CAF - WORKING PAPER #2023/11
Primera versión: December 21, 2023
Esta versión: December 21, 2023

Difference-in-Differences in Equilibrium: Evidence from Placed-Based Policies

Guillermo Alves[1] | **William H. Burton**[2] | **Sebastián Fleitas**[3]

[1]Economista Principal. Dirección de Investigaciones Socioeconómicas. CAF - Banco de esarrollo de América Latina y el Caribe. galves@caf.com

[2]KU Leuven and FWO

[3]PUC-Chile and CEPR.

Violations of the stable unit treatment value assumption (SUTVA) are a common threat to identification of the effects of policies causing the resorting of agents between treated and untreated groups. We show that in such contexts the difference-in-differences estimator can be decomposed into three effects (autarky, resorting and contamination). We also show that demand and supply elasticities are "sufficient statistics" for the relative size of these effects and that there exist a trade-off in terms of heterogeneity between SUTVA and parallel trends assumption violations. We illustrate our argument by studying a large placed-based tax break for the construction of residential housing in Uruguay. First, we obtain a series of difference-in-differences estimates of the effect of the policy on housing prices and show that they differ considerably depending on the degree of heterogeneity between subsidized and unsubsidized areas. Consistent with our conceptual framework, prices fall substantially when comparing heterogeneous areas, and very little or not at all when comparing similar areas. Second, we estimate a structural model of supply and demand for neighborhoods that rationalizes those different estimates and allows us to recover the three effects as well as the welfare impact of the policy. Overall, we find that SUTVA violations account for 25% of the effect on subsidized areas and lead to a sizable underestimation (24 p.p.) of the incidence of the tax break on subsidized areas.

KEYWORDS

Differences-in-differences, Place-based policy, Housing Prices

CAF - DOCUMENTO DE TRABAJO #2023/11

Primera versión: 21 de diciembre de 2023. Última versión: 21 de diciembre de 2023

Diferencias-en-diferencias en equilibrio: Evidencia desde las Políticas Basadas en el Lugar

Guillermo Alves[1] | **William H. Burton**[2] | **Sebastián Fleitas**[3]

[1]Economista Principal. Dirección de Investigaciones Socioeconómicas. CAF - Banco de esarrollo de América Latina y el Caribe. galves@caf.com

[2]KU Leuven and FWO

[3]PUC-Chile and CEPR.

Las violaciones del supuesto de no interferencia (SUTVA en inglés) dificultan la evaluación de políticas cuando estas causan movimiento de agentes entre grupos de tratamiento y control. En este trabajo se muestra que en esos casos el estimador de diferencias-en-diferencias se puede descomponer en tres efectos (autarquía, relocalización y contaminación), y que las elasticidades de demanda y oferta son "estadísticos suficientes" para medir la participación relativa de esos tres efectos y que existe un trade-off entre evitar violaciones de SUTVA y tener tendencias paralelas. El trabajo examina las implicancias empíricas de este argumento estudiando una política que otorga importantes exoneraciones tributarias a la construcción de viviendas en Uruguay. En primer lugar, se estiman una serie de diferencias-en-diferencias del impacto de la política y se encuentra que su magnitud varía considerablemente según el grado de heterogeneidad entre las áreas de tratamiento y control escogidas: los precios caen mucho más cuando se comparan áreas heterogéneas y poco o nada cuando se comparan áreas más similares. En segundo lugar, se estima un modelo de oferta y demanda de vivienda que racionaliza esas estimaciones, descompone el estimador de diferencias-en-diferencias en sus tres efectos, y mide el impacto de la política en el bienestar. Las violaciones de SUTVA suponen un 25 % del efecto en las áreas subsidiadas y generan que el estimador de diferencias-en-diferencias subestime la incidencia de la exoneración tributaria en 24 puntos porcentuales.

KEYWORDS

Diferencias-en-diferencias, Políticas basadas en el lugar, Vivienda

1 | INTRODUCTION

Non-experimental studies of policies causing the resorting of agents between treatment and control groups may suffer from violations of the crucial stable unit treatment value assumption (SUTVA) (Donaldson, 2015). Place-based interventions are prominent examples of these type of policies (Kline and Moretti, 2014b). Because these policies are usually not randomized, researchers rely on non-experimental methods - such as difference-in-differences (DiD) - to study their effects (Baum-Snow and Ferreira, 2015). Identifying the causal effects of these policies using these methodologies requires no violations of SUTVA. While most of the recent developments in the difference-in-differences methodology have focused on the parallel trends assumption and the staggered treatment, less attention has been paid to the SUTVA violations (Roth et al., 2023).

In this paper, we discuss the difference-in-differences estimator in the presence of violations to SUTVA in a context of resorting of agents between control and treatment groups. We show that in these contexts the difference-in-differences estimator can be decomposed into three effects. First, an "autarky effect" captures what would happen to treated areas if there were in isolation, and therefore no relocation effects existed. Second, a "resorting effect" captures the effect on treated areas caused by the inflow of agents into those areas. Third, a "control unit contamination" captures the effect on the control area caused by the outflow of agents away from this area.

We then apply these insights to the study of a place-based policy giving substantial tax breaks for housing development in lagging areas of Montevideo, the capital city of Uruguay. We start the analysis by using administrative data on the universe of housing transactions in Montevideo before and after the policy and estimating a series of difference-in-differences with housing prices as our dependent variable. We find three difference-in-differences results that are consistent with our conceptual framework. First, 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. Second, when we follow the common practice of using only observations close to the border, estimates are very small negatives or zeros. Third, consistent with the presence of contamination effects, the absolute magnitude of these border estimates increases with a measure of heterogeneity between both sides of the border and when we use control units located further away from the border.

With a simple linearization of a model of the supply and demand of housing in a city, we provide an analytical formula showing that the relative size of each of the three effects contained in the difference-in-differences estimator depends on the demand-side substitution patterns between neighborhoods as well as the supply elasticities of the neighborhoods. Importantly, 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 using of matching techniques (Neumark and Kolko, 2010; Chen et al., 2022).

Using our decomposition formula, we analyze three types of situations in which the assumptions on the network structure may or may not justify the implementation of a difference-in-differences approach. First, when the relocation of agents causes spillovers that are very local, it can be reasonably assumed that distant areas experience no spillovers. In those cases, identification of the effects of the policy can be achieved by comparing the treated area versus distant ones (Delgado and Florax, 2015; Clarke, 2017; Butts, 2021). A prominent example of this approach is Kline and Moretti (2014a), who drop neighboring counties from their control group in their evaluation of the impact of the Tennessee Valley Authority (TVA).

In many economic settings, the resorting of agents from untreated into treated areas implies that truly untreated areas may not exist, or may be hard to credibly detect and justify. 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 non-treated areas. Busso et al. (2013)'s study of Empowerment Zones constitute an example of this second type of situation in which difference-in-differences estimates can recover the effect of the policy.

A third type of situation occurs when the policy is large enough such that its effects extend to non-targeted areas. Consider, for instance, the case of common supply-side subsidies for housing construction which target entire neighborhoods in a given city, such as the Opportunity Zones program in the US. These policies redirect housing demand from non-subsidized into subsidized areas, potentially causing housing prices to fall in non-subsidized areas. This effect of the policy on non-subsidized areas constitutes a violation of SUTVA, and thus invalidates difference-in-differences designs.

We add an additional and final layer of structure to our analysis and further use our transaction data to estimate a structural model of the supply and demand of housing across Montevideo's neighborhoods. We model the demand for housing as the discrete choice problem of choosing a neighborhood within a city. 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 and Domınguez-Iino, 2019; Anagol et al., 2021) and across cities (Diamond, 2016; Alves, 2021). We estimate the price elasticity of the housing demand in a nested logit model using the introduction of the tax break to build a set of supply-shifting instruments. The housing supply in the model is characterized by a log-linear supply function for each neighborhood (Saiz, 2010; Baum-Snow and Han, 2023). We internally calibrate a common inverse supply elasticity for all neighborhoods with a procedure that is similar to Berger et al. (2022). This elasticity is the one matching the difference-in-difference object implied by our equilibrium model with our reduced form difference-in-differences estimate. As is common in the quantitative spatial literature, our main insights from the model arise from solving for a set of counterfactual equilibria (Ahlfeldt et al., 2015; Donaldson, 2017; Monte et al., 2018; Caliendo et al., 2019; Fajgelbaum et al., 2019).

We show that our model fits the data in terms of reproducing the parallel trends that we find in the descriptive analysis. Moreover, by solving for a series of counterfactual equilibria, we estimate directly the three additive effects behind the difference-in-differences estimator. We find that the "resorting effect" accounts for 40% of the "autarky effect" and that the "contamination effect" represents 25% of the total effect of the policy on the subsidized area, which is given by the sum of the autarky and resorting effects. Additionally, we show that the SUTVA violation causes the reduced-form difference-in-differences analysis to underestimate the share of the subsidy that reaches consumers by more than 24 percentage points. Once accounted for contamination, the incidence of of the subsidy is larger, and the difference in the estimated effect of the incidence of the subsidy amounts to around 30% of Uruguay's GDP per capita in the year the policy was introduced. Our methodological argument is thus quantitatively relevant in terms of policy implications.

Finally, we use the equilibrium counterfactuals from our estimated model to revisit the relationship we find in the reduced-form analysis between the heterogeneity across control and treated units and the size of the difference-in-differences estimate. Consistent with our decomposition formula, we confirm that contamination is positively correlated with both the homogeneity and the diversion ratios between those units. Importantly, this implies that the lower values of the reduced-form differences-in-difference estimates we obtained for more homogeneous units are caused by higher bias caused by contamination and are not just about treatment heterogeneity. This posits a note of caution to common research designs that try to maximize the comparability between control and treatment group.

Our paper contributes to three main strands of literature. First, we contribute to the literature on causal inference in urban and regional economics. In their comprehensive review of this literature, Baum-Snow and Ferreira (2015) include difference-in-differences as one of the main techniques for obtaining causal estimates. The authors highlight how the re-sorting of individuals between treatment and control areas constitutes a serious threat to identification in difference-in-differences designs in spatial settings. This threat can be seen as a special case of dealing with spatial spillovers in difference-in-differences settings, a topic that has received attention from several previous works (Clarke, 2017; James and Smith, 2020; Butts, 2021; Huber and Steinmayr, 2021; Myers and Lanahan, 2022).

Currently, successful identification of the effects of place-based policies with difference-in-differences designs in the presence of spillovers is restricted to two contexts. First, spatial spillovers can be handled by defining large enough treatment and control units such that spillovers are contained within those units (Feyrer et al., 2017; Huber and Steinmayr, 2021). Second, researchers may employ successive "donuts" or "rings" around the treatment area to flexibly capture the effect of the spillovers (James and Smith, 2020; Butts, 2021; Myers and Lanahan, 2022). As spillovers eventually fade away far enough from the treatment, the comparison of treated areas against those spillover-free areas yields an average treatment effect on the treated (Clarke, 2017). However, when policies are large enough, those spillover-free areas may not exist or may be hard to credibly find. Also, natural (sea, mountains) or man-made (parks, highways) barriers may restrict the construction of far-enough rings. We provide a methodological framework to empirically study the effects of place-based policies in such contexts.

Second, we contribute to the literature on the evaluation of place-based policies that subsidize the development of lagging areas. 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 and Zidar, 2016). We show that the existence of heterogeneous mobility patterns of agents across areas can generate wrong conclusions about the efficiency of placed-based policies when estimates are obtained by comparing only certain areas.

Third, we contribute to the burgeoning literature on the methodological improvement of difference-in-differences estimates (de Chaisemartin and D'Haultfœuille, 2021; Roth et al., 2023). Recently, there has been substantial progress in designs with multiple periods and variation in treatment timing (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Goodman-Bacon, 2021), potential violations in parallel trends (Rambachan and Roth, 2023; Roth and Sant'Anna, 2023), and improved inference (Ferman and Pinto, 2019). In their review of the state of the literature, Roth et al. (2023) include spillovers as one of the main areas for future research in this literature, with a special mention to spatial spillovers. We analyze a specific type of spatial spillover that we believe has high economic relevance. These are spillovers generated by the movement of economic agents across space in reaction to place-based policies. We stress the limitations of difference-in-differences designs in terms of recovering structural parameters of interest in those circumstances and show how structural methods can inform those estimates.

2 | DIFFERENCE-IN-DIFFERENCES IN EQUILIBRIUM

2.1 | Setup

We now study DiD estimates with a simple model of supply and demand for homogeneous housing units in the real state market across neighborhoods within a city. In the model, the demand side consists of households who decide if they want to buy a housing unit in a given neighborhood within the city or remain outside the city. The supply side is given by property-owners who choose the number of housing units they want to sell in each neighborhood. Both households and property-owners make a decision in each period which is independent of previous and future periods.

Neighborhood choice is the main focus of households when evaluating where to live in contemporary large cities characterized by sharp amenity differences between neighborhoods. There are two main determinants of households' discrete choice between neighborhoods. These are neighborhoods' housing prices and amenities. The latter captures both the relatively fixed aspects of the attractiveness of a neighborhood, such as the distance to the sea, and also the time-varying ones, such as crime.

For the supply side of the model, we assume that atomistic housing owners decide the number of housing units to sell in each neighborhood and period. Higher prices induces a higher supply of houses for sell and this relationship between prices and quantities offered is represented by an upwards sloping supply function.

2.2 | Decomposition of the DiD Estimator

At its very core, a difference-in-differences (DiD) estimator can be written as:

$$\hat{\beta}_{DiD} = (y^{Post}_{Treated} - y^{Pre}_{Treated}) - (y^{Post}_{NotTreated} - y^{Pre}_{NotTreated}) \tag{1}$$

with y denoting the variable of interest. The change in the untreated observations is used to compute changes over time, which is then subtracted from the change in the treated observations in order to identify the policy's effect. In situations in which the policy induces the resorting of agents between the treated and control groups, the DiD no longer identifies the effect of the policy (Baum-Snow and Ferreira, 2015). We next provide a decomposition of the DiD under those circumstances, which will help us rationalize the patterns observed in our DiD estimates seen in the previous section.

Without loss of generality, we highlight the various components underlying the DiD estimator using the real state market as an example. The outcome variable in our case is the price of housing, p^j_t, with j denoting neighborhood and t denoting time. In order to fix ideas, we start with two extreme cases of resorting of agents between the control and the treatment groups. First, we show the case of autarky, in which there is no substitution, and then we move to perfect substitution. After presenting these two extreme cases, we introduce a generalised decomposition for two neighbourhoods. Throughout the section we focus on demand-side resorting of households and thus abstract away from supply-side substitution. Supply-side substitution could easily be accommodated in this framework.

Figure 1 highlights an autarky situation in which consumers are only willing to consider housing in one particular neighborhood $j \in \{A, B\}$, but not the other. Implementing a supply-side subsidy in neighborhood A would first shift supply outwards in neighborhood A. Because of lower prices, demand for housing in neighborhood A expands. Neither demand nor supply are affected in neighborhood B and thus prices do not change in this

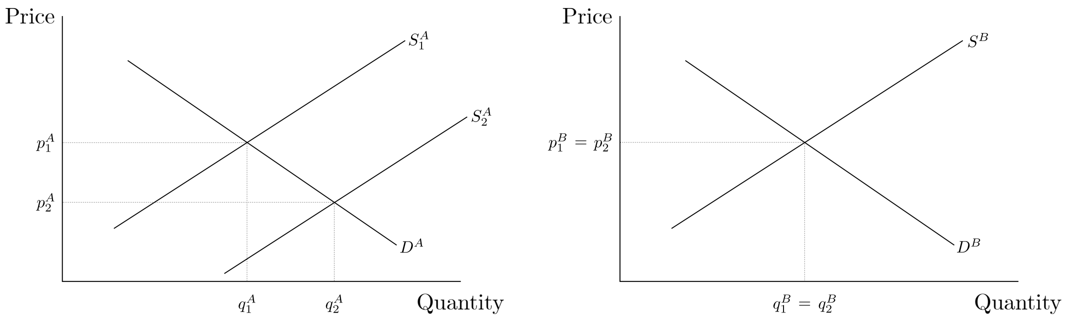

FIGURE 1 Neighborhoods A and B are independent products

neighborhood. The estimated DiD in this scenario is equal to the difference in prices between periods 2 and 1 in 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 does capture the effect of the policy on the targeted areas. We show next that this is not the case when there is resorting of agents between the two areas.

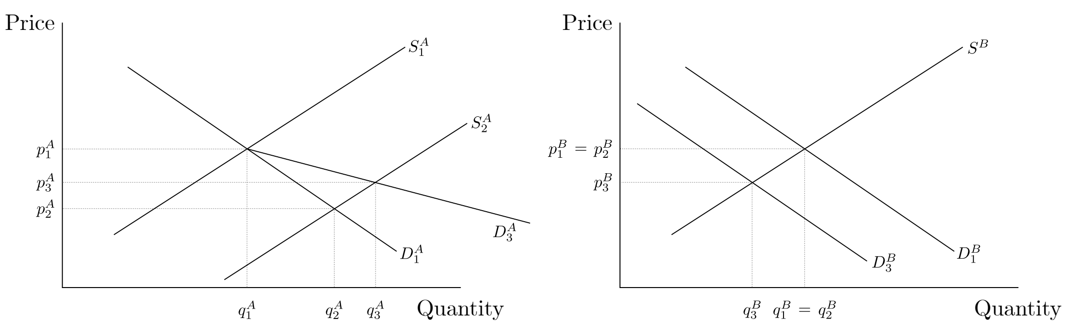

FIGURE 2 Neighborhoods A and B are perfect substitutes

Figure 2 highlights a situation in which consumers consider housing in different districts to be perfect substitutes. Consistent with this assumption of perfect substitutability, prices at $t = 1$ coincide between both neighborhoods. Again, a supply-side policy is enacted in district A, pushing housing prices p_1^A downwards until p_2^A. However, due to the assumed pattern of substitution, there is a new round of effect which we index as happening at $t = 3$. Prices in neighborhood A increase until p_3^A because consumers from district B are switching locations. Estimating the effect of the same policy using the DiD approach now yields the following:

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

We see that in the case of perfect substitution between different districts, the estimated DiD effect contains not only the autarky effect from before, but also the price increase due to higher demand for housing in the subsidized district A, as well as the price decrease in district B. As indicated in Equation 2, we call the additional demand effect in district A "resorting", and the price change in district B as "contamination". This last term "contaminates" the DiD effect because it causes that this estimation techniques does not recover the effect of the policy on the targeted areas, which is given by the sum of the other two terms in Equation 2. We next introduce a new formula that helps to understand which are determinants of the relative size of contamination relative to the other two effects. That relationship thus defines the relative size of the bias of the DiD estimate of the effect of the policy on the targeted area.

$$\hat{\beta}_{DiD} = \underbrace{(p_2^A - p_1^A)}_{\text{Autarky}} + \underbrace{(p_3^A - p_2^A)}_{\text{Resorting}} - \underbrace{(p_3^B - p_2^B)}_{\text{Contamination}} \tag{2}$$

2.3 | DiD Estimator from Supply and Demand

Having discussed two extreme versions of demand patterns, we now impose slightly more economic structure to understand the estimated DiD effect in terms of demand and supply. We specify demand for housing in a neighborhood j at a given vector of market prices $\mathbf{p}$ by $Q_D^j(\mathbf{p})$. The inverse housing supply is specified by $P_S^j(q^j)$. With two neighbourhoods, of which one is subsidized while the other is not, the estimated DiD effect can be expressed by the following approximation:

$$\beta_{DiD} \approx \underbrace{(p_2^A - p_1^A)}_{\text{Autarky shock}} \times \left[\underbrace{1}_{\text{Autarky term}} + \underbrace{\frac{\partial Q_D^A}{\partial p^A} \times \frac{\partial P_S^A}{\partial q^A}}_{\text{Resorting term}} - \underbrace{\frac{\partial Q_D^A}{\partial p^A} \times \frac{\partial P_S^B}{\partial Q_D^B} \times DR_{A,B}}_{\text{Contamination term}} \right] \tag{3}$$

Equation 3 highlights that any estimated DiD effect is actually a scaled version of the policy's effect in autarky. The scaling factor depends crucially on the responsiveness of demand and supply in the two neighborhoods, and also on the demand diversion ratio between the two. A full derivation of Equation 3 can be found in Appendix B.

In Equation 3, $\frac{\partial Q_D^A}{\partial p^A} \times \frac{\partial P_S^A}{\partial q^A}$ indicates how much of the "scaling" of the policy's effect in autarky is due to the resorting effect, and $-\frac{\partial Q_D^A}{\partial p^A} \times \frac{\partial P_S^B}{\partial q^B} \times DR_{A,B}$ how much is due to the contamination effect. The latter of these two effects 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 treated areas. Note that this effect increases linearly with respect to each of its three terms. Although this formula has the strong limitation of featuring only two areas, we later show that the linear relationship between contamination and the diversion ratio holds when we estimate and solve for a specific model with many neighborhoods.

The explicit formula for the role of the contamination effect in Equation 3 allows us to analyze some of the main identification strategies and assumptions followed by the previous literature. A first strand of the literature assumes that there is a sufficiently far away area such that it is unaffected by the policy and use the difference-in-differences estimate between the area of interest and this area. This strategy is often referred as "ring approach" and Kline and Moretti (2014a); Clarke (2017); Butts (2022) are some examples of relevant papers implementing it. Note that the formula in Equation 3 shows that this is analog to assume that the diversion ratio between the area of interest (A) and the control area (B) is zero ($DR_{A,B} = 0$). This assumption implies that the contamination effect is zero and thus the DiD does recover the true effect of the policy on targeted areas. A limitation of this strategy is that, when policies are "large", all areas could be in principle affected and it may be difficult to find an area in which $DR_{A,B} = 0$. The formula shows that when researchers have estimates of the demand for different neighbors, they can directly test this hypothesis of the existence of an unaffected area.

A second strand of the literature assumes that there is a large enough number of areas such that each of them is too small to have an effect on the rest through resorting. Examples of this strategy are Busso et al. (2013); Chen et al. (2022). The formula in Equation 3 shows that this assumption is analog to assume that $\frac{\partial Q_D^A}{\partial p^A} = 0$, implying that in these contexts the DiD captures only the autarky effect.

This discussion of some of the relevant methods in the literature through the lens of the formula already anticipated a key takeaway from Equation 3 for empirical work. The relative sizes of the three effects (autarky, resorting and contamination) can be estimated using only supply and demand elasticities. In this sense, the elasticities of supply and demand constitute "sufficient statistics" for the relative size of the effect (Saez, 2001; Chetty, 2009). Interestingly, because of the multiplicative way in which these elasticities enter the formula, researchers could discard contamination with only demand or (inverse) supply elasticity estimates if they happen to be close to zero. This situation is the one illustrated by the two strand of literatures above and can also happen in contexts with very elastic housing supply elasticities. The reverse, a situation where contamination is likely to be relevant and bias DiD estimates, occurs in contexts with relatively inelastic housing supplies Baum-Snow and Han (2023). Note that this discussion refers to the relative sizes of the three effects. In order to get the absolute value of those effects, one needs an estimate of the autarky term, which is a counterfactual object.

Finally, this formula can be generalized to having more than one subsidized and unsubsidized areas. In this case, the DiD estimator can be written as:

$$\begin{aligned}\hat{\beta}_{DiD} \approx\ & \underbrace{(p_2^A - p_1^A)}_{\beta^{AUT}_{DiD}} \\ & + \underbrace{\frac{\partial P_S^A}{\partial q^A} \times \left(\sum_{u \in US} \frac{\partial Q_D^A}{\partial p^u} \times \frac{\partial P_S^u}{\partial q^u} \times \left[\sum_{s \in S} \frac{\partial Q_D^u}{\partial p^s} \times (p_2^s - p_1^s) \right] \right)}_{\textit{Resorting from Unsubsidized Area(s)}} \\ & + \underbrace{\frac{\partial P_S^A}{\partial q^A} \times \left(\sum_{s \in S \setminus A} \frac{\partial Q_D^A}{\partial p^s} \times (p_2^s - p_1^s) \right)}_{\textit{Resorting from other Subsidized Area(s)}} \\ & - \underbrace{\frac{\partial P_S^B}{\partial q^B} \times \left(\sum_{s \in S} \frac{\partial Q_D^B}{\partial p^s} \times (p_2^s - p_1^s) \right)}_{\textit{Contamination from Reference Area}}\end{aligned} \tag{4}$$

The formula has similar terms than before, but it also has some differences. First, in this case, the resorting effect in equilibrium can be negative or positive. The resorting of agents can attenuate the autarky effect if the net effect is to bring people to area A, or increase it if the net effect is to send people to other subsidized areas (i.e. area A gains from the control area B, but loses to other subsidized areas). Second, the contamination effect is similar as before but now it is increased compared to the previous example when only one region receives the subsidy. The reason is that the unsubsidized area now changes more because people are leaving to all the subsidized areas. Finally, now the relative size of the effects cannot be computed anymore knowing only the supply and demand elasticities, since these objects are weighted by the original changes of prices induced by the policy. In this sense, the elasticities are not anymore a "sufficient statistic" for the relative size of the effects.

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 OZ tax breaks, which might be directed to commercial or residential development, LVIS tax breaks 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, especially when compared to US's Opportunity Zones. González-Pampillón (2022) estimates that LVIS tax benefits equaled 20% of the cost of the projects. The main component of those tax benefits is the total exemption from the country's corporate tax of 25%. Beyond this main component, LVIS projects were fully exempted from the value added tax on inputs, and units devoted to the rental market were partially exempted from the income and wealth taxes. Because these tax breaks were so large, 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. We thus 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 it less attractive to investors. Because those modifications would substantially change the impact of the policy on housing prices, we end our period of analysis in May 2014.

We focus our analysis on the impact of LVIS tax breaks in the department of Montevideo, which holds the homonymous 1.3 million capital city of Uruguay and concentrated 70% of LVIS projects (Berrutti, 2017). LVIS in Montevideo subsidized residential development in medium and low-income neighborhoods. The upper half of Figure 3 presents a map of the subsidized and unsubsidized areas in 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 Montevideo's urban area, including the central and older areas of the city as well as working-class neighborhoods.

Due to the generosity of its tax breaks, 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 of the policy. Another measure of the huge quantitative relevance of the policy is the total amount of investments benefited by LVIS tax cuts. González-Pampillón (2022) estimates that the total investment approved during the first five years of the law amounts 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 phase of construction, the first few LVIS projects only reached completion by 2013, and the first sales of LVIS properties occurred in 2014, with most sales being made in the following years (González-Pampillón, 2022). This timing is important because it implies that almost no LVIS projects and very few LVIS sales were completed during the period we study. This has two key implications for the context of our study. First, it motivates us to abstract from the positive externalities of LVIS projects documented by González-Pampillón (2022) and Borraz et al. (2021) for later years.[1] Being able to disregard positive spillovers of the policy on housing prices simplifies our analysis, and further reinforces our hypothesis of a negative effect of the policy on the housing prices of subsidized areas.[2] Second, since almost no LVIS projects and units were completed during our period of study, there is no contemporaneous shift in the number of units for sale. Thus the expected negative effect of the policy on prices fully operates through the capitalization of future lower construction costs into current housing prices. This hypothesis, which we robustly verify in the data, is supported by three elements. First, the generosity of the subsidy, estimated on 20% the cost. Second, the large relative size of the targeted area in terms of Montevideo's housing market. Third, that although there were no units built during our period of analysis, the number of applications, which are publicly available online, signaled agents that supply in the targeted neighborhoods was indeed going to expand substantially in the next few years.

The public data on developers' applications to obtain LVIS subsidies further 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 projects and housing units, respectively. This scenario of atomistic suppliers constitutes a fourth reason explaining the negative effect of the policy on prices and further motivates the perfectly-competitive assumption for the supply side in our model.

3.2 | Data

We use four sources of data. The most important one is the universe of housing transactions from the National Registry Office in Uruguay for the period 2010-2014. These data includes 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. So this database of registered housing transactions is representative of the highly formal housing market of Montevideo.

The transaction data further includes a unique property number, which allows us to match that database with 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 does not exist for the years we analyze, and thus we use the earliest dataset available, which corresponds to 2016. We drop the top and bottom percentile of the area and price distribution of the transaction dataset to avoid extreme values from affecting our estimates.

[1] These positive externalities have also been documented for LIHTC projects in the US (Baum-Snow and Marion, 2009; Diamond et al., 2018). Over time, the resorting of heterogeneous agents can significantly change urban amenities across the urban space Almagro and Domınguez-Iino (2019).

[2] Housing prices reflect future rents, and thus, in principle, these future rents could be positively affected by the spillovers of the new projects. However, González-Pampillón (2022) shows that spillover of new LVIS projects are highly localized and disappear after 200 meters. Based on this evidence, we argue that during our period of analysis, when few projects were constructed, it would be very hard for agents to anticipate the location and impact of future projects.

The third source of data is a geo-coded map of the areas subsidized by LVIS, similar to Figure 3. This geospatial data 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 national population census of 2010. These data provides census tracts' average years of education, which we use to define neighborhoods and nests, as we explain in the next subsection.

Table 1 presents summary statistics on the housing transaction data separately for subsidized and unsubsidized areas and before and after the introduction of the policy. The prices per squared meter and the sizes of houses are lower in the subsidized than in the unsubsidized areas. This is consistent with the policy subsidizing lagging areas in the city. Housing prices grow over time in all areas because our years of study coincide with a period of strong economic growth in Uruguay.

	Pre		Post	
	Subsidized	Unsubsidized	Subsidized	Unsubsidized
No. Obs.	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)

Note: Standard deviations are provided in parentheses.

TABLE 1 Housing prices and area by subsidy status in the pre and post periods

In numerous empirical exercises in this paper, we use a set of variables to control for housing characteristics. These control variables are obtained from the cadaster data except for distance to the coast, which we computed using the exact location of the transaction. The set of controls 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.

3.3 | Neighborhood definition

We complement our general decomposition formula with the estimation and computation of counterfactual equilibria of a specific model of the supply and demand of housing in Montevideo. This specific modeling of the equilibrium impacts of the tax break on housing prices follows a long tradition 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 partition of the space of the city into exclusive units. Because Montevideo is not divided in meaningful smaller administrative units, we partition the city into contiguous and homogeneous units using a spatial clustering algorithm. Throughout the paper we refer to the resulting units as neighborhoods.

We use the SKATER (Spatial 'K'luster Analysis by Tree Edge Removal) algorithm, which was developed by Martins et al. (2006) and has four convenient features for the problem at hand. Differently from regular, non-spatial clustering techniques, this algorithm guaranties 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 provide the empirical variation we need to estimate 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 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 subsidized and unsubsidized areas such that the whole area of each neighborhood falls in only one of those two categories. We indicate the algorithm to use tracts' average number of years of education from the 2010 population census to maximize units' homogeneity. We set a minimum of 10 transactions for the average number of monthly sales that neighborhoods should have and a target number of 50 neighborhoods on each area.

The spatial clustering algorithm gives us a total of 49 neighborhoods, 30 subsidized and 19 unsubsidized. In order to introduce richer substitution patterns in our discrete choice model, we further classify those 49 neighborhoods into the six nests of a nested logit model. This classification uses the same algorithm employed to define the neighborhoods except we do not require spatial contiguity for the resulting units and remove the constraint on the minimum number of monthly sales. Importantly, we allow the algorithm to form nest joining subsidized and unsubsidized neighborhoods freely. The results of this operation are presented in the lower half of Figure 3. Each of the six colors in that figure represent a different nest, the solid line represents the border of the policy, and the lighter lines show the borders of the neighborhoods.

4 | DIFFERENCE-IN-DIFFERENCES RESULTS

This section presents three sets of difference-in-differences (DiD) estimates of the effect of the policy. Consistent with our hypothesis on the subsidy having a negative effect on the prices of the targeted areas, the estimates in the three sets are consistently negative. However, their magnitude varies greatly depending on which units are included in the treatment and control groups. While some estimates imply large price reductions suggesting a highly beneficial impact of the tax break on consumers, others do no reject a zero impact, which would be consistent with landlords fully appropriating the subsidy. This could be a simple matter of heterogeneity of treatment effects. However, we show that our results are consistent with contamination being behind part of that variation. Following the framework in section Section 2, the presence of contamination introduces the possibility of bias in the estimates.

4.1 | Benchmark Difference-in-Differences

The general specification for our difference-in-differences regressions is the usual given by the following equation:

$$p_{ijt} = \gamma_j + \alpha_t + \beta \mathrm{Treat}_j \times \mathrm{Post}_t + f(X_{ijt}) + \epsilon_{ijt} \tag{5}$$

with p_{ijt} denoting the price per square meter of housing transaction i in neighborhood j at month t. Because each neighborhood is entirely treated or untreated, γ_j subsumes the usual Treat_j term. X_{ijt} is a vector of housing characteristics.

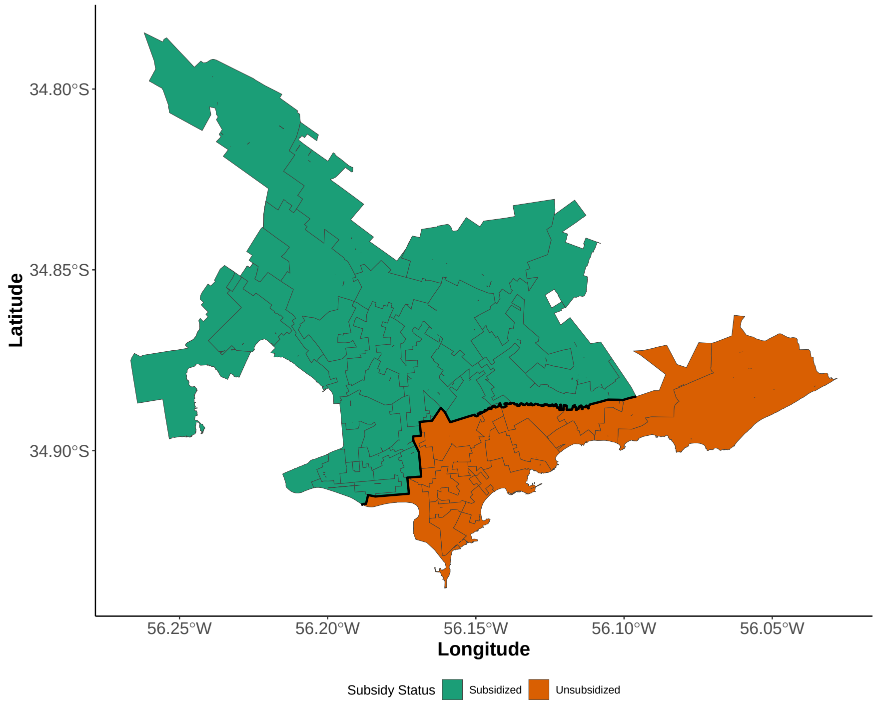

(a) Neighborhoods and subsidy status

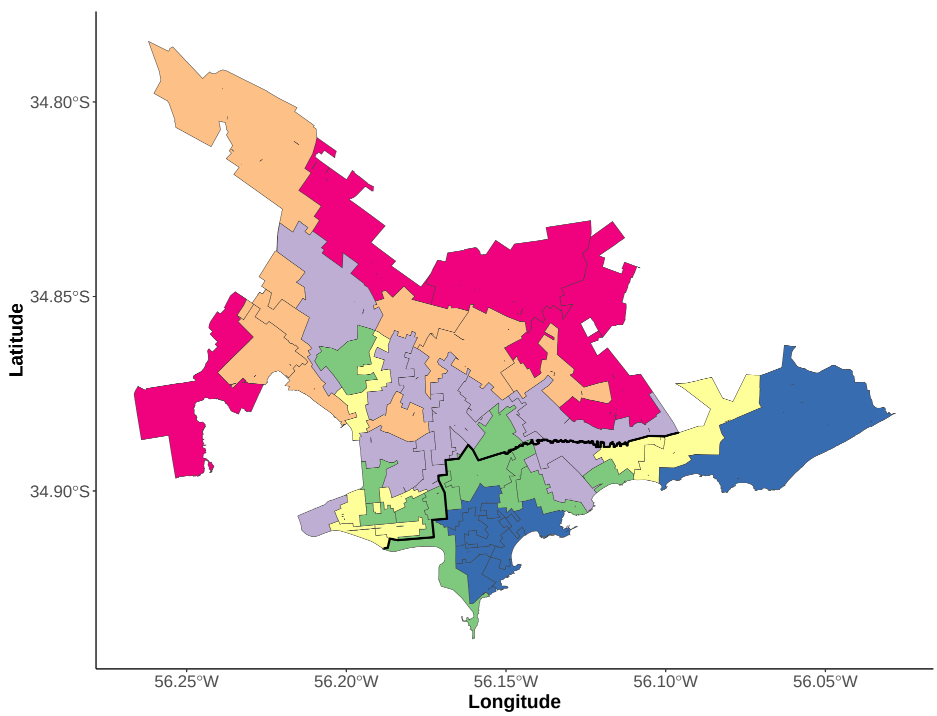

(b) Neighborhoods and nest classification

FIGURE 3 Urban Montevideo by Subsidy Status, Neighborhoods and Regions

	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

$^{*}\ldots p < 0.05$ $^{**}\ldots p < 0.01$ $^{***}\ldots p < 0.001$

Note: Standard errors are clustered at the neighborhood level.

Note: Polynomial of degree three used to control for housing characteristics.

TABLE 2 Difference-in-Differences Regressions

Columns 1 to 3 of Table 2 presents our first set of DiD estimates. The defining feature of this first set is that it considers all transactions in the city and using the basic DiD specification. Column 1 only has the three traditional DiD terms, namely those indicating treatment group, treatment timing, and the interaction of these two. The second column adds a third-order polynomial on the housing characteristics described in Subsection 3.2. These include built area, distance to the coast, construction year, and four variables measuring construction quality. The last column adds month-year and neighborhood fixed effects. These DiD estimates presented in Table 2 are complemented with graphical evidence in Figure A.2 and Figure A.5 in the Appendix.

The three first columns in Table 2 yield consistently negative estimates with a stable magnitude across the different specifications. This result is further confirmed graphically in Figure A.2 and Figure A.4 in Appendix A, which also show evidence of parallel pre-trends 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 | Difference-in-Differences with homogeneous units

A second set of estimates consists of implementing frequently used techniques that maximize the comparability between treated and control areas to minimize concerns about unobserved confounders (Baum-Snow and 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 resort more easily across similar areas, thus 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. All the estimates in this subsection are much lower than the ones in the previous section. This is consistent with these techniques introducing some 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 across the border of the policy (Neumark and Kolko, 2010; Chen et al., 2022). Estimates in Columns 4 to 6 of Table 2 follow this approach and compare the evolution of prices across subsidized and unsubsidized areas within a 500-meter buffer around the border. Figure A.1 in the Appendix provides a map of this buffer and Figure A.3 and Figure A.5 present the usual DiD graphs. The pre-policy price levels across both sides of the border in Figure A.3 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 the common border DiD design in this context would not be able to discard that the tax break had a null effect on the prices faced by consumers.

	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

TABLE 3 Difference-in-Differences Regressions - Extensions

Table 3 introduces additional techniques enhancing comparability between control and treated areas. The first row features DiD with propensity-score reweighting (A. Smith and E. Todd, 2005; Aker, 2010; Wang, 2013; Chen et al., 2022), the second implements a border regression discontinuity (Holmes, 1998; Black, 1999; Bayer et al., 2007; Turner et al., 2014), and the third one estimates a differences-in-discontinuities design (Grembi et al., 2016; Butts, 2023). 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 due to higher resorting between more homogeneous units.

4.3 | Difference-in-Differences with heterogeneous units

We show in two ways that DiD estimates of the impact of the tax break grow in absolute value as the heterogeneity between subsidized and unsubsidized areas increases. First, in Equation 5 we explicitly introduce heterogeneity by interacting the DiD term in the border specification with an index of price differences between both sides of the border. Figure A.6 in Appendix A illustrates how we compute this index. We start by defining a large number of equidistant points along the main border of the policy. Then draw a 500-meter circle around each of those points and compute the difference in the median price per-square meter between the transactions that are contained in that circle but are in opposite sides of the border (left panel of Figure A.6). As a result, each of the points along the border has a scalar value characterizing the heterogeneity in prices across the border at that point. The final step consists of attaching, to each housing transaction, a weighted average of those scalars, for which the transaction property lies within the respective 500-meter circles. The respective weights are the inverse of the distance between the transaction property and the applicable border points. We standardize the resulting heterogeneity index by subtracting its average and dividing it by its standard deviation.

The second column of Table A.1 in Appendix A presents the estimate on the interaction between the DiD term and the heterogeneity index. One standard deviation increase in border heterogeneity increases the absolute value 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 500m buffer. Figure A.7 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 values of the bottom half of the distribution of the heterogeneity index.

The second way in which we evaluate the relationship between DiD estimates and the heterogeneity between treatment and control consists in using the popular "ring approach" (Di Tella and Schargrodsky, 2004; Kline and Moretti, 2014a; Butts, 2022; Myers and Lanahan, 2022). This consists of using controls units which are further away from the border. If heterogeneity between treatment and control 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.

In certain contexts this ring approach can identify the true effect of the policy on treated areas as long as spillovers are zero after a certain distance from the border (Clarke, 2017; Butts, 2022; Myers and Lanahan, 2022). There is evidence that those distances can be quite large in some contexts. Clarke (2017) finds that externalities reach at least 30km away from the border of the policy and Myers and Lanahan (2022) establish the range with no externality as beyond the 40th or 60th percentile of their distribution of technological distance across firms and inventors. This requirement of no spillovers after a certain distance may not hold in other contexts because of two difficulties, which are present in our study and thus prevent us from recovering the true effect of the policy using the ring DiD estimates in Table 3. First, natural (sea, mountains) or human-made (park, highway) constraints may limit the distance after which one can define the control group. In our context, we study a coastal city, and the sea restricts the distance of the rings we can built. For instance, only 10% of our unsubsidized transactions are beyond 2,100 meters. This restriction imposed by the sea is clear in Figure A.8 in Appendix A. Second, as noted by Butts (2022), when policies are large-enough to induce sorting of agents across the whole city, spillover-free areas may well not exist.

The three sets of results in this section show that DiD estimates of the price effects of a place-based policy increase with the heterogeneity between the subsidized and unsubsidized areas chosen for comparison. This pattern is consistent with the framework introduced in Section 2. There should be less resorting in reaction to lower prices when subsidized and unsubsidized units are very heterogeneous, thus minimizing the contamination effect. Importantly, the framework emphasizes that this is not a problem of heterogeneous treatment effects but of biased estimates. We next complement these findings by solving for a specific estimated model that allows us to separately measure the contamination effect. Consistent with the reduced-form evidence in this section, we show that contamination does indeed correlate positively with both the degree of heterogeneity 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 | STRUCTURAL MODEL

In this section we introduce a specific model of real state transactions in a single city. The model is static and housing is homogeneous in quality and has a different price depending on the neighborhood where it is located. The demand-side of the model consists of a discrete-choice framework with households choosing the neighborhood in which they want to buy the generic unit of housing. The supply-side of the model consists of an upward sloping, log-linear housing supply for each neighborhood.

5.1 | Demand

Households make a discrete and exclusive choice regarding the neighborhood in which they are buying a house in Montevideo. This discrete set of geographical areas is complemented by an outside option consisting in buying a house in the localities belonging to the broader metro area of Montevideo. Potential buyers of a house compare the utility of their options using Equation 6, and choose the option that yields the highest indirect utility.

$$V_{ijt} = V(AM_{jt}, P_{jt}, \tilde{\epsilon}_{ijt}) \tag{6}$$

The first argument of the indirect utility function are the neighbourhood amenities 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 neighbourhood j at time t. $\tilde{\epsilon}_{ijt}$ denotes the unobserved preferences of consumer i at time t for neighborhood j.

We parameterize the indirect utility function with the following linear function:

$$V(AM_{jt}, P_{jt}, \epsilon_{ijt}) = A_j + B_t + \xi_{jt} - \alpha P_{jt} + \tilde{\epsilon}_{ijt} = \delta_{jt} + \tilde{\epsilon}_{ijt} \tag{7}$$

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 define $\tilde{\epsilon} = \zeta_{int} + (1-\sigma) \times \epsilon_{ijt}$, where σ with $0 < \sigma \leqslant 1$ is the nesting parameter. ζ_{int} is common to all products in nest n. We assume $\zeta_{int} + (1-\sigma) \times \epsilon_{ijt}$ follows an extreme value Type-1 distribution. Note that as σ approaches one, the within nest correlation of utility levels goes to 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. $\delta_{0t} = 0 \; \forall \; t$). As in (Berry, 1994), this structure yields a linear equation with which one can estimate the whole demand system. This is Equation 8, where s_j is the market share of area j in the whole market at time t and $\bar{s}_{j,n,t}$ is the market share of product j in nest n in period t.

$$\ln(s_{j,t}) - \ln(s_{0,t}) = \delta_{jt} = A_j + B_t + \xi_{jt} - \alpha P_{jt} + \sigma \ln(\bar{s}_{j,n,t}) \tag{8}$$

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:

$$P_{jt} = (1 - \tau_{jt}) * MC(Q_{jt}). \tag{9}$$

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 shocks to construction costs.

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

$$MC(Q_{jt}) = L_{jt} \times Q_{jt}^{\eta} \tag{10}$$

Applying logarithms to both sides of Equation 10, and combining the resulting expression with Equation 9 yields the inverse housing supply curve:

$$\ln P_{jt} = \ln L_{jt} + \ln(1 - \tau_{jt}) + \eta \ln Q_{jt} \tag{11}$$

5.3 | Parallel trends and contamination in this 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 differences-in-differences, we must check that it can satisfy parallel trends. We evaluate this by simulating a series of equilibria of the model with alternative parameters. We present the detail of those simulations in Appendix C and here summarize the two main conclusions we extract from that exercise.

The first conclusion is that our model allows for parallel trends despite being highly non-linear in both its supply and demand side. The second conclusion is that increasing the variance of the simulated parameters corresponding to neighborhood amenities (ξ_{jt}) 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 idiosyncratic changes in housing prices over time which reject 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 making neighborhoods very different imply lower degrees of resorting in reaction to the treatment, which means less contamination and bias of the DiD estimate.

6 | ESTIMATION

6.1 | Demand

We estimate our demand model on a dataset that has a single quantity and price for the each combination of neighborhood and month-year. Quantities are the number of transactions in each cell. In order to control for differences in housing quality across neighborhoods, prices are 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 the set of housing characteristics described in Section 3. Those characteristics include properties' age, area in squared meters, distance to the coast, and four variables from the cadaster describing construction quality.

The demand regressions estimate Equation 8. In these regressions, the A_j and B_t amenity terms are captured by neighborhood and time fixed effects and the time-varying amenities ξ_{jt} constitute the structural error. Since equilibrium prices and within-nest shares are correlated with amenities, OLS estimates in Table 4 are likely biased. We address this endogeneity by leverage 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 is benefited by 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 squared meters, and the share of that area in the total area of the nest.

Since time-varying neighborhood amenities are the structural error of our IV regressions, the identification assumption behind our set of instruments is that the tax break did not impact those amenities conditional on our set of fixed effects. This assumption deserves special attention given the abundant evidence on the impacts of new housing supply on neighborhood amenities (Baum-Snow and Marion, 2009; Rossi-Hansberg et al., 2010; Diamond et al., 2018), including evidence for the program we are studying (González-Pampillón, 2022; Borraz et al., 2021). Two elements from our context justify this assumption. First, as discussed in Section 3, we study a period in which the housing projects benefited by the subsidy had not yet been completed. Second, although in principle housing prices in our period could incorporate the future effect of new construction projects on improved amenities, this anticipation is limited because previous evidence shows that the price impacts of new housing projects are highly localized (González-Pampillón, 2022). It would be thus very hard for market participants to anticipate where these projects were going to be located and thus capitalize the resulting future amenities.

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); 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 share $s_{j|n}$ obtained in the second step. Note that these last two instruments are obtained in an equilibrium in which 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 neighborhoods 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 paragraphs. The final estimates, which are the ones considered 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.

TABLE 4 Demand Estimation

	$\ln(s_{jt}) - \ln(s_{0t})$			
	(1)	(2)	(3)	(4)
Price per 100 Square Meters (α)	0.02***	-0.00	-0.02***	-0.07***
	(0.00)	(0.00)	(0.01)	(0.01)
Within-Nest Log Market Share (σ)	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
Geography FE	-	Neighborhood	Neighborhood	Neighborhood
Time FE	-	Year x Month	Year x Month	Year x Month
K-P 1st stage F			0.71	21.6

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

6.2 | Supply

We calibrate the two parameters of Equation 11, one of them externally and the other internally. We calibrate τ_{jt} externally using González-Pampillón (2022)'s estimate on the LVIS subsidy representing 20% of the final housing price. We internally calibrate the inverse housing supply elasticity η with a matching procedure that mirrors the one applied by Berger et al. (2022) in their study of market power in the US labor market. We set η such that the DiD term in the structural equilibrium exactly matches its reduced-form counterpart of -181 from Section 4. In the next section we provide further details on the equilibrium computation and, in particular, how we obtain a that structural DiD term.

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 of 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 treatment and control units and positively correlated with diversion ratios.

7.1 | DiD decomposition and subsidy's incidence

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 neighborhoods' amenities and marginal costs, which we obtain as the residuals from the housing demand and supply equations, respectively. We focus our equilibrium comparisons in the period after the subsidy was introduced and evaluate the counterfactual equilibrium prices and quantities when the subsidy is set to zero. Analogously to the reduced-form DiD, the structural DiD is double difference between prices in the subsidized and unsubsidized areas with and without the subsidy.

Table 5 presents the decomposition and incidence results for the whole city. The second column has the results using 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 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 2. The ATT term is the difference in average equilibrium prices of the subsidized neighborhoods with and without the subsidy. Following our decomposition formula, the ATT is, in turn, the sum of the autarky and resorting terms. The autarky term is the change in the average equilibrium prices across subsidized areas due to the introduction of the subsidy but without allowing for resorting between neighborhoods. We then calculate the resorting term as the difference between the ATT and the autarky. This resorting effect in Table 5 is large, indicating that the reduction in housing prices in the subsidized areas would have been much higher if buyers had not reacted to the policy by resorting into these areas.

The contamination term is the most important one since it measures the difference between the DiD term and the ATT. This term can be thought of as the structural counterpart of the bias that the reduced-form estimate has as a measure of the impact of the policy. We obtain the contamination term as the difference in the average equilibrium prices of unsubsidized areas with and without the subsidy. The existence of a 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 areas.

TABLE 5 Decomposition of DiD Results Using Structural Model

	Reduced-Form	Structural
DiD	−181	−181
Did Decomposition:		
ATT		−242
Autarky		−404
Sorting		162
Contamination		−61
Contamination/ATT		25.2%
Incidence	55.5%	79.1%

The last row of Table 5 shows that the existence of substantial contamination has large implications in terms of the conclusions on the incidence of the policy one would get following either the reduced-form DiD (first column) or the structural ATT (second column). We calculate the incidence as the estimated effect on prices as a proportion of the subsidy.

While the true incidence is 79%, the one calculated using the reduced-form DiD would have been 24 percentage points lower. To obtain the amount of the subsidy for the structural incidence we apply the 20% rate over the price of the subsidized area that one would obtain by evaluating the equilibrium quantities on the unsubsidized inverse supply curve. For the reduced-form incidence we compute without the estimated model, thus mirroring what we would do as applied researchers without this model. In order to obtain an "unsubsidized equilibrium price for subsidized areas", to which we can apply the 20% rate, we add the observed average difference between the pre and post period in the prices of unsubsidized areas to the average price of the subsidized areas in the pre period.

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 have had an incidence of a 100%, and then all subsidy was translated to consumers, they would have saved 18,000 dollars. However, tax breaks typically are not entirely reflected on prices, and it is thus an important economic question to establish which share of the tax break reaches its potential beneficiaries. In our context, a researcher guided by the reduced form estimate of the incidence (55.5%) would have concluded that our consumer saved around 9,990 USD. However, once accounted for contamination, the incidence of 79.1% implies a saving of 14,238 USD. The difference between both estimates of the incidence is 2,934 dollars, which amounts to 29.8% of Uruguay's GDP per capita in the year the policy was introduced.

7.2 | Determinants of Contamination and Bias

The previous analysis showed that contamination can lead to wrong conclusions on the effect of a placed-based policy. In order 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 sorting, which in turn correlates with the similarity between control and treatment 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 placed-based policies may induce substantive resorting of agents between treatment and control areas.

Figure 4 presents the first piece of evidence on the positive correlation between contamination and demand-side sorting. It does so by plotting structural contamination as a share of ATT and the heterogeneity index introduced in Section 4 for every pair of subsidized and non-subsidized neighborhoods lying along the border of the policy. Going back to the reduced-form relationship between the border DiD estimate and the degree of heterogeneity across the border in Table 3 and Figure A.7, the results in Figure 4 indicate that contamination can explain why one may not discard that the policy had zero effects when comparing very homogeneous areas.

The second piece of evidence, presented in Figure 5, 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 numbers of pairs to estimate the regression equivalent of Figure 5 including a rich set of controls. Table A.2 in Appendix A shows robust and positive regression coefficients when controlling for none, either and both neighborhood and month $\times$ year fixed-effects.

Finally, our formula states that not only contamination but also ATT is correlated with the intensity of demand-side substitution. Since the resorting term is part of the ATT together

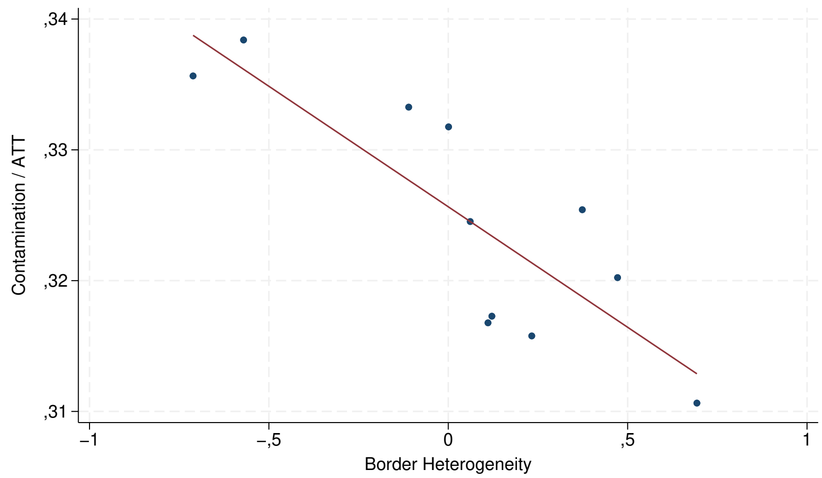

FIGURE 4 Contamination and border heterogeneity

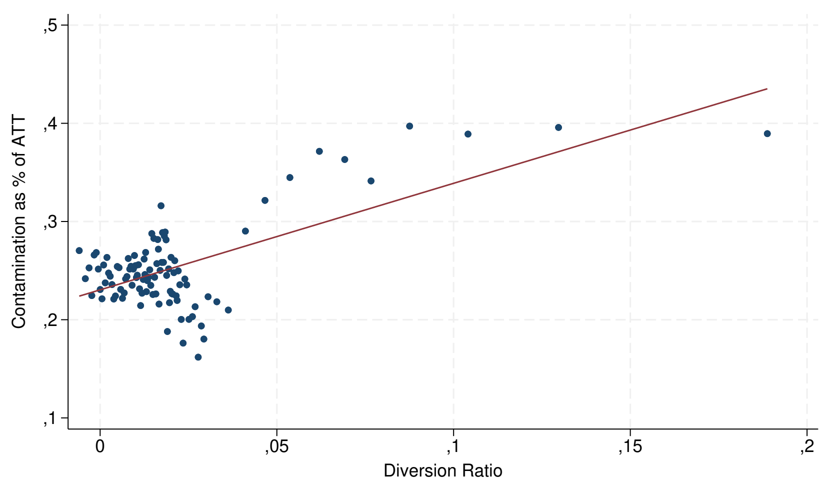

FIGURE 5 Contamination and diversion ratios

with autarky term, more of it would lead to lower DiD estimates of the impact of the subsidy. Similarly to Figure 4 above, Figure A.10 in Appendix A shows that the absolute value of the ATT effectively increases with the heterogeneity between 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 resorting, applied researchers focusing on very homogeneous areas would get systematically lower estimates. Second, and more substantive, the identification of substantive resorting affecting the ATT can be normatively relevant, since the higher prices caused by resorting may offset part of the benefits of the subsidy for incumbent households.

8 | CONCLUSION

The non-random assignment of place-based policies implies that their study requires the use of quasi-experimental methods, with difference-in-differences (DiD) being one of the most important. In this paper, we provide a framework to analyze when difference-in-differences estimates may or may not recover the effect of the policy. In contexts where place-based policies are large enough to affect non-targeted areas, reduced-form methods may not recover the actual effect of the policy because of SUTVA violations. We provide a structural framework to recover - in those contexts - the effects on quantities, prices, and welfare.

We illustrate the potential of our framework by analyzing a large tax break for housing development in lagging areas in Montevideo. We show that reduced-form difference-in-differences vary greatly depending on the spatial range of included treatment and control units. This variation, in turn, follows the pattern predicted by our framework: When the control and treated groups are more similar, the effect of the tax break on prices is lower. According to our framework, these heterogeneous results are not necessarily capturing an underlying heterogeneity in the effects of the policy but partly reflect a heterogeneity in the degree of demand spillovers and reference-group contamination across the different estimates.

We then present a new formula that shows that when a placed-based policy triggers relocation from "non-treated" into "treated" areas, the difference-in-differences estimator can be decomposed into three different effects, without being able to separately identify any of them. We then develop and estimate a structural demand and supply model of the market to disentangle these three effects.

In equilibrium counterfactuals we show that the reduced-form difference-in-differences estimates for the whole city substantially underestimate the benefit that consumers obtain from the tax break. The SUTVA violation accounts for about 25% of the total effect of the policy in the subsidized area when considering the whole city. We illustrate how this SUTVA violations can have serious consequences in terms of the welfare impacts of large place-based policies.

Thanks

We would like to thank Milena Almagro, Nathaniel Baum-Snow, Lorenzo Caliendo, Brantly Callaway, Jan De Loecker, Xavier D'Haultfoeuille, Jonathan Dingel, Tomás Domínguez-Iino, Geert Goeyvaerts and Frank Verboven for comments and suggestions. We thank Christian Valencia who provided excellent research assistance.

REFERENCES

A. Smith, J. and E. Todd, P. (2005) Does matching overcome LaLonde's critique of nonexperimental estimators? *Journal of Econometrics*, **125**, 305–353. URL: https://www.sciencedirect.com/science/article/pii/S030440760400082X.

Ahlfeldt, G. M., Redding, S. J., Sturm, D. M. and Wolf, N. (2015) The Economics of Density: Evidence From the Berlin Wall. *Econometrica*, **83**, 2127–2189. URL: https://ideas.repec.org/a/wly/emetrp/v83y2015ip2127-2189.html.

Aker, J. C. (2010) Information from Markets Near and Far: Mobile Phones and Agricultural Markets in Niger. *American Economic Journal: Applied Economics*, **2**, 46–59. URL: https://www.jstor.org/stable/25760219. Publisher: American Economic Association.

Almagro, M., Chyn, E. T. and Stuart, B. A. (2022) Urban Renewal and Inequality: Evidence from Chicago's Public Housing Demolitions. URL: https://papers.ssrn.com/abstract=4317066.

Almagro, M. and 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. URL: https://www.sciencedirect.com/science/article/pii/S0094119021000097.

Anagol, S., Ferreira, F. and Rexer, J. (2021) Estimating the Economic Value of Zoning Reform. URL: https://www.nber.org/papers/w29440.

Baum-Snow, N. and Ferreira, F. (2015) Causal inference in urban and regional economics. In *Handbook of regional and urban economics*, vol. 5, 3–68. Elsevier.

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

Baum-Snow, N. and Marion, J. (2009) The effects of low income housing tax credit developments on neighborhoods. *Journal of Public Economics*, **93**, 654–666. URL: https://ideas.repec.org/a/eee/pubeco/v93y2009i5-6p654-666.html.

Bayer, P., Ferreira, F. and McMillan, R. (2007) A Unified Framework for Measuring Preferences for Schools and Neighborhoods. *Journal of Political Economy*, **115**, 588–638. URL: https://ideas.repec.org/a/ucp/jpolec/v115y2007i4p588-638.html.

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

Berger, D., Herkenhoff, K. and Mongey, S. (2022) Labor Market Power. *American Economic Review*, **112**, 1147–1193. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20191521.

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

Berry, S. (1994) Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 242–262. URL: http://www.jstor.org/stable/2555829.

Black, S. (1999) Do Better Schools Matter? Parental Valuation of Elementary Education. *The Quarterly Journal of Economics*, **114**, 577–599. URL: https://econpapers.repec.org/article/oupqjecon/v_3a114_3ay_3a1999_3ai_3a2_3ap_3a577-599..htm. Publisher: Oxford University Press.

Borraz, F., Carozzi, F., Gonzalez-Pampillon, N. and Zipitria, L. (2021) Local Retail Prices, Product Varieties and Neighborhood Change. *American Economic Journal: Economic Policy*. URL: https://papers.ssrn.com/abstract=3928649.

Busso, M., Gregory, J. and Kline, P. (2013) Assessing the Incidence and Efficiency of a Prominent Place Based Policy. *American Economic Review*, **103**, 897–947. URL: https://www.aeaweb.org/articles?id=10.1257/aer.103.2.897.

Butts, K. (2021) Difference-in-Differences Estimation with Spatial Spillovers.

— (2022) JUE insight: Difference-in-differences with geocoded microdata. *Journal of Urban Economics*, 103493. URL: https://www.sciencedirect.com/science/article/pii/S0094119022000705.

— (2023) Geographic difference-in-discontinuities. *Applied Economics Letters*, **30**, 615–619. URL: https://doi.org/10.1080/13504851.2021.2005236. Publisher: Routledge _eprint: https://doi.org/10.1080/13504851.2021.2005236.

Caliendo, L., Dvorkin, M. and Parro, F. (2019) Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock. *Econometrica*, **87**, 741–835.

Callaway, B. and Sant'Anna, P. (2021) Difference-in-Differences with multiple time periods. *Journal of Econometrics*, **225**, 200–230. URL: https://www.sciencedirect.com/science/article/pii/S0304407620303948.

de Chaisemartin, C. and D'Haultfœuille, X. (2021) Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey. URL: https://papers.ssrn.com/abstract=3980758.

Chen, J., Glaeser, E. and Wessel, D. (2022) JUE Insight: The (non-)effect of opportunity zones on housing prices. *Journal of Urban Economics*, 103451. URL: https://www.sciencedirect.com/science/article/pii/S0094119022000286.

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

Clarke, D. (2017) Estimating Difference-in-Differences in the Presence of Spillovers. *Tech. Rep. 81604*, University Library of Munich, Germany. URL: https://ideas.repec.org/p/pra/mprapa/81604.html.

Delgado, M. S. and Florax, R. J. G. M. (2015) Difference-in-differences techniques for spatial data: Local autocorrelation and spatial interaction. *Economics Letters*, **137**, 123–126. URL: https://www.sciencedirect.com/science/article/pii/S0165176515004371.

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

Diamond, R. (2016) The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000. *American Economic Review*, **106**, 479–524. URL: http://www.aeaweb.org/articles.php?doi=10.1257/aer.20131706.

Diamond, R., McQuade, T. and Qian, F. (2018) Who Benefits from Rent Control? The Equilibrium Consequences of San Francisco's Rent Control Expansion.

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

— (2017) Railroads of the Raj: Estimating the impact of transportation infrastructure. *American Economic Review*. URL: https://scholar.google.com/citations?user=wyvHNI4AAAAJ&hl=ja.

Fajgelbaum, P. D., Morales, E., Suárez Serrato, J. C. and Zidar, O. (2019) State Taxes and Spatial Misallocation. *The Review of Economic Studies*, **86**, 333–376.

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

Feyrer, J., Mansur, E. T. and Sacerdote, B. (2017) Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution. *American Economic Review*, **107**, 1313–1334. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20151326.

González-Pampillón, N. (2022) Spillover effects from new housing supply. *Regional Science and Urban Economics*, **92**, 103759. URL: https://www.sciencedirect.com/science/article/pii/S0166046221001198.

Goodman-Bacon, A. (2021) Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, **225**, 254–277. URL: https://www.sciencedirect.com/science/article/pii/S0304407621001445.

Grembi, V., Nannicini, T. and Troiano, U. (2016) Do Fiscal Rules Matter? *American Economic Journal: Applied Economics*, **8**, 1–30. URL: https://www.aeaweb.org/articles?id=10.1257/app.20150076.

Holmes, T. J. (1998) The Effect of State Policies on the Location of Manufacturing: Evidence from State Borders. *Journal of Political Economy*, **106**, 667–705. URL: https://www.jstor.org/stable/10.1086/250026. Publisher: The University of Chicago Press.

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

James, A. G. and Smith, B. (2020) Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution: Comment. *American Economic Review*, **110**, 1905–1913. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20180888.

Kline, P. and 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**, 275–331. URL: http://qje.oxfordjournals.org/content/129/1/275.

— (2014b) People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs. *Annual Review of Economics*, **6**, 629–662. URL: https://ideas.repec.org/a/anr/reveco/v6y2014p629-662.html.

Martins, A., Neves, M., Câmara, G. and 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.

Monte, F., Redding, S. J. and Rossi-Hansberg, E. (2018) Commuting, Migration, and Local Employment Elasticities. *American Economic Review*, **108**, 3855–3890. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20151507.

Moretti, E. (2011) Local Labor Markets. *Handbook of Labor Economics*, Elsevier. URL: https://ideas.repec.org/h/eee/labchp/5-14.html.

Myers, K. R. and Lanahan, L. (2022) Estimating Spillovers from Publicly Funded R&D: Evidence from the US Department of Energy. *American Economic Review*, **112**, 2393–2423. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20210678.

Neumark, D. and Kolko, J. (2010) Do enterprise zones create jobs? Evidence from California's enterprise zone program. *Journal of Urban Economics*, **68**, 1–19. URL: https://www.sciencedirect.com/science/article/pii/S0094119010000033.

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

Rossi-Hansberg, E., Sarte, P. and Owens, R. (2010) Housing Externalities. *Journal of Political Economy*, **118**, 485–535. URL: https://www.jstor.org/stable/10.1086/653138.

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

Roth, J., Sant'Anna, P. H. C., Bilinski, A. and Poe, J. (2023) What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics*, **235**, 2218–2244. URL: https://www.sciencedirect.com/science/article/pii/S0304407623001318.

Saez, E. (2001) Using Elasticities to Derive Optimal Income Tax Rates. *The Review of Economic Studies*, **68**, 205–229. URL: https://doi.org/10.1111/1467-937X.00166.

Saiz, A. (2010) The Geographic Determinants of Housing Supply. *The Quarterly Journal of Economics*, **125**, 1253–1296. URL: http://qje.oxfordjournals.org/content/125/3/1253.

Serrato, J. and Zidar, O. (2016) Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms. *American Economic Review*, **106**, 2582–2624. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20141702.

Sun, L. and Abraham, S. (2021) Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, **225**, 175–199. URL: https://www.sciencedirect.com/science/article/pii/S030440762030378X.

Turner, M. A., Haughwout, A. and van der Klaauw, W. (2014) Land Use Regulation and Welfare. *Econometrica*, **82**, 1341–1403. URL: http://onlinelibrary.wiley.com/doi/10.3982/ECTA9823/abstract.

Wang, J. (2013) The economic impact of Special Economic Zones: Evidence from Chinese municipalities. *Journal of Development Economics*, **101**, 133–147. URL: https://www.sciencedirect.com/science/article/pii/S0304387812000934.

A | APPENDIX: FIGURES AND TABLES

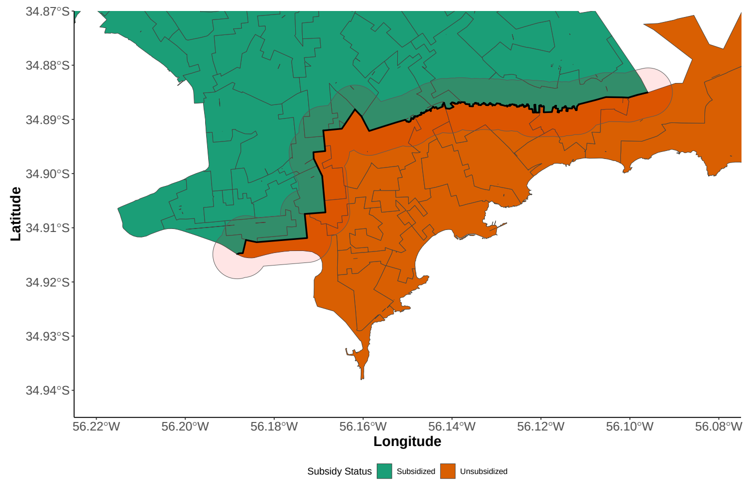

FIGURE A.1 Montevideo by Subsidy Status - 500m Buffer

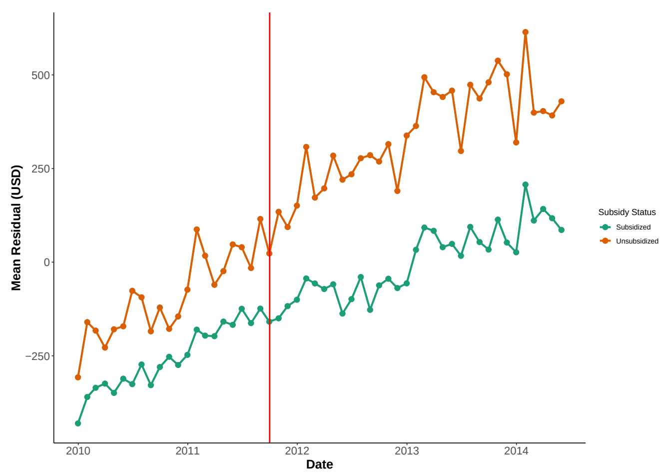

FIGURE A.2 Difference-in-Differences Pre-Trends over Time by Treatment Status

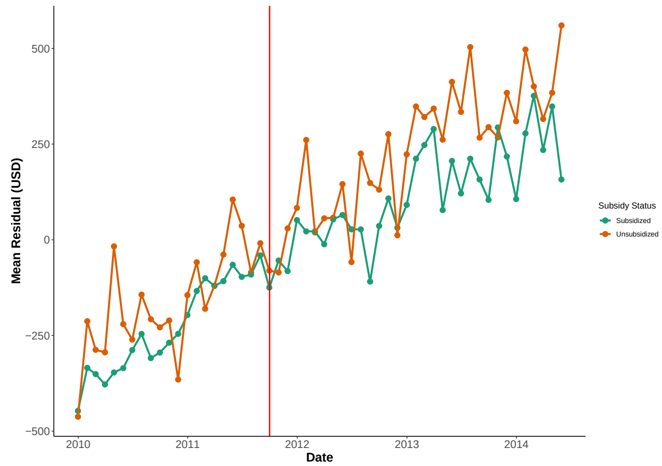

FIGURE A.3 Difference-in-Differences Pre-Trends over Time by Treatment Status - 500m Buffer

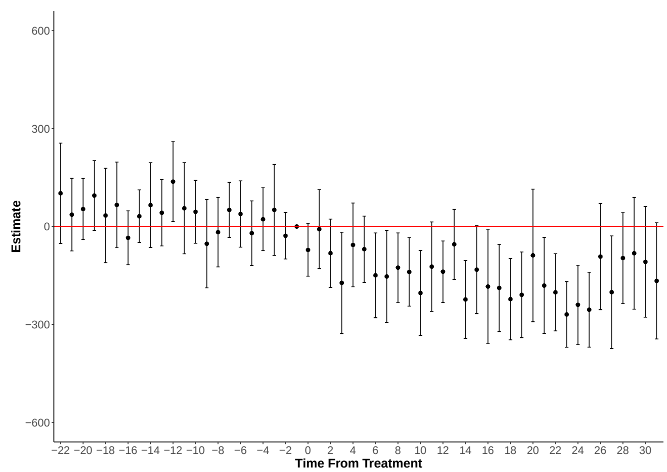

FIGURE A.4 Difference-in-Differences Pre-Trend Estimates - City-Wide

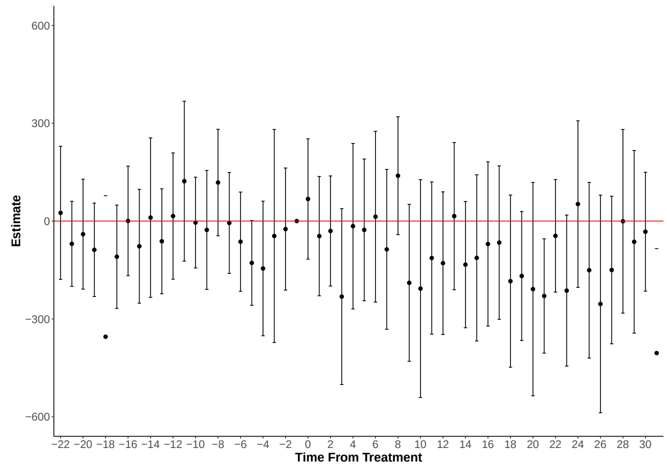

FIGURE A.5 Difference-in-Differences Pre-Trend Estimates - 500m Buffer

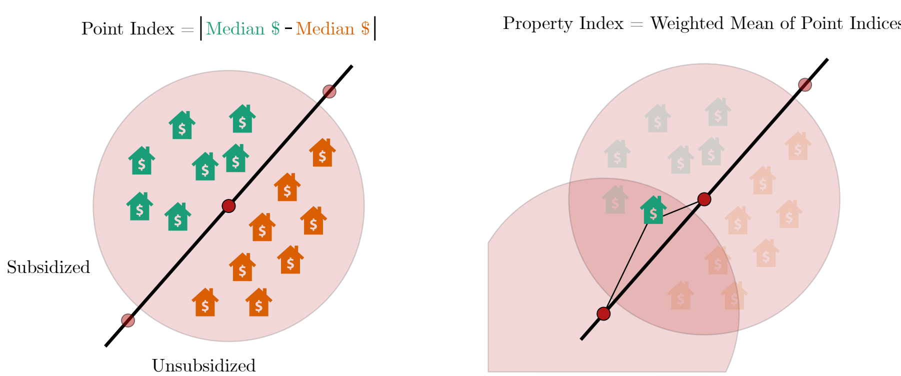

FIGURE A.6 How Border Z-Scores are Computed

	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

* . . . $p < 0.05$ ** . . . $p < 0.01$ *** . . . $p < 0.001$

Note: Standard errors are clustered at the neighborhood level.

Note: Polynomial of degree three used to control for housing characteristics.

TABLE A.1 DiD Regressions - USD per Square Meter with Heterogeneity

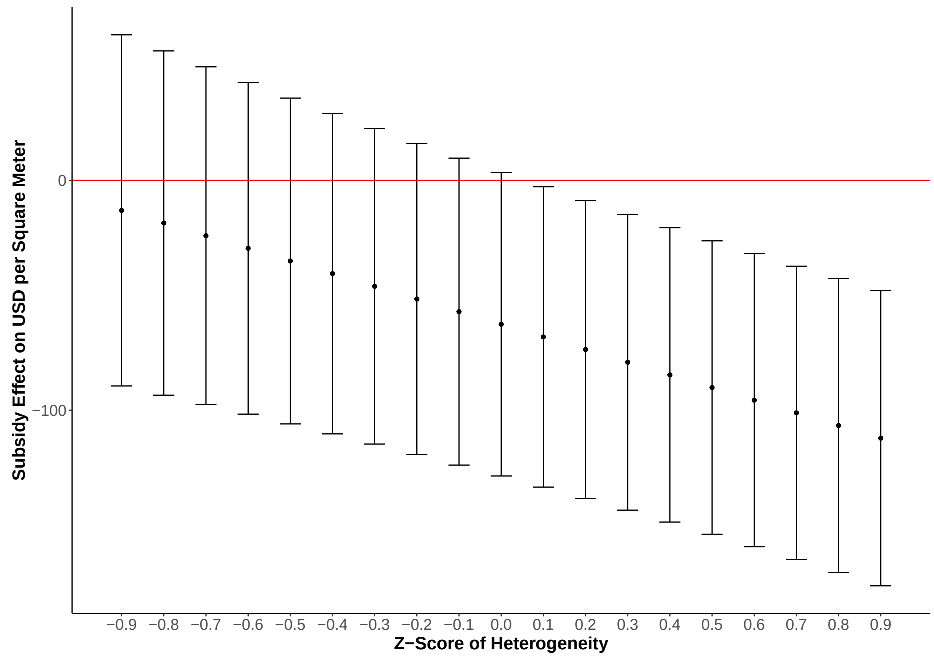

FIGURE A.7 Estimated Treatment Effect as a Function of Heterogeneity

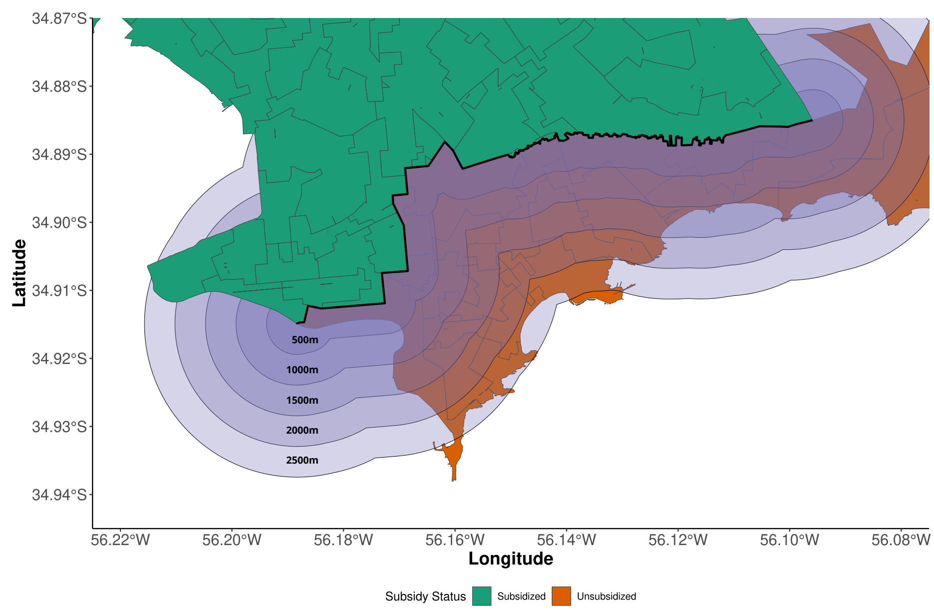

FIGURE A.8 Rings around the border of the policy: unsubsidized area

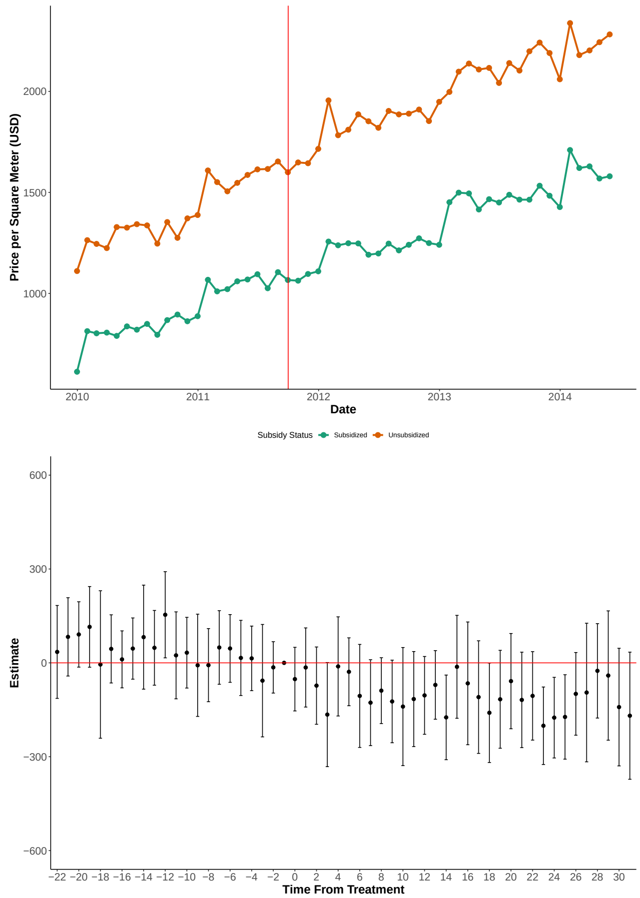

FIGURE A.9 Time Series Plot and Pre-Trends in Structural Model

TABLE A.2 Contamination and Diversion Ratio

	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
Neighborhood FE	No	Yes	No	Yes
Month_Year FE	No	No	Yes	Yes

Robust standard errors in parentheses

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

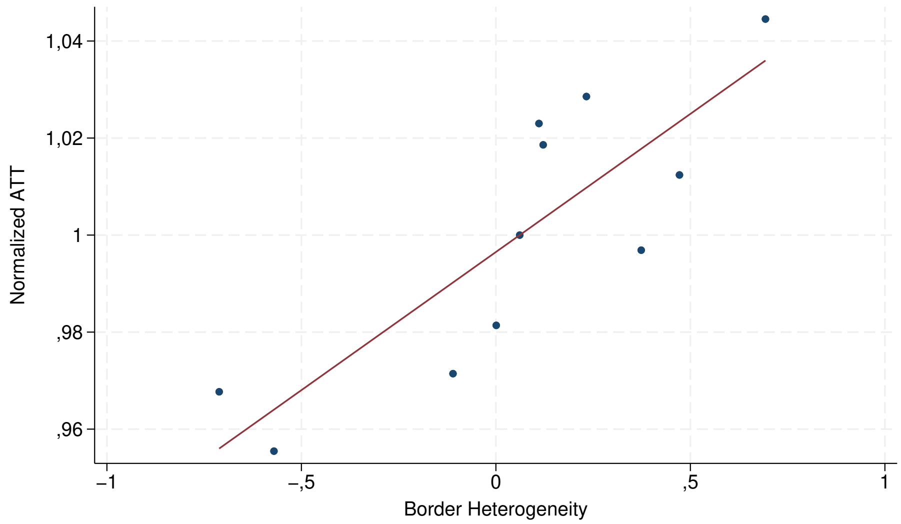

FIGURE A.10 ATT and border heterogeneity

B | APPENDIX: DERIVING THE DID DECOMPOSITION

We specify demand for housing in a particular district d at given prices $\mathbf{p}$ by $D^d(\mathbf{p})$. Supply is specified by $S^d(q^d)$. Please note that this illustration makes use of a linear approximation in both cases.

B.1 | One Subsidized Area

The shift in equilibrium housing quantity in district B in response to the initial policy-induced price change in district A is approximated in the following way:

$$q_3^B - q_2^B = \frac{\partial D^B}{\partial p^A} \times (p_2^A - p_1^A)$$

A similar statement can be made about the equilibrium housing quantity in district A.

$$q_3^A - q_2^A = \frac{\partial D^A}{\partial p^A} \times (p_2^A - p_1^A)$$

Relying on the assumption of full competition within each district, changes in equilibrium housing prices in both districts can be approximated.

$$p_3^A - p_2^A = \frac{\partial S^A}{\partial q^A} \times (q_3^A - q_2^A)$$

$$p_3^B - p_2^B = \frac{\partial S^B}{\partial q^B} \times (q_3^B - q_2^B)$$

Inserting the two earlier equations into the latter equations, we can express second-round equilibrium price changes as a function of demand and supply partial derivatives, as well as the initial policy-induced price change in district A.

$$p_3^A - p_2^A = \frac{\partial S^A}{\partial q^A} \times \frac{\partial D^A}{\partial p^A} \times (p_2^A - p_1^A)$$

$$p_3^B - p_2^B = \frac{\partial S^B}{\partial q^B} \times \frac{\partial D^B}{\partial p^A} \times (p_2^A - p_1^A)$$

Inserting these two expressions into the generalised version of the DiD estimator given in Equation 2, we arrive at Equation 3.

$$\begin{aligned}
\hat{\beta}_{DiD} &= (p_2^A - p_1^A) + (p_3^A - p_2^A) - (p_3^B - p_2^B) \\
&\approx (p_2^A - p_1^A) + \frac{\partial S^A}{\partial q^A} \times \frac{\partial D^A}{\partial p^A} \times (p_2^A - p_1^A) - \frac{\partial S^B}{\partial q^B} \times \frac{\partial D^B}{\partial p^A} \times (p_2^A - p_1^A) \\
&= (p_2^A - p_1^A) \times [1 + \frac{\partial S^A}{\partial q^A} \times \frac{\partial D^A}{\partial p^A} - \frac{\partial S^B}{\partial q^B} \times \frac{\partial D^B}{\partial p^A}] \\
&= (p_2^A - p_1^A) \times [1 + \frac{\partial D^A}{\partial p^A} \times (\frac{\partial S^A}{\partial q^A} - \frac{\partial S^B}{\partial q^B} \times DR_{A,B})]
\end{aligned}$$

with $DR_{A,B}$ being the diversion ratio between housing in district A and housing in district B.

B.2 | Two Subsidized Areas

Using again the notation from Subsection B.1, we now add a second subsidized district C.

$$dD^B = \frac{\partial D^B}{\partial p^A}(p_2^A - p_1^A) + \frac{\partial D^B}{\partial p^C}(p_2^C - p_1^C)$$

$$dD^A = \frac{\partial D^A}{\partial p^B} dP^B + \frac{\partial D^A}{\partial p^C}(p_2^C - p_1^C)$$

$$dD^C = \frac{\partial D^C}{\partial p^A}(p_2^A - p_1^A) + \frac{\partial D^C}{\partial p^B} dP^B$$

Using the supply equation, we can derive an expression for dP^B:

$$\begin{aligned}
dP^B &= \frac{\partial S^B}{\partial q^B} dD^B \\
&= \frac{\partial S^B}{\partial q^B} \times \left[\frac{\partial D^B}{\partial p^A}(p_2^A - p_1^A) + \frac{\partial D^B}{\partial p^C}(p_2^C - p_1^C)\right]
\end{aligned}$$

Using the same approach for the price change in district A, we get the following:

$$
\begin{aligned}
dP^A =& \frac{\partial S^A}{\partial q^A} dD^A \\
=& \frac{\partial S^A}{\partial q^A} \times \left[\frac{\partial D^A}{\partial p^B} dP^B + \frac{\partial D^A}{\partial p^C}(p_2^C - p_1^C)\right] \\
=& \frac{\partial S^A}{\partial q^A} \times \left[\frac{\partial D^A}{\partial p^B} \times \left(\frac{\partial S^B}{\partial q^B} \times \frac{\partial D^B}{\partial p^A} \times (p_2^A - p_1^A) + \frac{\partial S^B}{\partial q^B} \times \frac{\partial D^B}{\partial p^C} \times (p_2^C - p_1^C)\right) + \right. \\
& \left. \frac{\partial D^A}{\partial p^C}(p_2^C - p_1^C)\right]
\end{aligned}
$$

We can now re-write the DiD estimator:

$$
\begin{aligned}
\hat{\beta}_{DiD} =& (p_3^A - p_1^A) - (p_3^B - p_1^B) \\
=& (p_3^A - p_2^A) + (p_2^A - p_1^A) - (p_3^B - p_2^B) \\
\approx& (p_2^A - p_1^A) + \\
& \frac{\partial S^A}{\partial q^A} \times \left[\frac{\partial D^A}{\partial p^B} \times \left(\frac{\partial S^B}{\partial q^B} \times \frac{\partial D^B}{\partial p^A} \times (p_2^A - p_1^A) + \frac{\partial S^B}{\partial q^B} \times \frac{\partial D^B}{\partial p^C} \times (p_2^C - p_1^C)\right) + \right. \\
& \left. \frac{\partial D^A}{\partial p^C}(p_2^C - p_1^C)\right] + \\
& \frac{\partial S^B}{\partial q^B} \times \left[\frac{\partial D^B}{\partial p^A}(p_2^A - p_1^A) + \frac{\partial D^B}{\partial p^C}(p_2^C - p_1^C)\right]
\end{aligned}
$$

Please note:

1. The first term is the autarky effect.
2. The second term is the spillovers effects. In this case, the spillovers effect in equilibrium can be negative or positive. They are going to depend on the two exogenous changes. Spillovers can attenuate the autarky effect if the net effect is to bring people to A, or increase it if the net effect is to send people to C (A gains from B, but loses to C).
3. The third term is the contamination effect, and is similar to before but now it is increased compared to the previous example when only one region receives the subsidy (the reason is that B now changes because people are leaving to A but also because people are leaving to C)

C | APPENDIX: PARALLEL TRENDS WITHIN THE STRUCTURAL MODEL

C.0.1 | Parallel Trends vs. Contamination

Because we plan to use our model to decompose an structural version of our DiD estimate, it is important to first check that parallel trends are satisfied in our model. Roth and Sant'Anna (2023) has shown that functional form are one of the main challenges of parallel trends and our exercise certainly introduces a number of specific functional forms. We check under which conditions our model satisfies parallel trends by simulating a simplified version of the model presented above in which the supply-side is identical and the demand-side abstracts from the nested-logit component and thus features a simple logit demand system.

The endogenous variables of the model are P_{jt} and Q_{jt}, which are determined in equilibrium. For each value of the parameters discussed below, we perform 200 simulations, which are independent of each other. Each simulation features 10 products, out of which five are in the treatment group, and 20 time periods, distributed equally before and after the policy.

Variable	Parameters	
Base Heterogeneity	$\gamma_j \sim N(0, \sigma_j)$	$L_j \sim N(0, \sigma_j)$
Time Heterogeneity	$\gamma_t \sim N(0, \sigma_t)$	$L_t \sim N(0, \sigma_t)$
Idiosyncratic Heterogeneity	$\xi_{jt} \sim N(0, \sigma_{jt})$	$\tilde{\epsilon_{jt}} \sim N(0, \sigma_{jt})$
Inverse Supply Elasticity	$\eta = 0.3$	
Subsidies	$\tau_{jt} = 0.2$ if area j is subsidised in period t, else $\tau_{jt} = 0$.	

TABLE A.3 Simulation Setup - Random Variable Distributions

To understand the properties in terms of parallel trends and contamination, we simplify the setting focusing on three different 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). The distributions from which the exogenous variables are drawn and the other parameters of the simulation are presented in Table A.3.

The exercise allows us to extract three main takeaways. First, our model allows for parallel trends. We simulate the model for a specific set of parameters ($\sigma_j = 0.5, \sigma_t = 0.75, \sigma_{jt} = 0.25$) to show that, despite being very non-linear in both the demand and the supply side, our model can produce parallel trends between treated and untreated areas. The top graph in Figure A.11 suggests the presence of parallel trends in a typical DiD graph while the bottom graph in Figure A.11 presents the typical event study test for parallel trends in the literature.

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 $\{0.5, 1, 2\}$, while the other two variances (σ_t and σ_{jt}) can vary along a grid from 0 to 4 (in 0.25 increments).

The top graph of Figure A.12 presents the results. For each level of σ_j it represents the number of significant coefficients of a variable that indicates the time period interacted with the treatment status (as in the bottom graph in Figure A.11). We find that the violation of parallel trends is more related to higher levels of unobserved heterogeneity (σ_{jt}) than higher levels of time variance (σ_t). Additionally, when the level of baseline heterogeneity (σ_j) is large, then larger time heterogeneity σ_t also compromises the parallel trends.

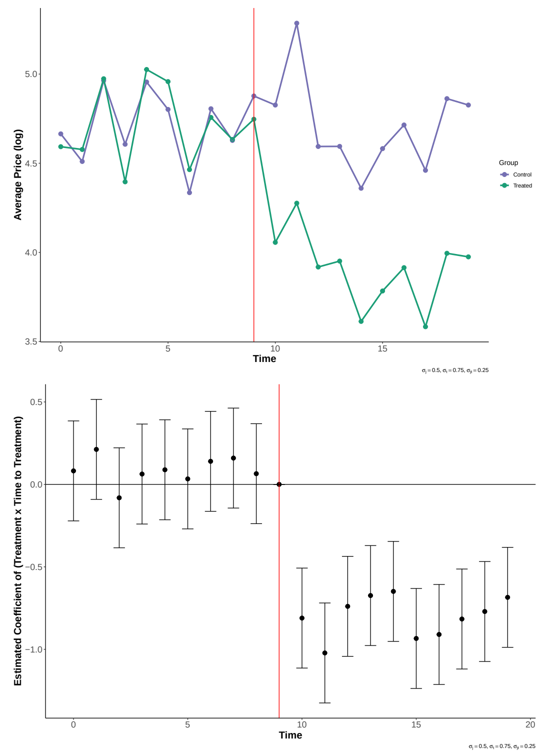

FIGURE A.11 Simulations for an specific draw of parameters ($\sigma_j = 0.5, \sigma_t = 0.75, \sigma_{jt} = 0.25$)

Finally, the third takeaway is that there is a trade-off between parallel trends violations and the contamination effect. The bottom graph of Figure A.12 presents the size of the contamination effect under these simulations. We graph the relative size of the equilibrium price effect on the unsubsidized areas compared to the price on these areas. Contamination is larger when the size of the base heterogeneity (σ_j) is relatively larger compared to the unexplained heterogeneity (σ_{jt}). Intuitively, the more similar the different locations, the higher the extent of reallocation from the control unit to the subsidized units. The opposite is true for the violations of parallel trends. The violation of parallel trends is created because different locations receive different shocks over time, so the mean utility they offer evolves differently over time and so does then the prices.

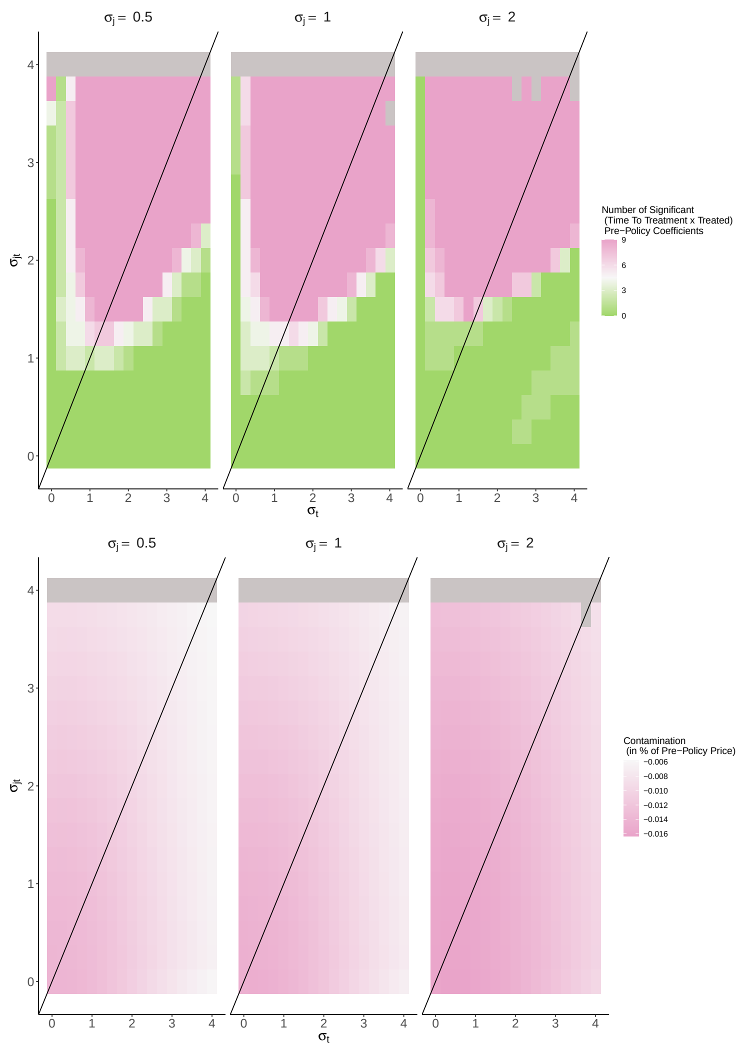

FIGURE A.12 Parallel Trends and Contamination Effects in Simulations