Designing Low Emission Zones: Theory and Evidence*

Nano Barahona
UC Berkeley & NBER

Natalia Fabra
CEMFI & CEPR

Imelda
Geneva Grad. Institute & CEPR

Juan-Pablo Montero
PUC-Chile & Aalto University

November 7, 2026

Abstract

Low Emission Zones (LEZs), which exclude polluting vehicles from parts of a city, are spreading. Yet their optimal design is poorly understood. Unlike price instruments, LEZs ration access inefficiently, imposing costs on inframarginal drivers of restricted vehicles. Welfare is quasi-convex in LEZ size, yielding bang-bang principles: LEZs should be as large as possible or not implemented at all; cleaner vehicles (e.g., hybrids) could be exempt as transition technologies but later restricted, which may require planner commitment. Using new-car purchases, we find Madrid's LEZ shifted buyers toward cleaner cars, but the surplus losses of switchers and inframarginal drivers exceeded emission savings.

Keywords: Air Pollution, Low Emission Zones, Quasi-Convex Welfare, Commitment.

JEL Codes: D12, H23, L98, Q53, R48.

*Emails: nanobk@berkeley.edu, natalia.fabra@cemfi.es, imelda@graduateinstitute.ch, and jmontero@uc.cl. Michael Duarte, José V. Figueroa, and Álvaro Oballe have provided excellent research assistance. Comments from Felipe Barbieri, Roberto Herrera, Cristian Huse, Emir Kamenica, Anna Russo, Reed Walker, Xavier Fageda, and seminar audiences at the IO+ Conference (University of Chicago), Aalto University, A Toxa Workshop, CEMFI, RIDGE, Simposio de Análisis Económico, Toulouse School of Economics, Universidad de Comillas, ULB, INRAE-CREST, HEC Montréal, MIT-Harvard, UC Berkeley, University of Chile, Yale University, the CEPR Women in Environmental Economics Workshop, and the MADBAR workshop are gratefully acknowledged. Fabra thanks support from the European Research Council (Advanced Grant ENERGY IN TRANSITION, GAP 101142583), and the Spanish Ministry of Science, Innovation and Universities / State Research Agency (MICIU/AEI/10.13039/501100011033). Montero thanks support from ANID-Fondecyt (Grant No. 1250007).

1 Introduction

Road transport is a leading source of greenhouse gases and local air pollutants. The associated health, economic, and societal consequences are substantial and well documented (OECD 2014; World Health Organization 2025). An increasingly popular policy tool to handle these external costs is the creation of Low Emission Zones (LEZs): urban areas that restrict access to vehicles that fail to meet certain environmental standards. Introduced in 2008 in several German cities, LEZs are now found all over Europe and in many cities around the world, including Shanghai, Beijing, Haifa, Seoul, and Jakarta.[1] Even some U.S. cities have begun experimenting with more targeted versions of LEZs, particularly for freight and last-mile delivery.

An often overlooked aspect of LEZs is the significant variation in their design, particularly concerning the geographic scope of the covered area and the types of restricted vehicles. For example, the LEZ in Berlin, implemented in 2008, spans 88 km^2, restricts some, but not all, gasoline and diesel cars, and imposes different geographic restrictions among its restricted cars. In contrast, the LEZ in Madrid, known as Madrid Central and implemented a decade later, covers a much smaller area, of just 5 km^2, but restricts all gasoline and diesel cars, making no geographic distinction among them.

Such wide variation raises important policy-design questions, which motivate this paper: What is the optimal LEZ size, and should some cars be restricted more, or later, than others? A central goal of a LEZ is to achieve sustained reductions in air pollution by creating incentives for drivers to switch to cleaner modes of transportation, whether cleaner cars or public transit (Ayuntamiento de Madrid 2018, arts. 22–23, p. 139).[2] Given this goal, we show that the answers to the above questions follow a bang-bang logic: a LEZ should be either as large as politically feasible or not implemented at all. Moreover, when less-polluting vehicles such as hybrids are to be restricted, it is optimal to restrict them in the same way as other restricted cars, but later.

Using Madrid's LEZ as both motivation and evidence, we initiate our analysis by assessing its impact on the adoption of hybrids and electric vehicles (EVs), which, along with liquefied petroleum gas or compressed natural gas vehicles (LPG/CNG), are exempt from the restriction. Looking at sales of new cars and comparing Madrid to other cities in Spain, we estimate a synthetic difference-in-differences model (Arkhangelsky et al. 2021) and show that Madrid's LEZ was effective in accelerating the adoption of hybrids and LPG/CNG vehicles, and to a lesser extent EVs, while also reducing the sales of new

[1] By 2024, 363 LEZs were active in 15 European countries, including 175 in Italy, 77 in Germany, and 36 in Spain (Delgado-Lindeman et al. 2025).

[2] This goal is consistent with EU policy, under which LEZs are an important instrument for cities seeking to comply with EU air-quality standards (Ayuntamiento de Madrid 2018, p. 134). It is also consistent with the findings of Bernardo, Fageda, and Flores-Fillol (2021), who, using data from several European LEZs over the period 2008–2016, find that LEZs reduced pollution but not congestion, which they attribute to extensive-margin adjustments toward cleaner cars rather than intensive-margin adjustments, i.e., reductions in vehicle use.

gasoline and diesel cars.

These findings raise policy questions that reach beyond Madrid's LEZ design. Was it wise to fully exempt LPG/CNG and hybrid cars? Two alternatives arise. A *spatial* alternative would restrict these cars less stringently than gasoline and diesel cars—barring them from a smaller zone, as in Berlin's tiered design. A *temporal* alternative would treat them as transition technologies: unrestricted for now, but with an announced timetable under which they later face the same restrictions as gasoline and diesel cars. Both alternatives would make EVs more attractive relative to them, but through different channels, and at different costs to their current owners and to drivers thinking of switching away from gasoline. So the natural question is: How do the two alternatives compare, and when is either preferable to Madrid's full exemption of LPG/CNG and hybrids?

We shed light on these and related questions with a simple dynamic model of three horizontally differentiated car types (gasoline, hybrids, EVs) that also differ in their emission rates. Hybrids—which represent any intermediate technology such as LPG/CNG—pollute less than gasoline cars but more than EVs. The model formalizes the bang-bang rationale: all restricted vehicles should be treated equally, and the LEZ should be either as large as possible or not implemented at all if it falls below a minimum scale. Within this rationale, hybrids may still qualify as a transition technology—kept as unrestricted as EVs for some time, and then as restricted as gasoline cars. This transition-technology design, however, may require commitment from the planner.

Unlike price-based instruments, such as taxes and tradable quotas, which operate under an efficient rationing rule, LEZs operate under an inefficient (e.g., proportional) rationing rule, leading to quasi-convex welfare functions.[3] To see this, imagine a social planner who introduces a small pollution tax on dirty vehicles. The welfare impact of this tax is confined to the marginal drivers who switch to cleaner cars in response to the tax. Importantly, social welfare is unaffected by the tax impact on inframarginal drivers—those who do not switch. Initially, the welfare loss to switchers from moving away from their preferred (i.e., no-intervention) car is smaller than the social welfare gain from reduced emissions. The planner finds it optimal to increase the tax up to its Pigouvian level, where the (concave) welfare function reaches its maximum level. Importantly, unless set far above its Pigouvian level, a tax always increases social welfare, no matter how small.[4]

Because LEZs involve inefficient rationing, they differ fundamentally from price-based instruments; under some conditions, even the best LEZ may be worse than having no LEZ. To see this, consider a small LEZ that restricts only gasoline cars. Unlike a tax, a LEZ imposes a welfare cost on every inframarginal driver who continues driving gasoline cars. These individuals will have to change their transport mode (e.g., to public tran-

[3]See Tirole (1988) for more on rationing rules.

[4]Shadow costs of public funds may affect the optimal tax in either direction according to a Ramsey formula but never eliminate it (Laffont and Tirole 1996).

sit), reroute around the LEZ, or cancel trips altogether. If the number of inframarginal individuals is large, the welfare losses borne by them can be large relative to the net welfare gains from the few marginal individuals who switch from gasoline cars to hybrids or EVs. This makes a small LEZ a potentially bad idea—unless it can be expanded above a minimum scale, large enough to leave few inframarginal drivers restricted.

This minimum scale is the LEZ size that achieves the same welfare as the no-intervention outcome. We show that when pollution harm is not too large, or substitution toward cleaner cars is weak, the minimum scale may not exist, so implementing no LEZ is the best option. In contrast, when pollution harm is relatively large, and substitution is strong, the minimum scale can be small, even zero, meaning that almost any LEZ would be beneficial. Even then, the quasi-convexity of the welfare function calls for the largest possible LEZ.

This bang-bang logic—to reduce the number of inframarginal consumers of restricted cars as much as possible—also explains why all restricted cars should be treated equally, i.e., subject to the largest politically feasible LEZ. We also show that, in many cases, hybrids should be treated as a *transition technology*: initially exempt from the restrictions, but subject to them at a later stage. Anticipating the declining value of hybrids, and provided the planner can commit to the announced future restrictions, some individuals may choose to switch away from hybrids before the restrictions take effect.[5] Consequently, fewer hybrids remain in circulation by the time the restriction arrives, reducing the cost of imposing it. In this sense, a LEZ with a transition technology works more like a price-based instrument: it affects consumer choices at the margin without destroying value for inframarginal consumers before the restriction becomes active.[6]

We conduct an empirical evaluation of Madrid's LEZ, focusing on new-car adoption and welfare. Because car prices and preferences cannot be projected far ahead, we estimate impacts around policy implementation and ask whether, at that time, the necessary conditions for even considering a LEZ were met. That is, we ask whether there was a minimum scale above which some LEZ design could improve welfare. We focus on households buying a new car, since that is where fleet composition changes: a second-hand purchase mostly reallocates the existing stock. Our welfare estimates account for the private costs of switching away from households' preferred options, the utility loss borne by those who continue to buy restricted cars, and the lifetime externalities of the chosen vehicles.

In the spirit of Berry, Levinsohn, and Pakes (1995), we estimate demand for new cars using registration data covering new cars sold in Spain from 2015 to 2019, together with

[5] Some switch to electric vehicles, while others move to gasoline cars; the latter case is less relevant, as gasoline cars are already restricted.

[6] Our definition of a transition technology differs from alternative notions, which are often based on the idea that electric vehicles remain too expensive or unfamiliar to serve as a viable outside option for many gasoline drivers. Under that view, hybrids should remain unrestricted until these concerns dissipate. While our definition is not in conflict with that view, it is fundamentally different.

their prices and characteristics. Consumers' preferences for cars, and their willingness to pay to avoid the LEZ, vary with income and other consumer characteristics. Consistent with our motivating evidence, we find that some consumers switch from gasoline and diesel cars toward hybrids, LPG/CNG cars, and EVs. But the resulting externality savings fall well short of the private losses: consumer surplus falls substantially more than the reduction in externality costs. The private cost of tightening the restriction exceeds the emission benefit throughout the range of intensities we evaluate, so a larger LEZ would not have helped; nor would extending the restrictions to LPG/CNG vehicles and hybrids, which raises the share of gasoline and diesel cars without materially increasing EV adoption.

These negative results could be reversed under a LEZ design that also restricts LPG/CNG and hybrids, if the following conditions are met: a substantially higher social cost of emissions combined with a larger restricted zone, or stronger substitution toward EVs. All this suggests that, at least in Madrid, the necessary conditions for considering a LEZ in the first place were not satisfied. It also points to a broader conclusion: for a LEZ to improve welfare, there must be fewer drivers buying gasoline and diesel cars, and the barriers to adopting cleaner ones must be lower. For EVs, these barriers include limited access to charging, short perceived driving range, and long charging times (Rosales-Tristancho et al. 2022).

It is important to point out that these negative results abstract from any local benefits that a LEZ may bring to residents and businesses within the restricted area, such as safer streets and improved public space. This raises the broader question of who ultimately determines LEZ design. This paper takes the stance of a planner who cares about overall welfare. A narrower perspective is that of a decision maker who cares mostly about the welfare of residents in the restricted area. A local constituency may evaluate a small LEZ quite differently, and possibly quite favorably, than such a planner. The wide variation in LEZ size that we observe—with Berlin and Madrid providing good examples—may partly reflect these different perspectives. Our contribution should be understood from the broader perspective: how best to use the LEZ instrument when the objective is overall welfare.[7]

Our paper contributes to four strands of the literature. First, we contribute to the vast literature on car markets. Closest to our work are the papers by Adda and Cooper (2000) and Gillingham et al. (2022), which pay explicit attention to both new and second-hand units and market dynamics. However, we depart from both papers in important ways. Adda and Cooper (2000) consider a single type of car, abstracting from any horizontal differentiation—both in preferences and pollution rates—which lies at the core of our framework. Gillingham et al. (2022), by contrast, offer a rich framework with various degrees of differentiation (in preferences and pollution) but, by focusing on sta-

[7] And in doing so, we completely abstract from the political-economy question of why cities opt for LEZs as opposed to more efficient instruments.

tionary equilibria, abstract from aggregate shocks that may evolve over time and generate transition dynamics. Unlike both papers, we study how a market with differentiated cars adjusts over time to a LEZ. This adjustment matters for policy design, especially when hybrids can be treated as a transition technology.[8]

Second, we contribute to the growing literature evaluating the performance of LEZ programs around the world.[9] Several papers document the positive health impacts of LEZs, driven by reductions in local air pollution (e.g., Klauber et al. 2024; Beshir and Fichera 2022; Margaryan 2021; Gehrsitz 2017), although these benefits come at the cost of restricted mobility (Sarmiento, Wägner, and Zaklan 2023). Others provide reduced-form evidence on the effects of LEZs on traffic, finding reductions in congestion within the restricted area (Galdon-Sanchez et al. 2023) that are often temporary, negative spillovers to adjacent areas (Tassinari 2024), or simply no effects (Bernardo, Fageda, and Flores-Fillol 2021; Garson et al. 2026).

Only a few papers study the effects of LEZs on the adoption of lower-emission vehicles, which are central to the policy-design questions at the core of our analysis. An early exception is Wolff (2014), who shows that German LEZs induced residents living close to restricted areas to increase their adoption of low-emission vehicles. More recently, also in the context of German LEZs, Aydin, Gehrsitz, and Traxler (2024) find that highly polluting vehicles were replaced by cleaner gasoline cars.[10] These empirical findings are consistent with the theoretical model we develop in this paper, as well as with the evidence on Madrid's LEZ that motivates it. However, none of these papers evaluate whether the effects of induced low-emission vehicle adoption are large enough to offset the welfare costs associated with restricted mobility and distorted choices.

Third, we contribute to the extensive literature comparing the performance of different policy instruments, typically categorized as either price-based (e.g., taxes, tradable quotas) or command-and-control (e.g., emissions and technology standards). While LEZs clearly fall within the command-and-control (CAC) category, they differ from other instruments in this group due to the bang-bang rationale underlying their optimal design. Under most standard CAC instruments, the "marginal-inframarginal" tradeoff at the core of a LEZ design does not arise. As a result, the welfare function remains concave, and the bang-bang rationale does not apply (see, e.g., Baumol and Oates 1988; Ellerman

[8] Our theory also connects to the work of Gavazza, Lizzeri, and Roketskiy (2014) on trade frictions in the second-hand market. We abstract from vertical differentiation, which is essential to generate trade in their model; nevertheless, second-hand-market trade may still emerge in our setting with the introduction of a LEZ or a tax.

[9] LEZs impose differential restrictions across vehicles, unlike some early programs that applied uniformly to all cars and proved counterproductive (e.g., Davis 2008; Gallego, Montero, and Salas 2013). Differential treatment also features in the vintage-specific restrictions studied by Barahona, Gallego, and Montero (2020).

[10] Other papers have analyzed the impacts of alternative instruments on fleet composition, including congestion pricing (e.g., Isaksen and Johansen 2021; Börjesson et al. 2012).

et al. 2000).[11]

Relatedly, Almagro et al. (2024), Durrmeyer and Martínez (2026), Ater et al. (2026), and Cook and Li (2025) quantify the welfare effects of different urban-transportation policies—road pricing, transit fares, and service frequencies in Chicago; tolls and license-plate bans in Paris; usage-based congestion pricing in Tel Aviv; and dynamically priced highway lanes in Seattle, respectively. All four papers hold the private vehicle fleet fixed and focus on short-run margins. Cook et al. (2025) find that, over its first 18 months, New York City's congestion pricing raised traffic speeds but had little effect on air quality. By contrast, we study a policy instrument whose central welfare tradeoff operates through a long-run margin: the adoption of lower-emission vehicles.

Finally, we contribute to the literature on time inconsistency in environmental policy, although through a distinctive channel: in our setting, the commitment problem does not arise under taxes—it arises only under a LEZ design that involves transition technologies. This stands in contrast with the standard framing (e.g., Laffont and Tirole 1996; Helm, Hepburn, and Mash 2003), where time inconsistency stems from the planner's ex post incentive to renege on ex ante (optimal) regulatory commitments—whether taxes or standards—once agents have invested in irreversible abatement technologies or R&D. Ours does not follow this expropriation logic but, rather, a pure discounting one: when hybrids are not so dirty, the initial announcement of their future restriction fails to induce enough switching toward cleaner options to reduce the number of inframarginal hybrid users sufficiently to justify restricting them ex post.[12]

The remainder of the paper is organized as follows. Section 2 provides background on Madrid Central, describes the data sources, and presents motivating evidence of its impact on green vehicle adoption. Section 3 develops a dynamic model of car choice and examines the implications for the optimal LEZ design. Section 4 presents the structural estimation, and Section 5 conducts the counterfactual analysis. Section 6 concludes. Additional results and proofs are relegated to the Appendix.

2 Madrid Central

The Madrid City Council implemented traffic restrictions on polluting vehicles, limiting access to a 472-hectare area (approximately 5 km^2) covering the historic center and the city's main commercial and leisure districts—known as Madrid Central (MC).[13] The

[11] One exception in CAC regulation that has some of our bang-bang flavor is the vintage-specific restrictions analyzed by Barahona, Gallego, and Montero (2020), where the optimal design involves a vintage threshold that separates cars subject to no restriction from those subject to the maximum feasible restriction.

[12] It also differs from other commitment problems involving durable goods (e.g., Coase 1972; Kahn 1986) in that the no-commitment LEZ solution need not be unique.

[13] See MC's map in Figure A.1 in the Appendix.

policy was announced in May 2018, and was implemented from November 2018 onward.[14]

The primary motivation for the regulation was to reduce nitrogen dioxide (NO_2) emissions, as Madrid had been persistently violating European Union air-quality limits. The European Commission had issued a final warning to Spain, and the city responded by introducing MC as part of an updated clean air plan. This paper examines whether MC induced the adoption of cleaner vehicles and whether the resulting environmental benefits were large enough to offset both the private welfare losses from shifting consumers away from their preferred vehicle choices and the disutility associated with restricted access for those who kept on buying polluting vehicles.

Under the 2018 MC regulation, vehicles were labeled according to their emissions characteristics. Vehicles with ZERO or ECO labels—covering qualifying LPG/CNG, hybrid, and electric vehicles—could enter the restricted zone freely.[15] For non-residents, conventional gasoline and diesel cars with an emissions label could enter only to park in an authorized public or private facility, while unlabeled cars required a specific exemption. Our empirical analysis reports outcomes for four vehicle groups: gasoline/diesel cars, including both labeled and unlabeled vehicles; exempt LPG (liquefied petroleum gas) and CNG (compressed natural gas) vehicles; hybrids, including plug-in and non-plug-in models; and electric vehicles (EVs), including battery-electric and fuel-cell vehicles. Residents of the MC area could continue to circulate with their existing vehicles regardless of their emissions label. Special access rules applied to residents' guests, people with reduced mobility, taxis, ride-hailing services, and delivery vehicles.

2.1 Data

We use several data sources. Vehicle registrations come from the Spanish Directorate-General for Traffic and cover every new and existing vehicle registered in Spain between 2015 and 2019.[16] For each record, we observe the vehicle's emissions standard, make-model, engine size, fuel type, powertrain, and tailpipe CO_2, along with the zip code where it is registered or moved to. These records distinguish three events: new-car registrations, second-hand transfers,[17] and scrappage.[18] We aggregate observations to the

[14] Formal enforcement, with €90 fines per violation, began on March 16, 2019. Until then, violators were notified by mail without penalties. While March 2019 is the relevant treatment date for analyzing car usage, the May 2018 announcement serves as the key reference for assessing impacts on vehicle (i.e., durable-good) adoption—the primary focus of this paper.

[15] ZERO covers battery-electric, fuel-cell, and extended-range electric vehicles, as well as plug-in hybrids with an electric range of at least 40 km. ECO covers non-plug-in hybrids, plug-in hybrids with shorter electric ranges, and natural gas or LPG vehicles, provided they meet the C-label emissions standards. For conventional gasoline and diesel cars, C covers gasoline vehicles meeting Euro 4–6 standards and diesel vehicles meeting Euro 6; B covers gasoline vehicles meeting Euro 3 and diesel vehicles meeting Euro 4–5. Vehicles that do not meet these requirements remain unlabeled.

[16] We restrict the sample to natural persons and exclude leasing from new registrations.

[17] We focus on incoming transfers between provinces, for which we observe the destination zip code where the car is registered; we observe only the province of origin.

[18] These are old cars no longer registered in Spain, either scrapped or moved abroad (e.g., to North Africa).

quarterly zip-code level, dropping the six zip codes that make up Madrid Central, and normalize registrations by zip-code population, expressing them per 10,000 residents.[19] Retail prices come from JATO's database. Demographic and income data come from MBR MarketData at the zip-code level for 2017. Finally, annual steel prices, built from daily New York Mercantile Exchange quotes, are used in Section 4 as an instrument for car prices.

Table 1 summarizes the sample. Madrid's zip codes are dense and comparatively rich: 2,092 households per square kilometer on average, €16,066 of purchasing power per capita, and 26% of households in the top income quintile, against 468 households per square kilometer, €11,917, and 12% across the rest of Spain. Madrid also records 3.43 hybrid and 0.15 EV registrations per 10,000 residents each quarter, against 1.07 and 0.02 in the rest of the country, respectively.

2.2 Comparing Madrid to similar zip codes

To evaluate Madrid Central, we compare Madrid's zip codes with similar zip codes outside Madrid province. Of the 10,520 zip codes outside Madrid province, we choose the closest on the following demographics: log household density, purchasing power per capita, the shares of households in each income quintile, and average household size. We select them using the K-nearest-neighbor matching method with replacement based on the Mahalanobis distance (Rubin 1980; Jann 2020). For each treated Madrid zip code $i \in \mathcal{T}$, the matched set $\mathcal{N}_K(i)$ is the subset S of the control pool $\mathcal{C}$ with $|S| = K$ that minimizes the Mahalanobis distance between i and the controls in S:

$$\mathcal{N}_K(i) = \arg \min_{\substack{S \subset \mathcal{C} \\ |S|=K}} \sum_{j \in S} (\mathbf{x}_i - \mathbf{x}_j)^\top \boldsymbol{\Sigma}^{-1} (\mathbf{x}_i - \mathbf{x}_j), \tag{1}$$

where $\mathbf{x}_i$ and $\mathbf{x}_j$ collect these variables for treated zip code i and control zip code j, and $\boldsymbol{\Sigma}$ is their covariance matrix. We consider four specifications with $K \in \{5, 10, 25, 100\}$ neighbors.

Columns (3)–(6) of Table 1 report the same statistics for the matched groups. The fewer neighbors we retain, the closer they are to Madrid: household density averages 780 at 100-nn and 1,499 at 5-nn, against 2,092 households per square kilometer in Madrid, and purchasing power per capita €13,306 and €14,804 against €16,066. The top-quintile share follows the same pattern. On registrations, Madrid records 3.43 hybrids and 0.15 EVs per 10,000 residents each quarter, compared with 1.88 and 0.06 in the five-neighbor group, and 29.18 second-hand transfers versus 46.10.

Figure 1 plots quarterly averages per 10,000 residents for Madrid and the 5- and 100-nearest-neighbor comparison groups. Panel (a) shows registrations of gasoline and

[19] Residents inside the LEZ are not subject to its restrictions, and the area is too small to serve as a control group, which is why we drop them.

Table 1: Summary statistics for Madrid and comparison zip codes

	(1) Madrid	(2) Non-Madrid	(3) 100-nn	(4) 25-nn	(5) 10-nn	(6) 5-nn
Panel A: New cars (per 10,000 residents)						
Total	43.19	34.00	36.96	38.68	39.14	39.48
	(42.57)	(45.11)	(41.67)	(38.59)	(36.84)	(34.44)
Gasoline/Diesel	37.78	32.47	35.11	36.59	36.88	37.13
	(38.12)	(43.50)	(40.06)	(37.08)	(35.28)	(32.81)
LPG	1.08	0.12	0.17	0.21	0.23	0.25
	(1.88)	(0.79)	(0.86)	(0.92)	(0.93)	(0.92)
Hybrids	3.43	1.07	1.37	1.62	1.77	1.88
	(5.15)	(3.77)	(3.86)	(3.87)	(3.84)	(3.84)
EVs	0.15	0.02	0.03	0.04	0.05	0.06
	(0.42)	(0.17)	(0.21)	(0.25)	(0.27)	(0.28)
Panel B: Scrappage (per 10,000 residents)						
Total	30.27	39.38	38.46	37.55	36.47	35.25
	(31.90)	(58.88)	(53.21)	(48.65)	(45.34)	(40.28)
Gasoline/Diesel	30.24	39.38	38.45	37.55	36.46	35.24
	(31.90)	(58.88)	(53.21)	(48.64)	(45.34)	(40.28)
Panel C: Second-hand transfers (per 10,000 residents)						
Total	29.18	57.69	54.50	51.32	48.31	46.10
	(51.52)	(77.48)	(69.60)	(63.55)	(58.65)	(54.15)
Gasoline/Diesel	28.83	57.46	54.23	51.03	48.01	45.78
	(51.39)	(77.33)	(69.43)	(63.35)	(58.45)	(53.86)
Panel D: Zip code characteristics						
Log household density	5.29	2.81	3.50	4.03	4.31	4.59
	(2.59)	(2.31)	(2.48)	(2.52)	(2.56)	(2.59)
Purchasing power (000s € per capita)	16.066	11.917	13.306	14.065	14.483	14.804
	(4.117)	(3.116)	(3.070)	(3.179)	(3.276)	(3.429)
Share HH quintile 1 (lowest)	0.16	0.28	0.23	0.20	0.19	0.18
	(0.08)	(0.12)	(0.09)	(0.08)	(0.08)	(0.08)
Share HH quintile 2	0.18	0.22	0.21	0.20	0.20	0.19
	(0.05)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)
Share HH quintile 3	0.20	0.20	0.21	0.21	0.20	0.20
	(0.03)	(0.04)	(0.02)	(0.02)	(0.02)	(0.02)
Share HH quintile 4	0.21	0.18	0.20	0.21	0.21	0.21
	(0.03)	(0.05)	(0.03)	(0.03)	(0.03)	(0.03)
Share HH quintile 5 (highest)	0.26	0.12	0.16	0.18	0.20	0.21
	(0.15)	(0.09)	(0.09)	(0.10)	(0.11)	(0.11)
Avg. household size	2.61	2.44	2.48	2.50	2.52	2.54
	(0.34)	(0.31)	(0.26)	(0.25)	(0.25)	(0.25)
Zip codes	290	10,520	5,496	2,811	1,571	948

Notes: The table presents summary statistics of quarterly car registrations per 10,000 residents per zip code, divided into three types: new cars, scrappage, and second-hand car transfers. Column (1) reports the statistics for zip codes in Madrid, column (2) reports those for all non-Madrid zip codes, while columns (3)–(6) report those for the 100, 25, 10, and 5 nearest neighbors (nn). We report the mean and standard deviation in parentheses, alongside the number of zip codes. LPG denotes LPG/CNG. Outcomes are winsorized at the 99.5th percentile over the full panel.

diesel vehicles, which fall faster in Madrid than in the comparison groups after the LEZ announcement. Panels (b), (c), and (d) show registrations of LPG/CNG vehicles, hybrid vehicles, and EVs, respectively, which rise faster in Madrid than in the control zip codes, even before the announcement. Panels (e) and (f) show scrappage and second-hand

transfers.

Across several of these outcomes, the pre-announcement trends for Madrid differ from those in the control zip codes, making a simple comparison of their paths hard to interpret as causal. To recover the causal impact of the LEZ on each of these outcomes, we use the synthetic difference-in-differences (SDID) estimator of Arkhangelsky et al. (2021). SDID relaxes the requirement that Madrid and the unweighted control zip codes follow parallel pre-trends. Instead, it constructs a synthetic control for the average zip code in Madrid by reweighting the matched untreated zip codes and time periods so that the pre-treatment path of the synthetic control closely matches that of Madrid. Under the weaker assumption that such a reweighted control group follows parallel trends with Madrid in the absence of the LEZ, the SDID method delivers a consistent DiD-type estimate of the policy effect.

The SDID estimator combines insights from standard difference-in-differences and synthetic control methods to estimate the average treatment effect $\hat{\tau}^{\text{sdid}}$ through a weighted two-way fixed effects regression:

$$(\hat{\tau}^{\text{sdid}}, \hat{\mu}, \hat{\alpha}, \hat{\beta}) = \arg\min_{\tau,\mu,\alpha,\beta} \left(\sum_{i=1}^{N} \sum_{t=1}^{T} \left(Y_{it} - \mu - \alpha_i - \beta_t - W_{it}\tau\right)^2 \hat{\omega}_i^{\text{sdid}} \hat{\lambda}_t^{\text{sdid}} \right), \qquad (2)$$

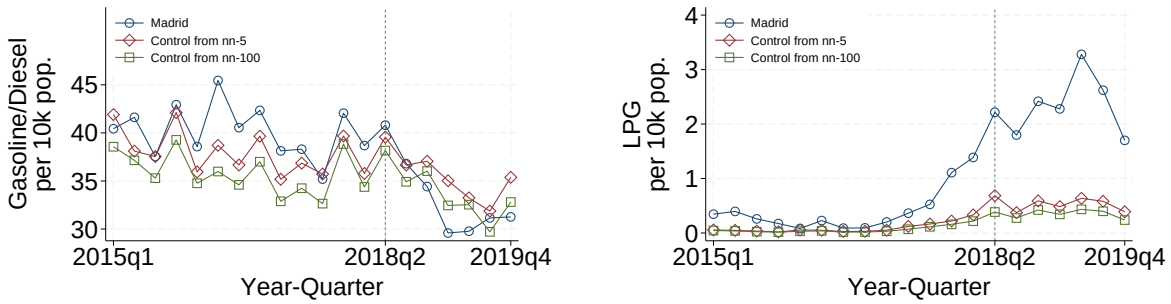
where Y_{it} denotes the outcome in zip code i and quarter t, measured per 10,000 residents. We consider new registrations of gasoline/diesel cars, LPG/CNG, hybrids, and EVs, as well as total second-hand transfers and total scrappage. The terms α_i and β_t represent zip code and quarter fixed effects; W_{it} equals 1 for Madrid province zip codes in post-treatment quarters and 0 otherwise.

The estimator employs two sets of data-driven weights: unit weights $\hat{\omega}_i^{\text{sdid}}$ reweight control units to match pre-treatment trends of treated units, while time weights $\hat{\lambda}_t^{\text{sdid}}$ balance pre- and post-treatment periods to improve counterfactual prediction. This dual weighting reduces reliance on strict parallel trends assumptions. We treat Madrid zip codes as the treated units. The donor pool is the matched set of equation (1). Our preferred specification—presented in Table 2, column (4)—uses 5 neighbors. Standard errors are bootstrapped using 100 replications.[20]

2.3 Impact on new-car registrations

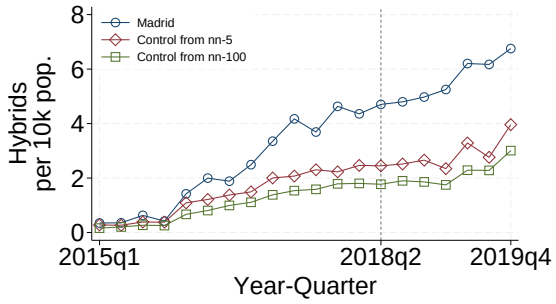
Figure 2 presents the core results from our SDID analysis on new-car registrations. The graphs display the quarterly trends for Madrid zip codes (the treated group) against their synthetic counterparts, which are weighted averages of control zip codes from the rest

[20] As a robustness check, we also add covariates as control variables following Arkhangelsky et al. (2021); results are virtually unchanged (Table A.1, column (5)), as expected for time-invariant controls. Our conclusions also hold under the placebo-based inference of their Algorithm 4, which reassigns treatment to donor-pool units and uses the dispersion of the resulting placebo estimates for inference (Appendix A, Table A.1).

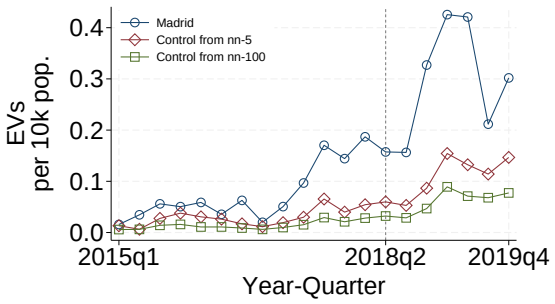


(a) New gasoline/diesel registrations

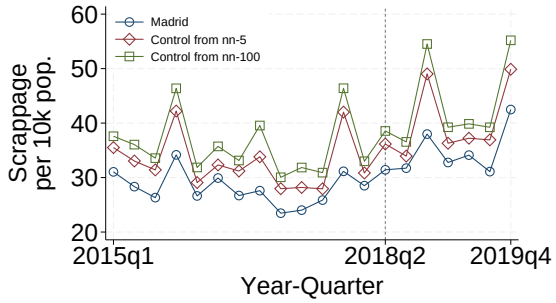
(b) New LPG/CNG registrations



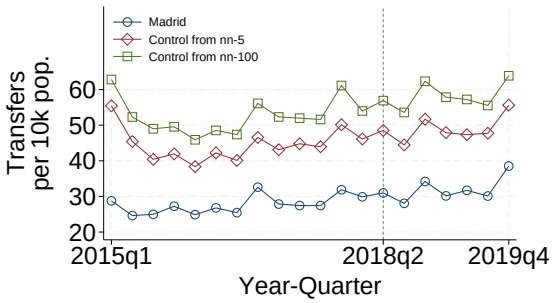
(c) New hybrid registrations



(d) New EV registrations



(e) Scrappage



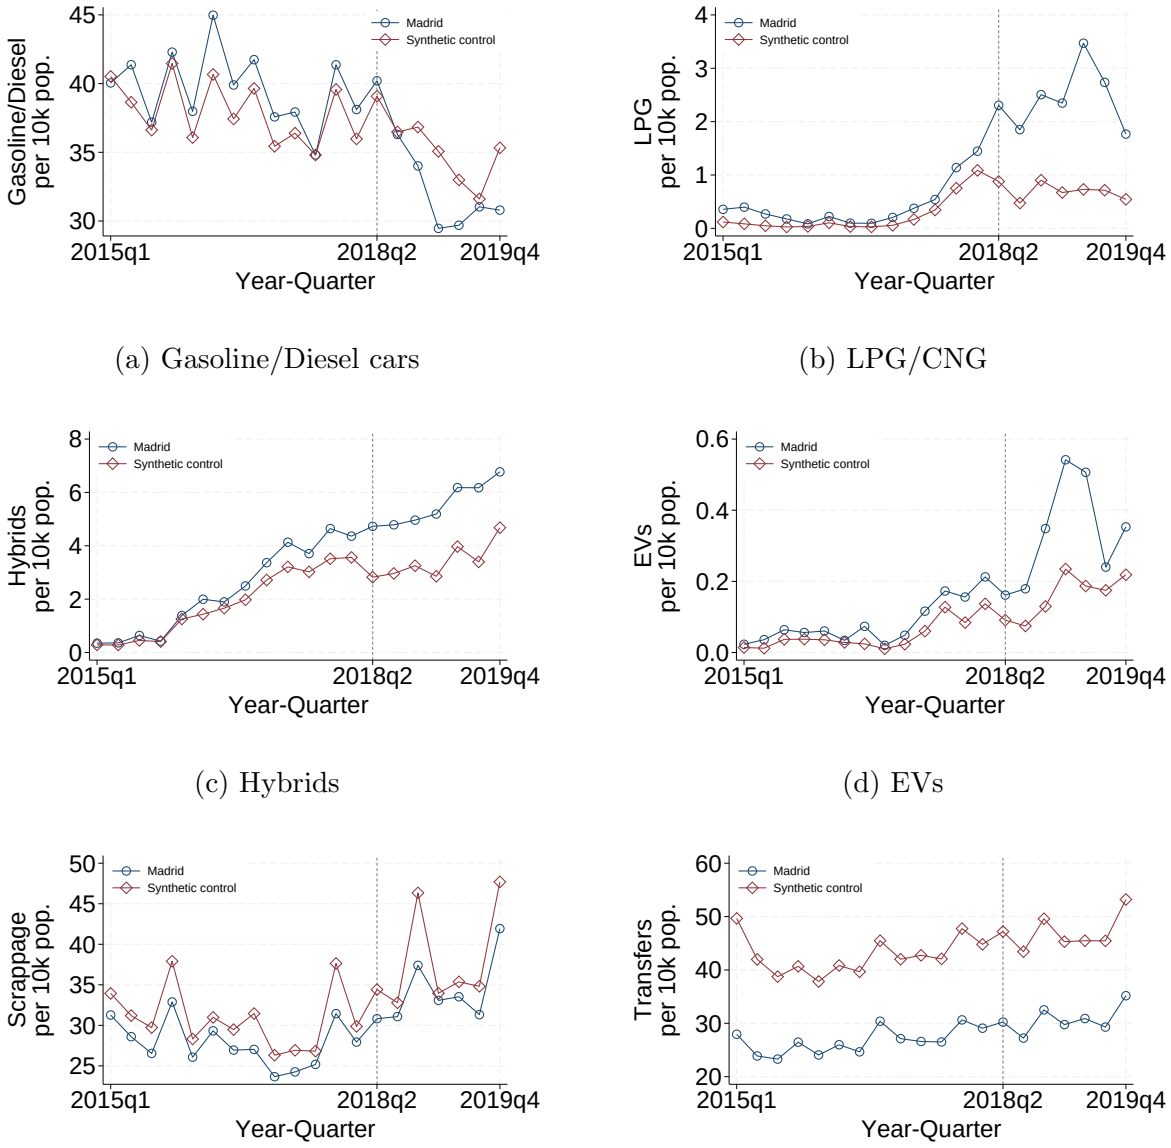
(f) Second-hand transfers

Figure 1: Quarterly average for new registrations, scrappage, and second-hand transfers

Notes: These figures show quarterly trends comparing Madrid with its 5- and 100-nearest-neighbor groups: (a) new gasoline/diesel vehicle registrations, (b) new LPG/CNG registrations, (c) new hybrid registrations, (d) new EV registrations, (e) vehicle scrappage, and (f) second-hand vehicle transfers. All outcomes are measured per 10,000 residents per zip code. The vertical line marks the Madrid Central announcement in May 2018.

of Spain. For all six outcomes—new gasoline/diesel cars, LPG/CNG, hybrids, EVs, scrappage, and second-hand transfers (which are mostly gasoline/diesel cars)—the synthetic control group tracks the trends in Madrid during the pre-treatment period. In contrast, after the announcement of Madrid Central, a clear divergence emerges for new registrations. On the one hand, registrations for LPG/CNG (Panel (b)), hybrids (Panel (c)), and EVs (Panel (d)) in Madrid rise sharply above their synthetic controls. On the other

hand, registrations for polluting gasoline/diesel cars (Panel (a)) fall relatively more in the treated zip codes than in the control ones.



(a) Gasoline/Diesel cars

(b) LPG/CNG

(c) Hybrids

(d) EVs

(e) Scrappage

(f) Second-hand transfers

Figure 2: Madrid zip codes vs. their synthetic zip codes

Notes: Panel (a) shows results for gasoline/diesel cars, Panels (b)–(d) for LPG/CNG, hybrids, and EVs, and Panels (e) and (f) for scrappage and second-hand transfers, using specification (4) in Table 2 (results are very similar across specifications). The vertical line marks the announcement of Madrid Central in May 2018. Panels show SDID estimates with synthetic controls. Treated units are matched zip codes in Madrid province, outside Madrid Central; controls are the synthetic-weight combinations of non-Madrid zip codes picked as the control group according to equation (1).

Table 2 quantifies the average post-treatment effects across four SDID specifications with different donor sets. The results confirm the visual evidence from the figures, showing that Madrid Central led to a significant reduction in new gasoline/diesel cars, alongside a statistically significant increase in new registrations of LPG/CNG vehicles, hybrids, and

EVs. In the preferred specification with 5-nearest-neighbor controls without covariates, column (4), the policy is associated with a quarterly reduction of 3.94 new gasoline/diesel registrations per 10,000 residents and an increase of about 1.49 LPG/CNG, 1.59 hybrid, and 0.14 EV registrations. The estimates are qualitatively similar across the alternative specifications using 100, 25, and 10 nearest neighbors.

Compared to registrations in the announcement quarter, the estimated effects correspond to a reduction of 10 percent for gasoline/diesel cars and to increases of 67 percent for LPG/CNG vehicles, 34 percent for hybrids, and 87 percent for EVs.

Table 2: SDID results: Impact on new-car registrations, transfers, and scrappage

	(1)	(2)	(3)	(4)
Specification	100-nn	25-nn	10-nn	5-nn
Panel A: Gasoline/Diesel				
Madrid × Post	-4.59***	-4.34***	-4.36***	-3.94***
	(0.54)	(0.55)	(0.68)	(0.66)
Panel B: LPG				
Madrid × Post	1.61***	1.50***	1.50***	1.49***
	(0.09)	(0.10)	(0.09)	(0.12)
Panel C: Hybrids				
Madrid × Post	2.09***	1.85***	1.69***	1.59***
	(0.15)	(0.18)	(0.18)	(0.20)
Panel D: EVs				
Madrid × Post	0.18***	0.17***	0.15***	0.14***
	(0.02)	(0.02)	(0.02)	(0.03)
Panel E: Second-hand transfers				
Madrid × Post	-1.67*	-1.16	-0.45	-0.58
	(1.00)	(1.07)	(0.98)	(1.11)
Panel F: Scrappage				
Madrid × Post	-0.77	-0.18	-0.55	-0.59
	(0.73)	(0.74)	(0.67)	(0.65)
Observations	115,700	62,000	37,160	24,680
Control zip codes	5,496	2,811	1,571	948

Notes: The table reports SDID estimates for quarterly new-car registrations (Panels A–D), second-hand transfers (Panel E), and scrappage (Panel F) per 10,000 residents at the zip-code level; standard errors are computed via 100 clustered bootstrap replications following Algorithm 2 in Arkhangelsky et al. (2021). Columns (1)–(4) use, respectively, 100, 25, 10, and 5 nearest neighbors. All specifications include unit and year-quarter fixed effects and SDID unit and time weights. Sample counts apply to all panels; control zip codes include the full matched pool. LPG denotes LPG/CNG. Stars denote significance at the 10 (*), 5 (**), and 1 (***) percent levels.

2.4 Impact on second-hand car transfers and scrappage

Panels (e) and (f) of Figure 2 show scrappage and second-hand transfers. Unlike new registrations, their gap to the synthetic control stays similar to where it was before the announcement. Table 2 confirms this: in the preferred specification, column (4), the quarterly effects are −0.58 second-hand transfers (Panel E) and −0.59 scrapped cars (Panel F) per 10,000 residents, neither statistically different from zero.

Taken together, our motivating evidence suggests that Madrid Central promoted the adoption of LPG/CNG and hybrid vehicles and—though to a lesser extent—electric vehicles, while reducing the registration of new gasoline and diesel cars. In contrast, the policy had no significant effect on the scrappage of existing polluting vehicles or on incoming second-hand transfers. As we discuss further in Section 4, these null effects indicate that it is appropriate to focus on the policy's impact on new-car registrations within Madrid, without concern that the effects spilled over into the second-hand or neighboring markets.

3 A Dynamic Model of Car Choice

Our motivating evidence provides a clear picture of the effects of Madrid's LEZ but does not allow us to infer the counterfactual outcomes under alternative policy designs—scenarios that are relevant beyond the specific case of Madrid. For example, is it optimal to exempt hybrids from the LEZ restrictions? Is it preferable to restrict hybrids, though with less stringent limits than those applied to gasoline and diesel vehicles—perhaps by imposing a larger restricted area for the latter, as implemented in Berlin? Alternatively, could the gradual introduction of restrictions on hybrids—announced in advance but applied with a delay—enhance the relative appeal of EVs without imposing immediate costs on hybrid owners?

While these alternative strategies might accelerate the transition toward electric mobility, whether they are socially optimal remains an open question, given the social costs imposed by the restrictions. To explore these issues, we develop a dynamic model of vehicle choice that allows us to evaluate the effects of different LEZ designs.

3.1 Preliminaries

Our simple dynamic model has infinitely lived drivers and three types of vehicles—electric (E), hybrid (H) and gasoline/diesel (G)—offered at three distinct and equidistant points on a unit circle. A unit mass of drivers is uniformly distributed on the circle. A driver located at $x \in [0, 1]$, who at the beginning of a given period holds a car of type $j = E, H, G$ and age $a = 1, 2, \ldots, T_j$, where T_j is the (endogenous) age of the oldest car

of type j in the market, derives a surplus of

$$u_a^j(x) = v_a^j - \gamma|x - x^j|, \tag{3}$$

where v_a^j is the value the individual obtains from driving the car during one period, which is decreasing with age (i.e., $v_a^j > v_{a+1}^j$), x^j marks car j's "location" (with $x^E = 0$, $x^H = 1/3$ and $x^G = 2/3$), and $\gamma > 0$ is the (per-period) horizontal-differentiation parameter (i.e., the transportation cost). We assume that the market is always fully covered, i.e., all drivers hold on to one vehicle in each period, and no more than one, in equilibrium. This boils down to assuming that values of v_a^j are sufficiently large.[21]

Vehicles also differ in how much they pollute each period. Given our focus on adoption (extensive margin), we assume that car use (intensive margin) is unaffected by a policy shock; hence, a vehicle's total emissions can be proxied by its emissions rate. This implies that the way to reduce pollution in our setting is by inducing drivers to switch to cleaner vehicles, not to drive less.[22] A gasoline vehicle of age a emits $e_a^G = e_a$ units of pollution per period, with $e_{a+1} \geq e_a$, while a hybrid vehicle of age a emits $e_a^H = \alpha e_a$ per period, with $\alpha \in (0, 1)$. Electric vehicles are emission-free, i.e., $e_a^E = 0$. The social cost of a unit of emission is normalized to 1, so emissions capture the social cost of pollution.

3.2 The no-intervention outcome

Consider first the case in which drivers face no environmental policies. Let U_a^j be the value (i.e., lifetime utility flow) of an individual located at x with a car of type j and age a. As in Adda and Cooper (2000), the set of Bellman equations that must hold in steady-state equilibrium is given by

$$\begin{aligned} U_1^j(x) &= u_1^j(x) - p_1^j + \delta U_2^j(x) \\ &\dots \\ U_a^j(x) &= u_a^j(x) + \delta U_{a+1}^j(x) \\ &\dots \\ U_{T_j}^j(x) &= u_{T_j}^j(x) + \delta U_1^j(x) + \delta z^j, \end{aligned} \tag{4}$$

[21] This "Salop" framework, with its uniform distribution of consumers and full coverage, is sufficient but not necessary for the results that follow. We adopt it to communicate our insights in the clearest possible way, at times with corner solutions (e.g., when no single unit of a particular type is sold in equilibrium). We adopt a more flexible (horizontal-differentiation) framework in the structural analysis following this section, where we allow for an arbitrary number of vehicles, of different characteristics.

[22] This assumption is consistent with the findings of Tassinari (2024) for the LEZ in Madrid and of Bernardo, Fageda, and Flores-Fillol (2021) and Garson et al. (2026) more generally. Nevertheless, we relax it in Appendix B, showing that results remain qualitatively unchanged when we let restricted cars be driven either more or less, proportionally (or inversely proportionally) to the LEZ's size. The increase in driving is consistent with rerouting, while the reduction in driving is consistent with a greater use of public transit.

where $\delta < 1$ is the discount factor, p_1^j is the price of a new (i.e., age 1) car, and z^j is its scrap value. Note that the car is scrapped at the beginning of $T_j + 1$ (or end of T_j).[23]

Given that individuals only differ in their horizontal preferences, there is no reason for them to trade their second-hand cars while in steady state. Individuals hold on to their cars until they are scrapped, at which point they acquire another one, either brand-new or second-hand. This implies that the price of second-hand cars adjusts so, in steady state, it satisfies the indifference condition

$$U_1^j(x) = U_a^j(x) - p_a^j, \tag{5}$$

for $a > 1$.[24] Solving the Bellman equations, we arrive at the following lemma.

Lemma 1 *The lifetime value of holding on to a car of type j for an individual located at x is given by*

$$U_1^j(x) = \frac{1}{1-\delta^{T_j}} \left[\sum_{a=1}^{T_j} \delta^{a-1} u_a^j(x) + \delta^{T_j} z^j - p_1^j \right],$$

where $u_a^j(x)$ is given by (3) and T_j is the age of the oldest type-j car in the market.

The terms in brackets correspond to the net present value of owning a type-j car, a cycle that repeats every T_j periods. For $T_j + 1$ to be the optimal scrapping age, age-T_j cars must not be scrapped and age-$T_j + 1$ cars must never be kept. These two optimality conditions respond to the values of the primitives (i.e., v_a^j, p_1^j, z^j, and δ). For example, it can be shown (see Appendix D) that if it is optimal to scrap cars at age 3, then the following two conditions must hold

$$v_2^j > \frac{1}{1+\delta} \left(v_1^j + \delta v_2^j + z^j - p_1^j \right) > v_3^j. \tag{6}$$

The inequality on the left says that individuals want to run their cars for more than one period, whereas the inequality on the right indicates that running the car for an additional (i.e., third) period is not worthwhile. In this case, it is better to scrap the car and get a brand-new one (or, equivalently, acquire a "one-period-old" car and run it only for one period).

Let us denote by $\tilde{x}^{jk}$ the location of drivers who are indifferent between owning vehicle j and $k \neq j$, i.e., $U_1^j(\tilde{x}^{jk}) = U_1^k(\tilde{x}^{jk})$. Drivers located between $\tilde{x}^{jk}$ and $\tilde{x}^{jl}$, with

[23] Note that for simplicity we are assuming that all cars survive until they are scrapped at the end of their lives, but we could easily let them exit the market at some exogenous rate due to crashes, fatal malfunctioning, etc., a rate that may vary with the vehicle's type and age. For example, if we denote by ξ_a^j the probability that a car of type j and age a survives to the next period, the corresponding Bellman equation would change to $U_a^j(x) = u_a^j(x) + \xi_a^j \delta U_{a+1}^j(x) + (1 - \xi_a^j)\delta(U_1^j(x) + z^j)$.

[24] This condition can be conveniently rewritten as

$$p_a = v_a + \delta U_{a+1} - U_1,$$

which is expression (3) in Adda and Cooper (2000).

$l \neq j, k$, will own a type-j vehicle, so from Lemma 1 we obtain that the total demand for new type-j vehicles is equal to

$$q_1^j = \int_{\tilde{x}^{jk}}^{\tilde{x}^{jl}} dx = \frac{1}{3} + \frac{1}{2\Gamma}\left(2(V_j + Z_j - P_j) - (V_k + Z_k - P_k) - (V_l + Z_l - P_l)\right), \quad (7)$$

for $j, k, l = E, H, G$ and $j \neq k, l$, and where $\Gamma = \gamma/(1-\delta)$, $V_j = \sum_{a=1}^{T_j} \delta^{a-1} v_a^j/(1-\delta^{T_j})$, $Z_j = \delta^{T_j} z^j/(1-\delta^{T_j})$, and $P_j = p_1^j/(1-\delta^{T_j})$. The demand for j is decreasing in its own (lifetime) price P_j and increasing in its (lifetime) service value V_j and (lifetime) scrap value Z_j. Moreover, because vehicle types are substitutes, demand for j is increasing in P_k and P_l, and decreasing in V_k, V_l, Z_k and Z_l.

Since equation (7) also corresponds to the per-period demand for second-hand cars, total welfare, in present value terms, is given by

$$W = \sum_{j=E,H,G} \left((V_j + Z_j - C_j - \mathcal{E}_j) q_1^j - \int_{\tilde{x}^{jk}}^{\tilde{x}^{jl}} \Gamma |x - x^j| dx \right), \quad (8)$$

for $j, k, l = E, H, G$ and $j \neq k, l$, and where $\mathcal{E}_j = \sum_{a=1}^{T_j} \delta^{a-1} e_a^j/(1-\delta^{T_j})$ is j's (lifetime) emissions harm, and $C_j = c^j/(1-\delta^{T_j})$ is its (lifetime) production cost, with c^j being the unit cost.[25]

3.3 Policy intervention: Low Emission Zones

We are interested in how Low Emission Zones influence the adoption of low-emission vehicles. Since our focus is on addressing environmental externalities, we assume vehicles are sold at cost (equivalent to assuming they are sold at a constant, uniform markup).[26]

3.3.1 LEZs vs. taxes

As a benchmark for evaluating the performance of LEZs, consider taxes.[27] Both types of interventions—taxes and LEZs—affect the market-equilibrium outcome, although in different ways. This follows from the demand expressions in (7). Taxes work by moving (lifetime) prices P_j away from C_j. For example, if each period a Pigouvian tax $\tau_a^j = e_a^j$

[25] Note that cars can have different lifetimes.

[26] Allowing car dealers to earn oligopoly rents in this way has no welfare implications and thus does not affect the optimal policy design. Introducing market power through other mechanisms, however, would not be neutral in this respect.

[27] Other price-based instruments could also be considered, such as subsidies for the purchase of new low-emission vehicles and for the earlier retirement of polluting ones. Note, however, that these alternative instruments do not replicate the work of taxes. Subsidies for the purchase of low-emission vehicles may fail to implement the first-best by extending their lifespans beyond the first-best level. Scrappage subsidies may help shorten the lifespan of existing gasoline cars in the short run, but in the long run, permanent scrappage subsidies could have the undesirable effect of making the purchase of new gasoline cars attractive to more individuals, despite their shorter lifespans.

is added to the registration of a vehicle of type j, then $P_j = C_j + \mathcal{E}_j$.[28]

A LEZ policy—which resembles a proportional rationing scheme—works differently by reducing the (lifetime) service value of the restricted car from v_a^j to $(1 - s_a^j)v_a^j$, where $s_a^j < 1$ is the size of the LEZ affecting vehicle j of age a.[29] The size of the LEZ should be interpreted broadly, not only in terms of its geographic coverage, but also in terms of the hours of the day and days of the week during which it is active.[30] One could argue that two different LEZs, of size s^G affecting gasoline cars and of size s^H affecting hybrids, could replicate the work of two different (lifetime) taxes, τ^G and τ^H. This is true in terms of market shares of new cars, as can be inferred from equation (7), but not in terms of welfare.

The key distinction lies in how LEZs and taxes affect consumers. LEZs work by limiting the use of restricted vehicles, reducing welfare not only for marginal consumers—who switch to a less preferred (but unrestricted) option—but also for inframarginal consumers who continue to hold on to restricted vehicles. In contrast, taxes influence decisions at the margin without destroying welfare for inframarginal users (aside from the tax payment itself).[31] As a result, the difference between taxes and LEZs boils down to efficient versus inefficient rationing. This difference is central to the design of a LEZ.[32]

To better understand how to arrive at the optimal LEZ design, we proceed sequentially by first studying the case of *short-lived* vehicles—lasting only one period, after which they are scrapped. This case serves to highlight the bang-bang logic of the optimal LEZ design. Yet, as with any intervention in a durable-good market, it might take time for the composition of the car fleet to adjust to its new (steady-state) equilibrium. Given this gradual adjustment, one can entertain the idea that the optimal LEZ design could also evolve alongside the fleet until reaching the new steady-state equilibrium. To analyze this, we next consider the case of *long-lived* vehicles, which allows uncovering potential dynamic components of the optimal LEZ design, i.e., the emergence of transition technologies and the importance of commitment.

Without loss of generality, we simplify the analysis by assuming that cars are symmetric except for their pollution levels: all cars have the same production cost c, provide the same service value v_a given their vintage, and have the same scrap value z. We also assume that pollution is age-independent, i.e., $e_a = e$ for all a.

[28] Note that T_j may have also changed (i.e., decreased) as a result of the tax.

[29] We adopt this multiplicative form here for tractability. In the next section, we adopt a more flexible form, for example, by allowing the LEZ impact to vary by income.

[30] Recall that our working assumption is that policies do not affect the intensive margin; that is, they do not influence driving intensity and therefore impact emissions solely through changes in vehicle adoption. Under our baseline assumption, welfare losses stem not from fewer trips but from less desirable ones—for example, longer detours to avoid the LEZ, or shopping trips redirected outside the LEZ instead of within it.

[31] We assume that the tax collection is rebated as a lump sum to individuals and thus does not distort overall efficiency.

[32] This difference would also extend to the intensive margin—specifically, driving behavior. Gasoline taxes discourage the least valued trips, leading to efficient rationing, while LEZs restrict all trips indiscriminately, regardless of their value, resembling proportional rationing.

3.3.2 Equilibrium with short-lived vehicles

We first consider *short-lived* vehicles, which last one period. Under the symmetry assumptions above, each vehicle type accounts for one-third of the market in the absence of intervention. Introducing a LEZ changes these shares depending on which vehicle types it restricts and how extensively it restricts each type.

A LEZ's design has two dimensions: its stringency (i.e., the type of cars that are restricted; either no cars, only gasoline cars, or both gasoline cars and hybrids), and its size, $s \in [0, \bar{s}]$, where $\bar{s}$ is the largest LEZ that is (politically) feasible to implement. We let $\bar{s} \leq s^{\max} \equiv \gamma/3v_1$, which is the size of a LEZ that would completely displace gasoline cars from the market.[33]

For now, we restrict attention to LEZ designs that treat all restricted vehicles equally, so $s^G = s^H = s$ if both gasoline cars and hybrids are restricted. We later allow restrictions to vary by vehicle type. Since EVs are pollution-free, it is never optimal to limit their access to the LEZ. Hence, without loss of generality, we set $v_1^E = v_1$.

The next proposition characterizes, for a given LEZ of size s, which vehicles (if any) should be restricted. Note that, due to its inefficient rationing, a LEZ design can never implement the first-best (FB) solution. Instead, Pigouvian taxes $\tau^j = e^j$ implement the FB solution by pricing cars at their social cost: $p_1^G = c + e$, $p_1^H = c + \alpha e$, and $p_1^E = c$.[34]

Proposition 1 (LEZ restrictions with short-lived vehicles) *Consider a LEZ of size $s \in [0, \gamma/3v_1]$. Define the emission thresholds*

$$e_G^{LEZ}(s) \equiv \frac{2\gamma - 3v_1 s}{3(2-\alpha)}, \qquad e_H^{LEZ} \equiv \begin{cases} \dfrac{2\gamma}{3(2\alpha-1)}, & \alpha \in (1/2, 1), \\ +\infty, & \alpha \in (0, 1/2], \end{cases}$$

so that $e_G^{LEZ}(s) < e_H^{LEZ}$ for all $\alpha \in (0,1)$. Then it is optimal to:

(i) restrict no vehicle if $e \leq e_G^{LEZ}(s)$;

(ii) restrict only gasoline cars if $e_G^{LEZ}(s) < e \leq e_H^{LEZ}$; and

(iii) restrict both gasoline and hybrid vehicles if $e > e_H^{LEZ}$.

[33] A LEZ that also restricts hybrids would completely displace both gasoline and hybrid cars if it were any larger than $\gamma/3v_1$, the size at which their drivers become indifferent between their car and an EV. It seems realistic, however, to restrict attention to LEZs not larger than $\gamma/3v_1$, or more generally, not larger than $\Gamma/3V$, as gasoline cars would be fully displaced. Looking at larger LEZs would only introduce additional cases without adding new insights.

[34] There is only one case, highly unrealistic and of little interest, where it could potentially replicate the FB: when the FB solution includes only electric vehicles in the market and authorities can deploy a sufficiently large LEZ, essentially banning gasoline cars and hybrids from circulation altogether.

The LEZ thresholds are depicted in Figure 3.[35] Restricting gasoline vehicles only entails a trade-off. On the one hand, gasoline-car users suffer a loss as they cannot access the restricted area. On the other hand, there is a reduction in emissions as consumers substitute hybrids and EVs for gasoline cars. Therefore, for a given LEZ size s, it is optimal to restrict gasoline cars when their emissions are high relative to the cost of reducing choice (proxied by γ). The critical emission threshold $e_G^{LEZ}(s)$ also depends on the size of the LEZ: a bigger LEZ reduces the market share of gasoline cars, thus reducing the mass of drivers who are restricted from entering the LEZ. Hence, for small s, the mass of (inframarginal) gasoline-car users is large, and it is not optimal to restrict them (nor hybrid users) unless their emissions are substantial. In this case, the welfare loss inflicted on inframarginal gasoline users is more than compensated by the emission-reduction gain from marginal gasoline users who switch to cleaner cars.

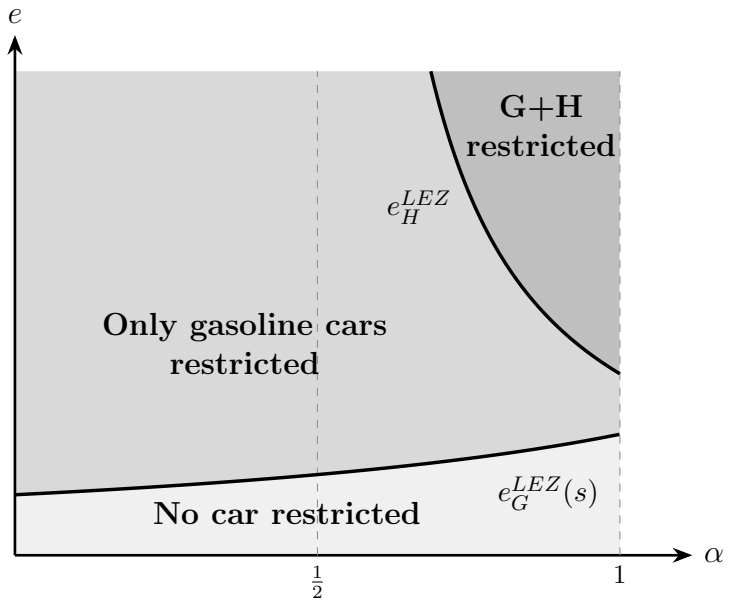


Figure 3: Optimal restrictions for a given LEZ size

Notes: The figure illustrates the optimal restrictions for a LEZ of given size s, as a function of the emissions rate e and the hybrid efficiency parameter $\alpha \in (0,1)$. The curve $e_G^{LEZ}(s)$ separates the region in which no car is restricted (below) from the one in which only gasoline cars are restricted (above); the curve e_H^{LEZ} separates the region in which only gasoline cars are restricted (below) from the one in which both gasoline and hybrid cars are restricted (above). The dashed vertical line at $\alpha = 1/2$ marks the asymptote of e_H^{LEZ}: for $\alpha \leq 1/2$, restricting hybrids is never optimal.

Similarly, restricting hybrid vehicles (in addition to gasoline cars) entails a trade-off

[35] Note that as $\alpha \to 1$, when $e \in \left(e_G^{LEZ}(s), e_H^{LEZ}\right)$, the optimal policy is to restrict only gasoline cars, even though hybrid and gasoline cars pollute (almost) the same. Both thresholds are continuous in α, and at $\alpha = 1$ they differ by $e_H^{LEZ} - e_G^{LEZ}(s) = v_1 s$. The reason is that the two restrictions achieve the same emission reduction (driven by the drivers who switch to EVs) but impose the usage loss on different masses of drivers. Once restricted, gasoline cars lose market share to both hybrids and EVs, so the restriction only binds on the remaining $1/3 - sv_1/\gamma$ gasoline drivers. Extending it to hybrids leaves total transportation costs unchanged but raises the mass of restricted drivers by a full $1/3$, as it now also binds on all hybrid drivers, including those who had switched from gasoline cars.

between the welfare loss of constraining hybrid users and the impact on emissions. However, unlike gasoline cars, restricting hybrids does not necessarily lead to lower emissions: some sales are shifted to cleaner electric vehicles, but others to more polluting gasoline cars as hybrids lose their advantage relative to gasoline cars. In particular, if hybrids are sufficiently clean relative to gasoline cars, i.e., if $\alpha < 1/2$, the second effect implies that net emissions would increase. Hence, in this case, it is never optimal to restrict hybrids from accessing the LEZ (as the threshold e_H^{LEZ} tends to infinity). Otherwise, if $\alpha > 1/2$, net emissions would decrease, making it optimal to restrict hybrids, but only when the emission reduction is high relative to the cost of limiting car choice (proxied by γ). The emission threshold does not depend on the size of the LEZ directly, given that, conditional on restricting gasoline cars (for which s has to be sufficiently large), the welfare gain from also restricting hybrids is proportional to s.

Having characterized the optimal access to a LEZ of a given size, we now characterize its optimal size.

Proposition 2 (Optimal LEZ with short-lived vehicles) *Let $\bar{s} \leq s^{\max} \equiv \gamma/3v_1$ be the largest (politically feasible) LEZ. There exists a threshold size $\tilde{s}$ such that it is optimal to have:*

(i) no LEZ if either $e \leq e_G^{LEZ}(s^{\max})$, or $e_G^{LEZ}(s^{\max}) < e \leq e_G^{LEZ}(0)$ and $\bar{s} < \tilde{s}$;

(ii) a LEZ of the maximum possible size $\bar{s}$ restricting gasoline vehicles only if either $e_G^{LEZ}(s^{\max}) < e \leq e_G^{LEZ}(0)$ and $\bar{s} > \tilde{s}$, or $e_G^{LEZ}(0) < e \leq e_H^{LEZ}$; and

(iii) a LEZ of the maximum possible size $\bar{s}$ restricting both gasoline vehicles and hybrids if $e > e_H^{LEZ}$.

As illustrated in Figure 4 and formally shown in the proof of the proposition, welfare is everywhere convex in the size of the LEZ. This contrasts sharply with welfare under taxes, which is everywhere concave. This difference explains why the optimum under taxes is interior—namely, Pigouvian taxes—whereas under a LEZ it follows a bang-bang logic: either implement no LEZ (i.e., $s = 0$) if it cannot be made large enough, or implement the largest feasible LEZ (i.e., $s = \bar{s}$).[36] In sum, intermediate-sized LEZs are never optimal.

This bang-bang logic reflects the fact that the larger the LEZ, the lower the sales of restricted vehicles. As a result, restricting access to these vehicles becomes increasingly attractive, since the associated utility loss affects fewer inframarginal users. Thus, a LEZ is socially beneficial if it is strict enough to shift purchase decisions away from polluting vehicles, even at the cost of reducing value for inframarginal consumers. If emissions are large enough, any gain on marginal sales compensates for the inframarginal loss. This

[36] Note that $\tilde{s} = 0$ for $e = e_G^{LEZ}(0)$, so if $e_G^{LEZ}(0) \leq e < e_H^{LEZ}$, any LEZ restricting gasoline cars leads to welfare gains.

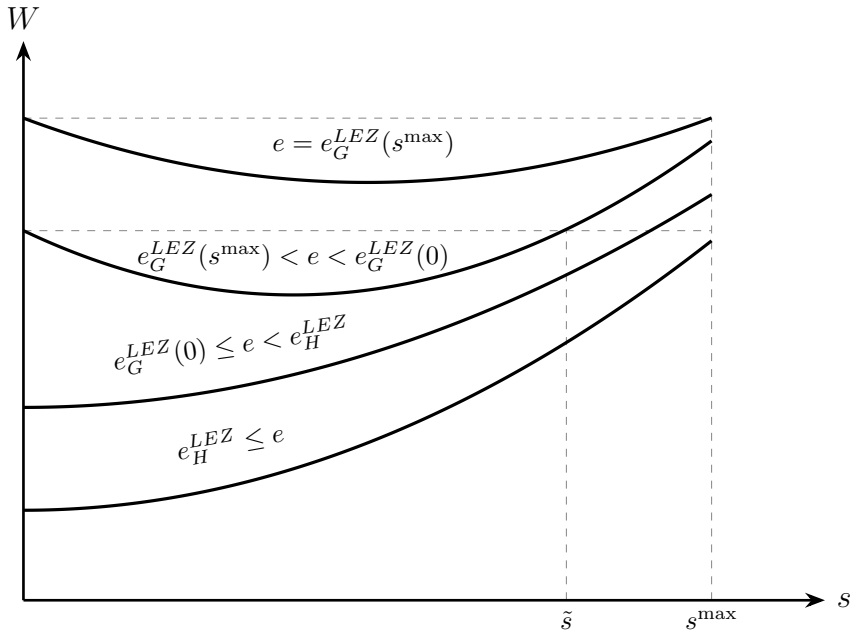


Figure 4: Optimal LEZ size

Notes: This figure depicts welfare as a function of the LEZ size s. Each curve corresponds to a different value of the emissions rate e, with higher values of e represented by lower curves. Welfare is everywhere convex in s: increasing s lowers the value of each restricted car, but as s grows, the share of restricted vehicles declines. Dots mark the two bang-bang candidates, $s = 0$ and $s = s^{\max}$. The dashed horizontal lines indicate the welfare levels at $s = 0$ for the top two curves; $\tilde{s}$ is the size at which the second curve returns to its $s = 0$ level.

explains why, when e is large enough, it is always optimal to create a LEZ and make it as large as possible. However, if emissions are not high enough (like the two upper curves in Figure 4), the LEZ has to be large enough so that the impact on the marginal sales compensates for the loss of inframarginal users, whose mass is lower the larger the size of the LEZ. The planner should therefore implement the largest feasible LEZ if it is large enough to increase welfare (part (ii) of the proposition), and otherwise introduce no LEZ (part (i)).

It is important to emphasize that the bang-bang result in Proposition 2 is not an artifact of the uniform distribution of consumer preferences along the circle. Although that distribution implies that welfare is convex in s over the entire relevant domain, the bang-bang logic only requires welfare to be quasi-convex in s. As we illustrate in Appendix E, this condition is consistent with a broader class of preference distributions. In the empirical analysis, where substitution patterns are generated by a discrete-choice model with generalized extreme value utility shocks, the relevant quasi-convexity condition is assessed numerically by evaluating welfare over a wide range of LEZ sizes.

One could extend our theory to assess whether the social planner could gain from a LEZ design that considers different sizes for hybrids and gasoline cars, say s^G and s^H. Convexity of the welfare function, however, implies that this is never optimal: hybrids

should be treated like either electric vehicles ($s^H = 0$) or gasoline cars ($s^H = s^G$).

Corollary 1 *If it is possible to set $s^H \in [0, s^G]$, the optimal LEZ involves either $s^H = s^G$ or $s^H = 0$.*

This result invites another question: in a dynamic environment, where the fleet adjusts gradually to a policy shock, could it be optimal to treat hybrids like electric vehicles for some time, after which they should be treated like gasoline cars? And, would it be credible to do so in the absence of commitment? Positive answers to these questions would make hybrids a transition technology according to our theory.[37] We explore this possibility next by extending the vehicles' lifetimes beyond one period.

3.3.3 Equilibrium with long-lived vehicles

Let $t = 0, 1, 2, \ldots, \infty$ denote time. At $t = 0$, the fleet is in its no-intervention steady state, as characterized in Section 3. Suppose that at $t = 1$, the planner may decide to intervene in the market by imposing a LEZ. Suppose also that all cars last for two periods regardless of their type, or whether a LEZ policy is in place. This means that they are scrapped at the end of the second period or the beginning of the third.[38]

In this two-period cycle, it will take at most three periods for the fleet composition to reach its new steady state after the introduction of a LEZ in period $t = 1$ (it will take only two periods if the LEZ design is time-invariant). This new steady state need not be stationary. Even in the absence of vertical differentiation, trade in the second-hand market may generate oscillations after the introduction of a LEZ. In particular, holders of second-hand gasoline cars with relatively weak preferences for them may prefer to sell immediately and buy a new unrestricted car, rather than wait until their gasoline vehicles are scrapped. To facilitate the exposition, we abstract from this oscillatory equilibrium in what follows by assuming trade frictions large enough to suppress second-hand-market trading. The frictions documented by Gavazza, Lizzeri, and Roketskiy (2014) motivate this assumption. More importantly, reducing trade frictions does not qualitatively change our results; it only alters the emission thresholds that characterize the optimal LEZ restrictions. See Appendix D for details.

In this context, the planner's welfare function for a LEZ $s^j_{a,t} \in [0, \bar{s}]$, denoting the restriction upon a vehicle $j = E, H, G$ of vintage $a = 1, 2$ in period $t \geq 1$, can be written as

[37] As we will see below, our view of transition technology differs sharply from that based on changes in the relative cost of electric vehicles over time.

[38] Formally, we are assuming that expression (6) holds at all times, with $v^j_a = (1 - s^j_a)v_a$. If anything, (6) shows that, assuming $s_a = s_{a+1} = s$, drivers would have incentives to hold on to their restricted vehicles longer, since $sv_1 + \delta sv_2 > (1 + \delta)sv_3$. The intuition is that new gasoline cars have become relatively more expensive given their lower quality throughout, so drivers want to replace them less frequently, if anything.

$$W = \sum_{t=1}^{\infty} \delta^{t-1} w_t = w_1 + \delta w_2 + \frac{\delta^2}{1-\delta} w_3, \tag{9}$$

where w_t is given by

$$\sum_{j=E,H,G} \left(\int_{\tilde{x}_{1,t}^{jk}}^{\tilde{x}_{1,t}^{jl}} (v_{1,t}^j - c - \gamma|x - x^j| - e^j) dx + \int_{\tilde{x}_{2,t}^{jk}}^{\tilde{x}_{2,t}^{jl}} (v_{2,t}^j + \delta z - \gamma|x - x^j| - e^j) dx \right), \tag{10}$$

$v_{a,t}^j = (1 - s_{a,t}^j) v_a$, and $\tilde{x}_{a,t}^{jk}$ is the location that separates individuals driving type-j vehicles from those driving type-k vehicles.

Because cars survive until they are scrapped, the location of the indifferent driver remains unchanged as the cars get older, i.e., $\tilde{x}_{2,t+1}^{jk} = \tilde{x}_{1,t}^{jk}$ for all $t \geq 1$. We also know that at $t = 3$ the fleet reaches its steady state, so $\tilde{x}_{1,t-1}^{jk} = \tilde{x}_{a,t}^{jk} = \tilde{x}_{a,t+1}^{jk}$ for all a and $t \geq 3$. To find $\tilde{x}_{1,t}^{jk}$, we use the set of Bellman equations (4) adjusted to account for the LEZ policy.[39]

Consider a LEZ policy introduced at $t = 1$. During this first period, second-hand cars affected by the policy will outnumber their new counterparts. This implies that, in equilibrium, the price of these second-hand cars must fall enough to make their drivers indifferent between them and new cars. If the policy is adjusted at $t = 2$—for example, by adding more restricted cars or extending its size—new and second-hand cars of each type will continue to be in different proportions. Only at $t = 3$ will new and second-hand cars of each type be present in equal numbers, marking the point at which the fleet reaches its new steady state.[40]

Thus, if at $t = 1$ the social planner commits to a time-invariant LEZ of size s, the short-lived analysis from Proposition 1 extends to this dynamic environment with only minor notational changes. For now, we assume that the planner can credibly commit to her announcement and make no policy adjustments at $t = 2$. Below, we discuss whether and when a lack of commitment may be a problem. To simplify notation from now on, we assume that all individuals, including the social planner, discount the future at $\delta = 1/2$. Our next proposition is the analogue of Proposition 1, extended to the case of vehicles with a two-period lifespan.

Proposition 3 (Time-invariant LEZ with long-lived vehicles) *Consider a time-invariant LEZ of size $s < s^{\max} \equiv \Gamma/3V = \gamma/(2v_1 + v_2)$ introduced at $t = 1$. There exist emission*

[39] Note that the price of second-hand units is still obtained from (5), even during the transition from one steady state to the next.

[40] Note that for individuals driving second-hand cars during the transition to the new steady state, the location $\tilde{x}_{2,1}^{jk}$, and also $\tilde{x}_{2,2}^{jk}$ if the LEZ design varies over time, does not necessarily correspond to the one that makes them indifferent between driving either type of car. Suppose that j is restricted at $t = 1$ but k is not. If at $t = 0$, individuals located at $\tilde{x}_{2,0}^{jk} = \tilde{x}_{1,0}^{jk}$ were indifferent between driving either type of car, at $t = 1$ individuals driving second-hand j cars are strictly worse off than their counterparts driving second-hand k cars. Despite this, individuals driving second-hand j cars will continue to drive them for one more period, after which they will switch to a new (or second-hand) k car.

thresholds e_G^{LEZ}*, decreasing in* s *and increasing in* α*, and* $e_H^{LEZ} > e_G^{LEZ}$*, independent of* s *and decreasing in* α*, such that it is optimal to:*

(i) restrict no vehicle if $e \leq e_G^{LEZ}$*;*

(ii) restrict only gasoline cars if $e_G^{LEZ} < e \leq e_H^{LEZ}$*; and*

(iii) restrict both gasoline and hybrid vehicles if $e > e_H^{LEZ}$*.*

In a dynamic environment, the planner has the option to commit to a LEZ that evolves over time. In principle, the LEZ can evolve in both stringency (i.e., the type of cars that are restricted) and size. It is not difficult to anticipate from Propositions 1 and 2 that the short-lived analysis regarding gasoline cars remains unchanged: if it is optimal to restrict gasoline cars from the very first day, it is optimal to restrict them to the maximum extent.

This parallel does not extend to hybrids. It may be optimal to treat them gradually (i.e., as a transition technology)—yet still within a bang-bang logic—by keeping them unrestricted like electric vehicles for some time, after which they become as restricted as gasoline cars.[41] As the next proposition establishes, in our dynamic environment where cars last for two periods, this translates into placing no restriction on hybrids during the first period, $t = 1$, and applying the gasoline-level restriction from the second period onward.

Proposition 4 (Hybrids as a transition technology) *Consider a LEZ of size* $s < s^{\max}$*, introduced at* $t = 1$*. There exist transition-technology (TT) emission thresholds* $e_{TTL}^{LEZ} < e_H^{LEZ}$ *and* $e_{TTU}^{LEZ} > e_H^{LEZ}$*, where* e_G^{LEZ} *and* e_H^{LEZ} *are defined as in Proposition 3, such that it is optimal to:*

(i) restrict no vehicle if $e \leq \min\{e_G^{LEZ}, e_{TTL}^{LEZ}\}$*;*

(ii) restrict only gasoline cars from $t = 1$ *onward if* $\min\{e_G^{LEZ}, e_{TTL}^{LEZ}\} < e \leq e_{TTL}^{LEZ}$*;*

(iii) restrict gasoline cars from $t = 1$ *onward, and treat hybrids as a transition technology (i.e., as electric cars during* $t = 1$*, and as gasoline cars from* $t = 2$ *onward) if* $e_{TTL}^{LEZ} < e < e_{TTU}^{LEZ}$*; and*

(iv) restrict both gasoline cars and hybrids from $t = 1$ *onward if* $e \geq e_{TTU}^{LEZ}$*.*

[41] In principle, gasoline cars could also be treated as a transition technology, with restrictions potentially phased in over a different horizon than for hybrids (although that would require extending the model to allow cars to last at least three periods). We rule out this possibility to avoid a proliferation of cases without adding much economic insight. This exclusion can also be justified on political-feasibility grounds: a LEZ that restricts no vehicle at the time of its introduction would be difficult to present as a meaningful policy intervention.

The top and middle rows in Figure 5 depict the relevant emission thresholds in Propositions 3 and 4, respectively. To convey intuition about the emergence of hybrids as a transition technology, consider two LEZ designs: a time-invariant design introduced at $t = 1$, with an emission threshold e_H^{LEZ}—determining whether hybrids should be restricted from $t = 1$ onward or not at all; and an alternative LEZ design, still announced at $t = 1$, with a lower emission threshold $e_{TTL}^{LEZ} < e_H^{LEZ}$ determining whether hybrid cars should be restricted from $t = 2$ onward or not at all.[42] According to our theory, it is optimal to treat hybrids as a transition technology whenever $e \in \left(e_{TTL}^{LEZ}, e_{TTU}^{LEZ}\right)$, that is, whenever e falls between the lower and upper transition-technology (TT) thresholds. Within this range, hybrids are not polluting enough to require immediate restriction, but they also fall short of being clean enough to remain unrestricted indefinitely.

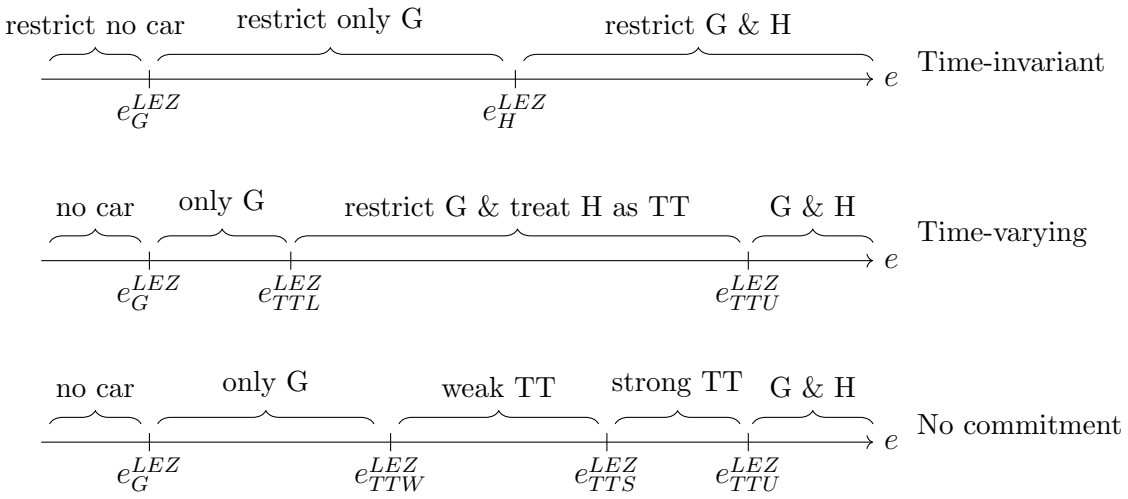


Figure 5: Emission thresholds for different LEZ designs

Notes: For given parameters, these diagrams show the thresholds for e corresponding to each type of LEZ solution under three cases: a time-invariant LEZ, a time-varying LEZ, and a setting in which the planner lacks commitment to implement the LEZ once it has been announced. For the time-varying case we only present the most relevant scenario, i.e., $e_G^{LEZ} \leq e_{TTL}^{LEZ}$.

Why does a transition technology emerge? Announcing at $t = 1$ that hybrids will be restricted in the future has immediate and future effects. It has an immediate impact on car choice: in anticipation of their lower future value, some individuals switch right away from hybrids, even if they are not yet restricted.[43] As a result, there will be fewer inframarginal consumers driving hybrids tomorrow—fewer than if hybrids were left unrestricted indefinitely—when the restriction becomes active. With fewer (second-hand) hybrids on the road tomorrow, the future restriction on hybrids becomes less costly to

[42] These are the only two relevant LEZ designs to consider. It is easy to see that announcing the restriction on hybrids to begin at $t = 3$ instead of $t = 2$ is strictly dominated by the other two designs.

[43] Some individuals switch to electric vehicles while others switch to gasoline cars. The shift to gasoline cars is less relevant, however, as they are already restricted. In fact, hybrids' market share goes down from $1/3 + (2v_1 + v_2)s/6\gamma$ to $1/3 + (2v_1 - v_2)s/6\gamma$, though it remains above $1/3$, the no-intervention share.

implement. The mere announcement that hybrids will be restricted in the future brings the LEZ instrument—at least with respect to hybrids and for a limited time (in our case, one period)—closer to a price instrument: it affects consumer choices at the margin without destroying value for inframarginal consumers, i.e., those who continue driving hybrids in $t = 1$.

Note that when $e_{TTL}^{LEZ} \geq e_{G}^{LEZ}$ the range $\left(e_{TTL}^{LEZ}, e_{TTU}^{LEZ}\right)$ in Proposition 4 is shrinking in s, thereby reducing the range within which hybrids qualify as a transition technology. The reason is that a larger LEZ leaves fewer inframarginal consumers affected by the LEZ to protect, reducing the need for a transition technology to ease their pain. In a way, the size of the LEZ and the option to treat hybrids as a transition technology act as substitutes.[44]

Our definition of a transition technology differs sharply from alternative notions, which are often based on the idea that electric vehicles are still too costly or unfamiliar to serve as a viable outside option for many gasoline drivers (Holland, Mansur, and Yates 2021). Under this alternative view, hybrids should remain unrestricted until these concerns dissipate. While our definition is not in conflict with that view, it is fundamentally different. For one, it is developed in a context in which primitives (i.e., costs, preferences, and pollution rates) do not change over time. For another, it applies specifically in the context of a LEZ, not of a FB intervention. Furthermore, these alternative notions do not have bite in a FB context—if anything, taxes on polluting cars should be even higher to account for potential learning-by-doing and network externalities associated with the adoption of electric vehicles. For the same reason, alternative notions should have no bite in a LEZ context either.[45]

Our notion of transition technology is also distinctive in its durable-good nature. As in other durable-good problems (e.g., Coase 1972), the social planner may need commitment to implement the LEZ design advanced in Proposition 4, at least for certain values of e.[46] It is clear that commitment is not required in case (i), when $e \leq e_{G}^{LEZ}$. Since no cars are restricted in this case, the world at $t = 2$ is not different than that at $t = 1$. Commitment is also unnecessary in cases (ii) and (iv). In case (ii), where only gasoline cars are restricted, the planner has no incentive to revise the LEZ design at $t = 2$ for two reasons. First, there are fewer second-hand gasoline cars, so fewer inframarginal drivers would be hurt if the restriction is maintained. Second, there are more second-hand hybrids, so extending the restriction to hybrids would now harm more inframarginal drivers. A similar logic applies to case (iv), when both gasoline and hybrid cars are restricted from $t = 1$ onward: by $t = 2$, there are fewer inframarginal drivers of either type, leaving no

[44] On the other hand, if s is sufficiently small and α sufficiently large—greater than $10/11$—such that $e_{TTL}^{LEZ} \leq e_{G}^{LEZ}$, then hybrids should be restricted whenever gasoline cars are, whether at $t = 2$ under a transition-technology mode, when $e \in \left(e_{TTL}^{LEZ}, e_{TTU}^{LEZ}\right)$, or at $t = 1$, when $e \geq e_{TTU}^{LEZ}$.

[45] It is important to emphasize that the bang-bang rationale behind our notion of transition technology is not an artifact of the assumption that cars last only two periods, nor of the assumption that a car's scrapping age does not respond to policy interventions.

[46] No such commitment is needed for the time-invariant design in Proposition 3, which is time consistent.

reason to lift the restriction on either group.[47]

The story is different in case (iii), when the planner announces a LEZ scheme in which hybrids are treated as a transition technology. In this case, the planner wants to restrict hybrids at $t = 2$, as promised, as long as

$$e \geq \frac{4\gamma(v_1 + v_2) + (2v_1 - v_2)v_2 s}{3(2v_1 + v_2)(2\alpha - 1)} \equiv e_{TTS}^{LEZ} \in \left(e_H^{LEZ}, e_{TTU}^{LEZ}\right), \tag{11}$$

where "TTS" stands for "strong" transition-technology threshold. Hence, the planner suffers from a time-inconsistency problem whenever $e < e_{TTS}^{LEZ}$. In this case, given the number of (second-hand) hybrids still on the road at $t = 2$, the planner no longer finds it optimal to impose the restriction she had committed to at $t = 1$: from today's vantage point, the discounted future benefit of the restriction no longer outweighs the cost to inframarginal hybrid drivers.

The planner may have a strong reputation for keeping her promises—as the monopolist in Ausubel and Deneckere (1989), for example—but what if she does not? Borrowing from the durable-good monopoly literature, one can think of different ways in which the planner could solve her time-inconsistency problem. One way to restore commitment, at least partially, is by placing some restriction upon hybrids in $t = 1$, e.g., by establishing an initially smaller LEZ from which hybrids are restricted, i.e., $s_{a,1}^H < s_{a,t>1}^H = s^G = s$. This further reduces the number of hybrids, thereby making it less costly to impose the restriction in period $t = 2$.

Proposition 5 (A no-commitment LEZ) *If the planner lacks commitment, there still exists a no-commitment transition-technology LEZ design that departs from the scheme in Proposition 4 as follows:*

(i) if $e \in \left[e_{TTL}^{LEZ}, e_{TTW}^{LEZ}\right)$, *where* $e_{TTW}^{LEZ} < e_{TTS}^{LEZ}$, *then remove any restriction imposed on hybrids (zone of no transition technologies);*

(ii) if $e \in \left[e_{TTW}^{LEZ}, e_{TTS}^{LEZ}\right)$, *then place a milder restriction on hybrids in* $t = 1$, *i.e.,* $s_{a,1}^H(e) < s^G = s$, *that is decreasing in* e, *while maintaining their restriction in* $t = 2$ *(zone of "weak" transition technologies); and*

(iii) if $e \in \left[e_{TTS}^{LEZ}, e_{TTU}^{LEZ}\right)$, *then maintain the restrictions imposed on hybrids (zone of "strong" transition technologies).*

As summarized in the bottom row of Figure 5, lack of commitment reduces the range for transition technologies to emerge in equilibrium but does not completely eliminate it. This resembles the partial commitment of the durable-good monopolist in Kahn

[47] As formally shown in the proof of Proposition 5, in case (ii), when $e < e_{TTL}^{LEZ}$, it is optimal to revise the policy at $t = 2$ and extend the restriction to hybrids if $e \geq e_{(ii)}^{LEZ} > e_H^{LEZ}$, which cannot happen. Similarly, in case (iv), when $e > e_{TTU}^{LEZ}$, it is optimal to revise the policy at $t = 2$ and lift the restriction on hybrids if $e \leq e_{(iv)}^{LEZ} < e_H^{LEZ}$, which cannot happen either.

(1986). But unlike in Kahn (1986)—or in Coase (1972) for that matter—where the no-commitment solution is unique, here the planner may end up worse off than in Proposition 5.[48]

To see this, suppose gasoline emissions e are slightly above the "strong" threshold e^{LEZ}_{TTS}. According to Proposition 5, it is an equilibrium for the planner to impose the restriction on hybrids at $t = 2$ if, at $t = 1$, drivers believe she will do so. However, there exists another equilibrium. It is also an equilibrium for the planner to renege on her announcement and impose no restriction on hybrids at $t = 2$ if, at $t = 1$, drivers expect her to renege and not impose the restriction at $t = 2$. In fact, this "more pessimistic" no-commitment equilibrium exists whenever

$$e \leq \frac{4\gamma(v_1 + v_2) + (2v_1 + v_2)v_2 s}{3(2v_1 + v_2)(2\alpha - 1)} \equiv e^{LEZ}_{TTS'} \in \left(e^{LEZ}_{TTS}, e^{LEZ}_{TTU}\right). \tag{12}$$

The existence of this alternative equilibrium further narrows the range in which transition technologies may emerge absent commitment, making the ability to commit even more valuable.

4 Empirical Analysis

The next two sections take the model to the data. We ask whether, at the time of implementation, Madrid satisfied the necessary conditions—in consumer preferences and vehicle externalities—for a LEZ to be worth considering at all, and, if not, what would have to change for those conditions to hold. We focus on the time-invariant optimal design of Proposition 2.[49] We estimate demand for new cars in Madrid and compute cars' externality costs to evaluate Madrid's LEZ against alternative policies and designs. At Madrid's estimated preferences and damages, the conditions for the LEZ to be welfare-improving fail. In Section 5, we ask how far each would have to move for them to hold.

We make two simplifications to keep the exercise tractable. First, we estimate impacts around the policy implementation, which avoids projecting car prices and consumer preferences far into the future. Second, we model only households in the market for a new car, since they are the relevant margin through which the policy can induce adoption of cleaner vehicles. This leaves out households outside the new-car market that drive old restricted cars. Since the LEZ makes them worse off, our estimates understate its welfare cost. The remaining margins—reduced driving, earlier scrappage, and second-hand

[48] Multiplicity arises in our context because the buyers of a potentially restricted vehicle never exit the market, as the valuation of their durable good depends on the planner's future actions.

[49] We do not take the time-varying designs to the data—hybrids as a transition technology, and the value of commitment (Propositions 4 and 5). They govern how a LEZ should be phased in, a question that arises only when a time-invariant LEZ is worth having in the first place (see Figure 5).

transfers—are also left out.[50]

4.1 Model

We extend the model of Section 3 by allowing for heterogeneous consumer preferences over car models, in addition to heterogeneous preferences over vehicle types (gasoline, diesel, LPG/CNG, hybrid, and electric), generating richer substitution patterns across products and letting us exploit granular price variation to identify demand elasticities in the spirit of Berry, Levinsohn, and Pakes (1995).[51] This empirical specification can be viewed as a higher-dimensional counterpart of the horizontal differentiation structure in our previous theoretical analysis, where, for tractability, consumers and vehicles are located in a low-dimensional Hotelling/Salop taste space. Here, horizontal differentiation is captured by heterogeneous preferences over many make–model–type combinations.

A product is a pair (m, j) of a make–model m and one of five types j: gasoline, diesel, LPG/CNG, hybrid, or electric. From Lemma 1, the lifetime utility that individual i derives from purchasing a new car of make–model m and type j in city c at quarter t can be written as

$$U^{imjct} = \frac{1}{1-\delta^{T_{mjct}}} \left[\sum_{a=1}^{T_{mjct}} \delta^{a-1} u_{imjcta} + \delta^{T_{mjct}} z_{mjct} \right] - \frac{p_{mjt}}{1-\delta^{T_{mjct}}}, \tag{13}$$

where p_{mjt} is the purchase price at time t, which is set at the national level, δ is the discount factor, T_{mjct} is the expected scrapping age, z_{mjct} is the expected scrap value, and u_{imjcta} is the per-period flow utility from driving the vehicle at age a.

An important assumption behind equation (13), which certainly does not hold in practice, is that all individuals are expected to hold on to their cars until their scrapping age, at which point they are expected to replace them with the same car. This is because, for tractability, our theory abstracts from vertical preferences. One way to relax this assumption is to allow new-car buyers to hold on to their cars for shorter periods of time, with the terminal age T_{mjct} representing this shorter holding period—which could even vary across consumers—and the terminal value z_{mjct} representing the resale price of the second-hand unit at that time. We do require, however, that both the holding period and the resale value be invariant to policy. Our motivating evidence supports this assumption: scrappage and second-hand transfers show no significant changes. This is also consistent with the theory. As explained in Appendix D, if a LEZ lowers the value of all affected vintages by roughly the same amount, the resale price and scrappage age

[50] Section 2 finds no significant effect on scrappage or on transfers out of Madrid, and Galdon-Sanchez et al. (2023) and Tassinari (2024) find no overall effect on Madrid's traffic beyond the first few months after implementation.

[51] This framework is also used by Miravete, Moral, and Thurk (2018) and Miravete and Moral (2025) to estimate car demand in the Spanish market. Their focus is on the impact of trade-related policies on car demand and its persistence over time.

of all affected second-hand units are virtually unchanged.

Another dimension absent from equation (13) is the possibility that preferences for cleaner cars may be changing over time. As long as these trends are not affected by the policy—unlike consumer choices, which are expected to respond to it—they can also be incorporated. Accounting for both shorter holding periods and policy-invariant preference trends, we reparameterize lifetime utility as follows:[52]

$$U^{imjct} = v_{imjct} - \alpha_i \, p_{mjt} - \phi_i \, \text{LEZ}_{jct} + \epsilon_{imjct}, \tag{14}$$

where α_i is the marginal utility of income, ϕ_i is the disutility from LEZ restrictions, LEZ_{jct} is an indicator equal to one for gasoline and diesel cars in Madrid during the post-policy period, and ϵ_{imjct} is an idiosyncratic taste draw. We assume that ϵ_{imjct} follows the generalized extreme value distribution of a nested logit with a single nest containing all inside goods, and denote the within-nest correlation by ρ. We normalize the value of the outside option—not buying a new car this quarter—to zero.

The term v_{imjct} collects all remaining utility components, which are unaffected by the policy:

$$v_{imjct} = \boldsymbol{\beta}_i' \mathbf{x}_j + \zeta \, b_{jct} + \eta_m + \eta_{cj} + \eta_{jt} + \xi_{mjct}, \tag{15}$$

where $\mathbf{x}_j$ contains indicators for the five vehicle types, with heterogeneous coefficients $\boldsymbol{\beta}_i$ defined in (16); η_m, η_{cj}, and η_{jt} denote make–model, city–type, and type–quarter fixed effects; b_{jct}, with coefficient ζ, is an indicator for diesel cars in Madrid from 2017 onwards;[53] and ξ_{mjct} is a product–market unobservable, observed by consumers but not by the econometrician.[54]

The heterogeneous parameters are defined as:

$$\begin{pmatrix} \log(\alpha_i) \\ \log(\phi_i) \\ \boldsymbol{\beta}_i \end{pmatrix} = \begin{pmatrix} \mu_\alpha + \pi_\alpha \, y_i + \sigma_\alpha \nu_i \\ \mu_\phi + \pi_\phi \, y_i + \psi_\phi^1 d_i + \psi_\phi^2 d_i^2 + \sigma_\phi \nu_i \\ \boldsymbol{\pi}_\beta \, y_i \end{pmatrix}, \tag{16}$$

where y_i denotes the logarithm of the average income of consumer i's zip code, measured in thousands of euros, d_i is the logarithm of its distance to Madrid Central, and $\nu_i \sim \mathcal{N}(0,1)$ captures unobserved preference heterogeneity. Adding rich individual-level heterogeneity matters for two reasons. First, heterogeneous preferences over car types generate richer substitution patterns. Second, and specific to our setting, heterogeneous LEZ costs allow

[52] As in the theory, we are assuming that the policy shock—whether a LEZ or a tax for that matter—is not anticipated by consumers.

[53] In December 2016, the mayors of Madrid, Paris, Athens, and Mexico City pledged at the C40 Mayors Summit in Mexico City to ban diesel vehicles from their cities by 2025. The indicator absorbs the resulting shift away from diesel, which could otherwise be confounded with the effect of the LEZ.

[54] Note that the discounting terms $1 - \delta^{T_{mjct}}$ and $\delta^{T_{mjct}}$ in equation (13) are absorbed by the shocks ϵ_{imjct} and ξ_{mjct} in (14) and (15), respectively. Since these shocks are unrelated to the policy, this reduction is valid as long as the holding horizon T_{mjct} is invariant to the policy intervention, as discussed above.

the inframarginal buyers who continue buying restricted cars to bear smaller losses than those who switched would have borne.

The demand model implies a money-metric measure of consumer surplus for potential new-car buyers. In city c and quarter t, average expected consumer surplus under a given LEZ regime is the nested-logit log-sum, deflated by the marginal utility of income:

$$\mathrm{CS}^{ct}(\mathrm{LEZ}) = \int w_i \frac{1}{\alpha_i} \log\left(1 + \left(\sum_j \sum_m \exp\left(\frac{v_{imjct} - \alpha_i p_{mjt} - \phi_i \,\mathrm{LEZ}_{jct}}{1-\rho}\right)\right)^{1-\rho}\right) di, \tag{17}$$

where w_i is the share of the market's households living in consumer i's zip code, so the weights integrate to one. The LEZ enters directly through $\phi_i \,\mathrm{LEZ}_{jct}$, as consumers derive less utility from driving restricted vehicles, both lowering their surplus and inducing substitution toward cleaner alternatives at the margin. We net the surplus in (17) against the vehicle externality costs detailed below in Section 4.2.4 when evaluating welfare in the counterfactuals of Section 5.

4.2 Estimation

4.2.1 Identification

We estimate the model by the generalized method of moments, following Berry, Levinsohn, and Pakes (1995).[55] Identification of demand requires instruments for prices and either instruments for quantities or micro data linking consumers' characteristics to their choices (Berry and Haile 2014, 2024). We use both. In our setting, we require an additional assumption to identify the disutility from the LEZ restrictions, analogous to the parallel-trends assumption of a difference-in-differences design. We discuss the instrument for prices, the disutility from the LEZ restrictions, the instruments for quantities, and the micro-moments in turn.

Instrument for prices. A key challenge in demand estimation is price endogeneity, as firms may set prices in response to unobserved demand shocks ξ_{mjct}. Following Li, Xiao, and Liu (2015), we instrument for prices with κ_{mjt}, the interaction between vehicle weight and the annual average steel price. This instrument can be interpreted as a shift-share design, where the common cost shock—the steel price—constitutes the shift, and vehicle weight provides predetermined, product-level exposure shares. Identification comes from differential price movements across vehicles with different weights in response to a common input cost shock. The identifying assumption is that, conditional on fixed effects, changes in unobserved demand shocks are not systematically correlated with vehicle weight, so that heavier and lighter vehicles do not experience differential demand trends that coin-

[55] We refer the reader to Conlon and Gortmaker (2020) and Conlon and Gortmaker (2025) for detailed guides on best practices, including those we adopt here.

cide with changes in steel prices. This assumption is analogous to the exogenous-shares condition in the shift-share literature (Goldsmith-Pinkham, Sorkin, and Swift 2020).

Disutility from LEZ restrictions. We identify the parameters governing the disutility from the LEZ restrictions from differential changes in demand across vehicle types, locations, and time. Identification relies on comparing changes in demand for restricted vehicles (gasoline and diesel) relative to unrestricted vehicles (LPG/CNG, hybrid, and electric), in Madrid relative to other cities, before and after the policy implementation. The identifying assumption is that, absent the LEZ, unobserved demand shocks would have evolved similarly across vehicle types in Madrid and other similar cities over time. Conditional on the fixed effects, any differential post-policy change in demand for gasoline and diesel vehicles in Madrid is therefore attributed to the LEZ restrictions. This identification strategy is analogous to a difference-in-differences design embedded in a discrete-choice demand model.

Instruments for quantities. We complement the previous variation with five additional instruments for quantities. Following Gandhi and Houde (2019), we build instruments that reflect the degree of differentiation of each product in its market. The first two instruments are the differences between a product's cost shifter, κ_{mjt}, and the average cost shifter of all products in its market and of the products of its type, $\frac{1}{|\mathcal{J}_{ct}|}\sum_{(m',j')\in\mathcal{J}_{ct}}(\kappa_{mjt}-\kappa_{m'j't})$ and $\frac{1}{|\mathcal{J}_{jct}|}\sum_{(m',j)\in\mathcal{J}_{jct}}(\kappa_{mjt}-\kappa_{m'jt})$, where $\mathcal{J}_{ct}$ is the set of products in market ct and $\mathcal{J}_{jct}$ its subset of type j.[56] We build two additional instruments by counting the number of products available in the market within one sample standard deviation, σ_x, of a product's characteristic, x_{mjct}, given by $\sum_{(m',j')\in\mathcal{J}_{ct}\setminus\{(m,j)\}}\mathbb{1}\{|x_{m'j'ct}-x_{mjct}|<\sigma_x\}$, for x given by car weight and power. An important assumption behind the validity of these two instruments is that unobserved demand shocks, ξ_{mjct}, have mean zero given the weight and power of all products in the market and the fixed effects. This assumption is standard in the literature (Berry and Haile 2014). Finally, we use a Madrid post-policy indicator as an additional instrument. Assuming that ξ_{mjct} does not explain the post-policy shift in aggregate new-car demand in Madrid relative to other cities, the model attributes that shift to the policy, which helps identify substitution between new-car purchases and the outside option.

Micro-moments. We use micro-moments to capture heterogeneity in price sensitivity, LEZ disutility, and car-type preferences. We divide consumers into two groups according to whether their zip code's average income exceeds €20,000. Following Petrin (2002), we match moments generated by the model to their empirical counterparts. For each income group, we match the purchase probability, the average price paid by buyers, and the shares of diesel, LPG/CNG, hybrid, and electric cars among buyers. We also match the share of restricted cars among buyers in each income group, separately in Madrid

[56] These instruments follow the formula for recentered instruments proposed by Borusyak, Bravo, and Hull (2025).

and the other cities and before and after implementation, and the share of unrestricted cars among Madrid buyers in each income group. The last four moments match the pre–post change in the log odds of a restricted purchase among Madrid buyers within distance bands of 0–3, 3–10, 10–20, and 20 km or more to MC. Appendix A, Table A.2, lists the 26 micro-moments. We estimate the preference parameters jointly using these micro-moments and the aggregate moment conditions.

4.2.2 Estimation sample

We estimate the model using the new-car registration data introduced in Section 2. As in that section, we exclude the zip codes inside Madrid's LEZ, whose residents are exempt from the restrictions. We aggregate the data to the make–model–type–city–quarter level and estimate over 2015–2019.[57] The estimation sample comprises 677 make–model×type products across 126 markets—the 7 provinces (Madrid, Barcelona, Valencia, Seville, Alicante, Málaga, and Murcia) over 18 quarters—for 46,603 product–market observations.[58] We set the size of each city–quarter market to 1.1 times the maximum number of sales observed across all sample years for that city and quarter-of-the-year.[59] The 7-city sample accounts for 48.2% of national new-car sales over this period, with Madrid alone at 14.8%. Table 3 summarizes the estimation sample.

Conventional internal-combustion-engine vehicles dominate the market: gasoline and diesel cars account for 93.5% of new-car sales, LPG/CNG for 0.9%, hybrids (Hs) for 5.4%, and electric vehicles (EVs) for 0.2%. Not buying a new car—the outside option—accounts for 14.6% of the market. Hybrids and EVs are more expensive: sales-weighted average prices are €25 thousand for a conventional combustion car, €20 thousand for an LPG/CNG car, €32 thousand for a hybrid, and €49 thousand for an electric vehicle. High-income zip codes account for 7.0% of the market–zip-code cells in our sample.

4.2.3 Demand estimates

Table 4, Panel A, reports the estimated demand parameters. They imply an average own-price elasticity of −5.27 (median −4.63), in line with previous studies for Spanish, European, and U.S. automobile markets.[60] Price sensitivity declines with income. The

[57] We exclude the two quarters between the announcement and implementation of Madrid Central.

[58] We work at the province level but refer to these units as "cities" throughout, for consistency with the theory of Section 3.

[59] We follow Miller and Weinberg (2017), who set market size to a fixed multiple of the maximum observed sales in each market. Two details differ: our multiplier, 1.1, is smaller than their 1.5; and we take the maximum within each city *and* quarter-of-the-year rather than within each region alone, so that market size tracks the pronounced seasonality of car sales. As they note, the nesting parameter limits the sensitivity of estimated elasticities to this assumption, which applies here too given our nested-logit structure.

[60] Berry, Levinsohn, and Pakes (1995) and Grieco, Murry, and Yurukoglu (2024) report own-price elasticities ranging from −3.1 to −6.8 for the U.S. market. For France, D'Haultfœuille, Durrmeyer, and Février (2019) report average elasticities between −3.9 and −6.4. For Spain, Miravete, Moral, and Thurk (2018) and Miravete and Moral (2025) report elasticities between about −2 and −7.

Table 3: Estimation sample: descriptive statistics

Panel A: Sample		
Product–market observations		46,603
Markets (city × quarter)		126
Cities		7
Quarters (2015–2019)		18
Make–models		360
Products (make–model × type)		677
Panel B: Market shares (%, avg. per market)		
	Of market	Of sales
Gasoline	43.1	50.5
Diesel	36.8	43.1
LPG	0.76	0.9
Hybrid (H)	4.59	5.4
Electric (EV)	0.16	0.2
Outside option (no purchase)	14.6	—
Panel C: Prices (000s €, sales-weighted mean / median)		
All products		25.0 / 23.4
Gasoline		21.5 / 20.3
Diesel		28.3 / 27.1
LPG		20.3 / 18.6
H		31.9 / 29.3
EV		48.9 / 39.6
Panel D: Consumers		
Market × zip code cells		30,420
High-income zip codes (%)		7.0
High-income households (%)		9.2

Notes: Natural-person new-car registrations, excluding leasing, over the 7-city sample, 2015–2019, excluding the second and third quarters of 2018 and zip codes inside Madrid's LEZ. Panel B reports market shares averaged across markets, both as a share of the whole market (summing to 100 with the outside option) and as a share of sales; Panel C reports sales-weighted prices in thousands of euros. LPG denotes LPG/CNG. In Panel D, "high-income" denotes zip codes with average income above €20,000.

estimated model reproduces the 26 micro-moments closely (Appendix A, Table A.2).

The ratio ϕ_i/α_i gives the willingness to pay (WTP) upon purchasing a new car to avoid LEZ restrictions, expressed in thousands of euros. Figure 6 shows WTP running from €4,052 at the 25th WTP percentile to €6,640 at the 75th, with a median of €5,190.

The nesting parameter is estimated at $\rho = 0.78$, which implies that consumers substitute mostly among cars rather than out of the new-car market, so that buyers displaced from a restricted vehicle tend to buy a different car instead.

Table 4: Estimation results

Panel A: Demand parameters						
Price (α)	μ_α	-2.268	π_α	-0.395	σ_α	0.098
		(0.190)		(0.064)		(0.126)
LEZ program (ϕ)	μ_ϕ	1.644	π_ϕ	-1.231	σ_ϕ	0.413
		(7.305)		(1.243)		(9.903)
	ψ_ϕ^1	0.193	ψ_ϕ^2	-0.055		
		(1.133)		(0.238)		
Nesting (ρ)	ρ	0.780				
		(0.038)				
Car types (β)	π_{β_G}	0.117	π_{β_D}	-0.039	π_{β_L}	0.166
		(0.178)		(0.185)		(0.230)
	π_{β_H}	0.332	$\pi_{\beta_{EV}}$	0.539		
		(0.188)		(0.305)		
Panel B: Other statistics						
Average own-price elasticity				-5.27		
WTP for LEZ avoidance (ϕ_i/α_i)		Q25: 4.05		Q50: 5.19		Q75: 6.64
Panel C: Lifetime externality cost (thousands EUR)						
Type		Mkt. Share (%)		Mean		Median
Gasoline		44.1		11.141		10.977
Diesel		16.6		12.493		12.261
LPG		5.9		10.934		10.668
Hybrid		14.1		9.072		8.680
EV		0.9		3.783		3.783

Notes: Panel A reports demand parameter estimates; prices, WTP, and the income levels inside logarithms are in thousands of euros. Vehicle-type coefficients are the components of $\boldsymbol{\pi}_\beta$ in (16), with L denoting LPG/CNG (LPG in Panel C). Panel C reports mean and median lifetime externality costs by vehicle type for Madrid in the post-policy period; EV costs consider a CO_2 rate of 5.7 g/km and local damages (accidents and local pollution, excluding noise); H costs consider their observed tailpipe CO_2 rates and, for local damages, a weighted average between gasoline and EV cost rates with a 63% gasoline share.

4.2.4 Externality costs

To value vehicle externalities, we follow Allcott et al. (2024) and aggregate global air pollution (i.e., CO_2 emissions), local air pollution, noise, and accidents over the vehicle's lifetime. Like them, we abstract from discounting, taking marginal damages to be constant in real terms over the vehicle's lifetime. Damages are estimated at the city level without accounting for potential differences in air-pollution exposure across neighborhoods. Also, we assume marginal damages to be invariant to the magnitude of the intervention, which is appropriate for global externalities but probably conservative for local externalities. Thus, the lifetime externality cost of product (m, j), in thousands of

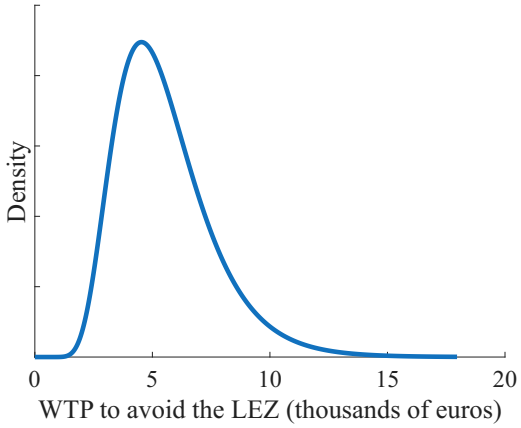


Figure 6: Distribution of WTP for LEZ avoidance

Notes: Distribution of willingness to pay for avoiding LEZ restrictions (ϕ_i/α_i, in thousands of euros) across Madrid zip codes and taste draws, weighted by households. Appendix Figure A.2 shows the distributions of α_i and ϕ_i.

euros, is

$$\text{Externality cost}_{mj} = \frac{1}{1{,}000}\left(\underbrace{e_{mj} \times L \times \frac{\text{SCC}}{1{,}000{,}000}}_{\text{CO}_2} + \underbrace{\ell_j \times \frac{L}{100}}_{\text{local damages}}\right), \tag{18}$$

where e_{mj} is the vehicle's observed tailpipe CO_2 rate in g/km, $L = 240{,}000$ km is the assumed lifetime mileage (Green NCAP 2022),[61] SCC is the social cost of carbon at €204 per ton (U.S. Environmental Protection Agency 2023), and ℓ_j is the local-damage rate (accidents, local pollution, and noise) in euros per 100 km, which we take as €2.05 for gasoline, €2.77 for diesel, €2.10 for LPG/CNG, and €1.46 for electric cars, from European Commission et al. (2020). Following Plötz et al. (2021), we assume that hybrids operate on gasoline for 63% of kilometers driven and on electricity for the rest, so $\ell_{\text{H}} = \frac{63}{100}\ell_{\text{gas}} + \left(1 - \frac{63}{100}\right)\ell_{\text{EV}} =$ €1.83. EVs have no tailpipe emissions, so we assign them the current carbon footprint of electricity generation in Spain, which gives an emissions rate of 5.7 g/km (30 g/kWh $\times$ 19 kWh/100 km; Red Eléctrica de España 2026).[62]

Panel C of Table 4 and Figure 7 report lifetime externality costs by vehicle type for Madrid in the post-policy period. Diesel exhibits the highest externality cost, at €12.49 thousand on average, followed by gasoline at €11.14 thousand, LPG/CNG at €10.93 thousand, hybrids at €9.07 thousand, and EVs at €3.78 thousand. We assume that a household choosing the outside option of not buying a new car avoids the total lifetime

[61] This is a fleet-average lifetime and abstracts from heterogeneity in vehicle survival: cars scrapped early accumulate fewer kilometers, and hence lower lifetime emissions, than this figure implies.

[62] This rate implicitly assumes that the vehicle is charged uniformly across hours; charging concentrated during the day (night), when the grid is cleaner (dirtier), would imply a lower (higher) emissions rate.

externality cost of the car it would otherwise have bought. In reality, however, this is true only if the household stops driving altogether, for example by switching to public transit. If instead the outside option is to postpone the purchase of a new car for some time, during which the household continues driving its current car, we would be overestimating the emission benefits of the LEZ.[63]

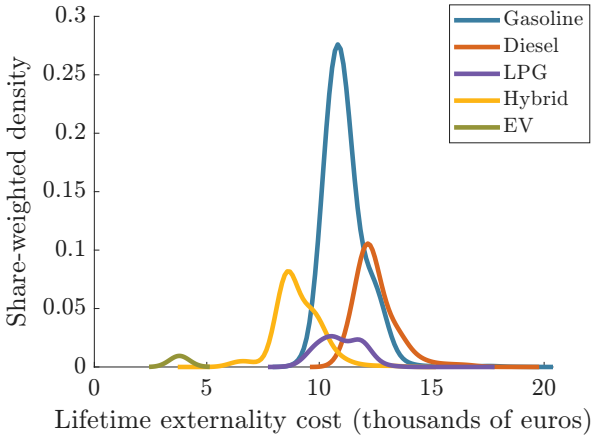


Figure 7: Distribution of lifetime externality costs by vehicle type

Notes: Kernel density estimates of the lifetime externality cost across make–model–type observations, weighted by market shares in Madrid post-policy. LPG denotes LPG/CNG.

4.3 Marginal welfare decomposition

Before turning to the full counterfactual analysis, we characterize the *marginal* welfare effect of the LEZ by taking derivatives of welfare with respect to the LEZ intensity. To ease exposition, we denote gasoline and diesel cars jointly by G. We measure the LEZ intensity by a scalar $\lambda \geq 0$ that multiplies the LEZ disutility of each restricted vehicle, so that a restricted car loses utility $\phi_i \lambda$ under a LEZ.[64] Given a LEZ's inefficient rationing, the derivative of welfare with respect to λ can be decomposed into two components. By the envelope theorem, the private cost is borne by the inframarginal drivers who continue to purchase restricted cars, while the social benefit comes from the marginal drivers who switch to cleaner alternatives. We first quantify where switchers go (the diversion ratios), and then combine both forces into a first-order welfare decomposition.

[63] For some households with multiple cars, taking the outside option could instead mean keeping fewer cars and using them more intensively.

[64] A larger restricted area, for example, raises the intensity because it increases detour distances and the share of a driver's trips affected by the restriction.

4.3.1 Diversion ratios

We compute diversion ratios from a marginal increase in LEZ intensity. The diversion ratio to type k is the share of the marginal loss in G demand that k absorbs,

$$DR_{G\to k} = \frac{\partial S_k/\partial\lambda}{-\,\partial S_G/\partial\lambda}, \tag{19}$$

where S_k and S_G are the market shares of type k and of restricted (i.e., G) cars, respectively.[65]

Table 5: Diversion of G demand: nested logit, SDID, and plain logit

	Nested logit ($\rho = 0.78$)	SDID	Plain logit ($\rho = 0$, $\propto$ shares)
To hybrid (H)	56.0	40.2	36.0
To electric (EV)	3.0	3.5	2.2
To LPG	26.9	37.8	15.2
To outside option	14.1	18.5	46.6

Notes: Model columns report the fraction of the marginal loss in gasoline and diesel (G) demand diverted to each alternative under a marginal tightening of the LEZ, computed for Madrid in the post-policy period. "Nested logit" uses the estimated $\rho = 0.78$; SDID uses column (4) of Table 2; "Plain logit" assumes homogeneous preferences and sets $\rho = 0$, so diversion is proportional to market shares. LPG denotes LPG/CNG.

We report the estimated diversion ratios in Table 5. A marginal tightening of the LEZ sends 56.0% of lost G demand to hybrids, 3.0% to EVs, and 26.9% to LPG/CNG, while only 14.1% leaves the new-car market. For comparison, the SDID estimates in column (4) of Table 2 imply that 40.2% of the fall in gasoline and diesel registrations reappears as hybrids, 3.5% as EVs, and 37.8% as LPG/CNG, while 18.5% leaves the new-car market.[66] For reference, the table also reports diversion implied by a plain logit with homogeneous preferences ($\rho = 0$). Its derivatives collapse to $\partial S_k/\partial\lambda = \phi\, S_k\, S_G$, so diversion is proportional to shares and the outside option absorbs 46.6%, compared with 18.5% in the SDID estimates and 14.1% in our model.

[65] Individual i's share of type k is $S_{k,i} = \sum_{(m,j)\in k} s_{imj}$, and the market share in (19) is $S_k = \int w_i\, S_{k,i}\, di$, with w_i being the weights in (17); $S_{G,i}$ and S_G are defined the same way over the restricted products G. Let $S_{\text{car},i} = \sum_{m,j} s_{imj}$ denote the share of all inside goods and $S_{G|\text{car},i} = \sum_{(m,j)\in G} s_{mj|\text{car},i}$ the share of G within the nest. The cross-derivatives are then $\partial S_{k,i}/\partial\lambda = \phi_i\, S_{k,i}\big[\rho\, S_{G|\text{car},i} + (1-\rho)\, S_{G,i}\big]/(1-\rho)$ for $k \in \{\text{LPG/CNG}, \text{H}, \text{EV}\}$ and $\partial S_{0,i}/\partial\lambda = \phi_i\, S_{0,i}\, S_{G,i}$ for the outside option $k = 0$, so the total loss in G demand is $-\partial S_{G,i}/\partial\lambda = \phi_i\, S_{G,i}\, \varsigma_i$, where ϕ_i is the LEZ disutility and $\varsigma_i \equiv \big[1 - \rho\, S_{G|\text{car},i} - (1-\rho)\, S_{G,i}\big]/(1-\rho)$ is the diversion intensity.

[66] The SDID diversion ratios divide the estimated effect of MC on each exempt type's registrations by the absolute value of its effect on gasoline and diesel registrations. The outside-option ratio is the residual.

4.3.2 The marginal welfare effect

Average consumer surplus (CS) is given by (17), now with the LEZ disutility of each restricted product scaled by λ. By the envelope theorem, the switching of marginal drivers has only a second-order effect on consumer surplus, so the marginal effect on surplus is given by the direct disutility borne by *inframarginal* restricted-car buyers:

$$\frac{\partial \,\mathrm{CS}}{\partial \lambda} = -\int w_i \,\frac{\phi_i}{\alpha_i}\, S_{G,i}\, di. \tag{20}$$

This term is negative and scales with the willingness-to-pay–weighted share of consumers who continue to buy restricted cars.

The emission benefit, by contrast, comes entirely from the switchers. Let $E_{k,i}$ be driver i's average lifetime externality cost among the cars of type k, weighted by their choice probabilities, with $E_{0,i} = 0$ for the outside option. Note that, because driver i's choice probabilities among inside goods satisfy the independence of irrelevant alternatives (IIA) and the LEZ lowers the utility of all restricted cars by the same amount, the restriction leaves the individual-level shares of cars within each type unchanged. As a result, $E_{k,i}$ does not depend on λ. The marginal reduction in externality costs is then

$$-\frac{\partial E}{\partial \lambda} = \int w_i \underbrace{\left(-\frac{\partial S_{G,i}}{\partial \lambda}\right)}_{\substack{\text{fall in demand}\\ \text{for restricted cars}}} \underbrace{\Delta E_i}_{\substack{\text{av. emission saving}\\ \text{per displaced car}}} di, \tag{21}$$

where $\Delta E_i \equiv \sum_{k\in\{\mathrm{LPG/CNG,H,EV,0}\}} DR_{G\to k,i}\left(E_{G,i} - E_{k,i}\right)$, with $DR_{G\to k,i}$ driver i's diversion ratio, defined as in (19) with their own shares. The fall in demand for restricted cars is $-\partial S_{G,i}/\partial\lambda = \phi_i\, S_{G,i}\,\varsigma_i$, where $\varsigma_i \equiv \left[1 - \rho\, S_{G|\mathrm{car},i} - (1-\rho)\, S_{G,i}\right]/(1-\rho)$ is the diversion intensity of footnote 65.

With welfare defined as $\bar{W} = \mathrm{CS} - E$, the marginal welfare effect of the LEZ's intensity is

$$\frac{\partial \bar{W}}{\partial \lambda} = \underbrace{-\int w_i \,\frac{\phi_i}{\alpha_i}\, S_{G,i}\, di}_{\text{private cost (inframarginal)}} + \underbrace{\int w_i\, \phi_i\, S_{G,i}\, \varsigma_i\, \Delta E_i\, di}_{\text{emission benefit (marginal switchers)}}. \tag{22}$$

At the current LEZ ($\lambda = 1$), the decomposition gives a private cost of €3.22 thousand and an emission benefit of €0.51 thousand per potential new-car buyer. This results in a negative net marginal effect of −€2.71 thousand.[67]

Equation (22) shows why the two sides sit so far apart. The cost scales with the drivers who do *not* switch, given by the inframarginal share, $S_{G,i}$, weighted by their

[67] This figure is consistent with the total welfare change of −€3.20 thousand computed in Section 5, which is the integral of the marginal effect from $\lambda = 0$ to 1. Due to the convexity of the welfare function, we expect the marginal effect at $\lambda = 1$ to be lower in magnitude than the integral.

willingness to pay, ϕ_i/α_i, while the benefit scales with those who *do*, given by $\phi_i S_{G,i}\varsigma_i$. This is the inefficiency highlighted in Section 3. To move a few marginal drivers to cleaner cars, the LEZ imposes a welfare cost on every inframarginal driver who keeps a restricted car. Quantitatively, the share of inframarginal drivers is large—at the current LEZ, 60.7% of potential buyers in Madrid still buy a gasoline or diesel car—and their average willingness to pay to avoid the LEZ is high (€5,305). On the other hand, the demand displaced at the margin is small, and it flows mostly to hybrids (56.0%) and LPG/CNG (26.9%), whose lifetime externality costs are not that different from those of gasoline and diesel cars.

The analysis so far evaluates the marginal effect at the current intensity ($\lambda = 1$). The theory of Section 3 shows that welfare is convex in LEZ intensity, so a larger LEZ might in principle reverse the sign of the marginal effect. We show next that, at our estimates, welfare decreases with intensity at every λ.

Note that, driver by driver, both terms of (22) are proportional to $\phi_i S_{G,i}$, the disutility borne by inframarginal drivers. Each util of that disutility costs driver i an amount $1/\alpha_i$ in money and displaces ς_i units of demand for restricted cars. Each displaced car saves ΔE_i in money. A euro of driver i's private cost therefore buys $\alpha_i\,\varsigma_i\,\Delta E_i$ euros of emission savings. Let $\tilde{w}_i \equiv w_i\,\phi_i\,S_{G,i}$. The marginal welfare effect in (22) is negative if and only if the LEZ buys less than one euro of emission savings per euro of private cost:

$$\frac{\int \tilde{w}_i\,\varsigma_i\,\Delta E_i\,di}{\int \tilde{w}_i\,\alpha_i^{-1}\,di} \;<\; 1. \tag{23}$$

At $\lambda = 1$, each euro of private cost buys 16 cents of emission savings.

Two results, proved in Appendix F, extend the inequality in equation (23) to every LEZ intensity. First, we show that $\Delta E_i(1)/(1-\rho) \geq \varsigma_i(\lambda)\,\Delta E_i(\lambda)$ for every driver i and intensity λ, where $\Delta E_i(1)$ is driver i's emission savings per displaced car at the current intensity (Appendix F.1). This implies that, at any intensity, a euro of driver i's private cost buys at most $\alpha_i\,\Delta E_i(1)/(1-\rho)$ euros of emission savings. Second, we show that, at our estimates, this ceiling is below one euro for virtually all drivers (Appendix F.2).[68] Condition (23) therefore holds at every $\lambda \geq 0$, so welfare decreases with LEZ intensity everywhere.

5 Counterfactual Analysis

As Section 3 makes clear, unlike a tax, which is always worth considering, a LEZ is worth considering only under certain conditions: when pollution damages are sufficiently

[68] Weighted by $\tilde{w}_i$ at $\lambda = 1$, its average is 55 cents, with a 99.9th percentile of 78 cents (Appendix Figure F.1). The few drivers whose ceiling exceeds one have a high LEZ disutility and leave restricted cars quickly, so at high intensities most of the weight falls on drivers whose ceiling is below one. Appendix F.2 shows that these exceptions cannot overturn the inequality of equation (23).

large or substitution toward cleaner cars is sufficiently strong. Section 4.3 showed that neither condition holds at our baseline estimates. This section puts magnitudes on that conclusion by examining how far vehicle external costs and consumer preferences must depart from their baseline estimates for a LEZ to deliver positive benefits.

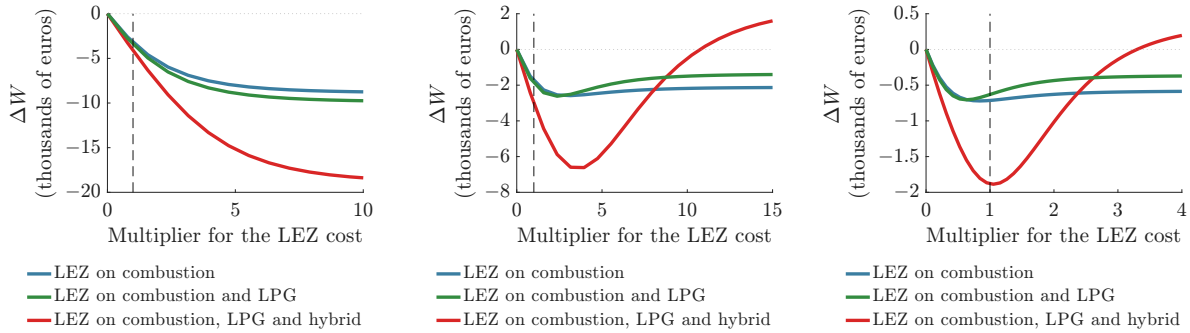
5.1 Welfare effects of LEZ policies

We evaluate five scenarios: (i) no LEZ; (ii) the current LEZ, which restricts gasoline and diesel vehicles; (iii) a LEZ that also restricts LPG/CNG; (iv) a LEZ that additionally restricts hybrids (Hs); and (v) as a benchmark, a Pigouvian emissions tax in place of the LEZ. Panel A of Table 6 reports market shares by vehicle type, externality costs, and the changes in consumer surplus and in total welfare relative to no LEZ,[69]

$$\Delta W = (\text{CS} - \text{E ext} + \text{Tax}) - (\text{CS} - \text{E ext})_0.$$

Madrid's LEZ successfully shifts buyers from gasoline and diesel cars toward unrestricted cars. However, consumer surplus falls by €3.68 thousand against a €0.48 thousand reduction in externality costs, for a net welfare change of −€3.20 thousand. Note that this calculation counts only the new-car market, so it is a lower bound on the overall loss.

Panel (a) of Figure 8 plots the welfare change as the LEZ intensity λ rises from 0 to 10, where $\lambda = 1$ is the current policy. Under the current LEZ design, welfare falls at every intensity, as Section 4.3 predicts.



(a) Baseline (b) 4× damages (c) 4× price sensitivity

Figure 8: Welfare under the LEZ, by policy intensity

Notes: Change in welfare as the LEZ disutility ϕ_i is scaled by a multiplier running from 0 to 10 in Panel (a), from 0 to 15 in Panel (b), and from 0 to 4 in Panel (c), where price sensitivity is scaled by 4; the vertical line marks the current policy. Panel (a) uses the estimated parameters, Panel (b) scales externality damages by 4, and Panel (c) scales α_i and ϕ_i by 4. Each panel shows three LEZ designs: restricting gasoline and diesel cars, also restricting LPG/CNG, and additionally restricting hybrids.

[69] Since consumer surplus is identified only up to an additive constant, we report its change relative to no LEZ.

Table 6: Counterfactual results

	Fleet	G	LPG	H	EV	ΔCS	E ext.	Tax	ΔW
Panel A: Baseline									
No LEZ	0.840	0.734	0.027	0.074	0.005	–	9.43	–	–
LEZ on combustion	0.817	0.607	0.059	0.141	0.009	-3.68	8.95	–	-3.20
LEZ on combustion and LPG	0.815	0.633	0.024	0.149	0.009	-3.91	8.93	–	-3.41
LEZ on combustion, LPG and hybrid	0.811	0.703	0.026	0.071	0.010	-4.48	9.06	–	-4.11
Pigouvian tax	0.783	0.650	0.026	0.093	0.013	-9.00	8.55	8.55	0.43
Panel B: Externality damages scaled by 4									
No LEZ	0.840	0.734	0.027	0.074	0.005	–	37.72	–	–
LEZ on combustion	0.817	0.607	0.059	0.141	0.009	-3.68	35.81	–	-1.77
LEZ on combustion and LPG	0.815	0.633	0.024	0.149	0.009	-3.91	35.72	–	-1.90
LEZ on combustion, LPG and hybrid	0.811	0.703	0.026	0.071	0.010	-4.48	36.24	–	-3.00
Pigouvian tax	0.566	0.284	0.014	0.121	0.147	-29.48	19.46	19.46	8.24
Panel C: Price sensitivity scaled by 4									
No LEZ	0.840	0.734	0.027	0.074	0.005	–	9.43	–	–
LEZ on combustion	0.776	0.211	0.185	0.360	0.019	-2.32	7.82	–	-0.71
LEZ on combustion and LPG	0.761	0.247	0.008	0.483	0.023	-2.66	7.40	–	-0.63
LEZ on combustion, LPG and hybrid	0.706	0.549	0.019	0.060	0.078	-3.91	7.40	–	-1.88
Pigouvian tax	0.574	0.302	0.015	0.128	0.129	-7.45	5.08	5.08	1.98

Notes: Monetary values for Madrid in thousands of euros per potential new-car buyer. Each panel reports the same five scenarios: no LEZ, the current LEZ restricting gasoline and diesel (combustion) cars, a LEZ also restricting LPG/CNG (LPG), a stricter LEZ additionally restricting hybrids, and a Pigouvian emissions tax. Fleet: share of potential buyers purchasing a new car. G, LPG, H, and EV: whole-market shares of gasoline and diesel cars combined, LPG/CNG cars, hybrids, and electric vehicles. ΔCS: change in consumer surplus relative to no LEZ within the same panel, per potential new-car buyer, including those who opt for the outside option. E ext: externality cost. Tax: revenue from the Pigouvian tax. ΔW: change in total welfare relative to no LEZ within the same panel. Panel B rescales damages only, so the shares in its no-LEZ and three LEZ rows are unchanged from Panel A; its Pigouvian-tax row differs, because the optimal tax rate scales with damages. Panel C scales α_i and ϕ_i by 4.

5.2 When could the LEZ work?

In this subsection, we explore conditions under which a LEZ could generate positive welfare gains. We discuss two changes that could make a LEZ more effective: higher externality damages and stronger substitution across cars.

5.2.1 Higher externality damages

We multiply externality damages by 4. Our estimates imply only a modest damage gap between gasoline cars and hybrids, so the emission savings of the LEZ are small. Higher damages increase the benefits of the LEZ. Because only the valuation of damages changes, the three LEZ rows in Panel B of Table 6 have the same market shares as in Panel A. At its current intensity, the welfare effect of the existing LEZ improves but remains negative, while the stricter designs still do worse.

Panel (b) of Figure 8 considers a more intense LEZ. As our theory predicts, welfare is quasi-convex in the LEZ's intensity. With higher damages, the curve becomes U-shaped.

A tighter restriction initially destroys surplus faster than it generates emission savings, but eventually the latter begins to dominate. A LEZ that only restricts gasoline and diesel cars reaches its minimum near −€2.59 thousand at an intensity of about 3 times the current intensity and then rises. The stricter design that also restricts hybrids crosses zero only at roughly 11 times the current intensity. This threshold shows how unlikely a LEZ is to generate welfare gains under Madrid's conditions, even before counting the losses imposed on owners of second-hand restricted cars.[70]

5.2.2 More price-elastic demand

The performance of a LEZ also depends on how much demand it displaces and on how much surplus that displacement destroys. Both are governed by $\alpha_i/(1-\rho)$, which sets the slope of the residual demand curve.[71] We scale it by 4, raising the mean own-price elasticity from −5.27 to −19.74, well outside the range previous studies report for automobile markets. In particular, we scale both α_i and ϕ_i by 4, so that willingness to pay to avoid the LEZ, ϕ_i/α_i, stays at its baseline estimate.[72]

Panel C of Table 6 reports the result. At the current intensity, all three LEZ designs still deliver negative net benefits. Since ϕ_i/α_i stays the same, a driver who keeps a restricted car loses as much as in the baseline counterfactual. However, many more drivers switch. At the estimated α_i, the LEZ leaves most of the gasoline and diesel fleet in place. At the scaled α_i, the same LEZ removes most gasoline and diesel cars from the fleet, yet the surplus those buyers give up is rather small.

Figure 8 sets the two exercises—higher externality damages and flatter demand—side by side. In both, welfare first falls and then rises, consistent with the bang-bang logic of Section 3. The design that also restricts hybrids eventually raises welfare in both, but the exercises differ in how intense the LEZ must be. More serious damages make each switch more valuable but leave substitution just as difficult, so this design raises welfare only beyond 11 times the current intensity. Flatter demand makes substitution cheaper, so the same design raises welfare at a much lower intensity.[73]

[70] Note that only in a scenario with welfare gains like this one would it be worth considering the benefits of treating hybrids as a transition technology.

[71] The same object governs substitution between cars. Raising the price of car k by one thousand euros moves the share of any other car j by $\alpha_i\, s_{ij}\left[\rho\, s_{k|\text{car},i} + (1-\rho)\, s_{ik}\right]/(1-\rho)$, the same bracket that governs diversion in footnote 65: a price change and a disutility change both shift utility.

[72] We re-invert the mean utilities so the no-LEZ fleet reproduces its estimated shares.

[73] Note that Panels (b) and (c) of Figure 8 show that restricting gasoline/diesel cars, LPG/CNG, and hybrids can generate positive welfare gains even when restricting gasoline/diesel cars alone does not. The analogous outcome for gasoline/diesel cars and hybrids is ruled out by our theory model with uniformly distributed consumers (see, for example, Propositions 1 and 2). This is explained by the richer substitution patterns allowed by the empirical demand system. Importantly, the same possibility can arise in our theoretical framework with a non-uniform distribution of consumers, while preserving the quasi-convexity of the welfare function. See Appendix E for an example.

6 Conclusions

We have shown, both theoretically and empirically, that LEZs differ fundamentally from price-based instruments (e.g., taxes, tradable quotas). A tax on polluting cars, no matter how small, is always socially beneficial as a way to induce drivers to switch toward cleaner vehicles, unless, of course, it is set too far above its Pigouvian level. A LEZ that restricts polluting cars, by contrast, may not be worth considering in the first place. It must be stringent enough to induce substantial switching while leaving sufficiently few inframarginal users affected by the restriction; otherwise, no LEZ may be preferable. Our estimates suggest that this was the case for Madrid's LEZ. Although the policy shifted some buyers toward cleaner cars, the surplus losses it generated far exceeded the gains from the emissions it saved. Our counterfactual analysis clarifies what it would take for a LEZ in Madrid to be welfare-improving: LPG/CNG and hybrids would also have to be restricted, externality damages would have to be much larger, or the residual demand for polluting vehicles sufficiently flat to generate stronger substitution toward cleaner alternatives.

These findings also speak to the role of combining LEZs with other instruments. In Weitzman's (1974) prices-versus-quantities framework, for example, combining instruments may be optimal because the planner faces a friction: policy cannot be continuously adjusted as new information arrives. That friction is absent from our framework. If price-based instruments are available, there is no intrinsic reason to combine them with an inefficient rationing instrument such as a LEZ, unless they cannot be set at their optimal levels. In such a case, prices can play a supporting role by reducing the number of inframarginal users. By making polluting vehicles less attractive, or cleaner vehicles easier to adopt, they may help satisfy the necessary conditions under which a LEZ becomes worth considering.

While our focus in this paper has been strictly on a LEZ's ability to induce drivers to switch toward cleaner vehicles, it is important to recognize that such policies may also generate local benefits for residents and businesses within the restricted area, including improvements in urban amenities. To the extent that these benefits are quantitatively important, the relatively limited geographic scope of Madrid's LEZ may have represented an adequate way to strike a balance between local welfare gains and city-wide welfare losses associated with restricting vehicle access. Quantifying this broader tradeoff, however, is beyond the scope of this paper.

References

Adda, Jérôme and Russell Cooper. 2000. "Balladurette and Juppette: A Discrete Analysis of Scrapping Subsidies." *Journal of Political Economy* 108 (4):778–806.

Allcott, Hunt, Reigner Kane, Maximilian S. Maydanchik, Joseph S. Shapiro, and Felix Tintelnot. 2024. "The Effects of 'Buy American': Electric Vehicles and the Inflation Reduction Act." NBER Working Paper 33032, National Bureau of Economic Research.

Almagro, Milena, Felipe Barbieri, Juan Camilo Castillo, Nathaniel G. Hickok, and Tobias Salz. 2024. "Optimal Urban Transportation Policy: Evidence from Chicago." NBER Working Paper 32185, National Bureau of Economic Research.

Arkhangelsky, Dmitry, Susan Athey, David A. Hirshberg, Guido W. Imbens, and Stefan Wager. 2021. "Synthetic Difference-in-Differences." *American Economic Review* 111 (12):4088–4118.

Ater, Itai, Brad Ross, Adi Shany, and Shoshana Vasserman. 2026. "Road Pricing under Heterogeneity." NBER Working Paper 35887, National Bureau of Economic Research.

Ausubel, Lawrence M. and Raymond J. Deneckere. 1989. "Reputation in Bargaining and Durable Goods Monopoly." *Econometrica* 57 (3):511–531.

Aydin, Eren, Markus Gehrsitz, and Christian Traxler. 2024. "Greener Fleet, Cleaner Air: How Low Emission Zones Reduce Pollution." IZA Discussion Papers 17144, IZA Network @ LISER.

Ayuntamiento de Madrid. 2018. "Ordenanza de Movilidad Sostenible." Boletín Oficial de la Comunidad de Madrid, núm. 253, pp. 130–251. Approved by the City Council on 5 October 2018.

Bagnoli, Mark and Ted Bergstrom. 2005. "Log-concave Probability and its Applications." *Economic Theory* 26 (2):445–469.

Barahona, Nano, Francisco Gallego, and Juan-Pablo Montero. 2020. "Vintage-Specific Driving Restrictions." *Review of Economic Studies* 87:1646–1682.

Baumol, William J. and Wallace E. Oates. 1988. *The Theory of Environmental Policy*. Cambridge University Press.

Bernardo, Valeria, Xavier Fageda, and Ricardo Flores-Fillol. 2021. "Pollution and Congestion in Urban Areas: The Effects of Low Emission Zones." *Economics of Transportation* 26–27:100221.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4):841–890.

Berry, Steven T. and Philip A. Haile. 2014. "Identification in Differentiated Products Markets Using Market Level Data." *Econometrica* 82 (5):1749–1797.

———. 2024. "Nonparametric Identification of Differentiated Products Demand Using Micro Data." *Econometrica* 92 (4):1135–1162.

Beshir, H. A. and E. Fichera. 2022. "And Breathe Normally: The Low Emission Zone impacts on health and well-being in England." Health, Econometrics and Data Group (HEDG) Working Papers 22/09, HEDG, c/o Department of Economics, University of York.

Borusyak, Kirill, Mauricio Caceres Bravo, and Peter Hull. 2025. "Estimating Demand with Recentered Instruments." *arXiv preprint arXiv:2504.04056* .

Börjesson, Maria, Jonas Eliasson, Muriel B. Hugosson, and Karin Brundell-Freij. 2012. "The Stockholm Congestion Charges—5 Years on. Effects, Acceptability and Lessons Learnt." *Transport Policy* 20:1–12.

Coase, Ronald H. 1972. "Durability and Monopoly." *The Journal of Law & Economics* 15 (1):143–149.

Conlon, Christopher and Jeff Gortmaker. 2020. "Best Practices for Differentiated Products Demand Estimation with PyBLP." *The RAND Journal of Economics* 51 (4):1108–1161.

———. 2025. "Incorporating Micro Data into Differentiated Products Demand Estimation with PyBLP." *Journal of Econometrics* 242:105926.

Cook, Cody, Aboudy Kreidieh, Shoshana Vasserman, Hunt Allcott, Neha Arora, Andrew Tomkins, Eray Turkel, and Freek van Sambeek. 2025. "The Network-Wide Effects of Congestion Pricing: Evidence from New York City." NBER Working Paper 33584, National Bureau of Economic Research.

Cook, Cody and Pearl Z. Li. 2025. "Value Pricing or Lexus Lanes? The Distributional Effects of Dynamic Tolling." Working Paper, Yale University.

Davis, Lucas W. 2008. "The Effect of Driving Restrictions on Air Quality in Mexico City." *Journal of Political Economy* 116 (1):38–81.

Delgado-Lindeman, Maira, Ruben Cordera, José Luis Moura, and Andrés Rodriguez. 2025. "Characteristics and Effects of Low Emission Zones in Europe: A Systematic Literature Review." *European Transport Research Review* 17:54.

D'Haultfœuille, Xavier, Isis Durrmeyer, and Philippe Février. 2019. "Automobile Prices in Market Equilibrium with Unobserved Price Discrimination." *Review of Economic Studies* 86 (5):2053–2096.

Durrmeyer, Isis and Nicolás Martínez. 2026. "The Welfare Consequences of Urban Traffic Regulations." Working Paper, Toulouse School of Economics.

Ellerman, A. Denny, Paul L. Joskow, Richard Schmalensee, Juan-Pablo Montero, and Elizabeth M. Bailey. 2000. *Markets for Clean Air: The U.S. Acid Rain Program.* Cambridge University Press.

European Commission, Directorate-General for Mobility and Transport, CE Delft, H. v. Essen, L. v. Wijngaarden, D. Sutter, C. Bieler, S. Maffii, D. Fiorello, F. Fermi, R. Parolin, A. Schroten, M. Brambilla, and K. El Beyrouty. 2020. *Handbook on the external costs of transport – Version 2019 – 1.1.* Publications Office.

Galdon-Sanchez, Jose Enrique, Ricard Gil, Felix Holub, and Guillermo 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* 21:1227–1267.

Gallego, Francisco, Juan-Pablo Montero, and Christian Salas. 2013. "The Effect of Transport Policies on Car Use: Evidence from Latin American Cities." *Journal of Public Economics* 107:47–62.

Gandhi, Amit and Jean-François Houde. 2019. "Measuring Substitution Patterns in Differentiated-Products Industries." NBER Working Paper 26375, National Bureau of Economic Research.

Garson, Jonathan, Clément Malgouyres, Maxime Tô, and Corentin Trevien. 2026. "The Impact of Low Emission Zones on Vehicles Fleet: Evidence from France." Documents de travail 2026-12, INSEE.

Gavazza, Alessandro, Alessandro Lizzeri, and Nikita Roketskiy. 2014. "A Quantitative Analysis of the Used-Car Market." *American Economic Review* 104 (11):3668–3700.

Gehrsitz, Markus. 2017. "The Effect of Low Emission Zones on Air Pollution and Infant Health." *Journal of Environmental Economics and Management* 83:121–144.

Gillingham, Kenneth, Fedor Iskhakov, Anders Munk-Nielsen, John Rust, and Bertel Schjerning. 2022. "Equilibrium Trade in Automobiles." *Journal of Political Economy* 130 (10):2534–2593.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik Instruments: What, When, Why, and How." *American Economic Review* 110 (8):2586–2624.

Green NCAP. 2022. "LCA: How Sustainable is Your Car?" Life-cycle assessment of passenger cars; assumes a vehicle lifetime of 16 years and 240,000 km.

Grieco, Paul L. E., Charles Murry, and Ali Yurukoglu. 2024. "The Evolution of Market Power in the U.S. Automobile Industry." *Quarterly Journal of Economics* 139 (2):1201–1253.

Helm, Dieter, Cameron Hepburn, and Richard Mash. 2003. "Credible Carbon Policy." *Oxford Review of Economic Policy* 19 (3):438–450.

Holland, Stephen P., Erin T. Mansur, and Andrew J. Yates. 2021. "The Electric Vehicle Transition and the Economics of Banning Gasoline Vehicles." *American Economic Journal: Economic Policy* 13 (3):316–44.

Isaksen, Elisabeth T. and Bjørn G. Johansen. 2021. "Congestion Pricing, Air Pollution, and Individual-Level Behavioral Responses." Memorandum 1/2021, Oslo University, Department of Economics.

Jann, Ben. 2020. "KMATCH: Stata module for multivariate-distance and propensity-score matching, including entropy balancing, inverse probability weighting, (coarsened) exact matching, and regression adjustment." Statistical Software Components S458346, Boston College Department of Economics.

Kahn, Charles. 1986. "The Durable Goods Monopolist and Consistency with Increasing Costs." *Econometrica* 54 (2):275–294.

Klauber, Hannah, Felix Holub, Nicolas Koch, Nico Pestel, Nolan Ritter, and Alexander Rohlf. 2024. "Killing Prescriptions Softly: Low Emission Zones and Child Health from Birth to School." *American Economic Journal: Economic Policy* 16 (2):220–48.

Laffont, Jean-Jacques and Jean Tirole. 1996. "Pollution permits and environmental innovation." *Journal of Public Economics* 62 (1):127–140. Proceedings of the Trans-Atlantic Public Economic Seminar on Market Failures and Public Policy.

Li, Shanjun, Junji Xiao, and Yimin Liu. 2015. "The Price Evolution in China's Automobile Market." *Journal of Economics & Management Strategy* 24 (4):786–810.

Margaryan, Shushanik. 2021. "Low Emission Zones and Population Health." *Journal of Health Economics* 76:102402.

Miller, Nathan H. and Matthew C. Weinberg. 2017. "Understanding the Price Effects of the MillerCoors Joint Venture." *Econometrica* 85 (6):1763–1791.

Miravete, Eugenio and Maria Moral. 2025. "The Long-Lasting Protection of a Defunct Industrial Policy." CEPR Discussion Papers 20525, Centre for Economic Policy Research.

Miravete, Eugenio J., Maria J. Moral, and Jeff Thurk. 2018. "Fuel taxation, emissions policy, and competitive advantage in the diffusion of European diesel automobiles." *RAND Journal of Economics* 49 (3):504–540.

OECD. 2014. *The Cost of Air Pollution: Health Impacts of Road Transport*. Paris: OECD Publishing.

Petrin, Amil. 2002. "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy* 110 (4):705–729.

Plötz, Patrick, Cornelius Moll, Georg Bieker, and Peter Mock. 2021. "From lab-to-road: real-world fuel consumption and CO_2 emissions of plug-in hybrid electric vehicles." *Environmental Research Letters* 16 (5):054078.

Red Eléctrica de España. 2026. "ESIOS: Sistema de Información del Operador del Sistema." URL https://www.esios.ree.es/. Electricity system data, accessed 20 June 2026.

Rosales-Tristancho, Abel, Raúl Brey, Ana F. Carazo, and J. Javier Brey. 2022. "Analysis of the barriers to the adoption of zero-emission vehicles in Spain." *Transportation Research Part A: Policy and Practice* 158:19–43.

Rubin, Donald B. 1980. "Bias Reduction Using Mahalanobis-Metric Matching." *Biometrics* 36 (2):293–298.

Sarmiento, Luis, Nicole Wägner, and Aleksandar Zaklan. 2023. "The Air Quality and Well-Being Effects of Low Emission Zones." *Journal of Public Economics* 227:105014.

Tassinari, Filippo. 2024. "Low Emission Zones and Traffic Congestion: Evidence from Madrid Central." *Transportation Research Part A: Policy and Practice* 185 (C).

Tirole, Jean. 1988. *The Theory of Industrial Organization*. The MIT Press.

U.S. Environmental Protection Agency. 2023. "Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances." Tech. rep., U.S. Environmental Protection Agency, Washington, DC.

Weitzman, Martin L. 1974. "Prices vs. Quantities." *The Review of Economic Studies* 41 (4):477–491.

Wolff, Hendrik. 2014. "Keep Your Clunker in the Suburb: Low-Emission Zones and Adoption of Green Vehicles." *The Economic Journal* 124 (578):F481–F512.

World Health Organization. 2025. "Transport: Sectoral Solutions for Air Pollution and Health." Tech. rep., World Health Organization.

Online Appendix for:

Designing Low Emission Zones: Theory and Evidence

(Not for publication)

Nano Barahona
UC Berkeley & NBER

Natalia Fabra
CEMFI & CEPR

Imelda
Geneva Grad. Institute & CEPR

Juan-Pablo Montero
PUC-Chile & Aalto University

November 7, 2026

A Additional figures and tables **1**

B When the LEZ affects driving **5**

C Proofs **7**
C.1 Proof of Lemma 1 . . . 7
C.2 Proof of Proposition 1 . . . 7
C.3 Proof of Proposition 2 . . . 8
C.4 Proof of Corollary 1 . . . 9
C.5 Proof of Proposition 3 . . . 10
C.6 Proof of Proposition 4 . . . 10
C.7 Proof of Proposition 5 . . . 13

D Additional results for the car-choice model **15**
D.1 Preliminaries . . . 15
D.2 No intervention . . . 16
D.3 Time-invariant LEZ with trade frictions . . . 17
D.4 Time-invariant LEZ with no frictions . . . 20
D.5 Gasoline cars as a transition technology . . . 21

E Quasi-convexity of the welfare function **23**
E.1 Quasi-convexity conditions 23
E.2 Illustrations 26
E.3 Reverse welfare ordering 28

F Marginal welfare effect decomposition **29**
F.1 The benefit of a util at every intensity 29
F.2 Evaluation at our estimates 31

Appendix A: Additional figures and tables



Figure A.1: Madrid Central map

Table A.1: SDID results: Impacts on new, second-hand, and scrappage registrations

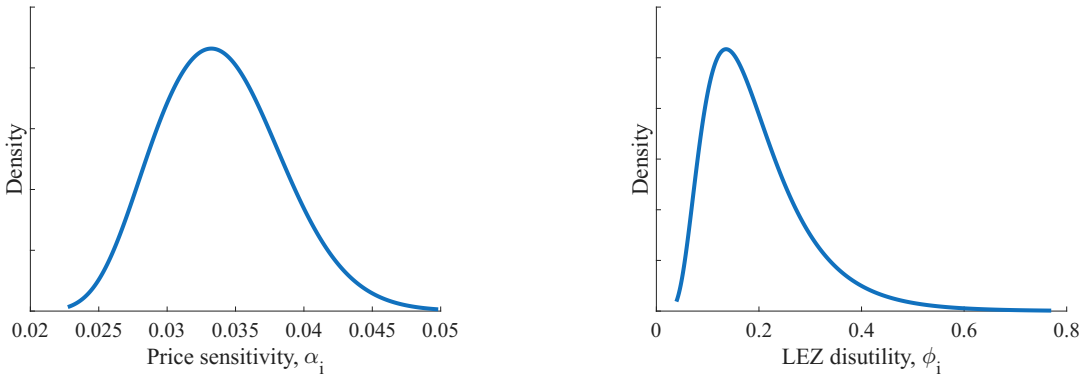
Specification	(1) 100-nn	(2) 25-nn	(3) 10-nn	(4) 5-nn	(5) 5-nn + cov.
Panel A: Gasoline/Diesel					
Madrid × Post	-4.59***	-4.34***	-4.36***	-3.94***	-3.94***
	(0.83)	(0.79)	(0.83)	(0.78)	(0.78)
Panel B: LPG					
Madrid × Post	1.61***	1.50***	1.50***	1.49***	1.49***
	(0.04)	(0.05)	(0.05)	(0.05)	(0.05)
Panel C: Hybrids					
Madrid × Post	2.09***	1.85***	1.69***	1.59***	1.59***
	(0.13)	(0.14)	(0.14)	(0.12)	(0.12)
Panel D: EVs					
Madrid × Post	0.18***	0.17***	0.15***	0.14***	0.14***
	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)
Panel E: Second-hand transfers					
Madrid × Post	-1.67	-1.16	-0.45	-0.58	-0.57
	(1.52)	(1.19)	(1.09)	(1.35)	(1.35)
Panel F: Scrappage					
Madrid × Post	-0.77	-0.18	-0.55	-0.59	-0.59
	(1.17)	(0.85)	(0.90)	(0.80)	(0.80)
Observations	115,700	62,000	37,160	24,680	24,680
Control zip codes	5,496	2,811	1,571	948	948

Notes: The table reports SDID estimates per 10,000 residents for quarterly new registrations (Panels A–D), second-hand transfers (Panel E), and scrappage (Panel F); standard errors are computed via 100 clustered placebo replications following Algorithm 4 in Arkhangelsky et al. (2021). Columns (1)–(4) use, respectively, 100, 25, 10, and 5 nearest neighbors without covariates; column (5) adds zip-code covariates (log households, log purchasing power per capita, average household size, and household income-quintile shares) to the 5-nearest-neighbor specification. All specifications include unit and year-quarter fixed effects and SDID unit and time weights. Sample counts apply to all panels; control zip codes include the full matched pool. LPG denotes LPG/CNG. Stars denote significance at the 10 (*), 5 (**), and 1 (***) percent levels.

Table A.2: Micro-moments in the data and in the estimated model

	Data		Model	
	Low income	High income	Low income	High income
Panel A: Purchases				
Purchase probability (%)	85.08	87.87	85.23	87.67
Average price paid (thousand €)	24.68	28.11	24.76	27.67
Panel B: Car types among buyers (%)				
Diesel	43.58	36.69	43.45	37.23
LPG/CNG	0.94	0.93	0.92	1.09
Hybrid	4.97	9.19	4.97	10.06
EV	0.15	0.44	0.15	0.57
Panel C: Restricted and unrestricted cars among buyers (%)				
Restricted, Madrid, pre-policy	93.74	89.93	93.68	89.56
Restricted, Madrid, post-policy	76.60	69.87	75.05	72.77
Restricted, other cities, pre-policy	96.35	93.06	96.41	92.14
Restricted, other cities, post-policy	90.28	85.95	90.55	82.05
Unrestricted, Madrid, both periods	10.80	15.27	11.19	14.92
Panel D: Change in the log odds of a restricted purchase, Madrid				
0–3 km	-1.20		-1.17	
3–10 km	-1.49		-1.39	
10–20 km	-1.47		-1.57	
20 km or more	-1.48		-1.54	

Notes: The table reports the 26 micro-moments used in estimation, in the data and at the estimated parameters. Low- and high-income consumers live in zip codes with average income up to and above €20,000, respectively. Panels A–C report moments for each income group: the purchase probability and the average price paid by buyers, the shares of car types among buyers, and the shares of restricted cars among buyers in Madrid and in the other cities, pre- and post-policy, together with the share of unrestricted cars among Madrid buyers over both periods. Panel D reports, for Madrid buyers in both income groups, the change in the log odds of a restricted purchase from the pre- to the post-policy period, by distance from the zip code to Madrid Central. The post-policy period starts in the fourth quarter of 2018, when Madrid Central was implemented.



(a) Price sensitivity, α_i (b) LEZ disutility, ϕ_i

Figure A.2: Distributions of price sensitivity and LEZ disutility

Notes: Distributions of α_i (Panel (a)) and ϕ_i (Panel (b)) in equation (16) across Madrid zip codes and taste draws, weighted by households, as in Figure 6, which shows their ratio. Both are in utility units: α_i per thousand euros of price, ϕ_i for a car subject to the LEZ restrictions.

Appendix B: When the LEZ affects driving

Consider *short-lived* vehicles where the LEZ not only limits access but also changes (reduces or increases) driving intensity. We might explain the reduction by greater use of public transit, while the increase might be due to more rerouting around the LEZ.

Specifically, for a LEZ of size s^j, emissions from a type-j vehicle are given by $(1-s^j)e^j$ (or $e^j/(1-s^j)$); that is, larger LEZs lead to less (more) driving and, therefore, lower (higher) emissions. Incorporating this expression into the welfare function and comparing the partial derivatives of welfare with respect to the LEZ sizes for gasoline and hybrid vehicles shows that it remains optimal to restrict gasoline cars first, as in the baseline model.

Here, we consider the case where driving is reduced. First, consider the case in which hybrids are unrestricted, $s^H = 0$, and assess the second-order derivative with respect to the LEZ size affecting gasoline cars, $s^G = s$,

$$\frac{\partial^2 W}{\partial s^2} = v_1 \frac{v_1 - 2e}{\gamma}.$$

And for $s^G = s^H = s$,

$$\frac{\partial^2 W}{\partial s^2} = v_1 \frac{v_1 - e\,(1+\alpha)}{\gamma}.$$

Since v_1 has to be large enough for cars to be welfare-improving and for drivers to be willing to buy them (full-coverage assumption), these two derivatives are positive, showing that, as in the baseline model, the welfare function is convex, leading to bang-bang solutions.

To assess whether to restrict gasoline cars or not, we compare welfare with and without restricting them. We use $W\left(s^H, s^G\right)$ to denote welfare as a function of the size of the LEZ for hybrids and gasoline. Equating $W\,(0,0) = W\,(0,s)$ (no LEZ against gasoline only) and solving for the emission threshold,

$$\widetilde{e}_G^{LEZ} = v_1 \frac{2\gamma - 3v_1 s}{2\gamma - 3v_1\alpha + 6v_1(1-s)}.$$

Compare this with our baseline threshold

$$\widetilde{e}_G^{LEZ} - e_G^{LEZ} = -\frac{2}{3}\frac{(2\gamma - 3v_1 s)(\gamma - 3v_1 s)}{(2-\alpha)\,(2\gamma - 3v_1\alpha + 6v_1(1-s))} < 0,$$

which is negative for $s < \frac{\gamma}{3v_1}$.

Note that if hybrids are not restricted and gasoline cars are, gasoline cars are not that harmful relative to the baseline as they are driven less. Hence, the emission threshold shifts in.

Likewise, equating $W(0,s) = W(s,s)$ (gasoline-only LEZ against both restricted) and solving for the emission threshold,

$$\widetilde{e}_H^{LEZ} = 2\gamma \frac{v_1}{2\gamma\alpha + 3v_1\alpha - 3v_1(1-s) + 3v_1\alpha(1-s)}.$$

Comparing it with the baseline threshold,

$$\widetilde{e}_H^{LEZ} - e_H^{LEZ} = -\frac{2}{3}\frac{\gamma}{2\alpha - 1}\frac{2\gamma\alpha + 3sv_1(1-\alpha)}{2\gamma\alpha + 3sv_1(1-\alpha) + 3v_1(2\alpha-1)},$$

which is negative in the relevant case where $\alpha > 1/2$. Again, the emission thresholds shift relative to the baseline, as hybrids become less harmful when their usage is restricted.

In sum, the baseline model's bang-bang structure remains intact, even when LEZ restrictions affect not only car values (extensive margin) but also driving intensity (intensive margin). This extension simply shifts the emission thresholds, making the LEZ policy more welfare-enhancing than in the baseline case.

Appendix C: Proofs

C.1 Proof of Lemma 1

Iterating the Bellman equations (4) from age $a = 1$ to age T_j and using the condition that, at the end of period T_j, the driver scraps the car and starts a new cycle, i.e., $U^j_{T_j}(x) = u^j_{T_j}(x) + \delta U^j_1(x) + \delta z^j$, one obtains

$$U^j_1(x) = \sum\nolimits_{a=1}^{T_j} \delta^{a-1} u^j_a(x) - p^j_1 + \delta^{T_j} z^j + \delta^{T_j} U^j_1(x).$$

Solving for $U^j_1(x)$ yields the expression in the lemma.

C.2 Proof of Proposition 1

Part (i). Let $v^G = (1 - s^G)v_1$ and $v^H = (1 - s^H)v_1$. Assuming for now that v^G and v^H are such that all models are sold in equilibrium, take the derivative of the welfare expression (8) with respect to v^G and v^H to obtain (recall that $v^E = v_1$)

$$\frac{\partial W}{\partial v^G} = \frac{1}{6\gamma}(6v^G - 3v^H - 3v_1 + 2\gamma - 3(2-\alpha)e) \tag{C.1}$$

$$\frac{\partial W}{\partial v^H} = \frac{1}{6\gamma}\left(6v^H - 3v^G - 3v_1 + 2\gamma - 3(2\alpha - 1)e\right) \tag{C.2}$$

respectively. These expressions show that welfare is convex in v^G and v^H, and that their difference

$$\frac{\partial W}{\partial v^G} - \frac{\partial W}{\partial v^H} = -\frac{3}{2\gamma}\left((1-\alpha)\,e + \left(v^H - v^G\right)\right) \tag{C.3}$$

is negative when evaluated at $v^H = v^G = v_1$. This indicates that any LEZ policy must start by restricting gasoline vehicles first.

To assess when this is the case, we compare welfare levels with and without restricting gasoline vehicles only by a LEZ of size s. Writing the welfare function as $W\left(s^H, s^G\right)$ for convenience, with $s^H = 0$ and $s^G = s$, we obtain

$$W(0,0) - W(0,s) = \frac{v_1 s}{6\gamma}\left(2\gamma - 3v_1 s - 3\left(2-\alpha\right)e\right)$$

which indicates that having no LEZ is optimal, as opposed to a LEZ of size s, when

$$e \leq e^{LEZ}_G(s) \equiv \frac{2\gamma}{3(2-\alpha)} - \frac{v_1 s}{2-\alpha}.$$

Part (ii). Suppose that $e > e^{LEZ}_G(s)$. To assess the benefit of adding hybrids to the LEZ restriction, we compare welfare levels with and without restricting them, conditional on

gasoline cars already being restricted. Simple algebra yields

$$W(0,s) - W(s,s) = \frac{v_1 s}{6\gamma}\left(2\gamma - 3(2\alpha - 1)e\right)$$

which indicates that adding hybrids to the LEZ is optimal when

$$e \geq e_H^{LEZ} \equiv \frac{2\gamma}{3(2\alpha - 1)}.$$

Note that both thresholds are continuous in α, and that at $\alpha = 1$ they differ by $e_H^{LEZ} - e_G^{LEZ}(s) = v_1 s > 0$: even when hybrids and gasoline cars pollute the same, there is a range of emission levels in which only gasoline cars should be restricted. To see why, consider $\alpha = 1$. Relative to no LEZ, restricting gasoline cars reduces emissions by $e\, s v_1/2\gamma$ (the mass of drivers who switch to EVs), imposes a usage loss $s v_1$ on the $1/3 - s v_1/\gamma$ remaining gasoline drivers, and raises transportation costs by $s^2 v_1^2/2\gamma$, for a net cost of $s v_1/3 - s^2 v_1^2/2\gamma$. Adding hybrids achieves the same emission reduction, $e\, s v_1/2\gamma$, and leaves total transportation costs unchanged (in both cases, two of the three indifference points shift by $s v_1/2\gamma$), but raises the mass of restricted drivers by exactly $1/3$, from $1/3 - s v_1/\gamma$ to $2/3 - s v_1/\gamma$, for a cost of $s v_1/3$. The difference in costs, $s^2 v_1^2/2\gamma$, divided by the emission gain per unit of e, $s v_1/2\gamma$, gives the gap $v_1 s$ between the two thresholds. For $\alpha < 1$, the emission gain from adding hybrids is smaller, as the $s v_1/2\gamma$ drivers who switch from hybrids back to gasoline cars raise emissions by $(1-\alpha)e$ each, while those who switch to EVs reduce them by αe; hence the term $(2\alpha - 1)e\, s v_1/2\gamma$ in $W(0,s) - W(s,s)$.

Also note that, for $\alpha > 1/2$, $e_H^{LEZ} > e_G^{LEZ}(s)$ for all s, as expected. Restricting hybrids pays only if $3(2\alpha - 1)e > 2\gamma$, which fails for every e when $\alpha \leq 1/2$, so the threshold e_H^{LEZ} tends to infinity.

Part (iii). It follows from parts (i) and (ii).

C.3 Proof of Proposition 2

Consider the threshold $\tilde{s}(e)$ given by $e_G^{LEZ}(\tilde{s}) = e$. We evaluate welfare $W(s^H, s^G)$ at various solutions for given s^H and s^G.

Part (i). Suppose $e < e_H^{LEZ}$. We know from Proposition 1 that it is optimal to set $s^H = 0$. Let $s^G = s$ and consider two cases. First, consider $e < e_G^{LEZ}(s^{\max}) = \frac{\gamma}{3(2-\alpha)} < e_G^{LEZ}(s)$. In this case, $W(0,0) > W(0,s)$ for any $s \leq s^{\max}$, so it is optimal not to create a LEZ, i.e., set $s = 0$. Second, consider $e_G^{LEZ}(s^{\max}) < e < e_H^{LEZ}$. Note that $\partial^2 W(0,s)/\partial s^2 = v_1^2/\gamma > 0$, so $W(0,s)$ is convex in s (see also the proof of Proposition 1). This implies that we need to compare welfare evaluated at both extremes: $W(0,0) - W(0,\bar{s}) > 0$ if

$$e < e_G^{LEZ}(\bar{s}) \equiv \frac{2\gamma}{3(2-\alpha)} - \frac{v_1}{2-\alpha}\bar{s},$$

which requires

$$\bar{s} < \frac{2\gamma - 3\,(2-\alpha)\,e}{3v_1} \equiv \tilde{s}.$$

Part (ii). By the same logic as above, if $e_G^{LEZ}(s^{\max}) < e < e_H^{LEZ}$ and $\bar{s} > \tilde{s}$ then $W\,(0,0) - W\,(0,\bar{s}) < 0$ so the optimal size of the LEZ is $\bar{s}$.

Part (iii). Last, if $e > e_H^{LEZ}$, conditionally on having a LEZ of size s, the optimal policy is to restrict both hybrids and gasoline vehicles (by Proposition 1). The welfare function $W\,(s,s)$ is convex in s, so the optimum is attained at a corner. From the proof of Proposition 1 above, we know that $W\,(0,\bar{s}) < W\,(\bar{s},\bar{s})$ whenever $e > e_H^{LEZ}$. Moreover, part (ii) with $e > e_H^{LEZ} > e_G^{LEZ}(0)$ implies that $W\,(0,0) < W\,(0,\bar{s})$. Combining both inequalities yields $W\,(\bar{s},\bar{s}) > W\,(0,0)$, so the largest feasible LEZ, $\bar{s}$, restricting both gasoline and hybrid vehicles is optimal.

C.4 Proof of Corollary 1

By contradiction. Suppose there exists some $\hat{s}^H \in (0, s^G)$ such that $W(\hat{s}^H, s^G) > \max\{W(0, s^G), W(s^G, s^G)\}$. Compute the derivative

$$\frac{\partial W(s^H, s^G)}{\partial s^H} = \frac{v_1(-2\gamma - 3e + 6\alpha e + 6v_1 s^H - 3v_1 s^G)}{6\gamma},$$

which is linear and increasing in s^H (with slope $v_1^2/\gamma > 0$). Hence $W(s^H, s^G)$ is convex in s^H and its maximum over $[0, s^G]$ is attained at one of the two endpoints, $s^H = 0$ or $s^H = s^G$, contradicting the supposition that an interior $\hat{s}^H \in (0, s^G)$ yields strictly higher welfare.

Equivalently, consider the two possible orderings of the endpoint values:

Case 1: $W(0, s^G) > W(s^G, s^G)$**.** Since W is convex in s^H, the derivative $\partial W/\partial s^H$ evaluated at $s^H = \hat{s}^H$ cannot be positive: otherwise, convexity would imply $W(s^G, s^G) > W(\hat{s}^H, s^G)$, contradicting $W(\hat{s}^H, s^G) > W(s^G, s^G)$. But then the derivative is nonpositive at $\hat{s}^H$ and, by convexity, also along the entire interval $[0, \hat{s}^H]$. Hence $W(0, s^G) \geq W(\hat{s}^H, s^G)$, contradicting $W(\hat{s}^H, s^G) > W(0, s^G)$.

Case 2: $W(s^G, s^G) > W(0, s^G)$**.** The argument is symmetric: convexity of W in s^H implies $W(\hat{s}^H, s^G) \leq \max\{W(0, s^G), W(s^G, s^G)\}$, again a contradiction.

Moreover, the negative cross-derivative

$$\frac{\partial^2 W(s^H, s^G)}{\partial s^H \partial s^G} = -\frac{v_1^2}{2\gamma} < 0$$

reinforces the bang-bang logic: a larger s^G reduces the marginal benefit of restricting hybrids, so the optimal s^H remains at a corner.

C.5 Proof of Proposition 3

It follows the proof of Proposition 4 while imposing $s_1^H = s_2^H = s$, leading to thresholds

$$e_G^{LEZ} \equiv \frac{4\gamma(v_1+v_2)-(2v_1+v_2)^2 s}{3(2v_1+v_2)(2-\alpha)} < \frac{4\gamma(v_1+v_2)}{3(2v_1+v_2)\,(2\alpha-1)} \equiv e_H^{LEZ}.$$

C.6 Proof of Proposition 4

The location of the indifferent consumer between the different types of cars will vary over time and with whether cars are new or second-hand. Denote by $\tilde{x}_{a,t}^{jk}$ the location of the consumer that at time $t \geq 1$ is indifferent between a car type $j = E, H, G$ of age $a = 1, 2$ and a car type $k \neq j$ of the same age a. Note that during the transition (i.e., $t < 3$), $\tilde{x}_{1,t}^{jk}$ is likely to differ from $\tilde{x}_{2,t}^{jk}$; what makes an individual indifferent between a new and a second-hand car of the same type is that the price of the second-hand unit adjusts accordingly, that is, according to (5).

Because cars survive until they are scrapped, we have that $\tilde{x}_{2,t+1}^{jk} = \tilde{x}_{1,t}^{jk}$ for all $t \geq 1$. We also know that at $t = 3$ the fleet reaches its steady state, so $\tilde{x}_{1,t-1}^{jk} = \tilde{x}_{a,t}^{jk} = \tilde{x}_{a,t+1}^{jk}$ for all a and $t \geq 3$. To find $\tilde{x}_{a,t}^{jk}$, we use the set of Bellman equations adjusted to account for a LEZ policy affecting gasoline cars and possibly hybrids.

A driver who buys a new car at $t = 1$ drives it during $t = 1$ and $t = 2$, scraps it at the end of $t = 2$, and at $t = 3$ chooses again among all car types, facing the steady-state restrictions. Let $V_3(x) \equiv \max_k U_1^k(x)$ denote the value of that choice, evaluated at the restrictions in place from $t = 3$ onward. Since $V_3(x)$ does not depend on the car bought at $t = 1$, every $t = 1$ choice has the same continuation $\delta^2 V_3(x)$.

For example, to find $\tilde{x}_{1,1}^{EH}$, consider first drivers of electric vehicles (recall that $x^E = 0$),

$$U_1^E(x) = v_1 - \gamma|x - x^E| - c + \delta U_2^E(x)$$

and

$$U_2^E(x) = v_2 - \gamma|x - x^E| + \delta z + \delta V_3(x).$$

Allowing hybrids to face different restrictions today ($t = 1$) and tomorrow ($t \geq 2$), we similarly have (recall that $x^H = 1/3$)

$$U_1^H(x) = (1 - s_1^H)v_1 - \gamma|x - x^H| - c + \delta U_2^H(x)$$

and

$$U_2^H(x) = (1 - s_2^H)v_2 - \gamma|x - x^H| + \delta z + \delta V_3(x),$$

where s_t^H is the size of the LEZ affecting hybrids in period t, with $s_2^H \geq s_1^H$ (recall that $s_t^H = s_2^H$ for all $t \geq 3$). Since the continuation $\delta^2 V_3(x)$ is common to both choices, $\tilde{x}_{1,1}^{EH}$

solves $U_1^E(x) = U_1^H(x)$, or

$$v_1 + \delta v_2 - (1+\delta)\gamma|x - x^E| = (1 - s_1^H)v_1 + \delta(1 - s_2^H)v_2 - (1+\delta)\gamma|x - x^H|,$$

so only the first ownership cycle matters for the choices made at $t = 1$.

To obtain $\tilde{x}_{1,2}^{EH}$, the location of the indifferent consumer between hybrids and EVs in period $t = 2$, we proceed in the same way. A driver who buys at $t = 2$ faces the restriction s_2^H during both periods of the car's life, and again faces a common continuation from $t = 4$ onward, so $\tilde{x}_{1,2}^{EH}$ solves the analogous condition with s_1^H replaced by s_2^H.

We skip the derivation of the remaining indifference locations, namely, $\tilde{x}_{1,1}^{EG}$, $\tilde{x}_{1,2}^{EG}$, $\tilde{x}_{1,1}^{HG}$ and $\tilde{x}_{1,2}^{HG}$, which follow the same steps. These indifference locations will be a function of the primitives and the LEZ design, which is given by the triplet (s^G, s_1^H, s_2^H). Note that since gasoline cars are restricted from $t = 1$ onward, $\tilde{x}_{1,1}^{EG} = \tilde{x}_{1,2}^{EG}$, reducing the number of indifference locations to the following five (after setting $\delta = 1/2$)

$$\begin{aligned}
\tilde{x}_{1,1}^{EH} &= \frac{1}{6\gamma}\left(\gamma + 2v_1 s_1^H + v_2 s_2^H\right), \\
\tilde{x}_{1,2}^{EH} &= \frac{1}{6\gamma}\left(\gamma + (2v_1 + v_2) s_2^H\right), \\
\tilde{x}_{1,1}^{HG} &= \frac{1}{6\gamma}\left(3\gamma + (2v_1 + v_2) s^G - 2v_1 s_1^H - v_2 s_2^H\right), \\
\tilde{x}_{1,2}^{HG} &= \frac{1}{6\gamma}\left(3\gamma + (2v_1 + v_2) s^G - (2v_1 + v_2) s_2^H\right), \text{ and} \\
\tilde{x}_{1,t}^{EG} &= \frac{1}{6\gamma}\left(5\gamma - (2v_1 + v_2) s^G\right).
\end{aligned}$$

Note that $\tilde{x}_{1,2}^{EH} - \tilde{x}_{1,1}^{EH} = v_1(s_2^H - s_1^H)/3\gamma \geq 0$: under a transition-technology (TT) design, some drivers who buy a hybrid at $t = 1$ switch to an EV when they replace it at $t = 3$; others switch to a gasoline car when they replace it at $t = 3$.

Plugging these indifference locations into the welfare function (9), we obtain an expression that can be conveniently written as a function of the LEZ design as follows,

$$W(s^G, s_1^H, s_2^H) = W_1(s^G, s_1^H, s_2^H) + \delta W_2(s^G, s_1^H, s_2^H) + \frac{\delta^2}{1-\delta} W_3(s^G, s_2^H).$$

Note that in writing this expression we are taking advantage of the fact that the indifference locations for new cars completely determine car ownership for second-hand units. Individuals who are found to be switching cars in response to the LEZ intervention, from a restricted to an unrestricted one, will only switch when their existing cars are ready to be scrapped, at the end of their second period of life, not before. It can be shown that these individuals are strictly better off driving their cars until their scrappage age rather than selling them earlier and switching to their new best option.

The rest of the proof follows that of Proposition 1. We compare welfare across

the four relevant designs: no LEZ, $(0,0,0)$; restricting only gasoline cars, $(s,0,0)$; the transition design, $(s,0,s)$; and restricting both gasoline cars and hybrids from $t=1$ onward, (s,s,s). All pairwise welfare differences are linear in e.

First, as in Proposition 3, restricting only gasoline cars yields higher welfare than no LEZ if and only if $e > e_G^{LEZ}$, with

$$W(s,0,0) - W(0,0,0) = \frac{(2-\alpha)(2v_1+v_2)s}{2\gamma}\left(e - e_G^{LEZ}\right).$$

Second, the transition design yields higher welfare than restricting only gasoline cars if and only if $e > \hat{e}_{TTL}^{LEZ}$, with

$$W(s,0,s) - W(s,0,0) = \frac{(2\alpha-1)(v_1+v_2)s}{2\gamma}\left(e - \hat{e}_{TTL}^{LEZ}\right)$$

where

$$\hat{e}_{TTL}^{LEZ} \equiv \frac{2\gamma(v_1+v_2) + v_1 v_2 s}{3(2\alpha-1)(v_1+v_2)}. \tag{C.4}$$

Third, summing the two expressions above, the transition design yields higher welfare than no LEZ if and only if $e > e_{TT0}^{LEZ}$, where

$$\begin{aligned} e_{TT0}^{LEZ} &\equiv \frac{6\gamma(v_1+v_2) - \left(4v_1^2 + 3v_1v_2 + v_2^2\right)s}{3\left(3v_1+v_2+\alpha v_2\right)} \\ &= \frac{(2-\alpha)(2v_1+v_2)\, e_G^{LEZ} + (2\alpha-1)(v_1+v_2)\,\hat{e}_{TTL}^{LEZ}}{(2-\alpha)(2v_1+v_2) + (2\alpha-1)(v_1+v_2)}. \end{aligned}$$

Since $\alpha \in (1/2, 1)$, both weights are positive, so e_{TT0}^{LEZ} lies between e_G^{LEZ} and $\hat{e}_{TTL}^{LEZ}$.

Two cases arise.

Case A: $\hat{e}_{TTL}^{LEZ} \geq e_G^{LEZ}$. Then $e_G^{LEZ} \leq e_{TT0}^{LEZ} \leq \hat{e}_{TTL}^{LEZ}$. For $e \leq e_G^{LEZ}$, no LEZ yields higher welfare than restricting only gasoline cars and than the TT design (as $e \leq e_{TT0}^{LEZ}$). For $e_G^{LEZ} < e \leq \hat{e}_{TTL}^{LEZ}$, restricting only gasoline cars yields higher welfare than both no LEZ and the TT design. For $e > \hat{e}_{TTL}^{LEZ}$, the TT design yields higher welfare than restricting only gasoline cars, which in turn yields higher welfare than no LEZ.

Case B: $\hat{e}_{TTL}^{LEZ} < e_G^{LEZ}$. Then $\hat{e}_{TTL}^{LEZ} < e_{TT0}^{LEZ} < e_G^{LEZ}$. For $e \leq e_{TT0}^{LEZ}$, no LEZ yields higher welfare than both the TT design and restricting only gasoline cars. For $e_{TT0}^{LEZ} < e \leq e_G^{LEZ}$, the TT design yields higher welfare than no LEZ, which in turn yields higher welfare than restricting only gasoline cars. For $e > e_G^{LEZ}$, the TT design yields higher welfare than restricting only gasoline cars, which in turn yields higher welfare than no LEZ. Hence, restricting only gasoline cars is never optimal.

Combining both cases, the lower threshold is

$$e_{TTL}^{LEZ} \equiv \begin{cases} \hat{e}_{TTL}^{LEZ}, & \text{if } \hat{e}_{TTL}^{LEZ} \geq e_{G}^{LEZ}, \\ e_{TT0}^{LEZ}, & \text{if } \hat{e}_{TTL}^{LEZ} < e_{G}^{LEZ}, \end{cases}$$

so that no vehicle is restricted if $e \leq \min\{e_G^{LEZ}, e_{TTL}^{LEZ}\}$, and only gasoline cars are restricted if $\min\{e_G^{LEZ}, e_{TTL}^{LEZ}\} < e \leq e_{TTL}^{LEZ}$, an interval that is empty in Case B. Since $\hat{e}_{TTL}^{LEZ}$ is increasing and e_G^{LEZ} is decreasing in s, Case B arises when s is sufficiently small, provided that $\hat{e}_{TTL}^{LEZ} < e_G^{LEZ}$ at $s = 0$, i.e., $\alpha > (6v_1 + 4v_2)/(6v_1 + 5v_2) > 10/11$. In both cases, $e_{TTL}^{LEZ} \leq \max\{e_G^{LEZ}, \hat{e}_{TTL}^{LEZ}\} < e_H^{LEZ}$.

Last, comparing the TT design with restricting both gasoline cars and hybrids from $t = 1$ onward, the latter yields higher welfare if and only if $e \geq e_{TTU}^{LEZ}$, with

$$e_{TTU}^{LEZ} \equiv \frac{2\gamma(v_1 + v_2) - v_1 v_2 s}{3v_1(2\alpha - 1)} > e_H^{LEZ},$$

which follows the same steps found in the proof of Proposition 1, so it is omitted. Since $e_{TTU}^{LEZ} > e_H^{LEZ} > e_{TTL}^{LEZ}$, the four cases in the proposition follow.

C.7 Proof of Proposition 5

At $t = 2$, taking as given the composition of the second-hand fleet inherited from $t = 1$, the planner re-optimizes over $s_2^H \in [0, s^G]$. By the same convexity argument used in the proof of Corollary 1, this optimization has a bang-bang solution: either $s_2^H = 0$ or $s_2^H = s^G$. We check each of the four cases in Proposition 4.

Case (i). In this case, no vehicle is restricted, so there is nothing to revise; the policy is trivially time consistent.

Case (ii). When only gasoline cars are restricted, the planner would revise the policy at $t = 2$ and extend the restriction to hybrids only if

$$e \geq \frac{4\gamma(v_1 + v_2) + (2v_1 + v_2)v_2 s}{3(2v_1 + v_2)(2\alpha - 1)} \equiv e_{(ii)}^{LEZ} > e_H^{LEZ}.$$

But this case applies only when $e \leq e_{TTL}^{LEZ} < e_H^{LEZ} < e_{(ii)}^{LEZ}$, so this revision never occurs.

Case (iii). In this case, hybrids are treated as a transition technology: unrestricted at $t = 1$ and restricted from $t = 2$ onward. The planner keeps her promise at $t = 2$ if and only if, given the (larger-than-steady-state) mass of second-hand hybrids, the marginal benefit of restricting hybrids exceeds the marginal cost. Solving this condition yields the threshold e_{TTS}^{LEZ} in (11). Thus, commitment is required whenever $e < e_{TTS}^{LEZ}$.

Case (iv). When both gasoline and hybrid cars are restricted from $t = 1$ onward, the

planner would revise the policy at $t = 2$ and lift the restriction on hybrids only if

$$e \leq \frac{4\gamma(v_1 + v_2) - (2v_1 + v_2)v_2 s}{3(2v_1 + v_2)(2\alpha - 1)} \equiv e^{LEZ}_{(iv)} < e^{LEZ}_{H}.$$

But this case applies only when $e \geq e^{LEZ}_{TTU} > e^{LEZ}_{H} > e^{LEZ}_{(iv)}$, so this revision never occurs either.

Hence, the only case in which the planner has an incentive to deviate from her $t = 1$ announcement is case (iii) of Proposition 4, for values of $e \in (e^{LEZ}_{TTL}, e^{LEZ}_{TTS})$. In the absence of commitment, the planner can partially restore credibility by imposing a smaller LEZ on hybrids at $t = 1$, $s^H_{a,1} < s^G$, which reduces the mass of second-hand hybrids on the road at $t = 2$ and thus the cost of restricting them. Given $s^H_{a,1}$, the planner keeps her promise at $t = 2$ if and only if

$$e \geq \frac{4\gamma(v_1 + v_2) + (2v_1 - v_2)v_2 s - 4v_1 v_2 s^H_{a,1}}{3(2v_1 + v_2)(2\alpha - 1)},$$

so the smallest credible restriction at $t = 1$ is

$$s^H_{a,1}(e) = \frac{3(2\alpha - 1)(2v_1 + v_2)}{4v_1 v_2} \left(e^{LEZ}_{TTS} - e\right),$$

which is decreasing in e and satisfies $s^H_{a,1}(e) \in (0, s^G)$ if and only if $e \in (e^{LEZ}_{(iv)}, e^{LEZ}_{TTS})$. At $e = e^{LEZ}_{(iv)} < e^{LEZ}_{H}$, this design coincides with restricting both types from $t = 1$ onward, which is dominated by restricting only gasoline cars; at $e = e^{LEZ}_{TTS} > \hat{e}^{LEZ}_{TTL}$, where the latter is given by (C.4), it coincides with the transition-technology design in Proposition 4, which dominates restricting only gasoline cars. The threshold e^{LEZ}_{TTW} is the value of e at which the weak transition design, with the smallest credible $s^H_{a,1}$, gives the same welfare as restricting gasoline only, i.e., at which $W\left(s, s^H_{a,1}(e), s\right) = W(s, 0, 0)$. The three cases in the proposition then follow by comparing e to the ordered thresholds $e^{LEZ}_{TTL} < e^{LEZ}_{TTW} < e^{LEZ}_{TTS}$.

Appendix D: Additional results for the car-choice model

D.1 Preliminaries

This appendix develops a simpler version of the car-choice model of Section 3, with two vehicle types: gasoline cars, G, and electric vehicles, E. The purpose is to show how prices and quantities adjust during the transition from the no-intervention equilibrium to the new post-intervention equilibrium, and to clarify the role of trade frictions, particularly in the emergence of transition technologies. For what we need to illustrate, nothing is lost by considering two types of cars rather than three.

There is a unit mass of consumers uniformly distributed on a Hotelling line $x \in [0, 1]$. Gasoline cars are located at $x = 0$ and electric vehicles at $x = 1$. All cars last for two periods and are then scrapped for value z. The flow value of a car of age $a = 1, 2, 3, \ldots$ is v_a, with $v_a > v_{a+1}$. As in the text, each gasoline car emits e, regardless of its age, while electric vehicles are pollution-free. Consumers discount the future with discount factor $\delta < 1$. To simplify the exposition, we assume that new gasoline cars and electric vehicles have the same production cost, $c_G = c_E \equiv c$, which is the price at which new cars are sold (c may include a constant markup).

A LEZ affects gasoline cars by reducing their age-specific flow value from v_a to $v_a - L_a$, where $L_a \geq 0$ is allowed to differ across vintages. In the text we use $L_a = sv_a$, but here we adopt a more general form, in part to see more precisely how a LEZ might affect prices of second-hand units and scrappage age.

The figure below illustrates the allocation. In the absence of intervention, the indifferent consumer is denoted by $\tilde{x}$. After the LEZ is introduced, the corresponding indifferent consumer in the new steady state can be denoted by $\hat{x} < \tilde{x}$. Consumers located between $\hat{x}$ and $\tilde{x}$ are central for the transition: absent trade, they continue holding their gasoline cars until scrappage; with frictionless trade, they may sell them immediately and switch to electric vehicles.

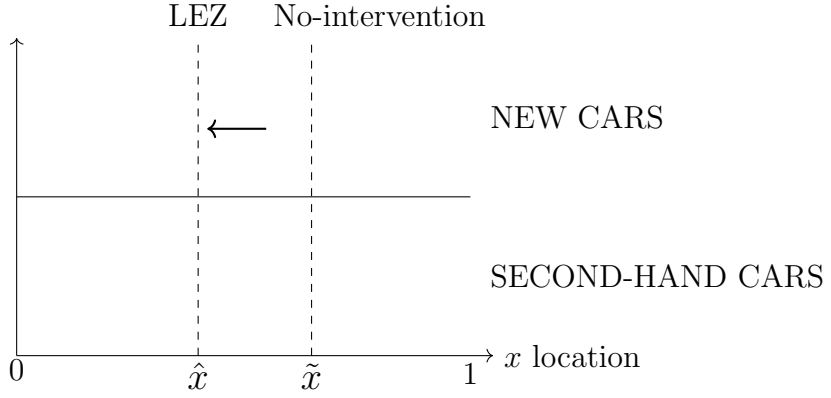


Figure D.1: Allocation of car types by location and intervention regime

Throughout this appendix, we will make two assumptions. The first is that the LEZ shifts some, but not all, consumers away from gasoline cars. This requires γ, the differentiation parameter, to be large enough relative to the LEZ impacts on gasoline drivers. Formally, we assume

$$\gamma > \frac{L_1 + \delta L_2}{1 + \delta}.$$

The second assumption ensures that the resale-friction cutoff is nontrivial, so that there exists some $\mu_c \in (0, 1)$ separating the no-trade and trade regimes. Formally, we assume that

$$L_2 < \frac{c + \delta z - \Delta v}{1 + \delta},$$

where $\Delta v \equiv v_1 - v_2 > 0$.

D.2 No intervention

In the no-intervention equilibrium, both new and second-hand vehicles of each type are unrestricted. Let $U_a^j(x)$ denote the continuation value of a consumer located at x who holds a type-j car of age a, for $j \in \{G, E\}$ and $a \in \{1, 2\}$.

For electric-vehicle owners,

$$U_1^E(x) = v_1 - \gamma(1 - x) - c + \delta U_2^E(x),$$

$$U_2^E(x) = v_2 - \gamma(1 - x) + \delta U_1^E(x) + \delta z.$$

For gasoline-car owners,

$$U_1^G(x) = v_1 - \gamma x - c + \delta U_2^G(x),$$

$$U_2^G(x) = v_2 - \gamma x + \delta U_1^G(x) + \delta z.$$

The no-intervention cutoff $\tilde{x}$ is defined by

$$U_1^G(\tilde{x}) = U_1^E(\tilde{x}).$$

Since $c_G = c_E = c$, symmetry implies $\tilde{x} = 1/2$. Thus one half of the market uses gasoline cars and the other half uses electric vehicles.

Second-hand prices clear the market by making consumers indifferent, within a vehicle type, between buying a new unit and buying a one-period-old unit. Thus,

$$U_1^j(x) = U_2^j(x) - \tilde{p}_2^j, \qquad j \in \{G, E\}.$$

Because the two technologies have the same cost,

$$\tilde{p}_2^G = \tilde{p}_2^E = \frac{c + \delta z - \Delta v}{1 + \delta} \equiv p_2^*.$$

In the no-intervention equilibrium, second-hand cars trade at the same price regardless of their type.

For cars to last exactly two periods, consumers must prefer using a one-period-old car rather than scrapping it immediately, but must prefer scrapping a two-period-old car rather than keeping it for a third period. Let v_3 denote the value of a car of age 3. The two-period scrappage condition is

$$v_2 > \frac{v_1 + \delta v_2 + z - c}{1 + \delta} > v_3. \tag{D.1}$$

The first inequality ensures that cars are held for more than one period. The second ensures that they are scrapped at the beginning of the third period.

D.3 Time-invariant LEZ with trade frictions

Let frictions in the second-hand market be such that no trade in second-hand units takes place. Below we will be precise about how trade frictions reduce trade in the second-hand market. At the beginning of period $t = 1$, the planner introduces a time-invariant LEZ that restricts gasoline cars. It is optimal for the planner to do so if

$$e > e_G^{ITF} \equiv \gamma \frac{L_1 + L_2}{L_1 + \delta L_2} - \frac{L_1 + \delta L_2}{2(1 + \delta)},$$

where the superscript *ITF* denotes time-invariant LEZ with trade frictions. A gasoline car of age a now delivers flow value $v_a - L_a$, while electric vehicles remain unrestricted.

In the post-intervention allocation, the cutoff $\hat{x}$ is defined by

$$U_1^G(\hat{x}) = U_1^E(\hat{x}).$$

Solving the two sets of Bellman equations gives the cutoff

$$\hat{x} = \frac{1}{2} - \frac{L_1 + \delta L_2}{2\gamma(1 + \delta)}. \tag{D.2}$$

Thus the LEZ shifts consumers away from gasoline cars and toward electric vehicles.

Second-hand prices again clear by making consumers indifferent, within each vehicle type, between buying a new unit and buying a one-period-old unit. For electric vehicles, nothing changes: $\hat{p}_2^E = p_2^*$. For gasoline cars, the LEZ changes the relative value of new

and second-hand units:

$$\hat{p}_2^G = \frac{c + \delta z - \Delta v - L_2 + L_1}{1+\delta} = p_2^* + \frac{L_1 - L_2}{1+\delta}.$$

This expression shows how the age profile of the LEZ loss affects used-car prices. If $L_1 = L_2$, the LEZ lowers the value of new and second-hand gasoline cars equally, and the second-hand gasoline price remains unchanged, $\hat{p}_2^G = p_2^*$. Interestingly, if $L_a = sv_a$, so $L_1 > L_2$, the LEZ reduces the value of new gasoline cars more than that of second-hand ones, raising the relative price of the latter.

The two-period scrappage condition for gasoline cars must also be adjusted. If the LEZ loss for a potential age-three gasoline car is L_3, then gasoline cars continue to be scrapped after two periods if

$$v_2 - L_2 > \frac{v_1 + \delta v_2 + z - c}{1+\delta} - \frac{L_1 + \delta L_2}{1+\delta} > v_3 - L_3. \tag{D.3}$$

When $L_1 = L_2 = L_3 = L$, this condition is exactly the no-intervention condition (D.1) shifted down by the common loss L, so the scrappage age is unchanged. By contrast, if $L_1 > L_2 > L_3$, as under a proportional loss $L_a = sv_a$, the right-hand inequality may fail if s is large enough: the LEZ can make it optimal to keep gasoline cars for longer, because the policy narrows the quality gap between new and old gasoline cars.

We now turn to the role of trade frictions during the transition. At $t = 1$, the economy inherits the no-intervention allocation. Thus consumers in the interval $(\hat{x}, \tilde{x})$ hold one-period-old gasoline cars even though, under the new LEZ, they would prefer to be electric-vehicle users.

Let second-hand-market frictions take the following reduced form. Buyers pay price p, but sellers receive only a fraction μp, with $\mu \in [0, 1]$. Any other division of transaction costs delivers the same results. The case $\mu = 1$ corresponds to a frictionless second-hand market, while lower values of μ capture larger resale frictions.

Consumers to the left of $\hat{x}$ are willing to buy a second-hand gasoline car up to the price $p = \hat{p}_2^G$. A consumer $x \in (\hat{x}, \tilde{x})$ who holds a one-period-old gasoline car at the time the LEZ is introduced compares two options. If they keep the gasoline car for one more period, they obtain

$$v_2 - L_2 - \gamma x + \delta z + \delta U_1^E(x).$$

If they sell the gasoline car immediately and switch to a new electric vehicle, they obtain

$$\mu \hat{p}_2^G + U_1^E(x).$$

Therefore, they sell if and only if

$$\mu \hat{p}_2^G + U_1^E(x) \geq v_2 - L_2 - \gamma x + \delta z + \delta U_1^E(x).$$

Equivalently, the set of early switchers is $[x_\mu, \tilde{x}]$, where

$$x_\mu = \hat{x} + \frac{(1-\mu)\hat{p}_2^G}{2\gamma} \leq \tilde{x}.$$

Hence, the mass of LEZ-induced early switchers is

$$m(\mu) = (\tilde{x} - x_\mu) = \left[\frac{L_1 + \delta L_2}{2\gamma(1+\delta)} - \frac{(1-\mu)\hat{p}_2^G}{2\gamma}\right] \geq 0.$$

Thus, frictions are large enough to suppress second-hand-market reallocation if and only if

$$m(\mu) = 0,$$

or equivalently, if

$$(1-\mu)\hat{p}_2^G \geq \frac{L_1 + \delta L_2}{1+\delta}.$$

This condition is the formal sense in which frictions need not be prohibitively large. They only need to make the resale wedge larger than the discounted average LEZ loss borne by a gasoline car over its remaining two-period cycle. In terms of μ, the condition can be written as

$$\mu \leq \mu_c \equiv 1 - \frac{L_1 + \delta L_2}{(1+\delta)\hat{p}_2^G}. \tag{D.4}$$

When $L_1 = L_2 = L$, we have $\hat{p}_2^G = p_2^*$, and the condition reduces to $(1-\mu)p_2^* \geq L$, or

$$\mu \leq 1 - \frac{L}{p_2^*}.$$

This is the condition we assume in Section 3. The frictions documented by Gavazza, Lizzeri, and Roketskiy (2014) range from 15% for one-year-old cars to almost 80% for cars twenty years old. If anything, the assumption is reinforced by the fact that we abstract from vertical differentiation. It would be hard to find drivers with strong preferences for gasoline cars who, when ready to buy a new model, would instead be willing to buy old second-hand units, say, cars more than ten years old.

Under condition (D.4), the consumers in $(\hat{x}, \tilde{x})$ do not sell their gasoline cars immediately. They instead keep them until scrappage and switch only when they replace their vehicles. Hence, the adjustment from the inherited no-intervention allocation to the post-LEZ allocation occurs through the gradual replacement of the vehicle stock, rather than through immediate reallocation in the second-hand market. It is important to emphasize that for our results in the text to follow, we just need a fraction of the consumers in $(\hat{x}, \tilde{x})$ to keep their cars until scrappage; in reality, those holding cars that are ten or more years old.

D.4 Time-invariant LEZ with no frictions

We now consider the opposite case, in which the second-hand market is frictionless, $\mu = 1$. Buyers and sellers face the same transaction price, and the resale wedge disappears.

The post-LEZ cutoff remains unchanged, given by (D.2). The second-hand prices that clear the market are the same as in the post-intervention steady state:

$$p_{2,t}^E = p_2^*,$$

and

$$p_{2,t}^G = \hat{p}_2^G = p_2^* + \frac{L_1 - L_2}{1+\delta},$$

for all $t \geq 1$.

With no resale wedge, all consumers in $(\hat{x}, \tilde{x})$ who hold one-period-old gasoline cars at the time of the LEZ introduction sell immediately and switch to new electric vehicles. The reason is that, at the buyer price $\hat{p}_2^G$, the seller receives the full price. Since $x_\mu = \hat{x}$ when $\mu = 1$, the mass of early switchers is

$$m(1) = \tilde{x} - \hat{x} = \frac{L_1 + \delta L_2}{2\gamma(1+\delta)}.$$

Let q_t denote new electric-vehicle purchases in period t. Since cars last two periods, the stock of electric vehicles in period t is the sum of new electric vehicles bought in period t and one-period-old electric vehicles bought in period $t-1$: $q_t + q_{t-1}$. Under the LEZ, the desired electric-vehicle stock share is

$$\alpha \equiv 1 - \hat{x} = \frac{1}{2} + \frac{L_1 + \delta L_2}{2\gamma(1+\delta)}.$$

Therefore, equilibrium requires

$$q_t + q_{t-1} = \alpha$$

for all $t \geq 1$. Since the economy starts from the no-intervention steady state, the initial mass of new electric-vehicle purchases is $q_0 = 1/4$. It follows that

$$q_1 = \alpha - q_0 = \frac{1}{4} + \frac{L_1 + \delta L_2}{2\gamma(1+\delta)},$$

$$q_2 = \alpha - q_1 = \frac{1}{4},$$

and, by iteration,

$$q_{2k+1} = \frac{1}{4} + \frac{L_1 + \delta L_2}{2\gamma(1+\delta)}, \qquad q_{2k} = \frac{1}{4},$$

for all $k \geq 0$. New gasoline-car purchases move in the opposite direction: $g_t = 1/2 - q_t$.

Thus, with a frictionless second-hand market, the LEZ generates an oscillatory transition in new-vehicle purchases. The stock of electric vehicles is constant at the post-LEZ level α, but its age composition alternates over time. In odd periods, the market absorbs an unusually large cohort of new electric vehicles; in the following period, that large cohort appears as a large supply of second-hand electric vehicles, reducing new electric-vehicle purchases. Prices need not oscillate. The cycle is driven by cohort dynamics and second-hand-market reallocation, not by cyclical price movements. This is not different from the cohort dynamics emphasized by Adda and Cooper (2000).

The two-period scrappage conditions remain those stated above. Electric vehicles continue to be scrapped after two periods if (D.1) holds, while gasoline cars continue to be scrapped after two periods if (D.3) holds. Provided these conditions hold, the frictionless equilibrium is well defined with a two-period vehicle lifetime.

D.5 Gasoline cars as a transition technology

The discussion above also clarifies why treating gasoline cars as a transition technology may depend critically on second-hand-market frictions. In the no-trade case, an announced future restriction can affect current new-car purchases while protecting current gasoline owners from an immediate loss. The policy therefore has a transition-technology logic: gasoline cars are left unrestricted initially, but consumers anticipate their lower future value and some switch toward electric vehicles before the restriction becomes active. When frictions are such that no trade in second-hand cars takes place, it is optimal for the planner to treat gasoline cars as a transition technology if $e \in (e_{TTL}^{VTF}, e_{TTU}^{VTF})$, where

$$\begin{aligned} e_{TTL}^{VTF} &\equiv \frac{2\gamma(1+\delta)(L_1+L_2) - L_1^2 - 2\delta L_1 L_2 - \delta L_2^2}{2(1+\delta)(L_1+L_2)} \\ < e_G^{ITF} &< \frac{2\gamma(1+\delta)(L_1+L_2) - L_1^2 - 2\delta L_1 L_2}{2(1+\delta)L_1} \equiv e_{TTU}^{VTF} \end{aligned}$$

and the superscript VTF stands for time-varying LEZ with trade frictions; otherwise, if $e \leq e_{TTL}^{VTF}$ it is optimal to never restrict them, while if $e \geq e_{TTU}^{VTF}$ it is optimal to restrict them from $t = 1$ onward.

This transition-technology logic weakens when second-hand-market frictions fall. To see this most clearly, consider the symmetric-loss case $L_1 = L_2 = L$. As shown above, the mass of consumers in $(\hat{x}, \tilde{x})$ who sell their second-hand gasoline car and switch immediately to an electric vehicle is

$$m(\mu) = \left(\frac{L - (1-\mu)p_2^*}{2\gamma} \right) \geq 0.$$

Since the total mass of consumers between the old and new cutoffs is

$$\tilde{x} - \hat{x} = \frac{L}{2\gamma},$$

the residual mass of gasoline owners who remain locked in because of resale frictions is

$$\ell(\mu) \equiv \frac{L}{2\gamma} - m(\mu) = \min\left\{\frac{L}{2\gamma}, \frac{(1-\mu)p_2^*}{2\gamma}\right\}.$$

This residual locked-in mass is the source of the transition-technology rationale. Delaying the restriction is valuable only to the extent that it avoids imposing immediate losses on gasoline owners who cannot otherwise reallocate through the second-hand market. As the second-hand market becomes frictionless, $\mu \to 1$, this mass converges to zero: $\lim_{\mu\to 1} \ell(\mu) = 0$. Thus, in the frictionless limit, gasoline owners who would be hurt by the restriction can sell their cars and switch immediately to electric vehicles. The announcement of a future restriction no longer needs to protect a locked-in stock of gasoline users; the second-hand market performs the reallocation instead. This implies that for any $\mu < 1$, some gasoline owners remain locked in, and the transition-technology logic remains.

It can also be shown that in the limit, as $\mu \to 1$, the cutoffs e_{TTL}^{VTF} and e_{TTU}^{VTF} converge to

$$e_G^{INF} = \frac{\gamma(1+\delta)(L_1+L_2)}{2(L_1+\delta L_2)} - \frac{L_1+\delta L_2}{2(1+\delta)},$$

where the superscript *INF* stands for time-invariant LEZ with no frictions. Not surprisingly, $e_G^{INF} < e_G^{ITF}$. With no frictions, the second-hand market immediately reallocates gasoline cars toward consumers less affected by the LEZ, reducing the welfare cost of introducing a LEZ in the first place.

Appendix E: Quasi-convexity of the welfare function

The first section of this appendix shows that the bang-bang logic in Proposition 2 extends beyond the uniform distribution of consumers around the circle. For simplicity, here we consider just the case in which the LEZ restricts gasoline cars but not hybrids. Extending the analysis to hybrids follows the same steps. We provide sufficient conditions on the distribution of consumers that ensure quasi-convexity of welfare in $s \in [0, \bar{s}]$. We then present a series of illustrations based on mixtures of normal distributions and a multiple-peaked density for which the bang-bang logic fails. Finally, we provide an example in which the welfare ordering implied by our model with uniformly distributed consumers is reversed: restricting both gasoline and hybrid cars can be welfare-enhancing even when restricting gasoline cars alone is not. Importantly, the bang-bang logic of LEZ design remains intact.

E.1 Quasi-convexity conditions

Consider the Salop model with its three short-lived vehicle types: electric vehicles E, hybrids H, and gasoline cars G. These technologies are located at 0, 1/3, and 2/3, respectively, on a circle of unit circumference. Consumers are located at $x \in [0, 1]$ according to a continuous distribution F, with density f. Recall that external damages are 0 for each electric vehicle, αe for each hybrid, and e for each gasoline car.

The consumer indifferent between electric and hybrid vehicles is located at $x_{EH} = 1/6$, between hybrids and gasoline cars is located at $x_{HG}(s) = 1/2 + sv/2\gamma$, and between gasoline cars and electric vehicles is located at $x_{GE}(s) = 5/6 - sv/2\gamma$. Thus,

$$q_E(s) = F(x_{EH}) + 1 - F(x_{GE}(s)),$$

$$q_H(s) = F(x_{HG}(s)) - F(x_{EH}),$$

and

$$q_G(s) = F(x_{GE}(s)) - F(x_{HG}(s)).$$

Social welfare is therefore given by

$$\begin{aligned} W(s) = &\int_0^{x_{EH}} u_E(x) f(x)\, dx + \int_{x_{EH}}^{x_{HG}(s)} u_H(x) f(x)\, dx \\ &+ \int_{x_{HG}(s)}^{x_{GE}(s)} u_G(x; s) f(x)\, dx + \int_{x_{GE}(s)}^{1} u_E(x) f(x)\, dx - \alpha e q_H(s) - e q_G(s), \end{aligned}$$

where $u_j(x)$ is the surplus a consumer located at x obtains from driving a type-j car. Differentiating with respect to s, and using the indifference conditions at $x_{HG}(s)$ and

$x_{GE}(s)$, gives

$$W'(s) = v\left[\frac{(1-\alpha)e}{2\gamma} f(x_{HG}(s)) + \frac{e}{2\gamma} f(x_{GE}(s)) - q_G(s)\right].$$

The term $-vq_G(s)$ is the direct utility loss imposed on consumers who continue to use gasoline cars. The first density term captures the externality benefit from marginal switches from gasoline to hybrid vehicles; the second captures the benefit from marginal switches from gasoline to electric vehicles.

The second derivative is

$$W''(s) = \frac{v^2}{2\gamma}\left[f(x_{HG}(s)) + f(x_{GE}(s)) + \frac{(1-\alpha)e}{2\gamma} f'(x_{HG}(s)) - \frac{e}{2\gamma} f'(x_{GE}(s))\right].$$

Therefore, welfare is convex in $s \in [0, \bar{s}]$ if

$$2\gamma\big[f(x_{HG}(s)) + f(x_{GE}(s))\big] + (1-\alpha)ef'(x_{HG}(s)) - ef'(x_{GE}(s)) \geq 0.$$

This condition, which clearly holds for the uniform distribution, is sufficient for a bang-bang solution, but it is stronger than necessary.

We now state a weaker sufficient condition. Let

$$J = [x_{HG}(0), x_{GE}(0)] = \left[\frac{1}{2}, \frac{5}{6}\right].$$

This interval contains every gasoline segment $[x_{HG}(s), x_{GE}(s)]$ for $s \in [0, \bar{s}]$, since $x_{HG}(s)$ increases and $x_{GE}(s)$ decreases with s. Assume that $f(x)$ is positive and continuously differentiable on a neighborhood of J, and that

$$\ell(x) \equiv \log f(x)$$

is concave on J, so $f(x)$ is log-concave. Equivalently,

$$\ell'(x) = \frac{f'(x)}{f(x)}$$

is nonincreasing on J. This restriction applies only to the region spanned by gasoline demand. It does not require the density to be uniform, symmetric, or globally log-concave around the circle.

Define

$$r_{HG}(s) \equiv \frac{f(x_{HG}(s))}{q_G(s)}, \qquad r_{GE}(s) \equiv \frac{f(x_{GE}(s))}{q_G(s)}.$$

These ratios measure the density at each margin relative to the mass of consumers who continue to choose gasoline cars.

Proposition. If $f(x)$ is positive and log-concave on J, then $W(s)$ is quasi-convex on $[0, \bar{s}]$.

Proof. Since

$$q'_G(s) = -\frac{v}{2\gamma}\big[f(x_{HG}(s)) + f(x_{GE}(s))\big],$$

we have, using $\ell'(x) = f'(x)/f(x)$, that

$$r'_{HG}(s) = \frac{vf(x_{HG}(s))}{2\gamma q_G(s)^2}\left[\ell'(x_{HG}(s))q_G(s) + f(x_{HG}(s)) + f(x_{GE}(s))\right].$$

Similarly,

$$r'_{GE}(s) = \frac{vf(x_{GE}(s))}{2\gamma q_G(s)^2}\left[-\ell'(x_{GE}(s))q_G(s) + f(x_{HG}(s)) + f(x_{GE}(s))\right].$$

Log-concavity of $f(x)$ implies that $\ell'(x)$ is nonincreasing. Therefore, for every $x \in [x_{HG}(s), x_{GE}(s)]$,

$$\ell'(x_{HG}(s)) \geq \ell'(x) \geq \ell'(x_{GE}(s)).$$

Multiplying by $f(x) > 0$ and integrating over $[x_{HG}(s), x_{GE}(s)]$ yields

$$\begin{aligned}\ell'(x_{HG}(s))q_G(s) &\geq \int_{x_{HG}(s)}^{x_{GE}(s)} \ell'(x)f(x)\,dx = \int_{x_{HG}(s)}^{x_{GE}(s)} f'(x)\,dx \\ &= f(x_{GE}(s)) - f(x_{HG}(s)) \geq \ell'(x_{GE}(s))q_G(s).\end{aligned}$$

Hence,

$$\ell'(x_{HG}(s))q_G(s) + f(x_{HG}(s)) + f(x_{GE}(s)) \geq 2f(x_{GE}(s)) > 0,$$

and

$$-\ell'(x_{GE}(s))q_G(s) + f(x_{HG}(s)) + f(x_{GE}(s)) \geq 2f(x_{HG}(s)) > 0.$$

It follows that

$$r'_{HG}(s) > 0, \qquad \text{and} \qquad r'_{GE}(s) > 0.$$

Using these ratios, the first derivative of welfare can be written as

$$W'(s) = vq_G(s)\big[R(s) - 1\big],$$

where

$$R(s) \equiv \frac{e}{2\gamma}\left[(1-\alpha)r_{HG}(s) + r_{GE}(s)\right].$$

Since both $r_{HG}(s)$ and $r_{GE}(s)$ are strictly increasing, we have $R'(s) > 0$. Therefore, $W'(s)$ can change sign at most once, and if it does, it changes from negative to positive. Hence welfare is either monotone or first decreases and then increases. In either case, for any

$s_0, s_1 \in [0, \bar{s}]$ and any $\lambda \in [0, 1]$,

$$W(\lambda s_1 + (1 - \lambda)s_0) \leq \max\{W(s_0), W(s_1)\}.$$

Thus, $W(s)$ is quasi-convex on $[0, \bar{s}]$.

E.2 Illustrations

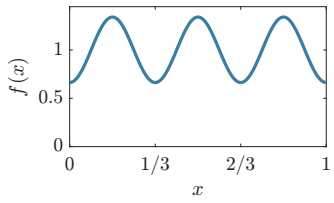
Log-concavity is a standard shape restriction satisfied by many common densities. Bagnoli and Bergstrom (2005), for example, have obtained log-concavity for distributions such as normal, logistic, extreme-value, exponential, and Laplace; Weibull, power function, gamma, and beta are log-concave for shape parameters of at least one. The examples below consider different combinations of normal distributions that generate quasi-convexity in the welfare function.

To start, suppose that $\gamma = 4$, $v = 1.7$, $e = 0.7$, $\alpha = 0$, and, as shown in Figure E.1, Panel (a), consumers are distributed along the circle according to a mixture of three normal densities, each centered halfway between two adjacent vehicle locations, with standard deviation $\sigma = 0.1$. This is an example in which consumers exhibit weak preferences for a particular type of vehicle. A large fraction of consumers are indifferent between two types of cars, all else equal. As shown in Figure E.1, Panel (b), the welfare function is convex everywhere in LEZ size s, even as gasoline cars exit the market, i.e., as s approaches $\gamma/3v \approx 0.78$.

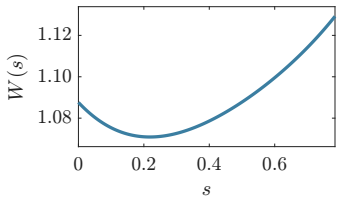
Consider now reducing σ by half, to 0.05. Now consumers exhibit even weaker preferences for a particular type; that is, most consumers are located halfway between two vehicle locations, as shown in Figure E.1, Panel (c). As shown in Figure E.1, Panel (d), welfare exhibits local concavity but continues to be quasi-convex in LEZ size.

We now move the centers of the density from the midpoints between technologies to the vehicle locations, while keeping $\sigma = 0.05$. In this case consumers exhibit very strong preferences for a particular vehicle type: most of the mass is concentrated around the electric, hybrid, and gasoline vehicle locations, as shown in Figure E.1, Panel (e). Welfare is convex throughout the relevant policy range. The optimal endpoint, however, depends on the magnitude of the pollution externality. With the externality set at its original level, $e = 0.7$, the optimal policy is not to implement a LEZ, as shown in Figure E.1, Panel (f). However, if the externality more than doubles, to $e = 1.5$, then, as indicated in Figure E.1, Panel (f), it is optimal to implement the largest possible LEZ, when gasoline cars are fully displaced from the market.

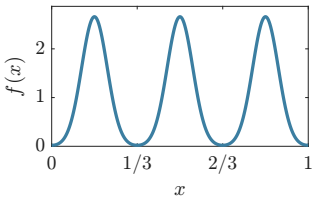
We close the subsection by illustrating what would be required, in terms of consumer preferences, for the welfare function to lose its quasi-convexity. When consumer preferences are captured by a multiple-peaked distribution like the one in Figure E.1, Panel (g), welfare need not be quasi-convex. As shown in Figure E.1, Panel (h), in this case, with



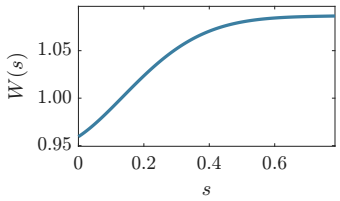
(a) Weak preferences for a particular vehicle type.



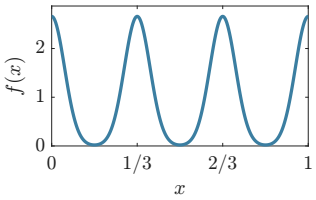
(b) Welfare is convex everywhere.



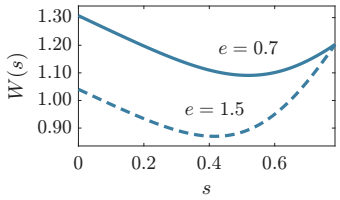
(c) Very weak preferences for a particular vehicle type.



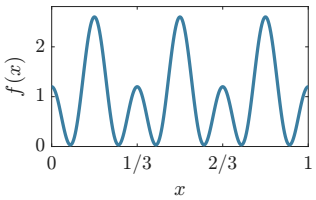
(d) Welfare is locally concave but globally quasi-convex.



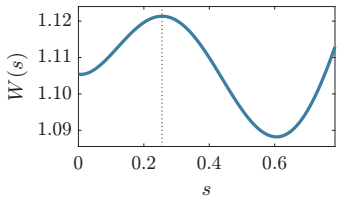
(e) Very strong preferences for a particular vehicle type.



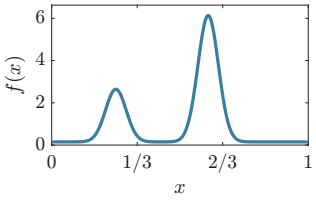
(f) The optimal LEZ policy is no LEZ if $e = 0.7$ and the largest feasible LEZ if $e = 1.5$.



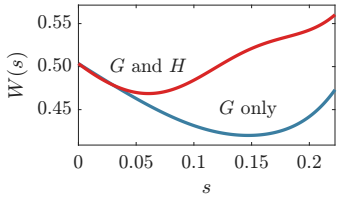
(g) Consumer preferences are distributed according to a multiple-peaked distribution.



(h) The optimal LEZ policy is not bang-bang.



(i) Non-uniform distribution of consumers.



(j) Restricting hybrids as well reverses the welfare ordering.

Figure E.1: Illustrations

Notes: Each row shows a density of consumers, $f(x)$, and the welfare it implies, $W(s)$, for LEZ sizes from 0 to $s^{\max} = \gamma/3v$; EVs, hybrids, and gasoline cars are located at $x = 0$, $1/3$, and $2/3$. Panels (a)–(h) restrict gasoline cars, with $\gamma = 4$, $v = 1.7$, $\alpha = 0$, and $e = 0.7$, except for the dashed line in Panel (f), where $e = 1.5$, and Panel (h), where $e = 0.5$ and the dotted line marks the optimal LEZ. Panels (i) and (j) use the example of Appendix E.3, in which the LEZ restricts gasoline cars only (G only) or hybrids as well (G and H).

$e = 0.5$, the optimal LEZ policy is no longer bang-bang; it is rather an intermediate-size LEZ.

E.3 Reverse welfare ordering

It is possible that restricting both gasoline/diesel cars and hybrids generates positive welfare gains, while restricting gasoline/diesel cars alone does not. This possibility is ruled out by our Salop model with uniformly distributed consumers, but it can arise under more general preference distributions while preserving the quasi-convexity of the welfare function.

Consider short-lived vehicles and let $\gamma = 1$, $v = 1.5$, $e = 1$, $\alpha = 0.94$, so the maximum feasible LEZ size is $s^{\max} = \gamma/3v = 2/9$. Suppose that consumers are distributed according to the mixture

$$f(x) = 0.15 + 0.25\, N(0.25, 0.04^2) + 0.60\, N(0.61, 0.04^2),$$

normalized over $[0, 1]$. Thus, most consumers are located near the gasoline car location, while a smaller mass is located between the electric and hybrid locations. Figure E.1, Panel (i) plots the corresponding density. Under no intervention, most consumers choose gasoline cars, 65%, and most of the remaining 35% choose hybrids. When only gasoline cars are restricted to the maximum extent possible (i.e., $s = s^{\max}$), substitution occurs primarily toward hybrids, which in this example are only slightly cleaner than gasoline cars. As a result, the environmental gains are not large enough to compensate for the private losses from the restriction, leading to a net welfare loss of 6%. By contrast, when both gasoline cars and hybrids are restricted to the maximum extent possible, the induced substitution toward electric vehicles is strong enough to generate a net welfare gain of 11%.

Importantly, as shown in Figure E.1, Panel (j), both welfare functions remain quasi-convex in LEZ size over the interval $s \in [0, s^{\max}]$. Hence, the bang-bang logic is preserved, even though the welfare ordering between a gasoline-only LEZ and a LEZ that also restricts hybrids differs from that of the symmetric model with uniformly distributed consumers.

Appendix F: Marginal welfare effect decomposition

In this appendix, we show that the marginal welfare effect of the LEZ's intensity is negative at every intensity $\lambda \geq 0$. We first show that $\Delta E_i(1)/(1-\rho) \geq \varsigma_i(\lambda)\,\Delta E_i(\lambda)$ for every driver at every intensity, by tracing how each driver's benefit of a util changes with λ. We then show that, at our estimates, $1/\alpha_i \geq \Delta E_i(1)/(1-\rho)$ for virtually all drivers, and that the few driver types for which it fails cannot overturn the sign.

F.1 The benefit of a util at every intensity

Let G and $\mathcal{G}$ denote the sets of restricted products (gasoline and diesel cars) and unrestricted products (LPG/CNG, hybrid, and electric cars). We consider the markets in which the restriction applies and suppress market subscripts. As in Section 4.3, every product in G loses utility $\phi_i\lambda$, with $\lambda = 1$ in Madrid Central. We split the sum inside the surplus formula (17) into the sums over G and $\mathcal{G}$,

$$\Lambda_i^G(\lambda) \equiv \sum_{(m,j)\in G} \exp\left(\frac{v_{imj} - \alpha_i p_{mj} - \phi_i\lambda}{1-\rho}\right), \qquad \Lambda_i^{\mathcal{G}} \equiv \sum_{(m,j)\in \mathcal{G}} \exp\left(\frac{v_{imj} - \alpha_i p_{mj}}{1-\rho}\right),$$

so that $\Lambda_i^G(\lambda) = \Lambda_i^G(0)\, e^{-\phi_i\lambda/(1-\rho)}$, while $\Lambda_i^{\mathcal{G}}$ does not depend on λ. With $\Lambda_i \equiv \Lambda_i^G + \Lambda_i^{\mathcal{G}}$, the probability that consumer i buys a car is $S_{\text{car},i} = \Lambda_i^{1-\rho}/(1+\Lambda_i^{1-\rho})$. Conditional on buying one, they buy a restricted car with probability $S_{G|\text{car},i} = \Lambda_i^G/\Lambda_i$. Both probabilities are the shares defined in footnote 65. The probabilities that they buy a restricted car, buy an unrestricted car, or buy no new car are then $S_{G,i} = S_{\text{car},i}S_{G|\text{car},i}$, $S_{\mathcal{G},i} = S_{\text{car},i}(1 - S_{G|\text{car},i})$, and $S_{0,i} = 1 - S_{\text{car},i}$, respectively.

By footnote 65, the demand that driver i diverts to each unrestricted type is proportional to its share. Among the drivers who switch away from restricted cars when λ increases, the shares going to unrestricted cars and to the outside option are given by the diversion ratios

$$DR_{\mathcal{G},i} = \frac{\rho\,(1 - S_{G|\text{car},i}) + (1-\rho)\,S_{\mathcal{G},i}}{(1-\rho)\,\varsigma_i}, \qquad DR_{0,i} = \frac{S_{0,i}}{\varsigma_i}.$$

The saving per displaced car in (21) is then $\Delta E_i = E_{G,i} - DR_{\mathcal{G},i}\,E_{\mathcal{G},i}$, where $E_{\mathcal{G},i}$ is driver i's average lifetime externality cost among unrestricted cars, weighted by their choice probabilities. Like $E_{G,i}$, it does not depend on λ. Collecting terms in (22), the marginal welfare effect is

$$\frac{\partial \bar{W}}{\partial \lambda} = \int w_i\,\phi_i\,S_{G,i}\left(H_i - \frac{1}{\alpha_i}\right) di, \qquad H_i \equiv \varsigma_i\,\Delta E_i, \tag{F.1}$$

where H_i is the emission benefit of a util, as in condition (23). We assume that $E_{G,i} \geq$

$E_{\not G,i} \geq 0$ for every driver, which holds at our estimates. Since $DR_{\not G,i} \leq 1$, the saving per displaced car is then nonnegative, $\Delta E_i \geq E_{G,i} - E_{\not G,i} \geq 0$.

Lemma F.1 *For each driver i, as λ increases: (i) the diversion intensity ς_i increases and stays below $1/(1-\rho)$; (ii) the diversion ratio to the outside option, $DR_{0,i}$, decreases, so the saving per displaced car, ΔE_i, cannot increase; (iii) the benefit of a util, H_i, increases. Consequently, $H_i(\lambda) \leq \Delta E_i(1)/(1-\rho)$ at every $\lambda \geq 0$, where $\Delta E_i(1)$ is the saving per displaced car at $\lambda = 1$.*

Parts (i) and (ii) pull the benefit of a util in opposite directions. As restricted cars become a smaller share of driver i's choices, each util of disutility displaces more of their restricted demand. A displaced buyer also becomes less likely to leave the market. Leaving saves the full damage of a restricted car, while switching to an unrestricted car saves only the difference in damages. Part (iii) shows that the first force dominates. Above the current intensity, parts (i) and (ii) bound the two factors of H_i. Below it, part (iii) keeps H_i below its value at $\lambda = 1$.

Proof. Raising λ lowers Λ_i^G, and hence Λ_i, while $\Lambda_i^{\not G}$ is unchanged. (i) The numerator of ς_i lies between zero and one, so $0 < \varsigma_i < 1/(1-\rho)$. Raising λ lowers $S_{\text{car},i}$ and $S_{G|\text{car},i} = \Lambda_i^G/\Lambda_i$, and hence $S_{G,i}$, so ς_i rises. (ii) Using $1 - S_{G|\text{car},i} = \Lambda_i^{\not G}/\Lambda_i$, $S_{\not G,i} = \Lambda_i^{\not G}/(\Lambda_i^\rho + \Lambda_i)$, and $S_{0,i} = 1/(1+\Lambda_i^{1-\rho})$, the ratio of the two diversion ratios is

$$\frac{DR_{\not G,i}}{DR_{0,i}} = \frac{\rho\,(1 - S_{G|\text{car},i}) + (1-\rho)\,S_{\not G,i}}{(1-\rho)\,S_{0,i}} = \frac{\Lambda_i^{\not G}\left(\rho\,\Lambda_i^{-1} + \Lambda_i^{-\rho}\right)}{1-\rho},$$

which rises as Λ_i falls. Since the two diversion ratios sum to one, $DR_{0,i}$ falls. Hence $\Delta E_i = E_{G,i} - DR_{\not G,i}\,E_{\not G,i}$ cannot rise. (iii) Rewrite

$$H_i = \frac{(E_{G,i} - E_{\not G,i})\left[\rho\,(1 - S_{G|\text{car},i}) + (1-\rho)\,S_{\not G,i}\right] + (1-\rho)\,E_{G,i}\,S_{0,i}}{1-\rho}.$$

As λ rises, Λ_i falls, which raises $1 - S_{G|\text{car},i}$, $S_{\not G,i}$, and $S_{0,i}$. The coefficients are nonnegative, so H_i increases. Finally, for $\lambda \geq 1$, parts (i) and (ii) and $\Delta E_i \geq 0$ give $H_i(\lambda) = \varsigma_i(\lambda)\,\Delta E_i(\lambda) \leq \Delta E_i(1)/(1-\rho)$. For $\lambda < 1$, parts (iii) and (i) give $H_i(\lambda) \leq H_i(1) \leq \Delta E_i(1)/(1-\rho)$. □

Proposition F.1 *Suppose that $1/\alpha_i \geq \Delta E_i(1)/(1-\rho)$ for every driver, with strict inequality for a set of drivers of positive weight. Then $\partial \bar{W}/\partial\lambda < 0$ at every $\lambda \geq 0$.*

Proof. By Lemma F.1 and this condition, $H_i \leq \Delta E_i(1)/(1-\rho) \leq 1/\alpha_i$ at every $\lambda \geq 0$, with the second inequality strict for a set of drivers of positive weight. Since $\phi_i S_{G,i} > 0$, the integrand in (F.1) is nonpositive for every driver and negative on that set. □

Proposition F.1 requires, for each driver, that the ceiling $\alpha_i\,\Delta E_i(1)/(1-\rho)$ be at most one. By Lemma F.1, a euro of driver i's private cost never buys more than this

ceiling in emission savings, at any intensity. Because the comparison holds driver by driver, it does not depend on how the composition of restricted-car buyers changes with λ.

F.2 Evaluation at our estimates

We evaluate the condition of Proposition F.1, $1/\alpha_i \geq \Delta E_i(1)/(1-\rho)$, for each driver in Madrid used in estimation and in the welfare calculations of Section 5. The assumption $E_{G,i} \geq E_{\not G,i} \geq 0$ holds for every driver. Figure F.1 plots the distribution of the ceiling $\alpha_i \, \Delta E_i(1)/(1-\rho)$, weighted by $\tilde{w}_i$. The condition holds wherever the ceiling is below one, and it fails only for drivers in the extreme upper tail of the taste shock ν_i. Because ν_i raises both α_i and ϕ_i, these drivers are so price sensitive that a util of disutility is worth little money to them. Their combined weight, as a share of potential buyers, is 7×10^{-21}.

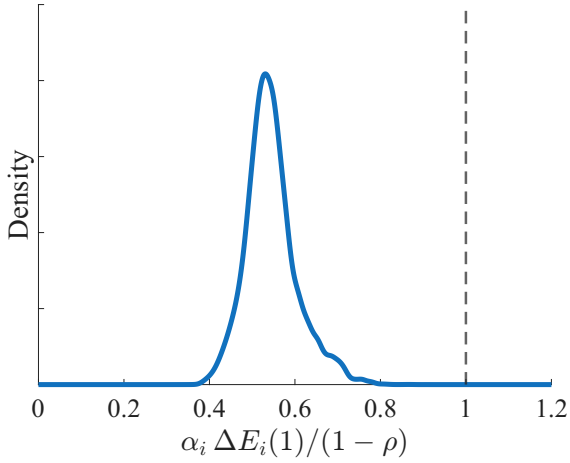


Figure F.1: Ceiling on emission savings per euro of private cost

Notes: Density of $\alpha_i \, \Delta E_i(1)/(1-\rho)$, the most emission savings, in euros, that a euro of driver i's private cost can buy at any intensity, across drivers in the Madrid markets after the policy, weighted by $\tilde{w}_i = w_i \, \phi_i \, S_{G,i}$ at the current intensity. The condition $1/\alpha_i \geq \Delta E_i(1)/(1-\rho)$ holds to the left of the dashed line at one euro. Gaussian kernel, with Silverman's bandwidth computed from the weighted distribution.

A negligible weight at the current intensity does not settle the sign at every intensity on its own. As λ grows, the restricted demand that remains comes increasingly from drivers with low ϕ_i, so exceptions with low ϕ_i could eventually dominate. The exceptions here have high ϕ_i, so their restricted demand vanishes rapidly as the LEZ tightens. The next proposition makes this precise.

Proposition F.2 *Let $h_i \equiv \Delta E_i(1)/(1-\rho) - 1/\alpha_i$, which is positive exactly for the drivers who violate the condition of Proposition F.1. Let $\phi_* \equiv \inf\{\phi_i : h_i > 0\}$ be the lowest LEZ disutility among these drivers. With $\bar{S}_{G,i} \equiv \Lambda_i^G(0)/\left[(\Lambda_i^{\not G})^\rho + \Lambda_i^{\not G}\right]$, which bounds driver i's*

restricted share at every intensity, define

$$P_+ \equiv \int_{h_i>0} w_i\,\phi_i\,h_i\,\bar{S}_{G,i}\,di, \qquad P_- \equiv \int_{h_i<0,\ \phi_i\le\phi_*} w_i\,\phi_i\,|h_i|\,S_{G,i}(0)\,di.$$

Then $\partial\bar{W}/\partial\lambda \le (P_+ - P_-)\,e^{-\phi_*\lambda/(1-\rho)}$ *for every* $\lambda \ge 0$. *The marginal welfare effect is therefore negative at every* $\lambda \ge 0$ *whenever* $P_+ < P_-$.

P_+ bounds the positive contribution of the drivers who violate the condition. P_- measures the negative contribution of the drivers who satisfy it and whose restricted demand fades no faster than that of any violator.

Proof. Since $S_{G,i} = S_{\text{car},i}S_{G|\text{car},i} = \Lambda_i^G/(\Lambda_i^\rho + \Lambda_i)$, $\Lambda_i^G(\lambda) = \Lambda_i^G(0)\,e^{-\phi_i\lambda/(1-\rho)}$, and Λ_i falls with λ without dropping below $\Lambda_i^{\mathcal{G}}$,

$$S_{G,i}(0)\,e^{-\phi_i\lambda/(1-\rho)} \;\le\; S_{G,i}(\lambda) \;\le\; \bar{S}_{G,i}\,e^{-\phi_i\lambda/(1-\rho)}. \tag{F.2}$$

By Lemma F.1, $\partial\bar{W}/\partial\lambda \le \int w_i\,\phi_i\,S_{G,i}\,h_i\,di$. For drivers with $h_i > 0$, apply the upper bound in (F.2) and $e^{-\phi_i\lambda/(1-\rho)} \le e^{-\phi_*\lambda/(1-\rho)}$. Drop the negative terms with $\phi_i > \phi_*$. For the remaining negative terms, apply the lower bound and $e^{-\phi_i\lambda/(1-\rho)} \ge e^{-\phi_*\lambda/(1-\rho)}$. □

At our estimates, $P_+/P_- \approx 10^{-19}$. Under the numerical integration used here, the marginal welfare effect is therefore negative at every $\lambda \ge 0$, so welfare decreases monotonically in the LEZ's intensity.