Difference-in-Differences in Equilibrium: Evidence from Place-Based Policies*

Guillermo Alves[†] William H. Burton[‡] Sebastian Fleitas[§]

September 17, 2024

Abstract

Violations of the stable unit treatment value assumption (SUTVA) pose a common challenge to identifying equilibrium effects of policies. We show that the difference-in-differences estimator can be decomposed into autarky, re-sorting, and contamination effects, and we provide a formula for calculating contamination using supply and demand partial derivatives. To illustrate, we estimate a structural demand and supply model to analyze a housing tax break in Uruguay. Our findings show that SUTVA violations account for 25% of the effect on subsidized areas, significantly underestimating the policy's incidence.

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, Jan De Loecker, Xavier D'Haultfoeuille, Jonathan Dingel, Tomás Domínguez-Iino, Sebastian Gallegos, Geert Goeyvaerts, Taylor Jaworski and Frank Verboven for comments and suggestions. We thank 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.
[‡]KU Leuven and FWO.
[§]PUC-Chile and CEPR.

1 Introduction

Non-experimental studies of policies causing the re-sorting of agents between treatment and control groups may suffer from violations of the stable unit treatment value assumption (SUTVA) (Donaldson, 2015). These policies, e.g. place-based policies, are usually not randomized and researchers rely on non-experimental methods - such as difference-in-differences (DiD) - to study their effects (Kline & Moretti, 2014b; Baum-Snow & Ferreira, 2015). Identifying the causal effects of these policies using such methodologies requires (among other assumptions) the non-violation of the SUTVA. While most of the recent developments in DiD methodology have focused on the parallel trends assumption as well as staggered treatment, less attention has been paid to SUTVA violations (Roth et al., 2023).

In this paper, we discuss the DiD estimator under SUTVA violations caused by equilibrium effects, i.e. the re-sorting of agents between control and treatment units. We proceed with three approaches. First, we estimate the effect of the policy intervention using the DiD estimator, and explore heterogeneity potentially correlated with the presence of equilibrium effects. Second, we decompose the DiD estimator and derive a sufficient statistics general formula to assess the degree of bias caused by those effects. Third, we quantify this bias by developing and estimating a structural model in the context of a place-based policy in Uruguay.

We aim to offer a bridge between these three different approaches by proposing guidelines for applied researchers on how to proceed under potential SUTVA violations. We show the value of studying heterogeneity in DiD estimates to evaluate whether the results are consistent with the presence of equilibrium effects, which create contamination in the estimates. Under no signs of contamination - e.g. by finding similar effects for units that are potentially more prone to contamination - the DiD can be a sensible approach to estimate the effects. However, if there are signs of contamination but the policy intervention (e.g. the subsidy) is a marginal change, then the sufficient statistics approach developed in this paper is a good framework to measure the size of contamination and guide the empirical strategy. Finally, if there are strong signs of contamination and the intervention is large, then a structural model is needed to consistently estimate the effect of the policy and assess the extent of contamination bias.

We begin by highlighting the nature of contamination, showing that in the presence of re-sorting, the DiD estimator can be decomposed into three effects. First, an "autarky effect" which captures what would happen to the treated area if it were isolated and therefore no relocation effects existed. Second, a "re-sorting effect" captures the effect on the treated area caused by the inflow of agents into this area. Third, a "control area contamination effect" captures the effect on the control areas caused by the outflow of agents from these areas. This control area contamination term is what prevents the DiD estimator from accurately measuring the average treatment effect on the treated

(ATT).

By linearizing a model of supply and demand for housing across neighborhoods in a city, we provide an analytical formula that approximates the DiD estimate of the introduction of a housing supply subsidy in some neighborhoods of that city. First, the formula shows that the DiD estimator is asymptotically biased. Second, the formula allows the researcher to know the sign of the bias under mild economic assumptions. Finally, the formula highlights that the relative size of each of the three effects contained in the DiD estimator depends on the demand-side substitution patterns between neighborhoods as well as the supply elasticities of the neighborhoods.

Our formula helps researchers both prior to as well as after conducting an empirical study, as supply and demand partial derivatives are "sufficient statistics" for the relative size of re-sorting and contamination (Saez, 2001; Chetty, 2009). Before conducting the study, if the researcher has an estimate of these demand and supply partials, the formula allows to compute an approximation of the relative size of the contamination effect (i.e. the bias of the estimator). After conducting the empirical study and obtaining a DiD estimate, the researcher can use those supply and demand partials to recover the actual magnitudes of all three effects.

We further generalize our formula to (the often typical) case in which more than one area is treated at the same time. This generalized formula includes indirect re-sorting and contamination effects created by other subsidized areas on the original areas of reference. We show that, in cases with more than one subsidized area, the DiD estimate suffers to a larger extent from the bias caused by the contamination effect.

Additionally, even if the researcher does not have estimates of demand and supply partials, the formula derived in this paper can offer some guidelines for applied work. More similar areas are likely to be closer demand-side substitutes and therefore be subject to the highest contamination effects. This contradicts the intuition behind choosing very similar units to define treatment and control groups in difference-in-differences designs, such as comparing areas across policy borders or employing matching techniques (Neumark & Kolko, 2010; Chen et al., 2022). By using simulations of a supply and demand model, we show that there is a trade-off between the parallel trends assumption and the SUTVA. When the areas are very similar in characteristics, the parallel trends assumption is satisfied more often but simultaneously the contamination effect is larger.

The decomposition formula also allows us to analyze the assumptions made in the implementation of the DiD approach in the previous literature on place-based policies. First, when the relocation of agents is very local, it can be reasonable to assume that distant areas are not affected by re-sorting. In these cases, identification can be achieved by comparing the treated area with distant ones (Delgado & Florax, 2015; Clarke, 2017; Butts, 2023a). A prominent example of this approach is Kline and Moretti (2014a), who exclude neighboring counties from their control group in their evaluation of the impact of the Tennessee Valley Authority (TVA).

Second, in many economic settings, the re-sorting of agents from untreated into treated areas implies that truly untreated areas may not exist or may be hard to credibly detect. In those contexts, researchers may still recover the impact of the policy under the assumption that all areas are small enough such that the mobility of agents does not affect prices and quantities in untreated areas. Busso et al. (2013)'s study of Empowerment Zones constitutes an example of this second type of situation in which DiD estimates can recover the effect of the policy.

We apply these insights to the study of a place-based policy that provides substantial tax breaks for housing development in lagging areas of Montevideo, the capital of Uruguay. We start the analysis by using administrative data on the universe of housing transactions in Montevideo before and after the policy to estimate a series of DiD regressions with housing prices as our dependent variable. The variation in these estimates is consistent with our conceptual framework of SUTVA violations and contamination bias. When using all housing transactions in the city, we find a large negative effect of the policy of around 18% of the average transaction price. However, when we follow the common practice of only using observations close to the border, where re-sorting is likely to be more pronounced, estimates are very small negatives or zeros. This is consistent with contamination caused by re-sorting having an attenuating role: prices fall on the unsubsidized side of the border as agents re-sort into the subsidized side. Also consistent with contamination effects, the absolute magnitude of these border estimates increases with a measure of heterogeneity across both sides of the border and also when we use control units located further away from the border.

We quantify the contamination using transaction data and a structural model of housing supply and demand across Montevideo's neighborhoods, in combination with model-free estimates. We model the demand for housing as a discrete choice problem of choosing a neighborhood within a city (Bayer et al., 2007; Almagro & Domınguez-Iino, 2019; Anagol et al., 2021).[1] We estimate the price elasticity of housing demand using a nested logit demand model and a set of supply-shifting instruments defined by the introduction of the tax break. The housing supply in the model is characterized by each neighborhood having a separate log-linear supply function (Saiz, 2010; Baum-Snow & Han, 2023). We use the model-free DiD estimate to calibrate our model. Specifically, we calibrate a common inverse supply elasticity for all neighborhoods by matching the reduced-form DiD estimate with its structural equivalent. We show that our model fits the data well in terms of reproducing the parallel trends that we observe.

By solving for a series of counterfactual equilibria, we compute the three additive effects of our decomposition formula for the DiD estimator. We find a"re-sorting effect" of 40% of the "autarky effect" and a "contamination effect" that represents 25%

[1]The application of discrete choice techniques to spatial settings was pioneered by Bayer et al. (2007) and has been applied to a variety of contexts, both within cities (Bayer et al., 2016; Almagro & Domınguez-Iino, 2019; Anagol et al., 2021) and across cities (Diamond, 2016; Alves, 2021).

of the ATT. The existence of substantial contamination implies that the reduced-form DiD approach underestimates the share of the subsidy that reaches consumers (i.e. the incidence of the policy) by 20 percentage points. This underestimation caused by contamination amounts to approximately 24% of Uruguay's GDP per capita in the year the policy was introduced.

Finally, we use the counterfactual equilibria to revisit the relationship we find in the reduced-form analysis between heterogeneity across control and treated units and the size of the DiD estimate. Consistent with our decomposition formula, we confirm that contamination is negatively correlated with our measure of heterogeneity between units, and positively correlated with the diversion ratios between these units. Importantly, this implies that the lower absolute values of the reduced-form DiD estimates obtained by comparing homogeneous units are effectively driven by contamination (i.e. a larger bias) and not just regular treatment heterogeneity. This calls for caution with the usual approach of maximizing the comparability between treatment and control groups.

Our paper contributes to three 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.[2]. Baum-Snow and Ferreira (2015) include DiD as one of the main identification methods in these settings and highlight the re-sorting of individuals between treatment and control areas as a serious threat to identification. This threat can be seen as a special case of dealing with spatial spillovers in DiD settings, a topic that has received attention from several previous works (Clarke, 2017; James & Smith, 2020; Banzhaf, 2021; Huber & Steinmayr, 2021; Myers & Lanahan, 2022; Butts, 2023a; Ding et al., 2023; Hollingsworth et al., 2024; Jardim et al., 2024). [3]

A more general literature discusses 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 three effects: on the targeted group without considering spillovers, spillovers on the targeted group, and spillovers on the non-targeted group. This decomposition into three terms was applied to DiD by Butts (2023a). We implement a similar decomposition and innovate on two

[2]Currently, successful identification using DiD in the presence of spatial spillovers is restricted to two contexts. First, the "aggregation approach" defines large enough units such that spillovers are contained within those units (Feyrer et al., 2017; Huber & Steinmayr, 2021). Second, the "ring approach" assumes spillovers fade away far enough from the treated area (Clarke, 2017; Butts, 2023a). However, natural (sea, mountains) or man-made (parks, highways) barriers may restrict how far one can go and, with large policies, spillover-free areas may not exist or may be difficult to identify.

[3]Other papers deal with SUTVA violations by deriving DiD equations using spatial quantitative models (Rudik et al., 2022; Hollingsworth et al., 2024). In this approach, the source of the spillover is explicitly modeled and included in the estimating equation, thus producing consistent estimates. This approach has been applied more generally to spillovers that are not necessarily spatial, as in Rotemberg (2019)'s analysis of the effect of subsidies to firms in India. Bachmann et al. (2023) also discuss the role of demand spillovers in a non-spatial DiD setting in their study of the 2015 Volkswagen emissions scandal.

fronts. First, we characterize the three terms using supply and demand partials. Second, we calculate the three terms with a structural model in a context with no spillover-free areas.

Second, we contribute to the literature on place-based policies targeting lagging areas by analyzing the contamination issue when estimating the effects. As highlighted by Kline and Moretti (2014b), evaluating the success of these programs requires going beyond their impact on specific variables and adopting a consistent equilibrium framework. One key lesson from spatial equilibrium models is that the efficiency impact of place-based policies depends on the degree by which the policy induces economic agents to relocate from untreated into treated areas (Moretti, 2011; Busso et al., 2013; Serrato & Zidar, 2016). We show that heterogeneous re-sorting can generate wrong conclusions about the efficiency of place-based policies when estimates are obtained by comparing only certain areas.

Third, our analysis of SUTVA violations in the DiD estimate contributes to the burgeoning literature on the methodological improvement of these estimates (de Chaisemartin & D'Haultfœuille, 2023; Roth et al., 2023). There has been substantial progress in designs with multiple periods and variation in treatment timing (Callaway & Sant'Anna, 2021; Goodman-Bacon, 2021; Sun & Abraham, 2021), potential violations in parallel trends (Rambachan & Roth, 2023; Roth & Sant'Anna, 2023), and improved inference (Ferman & Pinto, 2019). Roth et al. (2023) include spillovers as one of the main areas for future research in the DiD literature, with a special mention to spatial spillovers.

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 which never receives treatment, and one which 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 object of interest, namely 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 reaction to the introduction of a supply-side subsidy in the housing market of a city. Without loss of generality and under homogeneous effects, we assume 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:

$$\hat{\beta}_{DiD} = (p^A_{post} - p^A_{pre}) - (p^B_{post} - p^B_{pre}) \tag{3}$$

We now discuss the DiD formula in Equation 3 through the lens of a supply and demand model for homogeneous housing units across neighborhoods within a city. In the model, households on the demand side choose whether to buy a housing unit in one of two city neighborhoods or outside the city. There are two main determinants of households' discrete choice between neighborhoods: 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 neighborhood 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. This means that their actions are independent of both previous and future periods.

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.[4]

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

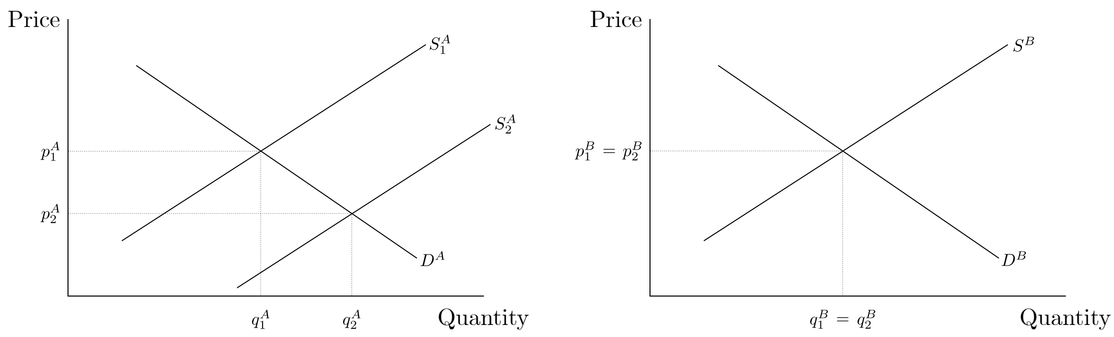

Figure 1 presents the autarky situation in which consumers do not reallocate between A and B, but may relocate between their respective neighborhood and the outside option. Implementing a supply-side subsidy in A would first result in an outward shift of the supply in this neighborhood.[5] Due to lower 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.[6] The estimated DiD in this scenario equals the difference in prices between periods 2 and 1 in the neighborhood A:

$$\hat{\beta}_{DiD}^{AUT} = (p_2^A - p_1^A) - (p_2^B - p_1^B) = p_2^A - p_1^A$$

Note that 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 highlights a situation in which consumers may relocate between the two neighborhoods. When the supply side subsidy is enacted in neighborhood A, housing prices drop from p_1^A to p_2^A. However, due to the assumed pattern of substitution, 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.[7] Both movements are due to re-sorting. Re-sorting increases prices in A

[4]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 reaction to a demand-side place-based policy could be analogously accommodated in the framework. As discussed in Section 3, supply-side re-sorting is not relevant in our empirical setting.

[5]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.

[6]We abstract away from other changes happening over time. Note that this is analog to assume "parallel trends" as in Roth et al. (2023).

[7]The new demand curve in A passes through the original (q_A^1, p_A^1) 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_A^1, capturing the re-sorting of agents away from B and into A in reaction to those lower prices.

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

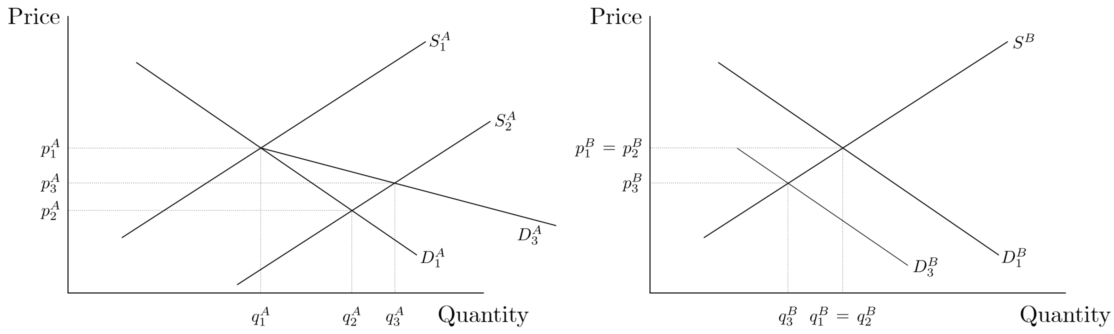

from p_2^A to p_3^A while reduces those in B from $p_1^B = p_2^B$ to p_3^B. Estimating the effect of the policy with DiD now yields the following:

$$
\begin{aligned}
\hat{\beta}_{DiD} &= (p_3^A - p_1^A) - (p_3^B - p_1^B) \\
&= (p_3^A - p_2^A + p_2^A - p_1^A) - (p_3^B - p_2^B + p_2^B - p_1^B) \\
&= (p_2^A - p_1^A) + (p_3^A - p_2^A) - (p_3^B - p_2^B)
\end{aligned}
$$

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.[8] As indicated in Equation 4, we refer to the additional effect in A as "re-sorting", and to the effect in B as "contamination". While in our context both re-sorting and contamination attenuate the autarky effect of the policy, the former is part of the "legitimate" effect of the policy on the targeted neighborhood while the latter "contaminates" the DiD estimate.

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

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).[9] 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

[8] 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).

[9] For example, Dix-Carneiro and Kovak (2017) state that their "methodology captures only relative effects across regions, as does the rest of the literature examining the regional or sectoral effects of trade".

on the targeted area. On the other hand, the contamination effect can be of interest by itself, as it shows the effect of the policy on non-targeted areas.

Going back to our example 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 while it actually reduced prices in both areas.

Finally, our market equilibrium setting makes clear that the issue of contamination is more general and does not only apply to the specific structure of the DiD. That is, alternative approaches that do not use the parallel trends assumption for identification would also suffer from contamination. Our framework highlights that the issue of contamination comes 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.

Next, we introduce an approximation formula for the DiD estimator that helps to understand the determinants of the relative sizes of both re-sorting and contamination. The relative size of contamination in that formula defines the relative size of the asymptotic bias of the DiD estimate.

2.3 DiD Decomposition with Supply and Demand Elasticities

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 quotient 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:

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

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

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

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 re-sorting. It captures the effect of people 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 households' moving decisions between subsidized and unsubsidized neighborhoods are very sensitive to prices 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 computed with the following formula derived in Appendix B:

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

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 indirect re-sorting, i.e. people reallocating 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 direct re-sorting. 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.

[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.

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 and contamination. First, Equations 4 to 6 show that the contamination effect biases the DiD estimate, and also provide the direction of that bias. For example, if the policy reduces prices in the subsidized areas but, through contamination, also reduces prices in the unsubsidized areas, then the DiD estimate (β_{DiD}) would be biased towards zero. Alternatively, the DiD estimate of a subsidy that increases jobs in one area by displacing jobs from another one would be upward biased.

Second, our formulas highlight the determinants of the bias. On one hand, contamination increases with the inverse elasticity of the housing supply of the unsubsidized area. This deserves special attention given the available evidence on neighborhood-level housing supplies being rather inelastic (Baum-Snow & Han, 2023). On the other hand, contamination increases with the diversion ratio between the two areas. Knowing these determinants can help the applied researcher choose better control areas in contexts of re-sorting.[13]

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. In other words, the researcher would be able to understand the relative importance of contamination even before starting the study. Additionally, after doing the study, the formula in Equation 5 allows researchers to combine their DiD estimate with supply and demand elasticities to calculate all the three terms in Equation 4. Critically, this allows one to recover the ATT without estimating a structural model.

Fourth, Equation 6 shows that with more than one subsidized area the ATT cannot be computed anymore knowing only the supply and demand elasticities. In this sense, the elasticities are not "sufficient statistics" anymore for the relative size of the three effects. To compute these effects, the researcher would need to know the effect of the policy in autarky in all the subsidized neighborhoods.

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 far away area unaffected by the policy and thus can be used as a "contamination-free control". This strategy is often referred to as the "ring approach" and Kline and Moretti (2014a), Clarke (2017), and Butts (2022)

[13] Even if researchers do not have actual estimates of the relevant elasticities, the literature provides rich proxies. 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 between the areas could help.

are examples of relevant papers implementing it. 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 individual 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 tax breaks we analyze were only directed at residential development. We refer to the policy by its familiar acronym in Spanish of "LVIS" (*Ley de Vivienda de Interés Social*). Although the name of the policy refers to the promotion of social housing, homes that benefited from the program did not have to be occupied by low-income households and could be freely sold at market prices.

Tax breaks in LVIS are quite large. González-Pampillón (2022) estimates that the total tax benefits equaled 20% of the cost of the projects. The main component of this policy was the complete exemption from the 22% value-added tax on inputs. LVIS projects were also fully exempted from the country's corporate tax of 25%, and units devoted to the rental market were partially exempted from both income and wealth taxes. Because these tax breaks were so extensive, we expect a negative effect of the policy on the price of housing in subsidized areas.

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, we end our analysis period in May 2014.

We study the impact of LVIS tax breaks in the department of Montevideo, which

[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).

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). Around half of the border coincides with one of the main avenues of the city, which has been historically the most important spatial division between high- and low-income neighborhoods in the city. The other half of the border is drawn across 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 its tax breaks implied that the policy had huge impacts 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 the first three years after. 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 implied that developers had to apply for tax benefits and obtain approval for their projects before beginning the construction phase. As a result of this application process plus the usual construction phase, the first few LVIS projects only reached completion in 2013. Subsequently, the first sales of LVIS properties occurred in 2014, while most of the sales were made in the following years (González-Pampillón, 2022). This timing implies that almost no LVIS projects and very few LVIS sales were completed during the period we study. Thus, our hypothesis on the negative effect of the policy on prices fully operates through the capitalization of future lower construction costs into current housing prices. Although there were almost no finished subsidized units during our period of analysis, the number of applications and approved projects was large. With these numbers being publicly available online and many projects being in their construction phase, it was common knowledge during our period of analysis that the supply in the targeted neighborhoods was going to expand substantially.

The 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 four sources of data. The 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 further includes a unique property number, which allows us to match each sale with its corresponding entry in the registry of the National Cadaster of Uruguay, our second main source of data. 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 percentile of the area and price distribution of the transaction dataset to avoid our estimates being affected by extreme values.

The third source of data 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. The fourth and last source of data is the 2010 national population census. These data provide census tracts' average years of education, which we use to define neighborhoods and nests, as we explain later in this section.

Table 1 presents the average number of transactions, their price per square meter, and their transacted area separately for the subsidized and unsubsidized sections of the city and distinguishes furthermore between before and after the introduction of the policy. Consistent with the policy subsidizing lagging areas, the prices of houses 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.

Throughout the paper, we use a set of housing characteristics as controls in various 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

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 starts, 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.

computed 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.

Our modeling of the equilibrium impacts of the tax break on housing prices follows a long tradition of using discrete choice techniques to study housing markets (Bayer et al., 2007; Diamond, 2016; Anagol et al., 2021; Almagro et al., 2022). These techniques require a division of the city's territory into exclusive units. Because Montevideo is not divided into administrative units that generate meaningful differences in taxation or public services provision, we partition the city ourselves into contiguous and homogeneous units using a spatial clustering algorithm (Martins et al., 2006). Throughout the paper, we refer to the resulting 49 units as neighborhoods. We further use the same algorithm to classify neighborhoods into homogeneous, not necessarily contiguous clusters which also represent the nests in our nested logit demand estimation. We explain the details of these two stages of spatial clustering in Appendix C. Panels a) and b) of Figure 3 provide a map of the neighborhoods and nests. In panel a), neighborhoods are classified according to their subsidized (30 units) or unsubsidized (19 units) status. Panel b) shows the classification of neighborhoods into nests.

Figure 3: Neighborhood Classification by Subsidy Status and Nest

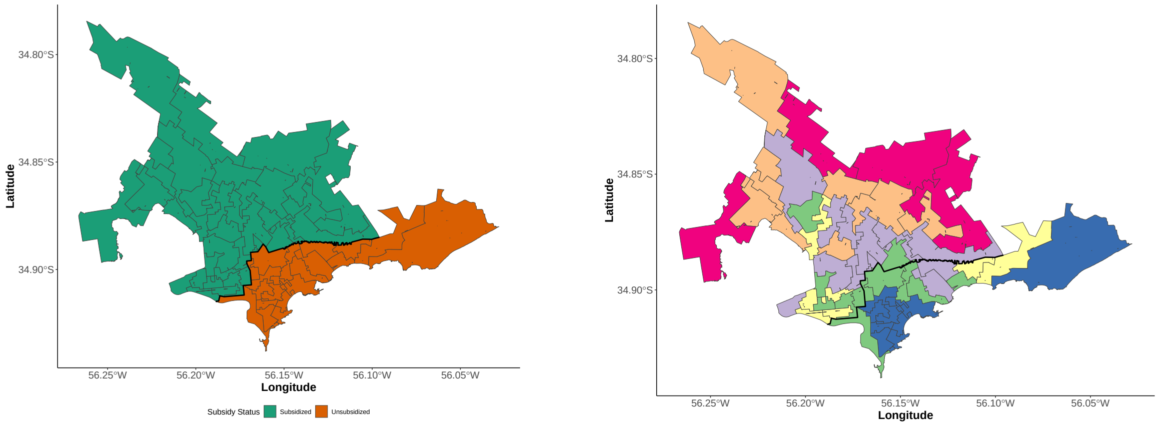

(a) Neighborhoods by Subsidy Status

(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

This section presents three sets of DiD estimates of the effect of the policy, with a specific subsection discussing each set. Overall, the results are consistent with both the subsidy having a negative effect on prices in the targeted areas and the presence of a bias created by contamination due to re-sorting. As discussed in Section 2, the presence of contamination introduces an attenuation bias in the DiD estimates. Our estimates are consistently negative, but their magnitude varies greatly depending on which neighborhoods are included in the subsidized and unsubsidized groups. While some estimates imply large price reductions suggesting a highly beneficial impact of the tax break on consumers, other estimates fail to reject a zero impact hypothesis, which would be consistent with landlords fully appropriating the subsidy. With the DiD approach, we cannot separately identify if these results are caused by a contamination bias, or by heterogeneity of the treatment effects. We will return to this point in Section 7 while discussing the counterfactuals using the structural model.

4.1 Benchmark Difference-in-Differences

The general specification for our DiD regressions is the canonical one, and it is given by the following equation:

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

with p_{ijt} denoting the price per square meter of housing transaction i in neighborhood j at month t. Because each neighborhood is completely subsidized or unsubsidized, the neighborhood fixed-effect γ_j subsumes the usual $Subsidy_j$ term. $f(X_{ijt})$ is a third-order polynomial on the set of housing characteristics mentioned in the previous section. These include the built area, the distance to the coast, the construction year, and four variables that measure the quality of construction.

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 indexes 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.

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 transactions in the city. Column 1 only has the three traditional DiD terms, namely those that indicate the subsidy status, the timing of the subsidy, and the interaction of these two. The second column adds the third-order polynomial on the set of housing characteristics. Column 3 adds month-year and neighborhood fixed effects. Figure A.4 and Figure A.6 in Appendix A provide the usual graphical evidence complementing these DiD estimates.

The first three columns in Table 2 show consistently negative estimates with a stable magnitude across the different specifications. 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 Difference-in-Differences 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). For instance, 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". However, as suggested by our analysis in Section 2, agents may re-sort more easily between similar areas, leading to larger contamination effects and more biased estimates. In our context, those agents would leave unsubsidized areas, depressing housing prices there, and causing the resulting DiD estimate to be biased towards zero. In fact, all the estimates in this subsection are significantly smaller in absolute value in comparison to the ones in the previous subsection. This pattern aligns with the notion that techniques emphasizing comparability may introduce greater bias due to contamination.

The first and most common technique to maximize comparability between "treated" 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). The estimates in Columns 4 to 6 of Table 2 follow this approach and compare the evolution of prices between subsidized and unsubsidized areas within 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 conducting 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 techniques that are usually implemented to 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 a difference-in-differences (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). Similarly to the border estimates discussed in the previous paragraph, all three point estimates in Columns 1 to 3 of Table 3 are much smaller in absolute value than the benchmark obtained for the whole city. Again, this is consistent with larger contamination caused by re-sorting between more homogeneous units.

Finally, we also estimate the effect using the difference-in-differences with the popular "ring approach" to construct the control group (Di Tella & Schargrodsky, 2004; Kline & Moretti, 2014a; Butts, 2022; Myers & Lanahan, 2022). This methodology consists of using "control areas" which are further away from the areas targeted by the policy. If the degree of heterogeneity between subsidized and unsubsidized grows with distance from the border, sorting and thus contamination should decrease, and according to our formula the DiD estimate should increase in absolute value. Columns 4, 5 and 6 in Table 3 present DiD estimates for 500-1000, 1000-1500 and 1500-2000 meter rings, respectively. These estimates grow in absolute value with the distance from the border, thus confirming the hypothesized pattern.

This ring approach can identify the true effect of the policy on subsidized areas as long as the spillovers (i.e. changes in prices due to re-sorting) are zero after a certain distance from the border (Clarke, 2017; Butts, 2022; Myers & Lanahan, 2022). There is evidence that these distances can be quite large in some contexts. Clarke (2017) finds that the spillovers of text messaging bans extend at least 30km from the policy border. Myers and Lanahan (2022) establish the range of no spillovers as being beyond the 40th or 60th percentile of their distribution of technological distance across inventors. This requirement of no spillovers after a certain distance may thus 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 the distance after which one can define the control group. In coastal cities, such as ours, water bodies restrict the ring distances one can build. For our case, this is shown in Figure A.8 in Appendix A. For instance, only 10% of our unsubsidized transactions are beyond 2,100 meters from the border. Second, as noted by Butts (2022), when policies are large enough to induce the re-sorting 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 between both sides of the border.

Figure A.9 in Appendix A illustrates how we compute the index that measures price differences across the border. We start by defining a large number of equidistant points along the border. Then draw a 500-meter circle around each of those points and compute the median price per square meter for each side of the border with the transactions falling within that circle (left panel of Figure A.9). Taking the difference between these two median prices within each circle yields a scalar value that measures the heterogeneity in housing prices across the border around that point. The final step consists of attaching, to each housing transaction, a weighted average of those scalars. In this final step we compute, for each transaction, a weighted average of the scalars along the border for which the sold property lies within the respective 500-meter circle around the point. The weights are the inverse of the distance between the transaction and the relevant border points. We standardize the resulting index by subtracting its average and dividing it by its standard deviation. We generically refer to this index as the "heterogeneity index" throughout the paper.

The second column of Table A.1 in Appendix A presents the estimate of the interaction between the DiD term and the heterogeneity index. Each additional standard deviation in the heterogeneity of the border increases the absolute value of the DiD estimate of the effect of the tax break by 55 USD per square meter. This is a huge magnitude given our benchmark DiD estimate of 181 USD for the whole city and a pre-policy average price of 1,112 USD in the 500-meter buffer. Figure A.10 in Appendix A plots the relationship between the DiD estimate and the border heterogeneity index implied by that estimate. Note that the 95% confidence interval for the DiD estimate in that figure includes the zero for all the values in the bottom half of the distribution of the heterogeneity index. This is actually the case of the example given in Figure 2 and the discussion in Subsection 2.2. The reason is that when the areas across the border are very similar, the DiD estimate can be zero since the policy does not affect the relative price of the two areas.

We complement these findings by solving for a specific estimated model that allows us to separately measure the contamination effect and the heterogeneous treatment effects. Using that model, we can test if 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 in the benchmark DiD estimate for this policy.

5 A Model for Quantifying Contamination

In this section we introduce a specific model of real-estate transactions across neighborhoods in a single city. The model is static and housing is assumed to be homogeneous. The demand side of the model consists of a discrete-choice framework with households choosing the neighborhood in which they want to buy a generic housing unit (GHU). Conditional on price, utility might differ across neighborhoods and time because of local exogenous amenities. The supply side of the model consists of property owners offering their property for sale according to an upward-sloping, log-linear housing supply function for each neighborhood.

5.1 Demand

Households make a discrete and exclusive choice regarding the neighborhood in which they are buying a 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. Potential buyers of a GHU compare the utility of their options using Equation 8, and choose the option that yields the highest indirect utility.

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

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 generic housing unit 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 (T1EV) distribution. Note that the within-nest correlation of utility levels goes to one as σ approaches one, and that for $\sigma = 0$ the within-nest correlation goes to zero and 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 with which one can estimate the whole demand system. This linear equation is given by Equation 10, where s_{jt} is the market share of area j in the whole market at time t and $\bar{s}_{jnt}$ is the market share of product j in nest n in period t.

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

5.2 Supply

Perfectly competitive agents sell a total of Q_{jt} generic housing units in neighborhood j at time t. 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 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. Given that our structural model introduces a number of specific functional forms, many of them non-linear, and we want to use this model to evaluate DiD estimates, we must check that the model can produce data that cannot reject the parallel trends assumption at common 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 the two main conclusions we extract from that exercise below.

The first conclusion from the simulation exercise is that our model is able to generate data that satisfy the parallel trends assumption for certain regions of the parameter space despite being highly non-linear in both its supply and demand side. 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. Intuitively, when neighborhoods experience large amenity shocks, this generates large changes in housing prices over time, thus rejecting any parallel trend test. 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 heterogenous and consumers re-sort less in reaction to the subsidy, which means less contamination and consequently a lower bias of the DiD estimate. These simulation results thus suggest that, in contexts of re-sorting, there could be a trade-off between satisfying the parallel trend assumption and having no SUTVA violations.

6 Estimation

6.1 Demand

We estimate our demand model on a dataset that has a single quantity and price for each combination of neighborhood and month-year. In order to control for differences in housing quality across neighborhoods, prices are the neighborhood $\times$ month-year fixed effects in a regression of transaction prices per square meter on those fixed effects plus a third-degree polynomial on the set of housing characteristics described in Section 3.

The demand regressions presented in Table 4 estimate Equation 10. In these regressions, the A_j and B_t amenity terms are captured by neighborhood and time fixed effects, respectively, and the time-varying, neighborhood-specific amenities ξ_{jt} constitute the error term. Since equilibrium prices and within-nest shares are correlated with these unobserved time-variant, neighborhood-specific amenities, OLS estimates in Table 4 are inconsistent. We address this endogeneity by leveraging the introduction of the LVIS policy 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.

Since time-varying, neighborhood-specific amenities are the error term of the regression, the identification assumption behind our set of instruments is that the tax break did not impact those 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 explain 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. However, the policy could have generated changes in amenities after the period we study, thus 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.

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.

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

The first OLS estimate of the price coefficient in Column 1 in Table 4 is positive, which is consistent with prices being positively correlated with neighborhood amenities. In Column 2 we add a rich set of neighborhood and month-year fixed-effects. These seem to remove part of the endogeneity, because the price estimate is still negative but much smaller, making it statistically indistinguishable from zero. Column 3 presents the estimates corresponding to the first of the three steps described in the previous paragraph. The final estimates, which are the ones we consider in the equilibrium counterfactuals in the next section, are presented in Column 4. These include a negative and significant coefficient for the price and a nested logit coefficient satisfying the restriction of being between 0 and 1. With these two estimated parameters, we then use Equation 10 to compute the ξ_{jt} terms, which we need for equilibrium computation.

6.2 Supply

Our inverse housing supply function in Equation 13 has two parameters, the percentage of the subsidy τ_{jt} and the inverse supply elasticity η. We externally calibrate the first one using González-Pampillón (2022)'s estimate of the LVIS subsidy representing 20% of the final housing price. We calibrate η such that there is an exact match between our benchmark reduced-form DiD estimate of -181 USD per square meter (Table 2 in Section 4) and its structural counterpart.[15]

The structural counterpart of the reduced-form DiD is the double difference in prices between subsidized and unsubsidized neighborhoods and between the equilibria of the model with and without the subsidy. Computing this structural double difference requires us to describe how we solve for the equilibrium of the model with and without the subsidy. We solve for the equilibrium of the model at the monthly level, thus mirroring the structure of our data. This procedure takes as inputs the IV-estimated demand parameters and the calibrated supply parameters presented in the previous section. It also uses 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 evaluate the counterfactual equilibrium prices and quantities when the subsidy is set to zero. Matching the structural double difference in prices with the -181 USD per square meter from the reduced-form DiD yields an inverse supply elasticity of $\eta = 0.33$.[16]

Finally, equilibrium prices computed from our structural model at the estimated parameters also exhibit parallel trends as we observe in the data and we use as an identification assumption in the DiD estimation. Figures A.11 and A.12 in Appendix A present the results for the equilibrium average housing prices by subsidy status and the monthly differences in housing prices between subsidized and subsidized areas.

Robustness We provide two robustness checks for our estimation of the inverse housing supply elasticity parameter. First, we internally calibrate η in the same way as in the benchmark but allowing the amenities to change between the equilibria with and without the subsidy, to reflect the evolution of the amenities that we observe in the data.

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

[16]This calibrated parameter implies a more elastic housing supply compared to available estimates (Saiz, 2010; Alves, 2021; Baum-Snow & Han, 2023). Note that ours is a monthly-level elasticity referring to property owners' decisions to sell their houses. This implies that we are looking at a very 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 compared to the selling decision we study.

Figure A.14 in Appendix A shows that amenities increase on average by 14.8% in both regions after the introduction of the subsidy. In the procedure described in the previous paragraph for computing the structural double difference in prices, the value of amenities was the same in the equilibria with and without the subsidy.[17] In this first robustness, we reduce amenities in all neighborhoods in the equilibrium without the subsidy such that amenities in the equilibrium with the subsidy are 14.8% higher. Calibrating the inverse supply using the structural double difference in prices from this robustness yields a value of 0.25, which is similar to our main calibration result.

Second, we obtain η by directly estimating Equation 13 instead of matching the reduced form moment. Since the residual of the supply curve can be correlated with the equilibrium quantities, we use a demand shifter as an instrument for the estimation. 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 presents the estimates of η under different specifications, including the instrumental variables estimates. Our preferred specification yields an estimate of $\eta = 0.29$ similar to our calibrated result.

7 Counterfactuals

In this section, we use the estimated model to solve for a set of counterfactual equilibria and achieve three goals. First, we decompose a structural equivalent of our DiD estimate into the three components presented in Section 2. This allows us to quantify the degree of contamination in this DiD estimate, which is indicative of the degree of bias in the benchmark reduced-form DiD estimate for the whole city. Second, we recover the incidence of the subsidy in terms of lower housing prices in the subsidized areas according to the model, and contrast it with the one obtained considering the benchmark reduced-form DiD estimate. 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 results of 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 has the reduced-form counterparts, when available. We have an equilibrium for each of the 32 months of the "post" period, so we report average results for all periods. Also, structural results for the whole city correspond to the average of all

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

neighborhoods. 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 introduction of the subsidy but without allowing for re-sorting between neighborhoods. We then calculate the re-sorting term as the difference between the ATT and the autarky. This re-sorting effect in Table 5 is large, indicating that the reduction in housing prices in the subsidized neighborhoods would have been much larger 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. Contamination can then be thought of as the structural counterpart of the bias of the reduced-form estimate. We obtain the contamination term as the difference in the average equilibrium prices of unsubsidized neighborhoods with and without the subsidy. The existence of contamination of around a quarter of the ATT in Table 5 indicates that the DiD term substantially underestimates the impact of the policy on the prices of the targeted neighborhoods.

Table 5: Decomposition of DiD Results Using the Structural Model

	Reduced-Form	Structural
DiD	-181	-181
ATT		-242
Autarky		-404
Re-Sorting		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, sorting, 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 re-sorting 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 last row of Table 5 shows that the presence of substantial contamination has implications for the conclusion regarding the incidence of the policy. The first column presents the incidence estimated using the reduced-form DiD with the contamination bias, while the second column shows the incidence from the structural approach. We calculate the incidence as the effect on the prices of the subsidized neighborhoods divided by the subsidy.[18] 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 by looking at the price faced by an average 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 between both estimates of the incidence is 3,589 USD, which amounts to 24.0% of Uruguay's GDP per capita in 2011, the year the policy was introduced.

Robustness As a robustness check of the share of contamination in our setting, 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. Contamination is similar to our benchmark result, varying from 9.7% to 22.6%. Thus, our main conclusion on the existence of large contamination in the setting under consideration is robust to these different methods of computing the inverse housing supply elasticity.

7.2 Determinants of Contamination and Bias

The previous analysis showed that contamination can lead to wrong conclusions on the effect of a place-based policy. To guide applied work in other contexts, it is useful to understand when contamination may matter more and thus lead to wrong conclusions. 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. In terms of guidance for applied work, 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 presents evidence of the positive correlation between contamination and demand-side 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.[19] Going back to the

[18] We obtain the amount of the subsidy by applying the 20% rate over the price that results from evaluating the unsubsidized inverse supply curve of the neighborhoods belonging to the subsidized area at those neighborhoods' quantities in the equilibrium with the subsidy.

[19] The heterogeneity index introduced in Section 4 assigns a scalar to each transaction. To obtain a

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 simple decomposition formula, the correlation does not only have the expected sign but it is also linear. Since we are looking at all neighborhoods and months, we have enough pairs to estimate the regression equivalent of the figure in panel b) of Figure 4, including a rich set of controls. Table A.3 in Appendix A shows robust and positive regression coefficients when controlling for none, either, and both neighborhood and month $\times$ year fixed effects.

Figure 4: Contamination, Border Heterogeneity and Diversion Ratios

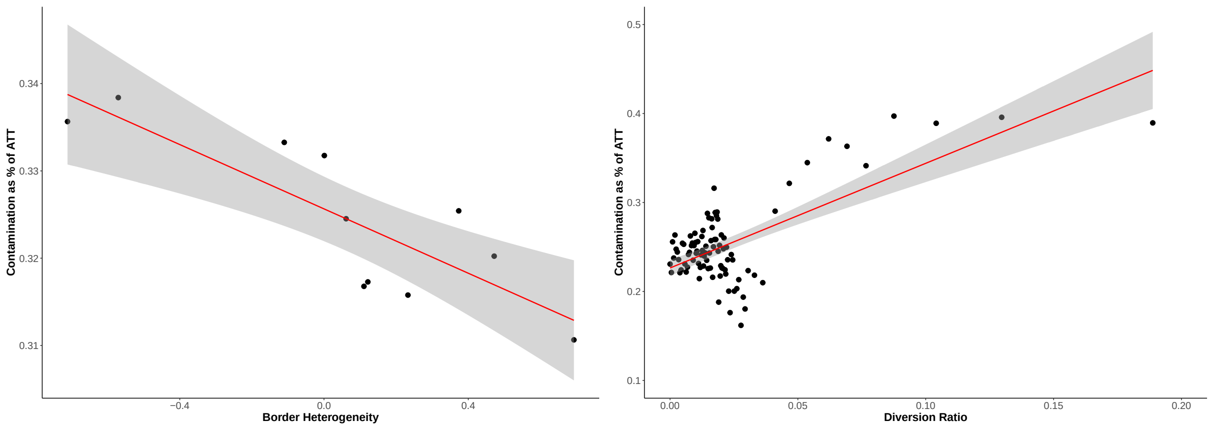

(a) Contamination and Border Heterogeneity (b) Contamination and Diversion Ratios

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 ATT is correlated with the intensity of demand-side substitution. Since the re-sorting 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

neighborhood pair-level index, we calculate the average value of the index for all transactions lying inside the area of the neighborhood pair.

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 may still matter 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 substantive, the identification of substantive re-sorting affecting the ATT can be normatively relevant, since the higher prices caused by re-sorting may offset part of the benefits of the subsidy for incumbent households.

7.3 The Role of Amenities

Finally, we discuss how the main conclusions of this section can be affected by the role of amenities. There is a consolidated literature quantifying the impact of new housing supply on amenities (Baum-Snow & Marion, 2009; Rossi-Hansberg et al., 2010; Diamond & McQuade, 2018). For the same policy we study, (González-Pampillón, 2022; Borraz et al., 2024) have found effects of the policy in the years following our period of study. Recently, there has been relevant progress on endogenizing those amenities in structural models (Almagro & Domınguez-Iino, 2019). In our model, amenities are allowed to vary over time but are not allowed to change in response to the policy. Figure A.14 presents the evolution of amenities in both the subsidized and unsubsidized areas. As discussed in Section 6.2, the estimates of η are also robust to the evolution of time-varying amenities.

Besides the robustness of our results to the previous issues, it is still relevant to understand the effect, on contamination, of an alternative situation in which changes in amenities would be correlated with the policy. Figure A.15 provides the intuition for such a situation. This figure is a variation of Figure 2 in which we let amenities change contemporaneously with the policy.

Figure A.15 highlights two additional forces when the relative value of the amenities in the subsidized neighborhood increases. On one hand, the magnitude of the ATT is further reduced as the subsidized area becomes more attractive. On the other hand, the magnitude of contamination increases, since the relative amenities of the non-subsidized area decrease and prices drop further.

In Table A.4 we approximate this counterfactual situation by computing the DiD, the ATT and contamination in counterfactual scenarios in which amenities in the subsidized areas deteriorate or improve with the policy. We show that a relative deterioration of amenities in the subsidized area of 20% can virtually make the contamination effect disappear. On the other hand, we show that a gentrification scenario, in which amenities increase more in the subsidized area, can make contamination substantially larger. In fact, a 20% improvement in amenities in subsidized neighborhoods would increase contamination from the 25.3% of the ATT that we find in our benchmark up to 76.1%.

8 Conclusion

Violations of the stable unit treatment value assumption (SUTVA) are a common threat to the identification of the equilibrium effects of policies. Because of the often non-random assignment of these policies, their study requires the use of quasi-experimental methods, with difference-in-differences being one of the most important. We discuss how estimates obtained by difference-in-differences may not recover the effect of policies in contexts where the re-sorting of agents changes the equilibrium outcomes of non-targeted units. Since place-based policies are one of the prominent examples of such programs, we illustrate how SUTVA violations can have serious consequences in terms of the welfare conclusions of large place-based interventions. We further 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.

We illustrate our methodological contribution by studying the impacts of a large place-based policy aimed at boosting housing construction in lagging areas of Montevideo, Uruguay. Because of our methodological focus, our study does not constitute a complete evaluation of the effects of this policy over time. 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.

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.

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. (2019). Location sorting and endogenous amenities: Evidence from amsterdam. *NYU, mimeograph*.

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

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.

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).

Banzhaf, H. S. (2021). Difference-in-Differences Hedonics. *Journal of Political Economy*.

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.

Bayer, P., McMillan, R., Murphy, A., & Timmins, C. (2016). A Dynamic Model of Demand for Houses and Neighborhoods. *Econometrica*, *84*(3), 893–942.

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.

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.

Callaway, B., & Sant'Anna, P. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230.

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.

de Chaisemartin, C., & D'Haultfœuille, X. (2023). Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: A survey. *The Econometrics Journal*, *26*(3), C1–C30.

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. *The 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*.

Ding, X., Bollinger, C., Clark, M., & Hoyt, W. H. (2023). Estimation of Welfare Effects in Hedonic Difference-in-Differences: The Case in School Redistricting.

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.

Ferman, B., & Pinto, C. (2019). Inference in Differences-in-Differences with Few Treated Groups and Heteroskedasticity. *The Review of Economics and Statistics*, *101*(3), 452–467.

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.

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

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, *225*(2), 254–277.

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. (2024). Economic Geography and Air Pollution Regulation in the United States.

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.

James, A. G., & Smith, B. (2020). Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution: Comment. *American Economic Review*, *110*(6), 1905–1913.

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. (2014a). 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.

Kline, P., & Moretti, E. (2014b). People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs. *Annual Review of Economics*, *6*(1), 629–662.

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

Martins, A., 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.

Moretti, E. (2011). *Local Labor Markets* (Handbook of Labor Economics). Elsevier.

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.

Rambachan, A., & Roth, J. (2023). A More Credible Approach to Parallel Trends. *Review of Economic Studies*.

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.

Serrato, J., & Zidar, O. (2016). Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms. *American Economic Review*, *106*(9), 2582–2624.

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

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, *225*(2), 175–199.

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.

A Appendix: Figures and Tables

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

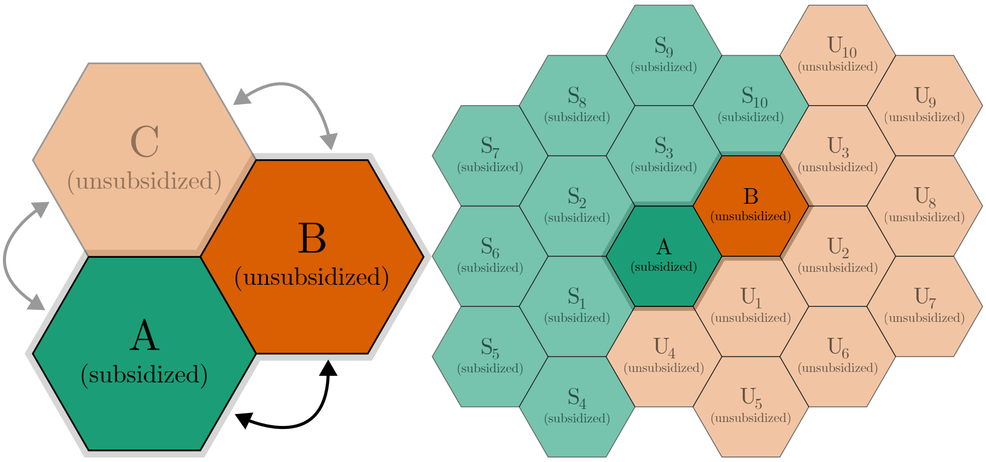

(a) Two Areas (b) Multiple Areas

Source: Authors' own illustration.

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

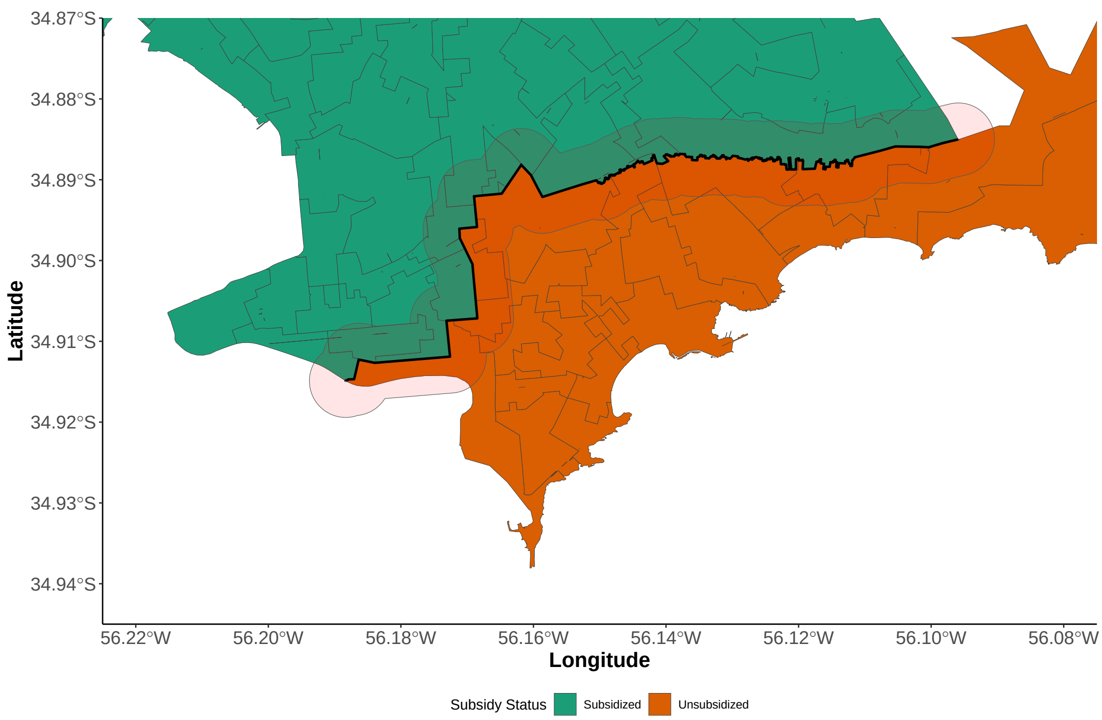

Source: Authors' own illustration using official shapefiles from the Geomatic Service of Uruguay.
Notes: The thicker line shows the border of the policy and the thinner lines the neighborhood limits. We defined neighborhoods using a spatial clustering algorithm, as explained in Appendix C. In panel a), 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

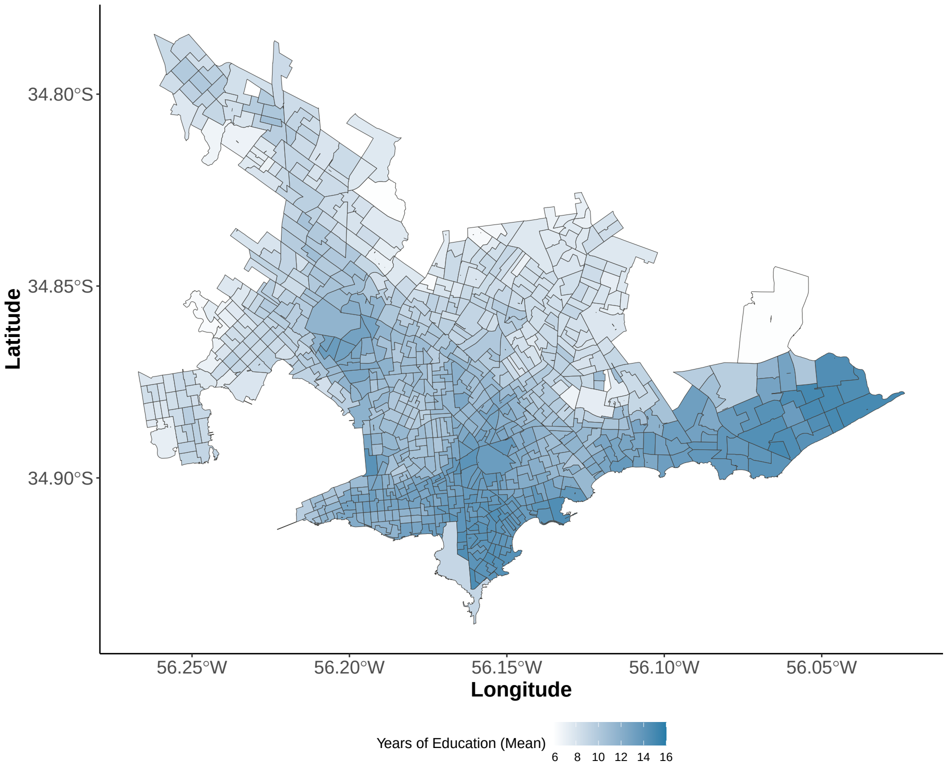

Source: Authors' illustration using official shapefiles from the Geomatic Service of Uruguay and microdata from the 2011 Uruguayan Census.
Notes: The tones of blue reflect the average years of education of the adult population living in each "segmento censal", an administrative unit comparable in size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

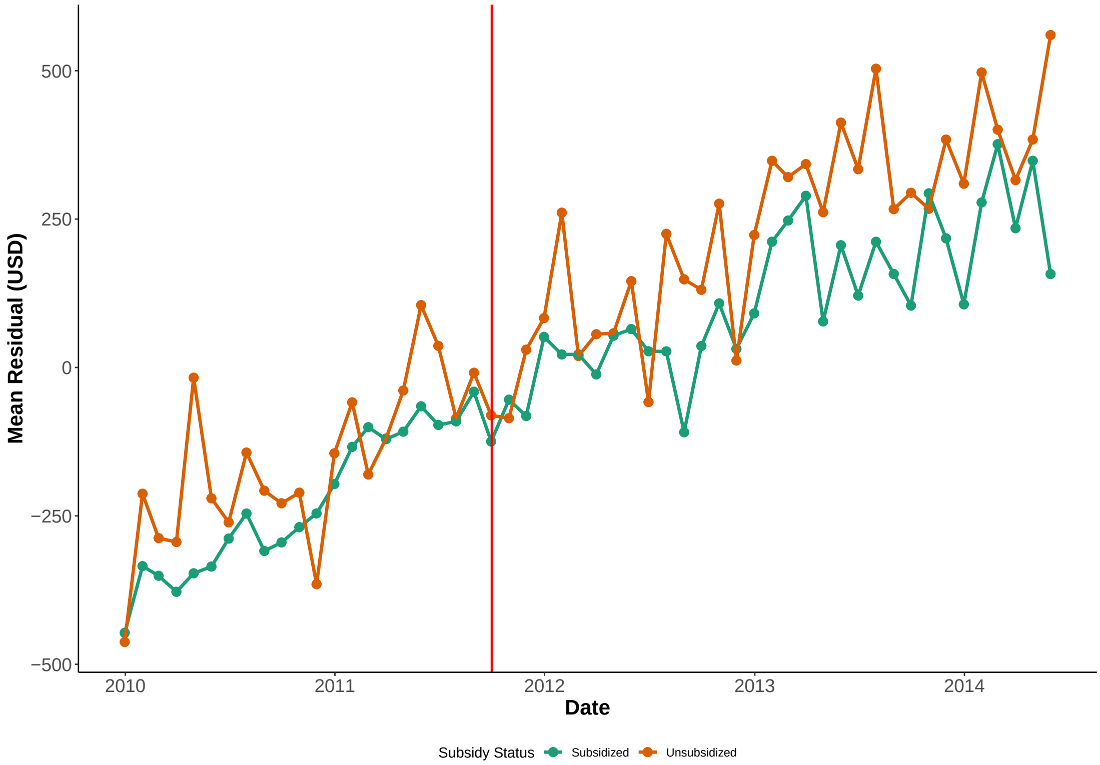

Source: Authors' calculations using housing transaction data from the National Registry Office in Uruguay and matched data on property characteristics from the country's national cadaster.
Notes: The graph plots, separately for transactions in the subsidized or unsubsidized areas, the average residualized 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

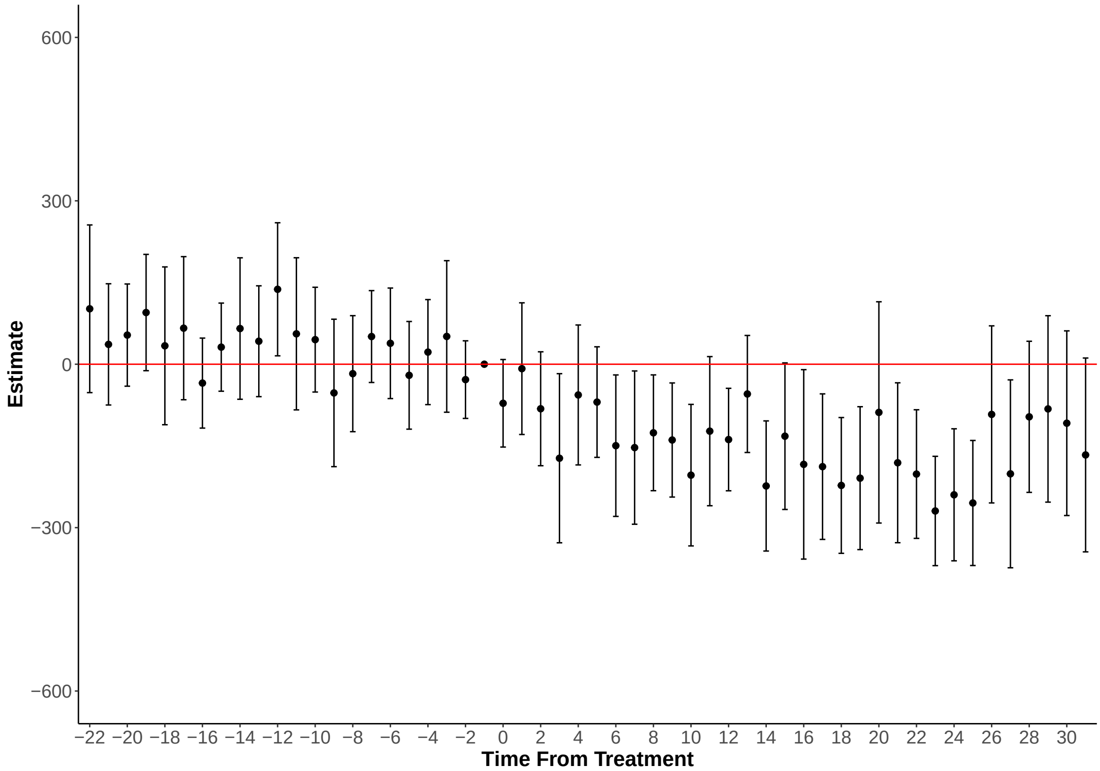

Source: Authors' calculations using housing transaction data from the National Registry Office in Uruguay and matched data on property characteristics from the country's national cadaster.
Notes: The graph plots the estimated coefficients, with their 95% confidence interval, of all year-month × subsidy 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

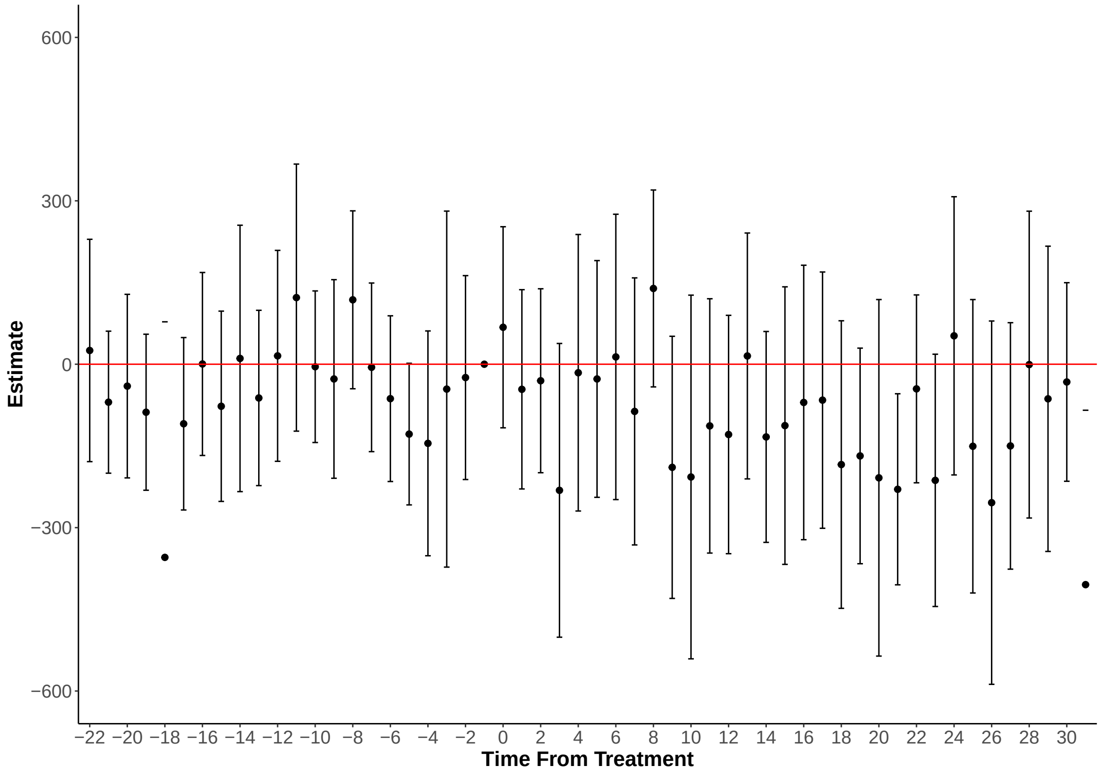

Source: Authors' calculations using housing transaction data from the National Registry Office in Uruguay and matched data on property characteristics from the country's national cadaster.
Notes: The graph plots the estimated coefficients, with their 95% confidence interval, of all year-month $\times$ 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

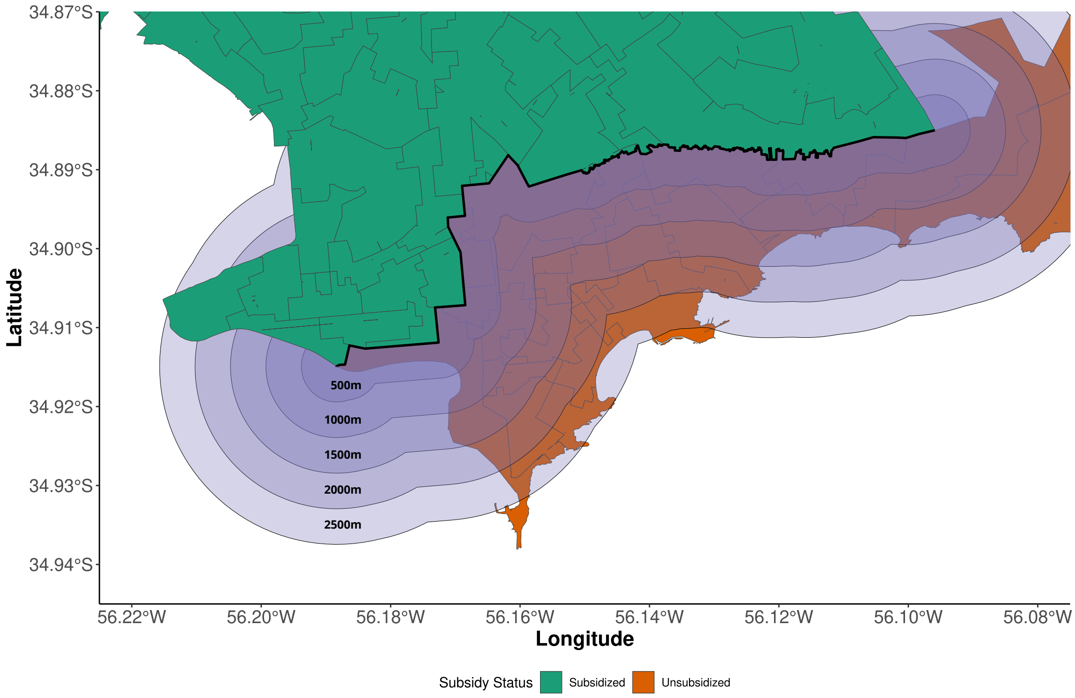

Source: Authors' own illustration using official shapefiles from the Geomatic Service of Uruguay.
Notes: The thicker line shows the border of the policy and the thinner lines the neighborhood limits. Each individual buffer covers the part of the unsubsidized area that is at most the distance indicated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

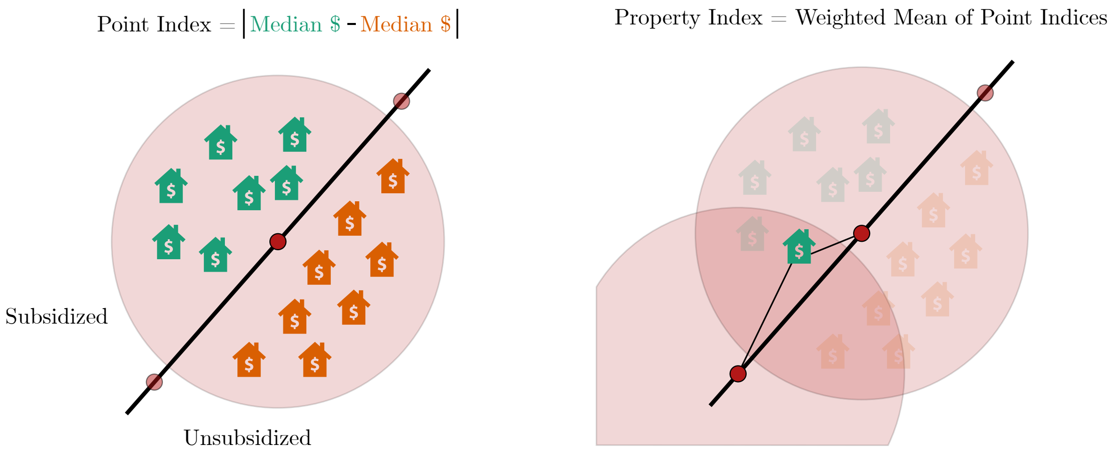

Source: Authors' own illustration.
Notes: The figure illustrates the method we use to compute a measure of heterogeneity along the border of the policy. The left panel shows how we compute the index of heterogeneity for a particular point on the policy border. The right panel shows how we aggregate point 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

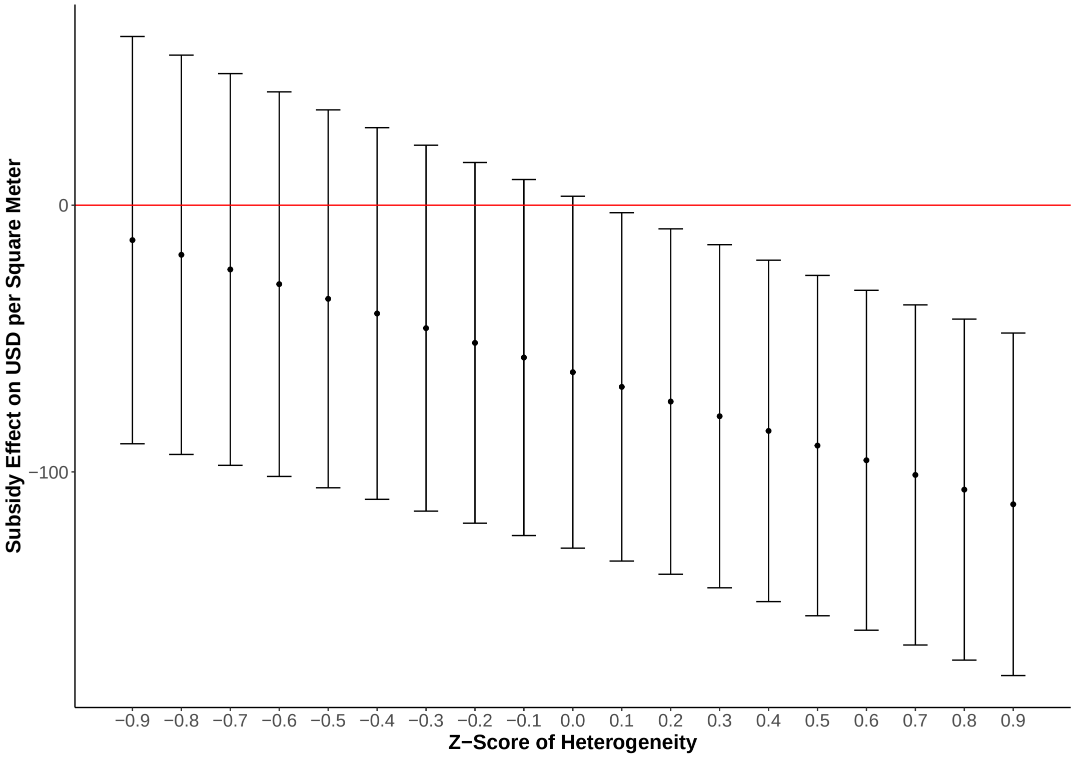

Source: Authors' calculations using housing transaction data from the National Registry Office in Uruguay and matched data on property characteristics from the country's national cadaster.
Notes: The graph plots the marginal effects, with their 95% confidence interval, for different values of the Z-score, 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

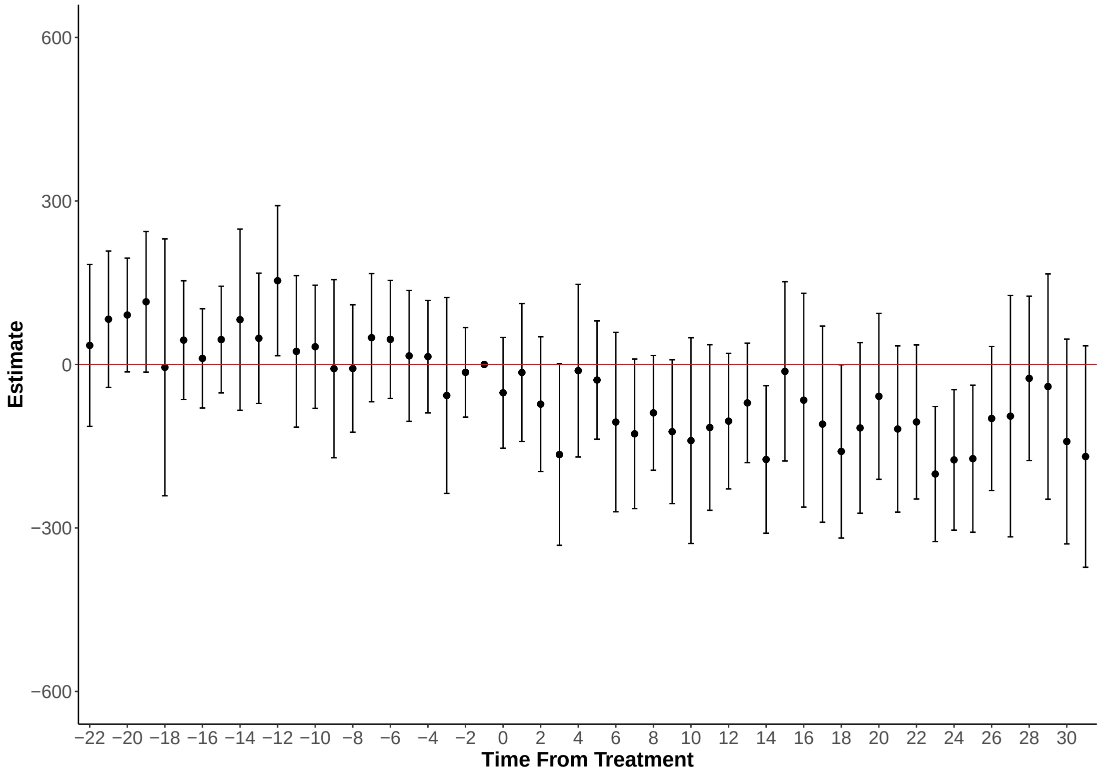

Source: Authors' calculations using counterfactual equilibrium exercises.
Notes: The graph plots the estimated coefficients, with their 95% confidence interval, of all year-month × subsidy dummies of a regression of equilibrium housing prices on month-year × subsidy dummies. The omitted fixed effect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

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

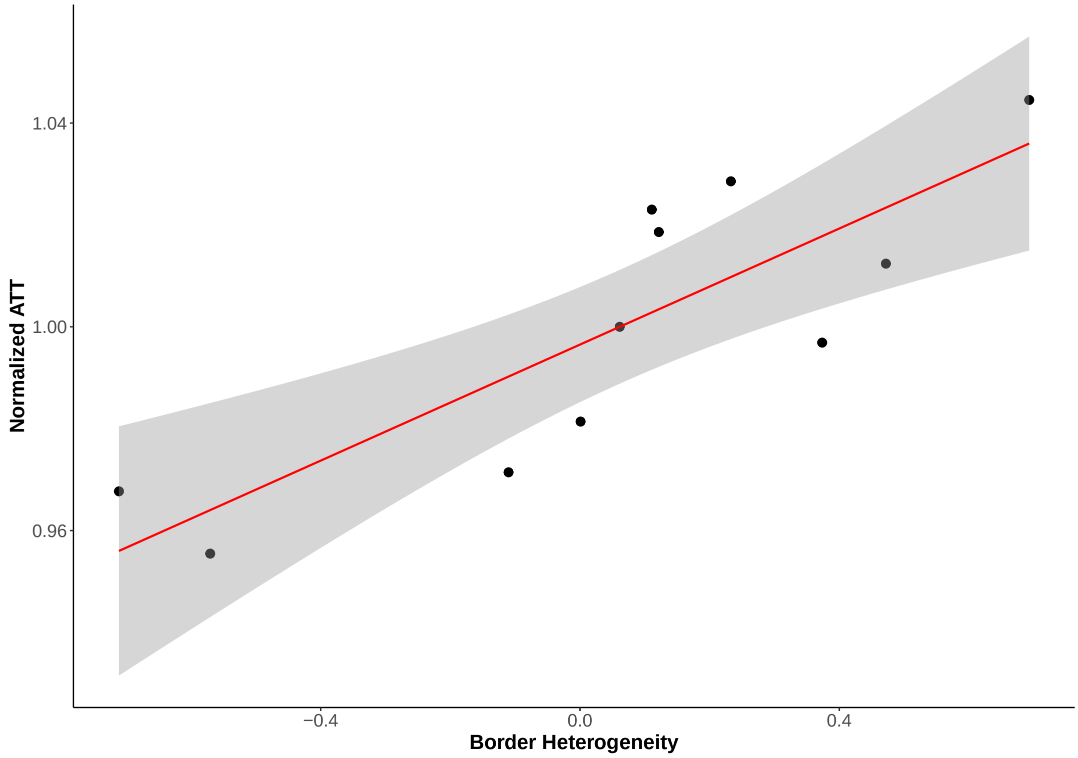

Source: Authors' calculations using counterfactual equilibrium exercises.
Notes: Each of the 13 dots in the figure represents a subsidized-unsubsidized neighborhood pair. These are all the neighborhood pairs lying across the border of the policy.Figure A.2 in the Appendix A provides a map of the neighborhoods 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

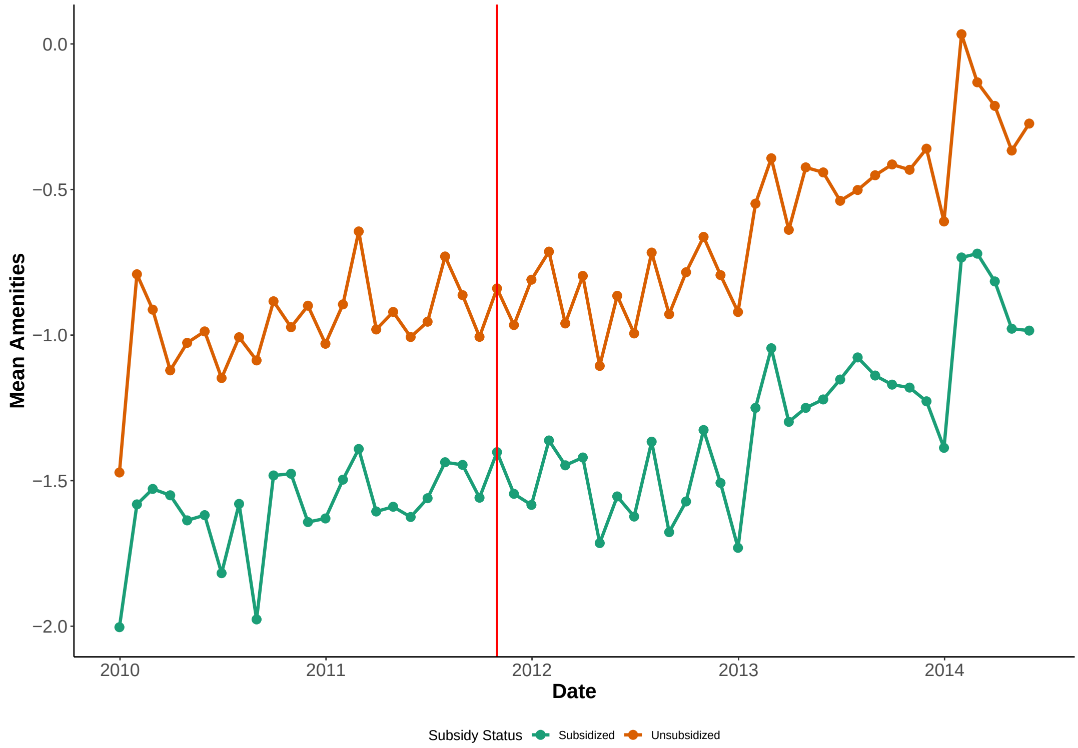

Source: Authors' calculations using counterfactual equilibrium exercises.
Notes: Amenities are obtained by removing the price effect from the mean utility of each product in each period. The individual lines represent the time series of the sales-weighted mean of these amenities by subsidy status.

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

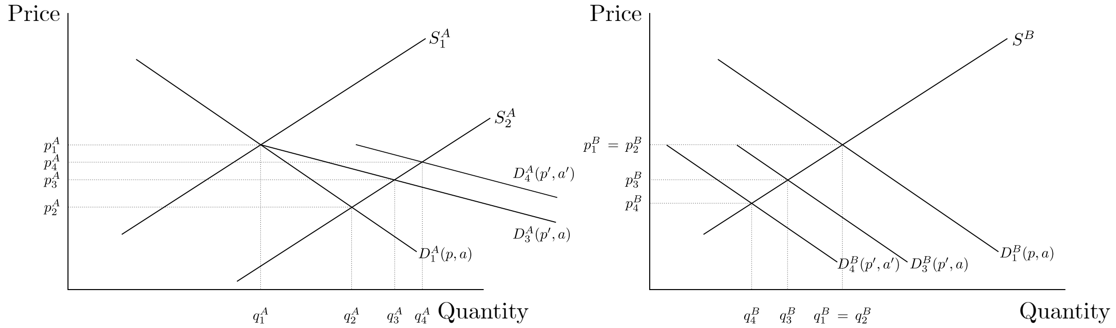

Table A.1: Difference-in-Differences Regressions - Heterogeneity

	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 × Month	Year × 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 (1)	ATT (2)	Contamination (3)	% Cont./ATT (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.[20] 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.

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

[20] 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 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.

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

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.

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

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 re-sorting" and "indirect contamination". The indirect re-sorting, i.e. the autarky change in C multiplied by the indirect re-sorting 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 unsubsized 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 re-sorting 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.

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

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.

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

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

C Appendix: Neighborhood Definition with a Spatial Clustering algorithm

We use the spatial clustering algorithm "SKATER" (Spatial 'K'luster Analysis by Tree Edge Removal), developed by Martins et al. (2006), which has four convenient features for the problem at hand. 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 in order 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 a total of 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.

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

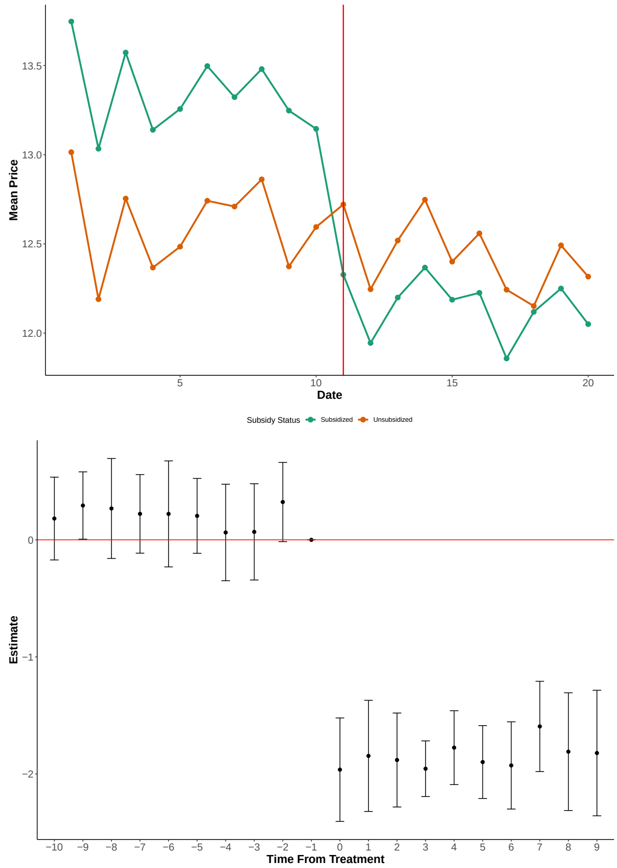

The second takeaway is to characterize under which type and size of heterogeneity our model rejects the parallel trends. To analyze this issue we perform simulations over several values of the heterogeneity parameters. In these simulations, the variance for j terms (σ_j) is limited to the set $\{0.5, 1..0\}$, 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 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, the third takeaway is that there is 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 $\sigma_j = 1$

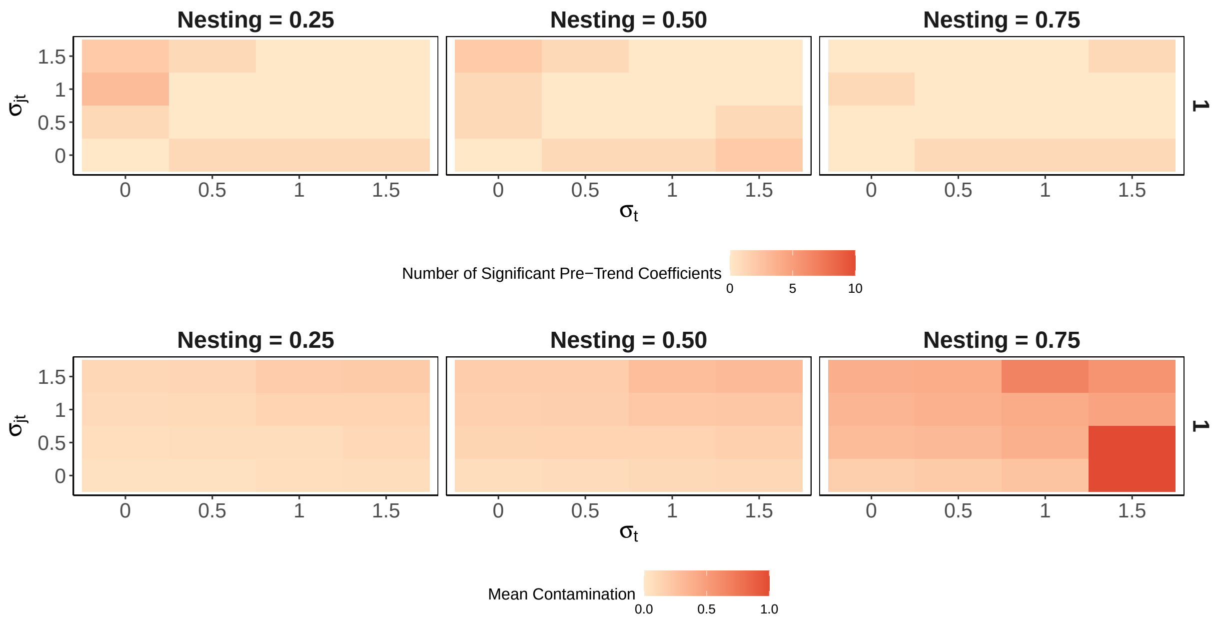

Figure D.3: Parallel Trends and Contamination Effects in Simulations for $\sigma_j = 0.5$