Paying for integration. The neighborhood impacts of a mixed-income housing demand voucher: evidence from Chile

Ignacio Aravena-Gonzalez*

July 31, 2026

Click here for the latest version

Abstract.
Despite the prominence of social integration in urban policy agendas, credible evidence on whether housing policies reshape neighborhood composition remains limited. This paper studies Chile's DS-19, a homeownership demand voucher that introduces an "integration bonus" through additional subsidies in mixed-income developments, most of which are located in peripheral areas. Using geocoded project data matched to Chile's Population Censuses between 2002 and 2024, I estimate difference-in-differences and propensity score matching models to address non-random development location and households preferences and to separately identify direct effects on treated census tracts and spillovers on nearby tracts defined by contiguity. The results indicate that DS-19 shifts neighborhood population dynamics by increasing the presence of higher-educated household heads in treated tracts, with substantially larger effects in metropolitan areas. Spillovers are concentrated in immediately adjacent tracts and are likewise stronger in large cities. Complementary evidence shows that DS-19 projects generate discrete changes in the built environment and local housing supply within treated blocks. Consistent with affordability pressures being more binding in large cities, these patterns suggest that pecuniary incentives paired with improved housing quality can induce localized demographic change and limited spillovers beyond the immediate surroundings.

Keywords: Affordable Housing, Housing Policy, Population Dynamics, Social integration.

*London School of Economics and Political Science. Corresponding author: i.a.aravena-gonzalez@lse.ac.uk

Acknowledgments: I gratefully acknowledge ANID (Agencia Nacional de Investigación y Desarrollo Chile) - Scholarship ID 72210405 for the support on this research. I thank Christian Hilber and Nancy Holman for their advice and guidance. Lastly, I thank Universidad de Los Andes, the CIT of Universidad Adolfo Ibáñez, and Transsa for their help with different datasets.

1 Introduction

Residential segregation and spatial inequality are central concerns for urban policy. A large body of research links exposure to distressed neighborhoods with lower economic mobility, worse health outcomes, and weaker access to jobs and opportunity (Chetty et al., 2016; Li et al., 2013; Miltenburg et al., 2018; Schwartz et al., 2006). These concerns have intensified as segregation has increased in many cities in Europe and the United States (Musterd et al., 2017; Watson, 2009), and as housing affordability pressures have amplified sorting along socioeconomic lines (Hilber & Mense, 2021). Yet, despite sustained policy interest in "social mix", identifying whether housing policies can causally reshape neighborhood composition, and whether any changes extend beyond the treated area, remains challenging (Cheshire, 2007).

This paper studies Chile's DS-19 program, a nationwide affordable homeownership policy introduced in 2015 that targets social inclusion as an explicit goal. Chile is a useful setting because its housing model relies heavily on market provision and demand-side subsidies, with a limited role for rental assistance compared to canonical social-housing systems (CNDU, 2014; Ducci, 1997). Historically, the interaction between fixed subsidies and developers' land-cost incentives has contributed to the peripheral concentration of low-income housing, often in areas with weaker amenities and public infrastructure (Agostini et al., 2016; Gil, 2018; Sabatini & Brain, 2008). DS-19 operates primarily as a one-time purchase subsidy for owner-occupiers, combined with project-level rules that require income mixing and minimum standards for the built environment. Beyond supporting the production of new units under regulated price and project requirements, it provides an additional monetary incentive for eligible buyers to purchase in mixed-income projects. In practice, the policy reduces the effective downpayment burden for middle-income households conditional on choosing a mixed-income development; hence, it pays for integration.

The program therefore speaks directly to a question that is less studied in the housing-policy literature. Classic relocation policies such as Moving to Opportunity (MTO) and HOPE-type interventions primarily aim to move disadvantaged households into higher-opportunity neighborhoods, often in rental markets and with additional supports (Boston, 2005; Chetty et al., 2016; Jeffrey Kling et al., 2007). By contrast, DS-19 asks whether a pecuniary incentive tied to homeownership and new construction can attract more advantaged households into historically disadvantaged areas, even when projects remain largely peripheral. This mechanism is closely related to evidence showing that households sort on local amenities and the housing product, and that new supply can trigger endogenous neighborhood change (Bayer et al., 2007; Diamond & Gaubert, 2022; Guerrieri et al., 2013).

The broader evidence on affordable housing and neighborhood change is mixed. In the United States, the Low-Income Housing Tax Credit (LIHTC) has been associated with changes in neighborhood composition and, in some contexts, with reduced local segregation (Freedman & Mcgavock, 2015; Horn & O'Regan, 2011). New developments can also raise nearby prices and shift local demographics through the replacement of older units and the introduction of higher-quality housing services (Diamond & Mcquade, 2019; Ooi & Le, 2013; Schwartz et al., 2006; Zahirovich-Herbert & Gibler, 2014). However, effects depend on treatment intensity and local conditions; place-based revitalization policies can fail to generate meaningful in-migration when the intervention is not large enough or when local constraints bind (González-Pampillón et al., 2020). These considerations are particularly relevant in metropolitan markets where housing scarcity and supply constraints are more severe (Carozzi et al., 2020; Hilber & Mense, 2021; Hilber & Turner, 2014).

A related but often under-measured margin is spatial spillovers, where amenities bundled with new projects may affect adjacent neighborhoods and not only the treated site. Much of the literature focuses on outcomes within the treated area, even though new projects can change nearby amenities (green space, local retail, perceived safety) and thereby affect surrounding neighborhoods (Asquith et al., 2019; Baum-Snow & Marion, 2009; Pennington, 2021). A nationwide program implemented across heterogeneous markets also allows an explicit test of whether impacts differ across city types and housing-market tightness.

I combine administrative information on DS-19 projects with Census outcomes observed in four waves (2002, 2012, 2017, 2024) at the census-tract level (*zonas censales* in Spanish). Treatment is staggered: DS-19 construction expands progressively within VIS-designated tracts, which yields variation in exposure timing and intensity. I estimate difference-in-differences models for treated tracts and for neighboring tracts within 1 km, classified by contiguity order, and I define potential controls outside a 2 km buffer to reduce contamination. To address developer site selection and baseline neighborhood heterogeneity, I use propensity score matching based on pre-treatment social and economic characteristics and policy-relevant locational attributes. Given the staggered rollout, I complement two-way fixed effects estimates with recent DiD estimators designed for staggered adoption (Callaway & Sant'Anna, 2021; Sun & Abraham, 2021).

The results show that DS-19 shifts neighborhood composition in treated tracts toward more educated households, with stronger impacts in metropolitan areas. The changes are larger for tertiary-educated household heads than for secondary education, increasing the tertiary-to-secondary ratio in treated tracts. Spillovers are more muted than direct effects but are detectable in the closest neighboring tracts in metropolitan markets, consistent with localized amenity effects and endogenous neighborhood change beyond the project footprint. In addition, I document patterns in housing-pressure outcomes (e.g., crowding and co-residence) that are consistent with non-trivial supply additions and population re-sorting.

This paper contributes to the literature in four ways. First, it evaluates a homeownership-based integration incentive that differs from canonical relocation programs and from rental assistance designs, helping to isolate how a pecuniary incentive tied to new construction can affect neighborhood sorting (Boston, 2005; Chetty et al., 2016). Second, it provides within-country evidence on how policy impacts vary with local market conditions, emphasizing the mediating role of metropolitan housing scarcity and constraints (Carozzi et al., 2020; Hilber & Mense, 2021; Hilber & Turner, 2014). Third, it quantifies spillovers using tract contiguity and links composition changes to built-environment and housing-pressure outcomes, adding a mechanism-based interpretation to reduced-form demographic effects (Asquith et al., 2019; Guerrieri et al., 2013; Pennington, 2021). Lastly, it contributes to a recent body of empirical literature connecting the impacts of housing and urban policies on population dynamics (Davis et al., 2021; González-Pampillón et al., 2020; Park et al., 2025)

Two limitations guide interpretation. The analysis identifies local treatment effects rather than general equilibrium impacts. In addition, DS-19 eligibility is defined by income thresholds, while income is not consistently comparable across Census waves at the same spatial scale. I therefore use educational attainment of household heads as the primary measure of socioeconomic composition and discuss its relationship to income-based eligibility in the data section and appendix.

The remainder of the paper is organized as follows. Section 2 describes Chile's housing system and DS-19's institutional design and develops testable predictions. Section 3 describes the data and the construction of treatment, spillover, and control groups. Section 4 presents the main results and heterogeneity by metropolitan status. Section 5 examines mechanisms using built-environment and housing-pressure outcomes. Section 6 reports robustness checks. Section 7 concludes and discusses policy implications.

2 Policy Background and Theoretical Predictions

2.1 Policy background on DS-19

Affordable housing became a policy priority in Chile after the 2010 earthquake and tsunami, both because of the reconstruction needs and because housing construction was seen as a relevant countercyclical lever (**BCN2015DecretoUrbanismo**). Around the same time, the country's first national urban development policy explicitly emphasized urban inequality and residential segregation as first-order challenges (CNDU, 2014). In this context, the Ministry of Housing and Urban Development (MINVU) launched a program initially regulated as DS-116 (**BCN2015DecretoUrbanismo**), later updated and continued under DS-19 (**BCN2016DecretoUrbanismo**).[1]

A key feature of this institutional setting is that Chile's housing assistance has historically relied on demand-side subsidies that support homeownership through private developers, while rental policy plays a secondary role. DS-19 follows this tradition but adds an explicit integration component. The program aims to reduce the housing deficit and promote mixed-income developments by increasing the effective purchasing power of eligible households when they buy units inside socially integrated projects (**BCN2015DecretoUrbanismo**).

DS-19 bundles two pre-existing subsidy instruments, DS-49 for very low-income house-

[1] Throughout the paper, I refer to the program as DS-19. When using DS-116 and DS-19 interchangeably, the intent is to denote continuity in the policy logic rather than a discrete policy break.

holds and DS-01 for middle-income households, and adds an integration bonus when the purchase occurs in a DS-19 project. Importantly, this additional subsidy is introduced without raising the relevant price ceilings, so the incentive operates through a larger subsidy for a given unit price. To illustrate, a middle-income household could receive a subsidy of UF 275 for a UF 2,200 dwelling inside a socially integrated project, compared to UF 100 in a comparable non-integrated purchase.[2] Similarly, a very low-income household could receive UF 900 toward a UF 1,100 unit inside a DS-19 project, compared to UF 700 outside these developments. These magnitudes are large relative to downpayment requirements and therefore potentially salient for location choice.

DS-19 is implemented through a centrally managed competitive selection process. Developers submit projects built on their own land, and MINVU ranks proposals within each region until the regional quota is met.[3] Only selected projects are scored and publicly listed after the selection round concludes, while non-selected proposals do not receive a disclosed score. Developers observe the scoring rubric in advance, which makes the contest predictable in structure but competitive in outcomes.

The program constrains both the unit mix and minimum standards. Units are grouped into three price segments, very low-income units with a ceiling of UF 1,100, low-income units between UF 1,100 and UF 1,400, and middle-income units between UF 1,400 and UF 2,200. A typical project allocates roughly 20% of units to very low-income households, 10% to low-income households, and 70% to middle-income households, although projects can receive additional points for including more low-income units. Developments are built as-of-right on developers' land, and the program does not grant zoning or density bonuses. Minimum unit sizes are set at 47 m^2 for houses and 52 m^2 for apartments.

Project ranking depends on location and amenity criteria that vary by local context. First, projects located in municipalities above a population threshold receive additional points. Second, the scoring system assigns points for proximity to basic amenities, conditional on city size. Criteria include distance thresholds to health facilities, parks, schools, and kindergartens, as well as access to public transport and the road network. A project

[2]UF is an inflation-indexed unit of account widely used in Chilean real estate markets. Its value in Chilean pesos varies over time by design.

[3]Regional quotas vary across years, reflecting housing needs and budget constraints.

must satisfy a minimum number of these criteria to be eligible. Third, projects can obtain additional points by providing on-site or adjacent amenities, such as commercial infrastructure, extra green areas, or playgrounds.[4] Finally, projects receive points based on the share of very low-income units and, in some cases, for including low-income units beyond the baseline requirement.

During the early implementation period, DS-19 scaled rapidly. However, new developments are concentrated around each region's main metropolitan area and, within them, largely located in peripheral rings. Figure 1 illustrates the spatial pattern for Greater Santiago, Greater Valparaíso, and Greater Concepción.

Note: The figure shows the location of DS-19 developments in yellow in the three largest metropolitan areas of Chile, Santiago, Valparaíso, and Concepción.

Figure 1: Location of DS-19 developments

Although many developments are peripheral, they differ substantially in micro-location and urban form. Some projects are single-family condominiums located on previously undeveloped plots at the urban fringe. Others are multifamily buildings, often mid-rise structures, inserted into denser and more consolidated neighborhoods. Figure 2 provides examples of these two configurations. In both cases, DS-19 developments tend to coexist with an older surrounding housing stock, a pattern that motivates the empirical strategy in Section 4, which tests whether observed neighborhood change aligns with DS-19 construction rather than with broader simultaneous redevelopment.

[4]These features are common among competitive proposals, but take-up varies across projects and locations.

Note: The figure shows two typical configurations of DS-19 projects, highlighted with red dashed lines. The left panel illustrates a peripheral single-family configuration, while the right panel illustrates a denser multifamily configuration in a more consolidated area.

Figure 2: Urban configurations of DS-19 developments

2.2 Related literature and theoretical predictions

This paper sits at the intersection of three literatures. The first studies neighborhood effects and the extent to which housing assistance can change households' locational choices and, through that channel, socioeconomic outcomes. The second evaluates whether subsidized housing construction and stock renewal can shift neighborhood composition and local prices. The third examines how housing market tightness and unaffordability shape demographic dynamics and the scope for policy-induced sorting. I use these strands to motivate a set of testable predictions for DS-19.

Vouchers, location choice, and neighborhood effects. A large body of work shows that where households live affects education, income, and health, but also that moving to higher-opportunity neighborhoods is difficult even when assistance exists. Classic evidence from mobility programs such as MTO and related interventions highlights that voucher recipients face search frictions, informational barriers, and supply constraints that limit moves to substantially better neighborhoods (Chetty et al., 2016; Jeffrey Kling et al., 2007). Recent work emphasizes that the design and generosity of housing assistance matters for the geography of take-up and for the opportunity set that households can realistically access. For example, reforms that better align subsidy ceilings with within-metro rent variation can expand access to higher-rent neighborhoods, precisely because payment rules interact with local price dispersion and landlord behavior (Park et al., 2025). This evidence is consistent with the DS-19 rationale, which increases the effective housing budget for eligible buyers conditional on choosing a socially integrated project, so the policy can relax affordability constraints and change location choices, particularly where price-to-income ratios are high.

This mechanism is especially relevant in settings where unaffordability is rising and households adjust fertility, household formation, and migration decisions in response to housing costs. A homeownership-focused subsidy that effectively reduces downpayment burdens can therefore plausibly affect who enters peripheral neighborhoods, and not only whether low-income households can afford a unit.

A second literature studies how subsidized construction and place-based interventions affect neighborhood composition and prices. Programs such as LIHTC and other subsidized developments often replace underutilized land or older stock, which can raise local amenities and change the demographic mix, sometimes reducing segregation and increasing local housing values (Diamond & Mcquade, 2019; Freedman & Mcgavock, 2015; Horn & O'Regan, 2011; Ooi & Le, 2013; Schwartz et al., 2006; Zahirovich-Herbert & Gibler, 2014). Related work highlights that higher-income or more educated households may have a higher willingness to pay for renovated stock and bundled amenities, which can generate endogenous compositional change (Bayer et al., 2007; Couture & Handbury, 2023; Diamond, 2016; Guerrieri et al., 2013). However, positive neighborhood change is not automatic. If the intervention is too small relative to the baseline disamenity, or if amenities are not salient for marginal movers, effects can be limited (González-Pampillón et al., 2020).

DS-19 differs from many canonical urban renewal policies because it does not rely on zoning benefits nor public investment, yet many projects remain in peripheral locations. This feature is central to the identification strategy and to interpretation. The relevant question is not whether the policy relocates low-income households into high-opportunity areas, but whether a pecuniary incentive combined with architectural and accessibility standards can attract middle-income buyers into peripheral, historically amenity-poor neighborhoods.

A third strand emphasizes that the incidence of demand or subsidy shocks depends on local supply conditions. Where land is scarce or regulation is tighter, subsidy-driven demand can translate more strongly into price responses and sorting (Carozzi et al., 2020; Hilber &

Mense, 2021; Hilber & Turner, 2014). This motivates a key dimension of heterogeneity in my setting. DS-19 is a national policy, but metropolitan areas face higher baseline prices and more binding affordability constraints than smaller cities. If the policy relaxes a binding budget constraint for educated middle-income households, compositional effects should be strongest in metropolitan markets.

Finally, spillovers are a natural object of interest because new projects can change perceived neighborhood quality beyond the parcel, through public goods (green space, safety, retail), reputational effects, and complementarities in housing demand (Baum-Snow & Marion, 2009; Diamond & Mcquade, 2019; Rossi-Hansberg et al., 2010). Yet the sign and spatial reach are ambiguous ex ante. Spillovers may be positive if amenities are non-rival and accessible to nearby residents, but limited if projects are gated or if amenities are largely internal. Spillovers may also attenuate quickly with distance, which motivates an empirical design based on contiguity rings rather than simple distance bins.

Guided by the discussion above, I derive the following predictions, which directly map into the empirical strategy and outcomes studied in Sections 4–6:

- Treated blocks should exhibit an increase in housing units and new construction intensity relative to the rest of the tract. Absent this built-environment change, any demographic shift would be harder to attribute to DS-19 rather than to re-sorting within a fixed stock.

- Treated census tracts should experience an increase in the share of more educated household heads relative to counterfactual neighborhoods. If the policy relaxes affordability constraints primarily for middle-income buyers, the relative increase should be larger for college graduates than for households with high-school education.

- Direct compositional effects should be stronger in metropolitan areas where baseline unaffordability is more binding and where the outside option of finding comparable units at similar prices is more limited. In smaller cities, the same subsidy increment may be inframarginal, leading to weaker sorting responses.

- Adjacent tracts should show smaller, attenuated effects relative to treated tracts. Effects should decay with contiguity order, with first-order neighbors more likely to benefit from accessible amenities or improved perceptions than second-order neighbors.

These predictions clarify why the empirical analysis separates treated tracts from contiguity rings, emphasizes heterogeneity by city size, and jointly studies demographic outcomes and built-environment measures.

3 Data

I build a tract-level panel by combining geocoded information on DS-19 developments with repeated cross-sections from the Chilean Population Census and additional spatial and housing-market datasets. The core unit of analysis is the census tract, equivalent to *Zona Censal*. I harmonize census geographies into a common tract identifier and aggregate block-level sources to the tract level when needed.

The main outcomes measure changes in neighborhood socio-economic composition using educational attainment. To study mechanisms, I also assemble tract-level measures of the built environment, local amenities, and housing-market conditions, including administrative cadaster records and appraisal data.

3.1 Data sources

The treatment data come from geolocated DS-19 project records that report project geometry, the number of subsidized units, and implementation timing. Because several developments can be phased within the same tract, I aggregate projects to the tract level and construct cohort indicators based on whether a tract had at least one DS-19 project completed before each census wave used in the analysis. To ensure correct timing, I cross-check projects using administrative completion information and keep only those plausibly delivered before the corresponding census fieldwork.

Socio-demographic outcomes are constructed using the 2002, 2012, 2017, and 2024 Population Censuses. The Census provides tract-level counts for individuals and households, including educational attainment, household size, and basic demographic structure. The

2012 Census was invalidated due to omission errors. Nevertheless, it has been used in academic work and remains informative for the spatial distribution of population characteristics at scale (Garreton et al., 2020). I therefore treat 2012 as an intermediate pre-period and place greater weight on design-based checks of pre-trends and placebo comparisons reported in Section **??**.

To proxy neighborhood quality at a fine spatial scale, I use the *IBT–UAI* index, a cross-section of normalized block-level indicators assembled from official sources. It includes measures related to safety, environment, access to public services, and housing conditions. I also build tract-level distances to key public goods used by the policy scoring system, including green areas, health facilities, schools, kindergartens, transit stops, and road access. These layers are geocoded from OpenStreetMap and processed in QGIS, then aggregated to the tract level using consistent rules.

I complement Census information with administrative records from the Chilean Internal Revenue Service (Servicio de Impuestos Internos (SII)). The cadaster provides block-level land-use and built-form measures that I use to construct tract-level proxies for available land, commercial share, and housing-stock intensity. To track changes in the built environment around the first program wave, I use cadaster vintages spanning the pre-policy period and the years immediately preceding the 2017 Census, which allows me to document differential changes in the housing stock within treated tracts. To proxy local price dynamics, I use SII housing sales.

3.2 Main outcomes

The primary outcomes are based on educational attainment, measured as tract-level shares and log ratios across education groups. Education is a standard proxy for permanent income and socio-economic status that is less sensitive to short-run income shocks (Garreton et al., 2020; González-Pampillón et al., 2020; Ruiz-Tagle & López M, 2014). The analysis focuses on low- and middle-income segments where education maps tightly into long-run earnings potential. To validate this interpretation in the Chilean context, Appendix 8.4 documents a strong and monotonic education–income gradient for household heads using CASEN 2017 under the complex survey design.

I classify household heads into two categories that match the Census definitions and the program's target population. The groups are secondary education, and university degree or higher. This partition captures salient socio-economic distinctions in Chile and aligns with evidence on social segmentation by education (Bargsted et al., 2020) as well as income shown in Appendix 8.4.

Given strong heterogeneity in housing-market tightness across space, I classify municipalities into city types and separately identify the three main metropolitan areas. This stratification is used to estimate heterogeneous effects and to interpret differences in sorting and spillovers across market environments.

3.3 Sample construction and spatial exposure

The empirical design distinguishes between treated tracts, adjacent tracts, and potential controls. Treated tracts are those containing at least one DS-19 development completed before a given census wave. Adjacent tracts are those within a 1 km neighborhood around treated tracts and are further classified by contiguity degree to capture decay in exposure. To reduce contamination, potential controls are drawn from tracts beyond a 2 km buffer from any treated tract and are selected using propensity score matching based on pre-treatment tract characteristics and the policy's location criteria. The spatial structure is shown in Figure 3

3.4 Descriptive statistics

Table 1 reports pre-treatment balance across treated tracts (VIS), adjacent tracts (ADJ), and the pool of potential controls (CTRL) before propensity score matching. The table shows substantial differences across groups in market conditions, city type, demographic composition, and baseline socio-economic composition. These gaps motivate the matching strategy and the emphasis on heterogeneity by city type in the main analysis.

First, DS-19 tracts are disproportionately located outside the largest metropolitan areas. Only 29% of VIS tracts are in metropolitan areas, compared with 60% in CTRL, while ADJ tracts lie in between (49%). This implies that a naive comparison between treated tracts

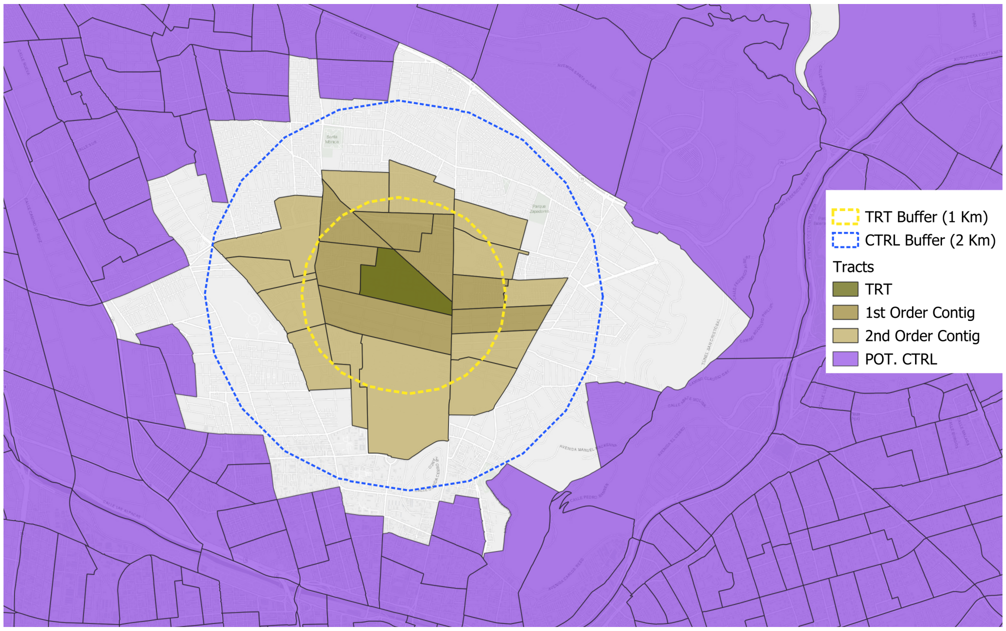

Note: The figure shows the strategy to define treatment and control groups. The green tract received a DS-19 development, the adjacent tracts are in decaying colours to denote different contiguity degrees. Then, after a 2 Km. buffer, all the remaining tracts are potential controls. Source: Google Street View.

Figure 3: Treated, adjacent and potential control tracts.

and the full set of potential controls partly compares different urban systems rather than treated and untreated neighborhoods within comparable markets.

Second, treated and adjacent tracts operate in markedly cheaper housing markets. Mean sale prices and UF/m^2 are substantially lower in VIS and ADJ than in CTRL (SMD around -0.5 to -0.6). Built-form differences are also stark: VIS tracts have a much larger share of residential plots (and fewer non-residential plots), consistent with DS-19 targeting predominantly residential, lower-price segments. These differences reinforce the concern that a ring-based control group without reweighting mixes neighborhoods with different land-use and market fundamentals.

Third, baseline population composition differs sharply. VIS tracts have a larger share of children (0–17) and a smaller share of older residents (60+), relative to CTRL. Educational composition also differs substantially: the share of individuals with secondary education or less is higher in VIS/ADJ (about 0.78) than in CTRL (0.675), while the share with technical education is lower in VIS/ADJ (about 0.22) than in CTRL (0.325). In addition, pre-policy changes between 2002 and 2012 indicate faster growth in VIS tracts in population, households, and educational composition. However, overall trends also show an increase in access to higher education, which is likely influenced by the expansion of student loans and various policies aimed at increasing access to higher education through student loans (Gaentzsch & Zapata-Román, 2020).

While these differences do not mechanically invalidate a difference-in-differences design, they make the case for (i) explicitly balancing observables through matching and (ii) conducting pre-trend tests and placebo exercises reported in Section **??**.

Overall, Table 1 shows that treated tracts are systematically different places: less metropolitan, cheaper, more residential, and with distinct demographic and baseline educational profiles. This is precisely why the empirical strategy combines a careful definition of exposure groups with propensity score matching and robustness checks that assess pre-trends and alternative specifications.

On another note, as shown in Figure 4, housing-unit trends also diverge after the policy. Treated blocks display a discrete increase in the number of dwellings relative to surrounding areas, consistent with DS-19 developments materially changing the local housing stock rather than reflecting only re-sorting within a fixed built environment.

Table 1: Pre-treatment balance before matching

	VIS	**ADJ**	**CTRL**	**Δ VIS–CTRL**	**Δ ADJ–CTRL**
Share Access Kindergarten	0.733	0.836	0.577	0.36	0.64
Share Access Bus Stop	0.609	0.774	0.637	-0.06	0.32
Share Access Parks	0.358	0.504	0.485	-0.29	0.04
Share Access Health Center	0.955	0.978	0.945	0.05	0.20
Share Access Schools	0.806	0.904	0.876	-0.25	0.11
Population	3202.513	3200.910	2718.586	0.36	0.36
Households	960.419	937.323	870.354	0.23	0.18
Metropolitan Area	0.292	0.492	0.603	-0.66	-0.22
Intermediate City	0.327	0.287	0.067	0.69	0.60
Mean Sale Price	1393.454	1378.462	2307.986	-0.54	-0.54
Mean Construction Area (m^2)	63.926	64.934	73.871	-0.39	-0.33
Mean Price (UF/m^2)	21.411	21.182	28.856	-0.63	-0.64
Share Plots w/ Housing	0.958	0.941	0.822	0.83	0.68
Share Plots w/ Commerce	0.036	0.049	0.056	-0.28	-0.09
pct_mujer	0.518	0.519	0.518	0.02	0.07
pct_0_17	0.287	0.262	0.236	0.87	0.44
Share Age 18–44	0.415	0.400	0.426	-0.16	-0.37
pct_45_59	0.177	0.187	0.181	-0.12	0.18
Share Age 60+	0.122	0.150	0.157	-0.58	-0.11
Share Female HH Head	0.358	0.382	0.380	-0.30	0.03
p_HSmax_pers	0.780	0.782	0.675	0.58	0.57
p_Tmin_pers	0.220	0.218	0.325	-0.58	-0.57
Δ Log Pop. (2002–12)	0.047	0.012	0.015	0.40	-0.05
Δ Log HH (2002–12)	0.056	0.020	0.027	0.37	-0.10
Δ Log Sec. Educ. (2002–12)	0.017	0.011	0.011	0.42	0.02
Δ Log Tech. Educ. (2002–12)	0.034	0.024	0.026	0.30	-0.13
Δ Log Univ. Educ. (2002–12)	0.011	0.004	0.004	0.39	-0.05

Notes: The table reports group means before matching. SMD denotes standardized mean differences relative to CTRL. Amenity variables are measured as tract-level shares satisfying the DS-19 access thresholds. Pre-period changes are computed between the 2002 and 2012 Census waves.

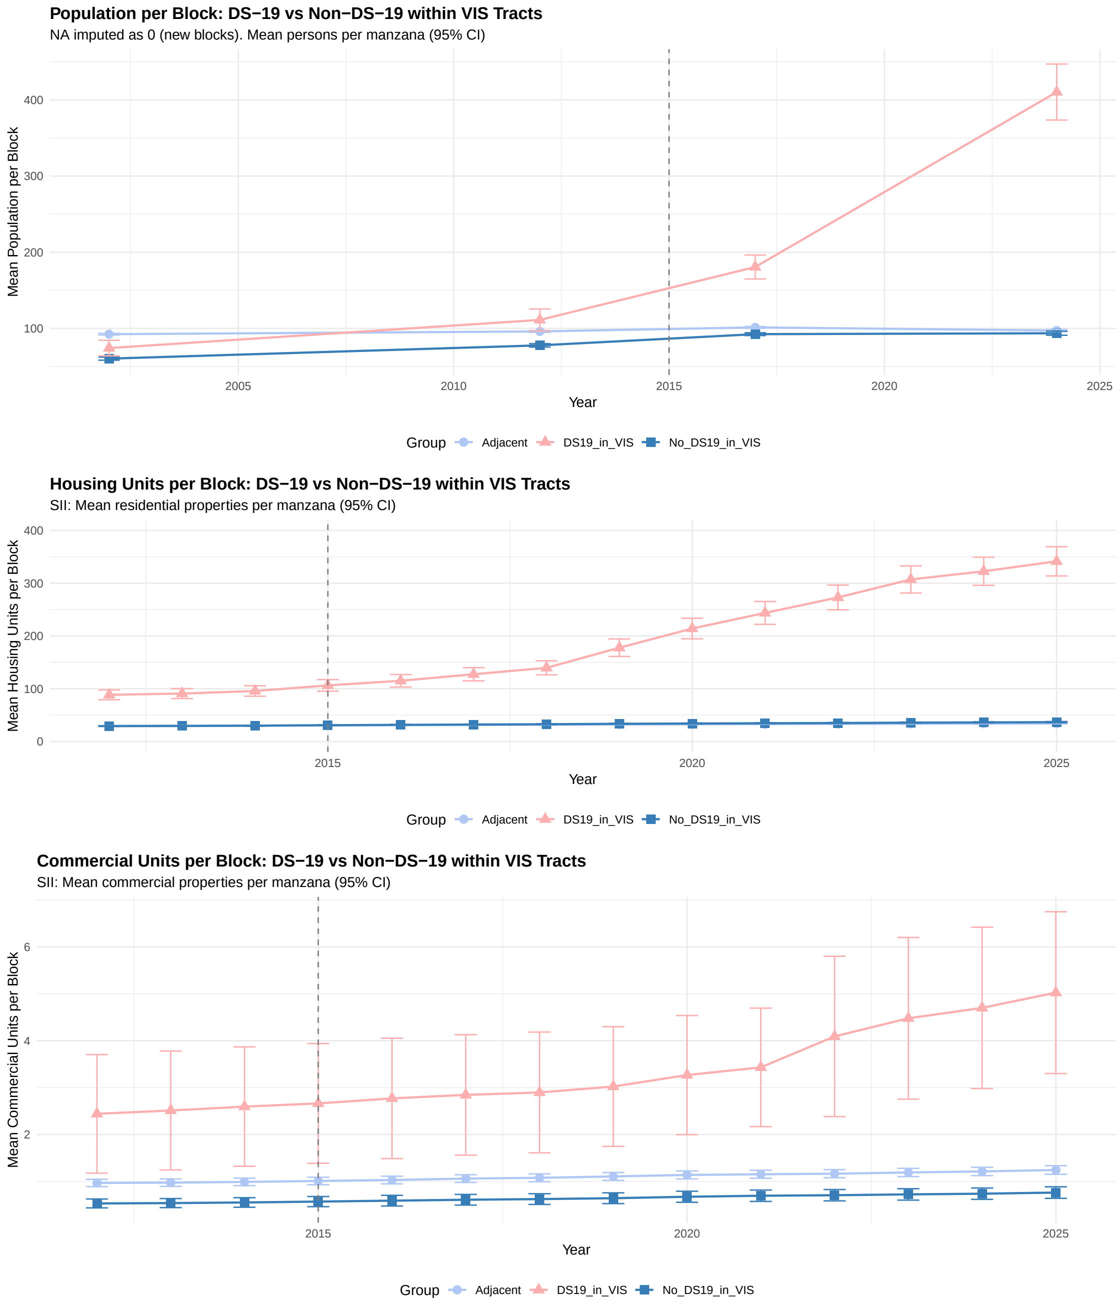

Figure 4: First stage mean values per block by group

4 Main Results

4.1 DS-19 as a local housing supply shock within treated tracts

Before turning to tract-level matched DiD on demographic composition, I document that DS-19 produces a sharp, localized change in the built environment within treated tracts. The purpose of this section is not to estimate tract-wide effects, but to show that the housing expansion in treated tracts is concentrated in the specific blocks where DS-19 projects are built. This evidence motivates interpreting subsequent demographic changes as responses to a supply shock rather than to broader tract-level redevelopment.

I use block-level administrative records from Chile's property tax cadaster, maintained by the Servicio de Impuestos Internos (SII), which is equivalent to the IRS in the US, to measure residential and commercial stock over time. I estimate a difference-in-differences model that compares blocks hosting a DS-19 project to other blocks within the same treated tract, using 2015 as the start of the policy period. In addition, I implement a placebo comparison that contrasts non-treated blocks inside treated tracts to blocks in adjacent tracts. This placebo is not intended to be a perfect counterfactual for treated blocks. Its role is to test how much housing and population change occurs outside the project footprint.

I estimate

$$Y_{bct} = \beta\big(Treat_b \times Post_{t\geq 2015}\big) + \alpha_b + \delta_{c\times t} + \varepsilon_{bct}, \tag{1}$$

where b indexes blocks (manzanas), c indexes municipalities (comunas), and t indexes time. The outcome Y_{bct} is alternatively the log of one plus residential stock, the log of one plus commercial stock, and the corresponding outcomes in levels. I also estimate the same model using total block population as the outcome.[5] In the within-tract specification, $Treat_b$ equals one for blocks that host a DS-19 project and zero for other blocks in the same treated tract. In the placebo specification, $Treat_b$ equals one for non-treated blocks in treated tracts and zero for blocks in adjacent tracts. The block fixed effect α_b absorbs all time-invariant block characteristics. The municipality-by-year fixed effects $\delta_{c\times t}$ absorb municipality-specific

[5] Due to data limitations, block-level census outcomes are available only for total population. Disaggregated outcomes, such as education composition, are measured at the tract level.

shocks and trends. Standard errors are clustered at the census tract level to account for spatial correlation in DS-19 project placement.

Table 2: First-stage evidence on housing stock and population at the block level

	Housing stock outcomes				Population outcomes	
	log(1+Res.)	log(1+Comm.)	Res. count	Comm. count	log(1+Pop.)	Pop. level
Panel A. DS-19 blocks vs other blocks within treated tracts						
DiD	0.7314***	0.2016***	120.6***	0.9582***	1.264***	180.4***
	(0.0408)	(0.0144)	(7.556)	(0.2345)	(0.1000)	(14.91)
Observations	205,047	205,047	205,047	205,047	71,812	71,812
R^2	0.9096	0.9422	0.7887	0.9550	0.5763	0.5555
Panel B. Placebo non-treated blocks in treated tracts vs adjacent tracts						
DiD	0.0228***	-0.0075**	1.102**	-0.1901***	0.7259***	20.01***
	(0.0085)	(0.0038)	(0.4501)	(0.0735)	(0.0579)	(1.805)
Observations	918,216	918,216	918,216	918,216	283,880	283,880
R^2	0.9188	0.9589	0.8944	0.9443	0.5507	0.6791

Notes. The unit of observation is the census block. In Panel A, treated blocks host a DS-19 project and controls are other blocks within the same treated tracts. In Panel B, treated blocks are non-treated blocks within treated tracts and controls are blocks in adjacent tracts. $Post_{t\geq 2015}$ equals one from 2015 onward. All regressions include block fixed effects and municipality-by-year fixed effects. Standard errors are clustered at the census tract level. In the placebo control group, the pre-policy mean is 29.94 residential properties and 0.98 commercial properties per block. For population, the pre-policy control mean is 69.05 within treated tracts and 94.24 in adjacent tracts. Residential and commercial stock are measured as the number of SII cadaster units classified by use. Significance levels are $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 2 shows that housing growth is concentrated in DS-19 blocks. Within treated tracts, blocks hosting a DS-19 project gain about 121 additional residential units as measured in the SII cadaster. Commercial stock also increases, but at a much smaller scale, by roughly one additional commercial unit per block. For reference, the mean numbers of residential and commercial units per control block in the pre-policy period are approximately 30 and 1, respectively. Population rises sharply in the same blocks, by about 180 residents per block in levels.

The placebo comparison shows much smaller changes outside the project footprint. Non-treated blocks inside treated tracts increase by about one residential unit relative to adjacent tracts and do not show a comparable increase in commercial stock. This pattern indicates that the bulk of housing growth within treated tracts is driven by DS-19 projects rather than by tract-wide redevelopment. The placebo comparison also shows population growth outside DS-19 blocks, but at a much smaller magnitude, around 20 residents per block, consistent with only a small amount of net housing expansion relative to the scale of DS-19 developments.

Taken together, the stock and population evidence support interpreting DS-19 as a local housing supply shock that mechanically brings new residents into the project blocks. This motivates the tract-level matched DiD that studies demographic composition and educational upgrading.

4.2 The impact of DS-19 on population dynamics

The previous section documents that DS-19 generates a sharp and localized increase in residential stock within treated tracts, concentrated in the blocks that host DS-19 projects. This section examines whether the local supply shock translates into changes in population levels and demographic composition at the census tract level. The key question is whether tracts that receive DS-19 projects experience differential demographic change relative to comparable tracts that do not receive such projects, and whether any effects extend to nearby tracts.

Throughout, I use Chilean census tracts as the primary neighborhood unit. These tracts are small administrative areas designed to be relatively homogeneous and are comparable to census tracts in other settings. The analysis uses repeated census waves, which in Chile correspond to decennial censuses and the most recent 2024 wave. Given that DS-19 begins in 2015, I treat 2017 and 2024 as post-policy waves.

Outcomes. To measure population dynamics, I use log counts as the primary outcomes. Log counts capture net changes in group size and avoid an ambiguity inherent to shares. A tract-level share can fall either because a group leaves or because other groups arrive, and both scenarios can generate similar share movements. I therefore study the log number of residents by educational attainment as a measure of inflows and outflows in levels. Overall, I separate between inhabitants with a high-school degree or tertiary education[6]. I complement these outcomes with composition and diversity measures that provide direct evidence on changes in mixing.

The baseline specification is

$$Y_{it} = \beta\big(TRT_i \times Post_t\big) + \alpha_i + \lambda_t + \varepsilon_{it}, \tag{2}$$

where i indexes census tracts and t indexes census waves. TRT_i is an indicator for treated tracts and, in separate estimations, for tracts adjacent to treated tracts. $Post_t$ equals one for post-policy waves, namely 2017 and 2024. The tract fixed effect α_i absorbs time-invariant tract characteristics, and the wave fixed effects λ_t absorb nationwide shocks and secular trends. In preferred specifications, I additionally allow for differential trends by city type using wave-by-city-type fixed effects. Standard errors are clustered at the tract level. Because DS-19 adoption is staggered, I use TWFE as a benchmark and report staggered-adoption estimators robust to heterogeneous treatment effects (Sun and Abraham; Callaway and Sant'Anna) in the robustness section.

4.2.1 Propensity score matching

As discussed in Section 3, DS-19 projects are not randomly located in space. Tracts that receive DS-19 developments differ from other tracts within the same city along dimensions that plausibly matter for both developer site selection and household location choices. These include baseline density, land availability, access to amenities, prior demographic trajectories, and local housing market conditions (Diamond & Mcquade, 2019; Freedman & Mcgavock, 2015; González-Pampillón et al., 2020; Won, 2022). Households also exhibit heterogeneous preferences over neighborhood attributes, which can shape sorting responses (Ahlfeldt et al., 2017; Baum-Snow & Marion, 2009).

To improve comparability, I construct matched control groups using propensity score matching based on pre-treatment covariates measured in 2012 and changes between 2002

[6]In Chile, there are 2 types of tertiary education: technicians that study two years in a local college and bachelor/professional degrees. Altought they are different in nature, their dynamics change significantly in the first decade of the 2000's since many local colleges started teaching professional degrees, making them inconsistent throuout time.

and 2012. The matching model includes policy-relevant location attributes, baseline levels, and pre-trends that proxy for developer and household selection. In particular, DS-19 assigns additional scoring to projects located in municipalities above a population threshold, which motivates including municipal size among the matching covariates.[7]. Furthermore, I also include mean outcomes to proxy for local real estate markets, such as housing prices, average unit size, and average UF per square meter transacted price.

I implement matching separately for three designs. The first compares tracts that host DS-19 projects to matched control tracts. The second and third focus on spillovers by constructing separate matched samples for first-order and second-order adjacent tracts, where adjacency is defined by tract contiguity. After matching, I estimate weighted difference-in-differences models using the matching weights.

The matched samples achieve strong balance, as summarized in Table 3. Across all three designs, the maximum post-match standardized mean difference is below 0.10 and the average post-match imbalance is approximately 0.03. Detailed balance diagnostics, including covariate density and love plots, are reported in Appendix **??**.

Table 3: Propensity score matching balance summary

	DS-19 tracts (VIS)	Adjacent 1 (RING1)	Adjacent 2 (RING2)
N treated (post-match)	321	970	454
N control (post-match)	514	598	502
ESS control	458.98	222.41	263.54
Max $\lvert SMD \rvert$ (post)	0.08	0.09	0.08
Mean $\lvert SMD \rvert$ (post)	0.03	0.03	0.03
N $\lvert SMD \rvert > 0.10$ (post)	0	0	0

Notes. The table reports balance diagnostics for the matched samples used in the tract-level DiD analysis. SMD denotes standardized mean differences. ESS is the effective sample size of the control group induced by matching weights. Detailed balance plots and covariate densities are reported in Appendix **??**.

[7] From a policy perspective, municipal population matters because DS-19 projects receive additional points when located in municipalities above 40,000 inhabitants.

4.2.2 Main effects on education composition

Table 4 reports tract-level DiD estimates for secondary-educated residents, tertiary-educated residents, and their ratio, separately by spatial exposure group and treatment cohort. The main pattern is that compositional change is concentrated in DS-19 tracts. Effects are larger for tracts treated by 2017 and smaller for tracts first treated between 2017 and 2024, consistent with longer exposure time for earlier cohorts and later-stage implementation for the 2024 cohort. Within DS-19 tracts, tertiary-educated inflows respond more strongly than secondary-educated inflows, which mechanically increases the tertiary-to-secondary ratio.

Table 4: DS-19 and tract-level population composition by cohort and spatial exposure

	DS-19 tracts (VIS)	Ring 1	Ring 2
Panel A. Cohort treated by 2017			
log(Secondary)	0.2803***	0.0227	0.0684
	(0.0562)	(0.0423)	(0.0447)
log(Tertiary)	0.4276***	0.0615	0.1171**
	(0.0654)	(0.0485)	(0.0535)
log(Tertiary / Secondary)	0.1473***	0.0388*	0.0487*
	(0.0331)	(0.0204)	(0.0295)
Panel B. Cohort first treated 2017–2024			
log(Secondary)	0.1207***	-0.0213	-0.0416*
	(0.0427)	(0.0309)	(0.0229)
log(Tertiary)	0.1979***	-0.0295	-0.0211
	(0.0504)	(0.0358)	(0.0246)
log(Tertiary / Secondary)	0.0772***	-0.0083	0.0206
	(0.0238)	(0.0153)	(0.0191)
Observations	3,340	6,272	3,824
Tract fixed effects	Yes	Yes	Yes
Wave fixed effects	Yes	Yes	Yes

Notes. The unit of observation is the census tract. Outcomes are log counts of residents by education group and the log ratio between tertiary and secondary counts. Panel A reports effects for tracts treated by 2017, while Panel B reports effects for tracts first treated between 2017 and 2024. All models are estimated on matched samples using propensity score weights. Standard errors are clustered at the tract level. Pre-policy matched-control means are 2,260 (Secondary) and 518 (Tertiary) for VIS, 2,223 and 588 for Ring 1, and 2,158 and 700 for Ring 2. The corresponding pre-policy means of the tertiary-to-secondary count ratio are 0.282, 0.339, and 0.433. Significance levels are $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Within DS-19 tracts, the estimates point to sizeable population inflows and educational upgrading. For the 2017 cohort, secondary-educated residents increase by 0.28 log points while tertiary-educated residents increase by 0.43 log points, implying changes of roughly 32% and 53%, respectively.[8] The log ratio rises by 0.15 log points, indicating that tertiary inflows outpace secondary inflows. Effects remain positive and precisely estimated for the 2024 cohort, but are smaller in magnitude, around 13% for secondary and 22% for tertiary, with an 8% increase in the ratio.

To anchor magnitudes in levels, I translate the log coefficients into implied changes using the pre-policy mean of the matched control group. In DS-19 tracts, the 2017 cohort implies an increase of about 32.4% in secondary-educated residents (roughly +731 residents per tract) and 53.4% in tertiary-educated residents (roughly +276 residents per tract), relative to a pre-policy mean of 2,260 and 518 residents, respectively. For the 2024 cohort, the implied magnitudes are smaller, about 12.8% (+290) for secondary and 21.9% (+113) for tertiary. These level translations indicate that the estimated upgrading reflects large absolute inflows, not only proportional shifts.

To assess dynamics and validate the identifying assumptions, Figure 5 plots event-study coefficients by cohort. In the baseline TWFE specification, some outcomes show borderline evidence of differential pre-trends under joint tests of pre-treatment leads. Once I add baseline 2012 controls and allow for city-type-specific trends, pre-treatment coefficients are close to zero and the joint pre-trends tests do not reject for any outcome. This matters in Chile because demographic trajectories differ sharply across metropolitan areas, intermediate cities, and the rest of the country. I further report staggered-adoption estimators robust to treatment-effect heterogeneity in the robustness section.

Spillovers are modest and not systematic. In first-order adjacent tracts, estimates are small and generally close to zero across cohorts. In second-order adjacent tracts, the 2017 cohort shows limited positive effects for tertiary education and the ratio, but these patterns do not persist for the 2024 cohort. Overall, the evidence indicates that DS-19 primarily induces change within treated tracts, with at most small and unstable re-sorting effects in nearby areas.

[8] Percent changes computed as $\exp(\beta) - 1$.

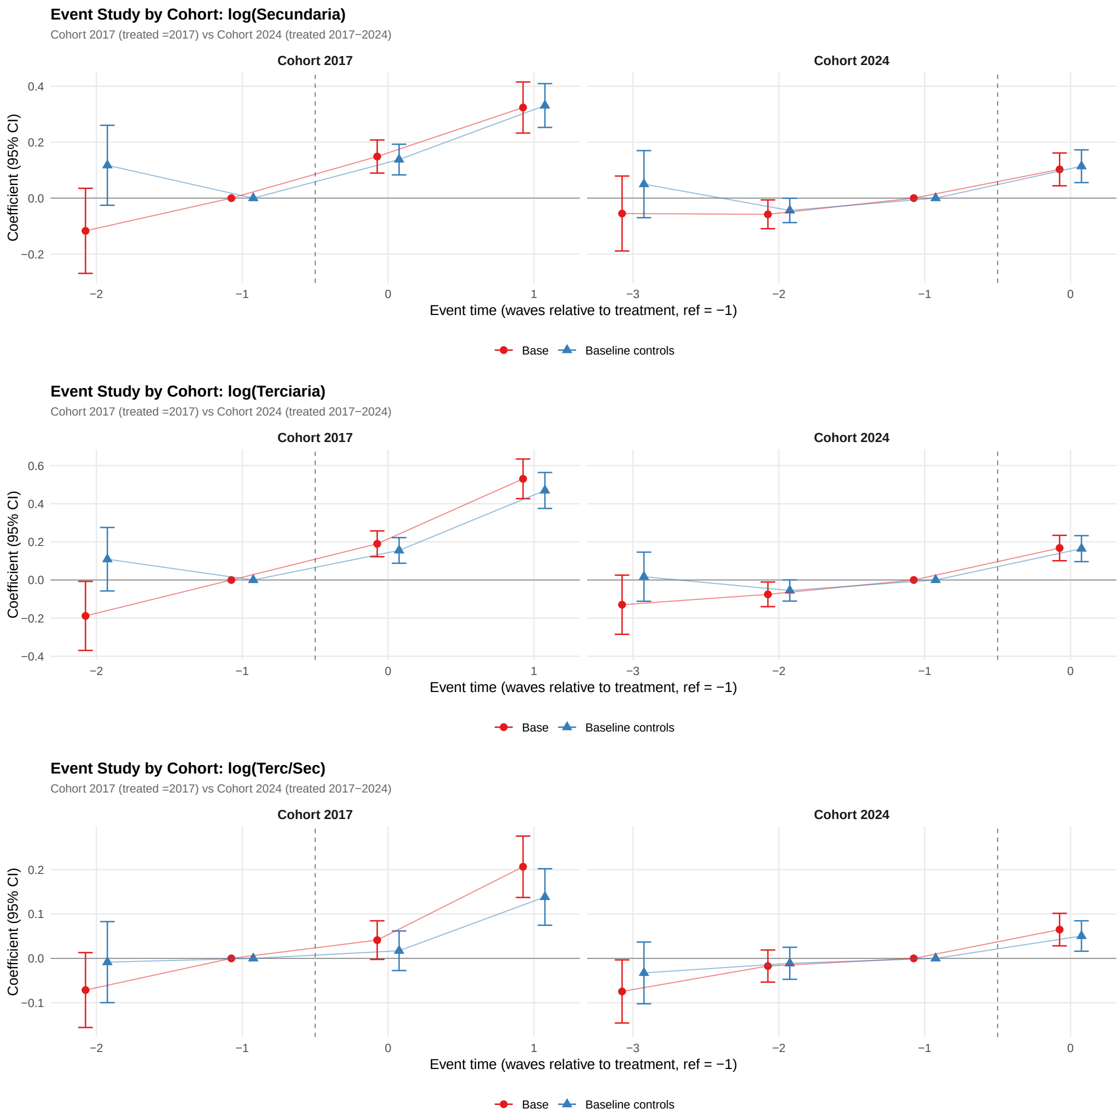

Figure 5: Event studies by cohort. Joint tests of pre-treatment coefficients in the baseline-controlled TWFE specification yield p-values of 0.66, 0.90, and 0.43 for log(Secondary), log(Tertiary), and log(Ratio), respectively.

5 Heterogeneity and mechanisms

Chile features stark differences in demographic trajectories across metropolitan areas, intermediate cities, and smaller urban markets. Because DS-19 projects were deployed nationwide and developers face different land and demand conditions across these city types, I next examine heterogeneity in the main effects and provide evidence on mechanisms that link the built-environment shock to population re-sorting.

5.1 Heterogeneity by metropolitan status

Chile's urban system is highly concentrated. A large share of the population lives in three metropolitan areas, Greater Santiago, Greater Valparaíso, and Greater Concepción, where housing demand is stronger and land-market constraints are more binding than in smaller urban markets. I therefore test whether DS-19 effects differ between metropolitan and non-metropolitan tracts. I define $Metro_i$ as an indicator for tracts located in these three metropolitan areas. The matched VIS sample spans tracts in Greater Santiago, Greater Valparaíso, Greater Concepción, intermediate cities, and the rest of the country, providing variation in baseline demographic trends and local housing-market tightness.

To estimate heterogeneity by cohort, I interact the cohort-specific DiD indicators with $Metro_i$ and estimate

$$\begin{aligned} Y_{it} = {} & \beta_{17} DiD_{it}^{2017} + \delta_{17}\big(DiD_{it}^{2017} \times Metro_i\big) + \beta_{24} DiD_{it}^{2024} \\ & + \delta_{24}\big(DiD_{it}^{2024} \times Metro_i\big) + \alpha_i + \lambda_{t\times c(i)} + X_{i,2012} \times \lambda_t + \varepsilon_{it}. \end{aligned} \tag{3}$$

where α_i are tract fixed effects, $\lambda_{t\times c(i)}$ are wave-by-city-type fixed effects, and $X_{i,2012} \times \lambda_t$ denotes baseline 2012 covariates interacted with census waves. All models are estimated on matched samples using propensity-score weights, and standard errors are clustered at the tract level.

Figure 6 reports the estimates by exposure group and cohort. In DS-19 tracts, effects are concentrated in metropolitan areas. Outside metro areas, point estimates are small and statistically indistinguishable from zero for both education groups and for the tertiary-to-secondary ratio. In contrast, metropolitan DS-19 tracts exhibit large increases in both secondary- and tertiary-educated residents, together with a clear rise in the tertiary-to-secondary ratio. This pattern holds for both cohorts, with larger magnitudes for the cohort treated by 2017, consistent with longer exposure time.

In adjacent tracts, the pattern differs sharply between metropolitan and non-metropolitan areas. In metropolitan areas, Ring 1 and Ring 2 tracts display modest positive effects in the levels of both education groups in several specifications, suggesting that demographic change is not fully confined to the treated tract and may extend to nearby areas. Outside metropolitan areas, adjacent tracts tend to exhibit negative effects in both education groups, consistent with within-city reallocation toward the DS-19 tract rather than net inflows at the broader neighborhood level. Notably, the ratio outcome in adjacent rings remains close to zero in both metro and non-metro areas, indicating that these changes mainly operate through levels rather than systematic shifts in composition outside the treated tract.

Taken together, the results suggest that in smaller urban markets DS-19 primarily reshuffles residents across nearby tracts, whereas in metropolitan areas the program is associated with net inflows that can extend modestly beyond the treated tract.

5.1.1 Effects by housing intensity

DS-19 construction rolled out progressively within VIS-designated tracts. By 2017, 94 had received DS-19 projects, while 321 had experienced any construction before 2024. Among tracts with construction by 2017, the average project size was about 59 units. By 2024, all VIS tracts had received DS-19 projects and average intensity reached about 382 units; many oh them were consecutive phases of the same development. This staggered rollout, combined with substantial variation in project size, allows me to test for dose-response effects of the housing-supply shock.

I extend the baseline DiD by interacting cohort-specific treatment indicators with the cumulative number of DS-19 units built in each tract by each wave (scaled by 100):

$$\begin{aligned} Y_{it} = {} & \beta_{17} DiD_{it}^{2017} + \gamma_{17}\big(DiD_{it}^{2017} \times Units_{it}/100\big) + \beta_{24} DiD_{it}^{2024} + \gamma_{24}\big(DiD_{it}^{2024} \times Units_{it}/100\big) \\ & + \alpha_i + \lambda_{t \times c(i)} + X_{i,2012} \times \lambda_t + \varepsilon_{it}, \end{aligned} \tag{4}$$

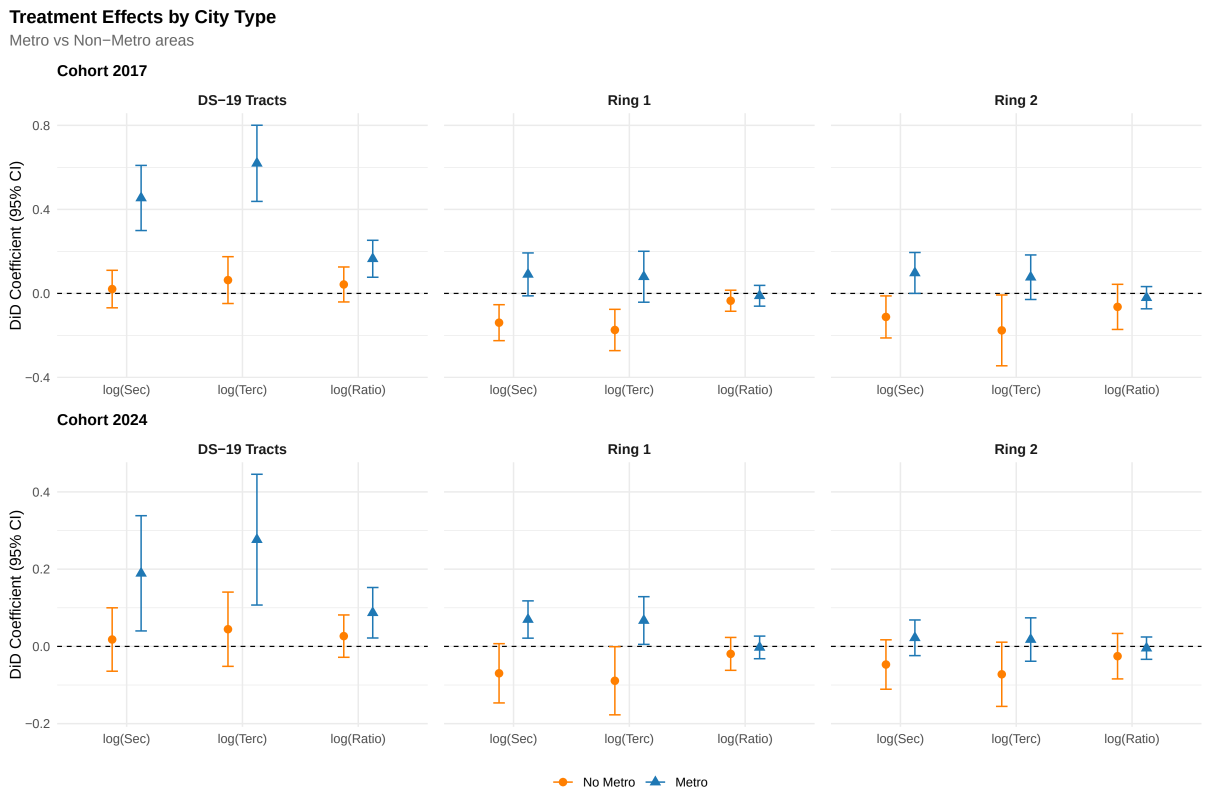

Figure 6: Treatment effects by metropolitan status and cohort. The figure reports tract-level DiD coefficients for log(Secondary), log(Tertiary), and log(Ratio), separately by exposure group and cohort. Models use propensity-score weights and include tract fixed effects, wave-by-city-type fixed effects, and baseline 2012 covariates interacted with census waves. Standard errors are clustered at the tract level.

where $Units_{it}$ is the cumulative number of DS-19 housing units built in tract i by wave t. I report implied effects evaluated at 0, 100, 200, and 300 units.

Figure 7 shows a clear dose-response pattern within DS-19 tracts. Effects increase approximately linearly with project size for both cohorts, and the slope is steeper for the cohort treated by 2017, consistent with longer exposure time. In contrast, estimates for Ring 1 and Ring 2 are comparatively flat and imprecise, indicating that intensity mainly matters within treated tracts rather than through strong re-sorting in adjacent areas. Overall, each additional 100 units increases estimated effects by about 0.02 and 0.07 log points, depending on the outcome

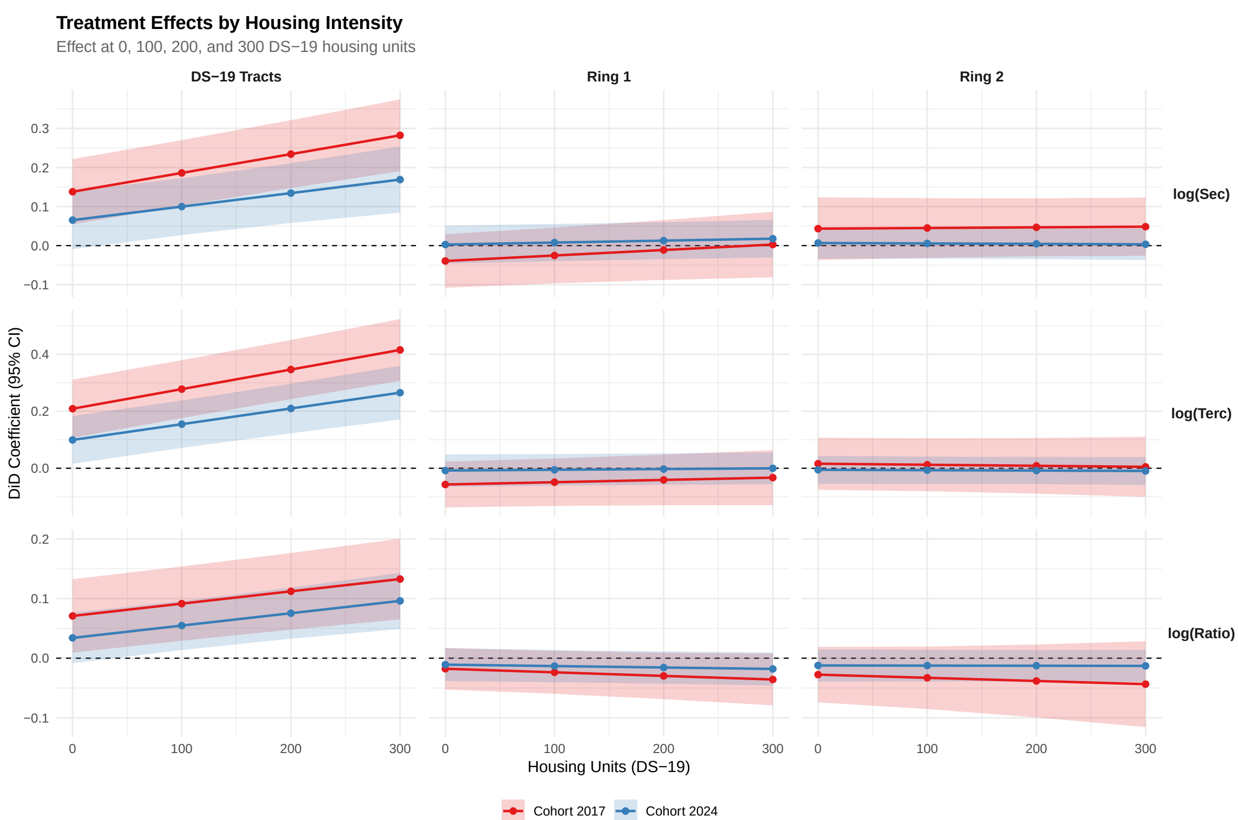

Figure 7: Treatment effects by housing intensity, evaluated at 0, 100, 200, and 300 DS-19 units.

5.1.2 Housing pressure and household composition

A key mechanism behind the education upgrading documented above is that DS-19 expands the local housing stock in locations facing different degrees of market tightness. If new supply relaxes binding constraints, one would expect reductions in housing stress indicators such as overcrowding and doubled-up households, together with demographic shifts consistent with household formation and residential re-sorting. I study five tract-level outcomes measured consistently across census waves: the overcrowding rate, the share of doubled-up households (allegados), the share of female-headed households, and the shares of residents aged 18–44 and 60+.

Figure 8 reports estimates for DS-19 tracts by treatment cohort and metropolitan status. Two patterns stand out. First, metropolitan DS-19 tracts experience a clear demographic re-sorting toward working-age residents. The share aged 18–44 rises by around 2–3 percentage points, while the share aged 60+ declines sharply, especially for the 2017 cohort. Female-headed households also increase, with larger effects in metropolitan areas, consistent with DS-19 attracting younger households and households targeted by the social-housing quota.

Second, housing pressure indicators respond differently across city types. In metropolitan areas, overcrowding declines for the early cohort, consistent with new supply easing housing constraints in tight markets. Outside metropolitan areas, overcrowding does not fall and slightly increases on average, suggesting that in smaller markets DS-19 operates primarily through local reallocation and compositional changes rather than through a strong relaxation of constraints. Effects on doubled-up households are negative for the early cohort and attenuate for the later cohort.

To keep the presentation focused on mechanisms where construction occurs, Figure 8 concentrates on DS-19 tracts. Estimates for adjacent rings and the full regression tables are reported in Appendix **??**.

6 Robustness checks

This section evaluates whether the main findings are sensitive to identification and specification choices. Two features of the setting motivate these checks. First, DS-19 exposure is staggered across space, with an early cohort already built by 2017 and a late cohort that accumulates exposure by 2024. Second, treatment effects can vary across cohorts and over time. Under staggered adoption and heterogeneous effects, conventional two-way fixed ef-

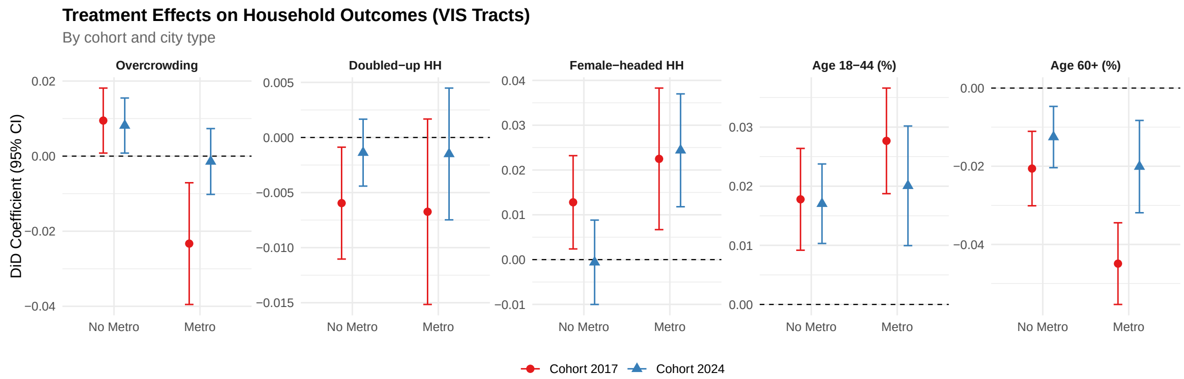

Figure 8: Housing pressure and household composition in DS-19 tracts (VIS), by cohort and metropolitan status. The figure plots tract-level DiD estimates for overcrowding, doubled-up households (allegados), the share of female-headed households, and the shares of residents aged 18–44 and 60+. Estimates are reported separately for the cohort treated by 2017 and the cohort first treated between 2017 and 2024, and by whether the tract belongs to a metropolitan area. All models are estimated on matched samples using propensity score weights, include tract fixed effects and wave-by-city-type fixed effects, and interact baseline 2012 covariates with wave dummies. Standard errors are clustered at the tract level; whiskers show 95% confidence intervals.

fects estimators may mix cohort-specific comparisons in a way that is difficult to interpret. I therefore complement the baseline specification with alternative estimators designed for staggered treatment timing and I probe the stability of the results under alternative weighting and control strategies.

6.1 Alternative estimators under staggered adoption

I re-estimate the dynamic effects using two approaches that are robust to staggered timing with heterogeneous treatment effects, the Sun–Abraham interaction-weighted event-study and the Callaway–Sant'Anna group-time average treatment effects. I implement both using the same baseline controls and fixed effects structure as in the main specification.

Figure 9 compares the event-study paths for the matched VIS sample across the baseline TWFE, Sun–Abraham, and Callaway–Sant'Anna estimators. The three methods deliver consistent patterns. Pre-treatment coefficients remain close to zero and post-treatment effects are positive and similar in magnitude across estimators for the three outcomes. This comparison supports the interpretation that the baseline estimates are not driven by estimatorspecific artifacts.

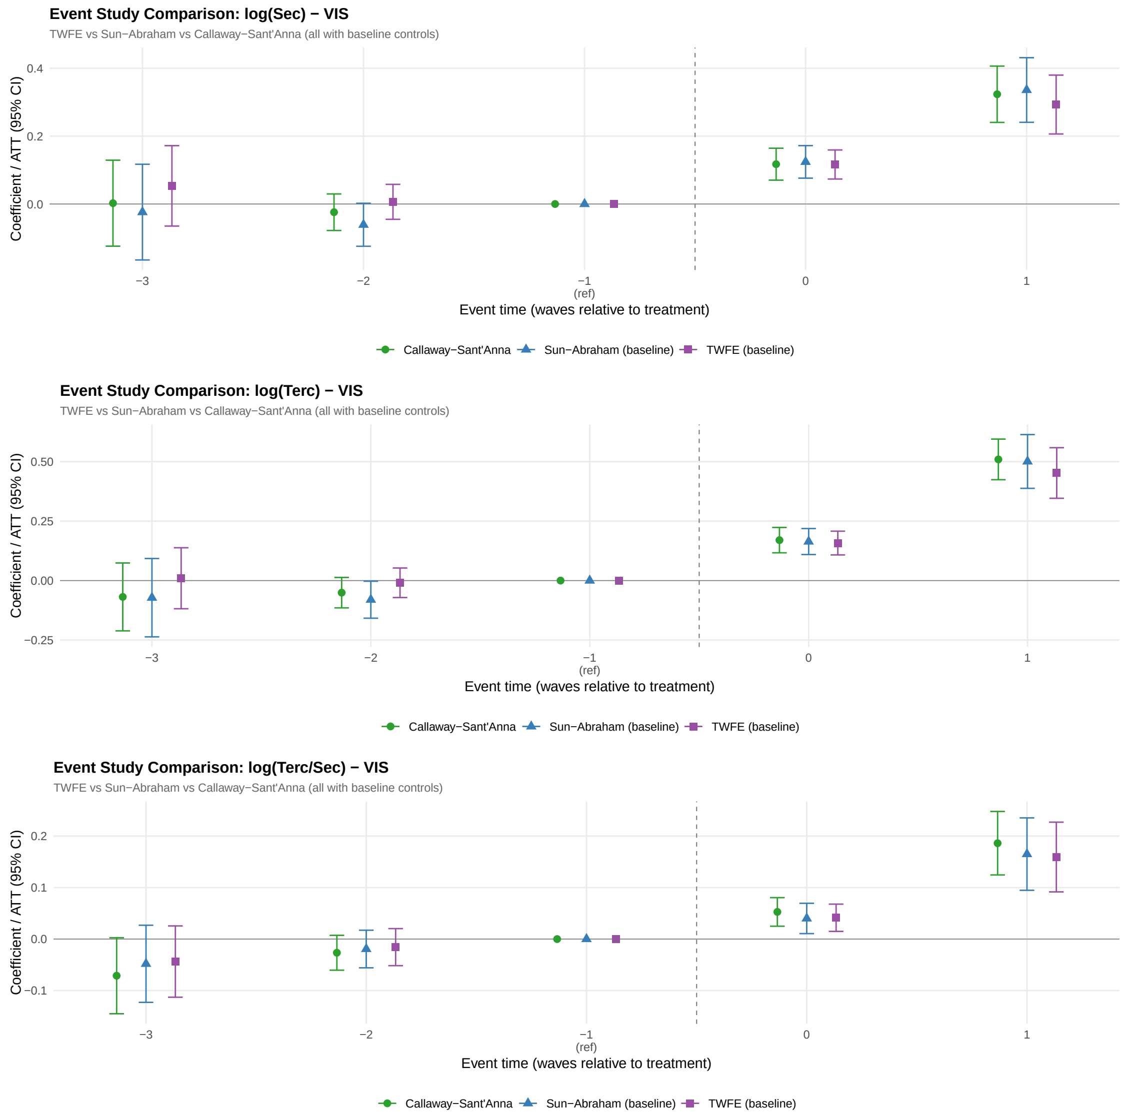

Figure 9: TWFE vs Sun–Abraham vs Callaway–Sant'Anna event studies in VIS tracts

Note. All series use the same baseline controls and fixed effects as in the main specification. Event time is measured in census waves relative to the cohort-specific treatment wave.

6.2 Pre-trends tests

As a complementary diagnostic, I test for differential pre-trends by conducting joint F-tests on the set of pre-treatment lead coefficients. Table A.1 reports the F-statistics and p-values for the baseline TWFE specification and the Sun–Abraham event-study. Across panels and outcomes, the null of zero pre-treatment leads is not rejected at conventional levels, which is consistent with parallel trends in the pre-policy period.

6.3 Placebo using the 2002–2012 pre-policy window

I also implement a placebo exercise that restricts the sample to the 2002 and 2012 waves and re-estimates the treatment specification in a period with no DS-19 construction. Table A.2 shows that placebo effects are generally small and statistically insignificant. One specification yields a marginally significant placebo estimate for secondary education in VIS tracts. This is not the main outcome of interest and it is not mirrored in tertiary education or in the tertiary-to-secondary ratio. I keep this result in the appendix to transparently document the diagnostic.

6.4 Alternative weighting and alternative specifications

Finally, I evaluate whether results depend on weighting choices and on the inclusion of baseline controls and city-trend controls. First, I compare the baseline PSM weighting to unweighted estimates and to unweighted estimates with baseline controls. Second, I estimate specifications that remove city-trend controls and specifications that remove baseline controls. Tables A.3 and A.4 show that the signs and relative magnitudes are stable. The primary interpretation is unchanged, with larger effects for tertiary education than for secondary education in VIS tracts and modest spillover patterns in adjacent tracts.

Taken together, these robustness checks support the baseline interpretation while making clear which diagnostic margins are weaker, particularly for secondary education in one placebo specification.

7 Conclusion and Policy Implications

This paper evaluates whether a large affordable-housing program in Chile, DS-19, alters neighborhood population dynamics beyond its direct contribution to the housing stock. DS-19 combines a pecuniary incentive to private developers with location-based eligibility and an inclusionary quota, and it was deployed nationwide starting in 2015. Using census waves spanning 2002–2024, I show that DS-19 generates a sharp and spatially concentrated increase in the residential stock within eligible tracts. I then test whether this localized supply shock translates into tract-level changes in population levels and composition relative to comparable tracts that did not receive DS-19 projects.

The central finding is that demographic change is concentrated in DS-19 tracts and is strongest for the earlier treatment cohort. In tracts treated by 2017, both secondary- and tertiary-educated populations rise substantially, with a stronger response among tertiary residents, which increases the tertiary-to-secondary ratio. Effects are positive but smaller for the cohort first treated between 2017 and 2024, consistent with shorter exposure time and a later phase of the program. Translating log effects into levels using pre-policy matched-control means indicates that the estimated upgrading reflects large absolute inflows, not only proportional shifts. These patterns support the interpretation that DS-19 changed who lives in treated neighborhoods, at least in the short run.

A second key result is the strong heterogeneity by metropolitan status and by spatial exposure. In DS-19 tracts, effects are close to zero outside large metropolitan areas, whereas metropolitan tracts exhibit large increases in both education groups and a clear rise in the tertiary-to-secondary ratio. This difference is consistent with Chile's housing-market geography. Greater Santiago, Greater Valparaíso, and Greater Concepción concentrate a large share of national population and housing scarcity, and they feature markedly different baseline demographic trajectories than intermediate cities and smaller urban markets. In such tighter environments, DS-19 appears to relax binding affordability and access constraints for middle-income households, leading to sizable re-sorting into treated tracts.

Spillovers to nearby tracts are present but nuanced and depend on city size. In metropolitan areas, nearby tracts exhibit small positive responses, while in smaller cities adjacent areas often move in the opposite direction. This pattern is consistent with two mechanisms that can coexist. First, treated tracts may draw households from nearby areas in smaller markets, producing a local reallocation with limited net inflows at the broader neighborhood scale. Second, in metropolitan areas DS-19 may generate broader neighborhood upgrading forces that extend beyond tract boundaries, potentially through changes in perceived neighborhood quality, new local demand, or amenity investment responding to the new resident mix. The evidence therefore points to predominantly within-tract change on average, with modest and context-dependent spillovers rather than a uniform gradient of effects across space.

Results on household structure and housing pressures help close the narrative on mechanisms and interpretation. DS-19 tracts experience systematic shifts in demographic composition toward working-age residents and away from older cohorts, alongside changes in housing-stress indicators such as overcrowding and co-residence. The direction and magnitude of these changes again differ between metropolitan and non-metropolitan settings, which is consistent with the idea that the same policy can operate through different margins depending on local market tightness, available land, and baseline demographic trends. Taken together, the findings suggest that DS-19 operates as a place-based policy that changes both the built environment and the sorting of households across neighborhoods, especially in Chile's largest urban labor markets.

These results have several policy implications. First, the effectiveness of demand-side or developer-side incentives for social mixing depends on local market conditions. In settings with binding affordability constraints, DS-19 can induce large compositional change in treated neighborhoods, while impacts are substantially smaller in less constrained markets. Second, the program's rollout speed and project scale matter. Longer exposure and higher unit intensity are associated with larger effects, which implies that implementation capacity and project size are central design levers, not merely administrative details. Third, the presence of context-dependent spillovers highlights that evaluating social-mixing policies only within treated tracts may miss meaningful redistribution across nearby neighborhoods, particularly in metropolitan areas. This strengthens the case for monitoring broader neighborhood effects and coordinating housing incentives with complementary policies on transport access, services, and land-use regulation.

This paper also has clear limitations that guide future research. The design estimates localized, short-run effects on population composition and cannot capture general-equilibrium adjustments across the city, nor can it directly separate in-migration from within-tract demographic change. Administrative links to mobility histories, housing tenure, and labor-market outcomes would allow sharper tests of whether observed re-sorting translates into sustained integration or economic mobility. A longer post-treatment horizon is also essential to assess persistence, displacement, and cohort maturation. Finally, future work could combine the tract-level evidence with project-level data on timing and unit mix to further clarify which program features drive the largest integration gains and under which market environments these gains are most likely to materialize.

References

Agostini, C. A., Hojman, D., Román, A., & Valenzuela, L. (2016). Segregación residencial de ingresos en el Gran Santiago, 1992-2002: una estimación robusta. *EURE (Santiago)*, *42*(127), 159–184. https://doi.org/10.4067/S0250-71612016000300007

Ahlfeldt, G. M., Moeller, K., Waights, S., & Wendland, N. (2017). Game of Zones: The Political Economy of Conservation Areas. *Economic Journal*, *127*(605), F421–F445. https://doi.org/10.1111/ECOJ.12454

Asquith, B. J., Mast, E., & Reed, D. (2019, December). *Supply Shock Versus Demand Shock: The Local Effects of New Housing in Low-Income Areas* (tech. rep.). W.E. Upjohn Institute. https://doi.org/10.17848/wp19-316

Bargsted, M., Espinoza, V., & Plaza, A. (2020). Pautas de homofilia en Chile. *Papers. Revista de Sociologia*, *105*(4), 583–612. https://doi.org/10.5565/rev/papers.2617

Baum-Snow, N., & Marion, J. (2009). The effects of low income housing tax credit developments on neighborhoods. *Journal of Public Economics*, *93*(5-6), 654–666. https://doi.org/10.1016/J.JPUBECO.2009.01.001

Bayer, P., Ferreira, F., & McMillan, R. (2007). A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy*, *115*(4), 588–638. https://doi.org/10.1086/522381

Boston, T. D. (2005). The Effects of Revitalization on Public Housing Residents: A Case Study of the Atlanta Housing Authority. *Journal of the American Planning Association*, *71*(4), 393–407. https://doi.org/10.1080/01944360508976710

Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230. https://doi.org/10.1016/J.JECONOM.2020.12.001

Carozzi, F., Hilber, C., & Yu, X. (2020). On the Economic Impacts of Mortgage Credit Expansion Policies: Evidence from Help to Buy.

Cheshire, P. (2007). *Segregated neighbourhoods and mixed communities A critical analysis* (tech. rep.). Joseph Rowntree Foundation.

Chetty, R., Hendren, N., & Katz, L. (2016). The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment. *American Economic Review*, *106*(4), 855–902. https://doi.org/10.1257/AER.20150572

CNDU. (2014). *Política Nacional de Desarrollo Urbano. Ciudades Sustentables y Calidad de Vida Hacia una Nueva Política Urbana para Chile* (tech. rep.). Comisión Nacional de Desarrollo Urbano.

Couture, V., & Handbury, J. (2023). Neighborhood Change, Gentrification, and the Urbanization of College Graduates. *Journal of Economic Perspectives*, *37*(2), 29–52. https://doi.org/10.1257/JEP.37.2.29

Davis, M. A., Gregory, J., Hartley, D. A., & Tan, K. T. K. (2021). Neighborhood effects and housing vouchers. *Quantitative Economics*, *12*(4), 1307–1346. https://doi.org/10.3982/QE1664;REQUESTEDJOURNAL:JOURNAL:17597331;WGROUP:STRING:PUBLICATION

Diamond, R. (2016). The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000. *American Economic Review*, *106*(3), 479–524. https://doi.org/10.1257/AER.20131706

Diamond, R., & Gaubert, C. (2022). Spatial Sorting and Inequality. *Annual Review of Economics*, *14*(Volume 14, 2022), 795–819. https://doi.org/10.1146/ANNUREV-ECONOMICS-051420-110839/CITE/REFWORKS

Diamond, R., & Mcquade, T. (2019). *Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development* (tech. rep.). https://web.archive.org/web

Ducci, M. E. (1997). *Chile: el lado obscuro de una política de vivienda exitosa* (tech. rep.).

Freedman, M., & Mcgavock, T. (2015). Low-Income Housing Development, Poverty Concentration, and Neighborhood Inequality. *Journal of Policy Analysis and Management*, *34*(4), 805–834. https://doi.org/10.1002/PAM.21856

Gaentzsch, A., & Zapata-Román, G. (2020). *Climbing the ladder: Determinants of access to and returns from higher education in Chile and Peru*, United Nations Research Institute for Social Development (UNRISD). www.unrisd.org

Garreton, M., Basauri, A., & Valenzuela, L. (2020). Exploring the correlation between city size and residential segregation: comparing Chilean cities with spatially unbiased indexes. *Environment and Urbanization*, *32*(2), 569–588. https://doi.org/10.1177/0956247820918983/ASSET/IMAGES/LARGE/10.1177{_}0956247820918983-FIG5.JPEG

Gil, D. (2018, May). Law and Inclusive Urban Development: Lessons from Chile's Enabling Markets Housing Policy Regime. https://papers.ssrn.com/abstract=3927296

González-Pampillón, N., Jofre-Monseny, J., & Viladecans-Marsal, E. (2020). Can urban renewal policies reverse neighborhood ethnic dynamics? *Journal of Economic Geography*, *20*(2), 419–457. https://doi.org/10.1093/JEG/LBZ011

Guerrieri, V., Hartley, D., & Hurst, E. (2013). Endogenous gentrification and housing price dynamics. *Journal of Public Economics*, *100*, 45–60. https://doi.org/10.1016/j.jpubeco.2013.02.001

Hilber, C., & Mense, A. (2021). *Why have house prices risen so much more than rents in superstar cities?*, Center for Economic Performance.

Hilber, C., & Turner, T. (2014). The Mortgage Interest Deduction and its Impact on Homeownership Decisions. *The Review of Economics and Statistics*, *96*(4), 618–637. https://doi.org/10.1162/REST{_}A{_}00427

Horn, K. M., & O'Regan, K. M. (2011). The low income housing tax credit and racial segregation. *Volume 21, Issue 3, Pages 443 - 473*, *21*(3), 443–473. https://doi.org/10.1080/10511482.2011.591536

Jeffrey Kling, B. R., Liebman, J. B., Katz, L. F., Beecroft, E., Feins, J., Goodson, B., Jacob, R., Kennedy, S., Orr, L., Patterson of Abt Associates, R., our collaborators Jeanne Brooks-Gunn, t., Del Conte Dickovick, A., Duncan, G., Leventhal, T., Ludwig, J., Psaty, B., Sanbonmatsu, L., & Whitaker, R. (2007). Experimental Analysis of Neighborhood Effects. *Econometrica*, *75*(1), 83–119. https://doi.org/10.1111/J.1468-0262.2007.00733.X

Li, H., Campbell, H., & Fernandez, S. (2013). Residential Segregation, Spatial Mismatch and Economic Growth across US Metropolitan Areas. *Urban Studies*, *50*(13), 2642–2660. https://doi.org/10.1177/0042098013477697/ASSET/IMAGES/LARGE/10.1177{_}0042098013477697-FIG1.JPEG

Miltenburg, E. M., van de Werfhorst, H. G., Musterd, S. (, & Tieskens, K. ((2018). Consequences of Forced Residential Relocation: Early Impacts of Urban Renewal Strategies on Forced Relocatees' Housing Opportunities and Socioeconomic Outcomes. *Housing Policy Debate*, *28*(4), 609–634. https://doi.org/10.1080/10511482.2018.1424722

Musterd, S., Marcińczak, S., van Ham, M., & Tammaru, T. (2017). Socioeconomic segregation in European capital cities. Increasing separation between poor and rich. *Urban Geography*, *38*(7), 1062–1083. https://doi.org/10.1080/02723638.2016.1228371

Ooi, J. T., & Le, T. T. (2013). The spillover effects of infill developments on local housing prices. *Regional Science and Urban Economics*, *43*(6), 850–861. https://doi.org/10.1016/j.regsciurbeco.2013.08.002

Park, J., Acolin, A., Aliprantis, D., Ba, B., Bayer, P., Black, S., Blanco, H., Calder-Wang, S., Coulson, E., Gould Ellen, I., Harari, M., Kwon, E., Moore, D., Ortego-Marti, V., Timmins, C., & Strochak, S. (2025). Housing Vouchers and Neighborhood Effects: Can Expanding Access to Opportunity Pay for Itself? *.

Pennington, K. (2021, August). *Does Building New Housing Cause Displacement?: The Supply and Demand Effects of Construction in* (tech. rep.).

Rossi-Hansberg, E., Sarte, P.-D., & Owens III, R. (2010). *Housing Externalities* (tech. rep. No. 3).

Ruiz-Tagle, J., & López M, E. (2014). El estudio de la segregación residencial en Santiago de Chile: revisión crítica de algunos problemas metodológicos y conceptuales. *EURE (Santiago)*, *40*(119), 25–48. https://doi.org/10.4067/S0250-71612014000100002

Sabatini, F., & Brain, I. (2008). La segregación, los guetos y la integración social urbana: mitos y claves. *EURE (Santiago)*, *34*(103), 5–26. https://doi.org/10.4067/S0250-71612008000300001

Schwartz, A. E., Ellen, I. G., Voicu, I., & Schill, M. H. (2006). The external effects of place-based subsidized housing. *Regional Science and Urban Economics*, *36*(6), 679–707. https://doi.org/10.1016/j.regsciurbeco.2006.04.002

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, *225*(2), 175–199. https://doi.org/10.1016/J.JECONOM.2020.09.006

Watson, T. (2009). Inequality and the measurement of residential segregation by income in American neighborhoods. *Review of Income and Wealth*, *55*(3), 820–844. https://doi.org/10.1111/J.1475-4991.2009.00346.X

Won, J. (2022). Exploring Spatial Clustering Over Time and Spillover Effects of the Low-Income Housing Tax Credit on Neighborhood-Level Income Segregation. *Urban Affairs Review*, *58*(3), 799–831. https://doi.org/10.1177/1078087420973436/ASSET/IMAGES/LARGE/10.1177{_}1078087420973436-FIG2.JPEG

Zahirovich-Herbert, V., & Gibler, K. M. (2014). The effect of new residential construction on housing prices. *Journal of Housing Economics*, *26*, 1–18. https://doi.org/10.1016/j.jhe.2014.06.003

8 Appendix

8.1 Propensity Score Matching

Figure A.1: Density plots

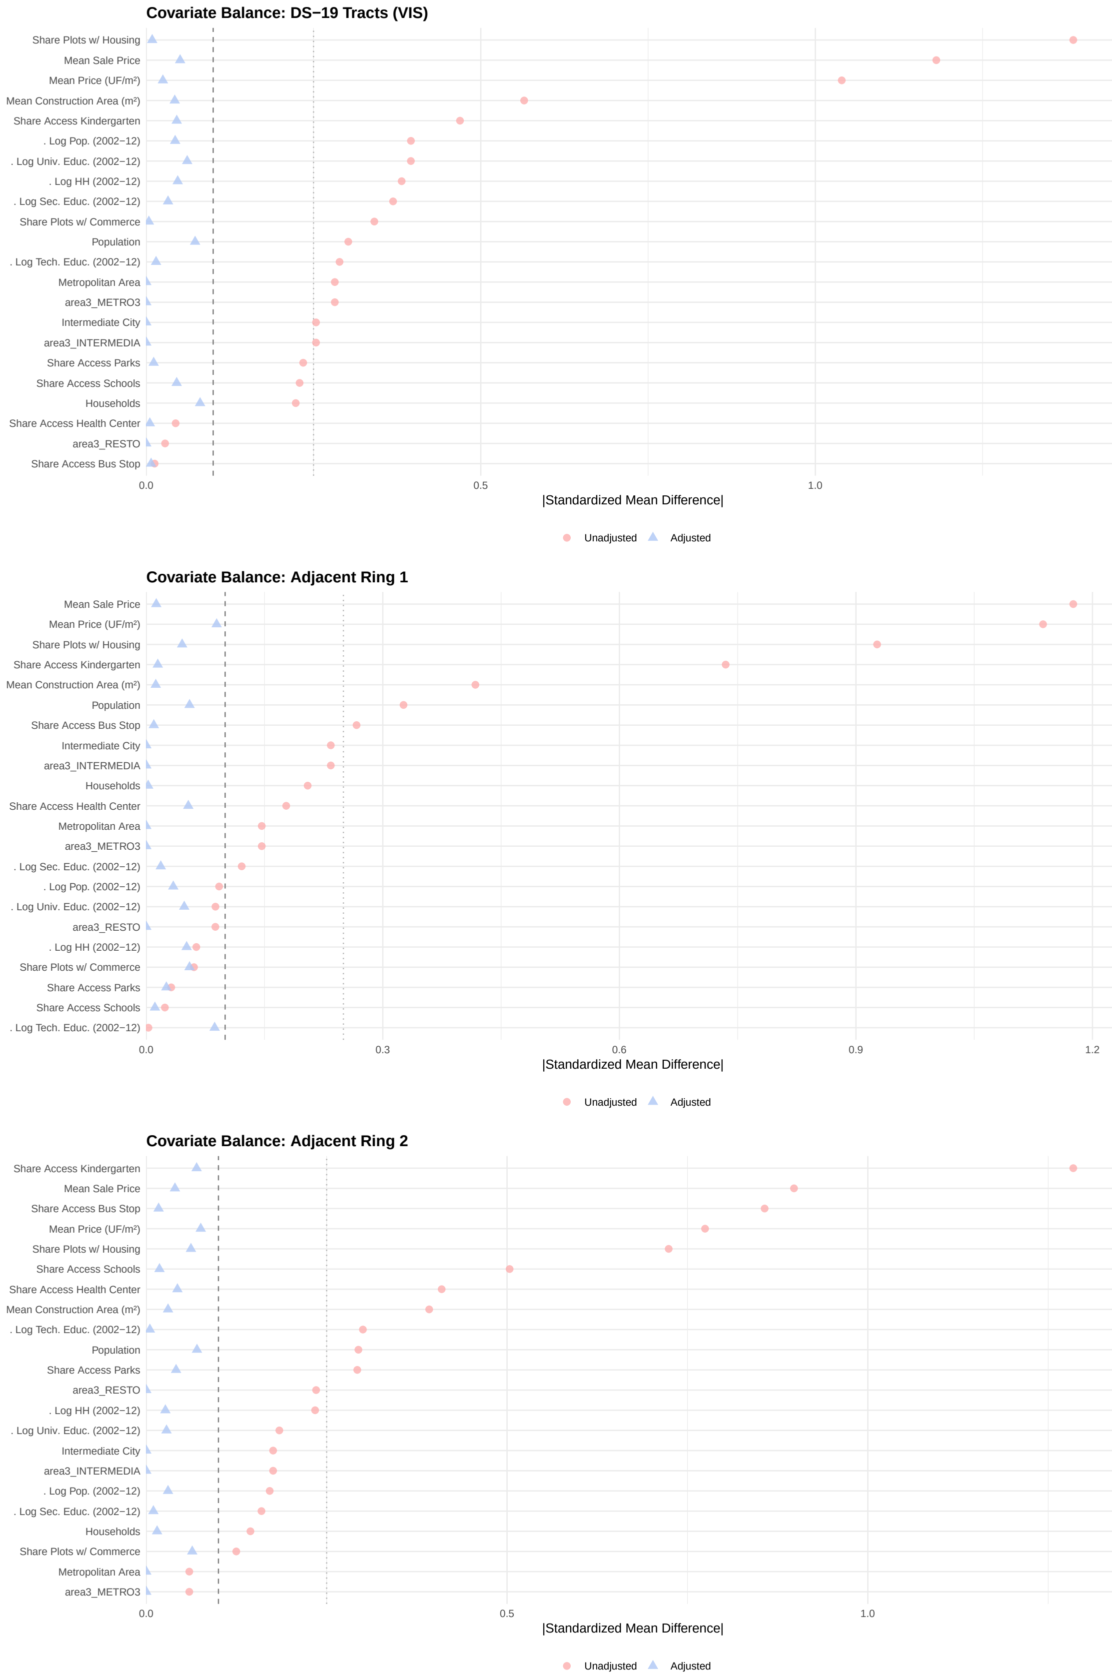

Figure A.2: SMD for each group

8.2 Tables

Table A.1: Pre-trends joint tests on lead coefficients

Panel	Outcome	Method	F-stat	p-value
VIS	ln_HSmax	TWFE	0.42	0.655
VIS	ln_Tmin	TWFE	0.10	0.904
VIS	ln_ratio	TWFE	0.84	0.431
VIS	ln_HSmax	Sun–Abraham	1.58	0.191
VIS	ln_Tmin	Sun–Abraham	1.42	0.235
VIS	ln_ratio	Sun–Abraham	0.77	0.513

Note. The table reports joint F-tests for the null that all pre-treatment lead coefficients are equal to zero.

Table A.2: Placebo estimates in the 2002–2012 pre-policy period

Panel	Outcome	Estimate	S.E.	Sig.
VIS	ln_HSmax	-0.0990	0.0528	*
VIS	ln_Tmin	-0.0764	0.0563	
VIS	ln_ratio	0.0226	0.0298	
VIS	ln_HSmax	-0.0217	0.0420	
VIS	ln_Tmin	-0.0547	0.0510	
RING1	ln_ratio	-0.0330	0.0228	
VIS	ln_HSmax	0.0178	0.0271	
VIS	ln_Tmin	0.0186	0.0312	
RING2	ln_ratio	0.0007	0.0198	

Note. Placebo estimates are obtained by restricting the sample to 2002 and 2012. The single marginally significant coefficient appears for secondary education in one VIS specification.

Table A.3: Alternative weighting schemes by cohort

Panel	Outcome	Weights	Cohort 2017	Cohort 2024
VIS	ln_ratio	PSM (baseline)	0.2803	0.1207
VIS	ln_ratio	Unweighted	0.3112	0.1473
VIS	ln_ratio	Unweighted + Controls	0.2536	0.1275
VIS	ln_Tmin	PSM (baseline)	0.4276	0.1979
VIS	ln_Tmin	Unweighted	0.4601	0.2267
VIS	ln_Tmin	Unweighted + Controls	0.3683	0.1789
VIS	ln_HSmax	PSM (baseline)	0.1473	0.0772
VIS	ln_HSmax	Unweighted	0.1489	0.0793
VIS	ln_HSmax	Unweighted + Controls	0.1147	0.0514
RING1	ln_ratio	PSM (baseline)	0.0227	-0.0212
RING1	ln_ratio	Unweighted	0.0448	-0.0035
RING1	ln_ratio	Unweighted + Controls	-0.0054	0.0198
RING2	ln_ratio	PSM (baseline)	0.0487	0.0206
RING2	ln_ratio	Unweighted	0.0393	0.0101
RING2	ln_ratio	Unweighted + Controls	-0.0134	-0.0111

Note. Entries report cohort-specific post-treatment coefficients under alternative weighting schemes.

Table A.4: Alternative specifications by cohort

Panel	Outcome	Specification	Cohort 2017	Cohort 2024
VIS	ln_ratio	Baseline + City Trends	0.2414	0.1166
VIS	ln_ratio	No City Trends	0.2743	0.1267
VIS	ln_ratio	No Controls	0.2803	0.1207
VIS	ln_Tmin	Baseline + City Trends	0.3532	0.1638
VIS	ln_Tmin	No City Trends	0.4287	0.2024
VIS	ln_Tmin	No Controls	0.4276	0.1979
VIS	ln_HSmax	Baseline + City Trends	0.1118	0.0472
VIS	ln_HSmax	No City Trends	0.1544	0.0757
VIS	ln_HSmax	No Controls	0.1473	0.0772
RING1	ln_ratio	Baseline + City Trends	0.0206	0.0083
RING1	ln_ratio	No City Trends	0.0277	0.0089
RING1	ln_ratio	No Controls	0.0227	-0.0212
RING1	ln_Tmin	Baseline + City Trends	0.0225	-0.0243
RING1	ln_Tmin	No City Trends	0.0654	-0.0052
RING1	ln_Tmin	No Controls	0.0615	-0.0296
RING1	ln_HSmax	Baseline + City Trends	0.0019	-0.0325
RING1	ln_HSmax	No City Trends	0.0377	-0.0141
RING1	ln_HSmax	No Controls	0.0388	-0.0083
RING2	ln_ratio	Baseline + City Trends	0.0206	0.0083
RING2	ln_ratio	No City Trends	0.0277	0.0089
RING2	ln_ratio	No Controls	0.0227	-0.0212
RING2	ln_Tmin	Baseline + City Trends	0.0225	-0.0243
RING2	ln_Tmin	No City Trends	0.0654	-0.0052
RING2	ln_Tmin	No Controls	0.0615	-0.0296
RING2	ln_HSmax	Baseline + City Trends	0.0019	-0.0325
RING2	ln_HSmax	No City Trends	0.0377	-0.0141
RING2	ln_HSmax	No Controls	0.0388	-0.0083

Note. The RING2 block reproduces the values reported by the script and should be verified because it matches RING1 exactly in the exported output.

8.3 Figures

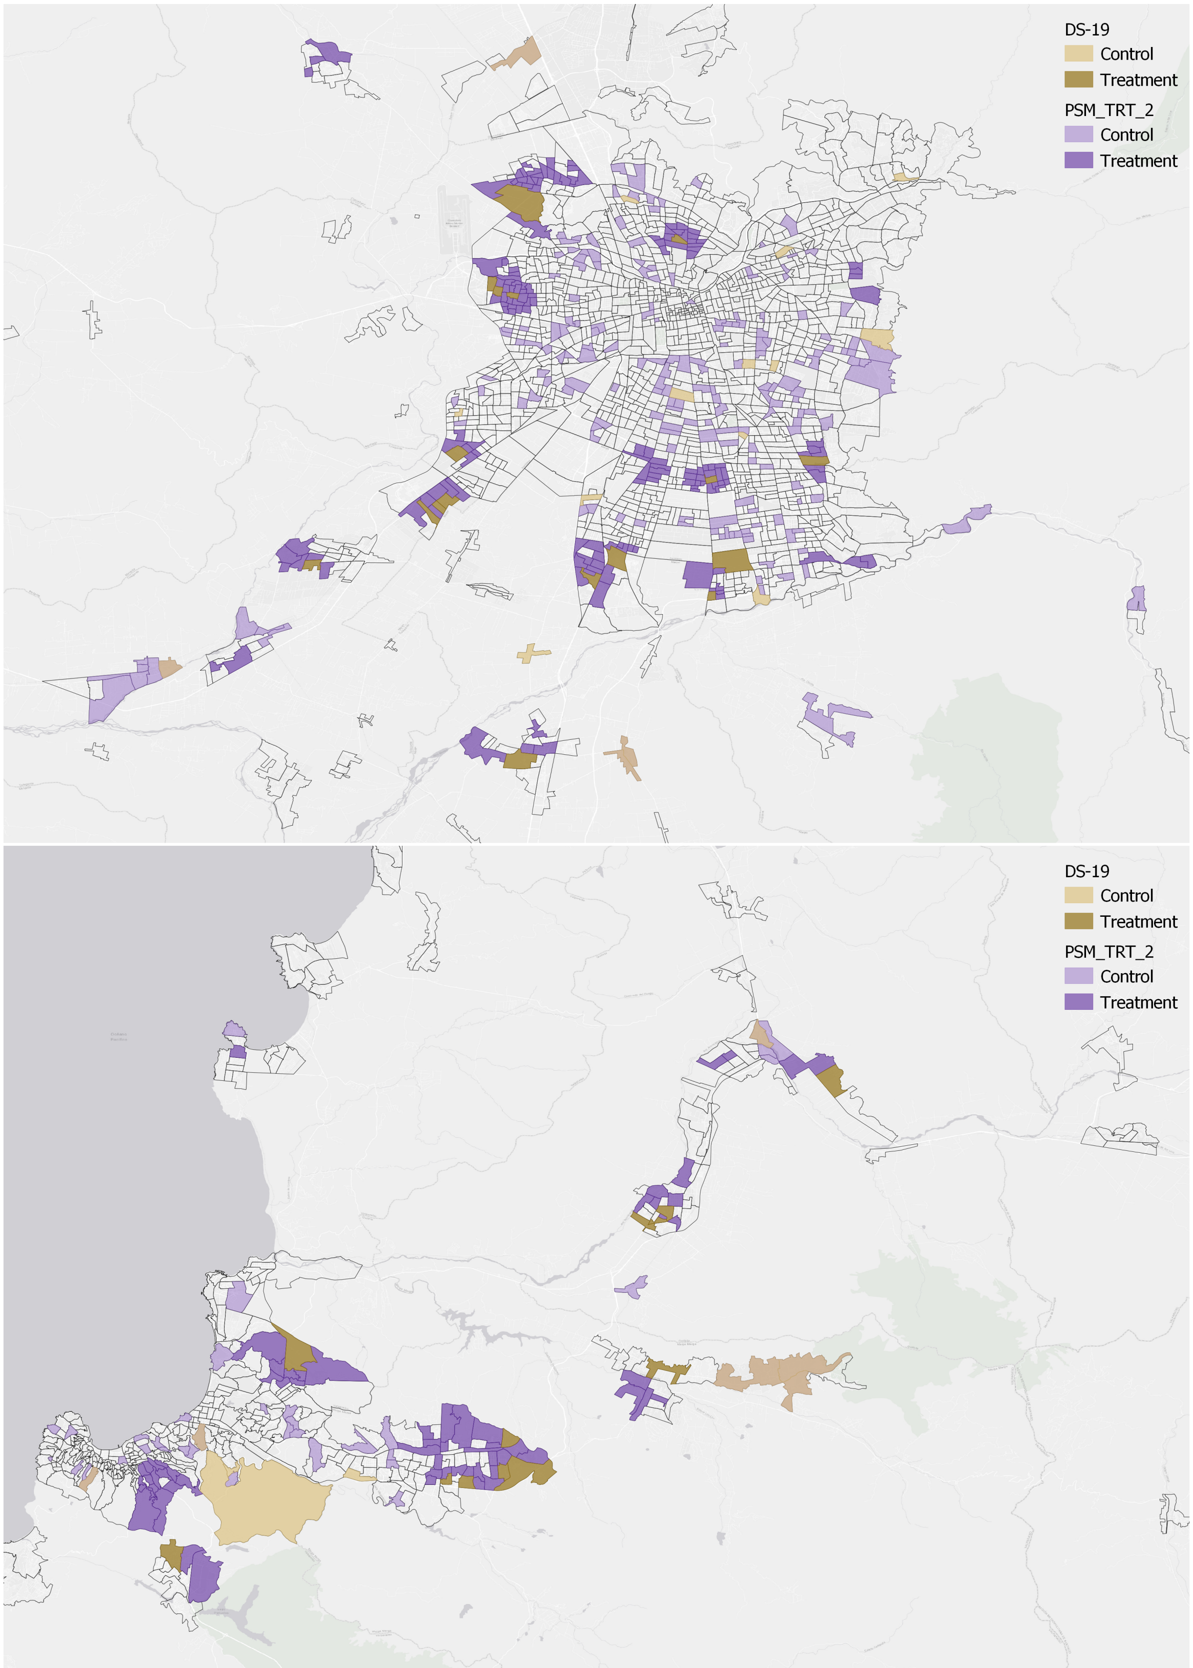

Figure A.3: PSM Results in Region Metropolitana and Valparaiso

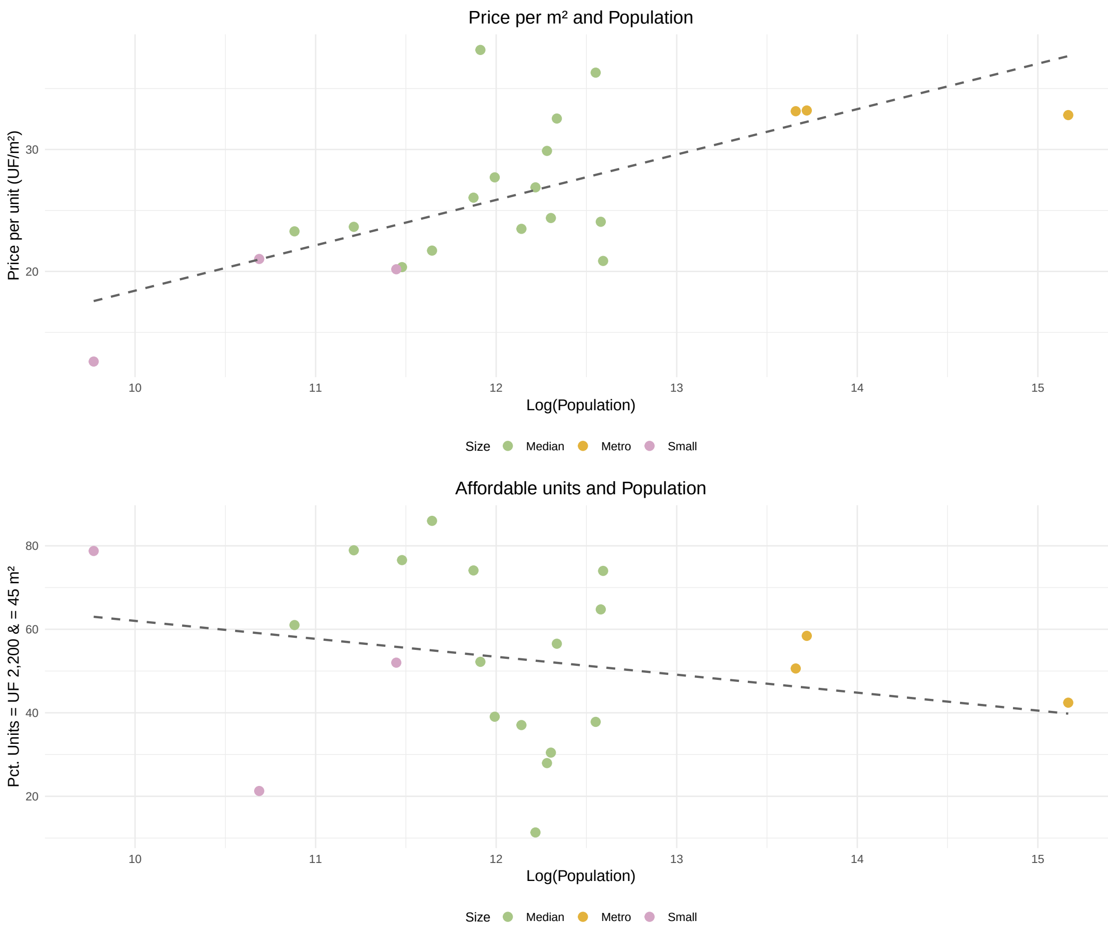

Figure A.4: Unit prices and probability of affordable units per city size

8.4 Education and income in CASEN 2017

The Population Census does not report income. In the main analysis, I therefore measure neighborhood socio-economic composition using educational attainment (shares by education level) at the census-tract level. This appendix validates that interpretation by documenting a strong and monotonic relationship between education and income among household heads using CASEN 2017, Chile's nationally representative household survey. I restrict the sample to household heads and focus on main-occupation labor income for employed heads. I report (i) descriptive income differences by education group, (ii) standard Mincer-type specifications relating log income to years of schooling and potential experience, and (iii) regressions using discrete education categories relative to secondary education. All OLS specifications use CASEN survey weights and account for the complex survey design, and the "full" specifications include demographics, job characteristics, and region fixed effects. I also report quantile regressions to show that the education gradient holds across the income distribution. Education categories in CASEN are constructed to match those used to compute census-tract educational shares in the main analysis.

Variables. The dependent variable is log(income), where income is main-occupation labor income (CLP) for employed household heads. Years of schooling and potential experience (and its square) follow the standard Mincer framework. "Demographics" include gender, marital status, an urban indicator, and an indigenous indicator. "Job controls" include formal-sector status, a large-firm indicator, and a public-sector indicator. "Region FE" are region fixed effects. Standard errors are reported in parentheses.

Table A.5: CASEN 2017: Mean Income by education level (household heads)

Ed. Level	n	Income (CLP)	log(income)	Age	Years schooling	Urban
Secondary	50,454	345,930	12.6	45.7	9.89	0.784
Technical	9,018	545,160	13.0	40.4	14.7	0.911
University	15,015	923,563	13.4	39.6	16.6	0.929
Postgraduate	1,768	1,738,890	14.1	41.9	19.1	0.950

Notes: Sample restricted to employed household heads in CASEN 2017. Income is main-occupation labor income (CLP). Statistics use survey weights. Education categories are defined to match those used to compute Census-tract educational shares in the main analysis.

Table A.6: CASEN 2017: Education–income gradient for household heads

	Panel A: Years of schooling		**Panel B: Education levels (ref: Secondary)**	
	(1) Basic	(2) Full	(3) Basic	(4) Full
Years of schooling	0.119***	0.100***		
	(0.002)	(0.002)		
Technical			0.421***	0.309***
			(0.012)	(0.009)
University			0.896***	0.767***
			(0.018)	(0.016)
Postgraduate			1.534***	1.318***
			(0.028)	(0.028)
Experience	0.011***	0.008***	0.035***	0.028***
	(0.001)	(0.001)	(0.002)	(0.001)
Experience^2	-0.0001***	-0.00003	-0.001***	-0.001***
	(0.00003)	(0.00002)	(0.00003)	(0.00002)
Female		-0.268***		-0.262***
		(0.006)		(0.006)
Married		0.104***		0.105***
		(0.008)		(0.008)
Urban		-0.026**		0.050***
		(0.013)		(0.013)
Indigenous		-0.041***		-0.050***
		(0.012)		(0.011)
Formal sector		0.307***		0.344***
		(0.012)		(0.012)
Large firm		0.147***		0.144***
		(0.008)		(0.008)
Public sector		0.147***		0.136***
		(0.010)		(0.010)
Observations	76,255	57,715	76,255	57,715
Demographics	No	Yes	No	Yes
Job controls	No	Yes	No	Yes
Region FE	No	Yes	No	Yes

Notes: Dependent variable is log main-occupation labor income of employed household heads. Panel A reports Mincer regressions using years of schooling; Panel B reports regressions using discrete education categories (reference: Secondary). "Full" includes demographics (female, married, urban, indigenous), job controls (formal sector, large firm, public sector), and region fixed effects. Estimates use CASEN survey weights and account for the complex survey design. Standard errors in parentheses. Education categories are defined to match those used to compute Census-tract educational shares in the main analysis. *Significance:* *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.7: CASEN 2017: Quantile regressions (education categories)

	Q25	Q50	Q75
Technical (vs Secondary)	0.323***	0.373***	0.480***
	(0.005)	(0.007)	(0.009)
University (vs Secondary)	0.613***	0.865***	1.064***
	(0.009)	(0.008)	(0.008)
Postgraduate (vs Secondary)	1.267***	1.451***	1.661***
	(0.024)	(0.018)	(0.023)
Experience	0.023***	0.021***	0.026***
	(0.001)	(0.001)	(0.001)
Experience^2	-0.0005***	-0.0004***	-0.0005***
	(0.00002)	(0.00001)	(0.00001)
Female	-0.325***	-0.230***	-0.305***
	(0.005)	(0.004)	(0.005)

Notes: Dependent variable is log main-occupation labor income of employed household heads. Coefficients are quantile regression estimates (Q25, Q50, Q75). Standard errors in parentheses. Education categories are defined to match those used to compute Census-tract educational shares in the main analysis. *Significance:* *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.