Difference-in-Differences When Units Are Substitutes: Evidence from Place-Based Policies*

Guillermo Alves† William H. Burton‡ Sebastian Fleitas§

November 26, 2025

Abstract

Comparing housing prices across policy borders is a common method for evaluating place-based interventions. However, re-sorting can alter prices in control areas, violating SUTVA. We analyze the difference-in-differences estimator in this setting and derive a sufficient statistics formula showing that bias increases with neighborhood substitutability and inelastic housing supply. Studying a housing tax break in Uruguay, we find that border estimates maximize substitutability and bias, failing to reject zero incidence. A structural model accounting for differential substitutability estimates an incidence of 80%. Common reduced-form alternatives, such as the ring and aggregation approaches, substantially reduce but do not eliminate bias.

Keywords: Place-Based Policies, Difference-in-Differences, SUTVA, Spillovers
JEL Codes: R31, R13, H22, L78

*We would like to thank Milena Almagro, Nathaniel Baum-Snow, Lorenzo Caliendo, Brantly Callaway, Chris Conlon, Geert Dhaene, Jan De Loecker, Xavier D'Haultfoeuille, Jonathan Dingel, Tomás Domínguez-Iino, Sebastian Gallegos, Geert Goeyvaerts, Taylor Jaworski, Jorge Pérez Pérez, Matthew Turner, and Frank Verboven for comments and suggestions. We thank Mauro Castiella and Christian Valencia, who provided excellent research assistance, and Chantal Boulay, Irene Chavione, Claudio Fernández, and María Laura Moratorio, who helped us access the data on transaction prices.

†CAF Development Bank of Latin America and the Caribbean.
‡Uber
§PUC-Chile and CEPR.

1 Introduction

Comparing outcomes across policy borders is a commonly used strategy for evaluating place-based policies (Black, 1999; Bayer et al., 2007; Neumark & Kolko, 2010; Turner et al., 2014; Chen et al., 2022). While such designs reduce the risk of bias from unobservables by selecting similar areas, they increase the likelihood of re-sorting spillovers from treatment into control areas, since more comparable areas tend to be closer substitutes. The presence of re-sorting as an equilibrium response to the policy violates the stable unit treatment value assumption (SUTVA), making many commonly used estimators inconsistent.[1] For example, relocation of agents in response to place-based policies or trade shocks can alter equilibrium prices in areas used as controls, biasing treatment effect estimates (Baum-Snow & Ferreira, 2015; Donaldson, 2015). Biases in price effects are particularly relevant, as they constitute the basis for evaluating the efficiency of taxes and subsidies, as well as for hedonic welfare analysis.

The purpose of this paper is to examine how re-sorting-induced SUTVA violations affect the estimation of housing price effects in a Difference-in-Differences (DiD) framework in the context of a place-based tax incentive. Using a general supply-and-demand spatial framework, as in Allen and Arkolakis (2023), we decompose the DiD estimator and derive a sufficient statistics formula to assess the bias (Saez, 2001; Chetty, 2009). We provide practical guidelines for researchers dealing with potential SUTVA violations due to re-sorting between treatment and control areas.

We begin by showing that, in the presence of re-sorting, the DiD estimator can be decomposed into three effects. First, an "autarky effect" captures what would happen to the treated area in isolation, absent relocation responses. Second, a "spillover-on-switchers effect" captures the equilibrium changes in the treated area caused by agents moving in or out. Third, a "contamination effect" reflects similar equilibrium adjustments in control areas. It is this contamination term that introduces bias in the DiD estimator, preventing it from accurately recovering the average treatment effect on the treated (ATT).

By linearizing the general supply and demand model, we provide an analytical formula that approximates the DiD estimate of introducing a housing supply subsidy in a subset of neighborhoods. First, the formula shows that the DiD estimator is asymptotically biased except in the extreme cases of perfectly elastic supply or perfectly inelastic demand in the control area. Second, it indicates the sign of the bias under mild economic assumptions. Finally, it highlights that the relevance of the three effects in the DiD es-

[1] A few estimation techniques explicitly recognize the potential source of bias due to spatial spillovers, but they may not be well suited to addressing re-sorting spillovers. The ring or donut approach assumes that spillovers decline with distance and vanish beyond a certain threshold (Clarke, 2017; Butts, 2022). This assumption is plausible for physical spillovers, such as pollution, but re-sorting spillovers may not decay monotonically with distance or have a clear spatial cutoff. Another approach involves aggregating spatial units to contain localized spillovers within broader treatment and control areas, but this method cannot account for spillovers between those larger units (Feyrer et al., 2017; Huber & Steinmayr, 2021).

timator depends on demand-side substitution patterns across neighborhoods and the supply elasticities of those neighborhoods.

Our formula has implications for applied work. More similar areas tend to be closer demand-side substitutes, making them more prone to contamination effects. This contradicts the common DiD practice of selecting highly similar treatment and control units, such as border comparisons or employing matching techniques (Neumark & Kolko, 2010; Chen et al., 2022). Conditional on having parallel pre-trends, applied researchers should prefer comparisons between less homogeneous areas when place-based policies induce substantive re-sorting. Using simulations of a specific supply and demand model, we show that in that model there is a trade-off between the parallel trends assumption and the SUTVA. When control and treatment areas are very similar, parallel trends hold more frequently, but contamination bias is more severe.

We apply these methodological insights to the study of a typical place-based policy that provides large tax breaks for housing development in lagging neighborhoods. We focus on a major program implemented in Montevideo, the capital of Uruguay. Our analysis begins by estimating a series of DiD regressions using administrative data on the universe of housing transactions in the city. We estimate alternative DiD specifications—including the border, ring, and aggregation approaches—comparing areas of varying similarity and, therefore, varying degrees of substitutability. We illustrate how heterogeneity analysis in DiD estimates can help detect bias. Under no signs of bias—e.g., similar effects across units with varying susceptibility to spillovers—DiD remains a viable identification strategy.

We document three heterogeneity patterns in our DiD estimates, all consistent with the presence of SUTVA violations due to re-sorting. First, when following an aggregation approach that uses all transactions in the city, we find a large negative effect of the policy, equivalent to around 18% of the average transaction price. In contrast, using only observations near the policy border yields very small negatives or zeros. Second, the absolute magnitude of these border estimates increases with price differences across different segments of the border, a proxy of substitutability. Third, estimates decrease when control units are located farther from the border, as in the ring approach. Importantly, we also show that these second and third patterns are inconsistent with our results being explained by policy-induced changes in amenities.

When evidence of contamination emerges and the policy represents a marginal change, our sufficient statistics formula allows researchers to approximate the bias. However, if there are strong signs of contamination and the intervention is large, a structural model becomes necessary to quantify the bias and recover the true policy effect. With this purpose, we develop a structural model that allows us to calculate the incidence of the tax break and benchmark the performance of the alternative reduced-form estimates.

We estimate a structural model of housing supply and demand across Montevideo's neighborhoods. Housing demand is modeled as a discrete choice over neighborhoods within a city (Bayer et al., 2007; Anagol et al., 2021; Almagro & Domınguez-Iino, 2025). We assume a nested-logit structure and estimate the price elasticity using a set of supply-shifting instruments generated by the tax break (Berry, 1994). Nests are formed by grouping neighborhoods with similar socioeconomic composition, based on census data, and validated using data on actual migration flows across neighborhoods. As in Adao et al. (2019) and Borusyak et al. (2022), the model's structure allows for consistent estimation by including regression terms (i.e., the within-nest shares) that explicitly account for spillovers. By explicitly accounting for spillovers and instrumenting for both the price and nested-logit terms, this strategy complies with Allen and Arkolakis (2023)'s gold standard for identification in spatial settings with spillovers. On the supply side, each neighborhood has a separate log-linear supply function (Saiz, 2010; Baum-Snow & Han, 2023). We use the model-free DiD estimate to calibrate a common inverse supply elasticity for all neighborhoods by matching that estimate with its structural equivalent. We show that our model fits the data well in terms of reproducing the parallel trends observed in the pre-treatment period.

We exactly decompose the DiD estimate into the three effects by solving for counterfactual equilibria of the estimated model. We find that the spillover-on-switchers effect amounts to 40% of the autarky effect, while contamination represents 25% of the ATT. This level of contamination implies that all reduced-form estimates substantially underestimate the share of the subsidy that reaches consumers, though to varying degrees, and in a way that is consistent with our framework. The aggregation approach exhibits the lowest bias and still underestimates the share of the subsidy that reaches consumers by 20 percentage points (60% versus 80%). In absolute terms, this bias amounts to a quarter of Uruguay's GDP per capita. As noted above, a border estimate cannot rule out the possibility that consumers did not benefit from the tax break. The ring approach quantifies the incidence at 30%.

Finally, we use the model to compute counterfactual equilibria to revisit the relationship we find in the reduced-form analysis between heterogeneity across control and treated units and the magnitude of the DiD estimate. Consistent with our decomposition formula, we confirm that contamination bias is negatively correlated with our measure of heterogeneity between units and positively correlated with diversion ratios. Importantly, this implies that the smaller absolute values of the reduced-form DiD estimates obtained by comparing more homogeneous units are effectively driven by contamination (i.e. greater bias) and not just regular treatment heterogeneity.

Related Literature. Our paper relates to four main strands of literature. First, we contribute to the literature on causal inference in spatial settings by making explicit the equilibrium nature of contamination and the trade-offs between comparability and contamination. Our decomposition formula illustrates two types of assumptions that have allowed consistent identification of the impact of place-based policies in the presence of re-sorting spillovers. Either the shock is small and local, such that distant areas remain unaffected by re-sorting and one can employ the ring approach (Kline & Moretti, 2014; Delgado & Florax, 2015; Clarke, 2017; Mayer & Trevien, 2017; Butts, 2023a) or units are infinitesimal such that re-sorting does not affect prices and quantities in untreated areas (Busso et al., 2013; Turner et al., 2014).[2]

Second, spatial re-sorting and its effects are at the heart of quantitative spatial models (Redding & Rossi-Hansberg, 2017). Previous work in this literature derives estimating equations that explicitly incorporate spillover effects and yield consistent estimates (Adao et al., 2019; Fu & Gregory, 2019; Borusyak et al., 2022; Rudik et al., 2022; Hollingsworth et al., 2025).[3] We follow the same approach to consistently estimate demand in our model and use counterfactual equilibria to assess the bias in reduced-form estimates and validate our approximation formula.

Third, we contribute to the general literature discussing identification under SUTVA violations. Sobel (2006) showed that comparing means between treatment and control under interference does not recover an ATT and instead yields the difference between two effects. Vazquez-Bare (2023) decomposed that difference in means into the three effects we study: the effect on the targeted group absent spillovers, spillovers on the targeted group, and spillovers on the non-targeted group. This decomposition into three terms was applied to the DiD context by Butts (2023a). We innovate on two fronts. First, we characterize the three terms using supply and demand partials.[4] Second, we compute the three terms with a structural model in a context with no spillover-free areas, which allows us to quantify the bias of different reduced-form estimators.

Fourth, beyond its methodological contribution, our paper contributes to the debate on the effectiveness of supply-side housing interventions to promote affordability (Been et al., 2019, 2025). Supply skeptics argue that new demand from other neighborhoods and improved amenities offset any downward pressure on prices from additional supply. We find instead that a large share of the generous tax break granted to developers is not offset by marginal cost increases or an increase in amenities and therefore is passed on to consumers through lower housing prices. This result adds to recent evidence showing that the negative effect of new supply on prices dominates the positive impacts on prices of improved amenities and immigration from nearby areas (Li, 2022; Asquith et al., 2023). Nevertheless, our finding that the spillover effect on switchers amounts to 40% of the autarky effect indicates that demand re-sorting exerts a sizable offsetting force on prices, helping to explain some of the political economy tensions surrounding the housing affordability debate.

[2]Jardim et al. (2024) informally suggest these two assumptions in their discussion of minimum wage effects under spatial spillovers.

[3]This approach has also been applied beyond spatial contexts. For example, Rotemberg (2019) studies spillovers of firm subsidies in India and Bachmann et al. (2023) analyze demand spillovers in the 2015 Volkswagen emissions scandal. Additionally, Fan and Yang (2025) decompose market equilibrium effects as differences of counterfactual scenarios.

[4]Munro et al. (2025) use price elasticities to recover consistent estimates of the direct effect of the policy as well as spillover effects in an experimental setting with agents facing a unique price in self-contained markets.

2 Difference-in-Differences in Equilibrium

2.1 SUTVA and Difference-in-Differences

The stable unit treatment value assumption (SUTVA) requires that the outcome of each unit does not depend on the treatment status of other units (Imbens & Rubin, 2015). This assumption allows one to write the potential outcome of every unit as effectively depending only on its assigned treatment status. In a canonical DiD framework with two periods ($t \in \{pre, post\}$) and discrete treatment ($D \in \{0, 1\}$), there are two types of units j. Namely, one that never receives treatment, and one that receives treatment only in the post-period. In this framework, the first type of units has a potential outcome $Y_{j,t}(0)$ and the second type has $Y_{j,t}(1)$. The causal estimand of interest is the average treatment effect on the treated (ATT) in the second period (Roth et al., 2023):

$$ATT = \beta = \mathbb{E}[Y_{j,post}(1) - Y_{j,post}(0)|D_j = 1] \tag{1}$$

The challenge to compute the object of interest β is that $Y_{j,post}(0)$ is not observed when $D_j = 1$. Under the assumptions of parallel trends and no anticipation, the DiD estimator surmounts this challenge by building a counterfactual for the never observed $\mathbb{E}[Y_{j,post}(0)|D_j = 1]$. This counterfactual is obtained by adding the average change in the outcomes of the untreated units between both periods to the baseline average for treated units:

$$\hat{\beta}_{DiD} = (\bar{Y}_{t=post,D=1} - \bar{Y}_{t=pre,D=1}) - (\bar{Y}_{t=post,D=0} - \bar{Y}_{t=pre,D=0}) \tag{2}$$

where $\bar{Y_{t,d}}$ is the sample mean in period t. When SUTVA is violated, for example, due to the re-sorting of agents between treatment and control, the DiD estimator fails to estimate the ATT. We discuss this case in the following subsection.

2.2 SUTVA Violations in a City-Wide Market Equilibrium

SUTVA violations can arise for several reasons, including network effects or market equilibria (Manski, 1993). We apply our discussion of DiD to SUTVA violations caused by demand-side re-sorting of agents in response to the introduction of a supply-side subsidy in a city's housing market. Without loss of generality and under homogeneous effects, we assume that there are only two neighborhoods and a generic area outside the city. One neighborhood A, with housing price p_t^A, receives the subsidy, and the other neighborhood B, with housing price p_t^B, does not. With housing prices as the outcome variable, Equation 2 can be written as:

$$\hat{\beta}_{DiD} = (p_{post}^A - p_{pre}^A) - (p_{post}^B - p_{pre}^B) \tag{3}$$

Analogously to the analysis of Allen and Arkolakis (2023) on the spatial equilibrium across cities, the within-city equilibria can be expressed with simple supply and demand equations. The demand side in our context consists of households choosing whether to buy a housing unit in one of the two neighborhoods or outside the city. There are two main determinants of that discrete choice: housing prices and amenities. These are denoted by the vectors $\mathbf{p_t}$ and $\mathbf{a_t}$, respectively. The demand function for housing in each neighborhood j is $D^j(\mathbf{p_t}, \mathbf{a_t})$.

The supply side is characterized by property owners who choose whether to sell their housing unit (new construction or existing unit) located in neighborhoods A or B. Higher prices induce a higher supply of housing units available for sale. This relationship is represented by an upward-sloping supply function, $S^j(q_t^j)$, with q_t^j denoting the quantity offered in neighborhood j at time t. We assume that both households and property owners make static decisions in each period.

We first examine the DiD estimator in the case of no re-sorting between neighborhoods A and B. We then examine the more general case with re-sorting. After presenting these two cases, we introduce a generalized decomposition for two neighborhoods, which we then extend to many neighborhoods.[5]

Figure 1: DiD with No Re-Sorting between Neighborhoods A and B

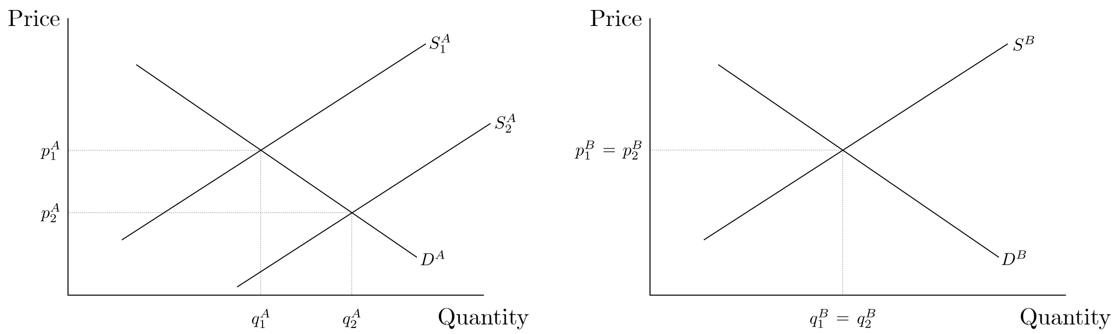

Figure 1 presents the autarky situation in which consumers do not relocate between A and B, but may relocate between each neighborhood and the outside option. Implementing a supply-side subsidy in A would first result in an outward shift of the supply in this neighborhood.[6] Due to lower prices, more households choose to live in A instead of outside the city, which explains the observed movement along the demand curve in A. Neither demand nor supply in neighborhood B are affected, and thus prices there do not change.[7] The estimated DiD in this scenario equals the difference in prices between periods 2 and 1 in the neighborhood A:

[5]Throughout the section, we focus on demand-side re-sorting of households, and thus abstract away from supply-side re-sorting. Supply-side re-sorting in response to a demand-side place-based policy could be analogously accommodated in the framework. As discussed in Section 3, we evaluate that supply-side re-sorting is not relevant in our empirical setting due to the existence of a very large number of small developers.

[6]Note that the policy affects the supply of all units—existing and new construction—because the future supply of new subsidized units depresses the prices of existing non-subsidized units. At any price, the supply would be larger once the subsidy is implemented.

[7]We abstract away from other changes happening over time. Note that this is analogous to assuming

$$\hat{\beta}^{AUT}_{DiD} = (p^A_2 - p^A_1) - (p^B_2 - p^B_1) = p^A_2 - p^A_1$$

In this situation of autarky, the DiD estimator correctly captures the effect of the subsidy on the targeted areas. Next, we show that this is not the case when agents re-sort between the two neighborhoods, as this violates SUTVA.

Figure 2: DiD with Re-Sorting between Neighborhoods A and B

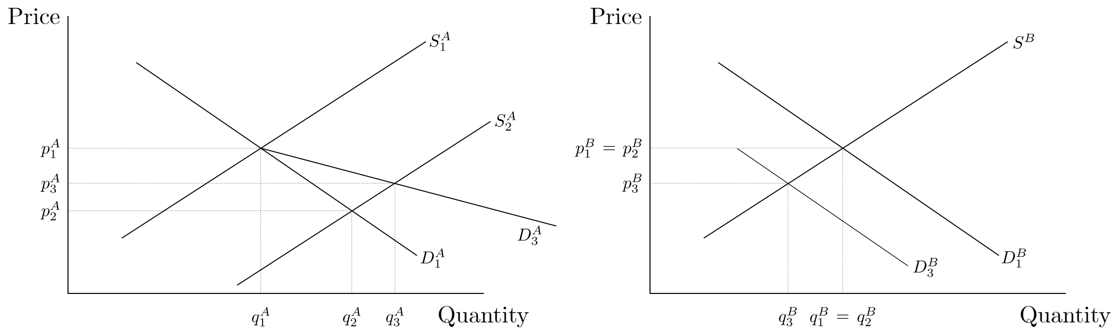

Figure 2 highlights a situation in which consumers may relocate between the two neighborhoods. When the supply side subsidy is introduced in neighborhood A, housing prices drop from p^A_1 to p^A_2. As in Allen and Arkolakis (2023), there is a new "round" of effects which we index as taking place at $t = 3$. Now, the demand curve rotates counterclockwise in neighborhood A, and shifts to the lower left in neighborhood B.[8] Both movements are due to re-sorting. Re-sorting increases prices in A from p^A_2 to p^A_3 while reduces those in B from $p^B_1 = p^B_2$ to p^B_3. Estimating the effect of the policy with DiD now yields the following:

$$\begin{aligned}
\hat{\beta}_{DiD} &= (p^A_3 - p^A_1) - (p^B_3 - p^B_1) \\
&= (p^A_3 - p^A_2 + p^A_2 - p^A_1) - (p^B_3 - p^B_2 + p^B_2 - p^B_1) \\
&= (p^A_2 - p^A_1) + (p^A_3 - p^A_2) - (p^B_3 - p^B_2)
\end{aligned}$$

With demand re-sorting between the two neighborhoods, the estimated DiD contains not only the autarky effect from before but also the price increase in A due to higher demand, as well as the price decrease in B due to the lower demand.[9] As indicated in Equation 4, we refer to the additional effect in A as "spillover-on-switchers", and to the effect in B as "contamination". While in our context both effects attenuate the autarky

"parallel trends" as in Roth et al. (2023). Those other changes may include improvements in amenities correlated with the policy, which we discuss in detail in Subsection 7.3.

[8]The new demand curve in A passes through the original (q^1_A, p^1_A) pair, reflecting that the amount of housing demanded would be the same at the original price, but yields higher demanded quantities for prices below p^1_A, capturing the re-sorting of agents away from B and into A in reaction to those lower prices.

[9]Using the "exposure mapping" notation from Aronow and Samii (2017), it is also possible to write these three terms as potential outcomes along the lines of proposition 2.1 in Butts (2023a).

effect of the policy, the former is part of the legitimate effect of the policy on the targeted neighborhood and the latter "contaminates" the DiD estimate.

$$\hat{\beta}_{DiD} = \underbrace{\underbrace{(p_2^A - p_1^A)}_{\text{Autarky}} + \underbrace{(p_3^A - p_2^A)}_{\text{Spillover-on-Switchers}}}_{\text{Treatment Effect on Subsidized Area}} - \underbrace{(p_3^B - p_2^B)}_{\text{Contamination}} \tag{4}$$

In this market equilibrium setting, the DiD estimate thus no longer recovers the ATT of the policy, which is given by the sum of the first two terms of Equation 4. As noted by Sobel (2006), differences in means under SUTVA violations recover the relative effect between treated and control units. This relative effect, given by the sum of the three effects (i.e. $\hat{\beta}_{DiD}$), could be of interest in some contexts. For example, when the researcher is interested in the effect of a policy on the outcome of one region relative to others, such as the distributional effects of trade shocks (Dix-Carneiro and Kovak, 2017). But even in these cases, there is great value in recovering the effects on different areas separately. On one hand, the ATT allows policymakers to understand the total effect of the policy on the targeted area. On the other hand, the contamination effect can be of interest by itself, as it shows an effect of the policy on non-targeted areas.

In Figure 2, the relative effect of the policy is zero because the price reduction is the same in both regions. More generally, in markets with the same fundamentals before the policy, the relative effect of the subsidy will always be zero. This is because equilibrium arbitrage would equalize prices after the policy, even when the policy is implemented only in one area. However, as in our example, the policy can reduce prices in both areas compared to the situation before the policy. So the DiD estimate would state that the policy had no (relative) effect when it reduced prices in both areas.

Finally, our market equilibrium framework implies that the issue of contamination in Equation 4 arises in a broader set of contexts than the one we study. First, it is not only specific to DiD methods. Contamination originates from the utilization of the control area equilibrium prices that already incorporate the effect of the policy (p_3^B). Therefore, other estimation approaches, such as regression discontinuity or propensity score matching, that also use p_3^B as a control, would produce estimates that suffer from contamination. Second, Equation 4 determines the sign of the bias by showing that contamination is subtracted from the ATT. In our context, the policy changes outcomes in treatment and control areas in the same direction (i.e., both ATT and contamination are negative), so contamination biases the estimate towards zero. In other very relevant contexts, such as studies of the impact of place-based policies on employment or firm creation within a city (Mayer et al., 2017), contamination may produce upward-biased estimates (i.e., ATT is positive and contamination negative).

2.3 DiD Decomposition with Supply and Demand Elasticities

In this subsection, we derive an approximation formula for the DiD estimator that helps to understand the determinants of the relative sizes of both the spillover-on-switchers and contamination effects. The relative size of contamination in that formula defines the relative size of the asymptotic bias of the DiD estimate.

We linearize the supply and demand model from above to express Equation 4 in terms of supply and demand elasticities. We start with two neighborhoods and one outside option, and then generalize to multiple subsidized neighborhoods.[10]

The case of one subsidized neighborhood. Define the inverse housing supply function as $P_S^j(q^j)$, and the diversion ratio $DR_{A,B}$ as the ratio between the change in demand for B and the change in the demand for A when the price of A changes.[11] As we show in Appendix B, the DiD estimator is approximately equal to:

$$\hat{\beta}_{\text{DiD}} \approx \underbrace{(p_2^A - p_1^A)}_{\text{Autarky in } A} \times \left[1 + \underbrace{\frac{\partial D^A}{\partial p^A} \times \frac{\partial P_S^A}{\partial q^A}}_{\text{Spillover-on-Switchers Scaling}} \underbrace{- \frac{\partial D^A}{\partial p^A} \times \frac{\partial P_S^B}{\partial q^B} \times DR_{A,B}}_{\text{Contamination Scaling}}\right] \tag{5}$$

Equation 5 highlights that the DiD estimate in a situation with re-sorting between subsidized and unsubsidized neighborhoods is a scaled version of the policy's effect in autarky. Intuitively, the scaling factors depend on the responsiveness of demand and supply in the two neighborhoods, increasing with the demand's sensitivity to prices and the supply-side responsiveness of prices to quantities.

The second term inside the square brackets in Equation 5 is the scaling factor due to the spillover-on-switchers. It captures the effect of households relocating to this area as a result of the subsidy. The last term inside the main bracket deserves special attention as it is the one causing the DiD estimator to be biased and unable to recover the true effect of the policy on the subsidized areas. This term increases linearly with respect to each of its three components: the partial of the demand in the subsidized neighborhood with respect to its own price, the partial of the inverse supply in the unsubsidized neighborhood with respect to its own quantity, and the diversion ratio between the two neighborhoods. Intuitively, the bias of the DiD estimator is higher when subsidized and unsubsidized neighborhoods are close substitutes and the supply curve in unsubsidized neighborhoods is more inelastic.

The case of multiple subsidized neighborhoods. The general formula still computes the DiD term between A and B but allows for re-sorting into A and B from all other neighborhoods.[12] In this general case, the DiD estimator can be approximately com-

[10]Figure A.1 in Appendix A presents a graphical representation of situations with a single versus multiple subsidized neighborhoods.

[11]The analytical definition of that diversion ratio thus is $DR_{A,B} = \frac{\partial D^B/\partial p^A}{\partial D^A/\partial p^A}$.

[12]Note that the case with one subsidized area and multiple unsubsidized areas is reflected in Equation 5. The right panel of Figure A.1 in Appendix A presents a graphical representation of the situation with multiple subsidized areas.

puted with the following formula derived in Appendix B:

$$\begin{aligned}\hat{\beta}_{DiD} \approx \underbrace{(p_2^A - p_1^A)}_{\text{Autarky in } A} \times \Bigg[1 + & \underbrace{\frac{\partial D^A}{\partial p^A} \times \frac{\partial P_S^A}{\partial q^A}}_{\text{Direct Spillover-on-Switchers Scaling}} - \underbrace{\frac{\partial D^A}{\partial p^A} \times \frac{\partial P_S^B}{\partial q^B} \times DR_{A,B}}_{\text{Direct Contamination Scaling}}\Bigg] \\ + \sum_{k \in \mathcal{K}} \underbrace{(p_2^k - p_1^k)}_{\text{Autarky in } k} \times \Bigg[& \underbrace{\frac{\partial D^k}{\partial p^k} \times \frac{\partial P_S^A}{\partial q^A} \times DR_{k,A}}_{\text{Indirect Spillover-on-Switchers Scaling}} - \underbrace{\frac{\partial D^k}{\partial p^k} \times \frac{\partial P_S^B}{\partial q^B} \times DR_{k,B}}_{\text{Indirect Contamination Scaling}}\Bigg]\end{aligned} \tag{6}$$

with $\mathcal{K}$ denoting the set of all subsidized neighborhoods excluding A.

Equation 6 has similar terms to before but also some differences. The first line is the same as in Equation 5. The second line includes two terms that capture the effects of the subsidy in all the other subsidized areas different from A. First, there is the indirect re-sorting spillover on switchers, i.e., households relocating away from neighborhood A into other subsidized neighborhoods. Since prices in other areas decrease, this indirect re-sorting moderates the price increase in A generated by the direct re-sorting spillover on switchers. Second, there is the indirect contamination effect. This captures the effect of the introduction of the subsidy in areas other than A on the price in neighborhood B. Therefore, the full contamination effect now is unequivocally larger than before. Overall, in the most typical case of more than one neighborhood being subsidized, the DiD estimator is even less accurate, suffering to a larger extent from the contamination effect.

2.4 Guidelines for Empirical Work

In this subsection, we discuss the main guidelines provided by the formulas above to researchers studying contexts with SUTVA violations due to re-sorting. First, Equations 4 to 6 show that the contamination effect biases the DiD estimate and provides the direction of the bias. Researchers can thus use the formula to know if contamination drives the estimate towards zero, as in our case, or may be a source of upward bias when the policy moves outcomes in opposite directions in treatment versus control areas.

Second, our formulas highlight the determinants of the bias. Knowing these determinants can help the applied researcher choose better control areas in contexts of re-sorting. The fact that contamination increases with the inverse housing supply elasticity of the unsubsidized area deserves special attention given the available evidence on neighborhood-level housing supplies being rather inelastic (Baum-Snow & Han, 2023). The literature already provides rich proxies of the relevant elasticities. On the supply side, Baum-Snow and Han (2023) show that housing supply is more elastic in places with more undeveloped land, flatter, and less regulated. On the demand side, researchers should look for control areas that consumers see as poor substitutes for the targeted areas. Data on relocation flows can help.

Third, in contexts with only one subsidized and one unsubsidized area (captured by Equation 5), the elasticities of supply and demand constitute sufficient statistics for the relative size of each effect (Saez, 2001; Chetty, 2009). That is, before doing the study, if the researcher has supply and demand elasticities from the literature, she would be able to compute the relative size of the effects. Additionally, after doing the study, the formula in Equation 5 allows to combine DiD estimates with supply and demand elasticities to calculate all three terms in Equation 4. For small changes around equilibrium, it is possible to recover the ATT without estimating a structural model.[13]

Fourth, Equation 6 shows that with more than one subsidized area, the ATT cannot be computed anymore knowing only the supply and demand elasticities. To compute the three effects, the researcher would need to know the effect of the policy in autarky in all the subsidized neighborhoods. In some contexts, including ours, researchers can have only one subsidized and one unsubsidized area following an aggregation approach.

Finally, the formulas help to review the assumptions that allowed the previous literature to identify the effect of place-based policies in contexts of re-sorting. One strand of literature assumes that there is a sufficiently faraway area unaffected by the policy, and thus can be used as a "contamination-free control". Kline and Moretti (2014), Clarke (2017), Mayer and Trevien (2017), and Butts (2022) are prominent examples implementing this ring or donut approach. Equation 5 shows that this approach requires that the diversion ratio between the area of interest (A) and the control area (B) is zero ($DR_{A,B} = 0$).[14] One important limitation of the ring approach is that when policies are "large" all areas could be affected. The formula shows that one can use demand estimates to directly test for the hypothesis of the existence of an unaffected area.

A second strand of the literature can be seen as assuming that there is a large enough number of areas such that each area is too small to affect the rest through re-sorting. Examples of this strategy are Busso et al. (2013) and Chen et al. (2022). The formula in Equation 5 shows that this is equivalent to assuming that $\frac{\partial D^A}{\partial p^A} = 0$, implying that in these contexts the DiD estimate captures only the autarky effect.

3 Institutional Context and Data

3.1 Institutional Context

The policy we analyze is a typical tax break for residential investment in lagging urban areas, similar to the Opportunity Zones (OZ) program in the US. In contrast to the OZ tax breaks, which might be directed to commercial or residential development, the one

[13]Munro et al. (2025) highlight how in certain settings, such as online marketplaces, researchers can estimate the necessary elasticities with small random price perturbations.

[14]Note that in principle we only need $\partial D^B / \partial p^A = 0$. However, we express this assumption in terms of the diversion ratio because it expresses the substitution in percentage terms and is typically used in the demand estimation literature (e.g. Conlon and Mortimer, 2021).

we analyze is only directed at residential development. We refer to the policy by its familiar Spanish acronym, "LVIS" (*Ley de Vivienda de Interés Social*). Although the name of the policy refers to the promotion of social housing, LVIS units did not have to be occupied by low-income households and could be freely sold at market prices during the period we analyze.

Tax breaks in LVIS are quite large. González-Pampillón (2022) estimates that the total tax benefits equal 20% of the cost of the projects. The main tax break is the complete exemption from the 22% value-added tax on inputs. LVIS projects are also fully exempt from the country's corporate tax of 25%, and units devoted to the rental market are partially exempt from income and wealth taxes.

The law that created LVIS was approved by the Uruguayan parliament in August 2011. Its implementation details, including the designation of the subsidized zones, were only defined in October of that year. Therefore, we take October 2011 as the starting date of the policy. The policy was substantially modified in June 2014, adding price ceilings and other restrictions that made the program less attractive to investors. Because these modifications substantially changed the potential impact of the policy on housing prices, our analysis ends in May 2014.

We study the impact of LVIS tax breaks in the department of Montevideo, which holds the homonymous 1.3 million capital city of Uruguay and concentrated 70% of LVIS projects during our period (Berrutti, 2017). LVIS in Montevideo subsidized residential development in low- and middle-income neighborhoods. The left panel of Figure 3 presents a map of subsidized and unsubsidized areas in the urban territory of the Montevideo department. The area without subsidies is located along the southeast coast of the city, by the Rio de la Plata river, and concentrates most of the middle- and high-income neighborhoods. The subsidized area covers almost three-quarters of urban Montevideo, including the central and older areas of the city as well as working-class neighborhoods.

The borders of the policy were defined jointly by the Ministry of Housing, the Ministry of Economics and Finance, and the local government of Montevideo with the explicit intention of excluding high-income neighborhoods from the subsidies (González-Pampillón, 2022; Borraz et al., 2024). About half of the border coincides with one of the main avenues of the city, which has been historically the most important spatial division between low- and high-income neighborhoods in Montevideo. The other half follows minor streets within homogeneous neighborhoods. In the paper, we exploit this contrast between low and high heterogeneity across different parts of the border to obtain DiD estimates corresponding to more or less intense re-sorting.

The generosity of the tax breaks implied that the policy had a huge impact on the location of residential investment in Montevideo. Berrutti (2017) shows that the share of the subsidized area in terms of square meters of construction permits went from around 20% before the policy to more than 60% in its first three years. Another measure of the huge quantitative relevance of the policy is provided by González-Pampillón (2022), who estimates that the total amount of investment that benefited from the tax break during the first five years of the policy amounted to 1.5% of the country's GDP.

The mechanics of the law required developers to apply for tax benefits and obtain project approval prior to starting construction. Due to this process, combined with the usual construction timelines, the first LVIS projects were only completed in 2013, and initial LVIS property sales occurred in 2014, with the majority of sales taking place in subsequent years (González-Pampillón, 2022). Consequently, almost no LVIS projects and very few LVIS sales were finalized during our period of analysis. This timing implies that we capture the capitalization of anticipated lower construction costs into current housing prices, rather than the direct effects from completed subsidized units or associated amenities. Since a significant number of LVIS projects were approved and under construction—and these figures were publicly available—, during our period of study it was widely known that housing supply in targeted neighborhoods would expand substantially. Importantly, many LVIS properties were sold "*en pozo*" (pre-sale), meaning that transactions began before units were finished. This suggests that, in addition to the expectations channel, the policy may have had a more immediate impact on the market through actual pre-sale activity.

The existence of public data on developers' applications to obtain the LVIS tax break allows us to characterize the new housing supply generated by the policy as being provided by highly atomistic producers. Of the 1,073 projects presented until October 2022, the average firm had 0.1% of the projects and 0.1% of the housing units. The maximum share attained by any single firm was 1.9% and 2.0% of the number of projects and housing units, respectively. This scenario of atomistic suppliers motivates the perfectly competitive assumption for the supply side in our model and reinforces our hypothesis of a negative effect of the policy on the housing prices of subsidized areas.

3.2 Data

We use five sources of data. The first and most important is the universe of housing transactions from the National Registry Office in Uruguay for the period 2010-2014. These data include the exact price and day for each housing sale as well as a measure of the area transacted. Uruguay is a high-income country according to the World Bank classification and has the lowest levels of informality in the region. This database of registered housing transactions is thus representative of the highly formal housing market of Montevideo.

The transaction data includes a unique property number, allowing us to match each sale with its corresponding entry in the National Cadaster of Uruguay registry, our second data source. This matching gives us the exact location of the parcel where the property is located and a set of housing characteristics, including the property area. We use this area from the cadaster when the area in the sales data is missing. The cadaster data do not exist for the years we analyze, and thus we use the earliest dataset available, which corresponds to 2016. We remove the top and bottom percentiles of the area and price distribution of the transaction dataset to avoid our estimates being affected by extreme values.

Table 1 presents descriptive statistics on the transaction data, separately for the subsidized and unsubsidized sections of the city, and distinguishing before and after the introduction of the policy. Consistent with the policy subsidizing lagging areas, prices are lower in the subsidized areas. Housing prices grow over time in all areas due to a context of strong economic growth in Uruguay during this period.

Table 1: Housing Prices and Area by Subsidy Status in the Pre and Post Periods

	Pre		Post	
	Subsidized	Unsubsidized	Subsidized	Unsubsidized
Number of Transactions	$10,035$	$6,793$	$13,112$	$8,861$
Mean Square Meter Price (USD/m^2)	701	$1,446$	955	$1,894$
	(505)	(675)	(680)	(874)
Mean Transaction Size (m^2)	125	96	123	91
	(136)	(105)	(134)	(99)

Source: Authors' calculations using housing transaction data from the National Registry Office in Uruguay.
Notes: Standard deviations are provided in parentheses. Calculations in the "Pre" supra column correspond to the period between January 2010, when our data start, until September 2011, the month before the policy's starting date. Calculations in the "Post" supra column correspond to the period beginning in October 2011 and ending in May 2014. The "subsidized" and "unsubsidized" columns indicate the area in which the transaction occurred. Figure 3 presents a map of those two areas.

Throughout the paper, we use a set of housing characteristics as controls in various regression exercises that have the price of housing as the dependent variable. These control variables are obtained from the cadaster data except for the distance to the coast, which we compute using the exact location of the transaction. The set of housing characteristics from the cadaster includes the age of the property as well as a set of categorical variables indicating construction category, construction condition, type of ceiling, and if there is ongoing construction work on the property.

The third data source is a geo-coded map of the areas subsidized by LVIS, similar to Figure 3. This map allows us to assign a subsidized or non-subsidized status to each housing transaction in the city, and to calculate the exact distance of those transactions to the borders of the policy.

Our fourth and fifth data sources are Uruguay's population census and the main household survey, the Encuesta Continua de Hogares. We use these to construct and validate the neighborhood and nest structure employed in the discrete-choice demand estimation. Since Montevideo lacks administrative units that reflect meaningful variation in taxation or public service provision, we partition the city into contiguous, homogeneous units by grouping similar census tracts. We apply the SKATER spatial clustering algorithm developed by Assunção et al. (2006), evaluating tract similarity with the average years of education of their adult population in the 2010 census. We first cluster contiguous tracts into neighborhoods and then group similar but not necessarily contiguous neighborhoods into nests. Panel (a) of Figure 3 displays the resulting 30 subsidized and 19 unsubsidized neighborhoods and panel (b) their grouping into nests. We provide further detail on this two-stage clustering process in Appendix C.

We employ household survey data on relocation patterns within Montevideo during the three years preceding the policy to validate the nest structure. Table C.1 (Appendix C) shows that migration flows between adjacent neighborhoods are 23% higher when both belong to the same nest. This difference remains statistically and roughly constant controlling for origin and destination neighborhood fixed effects.

Figure 3: Neighborhood Classification by Subsidy Status and Nest

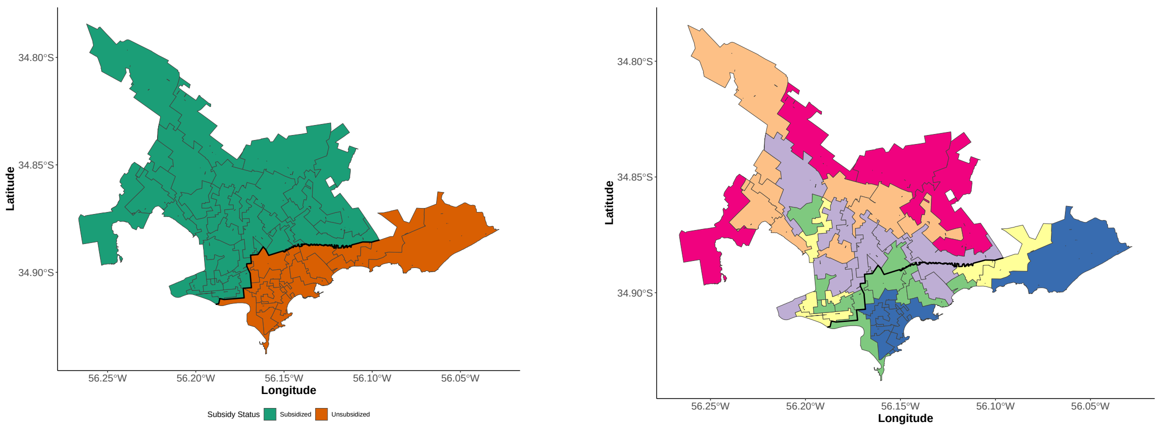

(a) Neighborhoods by Subsidy Status (b) Neighborhoods by Nest

Source: Authors' illustrations using official shapefiles from the Geomatic Service of Uruguay.
Notes: In panels a) and b), the thicker line shows the border of the policy and the thinner lines show the neighborhood limits. In panel b), the colors represent our grouping of neighborhoods into nests, which we use in the nested logit demand model. We defined neighborhoods and nests using a spatial clustering algorithm, as explained in Appendix C.

4 Difference-in-Differences Results

4.1 Benchmark Difference-in-Differences

We start by estimating the canonical DiD specification:

$$p_{ijt} = \gamma_j + \alpha_t + \beta Subsidy_j \times Post_t + f(X_{ijt}) + \epsilon_{ijt} \tag{7}$$

with p_{ijt} denoting the price per square meter of transaction i in neighborhood j at month t. Because each neighborhood is completely subsidized or unsubsidized, the neighborhood fixed-effect γ_j subsumes the $Subsidy_j$ term. $f(X_{ijt})$ is a third-order polynomial on the set of housing characteristics mentioned in the previous section.

Columns 1 to 3 of Table 2 present our first set of DiD estimates. The defining feature of this first set is that it implements the canonical DiD specification using all transac-

Table 2: Difference-in-Differences Regressions

	Dependent Variable:					
	USD per Square Meter					
	(1)	(2)	(3)	(4)	(5)	(6)
Subsidized × Post-Policy	-194^{***}	-178^{***}	-181^{***}	-1	-58^{*}	-61
	(31)	(26)	(27)	(52)	(32)	(38)
Housing Characteristics	-	✓	✓	-	✓	✓
Fixed Effect - Geography	Subsidized	Subsidized	Neighborhood	Subsidized	Subsidized	Neighborhood
Fixed Effect - Time	Post-Policy	Post-Policy	Year × Month	Post-Policy	Post-Policy	Year × Month
No. Obs	38,801	38,801	38,801	7,579	7,579	7,579
Data	City-Wide	City-Wide	City-Wide	500m Buffer	500m Buffer	500m Buffer
Pre-Policy Price per Square Meter	1,002	1,002	1,002	1,112	1,112	1,112

Source: Authors' calculations using housing transaction data from the National Registry Office in Uruguay and matched data on property characteristics from the country's national cadaster.
Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors are clustered at the neighborhood level and provided in parentheses. The "Housing Characteristics" controls consist of a polynomial of degree three on transaction area in square meters, building age in years, distance to the coast in meters, and indices of construction category, construction condition, type of ceiling, and if there is ongoing construction work on the property. The areas behind the "Subsidized" and "Neighborhood" fixed effects are shown in Figure 3. The 500-meter buffer restriction requires that the transaction is located less than 500 meters away from the border of the policy. This 500-meter buffer is shown in Figure A.2 in Appendix A.

tions in the city. Column 1 only has the three traditional DiD terms, the second column adds the controls, and the third adds month-year and neighborhood fixed effects. All three columns show negative estimates of very similar magnitude. This result is further confirmed graphically in Figure A.4 and Figure A.6 in Appendix A, which also show parallel pretrends between subsidized and unsubsidized areas. Our preferred estimate of -181 USD per square meter, in Column 3, is quite large, representing 18% of the average price per square meter before the policy.

4.2 Additional Estimates

A second set of estimates features commonly used techniques aimed at increasing the comparability between subsidized and unsubsidized areas to mitigate concerns regarding unobserved confounders (Baum-Snow & Ferreira, 2015; Chen et al., 2022). Consistent with our framework, all the estimates in this subsection are significantly smaller in absolute value than those in the previous one.

The first and most common technique to maximize comparability between treatment and control areas is to restrict the estimating sample to units located right along the border of the policy (Neumark & Kolko, 2010; Chen et al., 2022). In their evaluation of the employment impacts of Enterprise Zones in the US, Neumark and Kolko (2010) state that "the ideal control group consists of areas economically similar to enterprise zones but lacking enterprise zone designation". The estimates in Columns 4 to 6 of Table 2 follow this approach by restricting the sample to a 500-meter buffer around the border. Figure A.2 in Appendix A provides a map of this buffer, and Figure A.5 and Figure A.7 present the usual DiD graphs. The pre-policy price levels on both sides of the border in Figure A.5 indicate that both areas are indeed very similar. Our preferred point estimate, in Column 6, is -61 USD per square meter with a standard error of 38. Thus, a researcher employing a border DiD design in this context would not be able to reject the hypothesis that the tax break had a null effect on the prices faced by consumers.

Table 3: Difference-in-Differences Regressions - Extensions

	Dependent Variable:					
	USD per Square Meter					
	(1)	(2)	(3)	(4)	(5)	(6)
Subsidized × Post-Policy	-90^{***}	-112	-79^{*}	-84^{*}	-113^{*}	-121^{***}
	(32)	(75)	(45)	(45)	(57)	(36)
Housing Characteristics	✓	✓	✓	✓	✓	✓
Fixed Effect - Geography	Neighborhood	Neighborhood	Neighborhood	Neighborhood	Neighborhood	Neighborhood
Fixed Effect - Time	Year × Month	Year × Month	Year × Month	Year × Month	Year × Month	Year × Month
No. Obs	38,801	4,384	7,579	6,982	6,619	7,442
Data:						
Subsidized Area	All	0-500m	0-500m	0-500m	0-500m	0-500m
Unsubsidized Area	All	0-500m	0-500m	500-1000m	1000-1500m	1500-2000m
Estimation Method	DiD with PScore	RD	RD-DiD	Ring-DiD	Ring-DiD	Ring-DiD

Source: Authors' calculations using housing transaction data from the National Registry Office in Uruguay and matched data on property characteristics from the country's national cadaster.
Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors are clustered at the neighborhood level and provided in parentheses. The "Housing Characteristics" controls consist of a polynomial of degree three on transaction area in square meters, building age in years, distance to the coast in meters, and indexes of construction category, construction condition, type of ceiling, and if there is ongoing construction work on the property. The "DiD with PScore" is implemented by re-weighting observations in the unsubsidized areas with weights obtained from a probit model of receiving the subsidy. The characteristics used in that model coincide with the ones used as controls in all regressions. The RD and the RD-DiD are estimated with a second-degree polynomial on the distance to the border. The RD uses only data for the period after the subsidy was introduced. The "Subsidized" and "Unsubsidized Area" rows indicate the distance to the border of the policy required for each transaction to be considered in the regression. For instance, Columns (2) and (3) consider transactions located in the 500-meter buffer around the border, which is shown in Figure A.2 in Appendix A. Columns (4) to (6) consider the same 500-meter buffer for transactions in subsidized areas but different buffers for those in unsubsidized areas. These alternative buffers are drawn in Figure A.8 in the Appendix A.

The first three columns of Table 3 introduce three additional common techniques that enhance the comparability between subsidized and unsubsidized areas. The first column features DiD with propensity-score reweighting (A. Smith & E. Todd, 2005; Aker, 2010; Wang, 2013; Chen et al., 2022), the second implements a border regression discontinuity instead of DiD (Holmes, 1998; Black, 1999; Bayer et al., 2007; Turner et al., 2014), and the third one estimates a difference-in-discontinuities design (Grembi et al., 2016; Butts, 2023b). All of these estimates are much smaller in absolute value than the ones obtained for the whole city.

Finally, we follow the popular ring approach (Di Tella & Schargrodsky, 2004; Kline & Moretti, 2014; Butts, 2022; Myers & Lanahan, 2022). If the heterogeneity between subsidized and unsubsidized areas grows with distance from the border, re-sorting and thus contamination should decrease, and according to our formula the DiD estimate should increase in absolute value. This is indeed the observed pattern in Columns 4, 5, and 6 in Table 3, which present DiD estimates for 500-1000, 1000-1500, and 1500-2000 meter rings, respectively.

This ring approach can identify the true effect of the policy on subsidized areas as long as the spillovers are zero after a certain distance from the border (Clarke, 2017; Butts, 2022; Myers & Lanahan, 2022). This requirement may not hold in many contexts because of two difficulties, which are present in our study. First, natural (sea, mountains) or human-made (park, highway) constraints may limit how far from the border one can go.[15] This is often the case in coastal cities, such as ours. Only 10% of our unsubsidized transactions are beyond 2,100 meters from the border (See Figure A.8 in Appendix A). Second, as noted by Butts (2022), when policies are large enough to induce the re-sorting

[15]The spillover-free area can be quite far in some cases. For instance, Clarke (2017) finds that the spillovers of text messaging bans extend for at least 30km.

of agents throughout the entire city, spillover-free areas may well not exist.

4.3 Difference-in-Differences with Heterogeneous Effects

Finally, we explicitly introduce heterogeneity in our DiD estimator by interacting the DiD term in the border specification with an index of price differences across the border. Figure A.9 in Appendix A illustrates how we compute that index. We draw a large number of 500-meter circles centered along the border and compute the median price per square meter for each side of the border within each circle. We then assign each transaction the weighted average of the difference between the two medians for all circles that contain the transaction, with weights given by the inverse distance to the center of the circle. We standardize the index by subtracting its mean and dividing by its standard deviation.

The second column of Table A.1 in Appendix A presents the estimate of the interaction between the DiD term and the heterogeneity index. A standard deviation increase in the heterogeneity of the border increases the absolute value of the DiD estimate by 55 USD per square meter. This is a large magnitude given our estimate of 181 USD for the whole city. Figure A.10 in Appendix A plots the implied relationship between the DiD estimate and the border heterogeneity index and shows how the 95% confidence interval includes zero for the whole bottom half of the index distribution.

In Section 7.2 we further discuss these findings by computing—for our estimated model—separately the contamination effect and the heterogeneous treatment effects. Using our model, we show that contamination does indeed correlate positively with both the degree of homogeneity across the border and with diversion ratios. Recovering contamination for the whole city further allows us to quantify the level of bias of the alternative DiD estimates presented in this section.

5 A Model for Quantifying Contamination

5.1 Demand

In our model, households make a discrete and exclusive choice regarding the neighborhood in which they are buying a generic housing unit (GHU) in Montevideo. This discrete set of geographical areas is complemented by an outside option consisting of buying a GHU in the localities belonging to the broader metropolitan area of Montevideo. Households choose the option that yields the highest indirect utility using Equation 8.

$$V_{ijt} = V(AM_{jt}, P_{jt}, \tilde{\epsilon}_{ijt}) \tag{8}$$

The first argument of the indirect utility function is the neighborhood amenity term AM_{jt}. Examples of such could be time-invariant, such as distance to the coast or major public infrastructure, or time-variant, such as restaurants, shops, or public transportation schedules. The second argument, P_{jt}, is the price per square meter of a GHU in neighborhood j at time t. $\tilde{\epsilon}_{ijt}$ denotes the unobserved preferences of consumer i at time t for neighborhood j. We parameterize indirect utility with the following linear function:

$$V(AM_{jt}, P_{jt}, \epsilon_{ijt}) = A_j + B_t - \alpha P_{jt} + \xi_{jt} + \tilde{\epsilon}_{ijt} = \delta_{jt} + \tilde{\epsilon}_{ijt} \tag{9}$$

The parametrization allows amenities to vary exogenously over time. Amenities AM_{jt} are the sum of a fixed component A_j, a city-wide time-varying component B_t, and a term ξ_{jt} that varies over time at the neighborhood level and is unobservable to the econometrician. We use a nested logit model that allows for controlling for correlated unobserved heterogeneity across neighborhoods within nests. We define $\tilde{\epsilon} = \zeta_{int} + (1 - \sigma) \times \epsilon_{ijt}$, where σ with $0 < \sigma \leq 1$ is the nesting parameter. ζ_{int} is common to all products in nest n. We assume $\zeta_{int} + (1 - \sigma) \times \epsilon_{ijt}$ follows a Type-1 Extreme Value distribution. Note that the within-nest correlation goes to one as σ approaches one, and that for $\sigma = 0$ the within-nest correlation goes to zero, and in this case we return to the standard logit model.

The mean utility of the outside option is normalized to zero in every period (i.e. $\delta_{0t} = 0 \; \forall \, t$). Following Berry (1994), this structure yields a linear equation where s_{jt} is the market share of area j in the whole market and $\bar{s}_{jnt}$ is the market share of product j in nest n.

$$\ln(s_{jt}) - \ln(s_{0t}) = \delta_{jt} = A_j + B_t + \xi_{jt} - \alpha P_{jt} + \sigma \ln(\bar{s}_{jnt}) \tag{10}$$

Note that Equation 10 is derived from a model that accounts for relocation equilibrium effects and thus is not subject to the SUTVA violation problem of the reduced-form DiD. As originally pointed out by Berry (1994) and highlighted by Allen and Arkolakis (2023) for contexts with spatial spillovers, the consistent estimation of this equation requires instrumenting for both the price and the nested-logit market share.

5.2 Supply

Perfectly competitive agents sell a total of Q_{jt} GHUs in neighborhood j at time t.[16] The perfect competition assumption implies that housing prices—net of taxes—equal marginal costs:

$$P_{jt} = (1 - \tau_{jt}) * MC(Q_{jt}). \tag{11}$$

Marginal costs increase with the number of houses sold. This reflects that land is

[16] Perfect competition is the standard assumption in the literature modeling housing supply at the city or neighborhood levels (Ahlfeldt et al., 2015; Baum-Snow & Han, 2023; Almagro & Domınguez-Iino, 2025). In Section 3 we provide evidence on the highly atomistic nature of developers in our context.

fixed in each neighborhood and, as a result of this scarcity, it becomes more valuable with consumers' willingness to pay for living in the neighborhood. Marginal costs also have a fixed component L_{jt} capturing neighborhood-specific aspects, such as the total land available for housing construction, as well as city-level aspects, such as shocks to construction costs.

Following previous literature, we parameterize the marginal cost function with the following functional form (Saiz, 2010; Diamond, 2016; Baum-Snow & Han, 2023):

$$MC(Q_{jt}) = L_{jt} \times Q_{jt}^{\eta} \tag{12}$$

Applying logarithms to both sides of Equation 12, and combining the resulting expression with Equation 11 yields the inverse housing supply curve:

$$\ln P_{jt} = \ln L_{jt} + \ln(1 - \tau_{jt}) + \eta \ln Q_{jt} \tag{13}$$

5.3 Parallel Trends and Contamination in the Structural Model

Roth and Sant'Anna (2023) have shown that functional forms are one of the main challenges to parallel trends. Our structural model relies on a number of specific functional forms, many of which are non-linear. Since we use this model to evaluate DiD estimates, we need to verify that it can generate data where the parallel trends assumption is not rejected at typical sample sizes. We evaluate this by simulating a series of equilibria of the model with varying parameters.

We present the details of those simulations in Appendix D and summarize here the two main conclusions we obtain. The first conclusion is that our model is able to generate data that satisfy the parallel trends assumption for certain regions of the parameter space, even though it is highly non-linear in both its supply and demand sides. Figure D.1 presents an example of parallel trends and the DiD estimator for a set of simulations.[17]

The second conclusion is that increasing the idiosyncratic variation in neighborhood amenities over time leads to more violations of parallel trends but reduces the degree of contamination of the DiD estimate. Results for different relative values of the model parameters are presented in Figure D.2 and D.3. Intuitively, when neighborhoods experience amenity shocks that are large relative to the size of the other shocks, this generates large changes in relative housing prices over time and leads to a rejection of the parallel trend assumption. On the other hand,—as suggested by the decomposition formula in Section 2 and the reduced-form results in Section 4—those amenity shocks make neighborhoods more heterogeneous and consumers re-sort less in reaction to the subsidy. This means less contamination and consequently a lower bias of the DiD estimate. These simulation results thus suggest that, in contexts of re-sorting, there is a trade-off

[17] Equilibrium prices computed from our structural model at the estimated parameters also exhibit parallel trends. Figures A.11 and A.12 in Appendix A present this evidence.

between satisfying the parallel trend assumption and having no SUTVA violations for certain regions of the parameter space.

6 Estimation

6.1 Demand

We estimate Equation 10 with a dataset that has a single quantity and price for each combination of neighborhood and month-year. Those prices correspond to a homogeneous housing unit across neighborhoods, and we obtain them by estimating a regression that controls for a third-degree polynomial on the rich set of housing characteristics introduced in Section 3. We explain this procedure in detail in Appendix E.

The A_j and B_t amenity terms in Equation 10 are captured by neighborhood and time fixed effects, respectively, and the time-varying, neighborhood-specific amenities ξ_{jt} constitute the error term. This term is correlated with the equilibrium prices and within-nest shares, which makes OLS estimates of Equation 10 inconsistent. We address this endogeneity by leveraging the introduction of the tax break as a supply shifter to build a set of four instruments. The first one is identical to the DiD term and indicates if the neighborhood has benefited from the subsidy at time t. The other three instruments capture how the supply shifter differentially affects each nest. These are formed by interacting the DiD term with the number of other neighborhoods in the same nest receiving the subsidy, their area in square meters, and the share of that area in the total area of the nest.

The identification assumption behind our set of instruments is that the tax break did not impact the time-varying amenities conditional on the set of fixed effects. This assumption deserves special attention given the abundant evidence on the effects of new housing supply on neighborhood amenities (Baum-Snow & Marion, 2009; Rossi-Hansberg et al., 2010; Diamond & McQuade, 2018), including evidence for the program we are studying (González-Pampillón, 2022; Borraz et al., 2024).

As we discuss in Section 3, almost no LVIS projects were completed during the period we study. Because of this, we do not expect any direct impact of the subsidy on the attractiveness of neighborhoods during our period. The policy could have still generated changes in amenities after the period we study, impacting their present value and violating our exclusion restriction. González-Pampillón (2022) shows that new LVIS housing projects had a positive effect on housing prices after the period we study but that these effects were highly localized. This implies that the area benefited by the projects' spillovers constituted a very small share of the total subsidized area. That area was not only small but also mostly still undetermined during the period we study, thus making it very hard for agents to anticipate it.

Table 4: Demand Estimation

	Dependent Variable: $ln(s_{jt}) - ln(s_{0t})$			
	(1)	(2)	(3)	(4)
Price per 100 Square Meters	0.02***	-0.00	-0.02***	-0.07***
	(0.00)	(0.00)	(0.01)	(0.01)
Within-Nest Log Market Share	0.66***	1.00***	0.72***	0.69***
	(0.01)	(0.01)	(0.27)	(0.04)
Observations	2,646	2,646	2,646	2,646
Method	OLS	OLS	IV	Simulated IV
Fixed Effect - Geography	-	Neighborhood	Neighborhood	Neighborhood
Fixed Effect - Time	-	Year × Month	Year × Month	Year × Month
K-P 1st stage F			0.71	21.6

Source: Authors' calculations using housing transaction data from the National Registry Office in Uruguay and matched data on property characteristics from the country's national cadaster.

Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors are provided in parentheses. All four regressions estimate Equation 10 at the neighborhood × month-year level. The first independent variable, price per 100 square meters, is obtained as the neighborhood × month-year fixed effects in a regression of transactions prices per square meter on those fixed effects plus a third-degree polynomial on housing age, area in square meters, distance to the coast, and four variables from the cadaster describing construction quality. The IV regression in Column (3) has four instruments. The first is identical to the DiD term and indicates if the neighborhood is being subsidized at time t. The other three capture how the supply shifter differentially affects each nest. These are formed by interacting the DiD term with: the number of other neighborhoods in the same nest receiving the subsidy, their area in square meters, and the share of their area in the total area of the respective nest. The IV regression in Column (4) uses the same instruments of Column (3) plus two additional ones. These are the equilibrium price and within-nest log market share for each neighborhood × month-year combination in a simulated equilibrium of the estimated model. See Subsection 6.1 for more details on that simulation.

To strengthen the first stage, we implement a three-step IV approach (Bayer et al., 2007; Wong, 2013; Allen et al., 2020; Almagro et al., 2022). The first step consists of obtaining regular IV estimates using the four instruments described above. In the second step, we use these estimates to solve for the model's equilibrium when all time-varying parameters, including amenities, are set to zero. Finally, in the third step, we re-estimate demand, adding the equilibrium prices and nest shares obtained in the second step to the set of instruments. Note that the additional two instruments are obtained in an equilibrium in which, by construction, time-varying amenities are set to zero and thus are not affected by changes in neighborhoods' attractiveness.

The first OLS estimate of the price coefficient in Column 1 of Table 4 is positive, which is consistent with prices being positively correlated with neighborhood amenities. The neighborhood and month-year fixed effects seem to remove part of the endogeneity, because the estimate of the price coefficient in Column 2 is still negative but much smaller, making it statistically indistinguishable from zero. Columns 3 and 4 present the regular and three-step IV estimates, respectively. Column 4 shows a negative and significant coefficient for the price, and a nested logit term coefficient satisfying the restriction of being between 0 and 1.

6.2 Supply

Our inverse housing supply in Equation 13 has two parameters: the subsidy τ_{jt} and the inverse supply elasticity η. For the first parameter, we externally calibrate it using González-Pampillón (2022)'s estimate of the LVIS subsidy representing 20% of the final housing price. For the second parameter, in our main approach we use the DiD estimate as an additional moment to estimate the full supply and demand model. Specifically, we calibrate η such that there is an exact match between our benchmark reduced-form DiD estimate of -181 USD per square meter and its structural counterpart.[18]

The structural counterpart of the reduced-form DiD is the double difference in equilibrium prices between subsidized and unsubsidized neighborhoods and between the model with and without the subsidy. We solve for the model's equilibrium at the monthly level, thus mirroring the structure of our data, and taking the IV-estimated demand parameters and the calibrated supply parameters as inputs. The equilibrium computation also takes as inputs the amenities and marginal costs of the neighborhoods, which we obtain as the residuals from the housing demand and supply equations, respectively. We focus our equilibrium comparisons on the period after the subsidy was introduced and obtain counterfactual equilibrium prices by setting the subsidy to zero. Matching the structural double difference in prices with the reduced-form DiD yields an inverse supply elasticity of $\eta = 0.33$.[19]

Robustness. We provide two robustness checks for the estimation of the inverse housing supply elasticity. Later, in Section 7.1, we present the results of the counterfactuals using the parameter estimates from these robustness checks. For the first robustness check, we internally calibrate η in the same way but allow amenities to change between the equilibria with and without the subsidy, following their observed evolution. Figure A.14 in Appendix A shows that amenities increase on average by 14.8% in both regions after the introduction of the subsidy. Thus, in this first robustness, we let amenities grow by that magnitude between the equilibria without and with the subsidy.[20] The inverse supply elasticity obtained in this context is 0.25, very similar to the benchmark result.

In the second robustness check, we estimate η using Equation 13 and employ a demand shifter as an instrument. Specifically, we use the time-varying amenities (ξ_{jt}) as an instrument for the quantity in Equation 13. The identifying assumption is that these amenities are uncorrelated with the changes in the local construction costs. Table A.2

[18]This internal calibration procedure mirrors the one implemented by Berger et al. (2022) in their study of market power in the US labor market.

[19]This calibrated parameter implies a more elastic housing supply compared to available estimates (Saiz, 2010; Alves, 2021; Baum-Snow & Han, 2023). Ours is a monthly-level elasticity referring to property owners' decisions to sell their houses. This implies that we look at a short-term selling decision. In contrast, the available estimates in the literature are measured over two or three decades and focus on new housing units, which take more time to produce and sell.

[20]Since the demand model has an outside option, this parallel increase in time-varying amenities can still affect the results of the calibration.

presents the estimates of η under different specifications, including the instrumental variables estimates. Our preferred specification yields an estimate of $\eta = 0.29$, also very similar to our calibrated result.

7 Counterfactuals

We solve for a set of counterfactual equilibria of the estimated model to achieve three goals. First, we decompose a structural equivalent of our DiD estimate into the three components presented in Section 2. This gives us a measure of the bias in the benchmark reduced-form DiD estimate for the whole city. Second, we recover the incidence of the subsidy according to the model, and contrast it with the one obtained considering the alternative reduced-form DiD estimates. Third, we show that, as suggested by our decomposition formula in Section 2 and by the variety of reduced-form estimates in Section 4, neighborhood-level contamination is negatively correlated with the degree of heterogeneity between subsidized and unsubsidized areas, and positively correlated with diversion ratios.

7.1 DiD Decomposition and the Incidence of the Subsidy

Table 5 presents the decomposition of the DiD term and the incidence of the subsidy. The second column features the results obtained with the structural model and the first one their reduced-form counterparts, when available. Structural results are averages across the 32 months of the "post" period. The two DiD terms of the first row are identical by construction since we use this moment to calibrate the inverse housing supply elasticity parameter.

The five rows in the center of Table 5 present the decomposition of the DiD term following Equation 4. The ATT term is the difference in the average equilibrium prices of the subsidized neighborhoods with and without the subsidy. The autarky term is the change in average equilibrium prices across subsidized neighborhoods due to the subsidy without allowing for re-sorting between neighborhoods. We then calculate the spillover-on-switchers' term as the difference between the ATT and autarky terms. That effect is large, implying that the reduction in housing prices in the subsidized neighborhoods would have been much more pronounced if buyers had not reacted to the policy by re-sorting into these areas.

The contamination term is the most important since it measures the difference between the DiD term and the ATT. We quantify contamination as the difference in the average equilibrium prices of unsubsidized neighborhoods with and without the subsidy. The magnitude of contamination, of around a quarter of the ATT, indicates that the DiD term substantially underestimates the impact of the policy on the prices of the targeted neighborhoods.

The last row of Table 5 shows that the presence of contamination has implications for

Table 5: Decomposition of DiD Results Using the Structural Model

	Reduced-Form	Structural
DiD	-181	-181
ATT		-242
Autarky		-404
Spillover on switchers		162
Contamination		-61
Contamination/ATT		25.2%
Incidence	59.2%	79.1%

Source: Authors' calculations using counterfactual equilibrium exercises.
Notes: The DiD reduced-form estimate is taken from Column (3) of Table 2. The structural DiD coincides with the reduced-form estimate by construction. The structural ATT, autarky, spillover on switchers, and contamination terms are computed as the average difference in the equilibrium prices with and without the subsidy for different sets of neighborhoods in different counterfactuals. The ATT is computed for the regular counterfactual and subsidized neighborhoods only. The autarky term is also computed for subsidized neighborhoods but with a subsidized equilibrium in which households are not allowed to re-sort across neighborhoods. The spillover on switchers' term is the difference between ATT and autarky. The structural contamination considers the regular counterfactuals with and without the subsidy and is computed with unsubsidized neighborhoods only.

the conclusion regarding the incidence of the policy. We calculate the incidence as the effect on the prices of the subsidized neighborhoods divided by the subsidy.[21] While the incidence according to the structural model is 79%, the one calculated using the reduced-form DiD is 20 percentage points lower.

We illustrate the relevance of our incidence result with the average price faced by a consumer buying a housing unit in this city. The average price of houses in subsidized areas in the pre-period was 90,000 USD. If the subsidy had an incidence of 100%, implying that it was passed completely to consumers, each consumer would have saved 18,000 USD. However, tax breaks are typically not entirely reflected in prices, and it is therefore an important economic question to establish which share of the tax break reaches its intended beneficiaries. In our context, a researcher guided by the reduced-form estimate of the incidence (59.2%) would have concluded that our consumer saved around 10,649 USD. However, once contamination is considered, the incidence of 79.1% implies a savings of 14,238 USD. The difference of 3,589 USD amounts to 24.0% of Uruguay's GDP per capita in 2011, the year the policy was introduced.

Robustness. We compute the DiD decomposition under the alternative estimates of the inverse supply elasticity presented in Section 6.2. Table A.5 presents the results of this exercise. As indicated by the decomposition formula, contamination grows with the elasticity of housing prices to quantities. Contamination reaches 22.6% with IV-estimated elasticity ($\eta = 0.29$) and 9.7% when we calibrate the elasticity allowing amenities to grow over time ($\eta = 0.25$). Thus, our main conclusion on the existence of large contamination in the setting under consideration is robust to these different methods of

[21] We obtain the amount of the subsidy by applying the 20% rate over the price that results from evaluating the inverse supply curve (without the subsidy) of the subsidized neighborhoods at the quantities obtained in the equilibrium with the subsidy.

computing the inverse housing supply elasticity.

7.2 Determinants of Contamination and Bias

The previous analysis showed that contamination can lead to wrong conclusions about the effect of a place-based policy. To guide applied work in other contexts, it is useful to understand when contamination may matter more. We next show that the joint consideration of our decomposition formula, reduced-form estimates, and structural decomposition results consistently indicates that contamination increases with the intensity of demand-side re-sorting, which in turn correlates with the similarity between subsidized and unsubsidized areas. This implies that, conditional on having parallel pre-trends, applied researchers should prefer comparisons between less homogeneous areas when place-based policies may induce substantive re-sorting.

Panel a) of Figure 4 shows the positive correlation between contamination and demand-side re-sorting. We plot, for every pair of subsidized and non-subsidized neighborhoods along the border of the policy, the structural contamination as a share of ATT against the heterogeneity index introduced in Section 4. Going back to the reduced-form relationship between the border DiD estimate and the degree of heterogeneity across the border presented in both Table 3 and Figure A.10, the results in Figure 4 indicate that contamination can explain why one may not reject the hypothesis that the policy had zero effects when comparing very homogeneous areas.

The second piece of evidence, presented in panel b) of Figure 4, focuses on the whole city and shows how contamination is strongly and positively correlated with diversion ratios. Consistent with our decomposition formula, the correlation not only has the expected sign but it is also linear. Leveraging that we have the two terms for all neighborhoods and all months in the post period, we estimate the regression equivalent of the figure in panel b) of Figure 4, including a rich set of controls. The results in Table A.3 in Appendix A show a robust and positive relationship when controlling for none, either, and both neighborhood and month $\times$ year fixed effects.

Figure 4: Contamination, Border Heterogeneity and Diversion Ratios

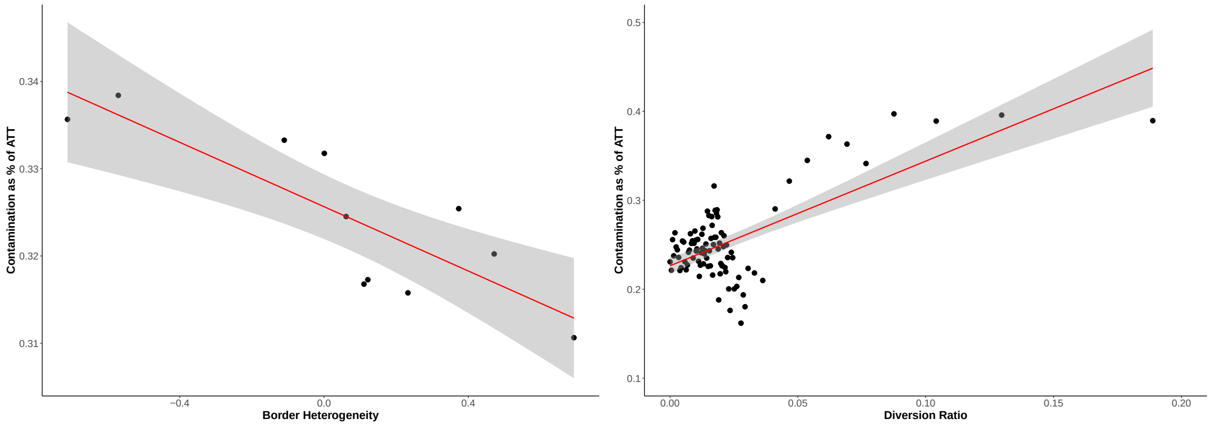

(a) Contamination and Border Heterogeneity (b) Contamination and Diversion Ratios

Source: Authors' calculations using counterfactual equilibrium exercises.
Notes: The y-axis in both panels presents contamination as a percentage of the ATT. Contamination is obtained as the difference in the housing prices of unsubsidized neighborhoods between the equilibria with and without the subsidy. The ATT is computed analogously but for subsidized neighborhoods. The straight red line represents the predicted value from a linear regression of the y-variable on the x-variable and the grey area is the 95% confidence interval for that prediction. In panel a), the 13 dots represent all the neighborhood pairs lying across the border of the policy. The x-axis shows the average heterogeneity index (introduced in Section 4) for the transactions belonging to the two neighborhoods of the pair. In panel b), the dots represent all the subsidized-unsubsidized neighborhood pairs. The x-axis shows the diversion ratio between the pair, calculated as the quotient between two partial derivatives, both derivatives taken with respect to the price of the subsidized member of the pair. The numerator of that quotient takes the partial of the demand of the unsubsidized member of the pair and the denominator the partial of the demand of the subsidized one.

Finally, our formula states that not only contamination but also the ATT are correlated with the intensity of demand-side substitution. Since the spillover-on-switchers term is part of the ATT, more of it would lead to lower DiD estimates of the impact of the subsidy. Similarly to Figure 4 above, Figure A.13 in Appendix A shows that the absolute value of the ATT effectively increases with the degree of heterogeneity between the neighborhoods across the border. Although this relationship is not relevant as a source of bias, it matters for applied work for two reasons. First, if ATT effects are heterogeneous due to re-sorting, applied researchers focusing on very homogeneous areas would get systematically lower estimates. Second, and more substantively, a large spillover-on-switchers effect can be normatively relevant when higher prices offset part of the benefits of the subsidy for incumbent households.

7.3 The Role of Amenities

We next examine how the paper's main conclusions are affected when allowing for changes in amenities in response to the policy. There is evidence that new construction can increase surrounding housing prices (Baum-Snow & Marion, 2009; Diamond & McQuade, 2018), including for the policy we study (González-Pampillón, 2022). In our benchmark model, amenities vary over time but do not respond endogenously to the policy.[22] As discussed in Section 6.2, Figure A.14 shows that amenities in our setting grow over time and that our estimates of η are robust to controlling for this growth.

Figure A.15 applies the graphical decomposition from Section 2 to illustrate how improved amenities in the subsidized areas would affect the DiD estimate. Two forces reduce the absolute value of the DiD. First, higher amenities in subsidized areas raise prices there, lowering the absolute value of the ATT. Second, higher amenities would induce a new round of relocation from unsubsidized into subsidized areas, further reducing prices in the latter and increasing contamination.

Table A.4 quantifies these effects. A gentrification scenario with a 20% amenity increase in the subsidized neighborhoods reduces the DiD term from -181 in our benchmark to -40.

These results have a first-order implication for interpreting our reduced-form estimates: a relatively low incidence could, in principle, reflect changes in amenities rather than contamination. However, two pieces of evidence allow us to refute this hypothesis. First, following the logic of Turner et al. (2014), positive spillovers from new construction should increase prices just across the border in non-subsidized areas. This implies lower contamination and higher incidence, exactly the opposite of what we observe. Second, DiD estimates should decline in absolute value as control areas farther from the border are considered, because positive spillovers from subsidized areas would dissipate with distance. Again, we find the opposite pattern.

8 Conclusion

Violations of the stable unit treatment value assumption (SUTVA) are a common threat to the identification of the effects of place-based policies. Because these policies are typically not randomly assigned, their analysis relies on quasi-experimental methods, with difference-in-differences being one of the most important. We discuss how difference-in-differences estimates may not recover the effect of policies in contexts where the resorting of agents changes the equilibrium outcomes of non-targeted units.

We illustrate how SUTVA violations can have serious consequences for the welfare conclusions of studies of large place-based interventions. We provide guidelines for applied work to detect contexts in which this might be more of a concern, and how to recover the true effect of the policy, subject to the availability of supply and demand elasticity estimates. When these estimates are not available and our guidelines detect a concern, we emphasize that, conditioning on parallel trends, researchers should avoid focusing on narrow comparisons of homogeneous units.

We illustrate our argument by studying a large place-based policy aimed at boosting housing construction in lagging areas of Montevideo, Uruguay. In part because of our

[22] Recent work has made progress in endogenizing amenities in structural models (Almagro & Domınguez-Iino, 2025).

methodological focus and restricted period, our study does not constitute a complete evaluation of the effects of this policy. Future work can adopt a longer perspective in which the policy may induce dynamic responses in housing supply, housing demand, and endogenous urban amenities, which are not present in our short-run analysis and can alter the overall conclusions on the impact of the policy.

References

A. Smith, J., & E. Todd, P. (2005). Does matching overcome LaLonde's critique of nonexperimental estimators? *Journal of Econometrics*, *125*(1), 305–353.

Adao, R., Arkolakis, C., & Esposito, F. (2019). *General equilibrium effects in space: Theory and measurement* (tech. rep.). National Bureau of Economic Research.

Ahlfeldt, G. M., Redding, S. J., Sturm, D. M., & Wolf, N. (2015). The Economics of Density: Evidence From the Berlin Wall. *Econometrica*, *83*, 2127–2189.

Aker, J. C. (2010). Information from Markets Near and Far: Mobile Phones and Agricultural Markets in Niger. *American Economic Journal: Applied Economics*, *2*(3), 46–59.

Allen, T., & Arkolakis, C. (2023). Economic Activity across Space: A Supply and Demand Approach. *Journal of Economic Perspectives*, *37*(2), 3–28.

Allen, T., Arkolakis, C., & Takahashi, Y. (2020). Universal Gravity. *Journal of Political Economy*.

Almagro, M., Chyn, E. T., & Stuart, B. A. (2022). Urban Renewal and Inequality: Evidence from Chicago's Public Housing Demolitions.

Almagro, M., & Domınguez-Iino, T. (2025). Location sorting and endogenous amenities: Evidence from amsterdam. *Econometrica*.

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

Anagol, S., Ferreira, F., & Rexer, J. (2021). Estimating the Economic Value of Zoning Reform.

Aronow, P. M., & Samii, C. (2017). Estimating Average Causal Effects Under General Interference, with Application to a Social Network Experiment. *The Annals of Applied Statistics*, *11*(4), 1912–1947.

Asquith, B. J., Mast, E., & Reed, D. (2023). Local Effects of Large New Apartment Buildings in Low-Income Areas. *The Review of Economics and Statistics*, *105*(2), 359–375.

Assunção, R., Neves, M., Câmara, G., & Da Costa Freitas, D. (2006). Efficient Regionalization Techniques for Socio-Economic Geographical Units Using Minimum Spanning Trees. *International Journal of Geographical Information Science*, *20*, 797–811.

Bachmann, R., Ehrlich, G., Fan, Y., Ruzic, D., & Leard, B. (2023). Firms and collective reputation: A Study of the Volkswagen Emissions Scandal. *Journal of the European Economic Association*, *21*(2).

Baum-Snow, N., & Ferreira, F. (2015). Causal inference in urban and regional economics. In *Handbook of regional and urban economics* (pp. 3–68). Elsevier.

Baum-Snow, N., & Han, L. (2023). The microgeography of housing supply. *Journal of Political Economy*.

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.

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.

Been, V., Ellen, I. G., & O'Regan, K. (2019). Supply Skepticism: Housing Supply and Affordability. *Housing Policy Debate*.

Been, V., Ellen, I. G., & O'Regan, K. (2025). Supply Skepticism Revisited. *Housing Policy Debate*, *35*(1), 96–113.

Berger, D., Herkenhoff, K., & Mongey, S. (2022). Labor Market Power. *American Economic Review*, *112*(4), 1147–1193.

Berrutti, F. (2017). Place-based subsidies and location decisions: The case of Uruguay. *Revista de Economía*, *24*(1).

Berry, S. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 242–262.

Black, S. (1999). Do Better Schools Matter? Parental Valuation of Elementary Education. *The Quarterly Journal of Economics*, *114*(2), 577–599.

Borraz, F., Carozzi, F., González- Pampillón, N., & Zipitría, L. (2024). Local Retail Prices, Product Variety, and Neighborhood Change. *American Economic Journal: Economic Policy*, *16*(1), 1–33.

Borusyak, K., Dix-Carneiro, R., & Kovak, B. (2022). Understanding Migration Responses to Local Shocks.

Busso, M., Gregory, J., & Kline, P. (2013). Assessing the Incidence and Efficiency of a Prominent Place Based Policy. *American Economic Review*, *103*(2), 897–947.

Butts, K. (2022). JUE insight: Difference-in-differences with geocoded microdata. *Journal of Urban Economics*, 103493.

Butts, K. (2023a). Difference-in-Differences with Spatial Spillovers.

Butts, K. (2023b). Geographic difference-in-discontinuities. *Applied Economics Letters*, *30*(5), 615–619.

Chen, J., Glaeser, E., & Wessel, D. (2022). JUE Insight: The (non-)effect of opportunity zones on housing prices. *Journal of Urban Economics*, 103451.

Chetty, R. (2009). Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods. *Annual Review of Economics*, *1*(1), 451–488.

Clarke, D. (2017). *Estimating Difference-in-Differences in the Presence of Spillovers* (tech. rep. No. 81604). University Library of Munich, Germany.

Conlon, C., & Mortimer, J. H. (2021). Empirical properties of diversion ratios. *The RAND Journal of Economics*, *52*(4), 693–726.

Delgado, M. S., & Florax, R. J. G. M. (2015). Difference-in-differences techniques for spatial data: Local autocorrelation and spatial interaction. *Economics Letters*, *137*, 123–126.

Di Tella, R., & Schargrodsky, E. (2004). Do Police Reduce Crime? Estimates Using the Allocation of Police Forces after a Terrorist Attack. *American Economic Review*, *94*(1), 115–133.

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.

Diamond, R., & McQuade, T. (2018). Who Wants Affordable Housing in their Backyard? An Equilibrium Analysis of Low Income Property Development. *Journal of Political Economy*.

Dix-Carneiro, R., & Kovak, B. K. (2017). Trade Liberalization and Regional Dynamics. *American Economic Review*, *107*(10), 2908–2946.

Donaldson, D. (2015). The Gains from Market Integration. *Annual Review of Economics*, *7*(1), 619–647.

Fan, Y., & Yang, C. (2025). Estimating Discrete Games with Many Firms and Many Decisions: An Application to Merger and Product Variety. *Journal of Political Economy*.

Feyrer, J., Mansur, E. T., & Sacerdote, B. (2017). Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution. *American Economic Review*, *107*(4), 1313–1334.

Fu, C., & Gregory, J. (2019). Estimation of an Equilibrium Model With Externalities: Post-Disaster Neighborhood Rebuilding. *Econometrica*, *87*(2), 387–421.

González-Pampillón, N. (2022). Spillover effects from new housing supply. *Regional Science and Urban Economics*, *92*, 103759.

Grembi, V., Nannicini, T., & Troiano, U. (2016). Do Fiscal Rules Matter? *American Economic Journal: Applied Economics*, *8*(3), 1–30.

Hollingsworth, A., Jaworski, T., Kitchens, C., & Rudik, I. (2025). Economic Geography and Air Pollution Regulation in the United States. *Journal of Political Economy. Microeconomics*.

Holmes, T. J. (1998). The Effect of State Policies on the Location of Manufacturing: Evidence from State Borders. *Journal of Political Economy*, *106*(4), 667–705.

Huber, M., & Steinmayr, A. (2021). A Framework for Separating Individual-Level Treatment Effects From Spillover Effects. *Journal of Business & Economic Statistics*, *39*(2), 422–436.

Imbens, G. W., & Rubin, D. B. (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press.

Jardim, E., Long, M. C., Plotnick, R., Vigdor, J., & Wiles, E. (2024). Local minimum wage laws, boundary discontinuity methods, and policy spillovers. *Journal of Public Economics*, *234*, 105131.

Kline, P., & Moretti, E. (2014). Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority. *The Quarterly Journal of Economics*, *129*(1), 275–331.

Li, X. (2022). Do new housing units in your backyard raise your rents? *Journal of Economic Geography*, *22*(6), 1309–1352.

Manski, C. F. (1993). Identification of Endogenous Social Effects: The Reflection Problem. *The Review of Economic Studies*, *60*(3), 531–542.

Mayer, T., Mayneris, F., & Py, L. (2017). The impact of Urban Enterprise Zones on establishment location decisions and labor market outcomes: Evidence from France. *Journal of Economic Geography*, *17*(4), 709–752.

Mayer, T., & Trevien, C. (2017). The impact of urban public transportation evidence from the Paris region. *Journal of Urban Economics*, *102*, 1–21.

Munro, E., Kuang, X., & Wager, S. (2025). Treatment Effects in Market Equilibrium. *American Economic Review*.

Myers, K. R., & Lanahan, L. (2022). Estimating Spillovers from Publicly Funded R&D: Evidence from the US Department of Energy. *American Economic Review*, *112*(7), 2393–2423.

Neumark, D., & Kolko, J. (2010). Do enterprise zones create jobs? Evidence from California's enterprise zone program. *Journal of Urban Economics*, *68*(1), 1–19.

Redding, S. J., & Rossi-Hansberg, E. A. (2017). Quantitative Spatial Economics. *Annual Review of Economics*, *9*(1).

Rossi-Hansberg, E., Sarte, P.-D., & Owens, R. (2010). Housing Externalities. *Journal of Political Economy*, *118*(3), 485–535.

Rotemberg, M. (2019). Equilibrium Effects of Firm Subsidies. *American Economic Review*, *109*(10), 3475–3513.

Roth, J., & Sant'Anna, P. (2023). When is parallel trends sensitive to functional form? *Econometrica*.

Roth, J., Sant'Anna, P. H. C., Bilinski, A., & Poe, J. (2023). What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics*, *235*(2), 2218–2244.

Rudik, I., Lyn, G., Tan, W., & Ortiz-Bobea, A. (2022). The Economic Effects of Climate Change in Dynamic Spatial Equilibrium. *Conference papers*.

Saez, E. (2001). Using Elasticities to Derive Optimal Income Tax Rates. *The Review of Economic Studies*, *68*(1), 205–229.

Saiz, A. (2010). The Geographic Determinants of Housing Supply. *The Quarterly Journal of Economics*, *125*(3), 1253–1296.

Sobel, M. E. (2006). What Do Randomized Studies of Housing Mobility Demonstrate? *Journal of the American Statistical Association*, *101*(476), 1398–1407.

Turner, M. A., Haughwout, A., & van der Klaauw, W. (2014). Land Use Regulation and Welfare. *Econometrica*, *82*(4), 1341–1403.

Vazquez-Bare, G. (2023). Identification and estimation of spillover effects in randomized experiments. *Journal of Econometrics*, *237*(1), 105237.

Wang, J. (2013). The economic impact of Special Economic Zones: Evidence from Chinese municipalities. *Journal of Development Economics*, *101*, 133–147.

Wong, M. (2013). Estimating Ethnic Preferences Using Ethnic Housing Quotas in Singapore. *The Review of Economic Studies*, *80*(3 (284)), 1178–1214.

A Appendix: Figures and Tables

Figure A.1: Visual Representation of Re-Sorting with Two or Multiple Areas

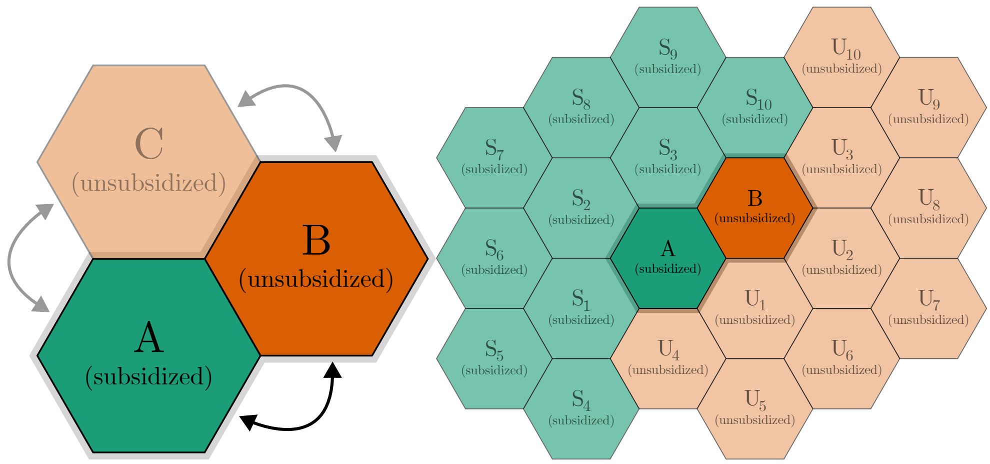

(a) Two Areas (b) Multiple Areas

Source: Authors' own illustration.

Figure A.2: Montevideo by Subsidy Status - 500m Buffer

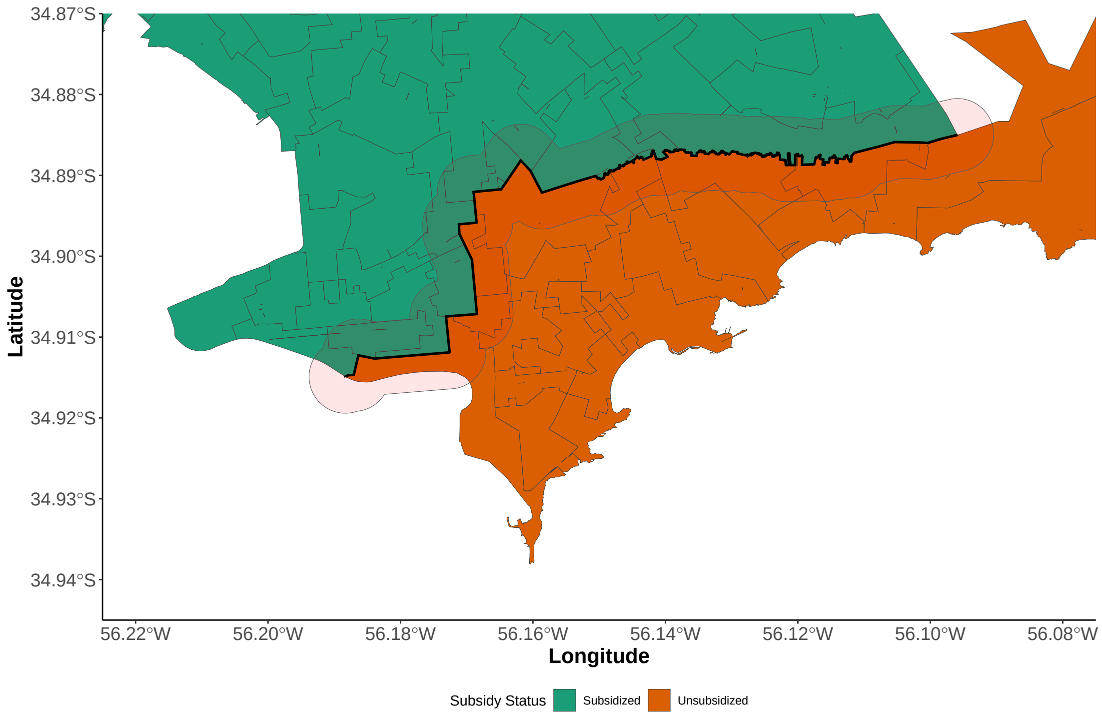

Source: Authors' own illustration using official shapefiles from the Geomatic Service of Uruguay.
Notes: The thicker line shows the border of the policy and the thinner lines the neighborhood limits. We defined neighborhoods using a spatial clustering algorithm, as explained in Appendix C. In panel a), the classification of neighborhoods into subsidized or unsubsidized follows the borders of the policy as defined in official government documents. The figure further displays a 500-meter buffer around the border of the policy.

Figure A.3: Average Years of Education by Census Tract

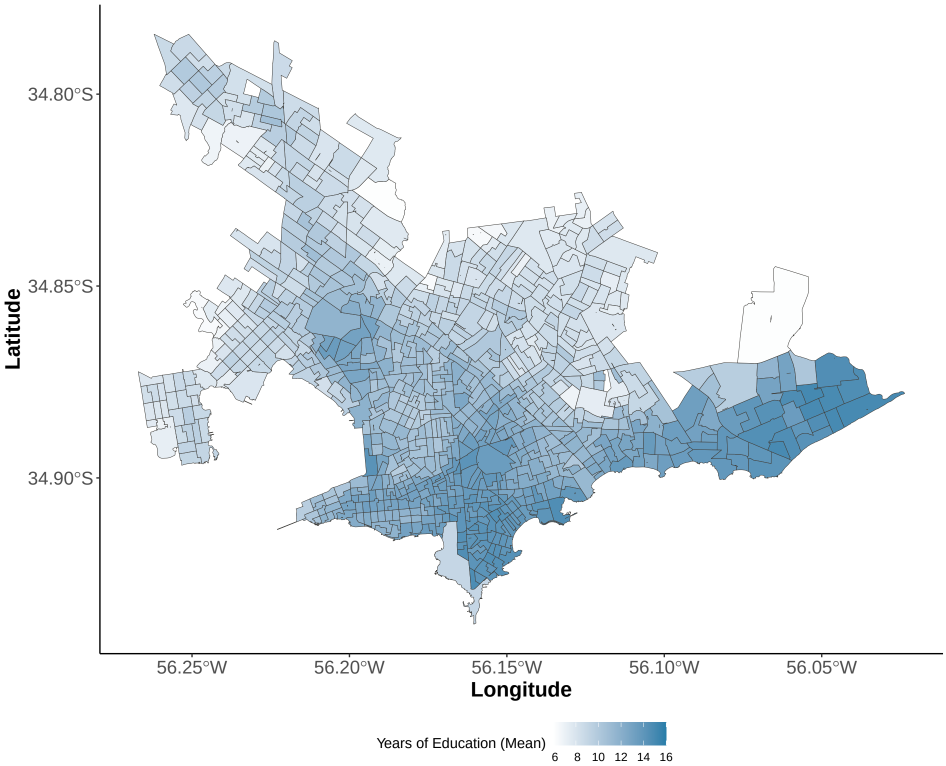

Source: Authors' illustration using official shapefiles from the Geomatic Service of Uruguay and microdata from the 2011 Uruguayan Census.
Notes: The tones of blue reflect the average years of education of the adult population living in each "segmento censal", an administrative unit comparable in size to a US census tract.

Figure A.4: Residualized Housing Prices by Subsidy Status - City-Wide

Source: Authors' calculations using housing transaction data from the National Registry Office in Uruguay and matched data on property characteristics from the country's national cadaster.
Notes: The graph plots, separately for transactions in the subsidized or unsubsidized areas, the average residualized price in each year-month. This residualized price is obtained as the residual of a regression of housing prices per square meter on a polynomial of degree three on transaction area in square meters, building age in years, distance to the coast in meters, and indexes of construction category, construction condition, type of ceiling, and if there is ongoing construction work on the property. The graph considers all housing transactions in the City of Montevideo.

Figure A.5: Residualized Housing Prices by Subsidy Status - 500m Buffer Across the Border

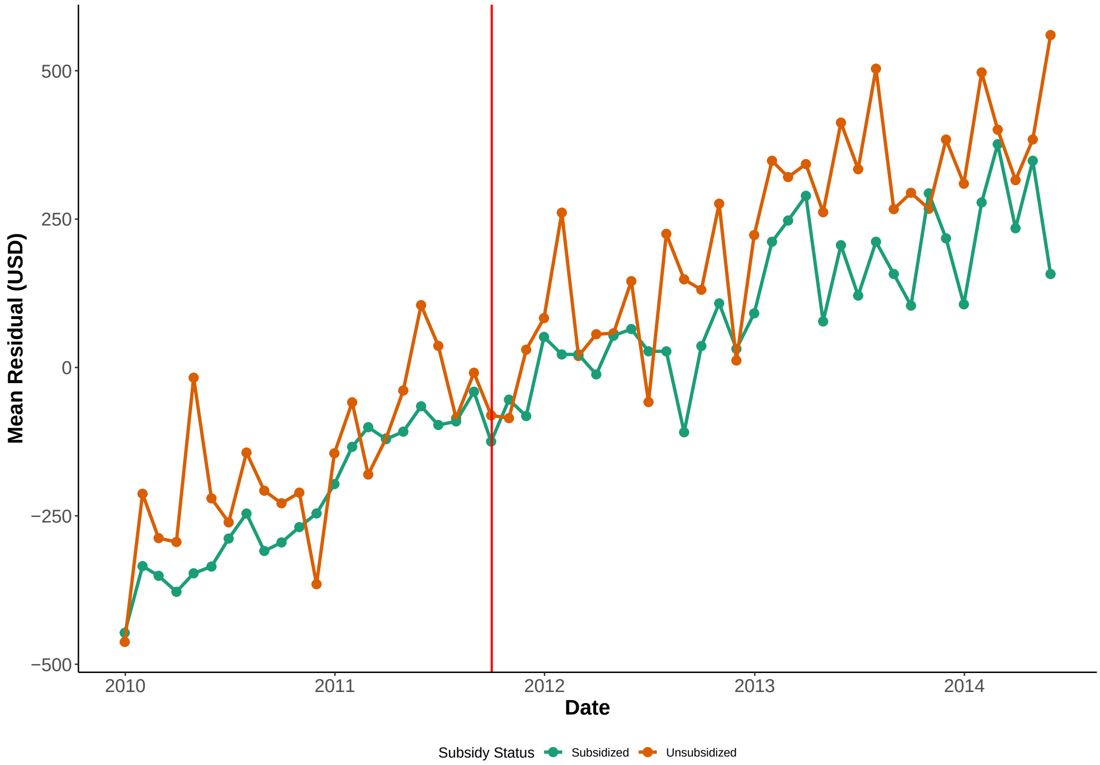

Source: Authors' calculations using housing transaction data from the National Registry Office in Uruguay and matched data on property characteristics from the country's national cadaster.
Notes: The graph plots, separately for transactions in the subsidized or unsubsidized areas, the average residualized price in each year-month. This residualized price is obtained as the residual of a regression of housing prices per square meter on a polynomial of degree three on transaction area in square meters, building age in years, distance to the coast in meters, and indexes of construction category, construction condition, type of ceiling, and if there is ongoing construction work on the property. The regression, and subsequently the graph, only considers transactions that are less than 500 meters away from the border of the policy. This 500-meter buffer is shown in Figure A.2 in Appendix A.

Figure A.6: Monthly Differences in Housing Prices Between Subsidized and Unsubsidized Areas Measured with respect to One Month Before the Starting Date of the Policy - City-Wide

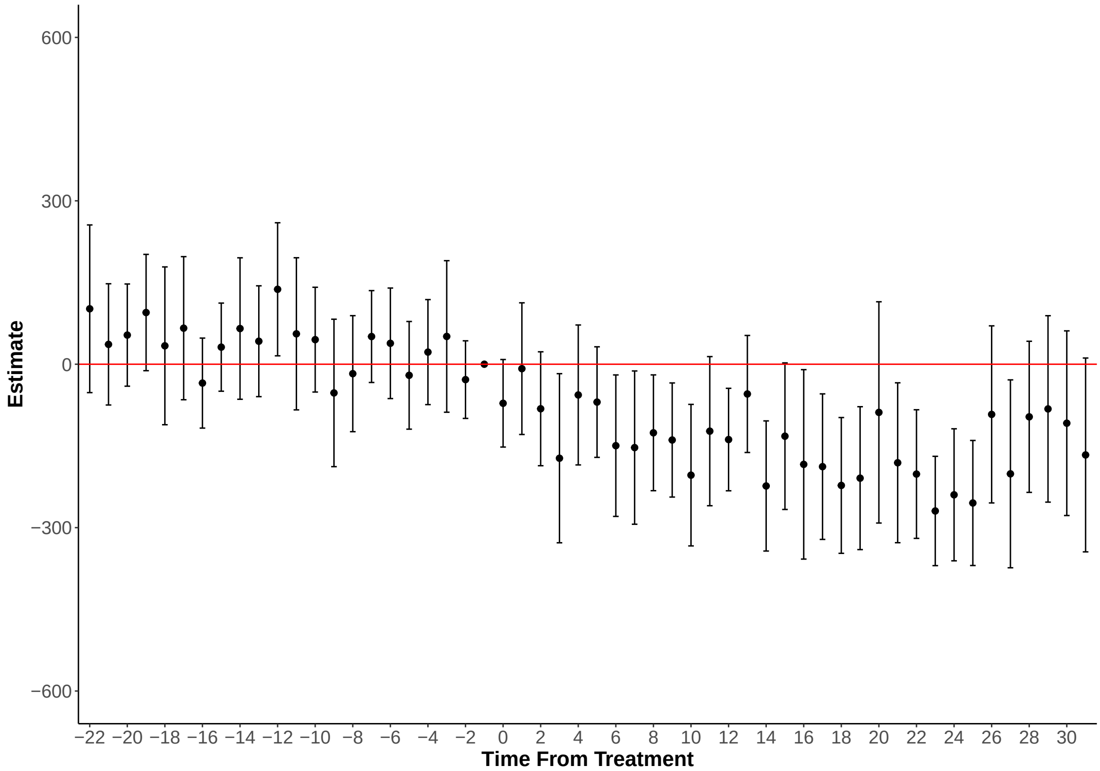

Source: Authors' calculations using housing transaction data from the National Registry Office in Uruguay and matched data on property characteristics from the country's national cadaster.
Notes: The graph plots the estimated coefficients, with their 95% confidence interval, of all year-month $\times$ subsidy dummies in a regression of housing prices per square meter on a polynomial of degree three on transaction area in square meters, building age in years, distance to the coast in meters, and indexes of construction category, construction condition, type of ceiling, and if there is ongoing construction work on the property. The regression, and consequently the graph, considers all housing transactions in the city. The omitted fixed effect is the month-year combination just before the starting date of the policy.

Figure A.7: Monthly Differences in Housing Prices Between Subsidized and Unsubsidized Areas Measured with respect to One Month Before the Starting Date of the Policy - 500m Buffer Across the Border

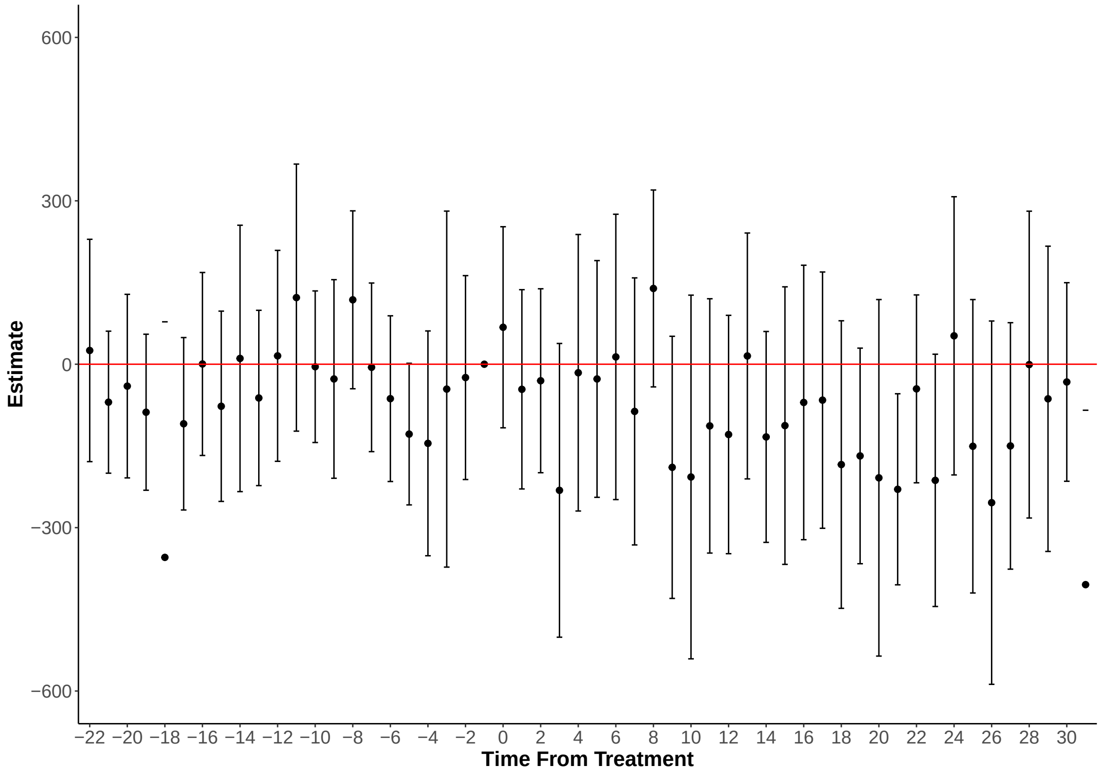

Source: Authors' calculations using housing transaction data from the National Registry Office in Uruguay and matched data on property characteristics from the country's national cadaster.
Notes: The graph plots the estimated coefficients, with their 95% confidence interval, of all year-month $\times$ subsidy dummies in a regression of housing prices per square meter on a polynomial of degree three on transaction area in square meters, building age in years, distance to the coast in meters, and indexes of construction category, construction condition, type of ceiling, and if there is ongoing construction work on the property. The regression, and subsequently the graph, only considers transactions that are less than 500 meters away from the border of the policy. This 500-meter buffer is shown in Figure A.2 in Appendix A. The omitted fixed effect is the month-year combination just before the starting date of the policy.

Figure A.8: Rings Around the Border of the Policy: Unsubsidized Area

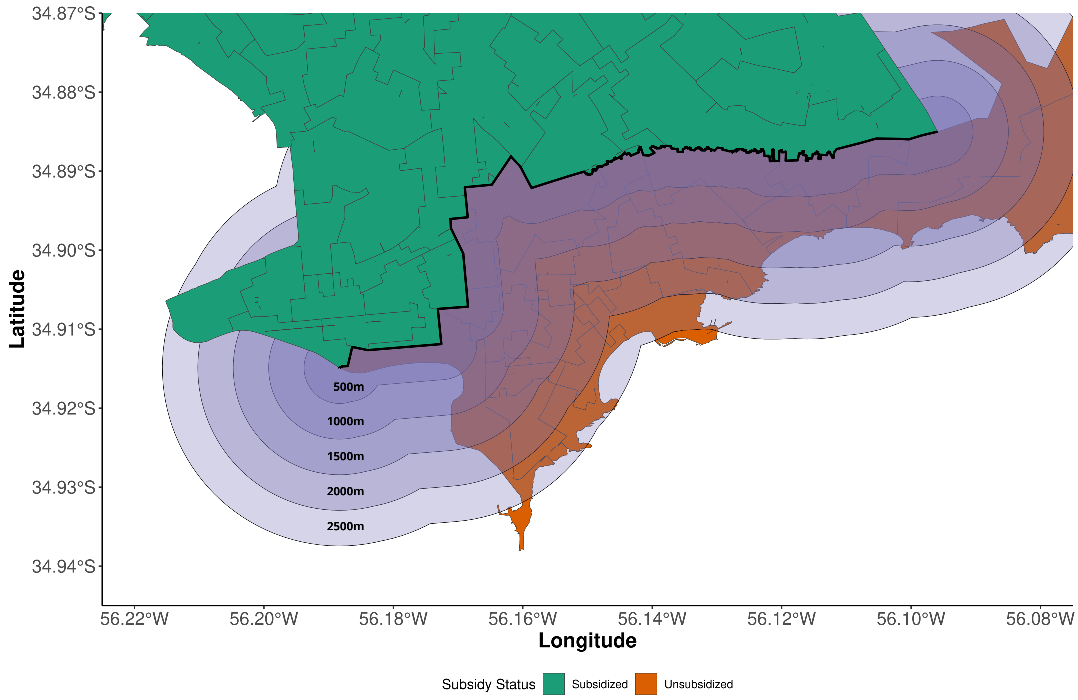

Source: Authors' own illustration using official shapefiles from the Geomatic Service of Uruguay.
Notes: The thicker line shows the border of the policy and the thinner lines the neighborhood limits. Each individual buffer covers the part of the unsubsidized area that is at most the distance indicated by the respective value in bold from the policy border. Larger buffer sizes naturally nest smaller buffer sizes.

Figure A.9: How Border Z-Scores are Computed

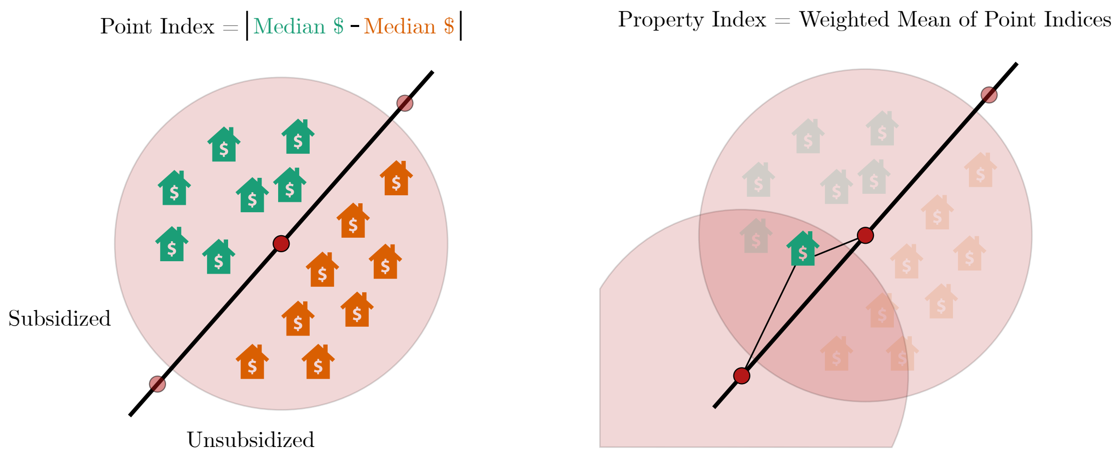

Source: Authors' own illustration.
Notes: The figure illustrates the method we use to compute a measure of heterogeneity along the border of the policy. The left panel shows how we compute the index of heterogeneity for a particular point on the policy border. The right panel shows how we aggregate point indices for individual properties. For more details on the calculation of this measure, see Section 4.

Figure A.10: Estimated Treatment Effect as a Function of Heterogeneity

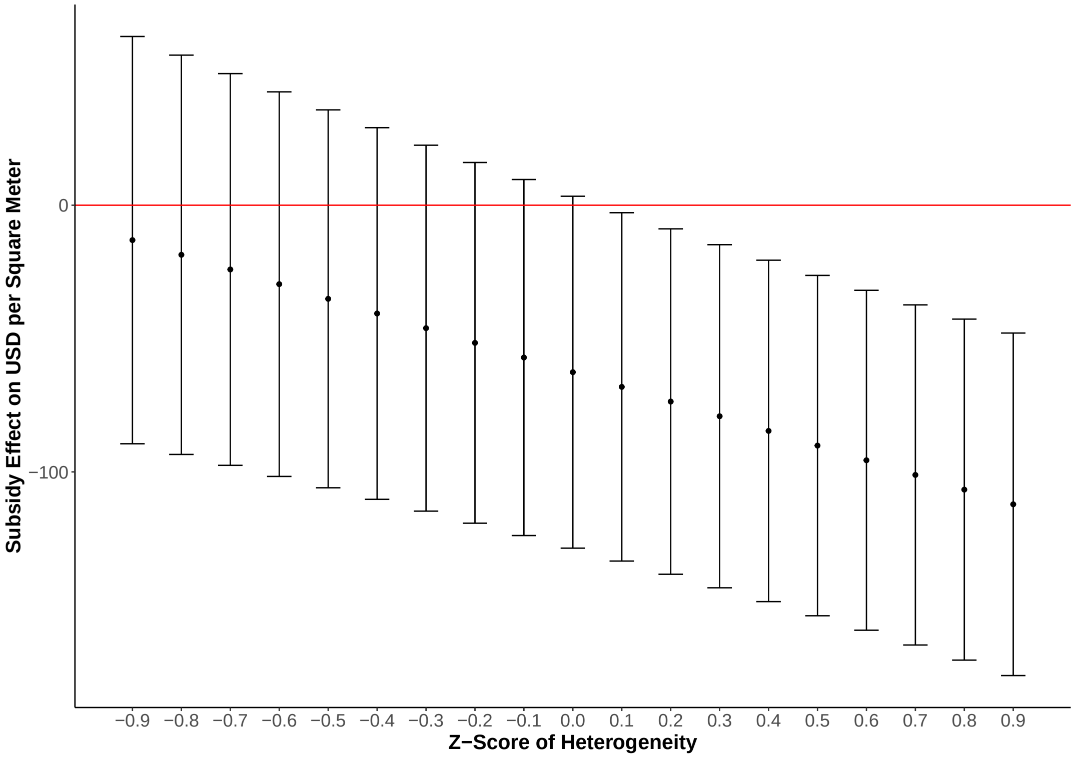

Source: Authors' calculations using housing transaction data from the National Registry Office in Uruguay and matched data on property characteristics from the country's national cadaster.
Notes: The graph plots the marginal effects, with their 95% confidence interval, for different values of the Z-score, of the interaction of that score with the difference-in-differences term in the regression estimated in Column (2) of Table A.1. This regression controls for neighborhood and year-month fixed effects polynomial of degree three on transaction area in square meters, building age in years, distance to the coast in meters, and indexes of construction category, construction condition, type of ceiling, and if there is ongoing construction work on the property. The regression is estimated using transactions located less than 500 meters away from the border of the policy. This 500-meter buffer is shown in Figure A.2 in Appendix A. The Z-score measures the average difference in housing prices between both sides of the border of the policy. For more details on the calculation of this measure, see Section 4.

Figure A.11: Average Housing Prices by Subsidy Status - Structural Model

Source: Authors' calculations using counterfactual equilibrium exercises.
Notes: The graph plots, separately for neighborhoods in the subsidized or unsubsidized areas, the average equilibrium prices for each year-month.

Figure A.12: Monthly Differences in Housing Prices Between Subsidized and Unsubsidized Areas Measured with respect to the Time Period One Month Before the Starting Date of the Policy - Structural Model

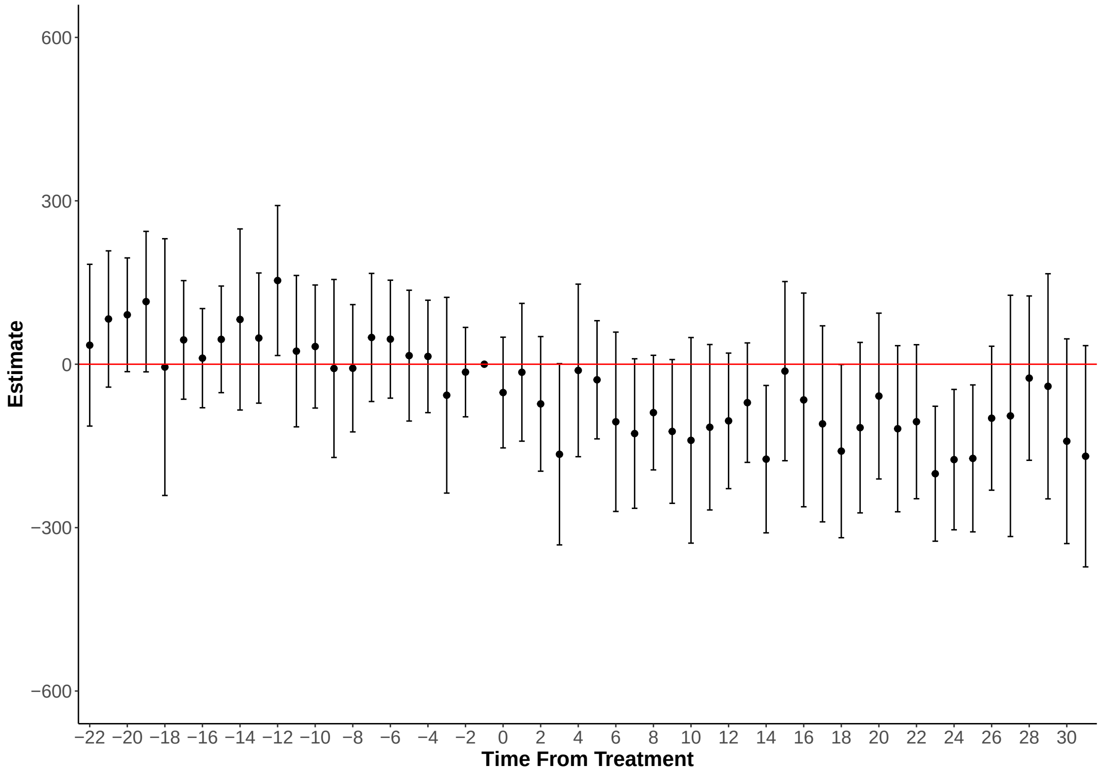

Source: Authors' calculations using counterfactual equilibrium exercises.
Notes: The graph plots the estimated coefficients, with their 95% confidence interval, of all year-month × subsidy dummies of a regression of equilibrium housing prices on month-year × subsidy dummies. The omitted fixed effect is the month-year combination just before the starting date of the policy.

Figure A.13: ATT and Border Heterogeneity - Structural Model

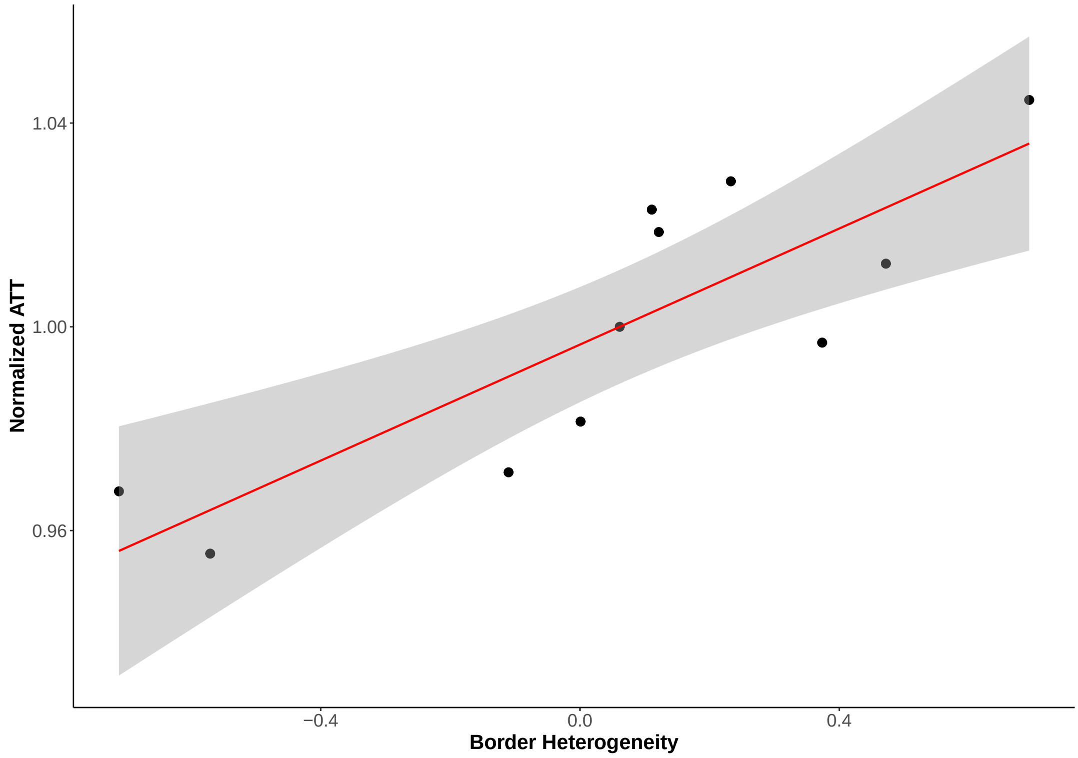

Source: Authors' calculations using counterfactual equilibrium exercises.
Notes: Each of the 13 dots in the figure represents a subsidized-unsubsidized neighborhood pair. These are all the neighborhood pairs lying across the border of the policy.Figure A.2 in the Appendix A provides a map of the neighborhoods with a focus on the border. The x-axis shows the diversion ratio. Using the estimated demand system presented in Table 4, the diversion ratio is calculated as the quotient between two partial derivatives, both of taken with respect to the price of the subsidized member of the pair. The numerator of that quotient takes the partial derivative of the demand of the unsubsidized member of the pair with respect to the price of the subsidized member and the denominator the partial derivative of the demand of the subsidized member with respect to its price. The y-axis presents the normalized ATT for the subsidized member of the pair. The ATT is obtained as the difference in the equilibrium housing prices in counterfactual scenarios with and without the subsidy for the subsidized member of each pair of neighborhoods. The normalization is performed by dividing by the average ATT across all neighborhoods. The straight red line represents the predicted value from a linear regression of the y-variable on the x-variable. The shaded grey area around it represents the 95% confidence interval around the predicted value.

Figure A.14: Evolution of Structural Amenities Across Subsidized and Unsubsidized Neighborhoods

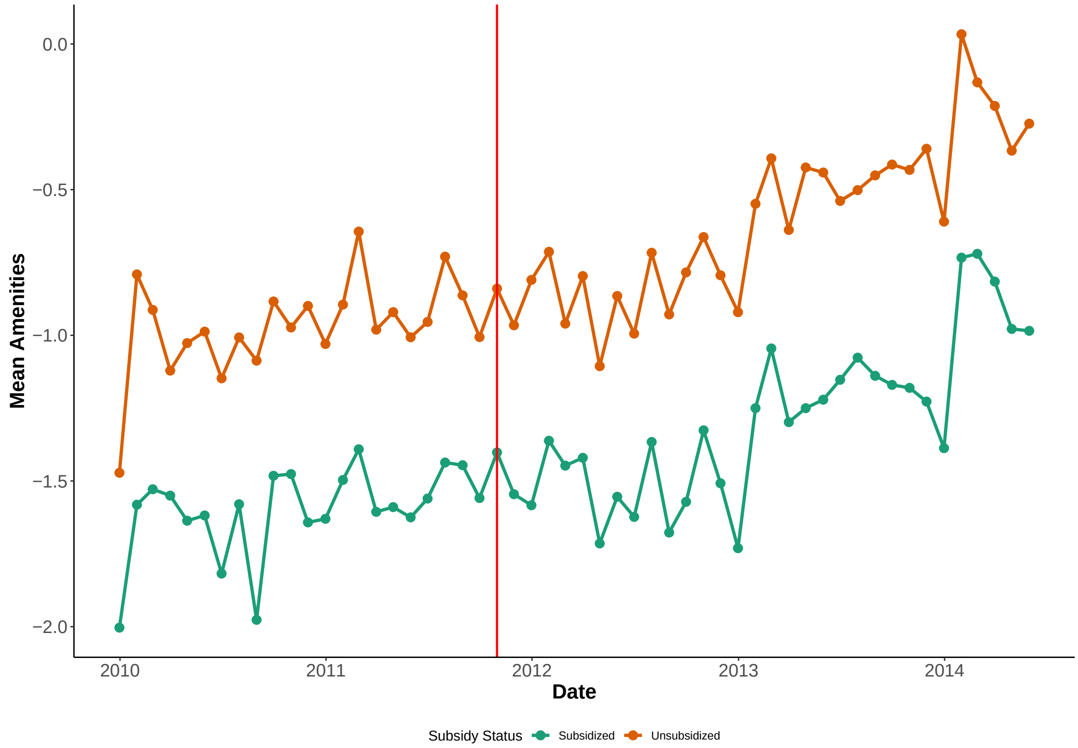

Source: Authors' calculations using counterfactual equilibrium exercises.
Notes: Amenities are obtained by removing the price effect from the mean utility of each product in each period. The individual lines represent the time series of the sales-weighted mean of these amenities by subsidy status.

Figure A.15: DiD with Re-Sorting between Neighborhoods A and B, and Improving Amenities in A

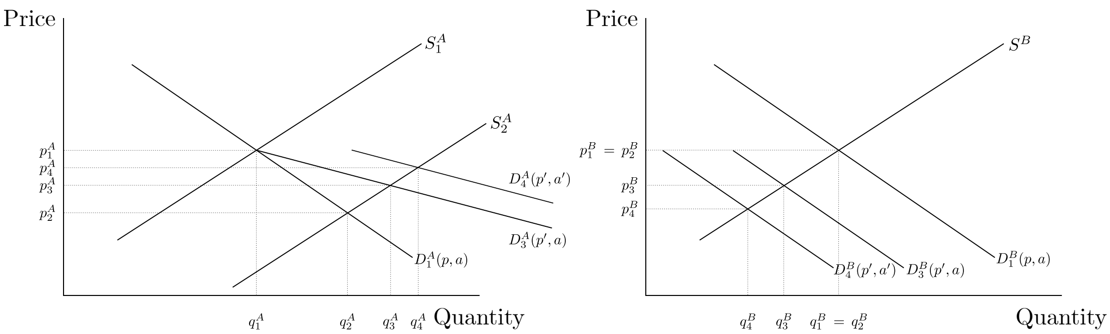

Notes: See section Section 2 for details on the notation and general logic behind the graphs. a and a' are two-element vectors featuring the amenity levels in areas A and B. The second vector has a higher value of amenities compared to the first one for area A and the same value for area B.

Table A.1: Difference-in-Differences Regressions - Heterogeneity

	Dependent Variable:	
	USD per Square Meter	
	(1)	(2)
Post × Treated	-61	-63
	(38)	(34)
Post × Treated × Z-Score	-	-55^{***}
		(14)
Housing Characteristics	✓	✓
Fixed Effect - Geography	Neighborhood	Neighborhood
Fixed Effect - Time	Year × Month	Year × Month
No. Obs	7,579	7,578
Data	500m Buffer	500m Buffer

Source: Authors' calculations using housing transaction data from the National Registry Office in Uruguay and matched data on property characteristics from the country's national cadaster.
Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors are clustered at the neighborhood level and provided in parentheses. The "Housing Characteristics" controls consist of a polynomial of degree three on transaction area in square meters, building age in years, distance to the coast in meters, and indexes of construction category, construction condition, type of ceiling, and if there is ongoing construction work on the property. The 500 meter buffer restriction requires that the transaction is located less than 500 meter away from the border of the policy. This 500 meter buffer is shown in Figure A.2 in Appendix A. The Z-score measures the average difference in housing prices between both sides of the border of the policy. For more detail on the calculation of this index see Section 4.

Table A.2: Supply Estimation

	Dependent Variable:			
	Logarithm of Price			
	(1)	(2)	(3)	(4)
Logarithm of Quantity	-0.007	-0.017**	2.115***	0.290***
	(0.014)	(0.008)	(0.170)	(0.031)
Observations	2,646	2,646	2,646	2,646
Method	OLS	OLS	IV	IV
Fixed Effect - Geography	-	Neighborhood	-	Neighborhood
Fixed Effect - Time	-	Year × Month	-	Year × Month

Source: Authors' calculations using housing transaction data from the National Registry Office in Uruguay and matched data on property characteristics from the country's national cadaster.
Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors are provided in parentheses. All four regressions estimate the inverse supply regression given by Equation 13. In that regression, observations are at the neighborhood × month-year level. In Columns (3) and (4), the instrumental variable used is the time-varying amenity ξ_{jt} from the demand regression.

Table A.3: Contamination and Diversion Ratio

	Dependent Variable:			
	Contamination			
	(1)	(2)	(3)	(4)
Diversion Ratio	2.57***	2.77***	2.51***	2.70***
	(0.07)	(0.08)	(0.06)	(0.07)
Observations	18,240	18,240	18,240	18,240
Fixed Effect - Geography	-	Neighborhood	-	Neighborhood
Fixed Effect - Time FE	-	-	Year $\times$ Month	Year $\times$ Month

Source: Authors' calculations using counterfactual equilibrium exercises.
Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors in parentheses. The four columns present the estimation results of a regression of contamination, measured in US dollars, on the diversion ratio. The observations in those regressions are all the possible pairs of subsidized-unsubsidized neighborhoods. Contamination is obtained as the difference in the equilibrium housing prices in counterfactual scenarios with and without the subsidy for unsubsidized neighborhoods. Using the estimated demand system presented in Table 4, the diversion ratio is calculated as the quotient between two partial derivatives, both of them taken with respect to the price of the subsidized member of the pair. The numerator of that quotient takes the partial of the demand of the unsubsidized member and the denominator takes the partial of the demand of the subsidized member.

Table A.4: Structural Decomposition when Amenities Change in the Subsidized Neighborhoods.

	DiD	ATT	Contamination	% Cont./ATT
	(1)	(2)	(3)	(4)
Benchmark	-181.1	-242.4	-61.4	25.3
Amenities deteriorate in subsidized neighborhoods:				
-5%	-215.9	-262.1	-46.2	17.6
-10%	-250.5	-282.0	-31.5	11.2
-15%	-284.9	-302.2	-17.3	5.7
-20%	-319.0	-322.6	-3.6	1.1
Amenities improve in subsidized neighborhoods:				
+5%	-146.0	-223.1	-77.0	34.5
+10%	-110.8	-204.0	-93.2	45.7
+15%	-75.4	-185.2	-109.8	59.3
+20%	-39.8	-166.7	-126.9	76.1

Source: Authors' calculations using counterfactual equilibrium exercises.
Notes: ATT is obtained as the difference in the average equilibrium housing prices across subsidized neighborhoods with and without the subsidy. Contamination is obtained analogously but for unsubsidized neighborhoods. DiD is ATT minus contamination.

Table A.5: Structural Decomposition under Different Methods of Obtaining the Inverse Supply Elasticity.

	η (1)	DiD (2)	ATT (3)	Cont. (4)	Cont./ATT (5)
Calibrated (Benchmark)	0.33	-181.1	-242.4	-61.4	25.3
Calibrated (Amenity Growth)	0.25	-181.1	-200.5	-19.4	9.7
Estimated	0.29	-179.3	-231.5	-52.2	22.6

Source: Authors' calculations using counterfactual equilibrium exercises.

Notes: ATT is obtained as the difference in average equilibrium housing prices in the subsidized neighborhoods with and without the subsidy. Contamination is computed analogously but for the unsubsidized neighborhoods. DiD is ATT minus contamination. In the "Benchmark" and "Estimation η" scenarios, amenities have the same value in the equilibrium with and without the subsidy. The scenario "Amenity Growth" introduces lower amenities in the pre-world counterfactual such that the difference in amenities between the pre and post-worlds equals the average change in the estimated amenities of the unsubsidized neighborhoods between the average of the whole post period and the month before the starting date of the policy.

B Appendix: Deriving the DiD Decomposition

The derivations for the approximation results of the generalized difference-in-differences (DiD) given in Equation 5 and Equation 6 are given below.

We specify demand for housing in a neighborhood j in time-period t at a given vector of market prices $\mathbf{p}_t$ to be given by $D^j(\mathbf{p}_t)$. The inverse housing supply function in neighborhood j in time-period t at quantity q_t^j is assumed to be given by $P_S^j(q_t^j)$. Inverse supply is thus only a function of within-neighborhood demand. Without loss of generality, we assume that the policy of interest is a housing construction subsidy implemented in neighborhood A while neighborhood B is not targeted by the policy.[23] The implied DiD empirical specification will always compare neighborhood A and neighborhood B.

Furthermore, we assume that equilibrium changes can be approximated by partial derivatives. We abstract away from any second- or higher-order effects. We start by assuming an exogenous shock (e.g. a subsidy) that moves the equilibrium to the new point (q_2, p_2) and then people react by re-sorting to the final equilibrium, through the demand effects. Please note that period $t = 1$ reflects the pre-policy equilibrium. Period $t = 2$ indicates the "artificial" time period in which the policy only affects the targeted neighborhood(s) in autarky. Period $t = 3$ is then the new post-policy equilibrium.

B.1 One Subsidized and One Unsubsidized

In reaction to the subsidy, the price in neighborhood A drop from p_1^A to p_2^A, with the corresponding change in quantities from q_1^A to q_2^A. In reaction to this exogenous change in (relative) prices, i.e. $(p_2^A - p_1^A)$, consumers in all neighborhoods re-evaluate their demand choices. The final change in equilibrium housing quantity in neighborhood A is given by Equation B.1, and in neighborhood B by Equation B.2.

$$q_3^A - q_2^A \approx \frac{\partial D^A}{\partial p^A} \times (p_2^A - p_1^A) \tag{B.1}$$

$$q_3^B - q_2^B \approx \frac{\partial D^B}{\partial p^A} \times (p_2^A - p_1^A) \tag{B.2}$$

Inserting these changes in equilibrium quantities into the local inverse housing supply equations, one can compute the changes in equilibrium prices.

$$\begin{aligned} p_3^A - p_2^A \approx & \frac{\partial P_S^A}{\partial q^A} \times (q_3^A - q_2^A) \\ \approx & \frac{\partial P_S^A}{\partial q^A} \times \frac{\partial D^A}{\partial p^A} \times (p_2^A - p_1^A) \end{aligned} \tag{B.3}$$

[23] In the traditional difference-in-differences (DiD) literature, neighborhood A would be considered the "treated unit" and neighborhood B would be the "control unit".

$$\begin{aligned} p_3^B - p_2^B \approx& \frac{\partial P_S^B}{\partial q^B} \times (q_3^B - q_2^B) \\ \approx& \frac{\partial P_S^B}{\partial q^B} \times \frac{\partial D^B}{\partial p^A} \times (p_2^A - p_1^A) \end{aligned} \tag{B.4}$$

Equation B.3 highlights three terms that determine the final price change in neighborhood A. First, it depends on the subsidy's "autarky" effect, i.e. $(p_2^A - p_1^A)$. Second, it is also determined by how price-sensitive housing demand in neighborhood A is with respect to the local price. Third, the responsiveness of local inverse supply also scales the change in final prices.

Similar to above, the size of the final price change in neighborhood B again depends on the subsidy's autarky effect in neighborhood A, and on the responsiveness of local inverse supply in neighborhood B. What, however, links the two neighborhoods is the partial derivative of demand for neighborhood B housing with respect to the price in neighborhood A. This partial derivative is a direct measure of demand substitution patterns between the two neighborhoods. If consumers do not consider these neighborhoods to be substitutes, this partial derivative is equal to zero. Thus, the local price neighborhood B does not change. If consumers, on the other hand, consider the two neighborhoods to be substitutes, this partial derivative is positive. The price in neighborhood B would then also change in reaction to the subsidy, despite the policy's scope being limited to neighborhood A.

Inserting these two expressions for final price changes into the generalised version of the DiD estimator given in Equation 4, we arrive at Equation 5.

$$\begin{aligned} \hat{\beta}_{DiD} =& (p_2^A - p_1^A) + (p_3^A - p_2^A) - (p_3^B - p_2^B) \\ \approx& (p_2^A - p_1^A) + \\ &+ (p_2^A - p_1^A) \times \frac{\partial P_S^A}{\partial q^A} \times \frac{\partial D^A}{\partial p^A} \\ &- (p_2^A - p_1^A) \times \frac{\partial P_S^B}{\partial q^B} \times \frac{\partial D^B}{\partial p^A} \\ \approx& (p_2^A - p_1^A) \times \left[1 + \frac{\partial P_S^A}{\partial q^A} \times \frac{\partial D^A}{\partial p^A} - \frac{\partial P_S^B}{\partial q^B} \times \frac{\partial D^B}{\partial p^A}\right] \\ \approx& \underbrace{(p_2^A - p_1^A)}_{\text{Autarky in } A} \times \left[1 + \underbrace{\frac{\partial D^A}{\partial p^A} \times \frac{\partial P_S^A}{\partial q^A}}_{\text{Spillover on the switchers Scaling}} - \underbrace{\frac{\partial D^A}{\partial p^A} \times \frac{\partial P_S^B}{\partial q^B} \times DR_{A,B}}_{\text{Contamination Scaling}}\right] \end{aligned} \tag{B.5}$$

with $DR_{A,B}$ being the diversion ratio between housing in neighborhood A and housing in neighborhood B. While the cross-price partial discussed previously is a non-normalized measure of substitutability between neighborhoods A and B, the diversion ratio is on the other hand a normalized measure of substitutability. It describes the ratio between the change in demand for neighborhood B and the change in the demand for neighborhood A when the price in A changes:

$$DR_{A,B} = \frac{\partial D^B / \partial p^A}{\partial D^A / \partial p^A} \tag{B.6}$$

B.2 Two Subsidized and One Unsubsidized

Building on the insights gained from Subsection B.1, we now add a third neighborhood C. Without loss of generality, we assume that neighborhood C is a neighborhood targeted by the policy and thus also subsidized.

Similar to before, the analysis starts with final changes in housing demand. The structure of Equation B.7 and others is very similar to above, with one exception. Because housing supply in neighborhood C is now also subsidized by the policy, an additional exogenous change in prices, i.e. $(p_2^C - p_1^C)$, needs to be accounted for when determining final demand changes.

$$q_3^B - q_2^B \approx \frac{\partial D^B}{\partial p^A} \times (p_2^A - p_1^A) + \frac{\partial D^B}{\partial p^C} \times (p_2^C - p_1^C) \tag{B.7}$$

$$q_3^A - q_2^A \approx \frac{\partial D^A}{\partial p^A} \times (p_2^A - p_1^A) + \frac{\partial D^A}{\partial p^C} \times (p_2^C - p_1^C) \tag{B.8}$$

$$q_3^C - q_2^C \approx \frac{\partial D^C}{\partial p^A} \times (p_2^A - p_1^A) + \frac{\partial D^C}{\partial p^C} \times (p_2^C - p_1^C) \tag{B.9}$$

Using the inverse supply equation for neighborhood B, one can derive an expression for the final price change in neighborhood B.

$$\begin{aligned} p_3^B - p_2^B &= \frac{\partial P_S^B}{\partial q^B} \times (q_3^B - q_2^B) \\ &= \frac{\partial P_S^B}{\partial q^B} \times \left(\frac{\partial D^B}{\partial p^A} \times (p_2^A - p_1^A) + \frac{\partial D^B}{\partial p^C} \times (p_2^C - p_1^C) \right) \end{aligned} \tag{B.10}$$

Using the same approach, we can derive an expression for $(p_3^A - p_2^A)$ using the inverse supply equation for neighborhood A.

$$\begin{aligned} p_3^A - p_2^A &= \frac{\partial P_S^A}{\partial q^A} \times (q_3^A - q_2^A) \\ &= \frac{\partial P_S^A}{\partial q^A} \times \left(\frac{\partial D^A}{\partial p^A} \times (p_2^A - p_1^A) + \frac{\partial D^A}{\partial p^C} \times (p_2^C - p_1^C) \right) \end{aligned} \tag{B.11}$$

Inserting these two expressions for final price changes into the generalised version of the DiD estimator given in Equation 4, one arrives at Equation B.12.

$$
\begin{aligned}
\hat{\beta}_{DiD} &= (p_3^A - p_1^A) - (p_3^B - p_1^B) \\
&= (p_3^A - p_2^A) + (p_2^A - p_1^A) - (p_3^B - p_2^B) \\
&\approx (p_2^A - p_1^A) \\
&\quad + (p_2^A - p_1^A) \times \left(\frac{\partial P_S^A}{\partial q^A} \times \frac{\partial D^A}{\partial p^A} - \frac{\partial P_S^B}{\partial q^B} \times \frac{\partial D^B}{\partial p^A} \right) \\
&\quad + (p_2^C - p_1^C) \times \frac{\partial P_S^A}{\partial q^A} \times \frac{\partial D^A}{\partial p^C} \\
&\quad - (p_2^C - p_1^C) \times \frac{\partial P_S^B}{\partial q^B} \times \frac{\partial D^B}{\partial p^C} \\
&\approx \underbrace{(p_2^A - p_1^A)}_{\text{Autarky in } A} \times \Bigg[1 + \underbrace{\frac{\partial D^A}{\partial p^A} \times \frac{\partial P_S^A}{\partial q^A}}_{\text{Direct Re-Sorting Scaling}} \underbrace{- \frac{\partial D^A}{\partial p^A} \times \frac{\partial P_S^B}{\partial q^B} \times DR_{A,B}}_{\text{Direct Contamination Scaling}} \Bigg] \\
&\quad + \underbrace{(p_2^C - p_1^C)}_{\text{Autarky in } C} \times \Bigg[\underbrace{\frac{\partial P_S^A}{\partial q^A} \times \frac{\partial D^A}{\partial p^C}}_{\text{Indirect Spillover on the switchers Scaling}} - \underbrace{\frac{\partial P_S^B}{\partial q^B} \times \frac{\partial D^B}{\partial p^C}}_{\text{Indirect Contamination Scaling}} \Bigg]
\end{aligned}
\tag{B.12}
$$

The final rewriting of the generalized version of the DiD estimator yields the same decomposition as in Subsection B.1 alongside one additional summand. The additional summand however has a very similar structure with a re-sorting term and a contamination term both scaling neighborhood C's autarky effect. Given that neighborhood C is not part of the implied DiD empirical specification which compares neighborhood A and neighborhood B, we refer to these terms as "indirect spillover on the switchers" and "indirect contamination". The former, i.e. the autarky change in C multiplied by the indirect spillover on the switchers scaling, captures the effect on the price in neighborhood A from people moving from A to C due to the subsidy-induced price decrease in the latter. This moderates the price increase in neighborhood A attributable to direct re-sorting. The indirect contamination, i.e. the autarky change in C multiplied by the indirect contamination scaling, captures the effect on the price in neighborhood B as people move from B to C due to the subsidy-induced price decrease in the latter. This increases the contamination in neighborhood B as prices fall even further there.

Nota Bene If neighborhood C were actually unsubsidized one can set $(p_2^C - p_1^C) = 0$, and thus the entire derivation is identical to the situation described in Subsection B.1.

B.3 Multiple Subsidized and Multiple Unsubsidized

Generalizing the results from Subsection B.2 to a setting with many subsidized and unsubsidized areas is straightforward. In Equation B.12 one can see that the effect of one additional subsidized neighborhood on the decomposed DiD estimator formula is one additional summand. On the other hand, as noted above, any additional unsubsidized neighborhood has no effect on the decomposition, as their effect is already captured in the direct spillover on the switchers term. Equation B.13 thus captures the generalization to many subsidized (and unsubsidized) neighborhoods. Please note that we are again only using neighborhoods A and B to decompose the DiD estimator.

$$\begin{aligned}\hat{\beta}_{DiD} \approx \underbrace{(p_2^A - p_1^A)}_{\text{Autarky in } A} \times \Bigg[1 + \underbrace{\frac{\partial D^A}{\partial p^A} \times \frac{\partial P_S^A}{\partial q^A}}_{\text{Direct Spillover-on-switchers Scaling}} - \underbrace{\frac{\partial D^A}{\partial p^A} \times \frac{\partial P_S^B}{\partial q^B} \times DR_{A,B}}_{\text{Direct Contamination Scaling}}\Bigg] \\ + \sum_{k \in \mathcal{K}} \underbrace{(p_2^k - p_1^k)}_{\text{Autarky in } k} \times \Bigg[\underbrace{\frac{\partial P_S^A}{\partial q^A} \times \frac{\partial D^A}{\partial p^k}}_{\text{Indirect Spillover-on-switchers Scaling}} - \underbrace{\frac{\partial P_S^B}{\partial q^B} \times \frac{\partial D^B}{\partial p^k}}_{\text{Indirect Contamination Scaling}} \Bigg]\end{aligned} \tag{B.13}$$

with $\mathcal{K}$ denoting the set of all neighborhoods subsidized by the policy of interest, excluding neighborhood A.

In the main text, we use Equation B.14. Equation B.14 is a simple re-writing of Equation B.13 in order to incorporate diversion ratios. Such reformulation allows for easier comparison with Equation 5.

$$\begin{aligned}\hat{\beta}_{DiD} \approx \underbrace{(p_2^A - p_1^A)}_{\text{Autarky in } A} \times \Bigg[1 + \underbrace{\frac{\partial D^A}{\partial p^A} \times \frac{\partial P_S^A}{\partial q^A}}_{\text{Direct Spillover-on-switchers Scaling}} - \underbrace{\frac{\partial D^A}{\partial p^A} \times \frac{\partial P_S^B}{\partial q^B} \times DR_{A,B}}_{\text{Direct Contamination Scaling}}\Bigg] \\ + \sum_{k \in \mathcal{K}} \underbrace{(p_2^k - p_1^k)}_{\text{Autarky in } k} \times \Bigg[\underbrace{\frac{\partial P_S^A}{\partial q^A} \times \frac{\partial D^k}{\partial p^k} \times DR_{k,A}}_{\text{Indirect Spillover-on-switchers Scaling}} - \underbrace{\frac{\partial P_S^B}{\partial q^B} \times \frac{\partial D^k}{\partial p^k} \times DR_{k,B}}_{\text{Indirect Contamination Scaling}} \Bigg]\end{aligned} \tag{B.14}$$

with $\mathcal{K}$ denoting the set of all neighborhoods subsidized by the policy of interest, excluding neighborhood A.

C Appendix: Neighborhood and Nest Definition

C.1 Spatial Clustering

We employ the spatial clustering algorithm "SKATER" (Spatial 'K'luster Analysis by Tree Edge Removal), developed by Assunção et al. (2006), because it has four convenient features for our context. First, unlike regular, non-spatial clustering techniques, this algorithm guarantees spatial contiguity of the resulting units. Second, it allows for the introduction of a constraint on the minimum number of observations each unit should have. We need this feature to make sure that each neighborhood has enough transactions to compute the average price and market shares we use in the estimation of the demand model. Third, the algorithm operates by maximizing the internal homogeneity of the resulting units in terms of a variable defined by the researcher. Finally, the procedure allows one to set a target number of units. This target has a lower priority in the functioning algorithm and may not be reached to satisfy the other constraints.

We apply the spatial clustering algorithm separately to the subsidized and unsubsidized sections of the city such that the entire area of each neighborhood falls into only one of those two categories. We indicate the algorithm to use the average number of years of education of the tracts from the 2011 population census to maximize the homogeneity of the units. Figure A.3 in Appendix A shows there are huge differences in years of education across Montevideo. This huge variation, together with the evidence about the sorting of households along education, makes this variable an ideal candidate for dividing the city into different units (Black, 1999; Bayer et al., 2007). We set a minimum of 10 transactions for the average number of monthly sales that each neighborhood should have, and a target of 50 subsidized and 50 unsubsidized neighborhoods.

The spatial clustering algorithm gives us 49 neighborhoods, 30 subsidized and 19 unsubsidized. We further classify those 49 neighborhoods into six groups, which are the nests of our nested logit model. For this second classification, we use the same algorithm as in the first one, except we do not require spatial contiguity for the resulting units, thus allowing the algorithm to join subsidized and unsubsidized neighborhoods in the same nest. The results of this operation are presented in the lower half of Figure 3. Each of the six colors in that figure represents a different nest, the solid line represents the border of the policy, and the lighter lines show the borders of the neighborhoods.

C.2 External Validity

The "Continuous Household Survey" (Encuesta Continua de Hogares" (ECH) in Spanish) provides data on intra-urban residential mobility in Montevideo for the years 2009, 2010 and 2011. Household movements are recorded at the barrio level, a statistical unit without substantive function in terms of public service provision. We transform these flows between barrios into flows between our neighborhoods with areal interpo-

Figure C.1: Matching Process

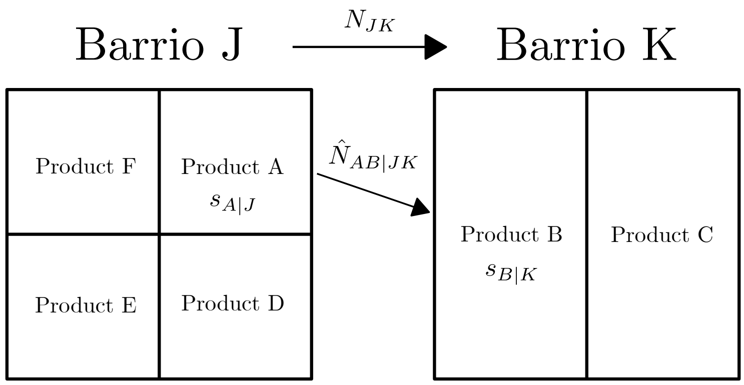

Source: Authors' illustration.
Notes: The graphic illustrates the process of translating the intra-urban residential mobility survey data from movement between barrios to movement between neighborhoods as defined by our structural model.

lation.

Figure C.1 illustrates the areal interpolation approach. The survey records aggregate flows of N_{JK} individuals moving from barrio J to barrio K. To reconcile these barrio-level flows with the neighborhood definitions employed in our structural model, we overlay the geospatial shapefile for barrios with the one for structural neighborhoods. This procedure enables us to partition each barrio into individual components, where each component lies entirely within a single structural neighborhood.

In the example shown, barrio J spans four structural neighborhoods — A, D, E, and F — while barrio K encompasses two — B and C. We denote the share of barrio J's area falling within neighborhood A as $s_{A|J}$, and define similar shares for all components. Assuming that movers are uniformly distributed within both the origin and destination barrios, we allocate flows proportionally. Specifically, the imputed number of individuals moving from neighborhood A to neighborhood B as part of the barrio-level movement from J to K is given below.

$$\hat{N}_{AB|JK} = N_{JK} \cdot s_{A|J} \cdot s_{B|K}$$

Aggregating across all barrio pairs yields an estimate of the total number of movers between any two structural neighborhoods:

$$\hat{N}_{AB} = \sum_{X,Y} \hat{N}_{AB|XY}$$

Once we recover the number of switchers for each structural neighborhood pair, we use this information to test whether individuals are more likely to move between neighborhoods within the same nest than across different nests. To do so, we estimate the following linear regression:

$$\hat{N}_{AB} = \delta_0 + \delta_1 * \mathbb{1}(\text{A \& B in Same Nest}) + o_A + d_B + \beta \cdot km_{AB} + \epsilon_{AB} \qquad \text{(C.1)}$$

Where $\hat{N}_{AB}$ denotes the number of movers from neighborhood A to neighborhood B. The indicator variable $\mathbb{1}$(A & B in Same Nest) equals one if neighborhoods A and B belong to the same nest. The coefficient δ_0 captures the baseline number of switchers between neighborhoods in different nests, while δ_1 measures the additional within-nest mobility. The specification includes origin and destination fixed effects, oA and d_B, as well as a control for the distance between neighborhoods, km_{AB}. Note the ordering of subscripts: the first index indicates the origin neighborhood, and the second the destination.

Table C.1: Flows across and within nests

	Dependent Variable:			
	Number of Switchers going from A to B			
	(1)	(2)	(3)	(4)
δ_0	461.3***	–	–	–
	(25.22)			
$\delta_1 \cdot \mathbb{1}$(A & B in Same Nest)	105.7**	97.53**	107.9**	90.74**
	(50.27)	(46.17)	(44.25)	(38.47)
$\beta \cdot km_{AB}$	−46.70***	−74.16***	−80.36***	−124.3***
	(5.372)	(9.670)	(9.417)	(14.78)
Fixed Effect - Geography		Origin Neighborhood	Destination Neighborhood	Origin Neighborhood + Destination Neighborhood
No. Obs	2,401	2,401	2,401	2,401
Data	City-Wide	City-Wide	City-Wide	City-Wide
Mean Number of Products In Same Nest	9.1	9.1	9.1	9.1
Mean Number of Products Outside Same Nest	39.9	39.9	39.9	39.9

Source: Authors' calculations using intra-urban residential mobility survey data from the National Institute of Statistics (INE) of Uruguay.
Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors are clustered at the (origin neighbourhood, destination neighbourhood)-level and provided in parentheses.

Estimates in Column 1 of Table C.1 indicate that residential flows between adjacent neighborhoods (i.e., those with $km_{AB} = 0$) within the same nest are 22.9 percent higher than flows between adjacent neighborhoods belonging to different nests. Moreover, being part of the same nest effectively offsets a spatial separation of up to 2.26 kilometers, meaning that same-nest neighborhoods can be further apart and still exhibit comparable flows to closer neighborhoods in different nests.

The positive, and statistically significant, coefficient for the within-nest indicator variable persists even after controlling for origin and destination fixed effects. These fixed effects absorb unobserved characteristics of the respective neighborhoods, such as differences in population size or inherent desirability. Taken together, these findings provide external validation for the nesting structure of neighborhoods generated by the algorithmic approach described above, suggesting that this classification meaningfully captures patterns of intra-urban residential mobility in Montevideo.

D Appendix: Trade-off between Parallel Trends and Contamination

We simulate alternative cities with different fundamentals (amenities and marginal costs) by introducing random variation in three types of shocks: a) the time invariant shocks that represent the "base heterogeneity" across locations (terms depending on j), b) the "time heterogeneity", which are time shocks that affect all locations at the same time (terms depending on t), c) the "idiosyncratic heterogeneity" shocks that vary by time and locations (terms depending on jt).

For the case of amenities (AM_{jt}) each of those three shocks is captured by a specific random variable, γ_j, γ_t and τ_{jt}, and we define $AM_{jt} = \gamma_j + \gamma_t + \tau_{jt}$. Analogously, for marginal costs (L_{jt}) we have $L_{jt} = L_j + L_t + \epsilon_{jt}$. Table D.1 presents the assumed distributions for the six random variables.

Table D.1: Simulation Setup - Random Variable Distributions

Variable	Parameters	
Base Heterogeneity	$\gamma_j \sim N(0, \sigma_j)$	$L_j \sim \log N(0, \sigma_j)$
Time Heterogeneity	$\gamma_t \sim N(0, \sigma_t)$	$L_t \sim \log N(0, \sigma_t)$
Idiosyncratic Heterogeneity	$\tau_{jt} \sim N(0, \sigma_{jt})$	$\epsilon_{jt} \sim \log N(0, \sigma_{jt})$

We extract three main takeaways from the simulation exercise. First, our model allows for parallel trends. We simulate the model for a specific set of parameters ($\sigma_j = 0.5, \sigma_t = 0.3, \sigma_{jt} = 0.2$) to show that, despite being very non-linear in both the demand and the supply side, our model can produce parallel trends between subsidized and unsubsidized areas. The top graph in Figure D.1 suggests the presence of parallel trends in a typical DiD graph, while the bottom graph in Figure D.1 presents the typical event study test for parallel trends in the literature.

Figure D.1: Simulations for A Specific Set of Parameters ($\sigma_j = 0.5, \sigma_t = 0.3, \sigma_{jt} = 0.2$) and Nesting Coefficient of 0.5

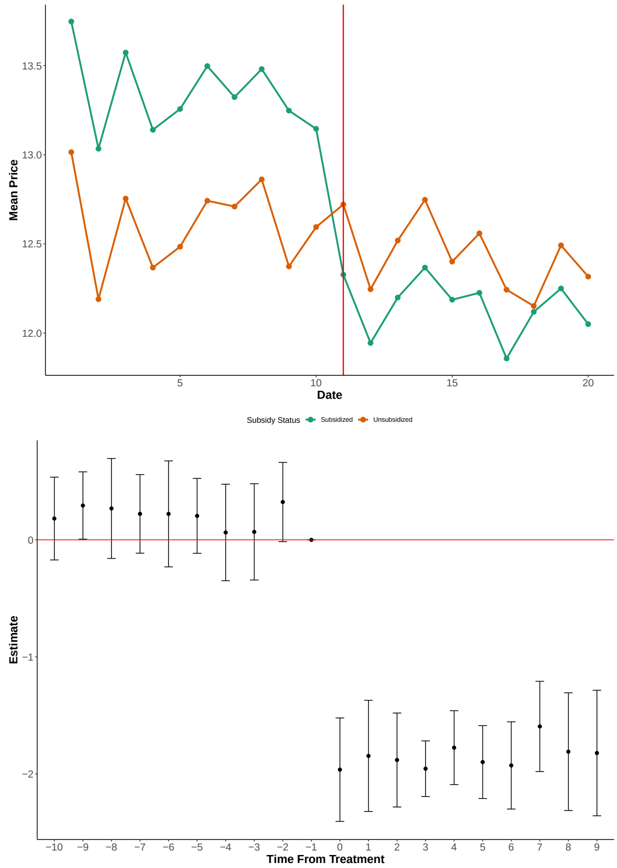

The second takeaway is to characterize under which type and size of heterogeneity our model rejects the parallel trends. To analyze this issue we perform simulations over several values of the heterogeneity parameters. In these simulations, the variance for j terms (σ_j) is limited to the set $\{0.5, 1..0\}$, while the other two variances (σ_t and σ_{jt}) can vary along a grid from 0 to 1.5 (in 0.5 increments).

Figure D.2 presents the results for $\sigma_j = 1$ and Figure D.3 shows the results for $\sigma_j = 0.5$. For each of the three levels of the nested logit nesting parameter (i.e. the plain σ in our model), the upper panel shows the number of significant coefficients in a regression of equilibrium prices on a set of interactions between time period and subsidy status and including neighborhood and time fixed effects. In all of the upper panels, the number of parallel trend violations is relatively small. They tend to occur when the variation in the jt dimension is large compared to the variation in t or, vice versa, when variation in t is large compared to variation in jt.

Finally, we highlight a trade-off between parallel trends violations and the contamination effect. The bottom graphs of the figures present the size of the contamination effect as $\%$ of the ATT in these simulations. In line with our theoretical predictions, contamination is higher when the substitutability of same-nest products is higher (as measured by higher nesting coefficients). In the lower panels, contamination is less with the lower values of the nest coefficient, but that comes at the cost of more violations in parallel trends in the upper panels.

Figure D.2: Parallel Trends and Contamination Effects in Simulations for $\sigma_j = 1$

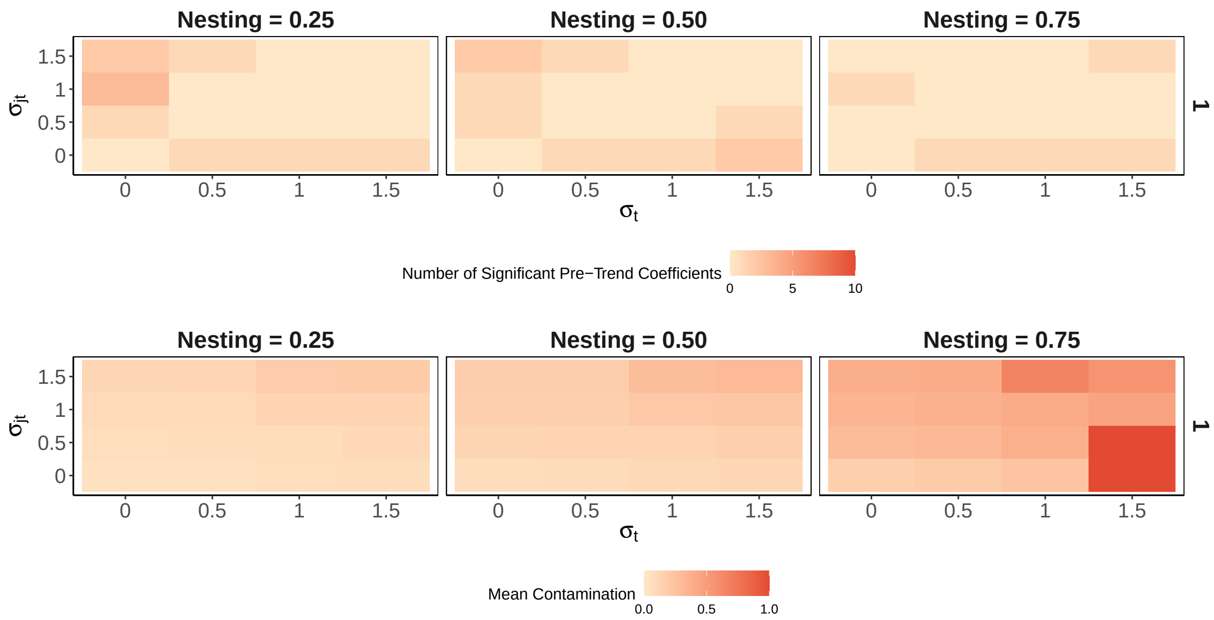

Figure D.3: Parallel Trends and Contamination Effects in Simulations for $\sigma_j = 0.5$

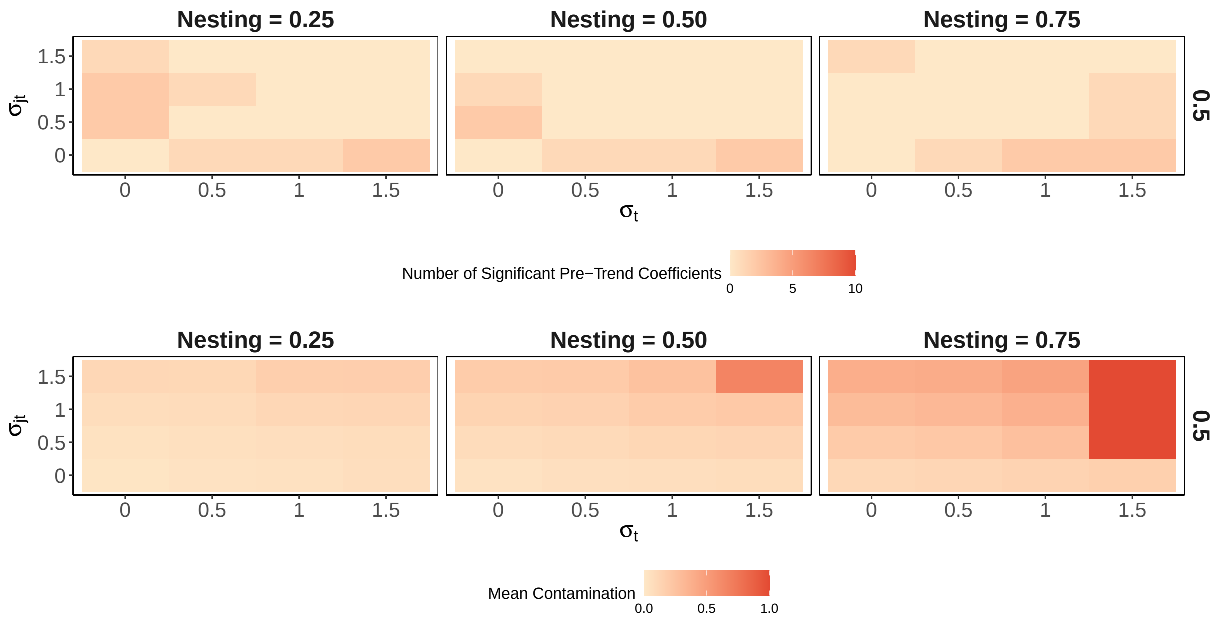

E Appendix: Construction of the Price of a Generic Housing Unit

In order to make housing comparable across space and time, we create a representative property, which we refer to as "generic housing unit" (GHU). This allows us to keep property characteristics constant and attribute any changes in price to an improvement in neighborhood amenities.

We arrive at the price P_{jt} we employ in our demand and supply estimation exercises in the paper by first estimating the following regression equation, which shares notation and many terms with Equation 7:

$$p_{ijt} = \gamma_{jt} + B * f(X_{ijt}) + \nu_{ijt} \tag{D.1}$$

with p_{ijt} denoting the price per square meter of housing transaction i in neighborhood j at month t. γ_{jt} is a vector of neighborhood and month-year fixed effects, and $f(X_{ijt})$ is a third-order polynomial evaluated on the set of housing characteristics X. Those characteristics are: transaction area in square meters, building age in years, distance to the coast in meters, and indexes of construction category, construction condition, type of ceiling, and whether there is ongoing construction work on the property.

After estimating Equation D.1, we build the price P_{jt} of the GHU in neighborhood j at time t as:

$$P_{jt} = \hat{\gamma}_{jt} + \hat{B} * f(\bar{X}_{ijt}) \tag{D.2}$$

where $\bar{X}_{ijt}$ is the vector with the average values of the covariates for the whole sample, and $\hat{\gamma}_{jt}$ and $\hat{B}$ are the estimated parameters from Equation D.1.