On the geography of vintage-specific restrictions

Carlos Fardella, Nano Barahona, Juan-Pablo Montero and Felipe Sepúlveda*

October 10, 2023

Abstract

Persistent air-pollution problems have led authorities in many cities around the world to impose limits on car use by means of vintage-specific restrictions or low-emission zones. Any vintage restriction must establish not only the cars that face a restriction but also its geographic area of application. As a result of the restriction, a fraction of restricted cars are exported outside the restricted area. Because restricted cars become cheaper, emissions in the restricted area could increase if exported cars remain too close to it. The extent to which such emissions leakage can occur crucially depends on transaction costs in the car market. We study this possibility with a model of the car market that allows for transaction costs and data from Santiago's 2017 vintage restriction. We fail to find emissions leakage, at least severe enough to undo the 2017 policy effects. Interestingly, transaction costs are shown to have a non-monotonic impact on emissions, and hence, on welfare.

Keywords: air quality, driving restrictions, low emission zones
JEL Classification: L62, Q53, R48

*Fardella (cafardella@uc.cl): FK Economics; Barahona (nanobk@berkeley.edu): Department of Economics, University of California, Berkeley; Montero (jmontero@uc.cl): Instituto de Economía, Pontificia Universidad Católica de Chile, Aalto University School of Business, and Sistemas Complejos de Ingeniería (ISCI); Sepúlveda (fasepu@ubc.ca): Department of Economics, University of British Columbia. We thank the Editor (Xavier Labandeira) and two anonymous referees for several comments. Montero also thanks Fondecyt (grant No. 1210252) and Economics for Energy for financial support.

1 Introduction

Persistent air-pollution problems have led authorities in many cities around the world to impose limits on car use by means of vintage-specific driving restrictions or low-emission zones. Unlike the first generation of driving restrictions (e.g., Eskeland and Feyzioglu 1997; Davis 2008; Gallego et al. 2013), which placed uniform restrictions upon all cars regardless of their pollution rates, vintage restrictions impose heavy limits on older, polluting vehicles and lighter or no limits on newer, cleaner ones.

The first vintage restriction was implemented in Santiago in 1992 when it reformed its existing restriction program to exempt all new cars equipped with a catalytic converter (a device that transforms toxic pollutants into less toxic gases) from the then one-day-a-week restriction. A few years later, Mexico City also reformed its existing driving restriction—better known as *Hoy No Circula*—to exempt any new car for their first eight years. Since then, the use of vintage restrictions has extended to many cities in Europe under the name of low-emission zones (LEZs). Sweden initiated the trend in 1996 followed by Germany in 2008. By now LEZs are found in most large cities in Europe, including London, Rome, and Madrid. LEZ programs have also been introduced in China; for example, in Beijing in 2009 and Nanchang in 2013.

Any vintage restriction is defined along three dimensions: (i) its technological extension, i.e., type of cars that are exempt from the restriction (e.g., electric and hybrid vehicles, petrol and diesel cars above certain vintage, etc); (ii) its temporal extension, i.e., hours of the day and days of the week when the restriction applies (e.g., one day a week every week of the year, every day during certain months of the year, etc), and (iii) its geographic extension, i.e., the area where the restriction applies (e.g., center district, entire city, etc).

Perhaps the only difference between the vintage restrictions in Santiago and Mexico City and the LEZ programs implemented in Europe and China pertains to dimension (ii). The former restrictions consider partial circulation bans (e.g., one or two days a week) while LEZs completely ban certain higher-emitting vehicles from entering the restricted zone. In either case one of the goals of these vintage restrictions is to accelerate the fleet turnover toward lower-emitting vehicles.

One aspect of these vintage restrictions that has received less attention in the literature concerns the interaction of dimensions (i), (ii) and (iii) and its connection to transaction costs in the car market. As a result of the restriction, a fraction of restricted cars are exported outside the restricted area. Because restricted cars become cheaper, emissions in the restricted area could increase if exported cars remain too close to it, contributing to its pollution.

This emissions-leakage possibility not only depends on (i), (ii) and (iii) but also on the transaction costs of moving cars across restricted and non-restricted areas. In a frictionless car market, emissions leakage is highly unlikely, if not impossible, because restricted cars would move far from the restricted zone, where they are more valuable (see, e.g., Barahona et al 2020).[1]

[1] In fact, if transactions costs are such that a large fraction of restricted cars remain in the periphery of the

The objective of this paper is to study this emissions-leakage possibility with a model of the car market and evidence from Santiago's 2017 vintage restriction. In 2017 the city of Santiago reformed its 1993 vintage restriction in an important way. It partially terminated the 1993 exemption that was placed upon all cars vintage 1993 and newer (note that by 2017 the 1993 vintage restriction did not play any role since there were virtually no cars vintage 1992 and older in circulation). From 2018 onward, it established a one-day-a-week restriction upon all cars vintage 2011 and older and none upon cars vintage 2012 and newer. Hybrid and electric vehicles, which today account for less than 1% of the fleet, remain exempt regardless of their vintage.

After a brief description of the Santiago's 2017 vintage restriction, our analysis begins in Section 2 with an evaluation of the impact of the restriction on the car market by looking at the evolution of (annual) vehicle circulation permits at the municipality level. We exploit the sharp discontinuity created by the 2017 restriction between restricted and non-restricted vintages to study the impact of the vintage restriction on the evolution of the car market in both restricted (i..e, Santiago) and non-restricted areas. We are particularly interested in understanding the exodus of restricted cars (vintage 2011 and older) from Santiago to nearby and more distant locations

There are three results to emphasize. The first result is that the 2017 vintage restriction had a positive effect on Santiago's fleet turnover toward cleaner cars. We find a sizable fraction of 2011 and older vehicles leaving Santiago because of the reform. Surprisingly, and this is our second result, the effects of the reform were mostly felt over Santiago's low- and middle-income municipalities. We find no evidence of households in high-income municipalities adjusting their purchasing decisions because of the reform.

A plausible explanation is that high-income households tend to own two or more vehicles, so they can more easily accommodate to the restriction without adjusting their purchasing decisions. In contrast, low and middle-income households own one car at most, so to continue driving everyday of the week they have no option but to upgrade their cars to accommodate to the restriction. A similar "no-effect at the top" is in Gallego et al (2013) for the *Hoy No Circula* program in Mexico City.

Our third result points to the presence of transaction costs from moving second-hand cars across restricted and non-restricted areas. We find that the fraction of restricted vehicles flowing to municipalities in Santiago's Periphery is not that different from the fraction of the these same vehicles flowing to more distant municipalities. This pattern can only be explained by the presence of transaction costs. On the one hand, the two non-restricted zones exhibit similar income levels and, on the other, restricted cars are necessarily more valuable in more distant municipalities, where they are never expected to be driven through the restricted zone. In a frictionless market, we would have expected a larger fraction of restricted cars flowing to the more distant locations.

restricted zone, pollution may increase in the restricted area if bus riders in the periphery switch from buses to restricted cars that now become cheaper for them.

The presence of transaction costs is potentially problematic. If a large fraction of the restricted cars leaving Santiago stay in localities near Santiago, this opens up the possibility of emissions leakage, particularly if many trips originated in these localities terminate in the restricted area anyway.

Whether this emissions leakage is severe enough to undo the effect of the reform (and eventually lead to more pollution than otherwise) is something we study in Sections 3 and 4 with the help of a model of the car market.[2] The model is presented in Section 3. It borrows extensively from Barahona et al (2020) except for two innovations: the presence of transaction costs and of a transition zone between restricted and unrestricted zones whose trips also contribute to the pollution problem afflicting the restricted zone.

In Section 4 we apply the model to study pollution and welfare impacts of Santiago's 2017 vintage reform. We study a more demanding restriction than the one actually implemented. In fact, we consider a restriction policy with a moving exemption threshold of 6 years of age, the threshold at the time of policy implementation (the actual policy kept the exemption threshold fixed at vintage 2012). This moving threshold allows us to evaluate policy impacts by simply comparing before and after steady-state outcomes. Besides being simpler to evaluate, the moving-threshold version can only make emissions leakage more likely. Even then, we fail to find evidence of emissions leakage, at least severe enough to undo the effect of the policy. Interestingly, we find transaction costs to have a non-monotonic impact on emissions, and hence, on welfare. Emissions are lowest (and welfare highest) for intermediate levels of transaction costs. According to our model, this is when the entry of new vehicles into the restricted zone reaches its highest level relative to the no-intervention benchmark.

Our work contributes to the expanding literature on driving restrictions (see, e.g., Eskeland and Feyzioglu 1997, Davis 2006, Gallego et al 2013, Viard and Fu 2015, Zhang et al 2017, Salgado and Mitnik 2021) and of vintage restrictions in particular (see, e.g., Wolff 2014, Barahona et al 2020, Galdon-Sanchez et al 2023, Fabra et al 2023). Nevertheless, our work separates from the latter for its focus on transaction costs and the possibility of emissions leakage. Wolf (1994) also study emissions leakage by looking at changes in pollution concentration, which can be problematic since cars are one of the many sources contributing to local pollution. Instead, we study emissions leakage by looking at movements in the car market across different zones.[3]

Our attention to transaction costs is not new either. Adda and Cooper (2000), Gavazza et al (2014), and Blundell et al (2022), for example, look at the intertemporal dimension of transaction costs, that is, at how often second-hand cars change hand.[4] Instead, we focus on its geographical dimension, that is, on how far vehicles go as they change hand.

[2]A necessary condition for leakage to be significant enough to lead to more emissions in the restricted zone is that after the reform many individuals in the periphery of the restricted zone switch from public transport to cars.

[3]There is also work on emissions leakage for stationary sources; see, for example, Martin et al (2014) and Ahlvik and Liski (2022).

[4]See also Liski et al (2023) for the impact of transaction costs on policy design in the car market using evidence from Finland.

The rest of the paper is organized as follows. The empirical analysis is in Section 2. We develop the model in Section 3. Its application to Santiago's 2017 vintage restriction is in Section 4. We conclude in Section 5.

2 Santiago's 2017 vintage restriction

The city of Santiago, Chile's capital and home to 40% of its population, suffers from long-standing air pollution problems, due partly to its geography—the fact it is surrounded by mountains—but also to a steady increase in car use. According to AirVisual (2018), Santiago is the twentieth most (local) air-polluted capital city in the world.[5] Efforts to control vehicle emissions date back to at least the mid 1980's, first in 1985 with a total prohibition on the import of used cars and then in the winter of 1986, with the introduction of a driving restriction program that placed a uniform restriction upon all cars regardless of their vintage or emissions rate. This restriction program has witnessed various reforms since then, the last of which occurred in 2017 when it partially terminated the 1993 exemption that was placed in 1992 upon all cars vintage 1993 and newer.

The 2017 reform established from 2018 onward a one-day-a-week restriction upon all cars vintage 2011 and older and none upon cars vintage 2012 and newer. Hybrid and electric vehicles remain exempt regardless of their vintage. The reason for choosing 2011 as the threshold vintage is the adoption of stricter emission standards for any new car entering the country in 2012 and later: Euro IV and Euro V norms for gasoline and diesel cars, respectively. Thus, the 2017 reform introduced a sharp discontinuity between the 2011 and 2012 vintages that we exploit here to illustrate the impact of a vintage restriction on the evolution of the car market in both restricted and non-restricted areas.[6] We are particularly interested in understanding the exodus of restricted cars (vintage 2011 and older) from the restricted area (the city of Santiago) to the periphery of the restricted area and more distant locations.

[5]Cars are major contributors of carbon monoxide (CO), hydrocarbons (HC), nitrogen oxides (NOx), and fine particles (PM2.5). HC and NOx are precursors to ground-level ozone (O3, also known as smog) and also contribute to the formation of PM2.5. At least in Santiago, vehicles are responsible for 30 and 36% of PM2.5 and O3 concentrations, respectively (Rizzi and De La Maza, 2017). These local pollutants, unlike global pollutants such as carbon dioxide (CO2), are characterized as having a local impact, at the city level, that lasts for a short time, sometimes only a few hours.

[6]Some readers may wonder whether 2012 models could have vehicle characteristics—in addition to those needed to comply with the new Euro norms—that make them more attractive to 2011 models after controlling for vintage. There are three reactions to this. As explained by Alé-Chilet et al (2023), one is that compliance with stricter environmental norms can sometimes make cars less, not more, attractive to drivers. The second reaction comes from the price analysis run by Bulnes (2018, p. 58). Using historic transaction prices for some popular models, he finds no price jump around the 2011/2012 discontinuity for prices collected before the 2017 reform. When looking at prices collected after the 2017, however, he does find evidence of a price jump around the 2011/2012 discontinuity: 2012 models are relatively more expensive than 2011 models after correcting for the vintage difference. And the third reaction comes from our own analysis of car fleet compositions across municipalities of different income levels. If 2012 models were equipped with desirable characteristics that 2011 models did not have, then we should find a slightly larger proportion of 2012 cars relative to 2011 cars in municipalities of higher income, where we should expect these characteristics to be better appreciated. We fail to find any difference.

2.1 Descriptive evidence

We start by looking at changes in fleet composition caused by the 2017 reform. In March every year, each car owner is required to obtain a circulation permit upon payment of an annual fee to her home municipality, so the database we use for such purpose comes from vehicle circulation permits at the municipality level collected by the National Statistics Bureau. Our permit-circulation database, which goes from 2007 to 2020, contains information for almost all the 346 municipalities in the country (only a few small and remotely located municipalities are missing from our database). These 346 municipalities are organized around 16 regions. The city of Santiago belongs to the Metropolitan Region (MR), which includes a total of 52 municipalities, 32 of which form the city of Santiago. For the purposes of our analysis, these 32 municipalities make the restricted area where the 2017 reforms applies.

As shown in Table 1, the municipalities in Santiago tend to be richer and with more cars than in the rest of the country. The table also identifies two areas that would play a role in our analysis of Section 4. Santiago's Periphery includes all the remaining municipalities that make the Metropolitan Region and the Neighboring Regions include the municipalities that are in Regions V and VI, which are the two regions that share a border with the Metropolitan Region. The 2017 reform is expected to cause an exodus of restricted cars from Santiago to municipalities located primarily in these two areas, i.e., Santiago's Periphery and Neighboring Regions. And since the car market in the Neighboring Regions is 63% larger than the market in Santiago's Periphery, in the absence of transaction costs we should expect a much larger fraction of the restricted cars exiting Santiago to end up in the Neighboring Regions and beyond.

Table 1: Characteristics of municipalities included in different areas by 2015

Area	Number	Population	Income	Vehicles	$\leq$ 2011	$\geq$ 2012
Santiago	32	164,080	\$493,099	0.164	0.102	0.062
		(102,006)	(\$369,985)	(0.127)	(0.073)	(0.056)
Santiago's Periphery	20	93,977	\$317,178	0.130	0.083	0.047
		(128,144)	(\$78,951)	(0.059)	(0.034)	(0.026)
Neighboring Regions	69	39,850	\$279,913	0.143	0.093	0.050
		(61,099)	(\$70,198)	(0.127)	(0.069)	(0.060)
Remaining Regions	216	40,719	\$253,945	0.106	0.073	0.034
		(59,763)	(\$70,447)	(0.115)	(0.079)	(0.048)
Total	337	56,872	\$289,414	0.122	0.081	0.041
		(81,620)	(\$155,800)	(0.118)	(0.075)	(0.051)

Notes: Based on information from the 2017 national Census and the 2015 circulation-permit database, this table contains characteristics at the municipality level for different areas of the country, including, number of municipalities in each area, average population, average monthly income per capita (in 2017 Chilean pesos), total number of vehicles per capita, number of restricted cars per capita (i.e., any model vintage 2011 or older), and number of no-restricted cars per capita (i.e., any model vintage 2012 or younger). Standard deviations are in parenthesis.

Figure 1 provides some evidence on the car exodus. The figure contrasts fleet compositions three years before (2014) and after (2020) the reform for restricted (municipalities in the city of Santiago) and non-restricted (municipalities in the rest of the country) areas. Darker bars correspond to pre-2012 models (i.e., 2011 and older), the ones subject to the restriction, and lighter bars correspond to post-2011 models. Although the effect is not particularly significant, before the 2017 reform Santiago exhibits a similar number of 2011 and 2012 vehicles than the rest of the country (Panels 1a and 1c). In contrast, in 2020 Santiago exhibits a larger fraction of 2012 vehicles relative to 2011 vehicles, almost 10,000 more (Panel 1b). For the rest of the country we observe the exact opposite, a fewer number of 2012 models relative to 2011 models (Panel 1d).

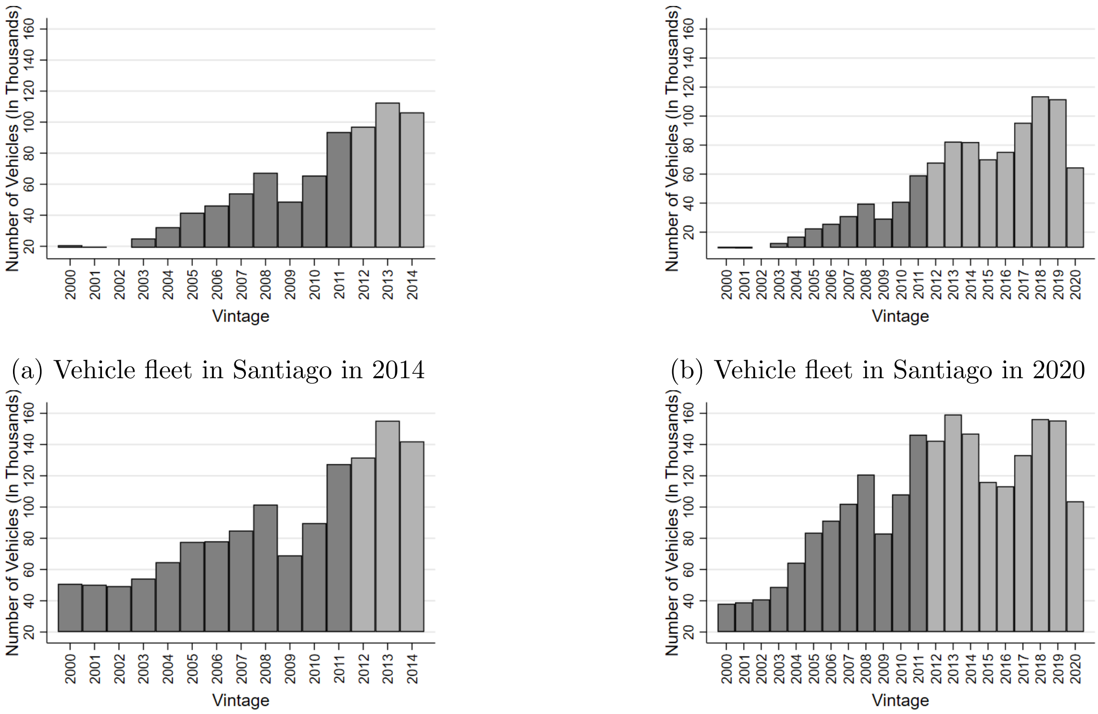

(a) Vehicle fleet in Santiago in 2014

(b) Vehicle fleet in Santiago in 2020

(c) Vehicle fleet in the rest of the country in 2014

(d) Vehicle fleet in the rest of the country in 2020

Figure 1: Evolution of vehicle fleet for Santiago and rest of the country

Notes: The figure depicts the vehicle fleet by vintage from 2000 to 2015 for Santiago and the rest of the country, contrasting years 2014 (pre-policy) and 2020 (post-policy). Darker bars denote restricted cars, vintage 2011 an older.

Without controlling for other variables, however, it is not obvious a priori how much of what we see in Figure 1 is due to the 2017 reform and how much is due to characteristics specific to Santiago that might affect car-purchasing decisions (e.g., higher average income in Santiago). The exercise shown in Figure 2 is a first attempt at addressing this concern. The figure plots changes in vehicle fleet as a function of income around the discontinuity 2011-2012 between 2014 and 2020 for different municipalities in the country. More precisely, it plots the following difference

$$\Delta Q^i \equiv \frac{q^i_{2011,2020}}{q^i_{2012,2020}} - \frac{q^i_{2011,2014}}{q^i_{2012,2014}} \tag{1}$$

where $q^i_{\tau,t}$ is number of cars vintage τ registered in municipality i in year t. If in some given municipality that difference turns out to be -0.5, it would mean that in 2014 there was one 2012 model for each 2011 model while in 2020 there was two 2012 models for each 2011 model.

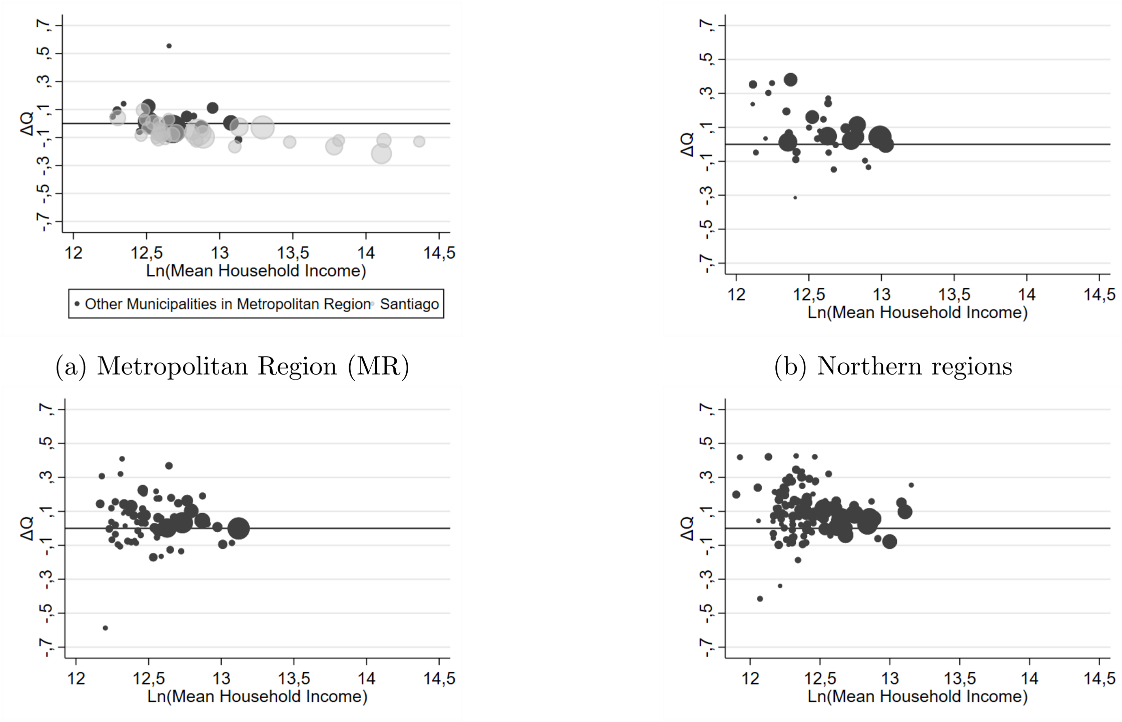

(a) Metropolitan Region (MR)

(b) Northern regions

(c) Neighboring regions (Regions V and VI)

(d) Southern regions

Figure 2: Change in vehicle ratio 2011/2012 between 2014 and 2020 by location and income

Notes: The figure illustrates the vehicle fleet as a function of income at the 2011-2012 discontinuity. It spans from 2014 to 2020 and includes municipalities in Santiago, Neighboring Regions, and further north and south.

Panel 2a contains all the municipalities in the Metropolitan Region, restricted (gray dots) and non-restricted (black dots) ones.[7] Both high and middle-income municipalities in Santiago exhibit a smaller fraction of restricted vintages than in non-restricted municipalities of comparable income. This difference does not extend to lower-income municipalities, which may suggest that some restricted cars that left higher income municipalities remained nevertheless in the restricted zone, but in lower-income hands (probably, these lower-income individuals saw an opportunity to buy these now cheaper cars). Similarly, the panel shows evidence of restricted cars leaving the restricted area and ending up in poorer municipalities outside Santiago. The other panels in the figure show that this pattern extends to municipalities in regions other than Santiago, those that border with the Metropolitan Region (Regions V and VI) and those further away (northern and southern regions).

While indicative of the impact of the 2017 reform on the car market, to have a better idea of its magnitude and scope we need to more formally isolate the effect of the reform from other

[7]Dots vary by size according to municipality population. This information was obtained from the 2017 national Census.

variables. We turn to that now.

2.2 Empirical analysis

To isolate the effects of the reform on the evolution of the fleet from other factors we start by running the following regression for each year between 2015 and 2020:[8]

$$\ln(q^i_{\tau,t}) = \beta_\tau DR_i + \omega_\tau \ln(INCOME_i) + \mu_\tau \ln(POP_i) + \delta_\tau + \varepsilon_i \tag{2}$$

where $q^i_{\tau,t}$ is defined as in (1), DR_i is a dummy variable equal to 1 if municipality i is in the restricted zone (i.e., city of Santiago), $INCOME_i$ is i' average income per capita, POP_i is i's population, δ_τ are vintage fixed effects, and ε_i is an error grouped at the municipality level. The coefficient of interest is β_τ, specifically the difference between β_{2011} and β_{2012}. It is expected that $\beta_{2011} < \beta_{2012}$, but only after the reform, when we should expect an exodus of old cars from the restricted municipalities.

Figure 3 plots the values of β_τ that we obtain from estimating (2). The figure shows that as years pass all coefficients are decreasing in value, so even after controlling for income and population, Santiago exhibits a faster fleet turnover. Moreover, we can see how after 2018 a statistically significant (at any usual level) discontinuity appears over the years 2011 and 2012, just as expected.[9]

A potential concern with the estimation of β_τ in (2) is that it can be biased by unobservable factors, unless the discontinuity in β_τ is only observed for vintages 2011 and 2012. It would be hard to imagine an unobservable factor that only affects substitution decisions for 2011 and 2012 cars, and not for any other pair of contiguous vintages. Accordingly, we proceed to run the following equation:

$$\begin{aligned}\ln(q^i_{\tau,t}) = {} & \alpha_\tau 1[t=2020] + \beta_\tau DR_i + \gamma_\tau 1[t=2020] \times DR_i + \omega_\tau \ln(INCOME_i) \\ & + \lambda_\tau 1[t=2020] \times \ln(INCOME_i) + \mu_\tau \ln(POP_i) \\ & + \phi_\tau 1[t=2020] \times \ln(POP_i) + \delta_\tau + \varepsilon_i\end{aligned} \tag{3}$$

where $t = \{2015, 2020\}$, and $1[t = 2020]$ is a dummy variable equal to 1 if the year is 2020, meaning post policy.[10] In this case, the coefficient of interest is γ_τ, where once again we expect $\gamma_{2011} < \gamma_{2012}$.

Figure 4 plots the values of γ_τ in (3). There is a clear, statistically significant discontinuity

[8] Looking at the underlying distribution of the data, in the Appendix we justify why this log-linear functional form describes the data better than letting the quantity to enter linearly. Yet, when we allow for the latter results do not qualitatively change in this and subsequent regressions.

[9] Although, this is not apparent at first glance from the figure, when performing a t-test both coefficients are statistically dissimilar. This is due to the covariance of both estimations.

[10] We use years 2020 and 2015 not only because they are at a symmetric distance from the treatment start but also because 2020 is the last year of our sample, where effects should be more pronounced as suggested by Figure 3. Similar results are obtained when we use 2019 instead of 2020. Note also that data from 2020 is by no means affected by the covid-19 pandemic since it is based on information collected by March of that year, right before the pandemic.

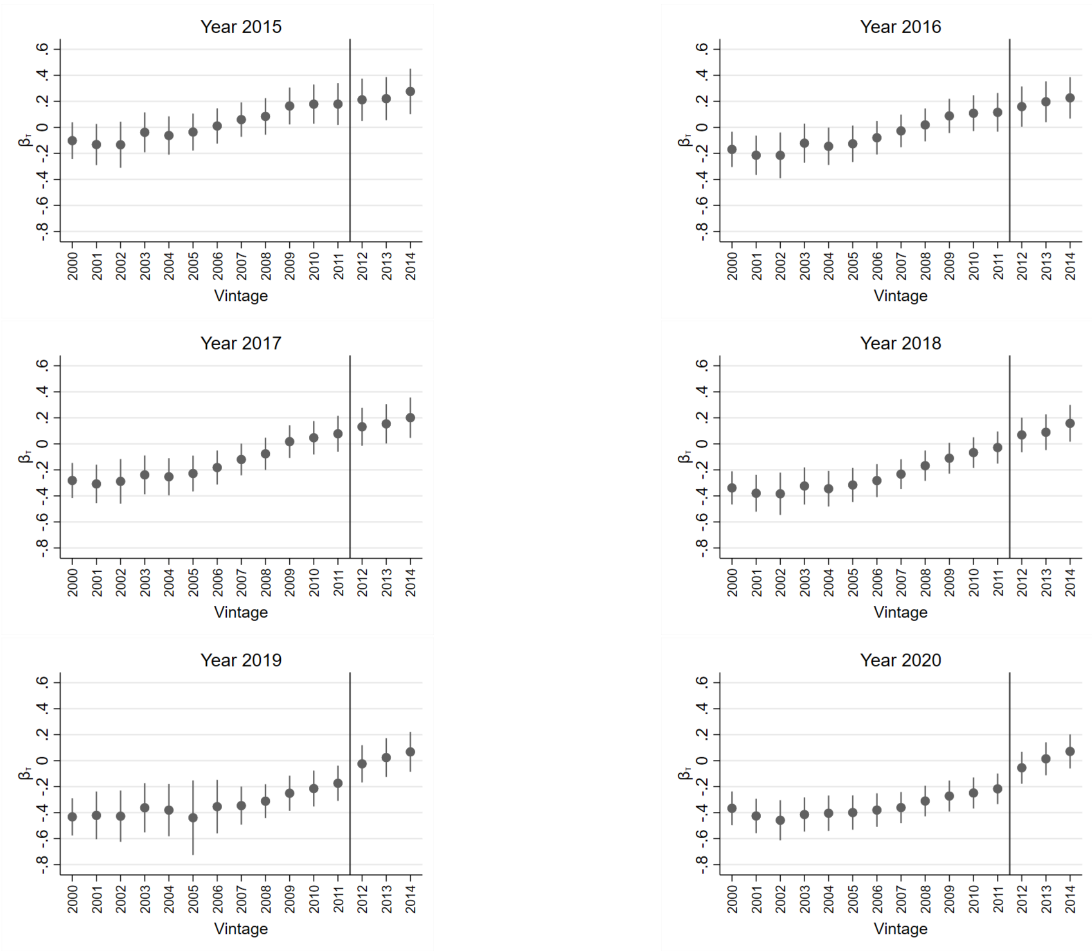

Figure 3: Changes in fleet composition around the 2011/2012 discontinuity: 2015-2020

Notes: The figure shows estimation results of coefficients β_τ in Regression (2) for each year. Confidence intervals are at 95%, clustering at the municipality level. The vertical line indicates the start of the reform.

around vintages 2011 and 2012, specifically, the restriction has reduced vintage 2011 by an extra 13 percentage points (pp) relative to the decrease of vehicles of vintage 2012 in the restricted area in the same period. The coefficients follow a U-shape. The restriction has larger effects for vintages close to the discontinuity. For older vintages, such as 2000 (whose coefficient is statistically equal to that of vintage 2012), there is no clear effect. This lack of effect for distant vintages is also in Barahona et al (2020).

In order to estimate the effect of income in changes in fleet composition, we extend (3) and

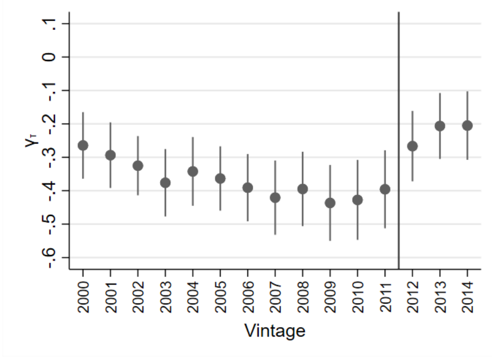

Figure 4: Impact of the 2017 reform by 2020

Notes: This figure shows estimation results of coefficients $\gamma_\tau 1[t = 2020]$ in Regression (3) for years 2015 and 2020. Confidence intervals are at 95%, clustering at the municipality level. The vertical line indicates the start of the reform.

run:

$$\begin{aligned} \ln(q^i_{\tau,t}) = {} & \alpha_\tau 1[t = 2020] + \psi^1_\tau DR_i \times High_i + \psi^2_\tau DR_i \times Medium_i + \psi^3_\tau DR_i \times Low_i \\ & + \beta^1_\tau 1[t = 2020] \times DR_i \times High_i + \beta^2_\tau 1[t = 2020] \times DR_i \times Medium_i \\ & + \beta^3_\tau 1[t = 2020] \times DR_i \times Low_i + \kappa^1_\tau High_i + \kappa^2_\tau Low_i + \omega_\tau \ln(INCOME_i) \\ & + \lambda_\tau 1[t = 2020] \times \ln(INCOME_i) + \mu_\tau \ln(POP_i) \\ & + \phi_\tau 1[t = 2020] \times \ln(POP_i) + \delta_\tau + \varepsilon_i \end{aligned} \tag{4}$$

where the difference between this regression and (3) is the inclusion of three income dummies: $High_i$ (above the 90th percentile) $Medium_i$ (between the 30th and 90th percentiles) and Low_i (below the 30th percentile).

The values of β^1_τ, β^2_τ, and β^3_τ in (4) are shown in Figure 5. According to the figure, the effects of the 2017 reform are felt mostly in low- and middle-income municipalities, which present a clear discontinuity over vintages 2011-2012. The reform appeared to have reduced the number of 2011 vehicles in these municipalities by 14.4 pp and 14.8 pp relative to the drop in 2012 vehicles, respectively. The figure also shows that although high-income municipalities in Santiago continue exporting a large number of older, restricted vehicles, this is not due to the reform. It is rather explained by their stronger preferences for newer models (i.e., higher income).

An explanation for why the reform appears to have had no effect on high-income municipalities is that most households in these municipalities own two or more vehicles, most of them young enough to not be affected by the restriction. Thus, having just one car of many affected by the restriction, if at all, can be easily accommodated at virtually no cost. In contrast, most households in low- and middle-income localities own one car at most, so for them the only option to accommodate to the restriction, provided their members want to continue driving everyday of the week, is to upgrade their cars. A similar "no-effect at the top" was found by

Gallego et al (2013) for the *Hoy No Circula* program in Mexico City.

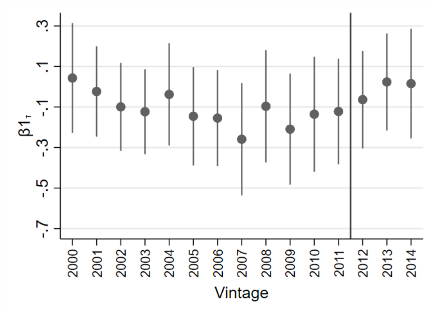

(a) Coefficients $\beta^1_\tau 1[t=2020] \times DR_i \times High_i$

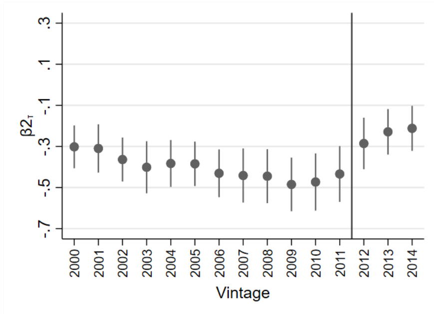

(b) Coefficients $\beta^2_\tau 1[t=2020] \times DR_i \times Medium_i$

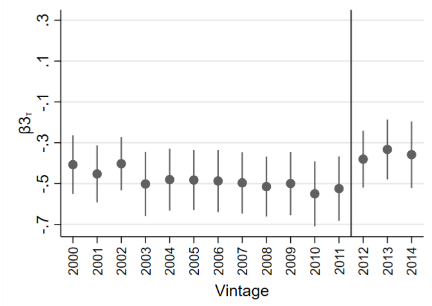

(c) Coefficients $\beta^3_\tau 1[t=2020] \times DR_i \times Low_i$

Figure 5: Impact of the 2017 reform for different income levels by 2020

Notes: This figure shows estimation results of coefficients β^1_τ, β^2_τ, and β^3_τ in Regression (4) for years 2015 and 2020. Confidence intervals are at 95%, clustering at the municipality level. The vertical line indicates the start of the reform.

Finally, in an attempt to test for the presence of transaction costs we estimate a variation of (3) by including an additional dummy to asses the extent to which municipalities outside Santiago (i.e., Santiago's Periphery) are importing the restricted vehicles that are leaving Santiago:

$$\begin{aligned} \ln(q^i_{\tau,t}) = \alpha_\tau 1[t=2020] + \beta_\tau DR_i + \gamma_\tau 1[t=2020] \times DR_i + \nu_\tau 1[MR_i=1, DR_i=0] \\ + \varphi_\tau 1[t=2020] \times 1[MR_i=1, DR_i=0] + \omega_\tau \ln(INCOME_i) \\ + \lambda_\tau 1[t=2020] \times \ln(INCOME_i) + \mu_\tau \ln(POP_i) \\ + \phi_\tau 1[t=2020] \times \ln(POP_i) + \delta_\tau + \varepsilon_i \end{aligned} \tag{5}$$

where the additional dummy $1[MR_i = 1, DR_i = 0]$ takes the value of 1 if municipality i belongs to Santiago's Periphery, i.e., belongs to the Metropolitan Region (MR) but is outside Santiago. The coefficient of interest for this regression is φ_τ, which represents the change in the quantity of each vintage in municipalities near Santiago post-policy compared to that of municipalities far from it.

Figure 6 shows the values of φ_τ and γ_τ in (5). As expected, the γ_τ coefficients in Panel 6a—those capturing policy effects in Santiago— are almost identical to those estimated in

(3). Perhaps surprisingly, the φ_τ coefficients in Panel 6a show no clear discontinuity around 2011-2012 vintages for municipalities near Santiago. This implies that the fraction of restricted vehicles flowing to municipalities in Santiago's Periphery is not that different from the fraction of the these same vehicles flowing to more distant municipalities, in Neighboring Regions.

This is clear evidence of the presence of transaction costs—preventing cars to be costlessly moved to more distant locations—for three reasons. First, as shown in Table 1, the average income of municipalities in Regions V and VI (which are the regions that share a border with the Metropolitan Region) is similar to the average income of municipalities near Santiago (those in Santiago's Periphery); see also Panels 2a and 2c. Second, the car market (measured by fleet size) is much larger in Regions V and VI than in Santiago's Periphery. And third, restricted cars are more valuable in more distant municipalities, since they are never expected to be driven through the restricted zone.

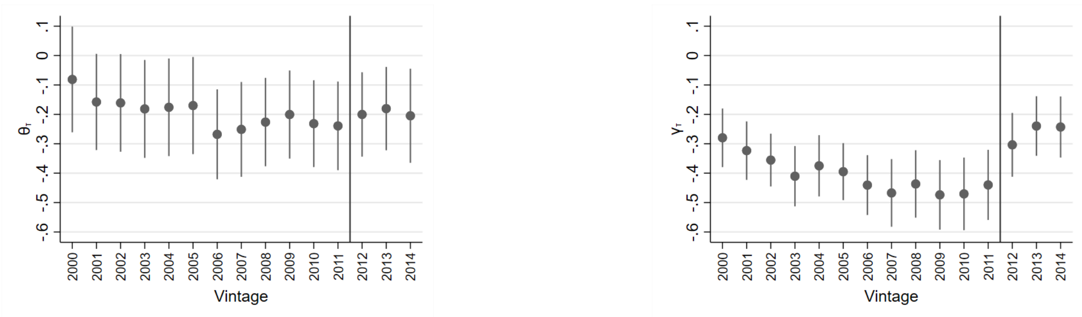

(a) Coefficients $\varphi_\tau 1[t = 2020] \times RM_i$

(b) Coefficients $\gamma_\tau 1[t = 2020] \times DR_i$

Figure 6: Testing for the presence of transaction costs

Notes: This figure shows estimation results of coefficients $\varphi_\tau 1[t = 2020] \times MR_i$, and $\gamma_\tau 1[t = 2020] \times DR_i$ in Regression (5) for years 2015 and 2020. Confidence intervals are at 95%, clustering at the municipality level. The vertical line indicates the start of the reform.

That a large fraction of the restricted cars leaving Santiago stay in localities near Santiago certainly points to the possibility of emissions leakage. According to the last Origin-Destination Survey, over 30% of the trips originated at the outskirt of Santiago terminate in the city of Santiago. Furthermore, the remaining 70% of trips that originate and terminate outside Santiago are far from harmless. According to estimates by the Ministry of the Environment, a good fraction of the emissions generated in these nearby localities travel to Santiago helped by winds that blow from coast to mountains (from west to east). Whether these forces are enough to severely undo the effect of the policy (and eventually lead to more pollution than otherwise) is something we cannot tell without the help of a model of the car market, something we turn to in Sections 3 and 4.

2.3 Pre-reform benchmark

Before we move to the modelling sections, we use the permit-registration data to construct a pre-policy benchmark that will serve to better tune some of the parameters of the model that we will use to estimate policy effects and the role of transaction costs. We run (5) for pre-policy years. The implicit assumption in our exercise is that before the policy, vehicles of the same age were equally likely to be traded outside Santiago regardless of the year in question. In other words, we should expect no difference, for example, between the fraction of 2008 vehicles leaving Santiago in year 2014 than the fraction of 2011 vehicles leaving in 2017.

Results of estimating (5) for pre-policy years are summarized in Figure 7. Panel 7a shows almost no difference in the flows of vehicles to Santiago's Periphery relative to that to more distant locations (Regions V and VI). Both areas have seen their fleets increase in similar proportions, except for middle-age vehicles of around 10-15 years old (vintages 2000 to 2005) that exhibit a lower increase in Santiago's Periphery than in the more distant locations, about 20 pp lower. Panel 7b, on the other hand, shows a clear export of old vehicles from Santiago of all vintages, of around 20% following a slight U-shape with a valley in 2002-2003 vintages.

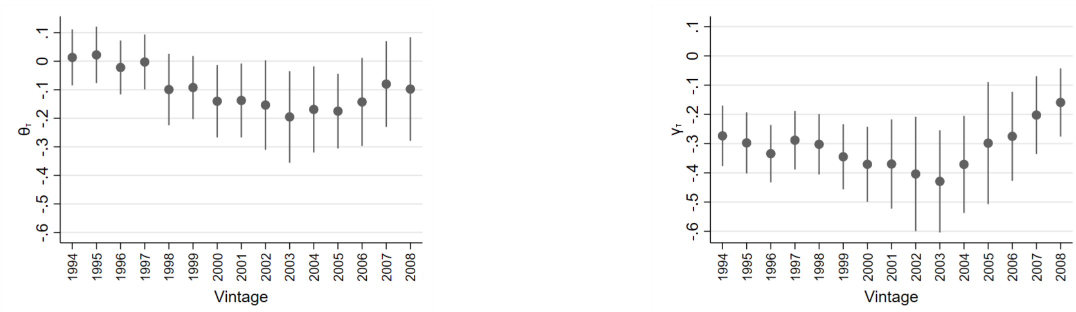

(a) Coefficients $\theta_\tau 1[t = 2014] \times RM_i$

(b) Coefficients $\gamma_\tau 1[t = 2014] \times DR_i$

Figure 7: Pre-reform benchmark

Notes: This figure shows estimation results of coefficients $\varphi_\tau 1[t = 2014] \times MR_i$, and $\gamma_\tau 1[t = 2014] \times DR_i$ of Regression 5 for years 2009 and 2014. Confidence intervals are at 95%, clustering at the municipality level. The vertical line indicates the start of the reform.

3 A model of the car market

3.1 Setup

The model follows very closely the one in Barahona et al (2020) except for two simplifications and two innovations. The first simplification is the focus on steady-state outcomes; two in particular, before and after the policy intervention. The second simplification is a focus on movements along the extensive margin only, on the type of cars individuals drive, and not along the intensive margin, on how much they drive. This simplification is without much loss since in this paper we will be exclusively looking at vintage restrictions, which only work through the extensive margin.[11]

[11] Barahona et al (2020) also study other policy instruments that do operate through the intensive margin such as fuel taxes.

The first innovation is that the car market now encompasses three different zones. There will be an inner, intermediate and outer zone. The inner zone is where pollution is more acute and the restriction effectively applies. In our case this would be the city of Santiago. Even if the restriction does not strictly falls over the entire city, in practice it does because most trips, if not all, must go through the actual restricted area at some point. This is particularly relevant for commuting trips to work or school which for most part have as final destination a location within the actual restricted zone.

The intermediate o transition zone includes areas at the outskirts of the inner zone with individuals that periodically need to travel to the inner zone, going through the actual restricted area. As a result, these individuals also suffer the impact of the restriction, but not as much as those living in the inner zone. Their trips are also expected to contribute to pollution in the inner zone, but not as much as those from individuals living in the inner zone. In our case, the intermediate zone would be composed by municipalities that fall outside the city of Santiago but still belong to the Metropolitan Region.

Finally, there is the outer zone with individuals that never travel to the inner zone. In our case this would be any municipality outside the Metropolitan Region, in particular those in Regions V and VI.

These three areas are denoted by restricted (r), intermediate/transition (t), and unrestricted (u). The total number of individuals living in these three areas is normalized to 1. We let $\beta^j \in (0,1)$ be the fraction living in area $j \in \{r,t,u\}$, so $\sum_j \beta^j = 1$.

The second innovation with respect to Barahona et al (2020) is the presence of transaction costs in the car market, that is, the costs of moving second-hand cars across the different zones. We leave the details for later when we explain how the car market clears. Part of the contribution of this paper is to understand the role of the relative size of the different areas and of transaction costs for emissions and welfare.

3.2 Vehicles, restrictions and individual preferences

To simplify the model we cluster vehicles of different ages in 4 groups denoted by $a \in \{0,1,2,3\}$. Each group includes cars of 6 contiguous ages or vintages. Thus, group $a = 0$ includes cars up to 5 years old, $a = 1$ between 6 and 11 years old, $a = 2$ between 12 and 17 years old, and $a = 3$ between 18 and 23 or more years old (very few cars last for more than 24 years). The probability that a car age a survives to the next age-group is $\zeta_a \leq 1$, where $\zeta_a \geq \zeta_{a+1}$ and $\zeta_3 = 0$ by construction.

The age of a car determines two important characteristics, its quality and level of restriction. The quality of a car is the same regardless of its location and equal to $s_a > s_{a+1}$. On the other hand, we denote by $R_a^j \in [0,1]$ the level of restriction faced by an age-a car in zone j, with $R = 1$ meaning no restriction and $R = 0$ full restriction. In our application R_a^j will be equal to 1 for all cars in the unrestricted zone ($j = u$) and less than 1 for old cars in the restricted zone ($j = r$). The restriction faced by cars in intermediate zone ($j = t$) will be somewhere in between.

Individuals do not own cars. They rent cars to competitive dealers, which buy them from car producers at price c.[12] In each zone there will be a continuum of individuals/drivers that vary in their willingness to pay for quality. A type-θ individual who lives in zone j and rents an age-a car for p_a^j obtains a utility of (since we will be focusing on steady-state outcomes we can omit any reference to time)

$$u(\theta) = \theta R_a^j s_a - p_a^j \tag{6}$$

where $\theta \in [0, \bar{\theta}^j]$ according to some cummulative distribution function $F^j(\theta)$. Note that p_a^j is zone-specific. This is because of transactions costs.

There will be a cutoff level θ_a^j that makes an individual in zone j indifferent between renting a car age $a \in \{0, 1, 2\}$ and age $a + 1$,

$$\theta_a^j R_a^j s_a - p_a^j = \theta_a^j R_{a+1}^j s_{a+1} - p_{a+1}^j$$

There will be also a cutoff level θ_3^j that makes an individual in zone j indifferent between renting an age-3 car and taking his/her outside option (e.g., taking public transport), which we normalize to zero, so

$$\theta_3^j R_3^j s_3 - p_3^j = 0 \tag{7}$$

3.3 Market clearing

Car dealers have the option to either rent their age-3 cars to drivers or scrap them (the residual value of an age-3 car that is rented instead of scrapped is zero). The scrappage value of a car is zone-specific and given by the (inverse) demand scrappage function

$$v^j(q_v^j) = \frac{b^j}{(q_v^j)^{\upsilon}} - l^j \leq \bar{v}^j$$

where $\bar{v}^j$, υ, b^j and l^j are parameters to be determined and q_v^j is the number of cars being scrapped.

As explained above, one of the innovation of the model is the presence of transaction costs. Dealers incur in an additional total cost of $\tau^{jk}(\Delta q_a^{jk})^2$ when moving a total of Δq_a^{jk} age-a cars from zone j to the contiguous zone $k \neq j$.[13]

Car dealers compete Bertrand with no horizontal differentiation.[14] Accordingly, in equilib-

[12]The assumption that individuals rent cars in each period instead of owning them is also in Barahona et al (2020). In their application of the model, however, each period corresponds to four years (in our application to five years), so in reality we are assuming that individuals own a car for 4/5 years before they trade them. This trading horizon is consistent, on average, with the trading horizons in Gavazza et al (2014).

[13]Marginally increasing transaction costs ensure positive vehicle flows between areas for a wide set of parameters values. Besides, this assumption is supported in practice. Not all vehicles are equally costly to be transferred across zones either for stock or geographic reasons. Some dealers are larger than others and with sales offices in the different zones.

[14]As explained in Barahona et al (2020), results will not change if we instead assume a constant mark up.

rium, car dealers must be indifferent between buying a new car and subsequently renting it in any of the three zones and, even more, between leaving it in a given zone and transferring it to another one. Since there will be always cars that remain in the same zone its entire life, the next break-even condition must hold in each area j

$$c = \sum_{a=0}^{3} (\delta\zeta_a)^a p_a^j \tag{8}$$

where δ is the discount factor that all agents use to discount the future.

Since a fraction of age-3 cars are scrapped in equilibrium (i.e., $q_v^j > 0$), car dealers must be indifferent between renting these vehicles one more time and receiving their scrapping value

$$p_3^j = v^j(q_v^j) \tag{9}$$

Moreover, at the margin dealers must be indifferent between moving an extra age-3 vehicle across adjacent areas and keeping it where it currently is. Thus, the price differential between two adjacent areas j and k for age-3 vehicles must satisfy the marginal condition

$$p_3^k - p_3^j = 2\tau^{jk}\Delta q_3^{jk} \tag{10}$$

where Δq_3^{jk} is the number of age-3 cars that cross the border of zones j and k. Note that if the flow of vehicles is from area j to k, so $\Delta q_3^{jk} > 0$, then $p_3^k > p_3^j$.

This same marginal condition must apply to younger second-hand cars, that is,

$$p_a^k - p_a^j = 2\tau^{jk}\Delta q_a^{jk} \tag{11}$$

for $a \in \{1, 2, 3\}$.

Since there is a probability ζ_a that an age-a vehicle survives, it must be true that in equilibrium the aggregate quantity of vehicles of each age-group must satisfy

$$q_0\zeta_0 = q_1\zeta_1 = q_2\zeta_2 = q_3 + q_v \tag{12}$$

where $q_v \geq 0$ is the total number of age-3 cars that are scrapped in equilibrium and q_a is the total number of vehicles age $a \in \{0, 1, 2, 3\}$ that are rented in equilibrium.[15]

Finally, as a function of the cutoff levels and rental prices of each zone, these quantities must satisfy the following clearing conditions

$$\sum_{j=r,t,u} \beta^j [F^j(\bar{\theta}^j) - F^j(\theta_0^j)] = q_0$$

[15]Note that individual preferences are such that only vehicles age $a = 3$ are scrapped in equilibrium.

and

$$\sum_{j=r,t,u} \beta^j [F^j(\theta_a^j) - F^j(\theta_{a+1}^j)] = q_{a+1}$$

for $a \in \{0, 1, 2\}$.

3.4 Pollution and welfare

Cars emit all sorts of pollutants, some with global effects (e.g., CO2) while other with local effects, i.e., effects at the city level and for much shorter periods of time (e.g., CO, HC, NOx). The focus of this paper is on local pollution. As documented by Barahona et al (2020) and more recently by Jacobsen et al (2023), cars emit more local pollution as they age. Therefore, we let e_a be an age-a car's emission per period in the absence of restriction ($R = 1$), with $e_{a+1} > e_a$. These emission levels control for the fact that older vehicles are run less often.

How harmful is a car's local emission to society depends not only on how often the car is used but also on where it is used. We assume that only the restricted zone suffers from air quality problems. We also assume that drivers who live in both the inner/restricted (r) and outer/unrestricted (u) zones use their cars in those zones only. The reality of drivers who live in the intermediate/transition (t) zone is slightly different. A fraction of their trips end up in the restricted zone, so they also contribute to its air quality problems. Thus, the pollution harm created by an age-a vehicle is given by $h^j R_a^j e_a$, where h^j denotes the car's contribution to pollution in the restricted zone, with $h^r > h^t > h^u = 0$.

In our context the social planner could restore the first-best outcome with rental taxes that vary by vintage and location, equal to the pollution externality, i.e., $h^j e_a$. In the absence of such taxes, the planner relies on vintage restrictions to, in principle, improve upon the no intervention outcome. Letting $\theta_{-1}^j \equiv \bar{\theta}^j$ to save on notation, the planner's welfare function in any given period can be written as

$$W = \sum_j \sum_a \beta^j \int_{\theta_a^j}^{\theta_{a-1}^j} R_a^j(\theta s_a - h^j e_a) dF^j(\theta) - c \sum_j \beta^j [1 - F^j(\theta_0^j)] + \sum_j S^j(q_v^j) \tag{13}$$

where $S^j(q_v^j)$ is the benefit of scrapping q_v^j cars in zone j, that is, $S^j(q_v^j) = \int_0^{q_v^j} v^j(x)dx$.[16]

The first term in (13) captures the social gain of driving a car. Note how the vintage restriction enters into the welfare function. As already documented by Barahona et al (2020), the socially optimal thing to do is to either place a complete restriction ($R = 0$) upon an age-a car in zone j when $\theta s_a - h^j e_a < 0$, or none ($R = 1$) when $\theta s_a - h^j e_a > 0$. The restriction in Santiago does not take this binary form, unlike LEZ programs that for the most part do. The remaining terms in (13) capture the cost of adding new cars into the market and the benefit of scrapping some old ones.

[16]Note that transaction costs have been omitted from (13). A good fraction of these costs take the form of information frictions, so in our competitive setting the way to estimate them would be to compare (13) with the welfare from the frictionless benchmark after accounting for changes in pollution harm.

4 Applying the model to Santiago's 2017 policy

In this section we adapt the model to study Santiago's 2017 policy. Our goal is not to produce a detail assessment of the policy. Among other things, this would require to compute transition dynamics which would be beyond the scope of the paper. Our goal is to use the model and data from Santiago's policy to illustrate the workings of these vintage restrictions, particularly the role of transaction costs and intermediate zones.

In this application, we will be comparing two steady-state outcomes: the pre-policy benchmark and the post-policy outcome. Given that the existing vintage restriction keeps the vintage threshold fixed at 2011, the post-policy steady-state outcome would only be reached when the policy is not longer relevant. Accordingly, for our analysis we consider a policy with a vintage threshold that adjust over time.[17] Given that at the time of its implementation in 2017, the policy exempted all cars 6 and less years old (vintage 2012 or later), we evaluate a policy with a moving exemption threshold of 6 years old. In terms of our model, this implies that all cars age $a = 0$ are exempt from the restriction. Since this policy is stricter than the actual policy, the possibility of emissions leakage appears more likely.

4.1 Parameter values and calibration

We make a number of simplifications. The cost of a new vehicle, c, is normalized to 1. This requires to set the quality of new cars to $s_0 = 0.41$, so as to maintain the same relative magnitudes between transport cost and surplus that we observe in Barahona et al (2020). We also adopt their linear decay rate: $s_a = 0.9^a$. In addition, we set the discount factor equal to one as well as the probabilities of survival: $\delta = \zeta_a = 1$ for all $a \in \{0, 1, 2, 3\}$.

As for car emissions, we normalize the emission of the older cars in the market to the unity, $e_3 = 1$, and follow the parameterization in Barahona et al (2020) for the other groups:[18] $e_a = 0.03 \exp(0.88(1+a))$. Pollution harm in the restricted zone is also borrowed from Barahona et al (2020) but in a "reduced form". In Barahona et al (2020), the no-intervention pollution cost is about one fifth of the benefit from driving. Using that ratio, we set $h^r = 0.87$ and let h^t to vary between 0 and h^r; more precisely, $h^t = \alpha h^r$, with $\alpha \in [0, 1]$.

With regard to the scrappage demand function $v^j(q_v^j)$, we let $b^j = 0.8$, $l^j = 1$, $\upsilon = 0.1$ and $\bar{v}^j = (p_2^j + p_3^j)/2$. The latter does not have equilibrium implications. It only enters in the welfare estimation. These parameters, along with the others, produce scrappage in equilibrium in all three zones and in levels consistent with what we observe in the data.

We do not have information to parameterize transaction costs (eventually it could be done with data on sales, which we do not have). Since one of the main goals of the paper is to understand the role of transaction costs on equilibrium outcomes, we simply normalize τ^{rt} to zero and let τ^{tu} to vary from zero to 1.5, which is high enough to significantly reduce trade.

[17] There are discussions to adopt an adjusting vintage threshold, as currently done in Mexico City.

[18] Note that Barahona et al (2020) consider older models, so we are only borrowing their relative pollution rates across models of different ages, not their absolute values.

The relative size of zone j, β^j, is proportional to its pre-reform fleet size, as recorded by annual circulation permits of 2015 (see the table below). It is important to mention that for the outer/unrestricted area, we only consider the two regions that share borders with Santiago's Metropolitan Region, that is, Regions V and VI.

Table 2: Size of different zones

	Number of cars	Relative Size
City of Santiago (Restricted zone)	920,784	0.60
Santiago's Periphery (Transition zone)	227,273	0.15
Neighboring Regions, V and VI (Unrestricted zone)	372,727	0.25

Note: Drawing from the 2015 circulation-permit database, this table presents the number of vehicles for Santiago, its peripheral municipalities, and neighboring regions. Additionally, it showcases their proportion in relation to the total number of vehicles in the three specified areas.

We assume that the distribution of consumers' willingness to pay for quality, θ, in all three areas follow an exponential function of the form $F^j(\theta) = 1 - \exp(-\lambda^j \theta)$ for $\theta \geq 0$ and where $(\lambda^r, \lambda^t, \lambda^u) = (0.5, 1.3, 1.3)$. These parameters capture the fact that average income in Santiago is higher than in the other two zones and that average income in these two other zones are comparable. Furthermore, this function and parameter values ensure pre-reform equilibrium outcomes to be consistent with those shown at the end of Section 2 (Pre-reform benchmark).

Finally, and following Barahona et al (2020), we let $R_0^r = 1$ and $R_a^r = 0.95$ for $a \in \{1, 2, 3\}$. For cars in the unrestricted zone we have that $R_a^u = 1$ for all a and for cars in the transition zone we adopt intermediate values, i.e, we let $R_0^t = 1$ and $R_a^t = 0.975$ for $a \in \{1, 2, 3\}$. We are ultimately interested in how pollution and welfare vary with transaction costs (τ^{tu}) and the contribution of the transition zone to pollution (α).

4.2 Pre-reform equilibrium

The pre-policy equilibrium outcome—prices, quantities and trade across zones—as a function of transaction costs is summarized in Figures 8, 9 and 10, respectively. Panels 8a–8d confirm that in the absence of transaction costs prices converge in the three zones. As we increase the cost of moving cars to (or from) the unrestricted zone interesting patterns emerge. The prices of new cars are now cheaper in the unrestricted zone. There are two reasons for that. One is that the demand for newer models (i.e., willingness to pay for higher quality) is lower in the unrestricted zone. The other is that given its smaller size, scrappage values in the unrestricted zone are relatively higher than in the other zones. A higher scrappage value necessarily implies, from (8), a lower rental price for new models in that same zone.

The equilibrium quantities depicted in Panels 9a–9d also exhibit interesting patterns as we increase transaction costs. For instance, there is a drop in the number of new vehicles (age-0) entering the inner zone while there is an increase in the outer zone. This is entirely explained by the interactions in the second-hand market. The presence of transaction costs makes the export of second-hand cars from higher-income zones to lower-income zones less profitable, ultimately

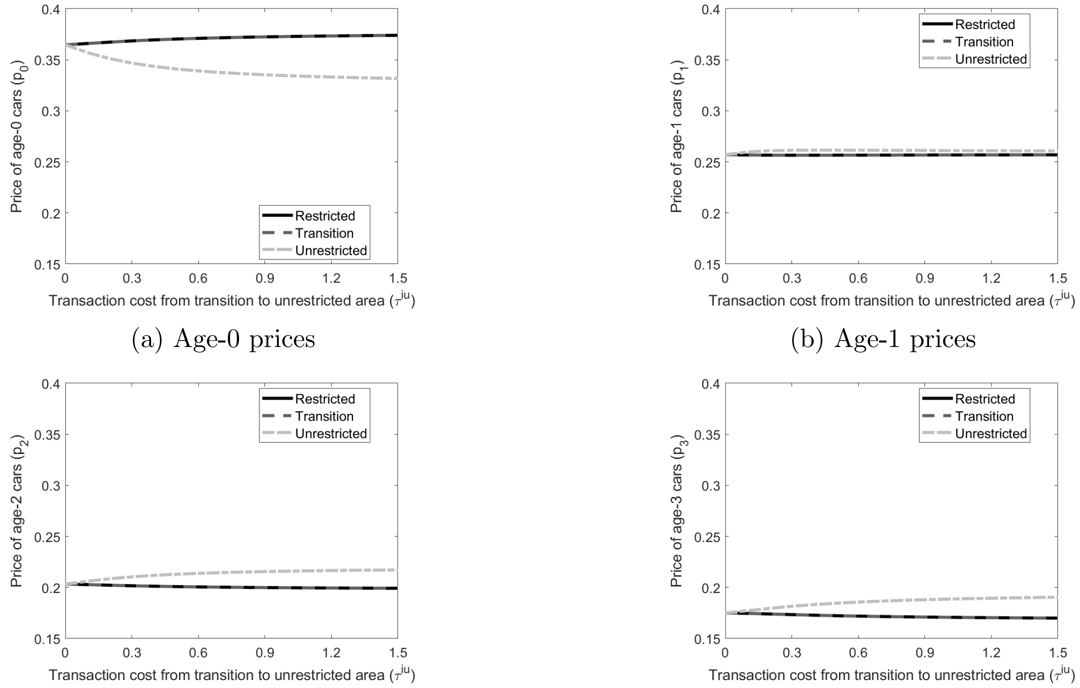

(a) Age-0 prices

(b) Age-1 prices

(c) Age-2 prices

(d) Age-3 prices

Figure 8: Pre-reform equilibrium prices for different models and zones

Note: This figure presents pre-policy equilibrium prices for each vehicle age and zone as a function of transaction costs.

depressing the number of new cars entering into the higher-income zones.

These export/import changes as a function of transaction costs can be better appreciated by looking at the Panels 10a–10c. Particularly noticeable is the non-monotonic flow of age-3 cars to the unrestricted and transition zones, which reaches its highest level for intermediate levels of transaction costs, around 0.3. As we will see shortly this non-monotonicity will have implications for emissions and welfare.

Results shown above can be made consistent with those at the end of Section 2 (pre-policy dynamics) if we believe transaction costs to be in a medium-low range, around 0.2. In fact, Panels 10a–10c show an export of cars of different ages—1, 2 and 3—from the restricted zone of about 20% (the solid line). These numbers are similar to those depicted in Panel 7b, indicating a drop of 20% on average for the different vintages; perhaps slightly more for cars around vintage 2002, equivalent to age-2 in our model.

4.3 Post-reform equilibrium

As mentioned above, our analysis focuses on comparing steady-state outcomes, before and after the intervention. The post-policy (steady-state) equilibrium outcome—prices, quantities and trade across zones—as a function of transaction costs is summarized in Figures 11, 12 and 13, respectively.

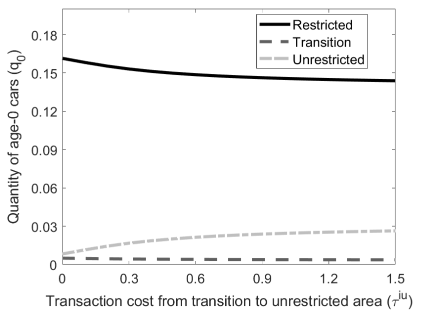

(a) Age-0 vehicles in different zones

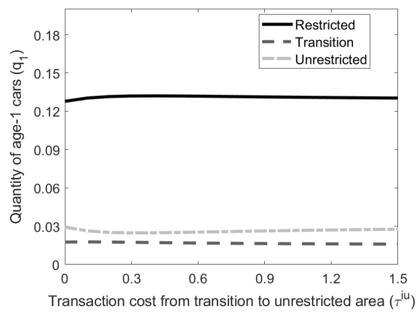

(b) Age-1 vehicles in different zones

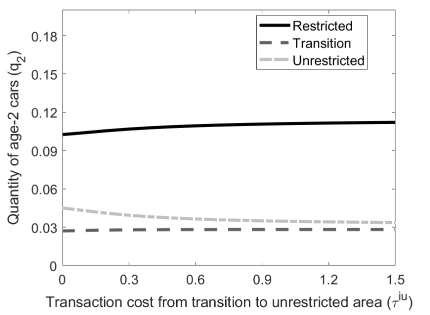

(c) Age-2 vehicles in different zones

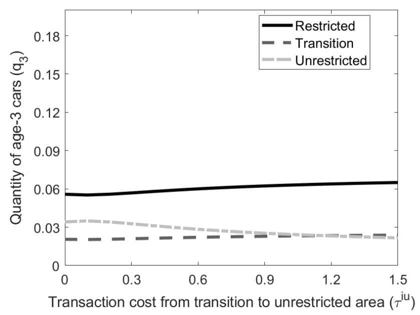

(d) Age-3 vehicles in different zones

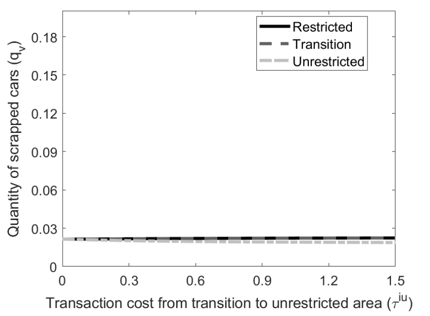

(e) Vehicles scrapped in different zones

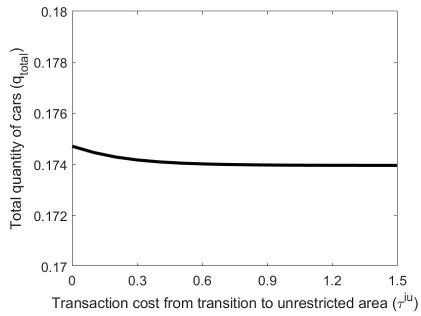

(f) Total number of cars

Figure 9: Pre-reform equilibrium quantities for different models and zones

Note: This figure presents pre-policy equilibrium quantities for each vehicle age and zone as a function of transaction costs.

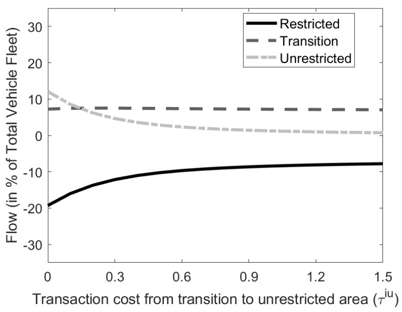

(a) Flow across zones of age-1 vehicles

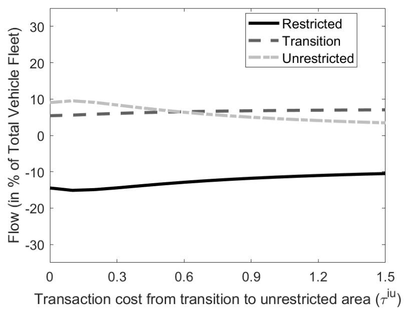

(b) Flow across zones of age-2 vehicles

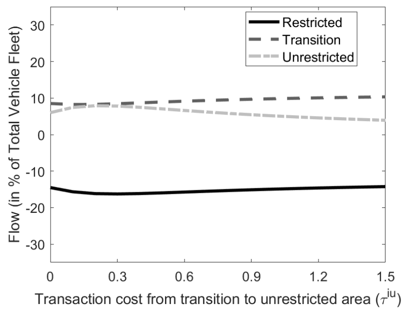

(c) Flow across zones of age-3 vehicles

Figure 10: Pre-reform equilibrium flows across zones for different models

Note: This figure presents pre-policy equilibrium flows across zones for each vehicle age as a function of transaction costs.

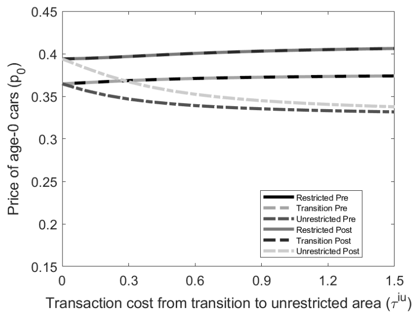

(a) Age-0 prices

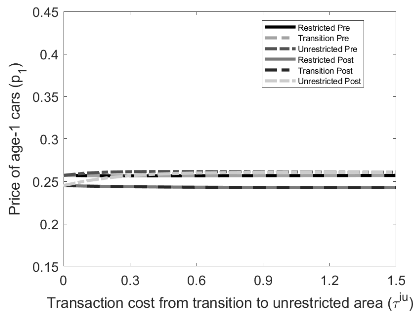

(b) Age-1 prices

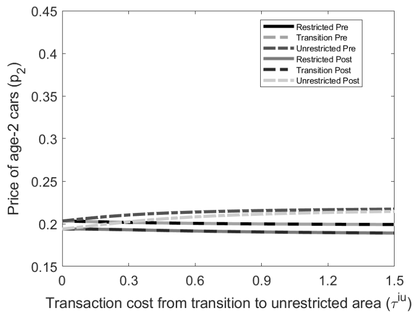

(c) Age-2 prices

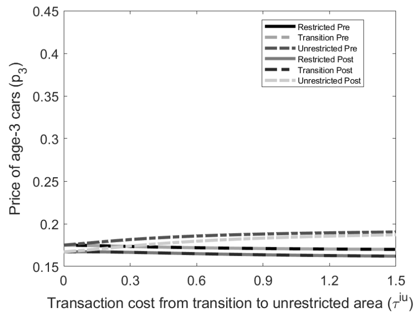

(d) Age-3 prices

Figure 11: Pre and post-reform equilibrium prices for different models and zones

Note: This figure presents pre and post-policy equilibrium prices for each vehicle age and zone as a function of transaction costs.

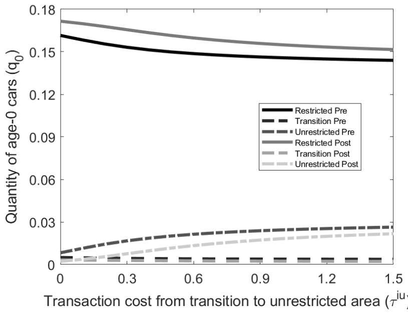

(a) Age-0 vehicles in different zones

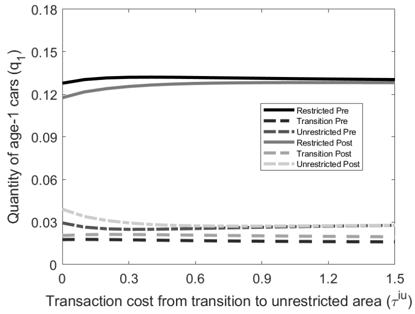

(b) Age-1 vehicles in different zones

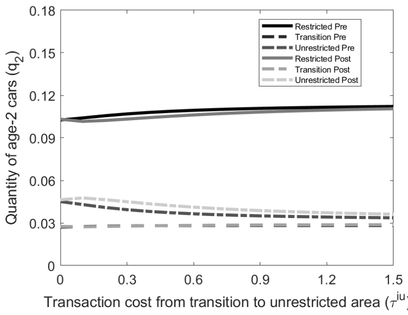

(c) Age-2 vehicles in different zones

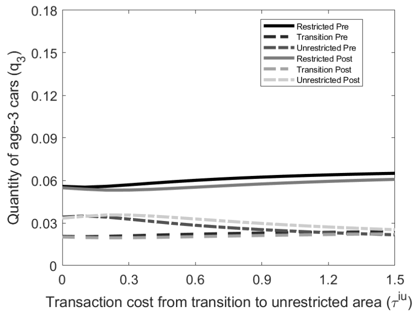

(d) Age-3 vehicles in different zones

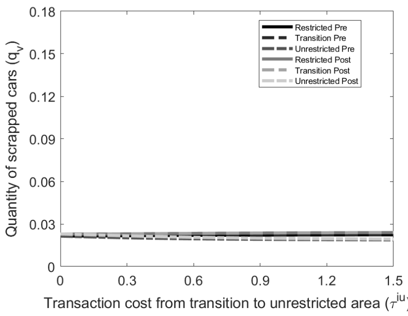

(e) Vehicles scrapped in different zones

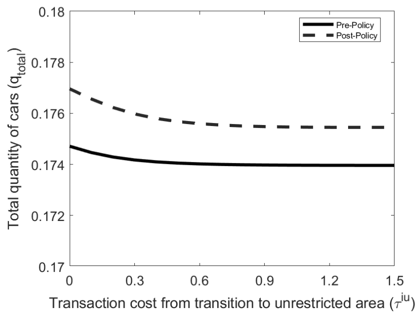

(f) Total number of cars

Figure 12: Pre and post-reform equilibrium quantities for different models and zones

Note: This figure presents pre and post-policy equilibrium quantities for each vehicle age and zone as a function of transaction costs.

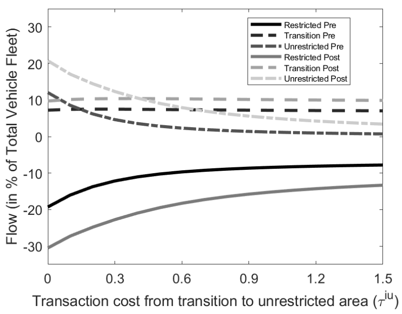

(a) Flow across zones of age-1 vehicles

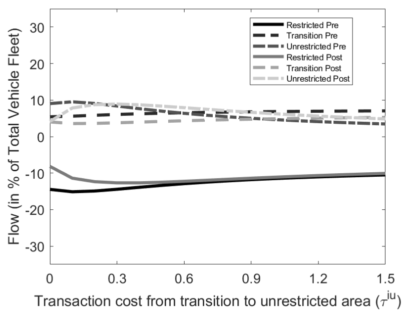

(b) Flow across zones of age-2 vehicles

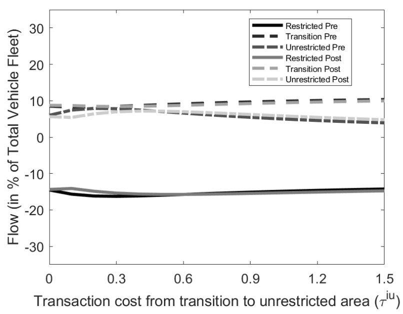

(c) Flow across zones of age-3 vehicles

Figure 13: Pre and post-reform equilibrium flows across zones for different models

Note: This figure presents pre and post-policy equilibrium flows across zones for each vehicle age as a function of transaction costs.

Perhaps the most interesting aspect of these figures is the non-monotonicity introduced by transaction costs on market outcomes. This is easy to see in Panels 13b and 13c, which show how the intervention have opposite effects on the import of age-2 and age-3 cars into the unrestricted zone depending on the level of transaction costs. For low level of transaction costs, the policy reduces the import of these cars into the unrestricted zone but for higher level of transactions costs it actually increases it relative to the pre-policy benchmark.

To explain this non-monotone pattern it helps to look at the price differential between the different models. The policy induces a large increase of new models (age-0) in the restricted zone that car dealers expect to export to the other zones as they age. As seen in Panel 13a, it is particularly large the increase in the export of age-1 cars from the restricted zone to the other two zones for low levels of transaction costs. This large increase in the flow of age-1 cars into the transition and unrestricted zones in turn explains the reduction in the flow of age-2 cars to these same zones: The presence of more age-1 cars has reduced the demand for age-2 cars in those zones. However, as transaction costs increase, the flow of age-1 cars to the unrestricted zone reduces, increasing the demand and flow of age-2 cars into that zone. At some point the flow of age-2 cars also decline with larger transaction costs because there are fewer new cars entering into the restricted zone in the first place, as shown in 12a. The latter figure also shows, although not as noticeably, a non-monotone impact of the policy on the increase of new cars into the restricted zone. This increase reaches its highest level for medium levels of transaction costs.

4.4 Pollution and welfare

We now present our estimates of the impact of the policy that we have considered so far, which is stricter than the actual policy, on both the level of pollution and welfare. Figure 14 summarizes the difference in pollution harm between the pre and post-policy outcomes as a function of transaction costs (τ^{tu}) and the contribution of the transition zone to pollution in the restricted area (α). We find no evidence of emissions leakage. We find savings in pollution costs in the range of 5 to 9%. Nonetheless, the possibility of leakage is present in the figure: pollution savings are lower as α increases.

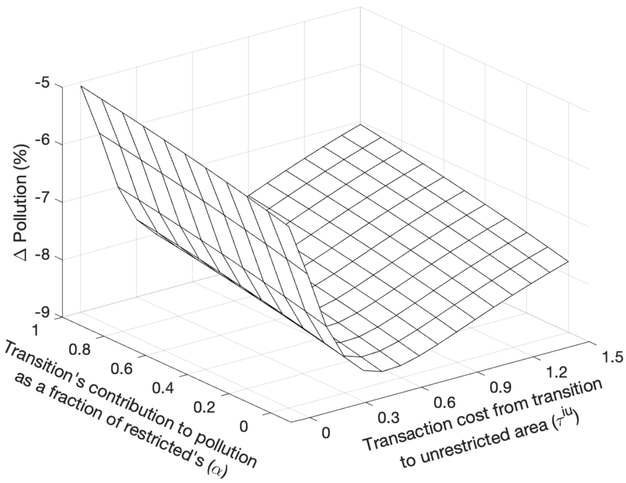

Figure 14: Pollution impact from policy intervention

Note: The picture depicts changes in pollution from pre to post-policy equilibrium levels as a function of transaction costs (τ^{iu}) and the relative contribution to pollution of the intermediate zone (α).

Another interesting aspect of the figure is the non-monotonicity of pollution savings as a function of transaction costs. Pollution savings reach their highest level for intermediate levels of transaction costs which is also when the increase in new cars into the restricted area is at its highest level, as we discussed earlier.

The reason pollution is highest under very low transaction costs is because many individuals in the different zones switch from public transport to (older) cars. As we increase transaction costs from these very low levels the prices of restricted cars increase, since it is more costly to move cars across zones, particularly to the unrestricted zone. These higher prices move low-income individuals back to public transport, leading to less pollution. At some point however, when transaction costs are sufficiently high, an increase in transaction costs lead to more, not less, pollution because restricted cars are now more likely to remain in the restricted zone and its periphery.

Any welfare gain from less pollution must be contrasted with the welfare loss from restricting the use of some cars and moving drivers to a less preferred driving option relative to one they chose in the pre-policy benchmark. Consistent with the pollution figure, Figure 15 also exhibits a non-monotonic pattern as a function of transaction costs. Welfare reaches its highest level for intermediate levels of transaction costs. Actually, welfare would fall below the pre-policy level in the absence of transaction costs and high values of α. According to the empirical evidence, we believe transaction costs to be around 0.2, which would lead to welfare gains between 1 and 2%.[19]

Another interesting aspect of the figure is that welfare increases with α. The reason for this is that when the intermediate zone was not contributing much to pollution (α small), restricting some of their trips to the restricted zone is socially inefficient. In fact, it could be the case that a car should face no restriction when used primarily, but not exclusively, in the transition zone; in other words, we can have a situation where $\theta s_a - h^r e_a < 0$ and $\theta s_a - h^t e_a > 0$.

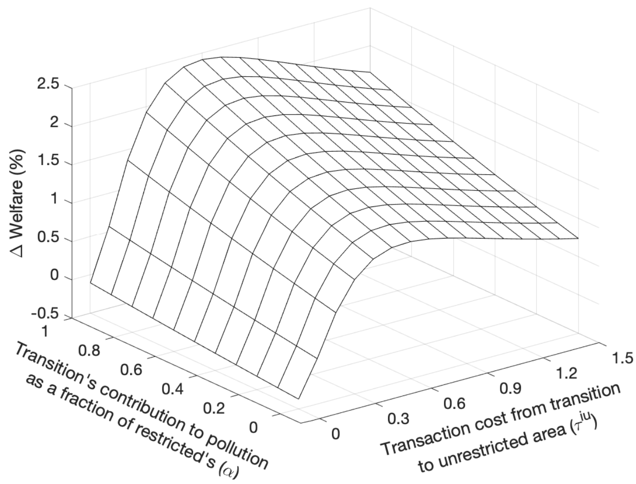

Figure 15: Welfare impact from policy intervention

Note: The picture depicts changes in total welfare from pre to post-policy equilibrium levels as a function of transaction costs (τ^{iu}) and the relative contribution to pollution of the intermediate zone (α).

5 Final remarks

We find empirical evidence that the Santiago's 2017 vintage restriction had a sizable effect on its fleet turnover, increasing the entry of new vehicles and at the same time accelerating the exodus of some middle-age, polluting vehicles to municipalities outside Santiago. Unfortunately, an important fraction these polluting vehicles remained in municipalities located in Santiago's Periphery. These are municipalities that were not affected by the restriction but contribute to Santiago's air quality problems, giving rise to the possibility of emissions leakage.

[19]These welfare gains are comparable to the ones in Barahona et al (2020) for a similar policy. In fact, their scenario 4 in Table 4 reports a welfare gain of 1.5%, the difference between 4,537 and 4,467.

We develop a model of the car market to test for this possibility. An important innovation of the model is the presence of transaction costs, i.e., the cost of moving second-hand cars to distant places. As far the model can tell, these costs may explain why an important fraction of the restricted cars that left Santiago remained in Santiago's periphery despite being more valuable in more distant locations. We fail to find evidence of emissions leakage, at least severe enough to undo the effect of the policy. Interestingly, we find transaction costs to have a non-monotonic impact on emissions, and hence, on welfare. Emissions are lowest (and welfare highest) for intermediate levels of transaction costs. This is when the entry of new vehicles into the restricted zone reaches its highest level relative to the no-intervention benchmark.

Given our focus on the possibility of emissions leakage and corresponding role of transaction costs, we leave some interesting aspects of these vintage restrictions unattended. For example, we are silent on the optimal design of these vintage policies. We do not question whether it was a sensible idea to exempt only cars that were 6 or less years old. Neither we address questions related to the optimal size of the restricted area and how this depends on transaction costs. We hope to address these and other related questions in future research.

References

[1] Adda, J., and R. Cooper (2000). Balladurette and juppette: A discrete analysis of scrapping subsidies, *Journal of Political Economy* 108, 778-806.

[2] Ahlvik, L., and M. Liski (2022), Global Externalities, Local Policies and Firm Selection, *Journal of the European Economic Association* 20, 1231-1275.

[3] AirVisual (2018).World air quality report. IQAir AirVisual, Staad, Switzerland (Available at https://www.airvisual. com/world-most-polluted-cities.)

[4] Alé-Chilet, J., C. Chen, J. Li, and M. Reynaert (2023), Colluding against environmental regulation, working paper, TSE.

[5] Barahona, N., F. Gallego, and J.-P. Montero (2020). Vintage-specific driving restrictions, *Review of Economic Studies* 87, 1646-1682.

[6] Blundell, R., R. Gu, S. Leth-Petersen, H. Low, and C. Meghir (2022), Durables and lemons: Private information and the market for cars, working paper.

[7] Bulnes, J. (2018), Santiago respira: Repercusiones de corto plazo de la nueva restricción vehicular permanente, Master thesis, Pontificia Universidad Católica de Chile.

[8] Davis, L. (2008), The effect of driving restrictions on air quality in Mexico City, *Journal of Political Economy* 116, 38-81.

[9] Eskeland, G. S. and T. Feyzioglu (1997), Rationing can backfire: The day without a car in Mexico City, *The World Bank Economic Review* 11, 383-408.

[10] Fabra, N., Imelda, and J.-P. Montero (2023), Transition technologies in the design of low-emission zones: Theory and evidence from Madrid Central, mimeo

[11] Gallego, F., J.-P. Montero, and C. Salas (2013), The effect of transport policies on car use: Evidence from Latin American cities, *Journal of Public Economics* 107, 47-62.

[12] Galdon-Sanchez, J.E., R. Gil, F. Holub, and G. Uriz-Uharte (2023), Social benefits and private costs of driving restriction policies: The impact of Madrid Central on congestion, pollution, and consumer spending, *Journal of the European Economic Association*, forthcoming.

[13] Gavazza, A., A. Lizzeri, and N. Roketskiy (2014). A quantitative analysis of the used-car market. *American Economic Review* 104, 3668-3700.

[14] Jacobsen, M., J. Sallee, J. Shapiro, and A. Van Benthem (2023), Regulating untaxable externalities: Are vehicle air pollution standards effective and efficient? *Quarterly Journal of Economics*, forthcoming.

[15] Liski, M., J.-P. Montero and W. Ogden (2023), Carbon-zero transition in car markets: Trade frictions and pollution, mimeo.

[16] Martin, R., M. Muûls, L.B. de Preux, and U. Wagner (2014), Industry compensation under relocation risk: A firm-level analysis of the EU Emissions Trading Scheme, *American Economic Review* 104, 2482-2508.

[17] Montero, J.-P., F. Sepúlveda, and L. Basso (2023), Pricing congestion to increase traffic: The case of Bogotá, working paper, Inter-American Development Bank.

[18] Rizzi, L., and C. De La Maza (2017), The external costs of private versus public road transport in the Metropolitan area of Santiago, Chile, *Transportation Research Part A: Policy and Practice* 98, 123-140.

[19] Salgado, E., and O. Mitnik (2021), Spatial and Time Spillovers of Driving Restrictions: Causal Evidence from Lima's Pico y Placa, working paper, Inter-American Development Bank.

[20] Viard, V. B., and S. Fu (2015), The effect of Beijing's driving restrictions on pollution and economic activity, *Journal of Public Economics* 125, 98-115.

[21] Wolff, H. (2014), Keep your clunker in the suburb: Low-emission zones and adoption of green vehicles, *Economic Journal* 124, F481-F512.

[22] Zhang, W., C.-Y.C. Lin Lawell, and V.I. Umanskaya (2017), The effects of license plate-based driving restrictions on air quality: Theory and empirical evidence, *Journal of Environmental Economics and Management* 82, 181-220.

Appendix A Underlying distribution of the data

In this Appendix we look at the underlying distribution of the data. We justify that the log-linear functional form that we use in our regressions describes the data better than letting the quantity to enter linearly. Panel (a) in Figure A1 shows the number of vehicles from vintage 2012 in municipalities in Santiago before and after the 2017 reform. As seen from the figure, all municipalities in Santiago are losing cars vintage 2012. Municipalities that had more cars in the pre-reform period lose more cars than municipalities with fewer cars. In Panel (b), we show the same figure but in logs rather than levels. Once we move to logs, the decrease in log-number of cars seems more even across municipalities. We take this as suggestion that a log-linear functional form describes the data better than a linear functional form. In Panels (c) and (d) of the figure, we show supplementary figures of the underlying distribution of the number of cars and the log-number of cars across municipalities in Santiago. As we can see, the distribution of the number of cars resembles a log-normal distribution and the distribution of the log-number of cars resembles a normal distribution.

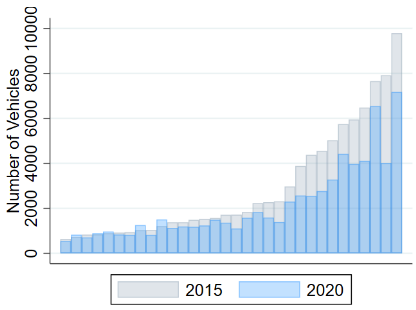

(a) Number of vehicles before and after the reform

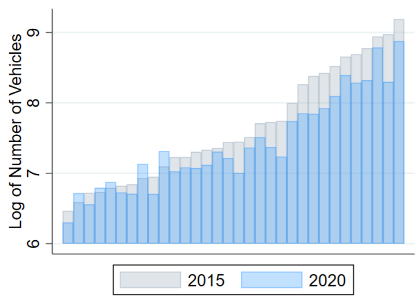

(b) Log-Number of vehicles before and after the reform

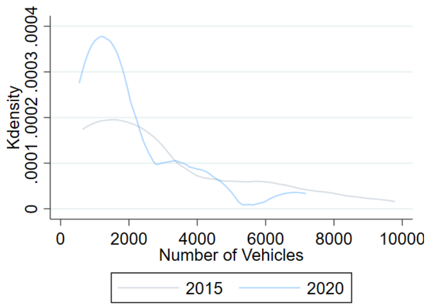

(c) Distribution of the number of vehicles before and after the reform

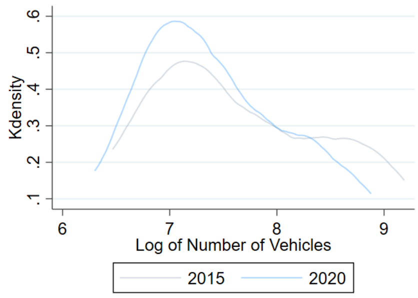

(d) Distribution of the log-number of vehicles before and after the reform

Figure A1: Distribution of 2012 vehicles in Santiago

Note: This figure portrays the distribution in terms of quantity and density of 2012 vehicles in Santiago's municipalities. It differentiates between pre-policy and post-policy periods, presenting data in both raw levels and logarithmic scales.