ϕ_i allows developers to produce more floor space from a given amount of land and capital.

$$H_i = \phi_i L_i^{1-\mu} K_i^{\mu}, \qquad 0 < \mu < 1, \qquad \phi_i > 0.$$

Developers take local conditions and prices as given and choose land and capital to maximize profits:

$$\max_{K_i, L_i \geq 0} \left\{ r_i \phi_i L_i^{1-\mu} K_i^{\mu} - P K_i - R_i L_i \right\}, \tag{22}$$

where r_i is the rental revenue received by developers per unit of floor space, R_i is the rental price per unit of land, and P is the common cost of capital across the city. Land available for development in each neighborhood, L_i, is held fixed across equilibria. Combining the optimal capital choice with the construction technology yields neighborhood floor-space supply:

$$H_i = L_i \, \phi_i^{1/(1-\mu)} \left(\frac{\mu r_i}{P} \right)^{\mu/(1-\mu)} \tag{23}$$

Holding land availability, the construction supply shifter, and the capital price fixed, higher floor-space rents encourage capital investment and increase floor-space supply. Conversely, lower rental returns reduce the incentive to build, other things equal. The elasticity of floor-space supply with respect to rents is $\mu/(1-\mu)$, which is common across neighborhoods.

Floor Space Market Clearing

Floor space in location i can be allocated to residential or commercial use. Let $\theta_i \in [0, 1]$ denote the share of floor space used commercially:

$$H_{iM} = \theta_i H_i,$$
$$H_{iR} = (1 - \theta_i) H_i.$$

Let r_{iR} and r_{iM} denote the rents paid by households and firms, respectively, and $\xi_i > 0$ is a wedge between commercial and residential rents. Let r_i denote the developer return. In locations where residential and commercial uses coexist, $r_i = r_{iR} = \frac{r_{iM}}{\xi_i}$, whereas in areas where only residential use is supplied, $r_i = r_{iR}$ and where only commercial use is supplied, $r_i = r_{iM}/\xi_i$. This implies that the following no-arbitrage condition between residential and commercial uses holds:

$$\begin{aligned} &\theta_i = 1 && \text{if } r_{iM} > \xi_i r_{iR}, \\ &\theta_i \in [0, 1] && \text{if } r_{iM} = \xi_i r_{iR}, \\ &\theta_i = 0 && \text{if } r_{iM} < \xi_i r_{iR}. \end{aligned} \tag{24}$$

Residential floor-space market clearing requires total rental payments to equal residents' expenditure on housing:

$$r_{iR} H_{iR} = \sum_{g \in \{L,H\}} (1 - \beta_g) \bar{w}_{ig} N_{Rig}. \tag{25}$$

Similarly, commercial floor-space market clearing requires total rental payments by firms in neighborhood i to equal their expenditure on commercial floor space:

$$r_{iM} H_{iM} = (1 - \alpha) y_i. \tag{26}$$

Equilibrium

Given worker and migration parameters $\{\beta_g, \eta_g, \varepsilon_g, \widetilde{T}_g, \zeta_g\}_{g \in \{L,H\}}$, production parameters $\alpha, \sigma, \{\alpha_{gj}\}, \lambda, \delta$, construction parameter μ, and parameters κ, χ, γ, consider an economy with location fundamentals $\{a_j, b_{ig}, \phi_i, \xi_i\}$, land endowments $\{L_i\}$, travel times $\{t_{ij}\}$, and disaster exposure $\{D_i\}$. The common capital price P, national populations, and utility outside Medellín are given. The final-good price is normalized to one.

An open-city equilibrium consists of wages $\{w_{jg}\}$, developer returns $\{r_i\}$, city populations $\{N_g\}$, expected utility levels $\{\bar{U}_g\}$, endogenous productivity $\{A_j\}$ and amenities $\{B_{ig}\}$, and the associated rents and allocation of workers, capital, and floor space such that:

(i) Workers optimize: consumption and housing demands satisfy (6), residence and

workplace choices satisfy (10) and (8), and residential populations satisfy (11). Amenities are consistent with neighborhood skill composition through (5). Expected utility is given by (14), and city populations satisfy the migration condition (16).

(ii) Firms and developers optimize: firms' factor demands satisfy the first-order conditions (19)–(20), with productivity satisfying (21). Developers solve (22), yielding floor-space supply (23), and the allocation between residential and commercial uses satisfies (24) where both uses are feasible.

(iii) Markets clear: workplace employment equals commuter headcounts in (12), and firms' effective-labor demand equals supply in (13). Residential and commercial floor-space markets satisfy (25) and (26). Aggregate land demand in each location equals the endowment L_i.

In the closed-city equilibrium, group populations N_g are fixed and the migration condition (16) is omitted.

7 Model Quantification and Results

I represent Medellín using 269 spatial locations: 248 residential neighborhoods and 21 workplace-only locations. Residential neighborhoods can also accommodate employment. The baseline equilibrium corresponds to the observed city in 2018, with PRIMED Phase I implemented. Residential population consists of adults aged 25 and above. I distinguish high skill residents, defined as those with college education or higher, and low skill otherwise. To quantify the model, I combine neighborhood-level information on residential population by skill, workplace employment by industry sector, residential and commercial property values, land area, commuting times, and disaster exposure. Residential population comes from the 2018 census, while the spatial distribution of employment comes from Chamber of Commerce records. I rescale workplace employment so that its citywide total matches the model's resident population. Cadastral property values anchor relative residential and commercial prices. Commuting times comes from Google Maps API, and neighborhood disaster exposure is measured as the long term average between 1993-2018 from the disaster records.

Conditional on the structural parameters described below, I invert the model to recover the local fundamentals that rationalize the observed equilibrium. These include fundamental amenities, productivity, and construction supply shifters. Fundamental amenities capture residential attractiveness after accounting for rents, access to employment, disaster exposure, and local skill composition. The recovered fundamentals allow the model to reproduce baseline residential population, workplace employment, and rents. They provide the starting point for the counterfactual exercises, in which I

change the program related amenity and disaster exposure components and solve for a new equilibrium. Additional calibration details are provided in Appendix D.2.

7.1 Parameter Estimation

Commuting elasticity ε_g. Taking logs to equation (8) and substituting $d_{ijt} = \exp(\kappa t_{ijt})$, I obtain the following gravity equation:

$$\ln \pi_{ij|ig,t} = \alpha_{ig,t} + \delta_{jg,t} + \underbrace{\nu_g}_{-\varepsilon_g \kappa} t_{ijt} + u_{ijg,t} \tag{27}$$

where $\pi_{ij|ig,t}$ is the probability that a worker in skill group g, conditional on residing in neighborhood i, works in neighborhood j. $\alpha_{ig,t}$ and $\delta_{jg,t}$ are origin-skill and destination-skill by year fixed effects. The origin-group-time fixed effects absorbs origin-side labor supply and group-specific origin characteristics. The destination-group-time fixed effect absorbs destination-side labor demand and workplace attractiveness. The parameter of interest is ν_g, the semi-elasticity of commuting shares with respect to travel time. In the model, this parameter equals $-\varepsilon_g \kappa$, where ε_g governs the elasticity of commuting choices and κ maps minutes of travel time into commuting costs. I set $\kappa = 0.01$, following previous results in the literature (Ahlfeldt et al., 2015; Tsivanidis, 2026). Thus, given κ, I identify $\varepsilon_g = -\nu_g/\kappa$ from the sensitivity of commuting flows to travel time. I estimate this equation using PPML to allow for zeros (Santos Silva & Tenreyro, 2006).

Results are presented in Table (6). I find $\varepsilon_L = 7.6$ and $\varepsilon_H = 5.6$. This suggests that low-skilled workers are more sensitive to travel time than high-skill workers. These results are consistent with previous findings across Latin American cities. For example, In Bogotá, Tsivanidis (2026) finds a commuting elasticity of 4 for low-skilled workers and 2.44 for high skills. These results also align with and (Zarate, 2024) who find a commuting elasticity of 4.4 for informal workers and 3.11 for formal workers in Mexico City. In Medellin, Khanna et al. (2025) also find that the illegitimate sector (presumably with low education background) have a commuting elasticity of 6.76, where as the legitimate sector has an elasticity of 3.86. I perform a series of robustness checks around this parameter in Appendix D.1, where I use travel time from Google Maps API, data aggregated at the Comuna level, trips departing to work at all hours or with all purposes other than other. I also estimate it separately for each survey year, and weighting travel time from the network using the share of reported transportation modes (car, public transportation) in each year. The results are broadly consistent with these estimates.

Disaster Disamenity γ. In the model, disaster exposure affects how households value residential locations. Population adjustment provides an observable implication of residential choices, which offers an empirical approximation to this mechanism. To do

Table 6: Commuting elasticity

	(1) Overall	(2) Low skill	(3) High skill
Travel time (min.), ν_g	-0.069^{***}	-0.076^{***}	-0.056^{***}
	(0.002)	(0.003)	(0.003)
Implied elasticity, ε_g	6.94	7.60	5.60
Observations	82,989	76,944	44,448
Origin × year FE	Yes	Yes	Yes
Destination × year FE	Yes	Yes	Yes

Notes: Direct home-to-work trips departing between 6-9AM. The regression uses data from the 2012 and 2017 Ordigin Destination Surveys and travel time computed with the network corresponding to that year. Average travel time gives car and public transportation equal weight. Each observation is a neighborhood-origin-neighborhood-destination-year cell. High skill is university/postgraduate education. Regressions in cols (2) and (3) are estimated separately by skill, allowing skill-specific fixed effects. The regression is estimated by PPML. The reported travel-time coefficient is ν_g, where $\nu_g = -\varepsilon_g \kappa$. Implied elasticities are $\varepsilon_g = -\nu_g/\kappa$, with $\kappa = 0.01$. Standard errors are clustered by the neighborhood origin-destination pair (i, j). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

this, I compare changes in log population between high and low susceptibility blocks in upgraded neighborhoods, relative to the control group [6]. In particular, I estimate equation 28.

$$\begin{aligned} Y_{isy} = {} & \beta_1 Upgraded_s + \beta_2 Post_y + \beta_3 \left(Post_y \times Upgraded_s\right) \\ & + \beta_4 High\ Risk_s + \beta_5 \left(Post_y \times High\ Risk_s\right) + \beta_6 \left(Upgraded_s \times High\ Risk_s\right) \\ & + \beta_7 \left(Post_y \times Upgraded_s \times High\ Risk_s\right) + \gamma_s + \psi_{c(s),y} + \epsilon_{isy}. \end{aligned} \tag{28}$$

Table 7 documents a smaller population response in blocks with high landslide susceptibility. The parameter of interest is $\beta_7 = -0.31$, which measures the difference in population responses between high and low susceptibility blocks. It indicates that the residential adjustment is lower in areas where risk is higher, which is in line with the expected theoretical sign of γ. Additionally, this is suggestive that the residential adjustment is lower in areas where risk is higher (proxy of disasters), which is consistent with residents' disutility for exposure as describe in 15. This categorical contrast does not by itself identify the amenity semi-elasticity per disaster per year. I set $\gamma = -0.31$ as a baseline approximation informed by this pattern. Appendix Table 27 reports results using alternative measures of landslide risk, which yield negative estimates of β_7 of similar magnitude. Because I measure this parameter imperfectly, I test the sensitivity of the model to different values.

[6] The structural equation that is derived from 10 allows to obtain an expression that allows to estimate the structural parameter γ. However, as I do not observe rents and income at baseline in these neighborhoods, I cannot estimate γ through this way. Instead, I argue that from changes in observed population in low VS high risk area I can empirically approximate it.

Table 7: Changes in Population by Risk Category

	Log population
Panel A. Landslide Susceptibility Index	
Post × Upgraded	0.357***
	(0.128)
Post × Upgraded × High landslide susceptibility	−0.311**
	(0.116)
R-squared	0.274
Observations	3,042
Neighborhood FE	Yes
Comuna × Year FE	Yes

Notes: Regressions use census-block-year observations from 1993, 2005, and 2018 and are weighted by log 1993 block population. Standard errors are clustered at the neighborhood level. Estimates correspond to equation (28). High susceptibility is defined relative to the median. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Share Parameters $\beta_g, \alpha, \alpha_{sg}$. I compute the share of expenditure on housing $(1 - \beta_g)$ using data from the Encuesta de Calidad de Vida (ECV). I find that $1 - \beta_L = 0.375$ and $1 - \beta_H = 0.318$ (See Appendix A.7). Using data from the GEIH survey and $\alpha = 0.8$, I construct the intensity weights of low and high skill labor used by each industry s for the CES labor aggregate, as suggested by Tsivanidis (2026). The corresponding bill share and shares α_{sg} are in Appendix Table 12.

Rest of parameters. The rest of the parameters are shown in Table 8.

7.2 Model validation

I examine its external validity by analyzing the correlation between the model's generated outcomes and observed data, unseen by the model. The results are shown in Appendix 28. There is a high correlation between floorspace generated by the model and the observed number of housing units from the census (Pearson=0.86 and R^2=0.74). This indicates that in the context of Medellín, the model is able to capture urban development under one single technology with one construction supply shifter, ϕ_i rationalizing local spatial heterogeneity. To validate amenities for high-skill residents generated by the model, I construct a composite index of observed amenities which includes several dimensions which residents care about: establishments, crime, school quality, noise, and gang ruling (for more details see Appendix D.6). The correlation between amenities for high-skill residents generated by the model (B_{iH}) and this composite index is also high (Pearson=0.42 and R^2=0.18). Then I compare the model's predicted average income with the data. Despite the high dispersion that income can have in medium-large city such as Medellin, the model has a good correlation (Pearson=0.81 and R^2=0.66).

Table 8: Parameters

Parameter	Description	Value	Source
Worker heterogeneity and location choice			
β_L	consumption share from low-skilled workers; housing share is $1-\beta_L$	0.625	ECV surveys
β_H	consumption share from high-skilled workers; housing share is $1-\beta_H$	0.682	ECV surveys
η_L	Fréchet shape parameter for residence-location shocks	1.92	Tsivanidis (2026)
η_H	Fréchet shape parameter for residence-location shocks	2.01	Tsivanidis (2026)
ε_L	Fréchet shape parameter for workplace match shocks	7.6	Estimated: Origin-Destination Surveys
ε_H	Fréchet shape parameter for workplace match shocks	5.6	Estimated: Origin-Destination Surveys
ζ_L, ζ_H	elasticity of city-versus-rest location odds with respect to relative utility	3	Bryan and Morten (2019)
κ	sensitivity of commuting disutility to travel time	0.01	Ahlfeldt et al. (2015); Tsivanidis (2026)
T_L	Mean workplace match productivity	1.0	Normalized
T_H	Mean workplace match productivity, $T_g = \Gamma\left(1-\frac{1}{\varepsilon_g}\right)\widetilde{T}_g^{1/\varepsilon_g}$	18.405	Calibrated to match wage premium
Amenities and climate risk			
χ	elasticity of amenities with respect to the high-skill residential share	1.002	Tsivanidis (2026)
γ	semi-elasticity of amenities with respect to disaster exposure	-0.31	Estimated
Floorspace supply			
μ	capital share in floorspace production	0.75	Ahlfeldt et al. (2015)
Production			
α	labor share in final-good production	0.8	Tsivanidis (2026)
α_{sL}	CES labor weight for skill group L	Table 12	GEIH
α_{sH}	CES labor weight for skill group H	Table 12	GEIH
σ	elasticity of substitution between low- and high-skill labor	1.42	Card (2009)
λ	elasticity of productivity with respect to spatially weighted employment density	0.01	Gómez and Zárate (2024)
δ	travel-time decay of production externalities	0.3585	Ahlfeldt et al. (2015)

Notes: ECV: Encuesta de Calidad de Vida; GEIH: Gran Encuesta Integrada de Hogares.

7.3 Welfare Effects from PRIMED

I evaluate the welfare effects of PRIMED Phase I by comparing Medellín's observed 2018 equilibrium with a counterfactual equilibrium in which the program is removed. Without PRIMED (NP), upgraded neighborhoods face higher natural disasters exposure due to no risk mitigation infrastructure and lower local fundamental amenities as a result of no improvement in roads, stairs, and other neighborhood level infrastructure. To recover the counterfactual exposure to climate events D_i^{NP}, I embed my annualized event study estimates from Table 2 in the model. My results imply that PRIMED reduced the annual number of natural disasters during heavy-rain episodes and the following week by 0.22. Let $\beta_D \equiv -(\beta_0^A + \beta_1^A) > 0$ denote the magnitude of this annualized effect. Removing the program's protection therefore increases disaster exposure, thus the counterfactual is $D_i^{\text{NP}} = D_i^{\text{P}} + \beta_D$ for neighborhoods upgraded in Phase I.

Conditional on this effect, I calibrate a common fundamental amenity multiplier, τ, by using as identified moment the population difference-in-differences result from Table 14 (Nakamura & Steinsson, 2018). Specifically, I set $b_{ig}^{\text{NP}} = b_{ig}^{\text{P}}/\tau$[7]in Phase I neighborhoods and choose τ so that the combined open city equilibrium population response to amenity improvements and disaster protection reproduces the estimated treated–control difference. This targeted moment corresponds to 1,357 additional residents in treated neighborhoods, which corresponds to 54.8% of the control group's mean population in 1993[8]. The calibrated multiplier is 1.0412, implying a 4.12% improvement in fundamental amenities. I then compute the welfare effects and decompose them into contributions from changes in amenities, disaster exposure, residential composition, income access, and rents.

Additionally to this, I also analyze how Medellin would like if the program would have been completed. To do this, I do a similar exercise to the one previously described. I assume that these neighborhoods would receive the same proportional improvement in fundamental amenities and the same absolute reduction in disaster exposure as Phase I neighborhoods. Starting from their fundamentals recovered from the observed 2018 equilibrium, I therefore set $b_{ig}^{\text{Full}} = \tau b_{ig}^{2018}$ and $D_i^{\text{Full}} = \max\{0, D_i^{2018} - \beta_D\}$ for eligible Phase II neighborhoods. Phase I remains implemented in this scenario. I solve the resulting equilibrium and compare the completed program with the same no-PRIMED counterfactual used to evaluate Phase I.

I present the welfare results of PRIMED in Table 9. These welfare gains do not

[7]The fundamental amenity improvement may also capture gains in local accessibility through improved roads and stairs. If upgrading reduces travel time by $\Delta t_i > 0$ minutes for residents of neighborhood i, regardless of workplace destination, then $d_{ij}^{\text{P}} = e^{\kappa(t_{ij} - \Delta t_i)}$. At the level of residential utility, this common reduction is equivalent to multiplying b_{ig} by $e^{\kappa \Delta t_i}$ and can therefore be absorbed into the recovered fundamental amenity improvement.

[8]I use neighborhood-level estimates because neighborhoods are the spatial units of the model. These estimates are also consistent with the block-level results.

account for the cost of implementing the program. The first two columns show the effects of having implemented only Phase I, under both closed and open city. The last two columns present the results if the program would have been completed.

Table 9: Welfare and spatial effects of PRIMED

	Phase I only		Full program: Phase I + II	
	Closed city	Open city	Closed city	Open city
Panel A. Welfare gains				
All residents (%)	1.96	0.66	6.09	2.08
Low skill (%)	2.15	0.71	6.66	2.22
High skill (%)	1.23	0.48	3.86	1.54
Panel B. Treated neighborhoods				
Neighborhoods reported	11	11	34	34
Additional residents	12,612	13,573	37,776	50,012
Population gain vs. no PRIMED (%)	21.93	24.00	18.97	25.55
Gain as share of 1993 population (%)	51.27	55.18	45.12	59.73
Residential rents (%)	4.98	5.35	4.45	5.63
Panel C. Untreated neighborhoods				
Neighborhoods	237	237	214	214
Population (%)	-0.96	0.92	-3.22	2.71
Residential rents (%)	-0.13	0.25	-0.47	0.78
Panel D. Medellín as a whole				
Residential neighborhoods in model	248	248	248	248
Population (%)	0.00	1.89	0.00	6.03
Labor income — low skill (%)	0.02	-0.23	0.06	-0.69
Labor income — high skill (%)	0.02	0.14	0.05	0.40
Residential rents (%)	0.12	0.51	0.27	1.51

Notes: All effects compare the indicated program with no PRIMED under the same city closure. Welfare is conditional Medellín utility before program costs. The baseline matches observed classified residents aged 25+ and relative cadastral prices. The full intervention covers 34 neighborhoods: 11 Phase I and 23 Phase II units classified as Phase II in the historical census analysis with blocks matched across 1993, 2005 and 2018. Rent changes are aggregated within the area indicated in each panel; nominal labor income changes are citywide. Both use geometric averages with fixed 2018 resident weights, calculated separately by skill group for income.

I find that only implementing Phase I infrastructure reforms led to positive city wide welfare gains in both closed and open city. Most welfare gains are captured by low skill residents, as they constitute the majority of the residents of these slums. By construction, the open-city model matches the population DID estimate. The model suggest that rents increased by nearly 5% in treated areas a consequence of these population changes, although the aggregate effects in the city are relatively modest. In the open city, welfare gains are lower because migration into the city increases demand for housing and labor supply, which pressures rents up by 5,4% in upgraded slums and 0.51% in all Medellin. There is also a small decrease in low skill income as there is higher labor supply. Figure 8 shows the spatial distribution of the changes in population, residential rents and amenities.

Table 10 examines the mechanisms behind these welfare gains. In the open city, the reduction in risk disaster largely contributes to this welfare gain, capturing about 49.4% of the gains. The improvement in local amenities explains 28.6%, whereas the increase in rents and fall in income due to the increase of residents in these areas is respectively-27.1% and -23.8%.

Table 10: Contributions to welfare gains from PRIMED Phase I

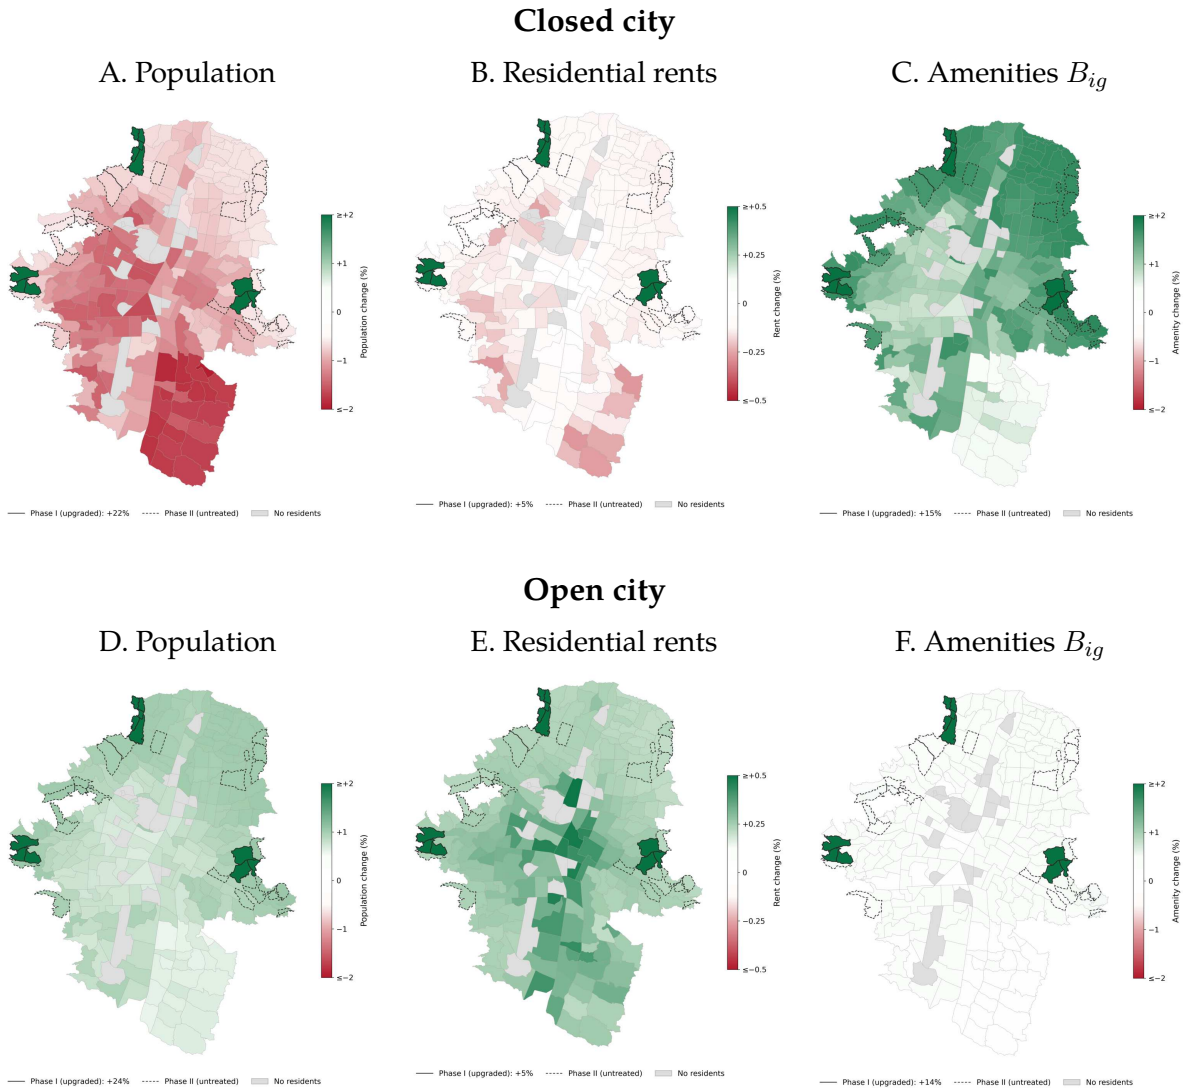
	Closed city		Open city	
Channel	pp	Share (%)	pp	Share (%)
Fundamental amenities	0.189	9.7	0.189	28.6
Disaster risk	0.328	16.7	0.326	49.4
High-skill share	1.465	74.7	0.481	72.9
Income access	0.018	0.9	−0.157	−23.8
Residential rent	−0.038	−2.0	−0.179	−27.1
Total welfare gain (%)	1.96	100.0	0.66	100.0

The welfare effects of having implemented the full program are larger than those from only implementing Phase I. They also indicate that low skill residents benefit more from these investments. Disaster protection accounts for 17.8% of welfare gains in the closed city and 50.9% in the open city (Appendix Table 36). Under this counterfactual, the results also indicate larger population increases in treated areas but also in the rest of the city. This increase leads to even higher increases in residential rents (1.51% in open city) and a slight decrease in low skill's income. Figure 9 shows the spatial distribution of the changes in population, residential rents and amenities.

Appendix Table 35 reports the Phase I counterfactual under alternative values of disaster disutility and the floor-space supply parameter. Across these specifications, PRIMED continues to generate welfare gains for Medellín in both the closed and open city. In the closed city, the welfare gain ranges from 1.16% when $\gamma = -0.1$ to 4.20% when $\gamma = -1$, compared with 1.96% under the baseline calibration. As expected, the contribution of disaster protection increases with $|\gamma|$, as residents experience greater disutility from disaster exposure.

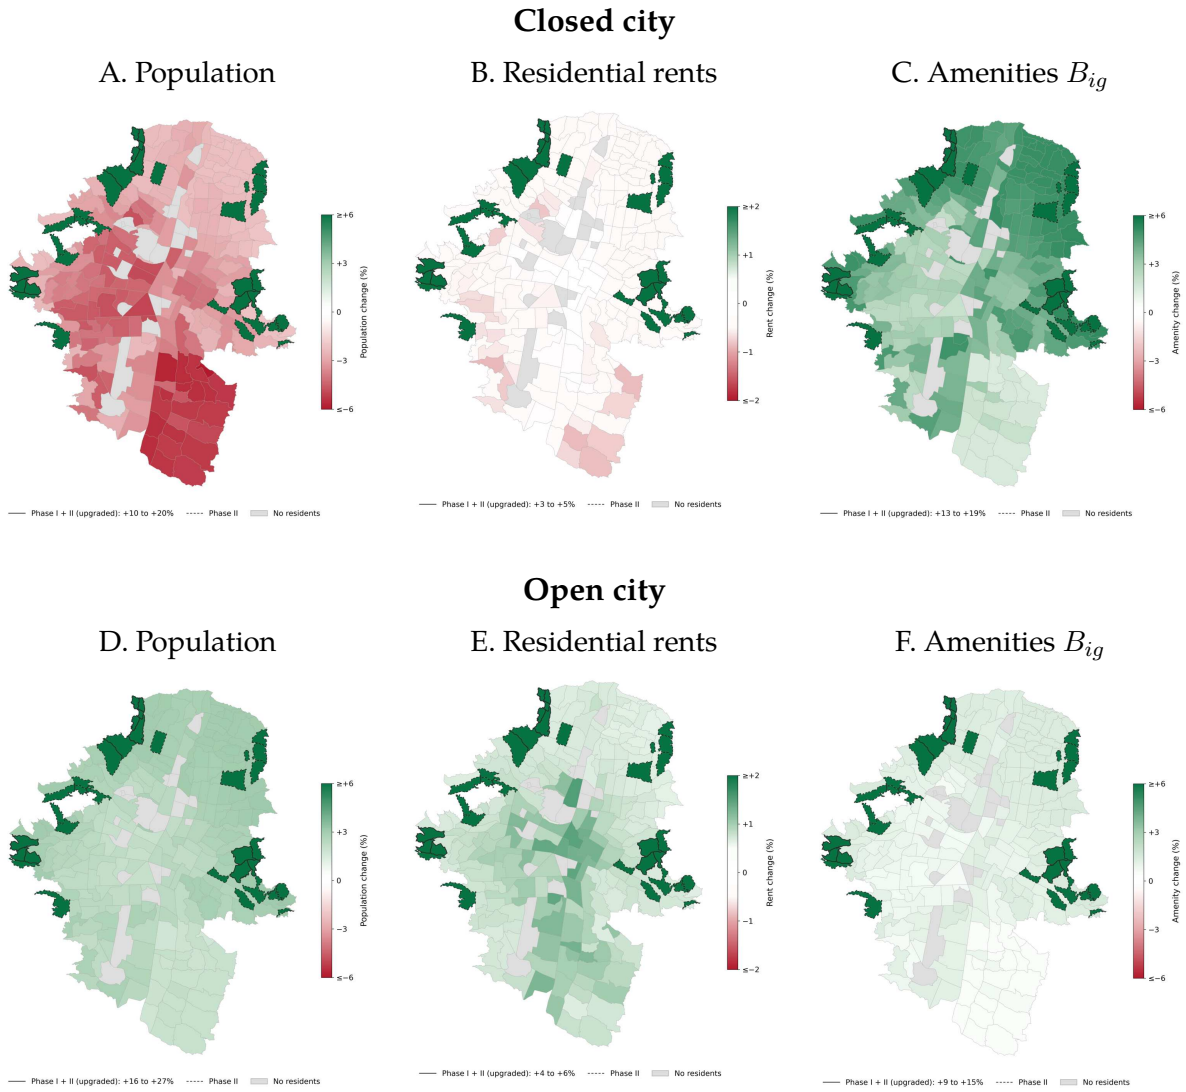
I also assess sensitivity to μ using housing-supply estimates for Rio de Janeiro as a benchmark. Guedes, Iachan, and Sant'Anna (2023) estimate housing supply in the presence of informality and report an inverse supply elasticity of 0.42 for Rio de Janeiro, equivalent to a supply elasticity of approximately 2.38. I set $\mu = 0.7$, which implies a floor-space supply elasticity of $\mu/(1-\mu) = 2.33$, compared with 3 under the baseline calibration. A lower supply elasticity implies larger increases in rents, which reduces welfare. The welfare gains remain close to the baseline results: 1.93% in the closed city and 0.63% in the open city, compared with 1.96% and 0.66%, respectively.

Figure 8: Effects of PRIMED Phase I



Notes: Percentage changes relative to no PRIMED under the indicated city closure. Colour scales are capped at the range of untreated neighborhoods; upgraded neighborhoods exceed the caps and their change is reported in the legend. Solid outlines denote Phase I neighborhoods; dashed outlines denote Phase II. Gray areas have no residents.

Figure 9: Effects of the full PRIMED program: Phase I + II



Notes: Percentage changes relative to no PRIMED under the indicated city closure. Colour scales are capped at the range of untreated neighborhoods; upgraded neighborhoods exceed the caps and their change is reported in the legend. Solid outlines denote Phase I neighborhoods; dashed outlines denote Phase II. Gray areas have no residents.

7.4 Climate change scenarios

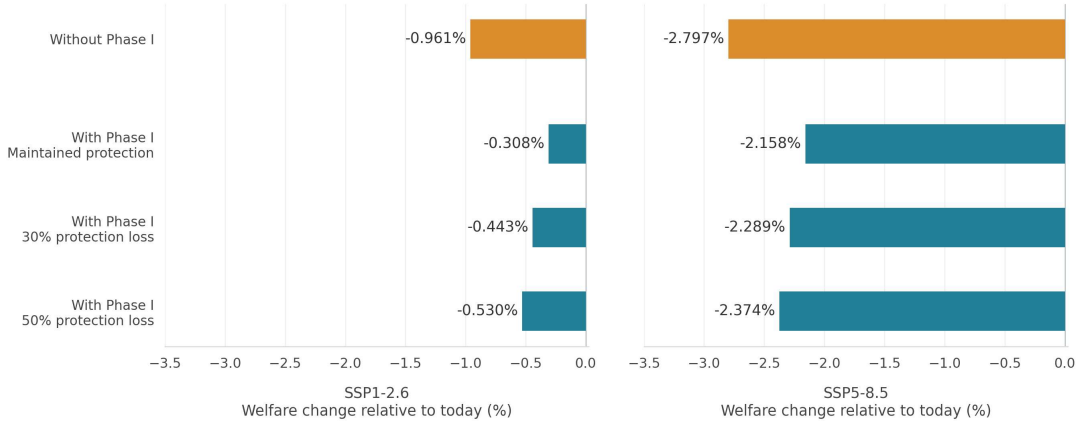
I use scientific projected precipitation data from the Colombian Meteorological Agency to project natural disasters under four different climate scenarios as indicated by the IPCC (IDEAM, 2024). These scenarios correspond to different trajectories of global warming under different emissions paths, and which are exogenous to Medellin as the city is a very small economy relative to the world. I project disasters across the city using Empirical Bayes method and historical data on natural disasters after 1993 and following the Delta Method from IPCC (2021). More details on this calculation can be found in Appendix D.7.

By 2081, the original PRIMED infrastructure would be more than eight decades old. For comparison, Colombia's design standard specifies a minimum service life of 75 years for most permanent retaining walls (INVÍAS & AIS, 2014). Given this long horizon, I examine how the welfare results change if the infrastructure loses part of its protective effect.

Figure 10 compares welfare with and without Phase I upgrading in 2081-2100, relative to today's equilibrium. I consider an open city, allowing migration as a form of adaptation to increasing disaster risk. The results indicate that in scenarios with low (SSP1-2.6) and high precipitation (SSP5-8.5), higher disasters will lead to welfare losses for the city. Phase I upgrading reduces these future welfare losses under both rainfall scenarios, but substantial losses remain even with maintained protection. Under the higher-rainfall scenario, welfare is 2.80% below today's level without Phase I, compared with 2.16% below today's level when its protection is maintained. When I assume a 30% or 50% loss of protection, welfare losses increase to 2.29% and 2.37%, respectively, narrowing the advantage of upgrading.

These results suggest that even if upgraded slums are coping relatively well with natural disasters today, as shown by the empirical results and PRIMED counterfactuals, by the end of century higher exposure can make upgrading insufficient to prevent future welfare losses.

Figure 10: Citywide welfare with and without Phase I upgrading, 2081-2100



Note: Open city, 2081-2100. Disasters are projected using Empirical Bayes, using a 2021-2040 delta-change following IPCC (2021). The 30% and 50% protection-loss cases are sensitivity assumptions.

8 Conclusion

This paper shows that slum upgrading can reduce exposure to climate events. The investments increased urban growth in lower-risk areas within upgraded neighborhoods, changing where new development took place. The benefits were not limited to the upgraded neighborhoods: the city as a whole gained from implementing this program. The main mechanism is an improvement in local amenities via infrastructure upgrading. When climate change is introduced in this analysis, I find that more frequent disasters lead to welfare losses, particularly when depreciation of the infrastructure is taken into account. This poses the trade-off between protecting today risky areas and locking them into future higher risk.

References

Ahlfeldt, G. M., Redding, S. J., Sturm, D. M., & Wolf, N. (2015). The economics of density: Evidence from the berlin wall. *Econometrica*, *83*(6), 2127–2189.

Aron-Dine, S. (2025). *Rebuild or Relocate? Recovery after Natural Disasters* (Tech. Rep.).

Balboni, C. (2025, 1). In harm's way? infrastructure investments and the persistence of coastal cities. *American Economic Review*, *115*, 77-116.

Betancur, J. J. (2007). Regularization of Informal Settlements in Medellin, Colombia. *Global Urban Development*, *3*(3), 1–3. Retrieved from https://www.globalurban.org/GUDMag07Vol3Iss1/Betancur.htm

Boustan, L. P., Kahn, M. E., Rhode, P. W., & Yanguas, M. L. (2020). The effect of natural disasters on economic activity in US counties: A century of data. *Journal of Urban Economics*, *118*.

Bryan, G., & Morten, M. (2019). The aggregate productivity effects of internal migration: Evidence from Indonesia. *Journal of Political Economy*, *127*(5), 2229–2268. doi: 10.1086/701810

Card, D. (2009). Immigration and inequality. *American Economic Review*, *99*(2), 1–21.

Combes, P.-P., Duranton, G., & Gobillon, L. (2021). The production function for housing: Evidence from france. *Journal of Political Economy*, *129*.

Deryugina, T. (2017). The fiscal cost of hurricanes: Disaster aid versus social insurance. *American Economic Journal: Economic Policy*, *9*(3), 168–198.

Diamond, R. (2016, 3). *The determinants and welfare implications of us workers' diverging location choices by skill: 1980-2000* (Vol. 106). American Economic Association.

FAA. (2022, March). *Community response to noise.* Federal Aviation Administration. Retrieved from https://www.faa.gov/regulations_policies/policy_guidance/noise/community (Updated March 29, 2022)

Funk, C., Peterson, P., Landsfeld, M., Pedreros, D., Verdin, J., Shukla, S., ... Michaelsen, J. (2015, 12). The climate hazards infrared precipitation with stations - a new environmental record for monitoring extremes. *Scientific Data*, *2*.

Galiani, S., Gertler, P. J., Undurraga, R., Cooper, R., Martínez, S., & Ross, A. (2017). Shelter from the storm: Upgrading housing infrastructure in Latin American slums. *Journal of Urban Economics*, *98*, 187–213.

Gertler, P., Gonzalez-Navarro, M., Undurraga, R., & Urrego, J. A. (2024). *In-situ upgrading or population relocation? direct impacts and spatial spillovers of slum renewal policies.*

Gómez, M. d. M., & Zárate, R. D. (2024). *Using quantitative spatial-equilibrium models to inform urban policymaking* (Project Report No. ZAF-21193). International Growth Centre. Retrieved from https://www.theigc.org/publications/using-quantitative-spatial-equilibrium-models-inform-urban-policymaking

Guedes, R., Iachan, F. S., & Sant'Anna, M. (2023). Housing supply in the presence of informality. *Regional Science and Urban Economics*, *99*, 103875. doi: 10.1016/j.regsciurbeco.2023.103875

Habitat, U. (1998, 7). *Best practices and local leadership programme (blp).* Retrieved from https://habitat.aq.upm.es/dubai/98/bp002.html

Habitat, U. (2000, 7). *Best practices and local leadership programme (blp).* Retrieved from https://habitat.aq.upm.es/dubai/00/bp540.html

Harari, M., & Wong, M. (2024). *Slum Upgrading and Long-run Urban Development: Evidence from Indonesia.*

He, C., Shi, P., Xie, D., & Zhao, Y. (2010). Improving the normalized difference built-up index to map urban built-up areas using a semiautomatic segmentation approach. *Remote Sensing Letters*, *1*(4), 213–221.

Hsiao, A. (2025). *Sea Level Rise and Urban Adaptation in Jakarta.*

IDEAM. (2024, December). *El ideam presenta nuevas proyecciones climáticas para colombia:*

Escenarios de cambio climático al 2100. Retrieved from https://www.ideam.gov.co/sala-de-prensa/noticia/el-ideam-presenta-nuevas-proyecciones-climaticas-para-colombia-escenarios-de-cambio-climatico-al

INVÍAS, & AIS. (2014). Norma colombiana de diseño de puentes: CCP-14 [Computer software manual]. Bogotá, Colombia. (Artículo 11.5.1, p. 11–6)

IPCC. (2021). *Climate change 2021: The physical science basis. contribution of working group i to the sixth assessment report of the intergovernmental panel on climate change*. Cambridge University Press. Retrieved from https://www.ipcc.ch/report/ar6/wg1/ doi: 10.1017/9781009157896

IPCC. (2023). *Ar6 synthesis report: Climate change 2023.* Panmao Zhai.

Kahn, M., & Ouazad, A. (2025). *Time consistent infrastructure investments: Optimal flood protection policies in spatial equilibrium* (Vol. 85).

Khanna, G., Medina, C., Nyshadham, A., Ramos-Menchelli, D., Tamayo, J., & Tiew, A. (2025, Dec). *Spatial mobility, economic opportunity, and crime* (Borradores de Economia No. 1334). Banco de la Republica de Colombia.

Li, D., Sun, L., Feng, K., Zhang, N., Yu, Y., Zhao, D., & Zhou, Y. (2025, 7). Disproportionate flood exposure for slum populations of the global south. *Nature Cities*, *2*, 626-638.

Marshall, R. J. (1991). Mapping disease and mortality rates using empirical bayes estimators. *Journal of the Royal Statistical Society Series C: Applied Statistics*, *40*(2), 283–294. doi: 10.2307/2347593

Marx, B., Stoker, T., & Suri, T. (2013). The economics of slums in the developing world. *Journal of Economic Perspectives*, *27*(4), 187–210.

Nakamura, E., & Steinsson, J. (2018). Identification in macroeconomics. *Journal of Economic Perspectives*, *32*(3), 59–86. doi: 10.1257/jep.32.3.59

NASA, & METI. (2019). *ASTER Global Digital Elevation Model (GDEM)*. (Version 3. Distributed by NASA EOSDIS Land Processes DAAC.)

Ozturk, U., Bozzolan, E., Holcombe, E. A., Shukla, R., Pianosi, F., & Wagener, T. (2022). How climate change and unplanned urban sprawl bring more landslides. *Nature*, *608*.

Santos Silva, J. M. C., & Tenreyro, S. (2006). The log of gravity. *The Review of Economics and Statistics*, *88*(4), 641–658. doi: 10.1162/rest.88.4.641

Tsivanidis, N. (2026, 2). Evaluating the impact of urban transit infrastructure: Evidence from bogotá's transmilenio. *American Economic Review*, *116*, 418-463.

UN-Habitat. (2021). *Housing, slums and informal settlements | urban indicators database.*

Zarate, D. (2024). *Spatial misallocation, informality, and transit improvements: Evidence from mexico city* (Tech. Rep.).

Zhang, X., Zhao, T., Xu, H., Liu, W., Wang, J., Chen, X., & Liu, L. (2024). Glc_fcs30d: the first global 30 m land-cover dynamics monitoring product with a fine classification system for the period from 1985 to 2022 generated using dense-time-series

landsat imagery and the continuous change-detection method. *Earth System Science Data*, *16*(3), 1353–1381.

Appendix

A Data **43**
- A.1 Precipitation . . . 43
  - A.1.1 CHIRPS (1981-2023) . . . 43
- A.2 Natural Disasters . . . 45
- A.3 Topographic characteristics . . . 46
- A.4 Landslide Susceptibility . . . 46
- A.5 Impervious Surface . . . 47
- A.6 Drainage Infrastructure . . . 49
- A.7 Housing Expenditure . . . 51
- A.8 Employment and high-skill bill shares . . . 51

B Rainfall Event-Study Robustness **53**

C Urban Outcomes Robustness **56**
- C.1 Population Size and Skill Composition . . . 56
  - C.1.1 Neighborhood-Level Population . . . 56
  - C.1.2 Alternative Control Group: Level II . . . 56
  - C.1.3 Conley-HAC Standard Errors . . . 58
- C.2 Density and Impervious Surface . . . 59
  - C.2.1 Alternative Control Group: Level II . . . 59
  - C.2.2 Unit Fixed Effects . . . 61
  - C.2.3 Conley-HAC Standard Errors . . . 61
  - C.2.4 Excluding Nearby Controls . . . 63
  - C.2.5 Potential Spillovers to Nearby Untreated Blocks . . . 65
- C.3 Housing Characteristics . . . 67
  - C.3.1 Alternative Control Group: Level II . . . 67
  - C.3.2 Conley-HAC Standard Errors . . . 68
- C.4 Additional Population Characteristics . . . 69
  - C.4.1 Population by Age and Sex . . . 69
  - C.4.2 Literacy and Migration Background . . . 70
- C.5 Changes in Population by Risk Category . . . 71

D Spatial Model Appendix **73**
- D.1 Commuting Elasticity . . . 73
- D.2 Supplemental details of calibration . . . 80
- D.3 Phase I parameter sensitivity . . . 81
- D.4 Full-program welfare decomposition . . . 82
- D.5 Model validation . . . 83
- D.6 Empirical composite index of amenities . . . 84
- D.7 Climate projections and counterfactual construction . . . 86

A Data

A.1 Precipitation

A.1.1 CHIRPS (1981-2023)

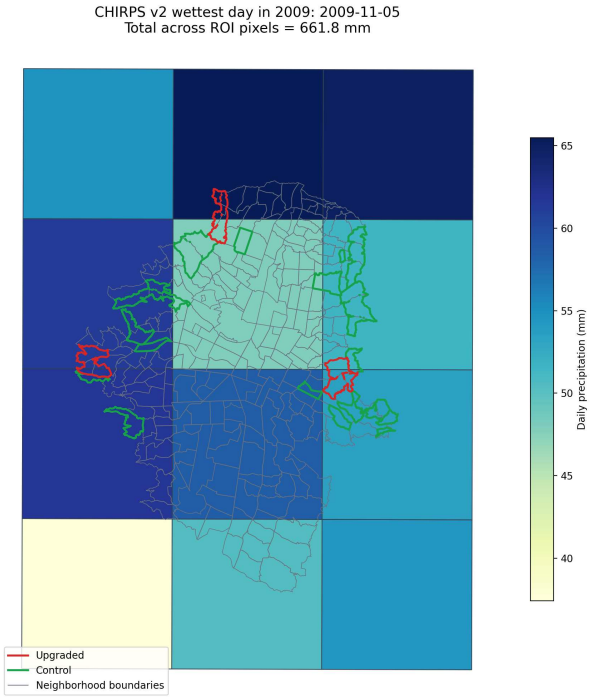
I use CHIRPS v2 (Climate Hazards Group InfraRed Precipitation with Station data) accessed via the Google Earth Engine public data catalog. CHIRPS is a quasi-global gridded rainfall dataset covering land areas from $50°$S to $50°$N. Available from 1981 to the near-present at daily, pentadal, and monthly frequencies, it features a high native spatial resolution of $0.05° \times 0.05°$, approximately 5.5 km at the equator (Funk et al., 2015). CHIRPS blends three core sources of information: a high-resolution long-term precipitation climatology (CHPclim), satellite-based thermal infrared (TIR) cloud-top temperature anomalies used to estimate rainfall patterns, and in-situ rain-gauge station observations. The satellite component recovers spatial rainfall distribution in regions with sparse monitoring networks, while the station data calibrate and anchor these estimates to ground reality. The resulting data variable represents daily rainfall depth in millimeters (mm) per grid cell.

I assign CHIRPS daily precipitation to neighborhood polygons using the CHIRPS grid cells that intersect each neighborhood. Because precipitation is an intensive variable, I do not allocate or split rainfall totals across neighborhoods. Instead, each neighborhood receives the rainfall value of the relevant CHIRPS cell or cells. For the main specification, I use a dominant-cell assignment. I reconstruct the CHIRPS grid cells from the daily pixel centroids, intersect them with the neighborhood polygons, and compute overlap areas in an equal-area projection. If a neighborhood intersects only one CHIRPS cell, it receives that cell's daily precipitation. If it intersects multiple cells, it receives the precipitation value from the CHIRPS cell with the largest area overlap. As a robustness measure, I also construct an overlap-weighted full-cell version. In that version, neighborhoods intersecting multiple CHIRPS cells receive the weighted average of the intersecting cell precipitation values, where weights are the share of the neighborhood area falling in each cell.

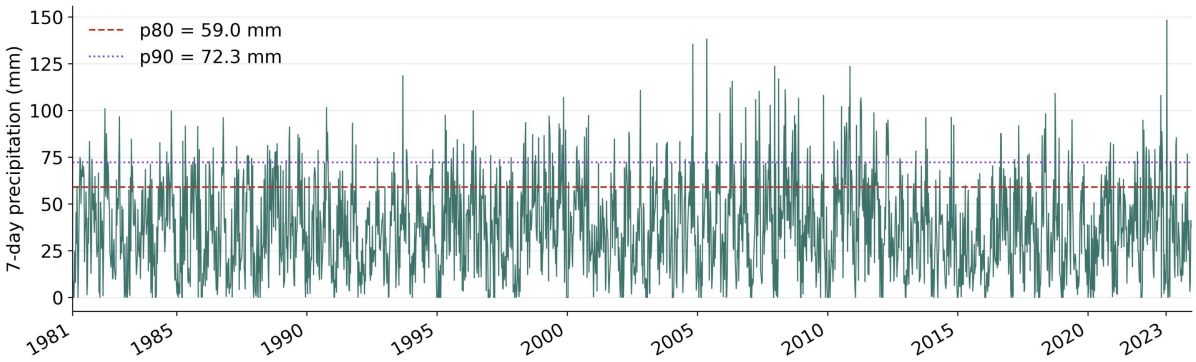
To compute heavy-rain events, I aggregate daily precipitation over a window of 7 days and classify it as heavy-rain if it is greater than the 80th percentile of the 7-day distribution of precipitation between 1981-2023, which is equivalent to 59.0mm. As robustness, I also construct heavy-rain episodes for three and five days and alternative thresholds of the distribution (p85,p90).

Figure 11: Precipitation data from CHIRPS v2

Panel A. Spatial coverage



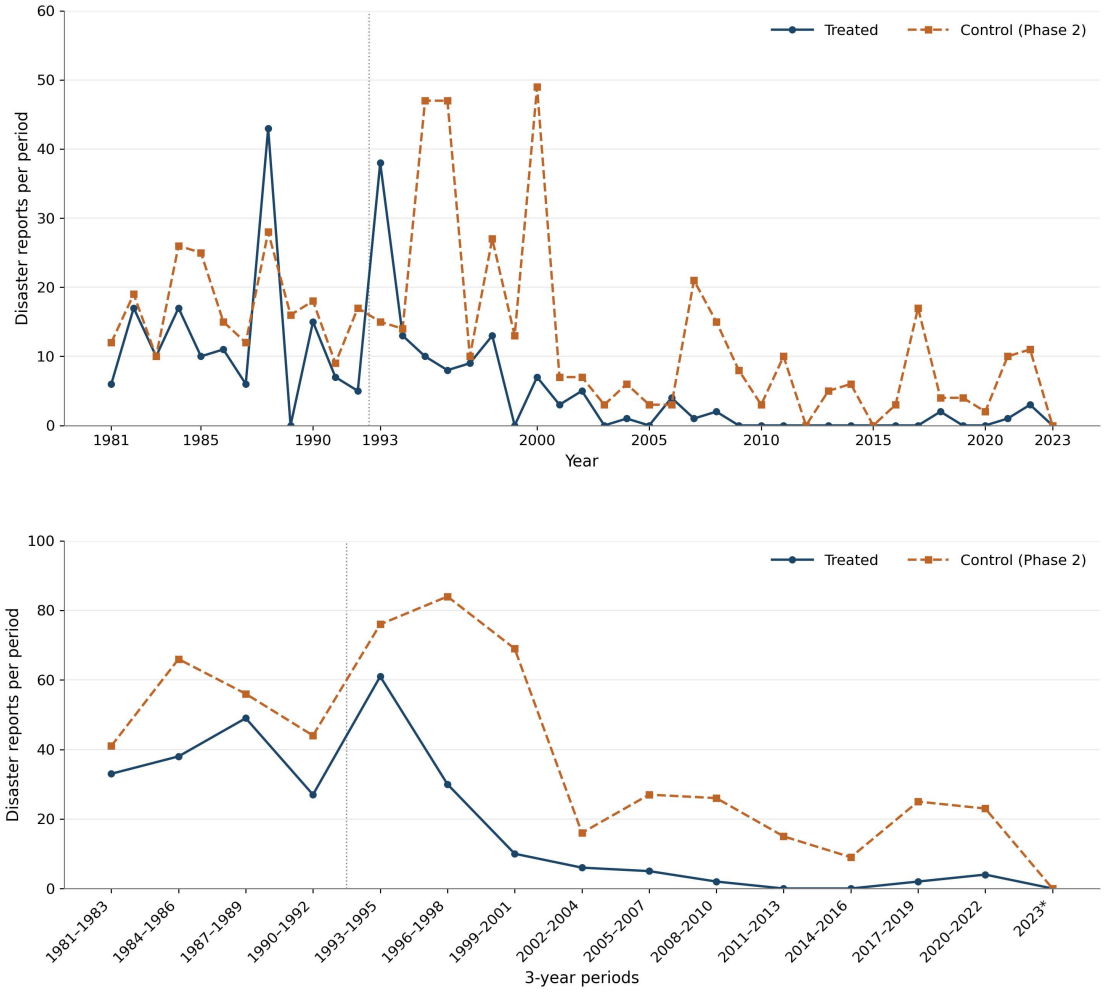
Panel B. 7-day precipitation



Note: Precipitation is taken from CHIRPS V2. Precipitation is assigned to neighborhoods based on the dominant-cell principle. I aggregate precipitation over windows of seven days.

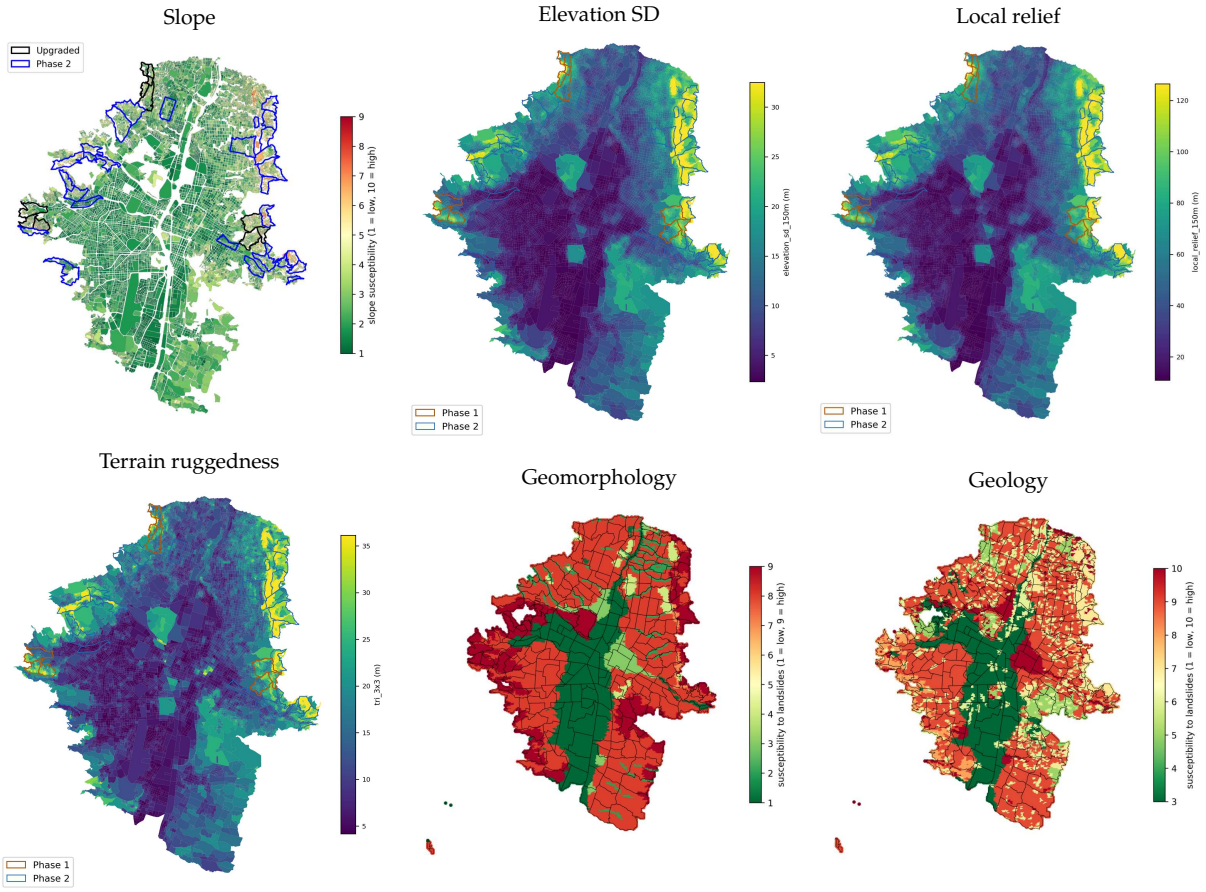
A.2 Natural Disasters

Figure 12: Natural disasters in upgraded and Phase II neighborhoods (1981–2023)



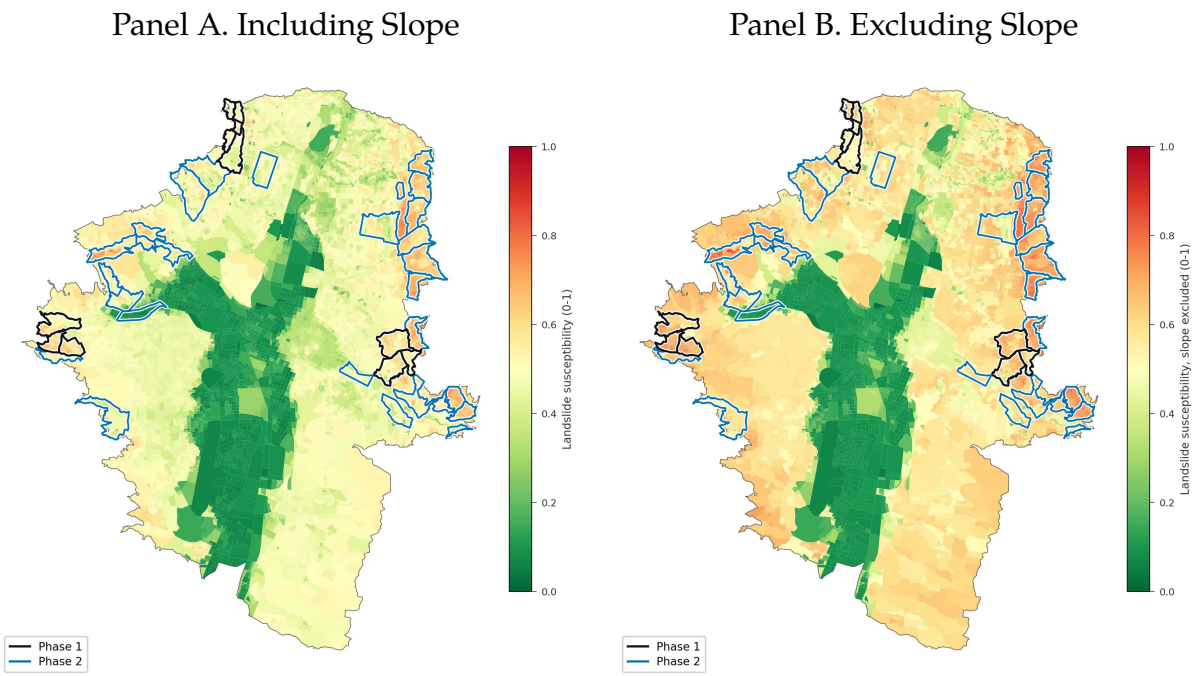
Note: Natural disasters between 1981–1998 are taken from DesInventar (United Nations) and PRIMED (historical archives of disasters in upgraded areas). Natural disasters between 1998–2023 are taken from UNGRD (reports on emergencies).

A.3 Topographic characteristics



A.4 Landslide Susceptibility

Figure 13: Landslide Susceptibility Indices

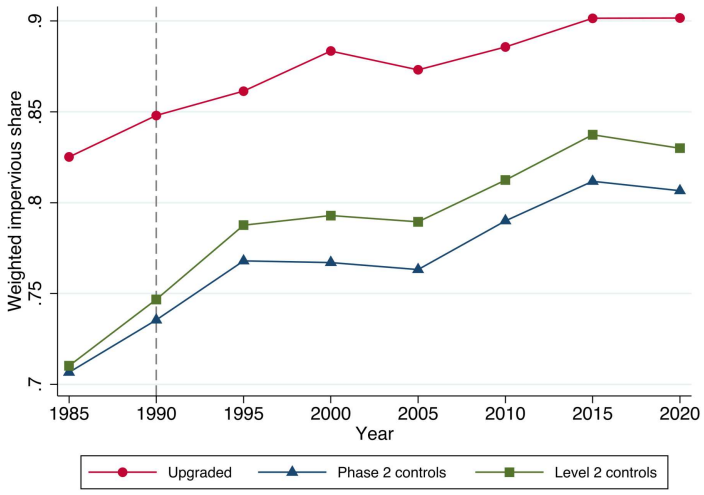


Note: Panel A presents the block-level index based on geology, geomorphology, slope, and terrain ruggedness. Panel B excludes slope. Both indices are scaled from zero to one. **Source:** Author's calculations using 2005 census-block geometries.

A.5 Impervious Surface

I use data from the catalog developed by Zhang et al. (2024) to measure impervious land cover between 1985 and 2020. I define an indicator equal to one when a 30 m $\times$ 30 m cell is classified as "impervious surface" and zero otherwise. This category captures artificial land covers such as building rooftops, roads, and industrial facilities, while excluding urban green spaces and water bodies. The measure describes horizontal land coverage and contains no information on building height or floor space. The data is available at https://zenodo.org/records/8239305

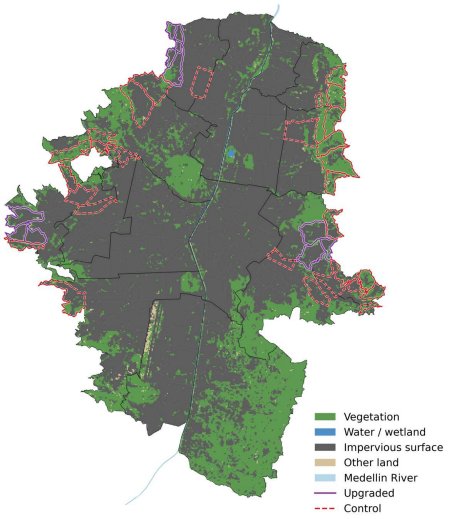
Figure 14: Impervious Surface Over Time by Control Group



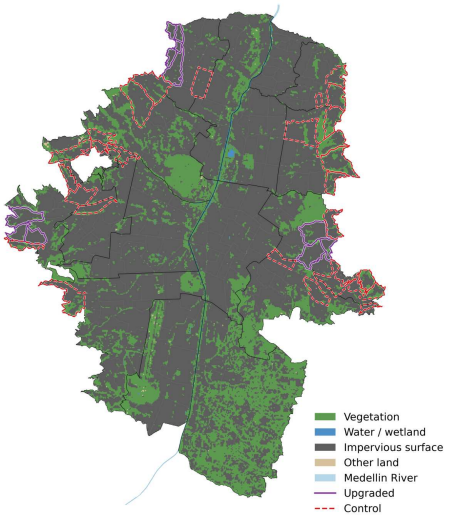
Note: Figure shows the evolution of the impervious-surface indicator which is equal to one when a 30-meter year-cell is classified as impervious in the Global Land Cover data with fine classification (Zhang et al., 2024). The series is weighted by log of 1993 neighborhood population.

Figure 15: Land Cover in Medellin

Panel A. 1985



Panel B. 2020



Source: Data at the 30-meter resolution from (Zhang et al., 2024)

A.6 Drainage Infrastructure

I use georeferenced drainage and sewerage data from Medellín's utility company, Empresas Públicas de Medellín (EPM), available through its geoportal. These data describe one component of the broader protective works, which also included slope stabilization and retaining walls.

Estimated full-flow capacity (L/s) uses EPM pipe diameter and slope in Manning's equation with a common $n = 0.013$ (EPM, 2013). Large pipes have diameter ≥900 mm.

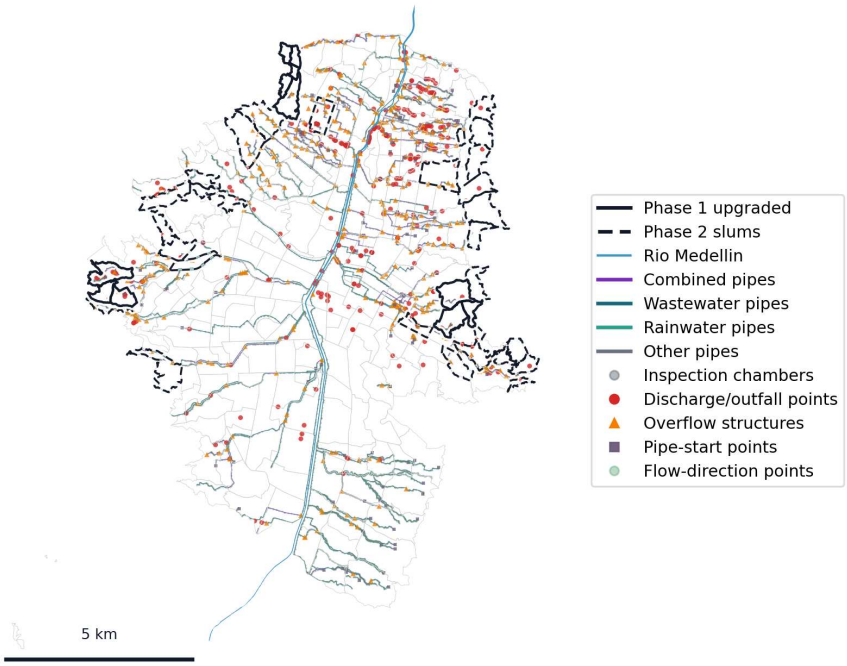
Table 11: Mapped pipe infrastructure by treatment status

	Phase I upgraded	Phase II comparison
Mean length of pipes ≥900 mm (m)	98.2	6.8
Mean share of pipe length ≥900 mm (%)	15.52	0.48
Mean estimated pipe capacity (L/s)	2,220	878
Median neighborhood capacity (L/s)	1,244	584
Mean total pipe density (km/km^2)	1.62	2.20
Mean inspection-chamber density (count/km^2)	79.5	82.1

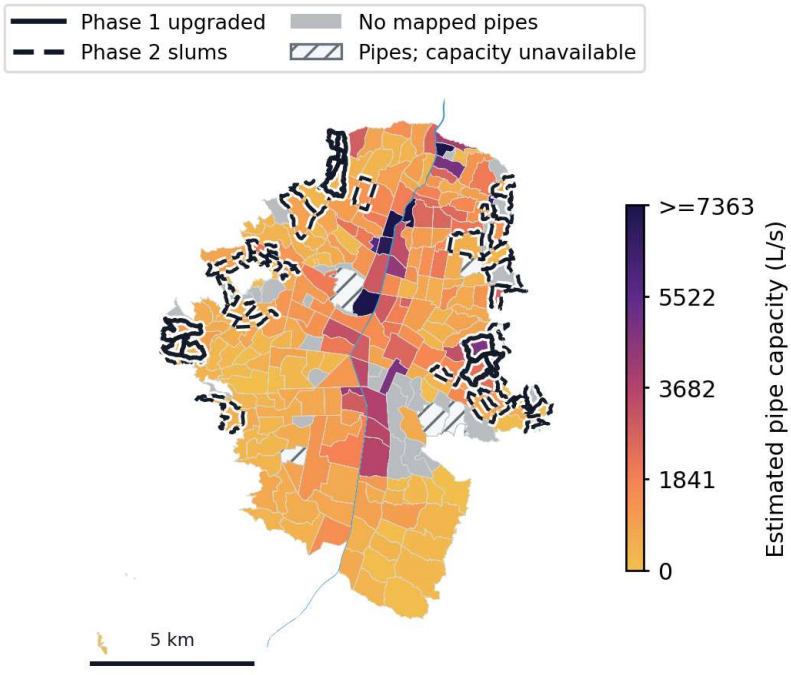
Note: Within each neighborhood, estimated pipe capacity is weighted by pipe length. Capacity rows use 9 Phase I and 23 Phase II neighborhoods; other rows use 11 and 28, respectively.

Figure 16: Drainage infrastructure in Medellín

Panel A. Mapped drainage and sewerage network



Panel B. Estimated neighborhood pipe capacity (L/s)



Note: Data from EPM's *Tubería Alcantarillado* main-network layer and related point layers. Capacity in B matches Table 11. Solid outlines indicate Phase I; dashed outlines indicate Phase II. Gray means no pipes mapped in the EPM layer; hatching marks mapped pipes with unestimated capacity. Capacity colors are capped at the citywide 98th percentile.

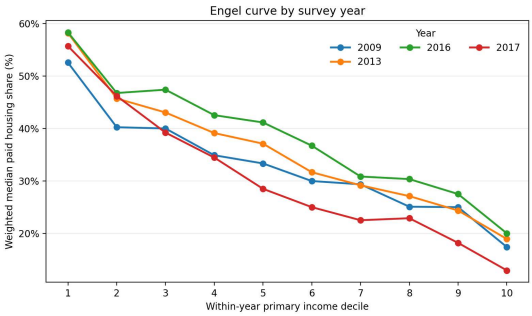
A.7 Housing Expenditure

I obtained data on housing expenditure and income from the Encuesta de Calidad de Vida available at https://www.medellin.gov.co/es/centro-documental/encuesta-calidad-de-vida/. This data is available since 2004 to present. However, I use data until 2017 because this was the last year for which the "neighborhood" question is publicly available and that is closer to the 2018 census population. To compute the average share of income allocated to housing, I use data from the years 2009, 2013, 2016 and 2017. These years have more reliable data in terms of income, than other years in between.

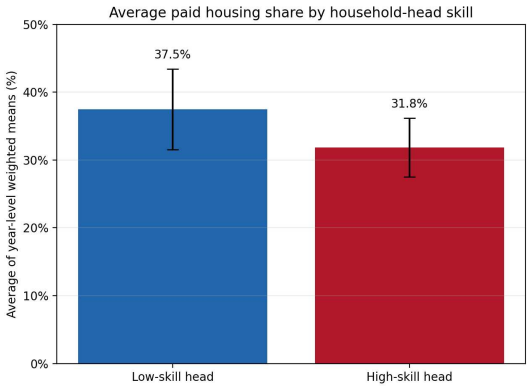
My main quantitative results use the average share of expenditure on housing.

Figure 17: Housing Expenditure

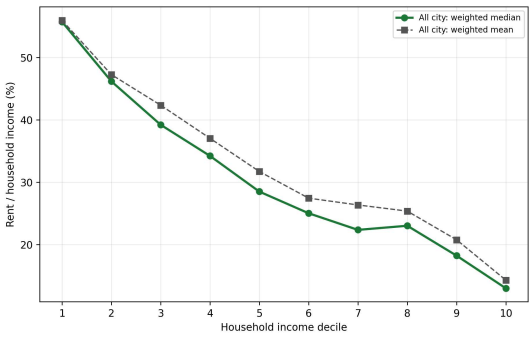
Panel A. 2009–2017, Income Deciles



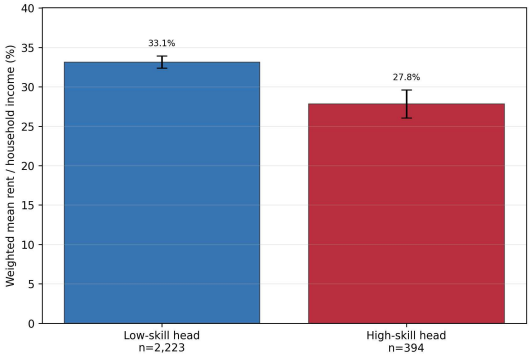
Panel B. 2009–2017, Head Skill



Panel C. 2017, Income Deciles



Panel D. 2017, Head Skill



Source: Encuesta de Calidad de Vida.

A.8 Employment and high-skill bill shares

High-skill wage-bill shares by sector are constructed from the 2015-2018 Gran Encuesta Integrada de Hogares (GEIH) microdata produced by DANE. The sample is restricted to employed individuals in Medellin with non-missing labor income, education status, and industry classification. For each industry and year, total wage bill is computed as the sum of labor income multiplied by the GEIH expansion factor, and the high-skill

wage bill is computed analogously for workers with completed college education. The sector-level high-skill bill share is the average across 2015-2018 of the yearly ratio of high-skill wage bill to total wage bill. Industries are classified using DANE's Colombian adaptation of the International Standard Industrial Classification of All Economic Activities (ISIC), Revision 4. The resulting shares are matched to the employment-sector data over the 20 industry sections used in the model. The employment data by location comes from confidential records from the Chamber of Commerce of Medellin. The data is at the establishment level and it contains the number of employees for all registered firms in the city, regardless of their employment size or contributions to social security.

To calibrate the sector-specific skill intensity weights, I assume that commercial floorspace accounts for 20% of sectoral expenditure, leaving 80% for labor ($\alpha = 0.8$). This assumption is standard in the literature. For each sector s, the low-skill and high-skill labor shares are then constructed as

$$\alpha_{sL} = 0.8\left(1 - \text{HighSkillBillShare}_s\right), \quad \text{and} \quad \alpha_{sH} = 0.8 \cdot \text{HighSkillBillShare}_s.$$

By construction, $\alpha_{sL} + \alpha_{sH} = 0.8$ in every sector. These weights are averaged using each location industry employment shares to obtain α_{gj}.

Table 12: Employment data and labor skill intensity share by industry sector

Sector name	Number of employees	High-Skill Bill share	α_{sL}	α_{sH}
Agriculture, forestry and fishing	41,699	29.6%	0.563	0.237
Mining and quarrying	9,556	61.9%	0.305	0.495
Manufacturing	167,015	24.0%	0.608	0.192
Electricity, gas, steam and air conditioning	3,786	74.4%	0.205	0.595
Water supply, sewerage, waste management	2,462	33.2%	0.534	0.266
Construction	100,830	20.8%	0.633	0.167
Wholesale and retail trade	153,779	18.6%	0.651	0.149
Transportation and storage	47,486	12.5%	0.700	0.100
Accommodation and food service	37,882	10.1%	0.719	0.081
Information and communication	49,678	62.6%	0.300	0.500
Financial and insurance activities	47,475	65.9%	0.273	0.527
Real estate activities	18,540	20.0%	0.640	0.160
Professional, scientific and technical activities	60,354	70.9%	0.233	0.567
Administrative and support service activities	175,108	26.8%	0.586	0.214
Public administration and defence	1,545	62.1%	0.303	0.497
Education	7,392	79.0%	0.168	0.632
Human health and social work	45,369	62.3%	0.302	0.498
Arts, entertainment and recreation	11,564	28.9%	0.569	0.231
Other service activities	5,119	21.2%	0.630	0.170
Household activities as employers	233	0.8%	0.793	0.007

Notes: High skill corresponds to college educated workers. Employment data comes from the Chamber of Commerce in Medellin and college bill-share are built using GEIH survey from DANE.

B Rainfall Event-Study Robustness

I assess the sensitivity of the heavy-rainfall event-study results to the definition of rainfall episodes, the choice of comparison slums, and spatial correlation in the errors. The main specification uses seven-day windows and rainfall above the 80th percentile. Table 13 summarizes the estimates at the time of heavy rain ($t = 0$), in the following window ($t = 1$), and their sum.

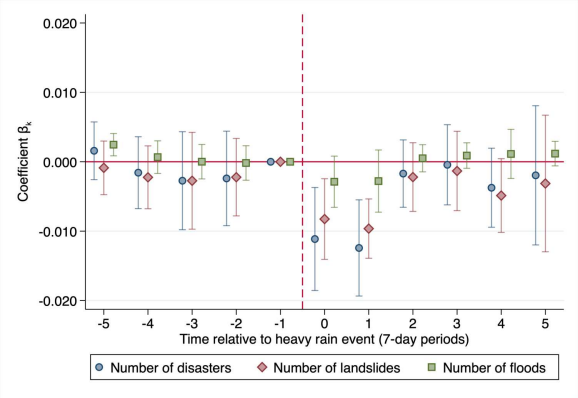
Table 13: Robustness of the heavy-rain event study

	Disasters			Landslides				
	$t=0$	$t=1$	$t=0+t=1$	$t=0$	$t=1$	$t=0+t=1$	Obs.	Slums
A. Heavy-rain definition								
7-day window, P80 (main)	−0.0111***	−0.0124***	−0.0236***	−0.0083***	−0.0096***	−0.0179***	83,028	37
	(0.0037)	(0.0034)	(0.0064)	(0.0029)	(0.0021)	(0.0046)		
7-day window, P85	−0.0116**	−0.0185***	−0.0302***	−0.0072*	−0.0136***	−0.0208***	83,028	37
	(0.0044)	(0.0052)	(0.0089)	(0.0036)	(0.0030)	(0.0060)		
3-day window, P80	−0.0046***	−0.0047**	−0.0092***	−0.0031**	−0.0034*	−0.0064**	193,695	37
	(0.0015)	(0.0019)	(0.0032)	(0.0012)	(0.0018)	(0.0029)		
5-day window, P80	−0.0099***	−0.0023	−0.0122***	−0.0071***	−0.0025	−0.0095**	116,217	37
	(0.0022)	(0.0027)	(0.0042)	(0.0022)	(0.0021)	(0.0037)		
B. Control group								
Drop the 4 Phase II controls adjacent to upgraded slums	−0.0094***	−0.0098***	−0.0192***	−0.0073***	−0.0084***	−0.0157***	74,052	33
	(0.0031)	(0.0025)	(0.0048)	(0.0026)	(0.0019)	(0.0041)		
Level 2 controls only	−0.0178***	−0.0173***	−0.0351***	−0.0127***	−0.0122***	−0.0249***	49,368	22
	(0.0051)	(0.0058)	(0.0106)	(0.0040)	(0.0031)	(0.0067)		
C. Spillovers to adjacent neighborhoods								
Adjacent controls vs. non-adjacent controls	0.0191	0.0183	0.0374	0.0157	0.0088*	0.0245	58,344	26
	(0.0141)	(0.0135)	(0.0270)	(0.0098)	(0.0050)	(0.0147)		
Bordering neighborhoods vs. non-adjacent controls	−0.0012	−0.0039	−0.0051	0.0020	−0.0013	0.0007	89,760	40
	(0.0050)	(0.0040)	(0.0082)	(0.0037)	(0.0019)	(0.0051)		
D. Spatial correlation (Conley standard errors)								
Conley, 1 km	−0.0111***	−0.0124***	−0.0236***	−0.0083**	−0.0096***	−0.0179***	83,028	37
	(0.0042)	(0.0037)	(0.0065)	(0.0037)	(0.0031)	(0.0058)		
Conley, 2 km	−0.0111***	−0.0124***	−0.0236***	−0.0083**	−0.0096***	−0.0179***	83,028	37
	(0.0042)	(0.0038)	(0.0066)	(0.0038)	(0.0031)	(0.0059)		

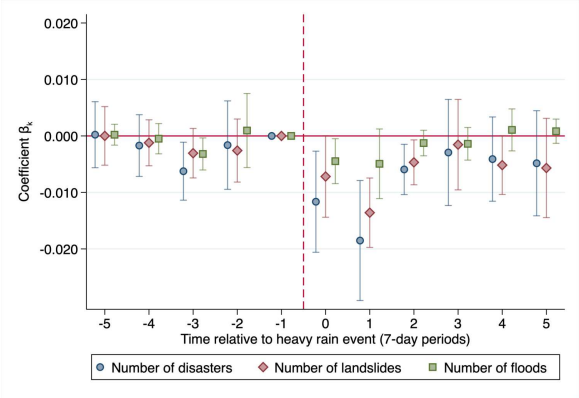
Notes: Coefficients on $Upgraded_s \times Post_t \times \mathbf{1}[k]$ at $k = 0$, $k = 1$, and their sum. Seven-day windows and the 80th-percentile rainfall threshold unless indicated. Panel C reports β_k from Equation (29), estimated on non-upgraded slums only. Panels A–C report standard errors clustered at the slum level. Panel D reports Conley standard errors with 1 and 2 km distance cutoffs and a Bartlett kernel. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure 18: Heavy-rain definition

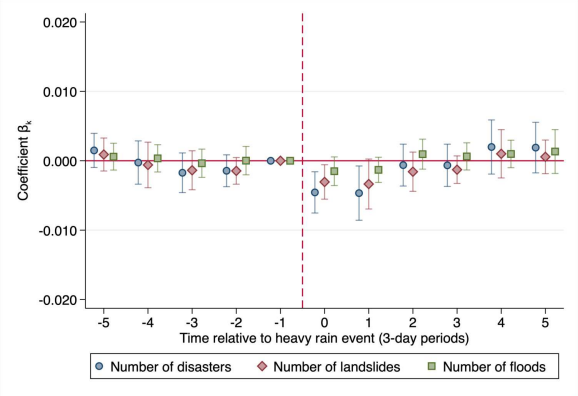
Panel A. 7-day, P80 (main)



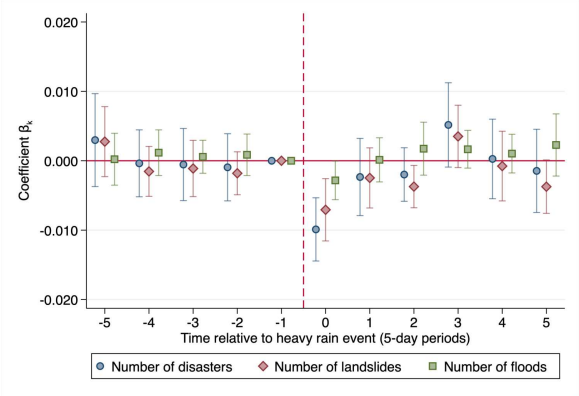
Panel B. 7-day, P85



Panel C. 3-day, P80



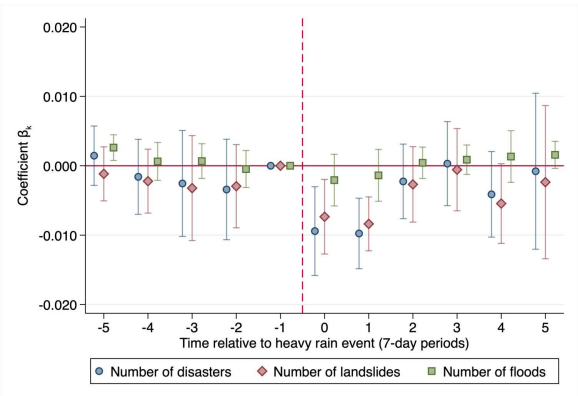
Panel D. 5-day, P80



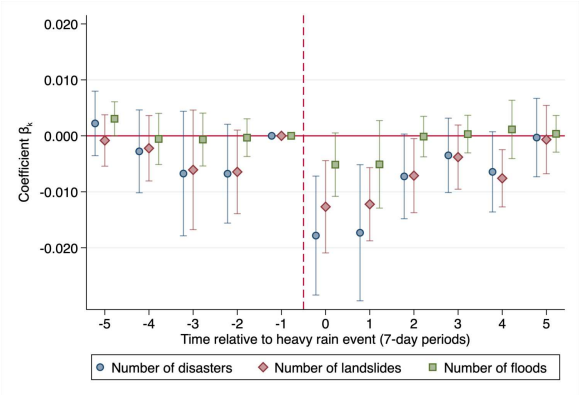
Notes: Window length and percentile threshold as indicated in each panel. 95% confidence intervals; standard errors clustered at the slum level.

Figure 19: Control group

Panel A. Dropping the 4 Phase II controls adjacent to upgraded slums



Panel B. Level II controls only (11)



Notes: Seven-day windows, P80. 95% confidence intervals; standard errors clustered at the slum level.

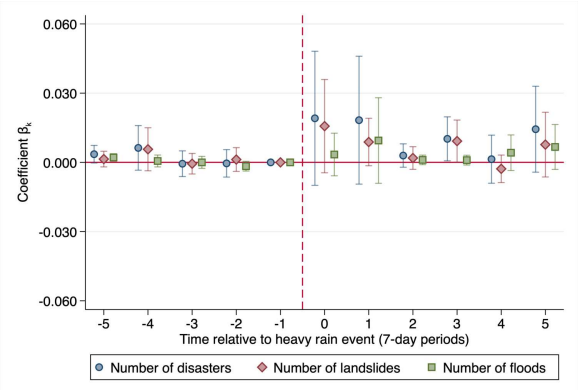
Spillovers. Among slums that were not upgraded, I estimate

$$y_{sw} = \alpha_s + \lambda_{t(w)} + \sum_{k \neq -1} \theta_k \, \mathbf{1}[K_{sw} = k] + \gamma \left(Adjacent_s \times Post_w \right) + \sum_{k \neq -1} \beta_k \big(Adjacent_s \times Post_w \times \mathbf{1}[K_{sw} = k] \big) + \varepsilon_{sw}, \tag{29}$$

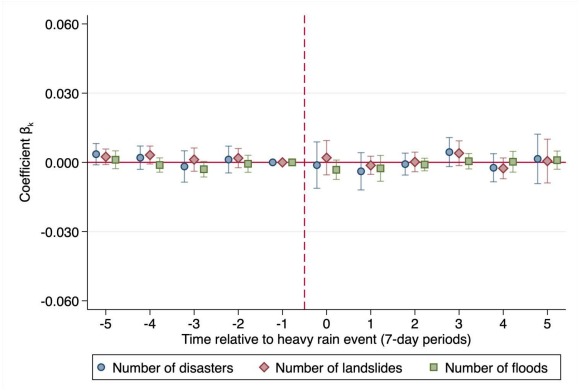
where $Adjacent_s$ equals one if slum s shares a border with an upgraded slum. Panel A uses the 26 Phase II controls (4 adjacent); Panel B adds the 14 neighborhoods outside the Phase II sample that border an upgraded slum (18 adjacent versus 22 non-adjacent controls).

Figure 20: Spillovers to neighborhoods adjacent to upgraded slums

Panel A. Adjacent controls (4) vs. non-adjacent controls (22)



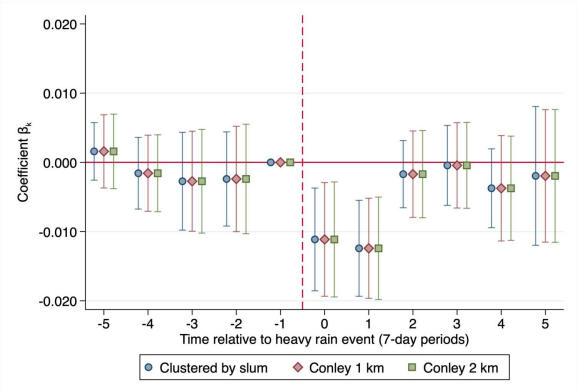
Panel B. Bordering neighborhoods (18) vs. non-adjacent controls (22)



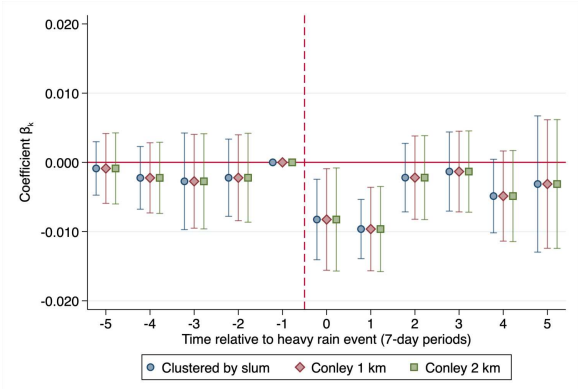
Notes: Seven-day windows, P80. Coefficients β_k from Equation (29); upgraded slums excluded. 95% confidence intervals; standard errors clustered at the slum level.

Figure 21: Spatial correlation: Conley standard errors

Panel A. Number of disasters



Panel B. Number of landslides



Notes: Seven-day windows, P80. Same point estimates. 95% confidence intervals based on standard errors clustered at the slum level and on Conley standard errors with 1 and 2 km distance cutoffs and a Bartlett kernel.

C Urban Outcomes Robustness

C.1 Population Size and Skill Composition

C.1.1 Neighborhood-Level Population

Table 14: Effects of Slum Upgrading on Neighborhood Population: Phase 2 Controls

	All residents	Adults 25+	Low skill	High skill
Post × Upgraded	2,563.6***	1,357.5***	1,348.7***	8.8
	(881.1)	(470.1)	(456.5)	(100.8)
Mean, controls (1993)	5,239.4	2,478.5	2,205.7	272.8
Mean, upgraded (1993)	5,260.0	2,236.1	1,988.2	247.9
R^2	0.972	0.964	0.963	0.933
Neighborhood FE	Yes			
Comuna × year FE	Yes			
Weights: log population (1993)	Yes			
SE clustered by neighborhood	Yes			
Neighborhoods (upgraded/control)	33 (11/22)			
Observations	99			

Notes: Regressions use neighborhood-year observations from 1993, 2005, and 2018, weighted by log 1993 neighborhood population. Outcomes are resident counts. Adults aged 25+ have classified education and equal the sum of low- and high-skill residents; high skill denotes college education. The sample includes 11 upgraded and 22 control neighborhoods. Standard errors are clustered at the neighborhood level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.1.2 Alternative Control Group: Level II

These specifications restrict the control group to Level II neighborhoods scheduled for Phase 2. The main text uses all Phase 2 controls; the treatment group is unchanged.

Table 15: Effects of Slum Upgrading on Population Size and Skill Composition: Level II Controls

Control group:	Level II				
	(1) All residents	(2) Adults 25+	(3) Low-skill adults	(4) High-skill adults	(5) High-skill share
Post × Upgraded	81.503***	27.568***	22.628***	4.940**	-0.005
	(25.225)	(9.710)	(7.910)	(2.323)	(0.010)
Control mean (1993)	191.239	88.239	79.622	8.617	0.111
R^2	0.229	0.293	0.313	0.111	0.298
Observations	2,289	2,289	2,289	2,289	2,289
Neighborhood FE	Yes	Yes	Yes	Yes	Yes
Comuna × year FE	Yes	Yes	Yes	Yes	Yes

Notes: Regressions use census-block-year observations from 1993, 2005, and 2018 and are weighted by log 1993 block population. Standard errors are clustered at the neighborhood level. Columns (2)–(4) count residents aged 25+ with classified education; high skill denotes college education. Column (5) is the high-skill share among classified adults aged 25+. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.1.3 Conley-HAC Standard Errors

Table 16: Population Composition with Conley-HAC Standard Errors

	(1) All residents	(2) Adults 25+	(3) Low-skill adults	(4) High-skill adults	(5) High-skill share
Panel A: Phase 2 Control Group					
Post × Upgraded	49.075***	18.011**	17.484**	0.527	-0.005
	(16.764)	(8.347)	(6.958)	(2.494)	(0.012)
Control mean (1993)	175.412	83.461	74.065	9.397	0.125
N	3042	3042	3042	3042	3042
R-squared	0.228	0.285	0.305	0.150	0.354
Panel B: Level II Control Group					
Post × Upgraded	81.503***	27.568**	22.628**	4.940**	-0.005
	(22.535)	(11.114)	(9.882)	(1.983)	(0.008)
Control mean (1993)	191.239	88.239	79.622	8.617	0.111
N	2289	2289	2289	2289	2289
R-squared	0.229	0.293	0.313	0.111	0.298
Neighborhood FE	Yes	Yes	Yes	Yes	Yes
Comuna-by-year FE	Yes	Yes	Yes	Yes	Yes

Notes: Parentheses contain Conley-HAC standard errors using a Bartlett spatial kernel with a 0.5 km bandwidth and a temporal correction across census waves. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.2 Density and Impervious Surface

C.2.1 Alternative Control Group: Level II

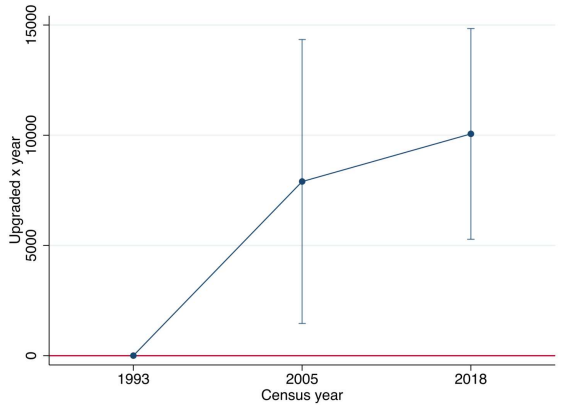
Table 17: Effects of Slum Upgrading on Density and Impervious Surface: Level II Controls

Control group:	Level II slums				
	Density levels		Density logs		Land cover
	(1) Housing per km^2	(2) Population per km^2	(3) Log housing per km^2	(4) Log population per km^2	(5) Impervious surface = 1
Post × Upgraded	8,980***	33,477***	0.395***	0.386**	0.064**
	(2,629)	(9,432)	(0.130)	(0.142)	(0.029)
Control mean (baseline year)	15,511	68,370	9.438	10.912	0.747
Neighborhood FE	Yes	Yes	Yes	Yes	Yes
Comuna-by-year FE	Yes	Yes	Yes	Yes	Yes
Unit of observation	Block-year				30m cell-year
Sample years	1993, 2005, 2018				1985–2020
Observations	2,289	2,289	2,289	2,289	63,448
R-squared	0.232	0.136	0.312	0.182	0.147

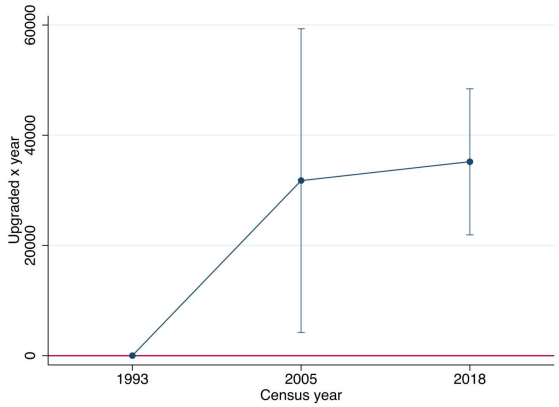
Notes: Columns (1)–(4) use census-block-year observations from 1993, 2005, and 2018, weighted by log 1993 block population. Column (5) uses 30 m × 30 m cell-year observations at five-year intervals over 1985–2020, weighted by log 1993 neighborhood population; its baseline control mean is for 1990. Column (5) equals one if the cell is classified as impervious surface and zero otherwise; it measures land coverage, not building height. Standard errors are clustered at the neighborhood level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure 22: Urban Growth Dynamics: Level II Controls

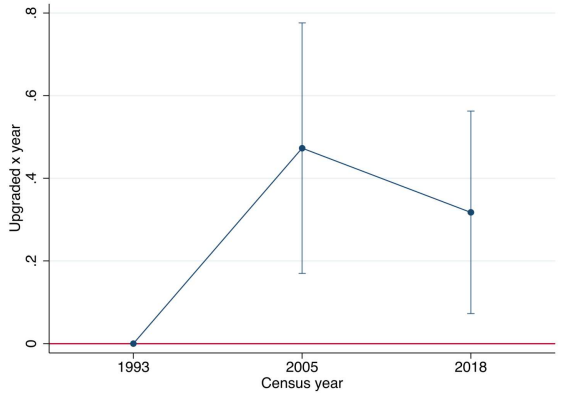
Panel A. Housing Density



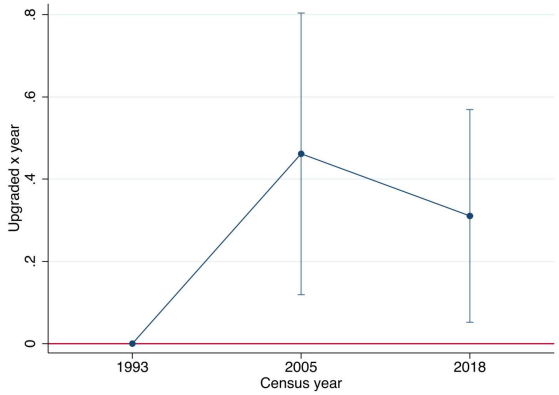
Panel B. Population Density



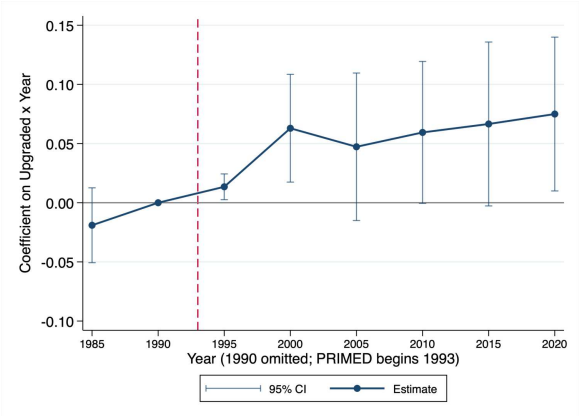
Panel C. Log Housing Density



Panel D. Log Population Density



Panel E. Impervious Surface



Note: Panels A–D use census-block-year observations. Panel E uses 30 m × 30 m cell-year observations. Standard errors are clustered at the neighborhood level. 95% confidence intervals shown.

C.2.2 Unit Fixed Effects

Table 18: Urban Growth with Block and Cell Fixed Effects: Level II Controls

Control group:	Level II slums				
	Density levels		Density logs		Built-up surface
	(1) Housing per km^2	(2) Population per km^2	(3) Log housing per km^2	(4) Log population per km^2	(5) Built-up surface = 1
Post × Upgraded	8,980***	33,477***	0.395***	0.386**	0.064**
	(2,629)	(9,432)	(0.130)	(0.142)	(0.029)
Control mean (baseline year)	15,511	68,370	9.438	10.912	0.747
Unit FE	Block	Block	Block	Block	Cell
Comuna-by-year FE	Yes	Yes	Yes	Yes	Yes
Unit of observation	Block-year				30m cell-year
Sample years	1993, 2005, 2018				1985–2020
Observations	2,289	2,289	2,289	2,289	63,448
R-squared	0.706	0.711	0.769	0.699	0.758

Notes: Columns (1)–(4) absorb census-block fixed effects; column (5) absorbs 30 m cell fixed effects. Census outcomes are observed in 1993, 2005, and 2018; the impervious-surface panel is observed at five-year intervals over 1985–2020. Regressions are weighted by baseline 1993 log population. Standard errors are clustered at the neighborhood level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.2.3 Conley-HAC Standard Errors

Table 19: Urban Growth with Conley-HAC Standard Errors: Phase 2 Control

Control group:	Phase 2					
	(1) Log housing per km^2	(2) Log population per km^2	(3) Impervious surface = 1	(4) Log housing per km^2	(5) Log population per km^2	(6) Impervious surface = 1
Post × Upgraded	0.256***	0.194**	0.009	0.256***	0.194**	0.009
	(0.082)	(0.090)	(0.024)	(0.075)	(0.084)	(0.024)
Control mean (baseline year)	9.396	10.843	0.735	9.396	10.843	0.735
Neighborhood FE	Yes	Yes	Yes	Yes	Yes	Yes
Comuna-by-year FE	Yes	Yes	Yes	Yes	Yes	Yes
Unit of observation	Block-year	Block-year	30m cell-year	Block-year	Block-year	30m cell-year
Sample years	1993, 2005, 2018	1993, 2005, 2018	1985–2020	1993, 2005, 2018	1993, 2005, 2018	1985–2020
Serial correlation	No	No	No	Yes	Yes	Yes
Observations	3,042	3,042	98,160	3,042	3,042	98,160
R-squared	0.326	0.218	0.160	0.326	0.218	0.160

Notes: Parentheses contain Conley-HAC standard errors using a Bartlett spatial kernel with a 0.5 km bandwidth. Columns (1)–(3) allow spatial correlation; columns (4)–(6) additionally allow serial correlation across census waves or adjacent five-year periods. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 20: Urban Growth with Conley-HAC Standard Errors: Level II Control

Control group:	Level II slums					
	(1) Log housing per km^2	(2) Log population per km^2	(3) Impervious surface = 1	(4) Log housing per km^2	(5) Log population per km^2	(6) Impervious surface = 1
Post × Upgraded	0.395***	0.386***	0.064**	0.395***	0.386***	0.064**
	(0.104)	(0.110)	(0.028)	(0.095)	(0.103)	(0.028)
Control mean (baseline year)	9.438	10.912	0.747	9.438	10.912	0.747
Neighborhood FE	Yes	Yes	Yes	Yes	Yes	Yes
Comuna-by-year FE	Yes	Yes	Yes	Yes	Yes	Yes
Unit of observation	Block-year	Block-year	30m cell-year	Block-year	Block-year	30m cell-year
Sample years	1993, 2005, 2018	1993, 2005, 2018	1985–2020	1993, 2005, 2018	1993, 2005, 2018	1985–2020
Serial correlation	No	No	No	Yes	Yes	Yes
Observations	2,289	2,289	63,448	2,289	2,289	63,448
R-squared	0.312	0.182	0.147	0.312	0.182	0.147

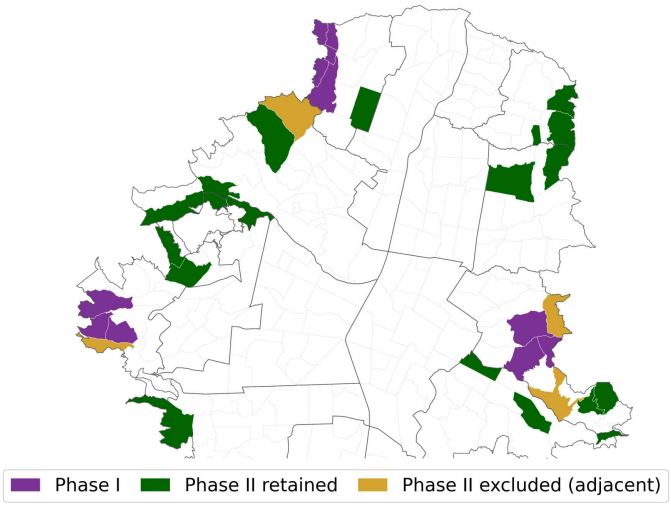
Notes: Parentheses contain Conley-HAC standard errors using a Bartlett spatial kernel with a 0.5 km bandwidth. Columns (1)–(3) allow spatial correlation; columns (4)–(6) additionally allow serial correlation across census waves or adjacent five-year periods. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.2.4 Excluding Nearby Controls

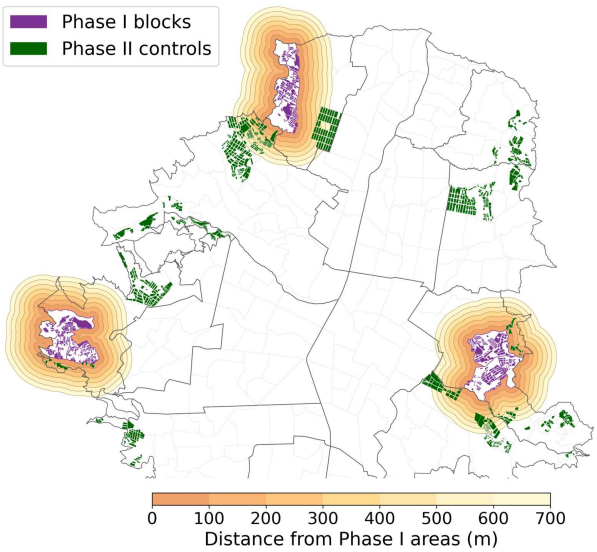
Figure 23 illustrates two alternative sample restrictions for robustness checks: excluding adjacent Phase II neighborhoods (Panel A) and excluding control blocks near upgraded areas (Panel B). Table 21 reports the corresponding estimates for resident counts and housing density.

Figure 23: Robustness to Excluding Nearby Controls

Panel A. Excluding Adjacent Phase II Neighborhoods



Panel B. Excluding Control Blocks by Distance



Note: Panel A excludes four Phase II neighborhoods that share a boundary with Phase I areas and retains 19. All control blocks in the four adjacent neighborhoods are excluded. Panel B shows buffers in 100 m increments up to 700 m. For each cutoff, the Phase II control blocks shown in green are excluded if their polygons intersect the corresponding buffer. Treated Phase I blocks remain in both checks.

Table 21: Effects of Slum Upgrading after Excluding Nearby Controls

			Control blocks excluded within (m)						
	Main specification	Excluded adjacent slums	100	200	300	400	500	600	700
Panel A. Number of residents									
Residents per block	49.08**	51.34**	54.18**	57.25***	57.25***	56.60**	51.74**	45.63*	31.83
	(21.30)	(22.56)	(20.80)	(20.71)	(20.71)	(20.96)	(21.63)	(23.05)	(18.81)
Control mean (1993)	175.41	183.68	176.72	177.52	176.88	178.21	175.94	172.07	168.39
Panel B. Housing density									
Housing units per km^2	6,383.51***	5,949.44**	6,360.20***	6,414.48***	6,414.48***	6,327.00***	6,015.27**	5,155.03**	3,968.07*
	(2,086.62)	(2,422.73)	(2,218.41)	(2,236.25)	(2,236.36)	(2,247.21)	(2,289.71)	(2,301.77)	(2,000.19)
Control mean (1993)	15,546.85	15,812.26	15,708.95	15,814.72	15,965.70	15,888.74	15,794.87	15,691.23	15,532.17
Sample and specification									
Block-year observations	3,042	2,697	2,982	2,952	2,916	2,835	2,724	2,601	2,490
Treated blocks	366	366	366	366	366	366	366	366	366
Control blocks	648	533	628	618	606	579	542	501	464
Neighborhood fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Comuna-by-year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: All columns compare upgraded Phase I blocks with Phase II controls. The main specification retains all eligible controls. The adjacency restriction excludes all control blocks in Phase II neighborhoods sharing a boundary with Phase I areas (Figure 23, Panel A). Each distance column separately excludes control blocks whose polygons intersect the corresponding buffer around Phase I areas (Panel B). Treated blocks remain in every specification. Regressions use log 1993 block population weights. Parentheses contain standard errors clustered by neighborhood. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.2.5 Potential Spillovers to Nearby Untreated Blocks

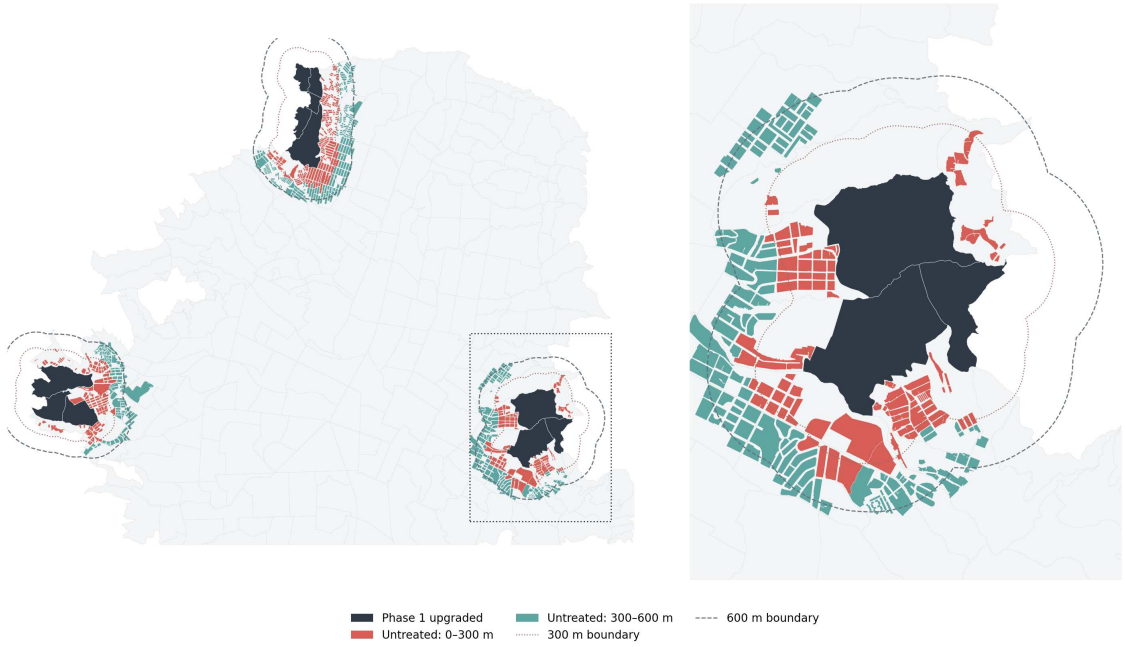
To examine potential spillovers, I compare changes in resident counts and housing density among untreated blocks at different distances from upgraded areas. I estimate

$$Y_{bt} = \alpha_{n(b)} + \lambda_{c(b),t} + \gamma Near_b + \delta Post_t + \beta(Near_b \times Post_t) + \varepsilon_{bt}, \tag{30}$$

where $Near_b$ indicates a block within 300 m of an upgraded area, and $Post_t$ equals one in 2005 and 2018, with 1993 as the baseline. The terms $\alpha_{n(b)}$ and $\lambda_{c(b),t}$ denote neighborhood and comuna-by-year fixed effects. The coefficient β measures the additional change in the near band relative to the reference band.

Table 22 compares untreated blocks within 0–300 m of upgraded areas with reference bands starting at 300 m and extending from 600 m to 1 km. The sample includes all eligible untreated blocks, whether or not they belong to Phase II. Figure 24 maps the 300–600 m comparison. With that reference band, the estimated differences in changes are 13.3 residents per block and 369.2 housing units per km^2. Neither differential is statistically significant across any of the reference bands.

Figure 24: Nearby Untreated Blocks Used in the Spillover Comparison



Note: Dark polygons mark upgraded Phase I areas. Red blocks form the 0–300 m group; teal blocks form the 300–600 m reference group. Colored blocks belong to the estimation sample in the first column of Table 22. The right-hand panel enlarges the boxed area. Distance-band membership uses block-polygon intersection.

Table 22: Potential Spillovers with Alternative Reference Bands

Reference band (m)	300–600	300–700	300–800	300–900	300–1,000
Residents per block					
Near × Post	13.3	13.6	10.4	8.6	6.8
	(13.1)	(13.8)	(14.9)	(14.9)	(15.2)
p-value	0.319	0.331	0.488	0.568	0.654
Housing units per km^2					
Near × Post	369.2	890.0	-113.5	59.9	66.7
	(1,309.5)	(1,280.5)	(1,731.9)	(1,656.5)	(1,601.8)
p-value	0.780	0.491	0.948	0.971	0.967
Observations (block-years)	2,457	2,859	3,267	3,750	4,284
Reference blocks	434	568	704	865	1,043
Neighborhood clusters	34	37	43	51	57
Neighborhood FE	Yes	Yes	Yes	Yes	Yes
Comuna × year FE	Yes	Yes	Yes	Yes	Yes

Notes: Estimates of β in Equation (30) for all eligible untreated blocks, including Phase II and other untreated areas. Upgraded Phase I blocks are excluded. The near group is fixed at 0–300 m in every column. Reference bands start at 300 m and extend to the outer distance shown in each column heading. Distance bands use block-polygon intersection with buffers around Phase I areas; blocks intersecting the 300 m buffer enter the near group. Regressions use census-block-year observations from 1993, 2005, and 2018, weighted by log 1993 block population. Parentheses contain standard errors clustered by neighborhood.

C.3 Housing Characteristics

C.3.1 Alternative Control Group: Level II

Table 23: Effects of Slum Upgrading on Housing Characteristics: Level II Controls

Control Group:	Level II slums			
	(1)	(2)	(3)	(4)
	Water access	Sewerage access	Avg. rooms	Brick walls
Post × Upgraded	0.006	0.015	0.287*	0.079
	(0.006)	(0.021)	(0.165)	(0.057)
Control mean (1993)	0.956	0.884	3.239	0.810
Neighborhood FE	Yes	Yes	Yes	Yes
Comuna-by-year FE	Yes	Yes	Yes	Yes
N	2289	2289	1526	1526
R-squared	0.259	0.445	0.422	0.149

Notes: Regressions use census-block-year observations from 1993, 2005, and 2018 and are weighted by log 1993 block population. Standard errors are clustered at the neighborhood level. Columns (4) and (5) have fewer observations because questions on rooms and wall materials were asked only in 1993 and 2018. Water, sewerage, and brick-wall outcomes are shares. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.3.2 Conley-HAC Standard Errors

Table 24: Housing Characteristics with Conley-HAC Standard Errors

	(1) Water access	(2) Sewerage access	(3) Avg. rooms	(4) Brick walls
Panel A: Phase 2 Control Group				
Post × Upgraded	-0.017	-0.028	0.196	0.023
	(0.011)	(0.019)	(0.121)	(0.031)
Control mean (1993)	0.956	0.887	3.389	0.817
N	3042	3042	2028	2028
R-squared	0.275	0.421	0.422	0.314
Panel B: Level II Control Group				
Post × Upgraded	0.006	0.015	0.287***	0.079*
	(0.006)	(0.018)	(0.102)	(0.041)
Control mean (1993)	0.956	0.884	3.239	0.810
N	2289	2289	1526	1526
R-squared	0.259	0.445	0.422	0.149
Neighborhood FE	Yes	Yes	Yes	Yes
Comuna-by-year FE	Yes	Yes	Yes	Yes

Notes: Parentheses contain Conley-HAC standard errors using a Bartlett spatial kernel with a 0.5 km bandwidth and a temporal correction across census waves. Columns (3) and (4) use the 1993 and 2018 census waves. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.4 Additional Population Characteristics

C.4.1 Population by Age and Sex

Table 25: Effects of Slum Upgrading on Population by Age and Sex

	(1) Ages 0–14	(2) Ages 15–64	(3) Ages 65+	(4) Female residents
Panel A. Phase 2 controls				
Post × Upgraded	18.000*	34.648**	-3.573***	26.243**
	(8.856)	(13.341)	(1.252)	(11.658)
Control mean (1993)	57.181	111.404	6.827	92.943
R^2	0.209	0.249	0.394	0.229
Observations	3,042	3,042	3,042	3,042
Panel B. Level II controls				
Post × Upgraded	33.261***	53.401***	-5.159***	44.685***
	(9.505)	(16.801)	(1.406)	(13.340)
Control mean (1993)	64.244	120.111	6.884	101.300
R^2	0.192	0.254	0.413	0.232
Observations	2,289	2,289	2,289	2,289
Neighborhood FE	Yes	Yes	Yes	Yes
Comuna × year FE	Yes	Yes	Yes	Yes

Notes: Regressions use census-block-year observations from 1993, 2005, and 2018 and are weighted by log 1993 block population. Standard errors are clustered at the neighborhood level. Outcomes are resident counts. Female residents include all ages and overlap with the age categories. Panel A uses Phase 2 controls; Panel B uses Level II controls. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.4.2 Literacy and Migration Background

Table 26: Effects of Slum Upgrading on Literacy and Migration Background

	(1) Literacy rate	(2) Born in Medellin	(3) Lived in Medellin five years earlier
Panel A. Phase 2 controls			
Post × Upgraded	0.012	-0.006	0.017
	(0.010)	(0.022)	(0.012)
Control mean (1993)	0.897	0.539	0.843
R^2	0.422	0.275	0.272
Observations	3,042	3,042	3,042
Panel B. Level II controls			
Post × Upgraded	0.014	-0.011	0.020*
	(0.011)	(0.022)	(0.011)
Control mean (1993)	0.884	0.520	0.829
R^2	0.423	0.281	0.269
Observations	2,289	2,289	2,289
Neighborhood FE	Yes	Yes	Yes
Comuna × year FE	Yes	Yes	Yes

Notes: Regressions use census-block-year observations from 1993, 2005, and 2018 and are weighted by log 1993 block population. Standard errors are clustered at the neighborhood level. All outcomes are shares. Having lived in Medellin five years earlier includes residents who remained in the same neighborhood and does not identify intraurban migration. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.5 Changes in Population by Risk Category

Table 27: Changes in Population by Risk Category: Phase II Controls

	Log population
Panel A. Landslide Susceptibility Index	
Post × Upgraded	0.357^{***}
	(0.128)
Post × Upgraded × High landslide susceptibility	-0.311^{**}
	(0.116)
R-squared	0.274
Panel B. Index excluding slope	
Post × Upgraded	0.358^{**}
	(0.134)
Post × Upgraded × High landslide susceptibility	-0.291^{**}
	(0.122)
R-squared	0.270
Panel C. Index excluding slope, with slope controls	
Post × Upgraded	0.274^{**}
	(0.127)
Post × Upgraded × High landslide susceptibility	-0.238^{**}
	(0.112)
R-squared	0.275
Panel D. Slope above $15°$	
Post × Upgraded	0.230^{*}
	(0.113)
Post × Upgraded × (Slope $> 15°$)	-0.250^{**}
	(0.117)
R-squared	0.271
Observations	3,042
Neighborhood FE	Yes
Comuna × Year FE	Yes

Notes: Regressions use census-block-year observations from 1993, 2005, and 2018 and are weighted by log 1993 block population. Standard errors are clustered at the neighborhood level. Estimates correspond to equation (28). High susceptibility is defined relative to the median. Panel C additionally controls for $Slope$, $Post \times Slope$, and $Upgraded \times Slope$. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 28: Changes in Population by Risk Category: Level II Controls

	Log population
Panel A. Landslide Susceptibility Index	
Post × Upgraded	0.479^{***}
	(0.113)
Post × Upgraded × High landslide susceptibility	-0.261^{*}
	(0.126)
R-squared	0.249
Panel B. Index excluding slope	
Post × Upgraded	0.494^{***}
	(0.114)
Post × Upgraded × High landslide susceptibility	-0.276^{**}
	(0.108)
R-squared	0.250
Panel C. Index excluding slope, with slope controls	
Post × Upgraded	0.439^{***}
	(0.116)
Post × Upgraded × High landslide susceptibility	-0.247^{**}
	(0.103)
R-squared	0.254
Panel D. Slope above $15°$	
Post × Upgraded	0.366^{***}
	(0.129)
Post × Upgraded × (Slope $> 15°$)	-0.105
	(0.104)
R-squared	0.247
Observations	2,289
Neighborhood FE	Yes
Comuna × Year FE	Yes

Notes: Regressions use census-block-year observations from 1993, 2005, and 2018 and are weighted by log 1993 block population. Standard errors are clustered at the neighborhood level. Estimates correspond to equation (28). High susceptibility is defined relative to the median. Panel C additionally controls for $Slope$, $Post \times Slope$, and $Upgraded \times Slope$. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

D Spatial Model Appendix

D.1 Commuting Elasticity

Table 29: Commuting elasticity using Google Maps travel time: 2017

	(1) Overall	(2) Low skill	(3) High skill
Travel time (min.)	-0.057***	-0.063***	-0.043***
	(0.003)	(0.003)	(0.004)
Implied elasticity	5.73	6.27	4.32
Observations	41,289	36,838	20,851
Origin FE	Yes	Yes	Yes
Destination FE	Yes	Yes	Yes

Notes: Direct home-to-work trips by adults aged 18+ departing between 6–9AM from the 2017 Origin-Destination Survey; one first work trip per person. Travel time is the average of car and public-transport times from the Google Maps API, with requested service dates of March 31 and July 3, 2026, respectively. These times are proxies for 2017 conditions. Each observation is a neighborhood-origin–neighborhood-destination cell. Within-neighborhood trips are excluded, and zero surveyed flows are included on feasible routes. High skill is university/postgraduate education. PPML regressions include origin and destination fixed effects and are estimated separately by skill in cols (2) and (3). The reported coefficient is ν_g. Implied elasticities are $\varepsilon_g = -\nu_g/\kappa$, with $\kappa = 0.01$. Standard errors are clustered by neighborhood origin-destination pair. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 30: Commuting elasticity at the Comuna level

	(1) Overall	(2) Low skill	(3) High skill
A. PPML: all commune-pair cells			
Travel time (min.), ν_g	-0.068^{***}	-0.073^{***}	-0.058^{***}
	(0.004)	(0.004)	(0.003)
Implied elasticity, ε_g	6.80	7.27	5.80
Observations	512	512	512
B. PPML: positive commune flows			
Travel time (min.), ν_g	-0.065^{***}	-0.068^{***}	-0.053^{***}
	(0.003)	(0.004)	(0.003)
Implied elasticity, ε_g	6.48	6.78	5.30
Observations	451	427	327
C. OLS: log positive commune flows			
Travel time (min.), ν_g	-0.062^{***}	-0.061^{***}	-0.048^{***}
	(0.004)	(0.005)	(0.004)
Implied elasticity, ε_g	6.20	6.05	4.77
Observations	451	427	327
Origin commune × year FE	Yes	Yes	Yes
Destination commune × year FE	Yes	Yes	Yes

Notes: Direct home-to-work trips from the 2012 and 2017 Origin-Destination Surveys, aggregated to the comuna level. Travel time is computed with the network corresponding to each year. Each observation is a comuna-origin–comuna-destination–year cell. Panel A includes all cells; Panels B and C include positive flows only. Panels A and B are estimated by PPML and Panel C by OLS. The reported coefficient is ν_g. Implied elasticities are $\varepsilon_g = -\nu_g/\kappa$, with $\kappa = 0.01$. Standard errors are clustered by comuna origin-destination pair. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 31: Commuting elasticity across trip definitions

	(1) Overall	(2) Low skill	(3) High skill
A. Main: home-to-work trips, departures 06:00–09:00			
Travel time (min.), ν_g	-0.069^{***}	-0.076^{***}	-0.056^{***}
	(0.002)	(0.003)	(0.003)
Implied elasticity, ε_g	6.94	7.60	5.60
Observations	82,989	76,944	44,448
B. Direct home-to-work trips, all departure hours			
Travel time (min.), ν_g	-0.069^{***}	-0.075^{***}	-0.057^{***}
	(0.002)	(0.002)	(0.003)
Implied elasticity, ε_g	6.88	7.50	5.67
Observations	91,563	87,391	56,037
C. All trips, all purposes and departure hours			
Travel time (min.), ν_g	-0.099^{***}	-0.102^{***}	-0.070^{***}
	(0.001)	(0.001)	(0.002)
Implied elasticity, ε_g	9.95	10.19	6.95
Observations	104,032	104,032	88,045
Origin × year FE	Yes	Yes	Yes
Destination × year FE	Yes	Yes	Yes

Notes: Data are from the 2012 and 2017 Origin-Destination Surveys, and travel time is computed with the network corresponding to each year. Panel A (Main) repeats Table 6, using direct home-to-work trips departing from 06:00 (inclusive) to 09:00 (exclusive), Panel B uses direct home-to-work trips at all hours, and Panel C uses all trips. Panels A and B retain the first work trip per adult (18+) whose origin neighborhood matches home. Each observation is a neighborhood-origin–neighborhood-destination–year cell. Regressions are estimated by PPML on unexpanded trip counts, including zero-flow cells. Average network time gives car and public transportation equal weight and is constructed without a departure-time window. The reported coefficient is ν_g. Implied elasticities are $\varepsilon_g = -\nu_g/\kappa$, with $\kappa = 0.01$. Standard errors are clustered by neighborhood origin-destination pair. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 32: Commuting elasticities by survey year

	(1) Overall	(2) Low skill	(3) High skill
A. 2012			
Travel time (min.), ν_g	-0.069^{***}	-0.076^{***}	-0.055^{***}
	(0.003)	(0.003)	(0.005)
Implied elasticity, ε_g	6.88	7.65	5.47
Observations	45,365	44,106	27,874
B. 2017			
Travel time (min.), ν_g	-0.066^{***}	-0.071^{***}	-0.055^{***}
	(0.003)	(0.003)	(0.004)
Implied elasticity, ε_g	6.65	7.14	5.53
Observations	46,198	43,285	27,789
Origin FE	Yes	Yes	Yes
Destination FE	Yes	Yes	Yes

Notes: Direct home-to-work trips at all departure times, with each sampled trip weighted by its expansion factor. Panels A and B use the 2012 and 2017 Origin-Destination Surveys separately. The 2017 expansion factors are approximate weights assigned by origin comuna. Travel time is computed with the network corresponding to each year and averages car and public transportation with equal weights. Each observation is a neighborhood origin-destination cell. PPML regressions include origin and destination fixed effects and zero-flow cells. The reported coefficient is ν_g. Implied elasticities are $\varepsilon_g = -\nu_g/\kappa$, with $\kappa = 0.01$. Standard errors are clustered by neighborhood origin-destination pair. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 33: Network travel time weighted by survey mode shares

	(1) Overall	(2) Low skill	(3) High skill
A. Citywide mode shares by year			
Travel time (min.), ν_g	-0.053^{***}	-0.058^{***}	-0.044^{***}
	(0.001)	(0.002)	(0.002)
Implied elasticity, ε_g	5.30	5.78	4.37
Observations	91,563	87,391	56,037
B. Origin–destination comuna mode shares by year			
Travel time (min.), ν_g	-0.049^{***}	-0.051^{***}	-0.047^{***}
	(0.001)	(0.002)	(0.002)
Implied elasticity, ε_g	4.92	5.07	4.66
Observations	91,563	87,391	56,037
Origin × year FE	Yes	Yes	Yes
Destination × year FE	Yes	Yes	Yes

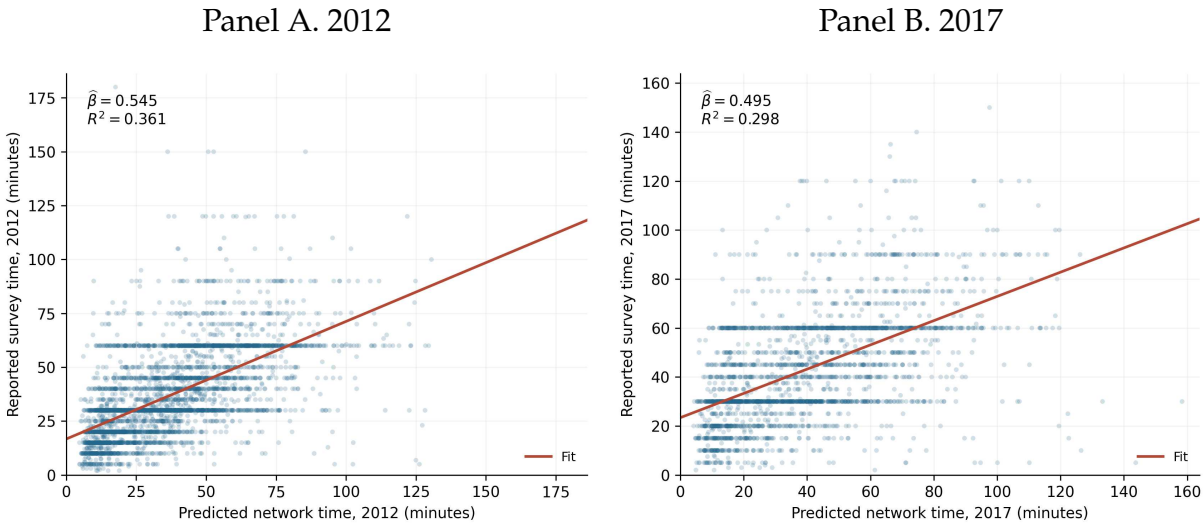
Notes: Direct home-to-work trips at all hours from the 2012 and 2017 Origin-Destination Surveys. Panel A weights car and public-transport network times using the share of declared tranportation mode in each year at the city level. Panels B weights travel time using the share of each tranportation model declared in each survey at the comuna level. Each observation is a neighborhood-origin–neighborhood-destination–year cell. Regressions are estimated by PPML. The reported coefficient is ν_g. Implied elasticities are $\varepsilon_g = -\nu_g/\kappa$, with $\kappa = 0.01$. Standard errors are clustered by neighborhood origin-destination pair. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 34: Network travel-time assumptions by mode

	Speed (km/h)		Time added (minutes)			
Mode	2012	2017	Wait	Dwell	Other	Network availability
Car (route mean)	30.6	30.6	0	0	3	Both years
Walking/access	5	5	0	0	0	Both years
Ordinary bus	16	16	7.5	0	0	Both years
Metro	40	40	3	0.25	0	Lines A and B
Metrocable	18	18	2	0	0	K, J, L in 2012; adds H in 2017
Metroplus	16	16	5	0.5	0	Line 1 in 2012; adds Line 2 in 2017
Tram	–	16	4	0.5	0	Ayacucho T from 2016

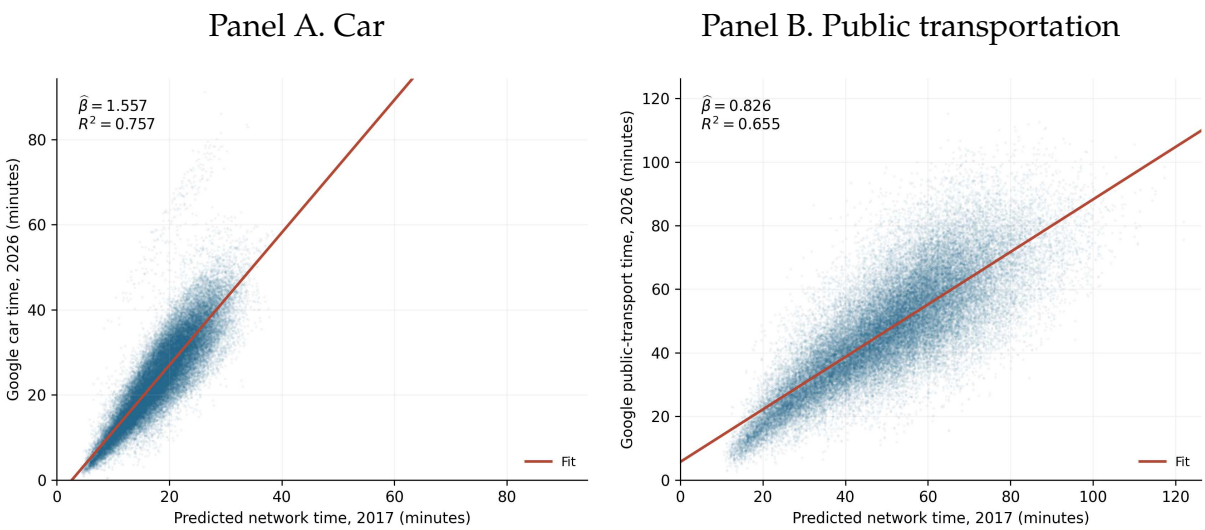
Notes: The table reports the speeds and time components used to construct the 2012 and 2017 travel-time networks. The 2017 network adds Line 2, Tram T, and Cable H. A dash indicates that the service was unavailable in that year.

Figure 25: Reported survey time and network time



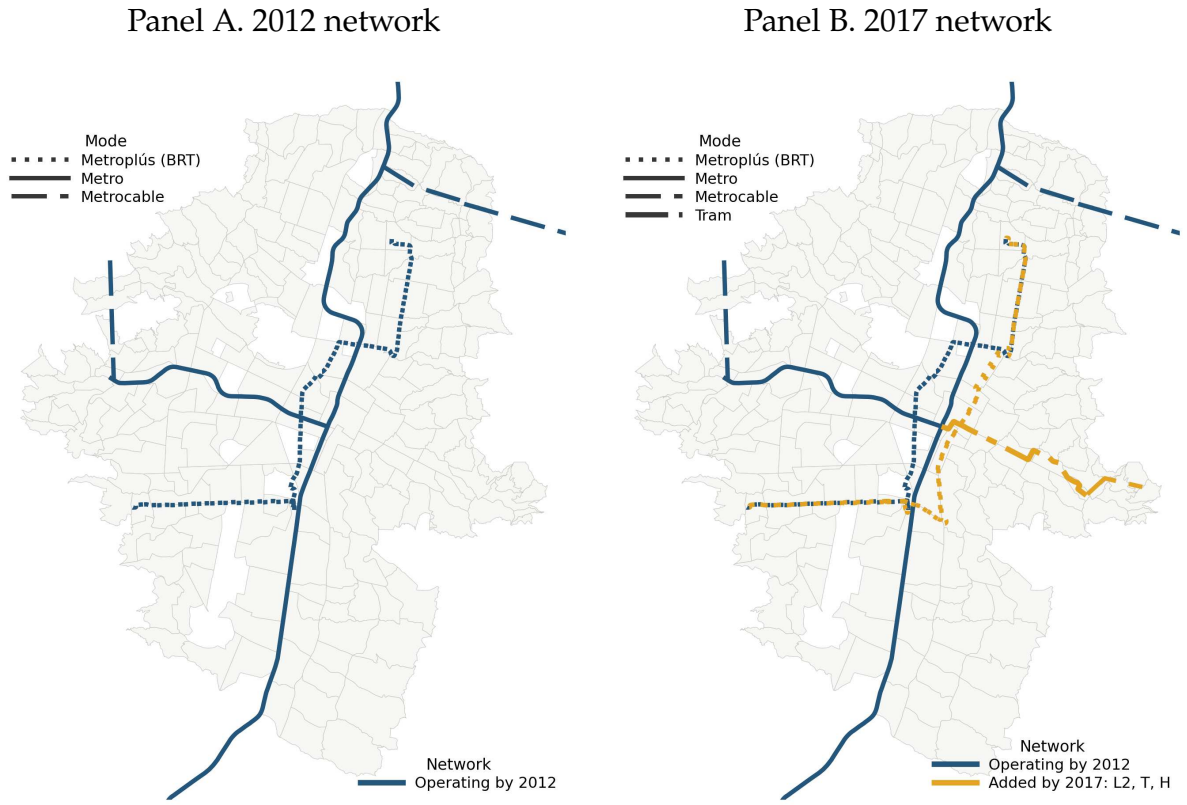
Note: Each point is an observed neighborhood origin-destination-chosen-mode cell. Network times use the network dated to the corresponding survey year. The displayed coefficient is the slope from a levels regression with an intercept.

Figure 26: Google 2026 time and 2017 network time



Note: Each point is a neighborhood pair with both a Google route and a feasible network route. The car comparison has 54,086 pairs. The public comparison has 57,631 pairs and uses the faster available complete ordinary-bus or integrated Metro/Metroplús/Metrocable/tram journey. The displayed coefficient is the slope from a levels regression with an intercept.

Figure 27: Historically dated mass-transit networks



Note: The maps use the city's mass-transit line geometry and the opening-year rule used to construct the dated routing networks. Relative to 2012, the 2017 network adds Metroplús Line 2, Tram T, and Cable H. Lines M, O, and P opened after 2017 and are excluded. The road, walking, and ordinary-bus graphs are held fixed across the two years and are omitted from the maps to make the historical change visible.

D.2 Supplemental details of calibration

Income. To construct income, I normalize $T_L = 1$ and jointly solve for T_H and workplace wages using labor-market clearing and the observed distribution of residents and employment. I choose T_H to reproduce the monthly wage premium, measured as the geometric mean of the 2015–2018 annual ratios of survey-weighted mean monthly labor earnings for high- and low-skill workers from the *Gran Encuesta Integrada de Hogares* (*GEIH*). This target equals 3.227 and yields $T_H = 18.405$. Using the recovered workplace wages and workplace-choice probabilities, I compute mean labor income in each neighborhood and skill group from equation (9).

Floor space and construction. Given recovered incomes and firm wage bills, I obtain residential and commercial floor-space demand by dividing household housing expenditure and firms' floor-space expenditure by the corresponding cadastral rents. Total demand is $H_i^D = H_{iR}^D + H_{iM}^D$. Imposing market clearing, $H_i^S = H_i^D$, I recover ϕ_i as the local construction supply shifter that rationalizes demand-implied floor space, given land, rents, and the common capital price:

$$\phi_i = \left(\frac{H_i^D}{L_i}\right)^{1-\mu} \left(\frac{P}{\mu r_i}\right)^{\mu}. \tag{31}$$

I determine the baseline common capital price P by normalizing the geometric mean of ϕ_i across all locations to one.

Productivity and amenities. I recover productivity fundamentals from firms' production conditions after accounting for agglomeration. I recover skill specific amenity fundamentals as a residual from observed residential shares after accounting for income access, rents, disaster risk, and residential composition. For each skill group, I normalize the geometric mean of amenity fundamentals across residential locations to one.

D.3 Phase I parameter sensitivity

Table 35: Phase I: sensitivity to disaster disutility and floor-space supply

	Baseline		$\gamma = -0.1$		$\gamma = -0.5$		$\gamma = -1$		$\mu = 0.7$	
	Closed	Open	Closed	Open	Closed	Open	Closed	Open	Closed	Open
Panel A. Welfare gains										
Fundamental amenities (pp)	0.189	0.189	0.197	0.196	0.183	0.182	0.170	0.168	0.190	0.189
Disaster risk (pp)	0.328	0.326	0.110	0.109	0.511	0.507	0.935	0.925	0.328	0.327
Other channels (pp)	1.445	0.145	0.856	0.086	1.941	0.195	3.091	0.309	1.410	0.112
Total welfare gain (%)	1.962	0.660	1.162	0.391	2.635	0.884	4.195	1.403	1.928	0.628
Disaster share (%)	16.691	49.416	9.446	27.960	19.373	57.361	22.279	65.970	17.027	52.019
Panel B. Phase I neighborhoods: removal effects										
Residents (%)	-17.986	-19.356	-10.681	-11.571	-24.096	-25.790	-38.149	-40.316	-17.608	-18.921
Residential rents (%)	-4.746	-5.079	-2.738	-2.940	-6.523	-6.962	-11.024	-11.686	-5.531	-5.903
Residential floorspace (%)	-14.105	-14.991	-8.295	-8.859	-19.060	-20.176	-30.845	-32.345	-12.996	-13.779
Panel C. Citywide removal effects										
Residents (%)	0.000	-1.852	0.000	-1.105	0.000	-2.473	0.000	-3.885	0.000	-1.764
Residential rents (%)	-0.121	-0.507	-0.066	-0.296	-0.174	-0.690	-0.328	-1.142	-0.143	-0.576

Notes: Panel A decomposes citywide welfare gains from Phase I relative to no Phase I using five-factor Shapley accounting. Other channels sums the signed contributions from high-skill resident share, income access, and residential rents. Contributions in percentage points sum to total welfare gains before rounding; disaster share is the disaster-risk contribution divided by the total gain. Panels B and C report removal effects relative to the observed city with Phase I. Residents are classified adults aged 25 and above. Rents per unit of residential floor space use fixed observed resident weights. Baseline and γ alternatives use $\mu = 0.75$; the lower-supply case uses $\mu = 0.7$ and $\gamma = -0.31$. Baseline fundamentals are re-inverted; other structural parameters and policy shocks are held fixed.

D.4 Full-program welfare decomposition

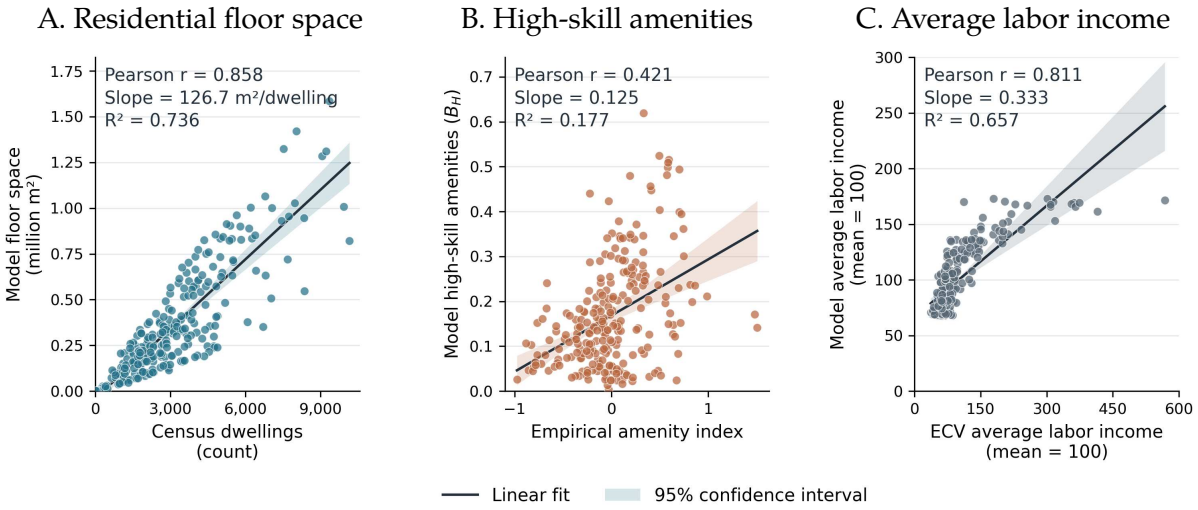
Table 36: Contributions to welfare gains from the full PRIMED program

	Closed city		Open city	
Channel	pp	Share (%)	pp	Share (%)
Fundamental amenities	0.662	10.9	0.648	31.1
Disaster risk	1.082	17.8	1.059	50.9
High-skill share	4.402	72.3	1.415	67.9
Income access	0.062	1.0	−0.476	−22.8
Residential rent	−0.119	−1.9	−0.563	−27.1
Total welfare gain (%)	6.09	100.0	2.08	100.0

Notes: The completed program includes Phases I and II and is compared with no PRIMED under each city closure. Contributions use the same five-factor Shapley accounting as Table 10. Columns report percentage-point contributions and shares of the net welfare gain. Contributions sum to total welfare gains before rounding.

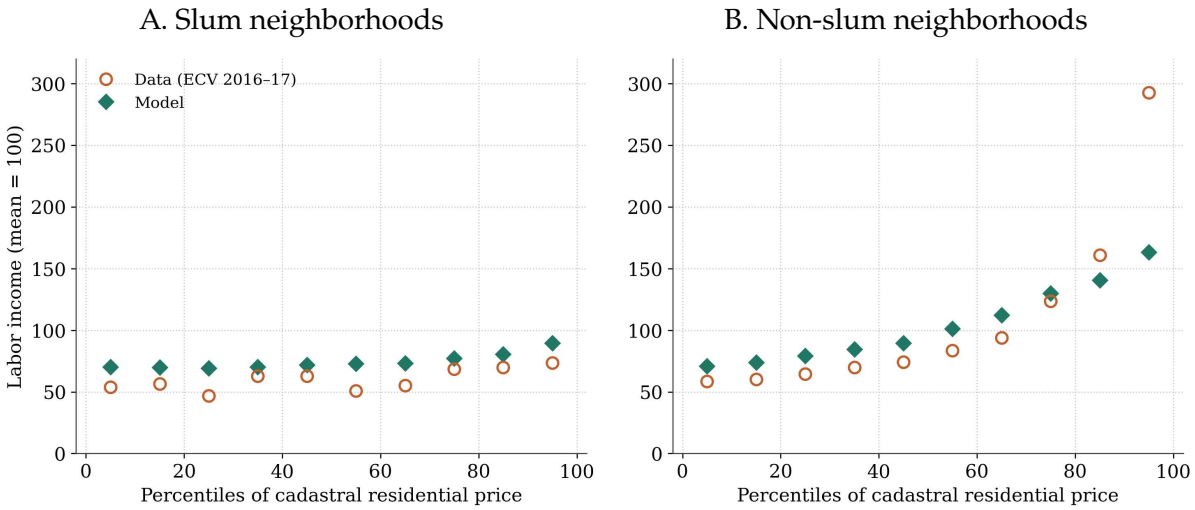
D.5 Model validation

Figure 28: External validation of the model



Notes: Each point is a neighborhood. A: 248 residential neighborhoods; floor space is inferred from expenditure and anchored prices. The displayed slope is in m^2 per census dwelling. B: total high-skill amenities B_H in levels, for the same 248 neighborhoods. The empirical amenity index built as described in subsection D.6. C: all-resident income for 232 eligible neighborhoods in the pooled 2016-2017 Encuesta de Calidad de Vida Lines are descriptive OLS; shading is the pointwise 95% confidence interval for the fitted mean.

Figure 29: Labor income by cadastral residential-price decile



Notes: Cadastral residential-price deciles are computed separately within slum and non-slum neighborhoods. Each point reports the mean labor-income index within a decile, weighted by the number of residents aged 25 and older. Each income series is scaled so that its simple average across neighborhoods equals 100.

D.6 Empirical composite index of amenities

I construct an empirical index to assess whether the model's recovered amenities for high-skill residents align with observed neighborhood conditions. The index combines five dimensions: establishments, schools, safety, governance, and quietness. I standardize each component, orient it so that higher values indicate more favorable conditions, and take their simple mean. This gives each dimension an equal positive weight.[9] The index is available for 249 neighborhoods; 248 enter the model validation.

Establishments. I use data on establishments from Overture Places. I exclude establishments identified as permanently closed. I calculate densities per square kilometer for ten service categories, standardize each density, average the ten scores, and standardize that average. Thus, categories with more establishments do not mechanically receive greater weight.

Schools. I use 2019 Saber 11 results and municipal school locations to measure the density of school branches in the top decile of Medellín's test-score distribution. The component is the standardized density, retaining zeros where no such school is present. It measures access to schools with high observed performance.

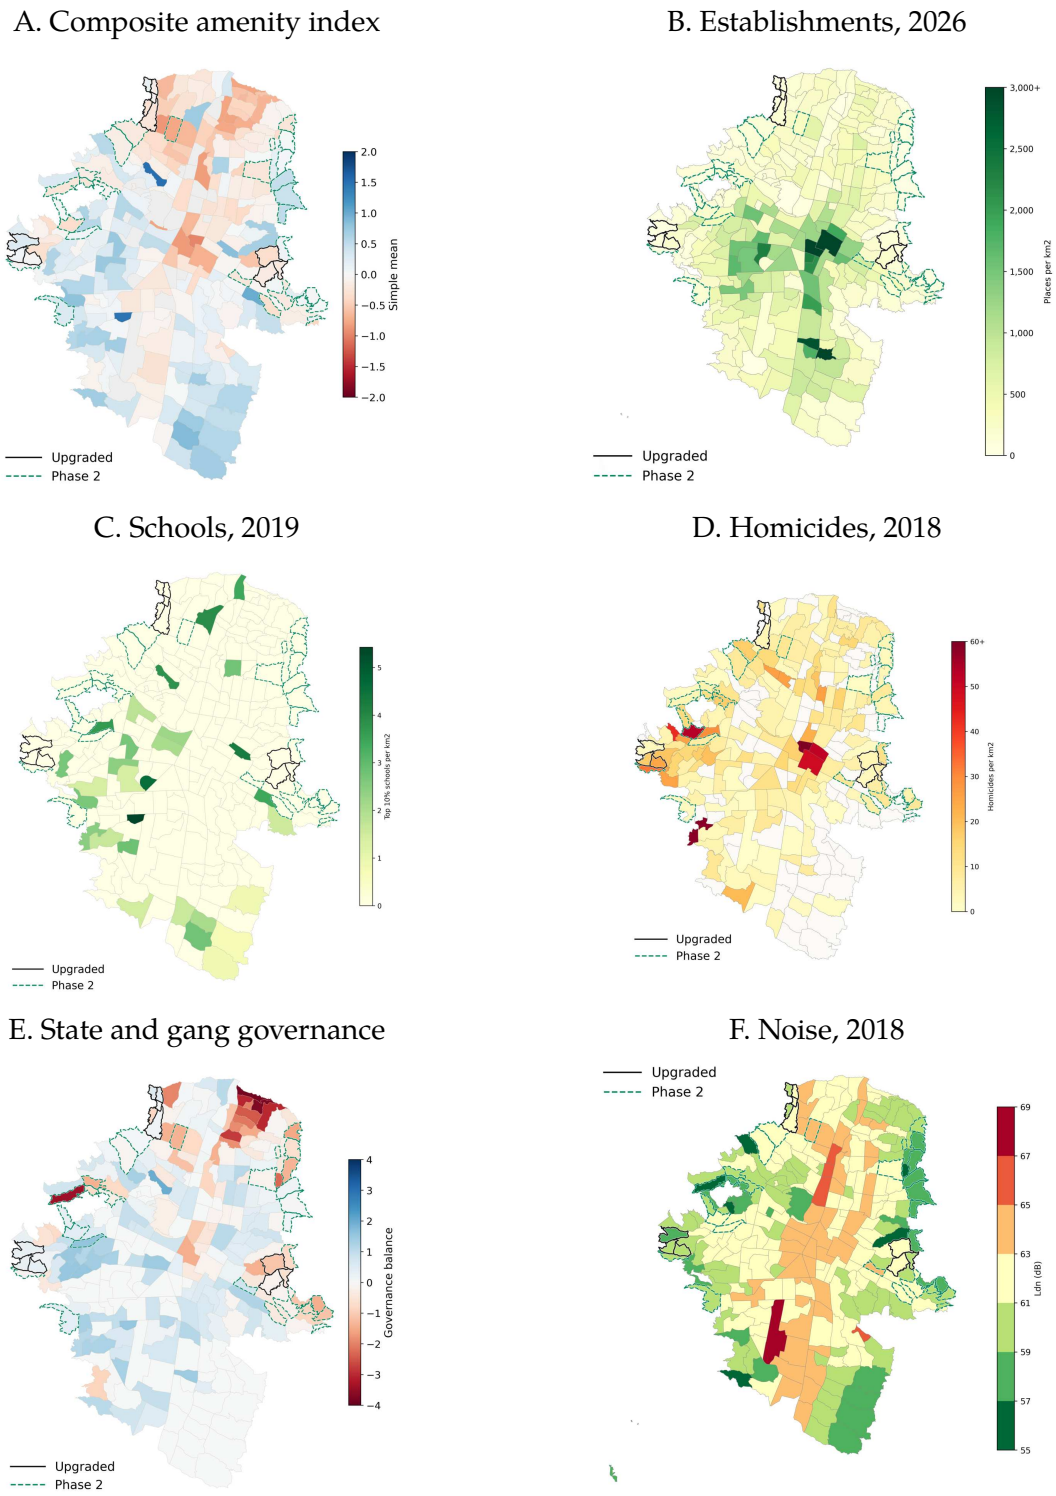
Safety. I use MEData/SISC records for 2018 covering homicides and seven other measures: domestic-violence protection requests, nonfatal injury, commercial theft, personal theft, car theft, motorcycle theft, and extortion. I standardize and reverse each density, average the eight scores, and standardize the average. Higher values indicate less recorded crime.

State and gang governance. I use Blattman et al. (2025) replication data linked to the study's 2019 boundaries. The component combines stronger state governance relative to gang governance with lower security-fee payments. I transfer these measures using area weights and assign a neutral score of zero where coverage is below 50% or either measure is unavailable.

Quietness. I use AMVA's 2018 noise maps. I combine daytime and nighttime noise levels for each neighborhood, applying a 10 dB nighttime penalty to account for greater sensitivity to noise at night (FAA, 2022). I standardize this measure and reverse its sign, so higher values indicate quieter neighborhoods.

[9] Principal components analysis selects weights to maximize explained variance. Here, the first component assigns negative weights to establishments, schools, and governance. I prefer equal weights because they better reflect the characteristics of neighborhoods with more amenities, including access to services and schools, greater safety, and less noise.

Figure 30: Composite index and sources of neighborhood amenities



Notes: A averages all five components. B shows total establishment density; the index averages ten standardized category densities. D shows homicides, one of eight safety inputs. E includes neutral governance scores where coverage is limited. F shows noise before sign reversal. Color scales differ. Solid outlines indicate Phase I; dashed outlines indicate Phase II.

Sources: Overture Places; ICFES and Medellín school locations; MEData/SISC; Blattman et al. replication data; AMVA.

D.7 Climate projections and counterfactual construction

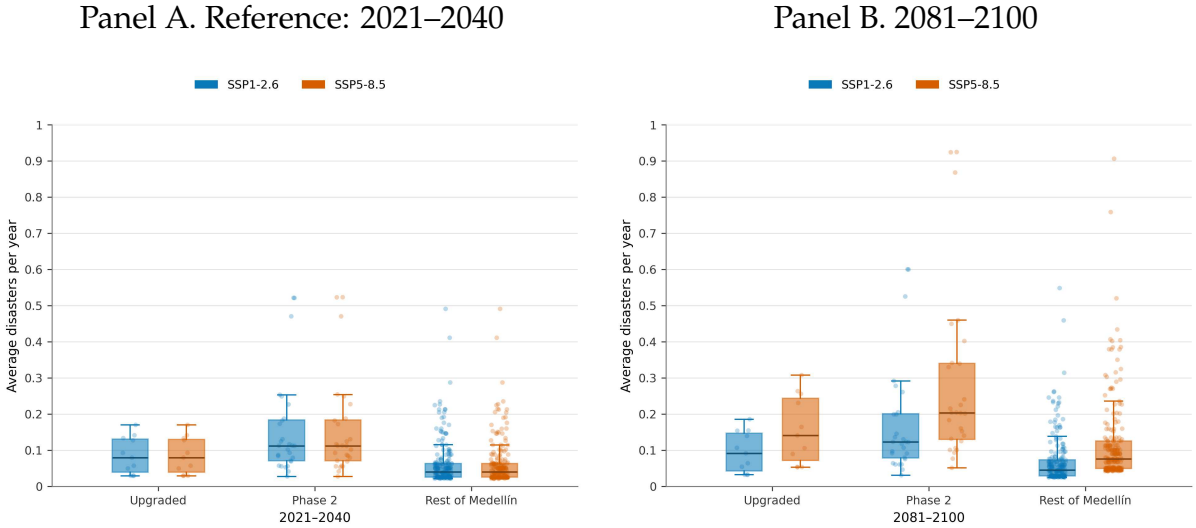
The precipitation inputs are the Colombian meteorological agency's (IDEAM) 2021–2100 projections for SSP1-2.6, SSP2-4.5, SSP3-7.0, and SSP5-8.5 (IDEAM, 2024). I combine projected heavy-rain exposure with the historical disaster rate per qualifying rainfall window, estimated over 1994–2023, to obtain annual neighborhood disaster-frequency predictions. In upgraded neighborhoods, these rates reflect conditions during and after Phase I. I use empirical-Bayes estimates of these rates (Marshall, 1991). The reported horizon is 2081–2100, using the corresponding SSP's modeled 2021–2040 frequency as its reference.

Following the delta-change approach described in IPCC (2021), I add the change in predicted frequency to the disaster frequency in the calibrated economy:

$$D^{\text{future}}_{i,m,s,h} = D^{\text{current}}_{i} + \left(\widehat{D}_{i,m,s,h} - \widehat{D}_{i,m,s,2021\text{–}2040} \right), \tag{32}$$

where i denotes the neighborhood, m the disaster-prediction method, s the SSP, and h the future horizon. The parenthesized term is an annual-frequency increment. The 2021–2040 reference is modeled, not an observed historical baseline. This applies the delta-change principle to a derived disaster measure.

Figure 31: Projected annual disaster frequency by SSP



Notes: Empirical-Bayes predictions of annual flood-plus-landslide counts by neighborhood, on a linear scale. Panels A and B show the modeled 2021–2040 reference and 2081–2100 predictions. Their neighborhood-level difference enters equation (32). Dots denote neighborhoods; boxes show the interquartile range and median.

Table 37: Projected citywide welfare changes, 2081–2100

Method	SSP	Closed city		Open city	
		Maintained	50% loss	Maintained	50% loss
Empirical Bayes	SSP1-2.6	-0.697	-1.347	-0.308	-0.530
Empirical Bayes	SSP2-4.5	-1.637	-2.280	-0.748	-0.968
Empirical Bayes	SSP3-7.0	-2.856	-3.490	-1.284	-1.503
Empirical Bayes	SSP5-8.5	-4.857	-5.476	-2.158	-2.374

Notes: Percent change in citywide expected utility relative to the observed Phase I equilibrium. Phase I remains in place in every scenario shown. Each SSP is shown under closed- and open-city assumptions. Low- and high-skill contributions use fixed observed skill weights.