Income and Energy Tax Pass-through: Evidence from Gas Stations

Samuel Stolper

August 2024

Abstract

The distributional impact of energy taxation depends on not just relative consumption but also relative price changes. I illustrate this by estimating the pass-through of tax changes in Spain's retail automotive fuel market, using a new difference-in-differences method (de Chaisemartin and D'Haultfoeuille, 2024a) that avoids key issues with two-way fixed effects models. I find that gas stations passed through regional diesel tax increases from 2010-2013 at a rate of roughly 91 percent overall, and at significantly higher rates in municipalities with higher median income. Across a range of specifications, pass-through is 4-7 percentage points higher in the top quintile than in the bottom, and at its lowest among stations in the smallest municipalities which lack income data. Given the empirical relationship of fuel consumption and wealth in Spain, the sign and magnitude of this pass-through/income correlation are consistent with a modest offsetting of overall regressive consumer surplus impacts of the tax.

Keywords: pass-through; distributional equity; fuel tax; energy policy
JEL Codes: Q48; H22; D30; C33

Introduction

Taxation is a fundamental instrument of energy and environmental policy, used widely in transportation, electricity, and other emissions-intensive markets. Economists tend to support such taxes due to the gains in total social welfare from "pricing in" externalities. The *distribution* of welfare impacts of taxation, however, is not guaranteed to be progressive. The usual first step in assessing the fairness of energy taxation (or that of any other good) is to examine how household expenditure on the taxed form(s) of energy, as a proportion of one's total expenditure, varies across the wealth (i.e., income, or better yet total expenditure; Poterba, 1991) distribution. Gross of any redistribution of tax revenues, energy taxes are then judged as regressive when the poor devote a larger portion of their total expenditure to the taxed energy relative to to the rich, and progressive when the poor devote a *smaller* portion. Empirically, this relationship varies across countries and types of energy: for instance, globally, electricity taxes appear uniformly regressive, while transportation fuel taxes sometimes are progressive (Flues, F. and Thomas, A., 2015; Pizer and Sexton, 2020).

This method for judging fairness relies primarily on "relative quantities" – that is, how much of the taxed good in question is consumed by the rich relative to the poor. Yet it is not just relative quantities that dictate regressivity; it is also relative price impacts. The first-order approximation of a tax's effect on consumer surplus is a function of both consumption *and* price changes. That is, *pass-through* of the tax – from its point of statutory incidence, to prices paid by households – has the potential to be a significant determinant of who bears the burden of energy taxes. Distributional welfare analyses of commodity taxation in general assume that pass-through of taxes is uniform across households (Gruber and Koszegi, 2004; Bento et al., 2009; Horowitz et al., 2017). But recent research has uncovered an abundance of heterogeneity in energy tax pass-through rates; in automotive fuel markets, for example, such rates have been found to vary as a function of market power (Doyle Jr. and Samphantharak, 2008; Genakos and Pagliero, 2022), vertical structure (Bajo-Buenestado and Borrella-Mas, 2022), and supply elasticity (Marion and Muehlegger, 2011; Kilian and Zhou, 2024). Only one existing study of energy tax pass-through focuses on its relationship with wealth (Harju et al., 2022), and it finds higher pass-through rates of a Finnish auto fuel tax hike in zip codes with lower average income.

This paper provides a second data point on the empirical relationship between energy tax pass-through and income, by studying tax changes in Spain's retail automotive fuel market. I collect daily prices of auto diesel – the dominant auto fuel in Spain – at nearly 10,000 gas stations across the country between 2007 and 2013, made available through an informational mandate unveiled in January 2007 by Spain's Ministry of Energy. I combine these data with panel-varying regional fuel taxes and gas station attributes, including municipality-level income data, and use a difference-in-differences (DD) strategy to estimate pass-through nationally and in subgroups of interest.

From an estimation standpoint, the setting is similar to those of other fuel tax pass-through studies (Alm et al., 2009; Marion and Muehlegger, 2011; Kilian and Zhou, 2024) and common in observational economics studies generally (see, e.g., de Chaisemartin and D'Haultfoeuille, 2024a) – the "treatment" is a series of (fourteen) staggered, area-specific tax changes of differing magnitudes. A growing econometrics literature (de Chaisemartin and D'Haultfoeuille, 2020; Sun and Abraham, 2021; Callaway and Sant'Anna, 2021; Goodman-Bacon, 2021; Borusyak et al., 2024) has shown that the conventional identification strategy in this context among economists, two-way fixed effects (TWFE) regression, is poorly suited to the task of treatment effect estimation when treatment effects vary over time and across space. Among several new related methods for credible estimation of treatment effects via "staggered difference-in-differences", the method of de Chaisemartin and D'Haultfoeuille (2024a) (henceforth, dCDH) is uniquely suited to my setting of non-binary treatments. I thus use dCDH as my preferred method of pass-through estimation, while comparing its results to those of the more traditional TWFE. I present event study results as well as point estimates of cumulative average pass-through for different subsamples of gas stations, focusing primarily on income heterogeneity.

In the national event study, price differences between treatment and control group stations are flat for the seven months prior to a tax change, providing support for the parallel trends assumption. In the preferred dCDH specification, I find that national average pass-through of regional diesel tax changes into prices at the pump is approximately 91.4 in the first year following a tax hike. The pass-through rate reaches its rough peak quickly – in the first month – and remains there for at least eighteen months (the maximum analysis window I employ). It is statistically significantly different from fully 100 percent, or "one-for-one", pass-through.

I then re-estimate event studies in subsamples defined by quintile of the municipality median income distribution, as well as in the group of gas stations in municipalities with population below the threshold (5,000) for inclusion in the income data. Prior price differences between treatment and control groups again show no evidence of non-parallel trends. Across a range of specifications, pass-through rates are at their lowest in the bottom income quintile and in the smallest towns lacking income data, while being higher and generally similar across the top four quintiles. The difference in average pass-through rate in the top quintile vs. the bottom is in the range of 4-7 percentage points; in the preferred specification for inference, pass-through averages 87.6 percent in the bottom quintile and 94.6 percent in the top, and the difference in rates is statistically significant at the five percent level. The qualitative result of rising pass-through in income is robust to the choice of whether to weight observations (in the preferred specification, I weight by station-year sales volume), as well as the cross-sectional unit of observation (station, municipality, or region-by-income-quintile).

I use these results and other suggestive analyses to discuss three extensions: implications for distributional welfare assessment; potential causal channels; and non-tax fuel cost pass-through. I show that proportional fuel consumption (per dollar expenditure) drops in household expenditure decile, making fuel taxes broadly regressive; progressive pass-through (i.e., rising in wealth) of the magnitude I find here offsets a modest amount of that regressivity. I also show that station branding patterns, which are themselves strongly predictive of pass-through rates, do not explain away the income/pass-through relationship, nor do other observable station and local attributes in a two way fixed effects setup. Lastly, I show that the income/pass-through relationship is much weaker when it comes to crude oil prices, which is consistent with a literature that finds more muted consumer response to auto fuel tax changes than upstream fuel cost changes (Davis and Kilian, 2014; Li et al., 2014).

This research contributes to a long economics literature related to pass-through (see, e.g., Jenkin, 1872; Weyl and Fabinger, 2013) that spans traditional sub-fields such as public finance (Poterba, 1996), industrial organization (Bonnet et al., 2013), and international economics (Cavallo et al., 2021).

Most specifically, my findings relate to the literature concerned with estimating the pass-through of energy taxes (Marion and Muehlegger, 2011; Fabra and Reguant, 2014), subsidies (Lade and Bushnell, 2022; Muehlegger and Rapson, 2022), and input costs (Borenstein et al., 1997; Ganapati et al., 2020; Muehlegger and Sweeney, 2022). In empirically documenting variation in fuel tax pass-through with income, I am highlighting a largely-ignored dimension of potential heterogeneity that has implications for the fairness of policy. In applying the de Chaisemartin and D'Haultfoeuille (2024a) method of DD estimation, I provide an early example of new techniques for improved panel estimation of treatment effects (alongside, e.g., Rico-Straffon et al., 2023) that are well-suited to pass-through estimation.

My findings also relate to the literature devoted to distributional welfare impacts of policy. Studies which focus on how progressive or regressive a cost change is tend not to allow for heterogeneous markup adjustment by firms, instead choosing a single pass-through rate to apply throughout the analysis. This practice is prevalent in the energy tax literature (Metcalf, 1999; West, 2004; Bento et al., 2009; Mathur and Morris, 2012) but is also employed in studies of the U.S. sales tax (Caspersen and Metcalf, 1994) and cigarette taxes (Gruber and Koszegi, 2004). My results suggest that assessment of pass-through heterogeneity across the wealth spectrum may be a worthwhile exercise before proceeding with a "uniform pass-through" assumption in welfare analysis.

The rest of this paper is organized as follows: Section 1 reviews pass-through theory to illustrate how it is possible for pass-through to vary with income; Section 2 describes the empirical context, Spain's retail auto fuel market, and all data used; Section 3 explains the estimation methods used; Section 4 presents the main results; Section 5 interprets the results with regard to underlying mechanisms, distributional welfare, and existing related findings; Section 6 concludes.

1 Pass-through and welfare

Pass-through is "the diffusion throughout the community of economic changes which primarily affect some particular branch of production or consumption" (Marshall, 1890). Those changes can be costs, such as input costs or taxes, physically imposed on one part of a supply chain and passed through to others; and they can similarly be benefits, such as government subsidies for specific retail investments. Pass-through has broad value across economics because it is both an outcome of interest in its own right and a sort of "sufficient statistic" (Chetty, 2009) for answering questions about market structure and policy impacts (Jaffe and Weyl, 2013; Atkin and Donaldson, 2015; Pless and van Benthem, 2019).

In perfect competition, pass-through is entirely a function of elasticities of supply and demand. Equation 1 provides the mathematical definition:

$$\frac{dp_c}{dc} = \frac{\epsilon_S}{\epsilon_S - \epsilon_D} = \frac{1}{1 - \frac{\epsilon_D}{\epsilon_S}} \tag{1}$$

Pass-through of cost c to retail price p_c rises in the supply elasticity (ϵ_S) and falls in the absolute demand elasticity (ϵ_D). In the polar cases of either perfectly elastic supply ($\epsilon_S \to +\infty$) or perfectly *in*elastic demand ($\epsilon_D \to 0$), pass-through rates are identically 100%.

In *im*perfect competition, pass-through varies with not just the first derivative (elasticity) but also the second (convexity). Consider the formula for pass-through in monopoly with constant marginal costs c:

$$\frac{dp_m}{dc} = \frac{\frac{\partial p(q_m)}{\partial q_m}}{2\frac{\partial p(q_m)}{\partial q_m} + q_m \frac{\partial^2 p(q_m)}{\partial q_m^2}} \tag{2}$$

Equation 2 shows that monopoly pass-through can, in principle, be higher or lower than perfectly competitive pass-through – it depends on the demand convexity parameter $\frac{\partial^2 p(q_m)}{\partial q_m^2}$ (Seade 1985).[1] Under perfect competition, constant marginal cost (i.e., perfectly elastic supply) guarantees fully 100 percent pass-through. Under monopoly with linear demand, $\frac{\partial^2 p(q_m)}{\partial q_m^2} = 0$ and pass-through simplifies to a constant 50 percent. If, however, $\frac{\partial^2 p(q_m)}{\partial q_m^2} > 0$, then market power increases pass-through relative to that of perfect competition. If $\frac{\partial^2 p(q_m)}{\partial q_m^2}$ is positive and sufficiently large, then pass-through can exceed 100 percent.

While the relationship between pass-through and market power has received much attention, the relationship between pass-through and wealth has not. Demand elasticities are a fundamental determinant of pass-through, and it is plausible that they themselves vary with wealth. In retail auto fuel markets, for instance, it could be that demand is more inelastic in poorer areas because the poor have more fundamental, non-discretionary uses of auto fuel; or it could be that the *wealthy* have more inelastic demand because of a wealth-induced price insensitivity. Empirically, retail auto fuel (absolute) demand elasticities have been found to both fall with income (Hughes et al., 2008; Kayser, 2000) and rise with income (Kilian and Zhou, 2024; Houde, 2012). My main findings are consistent with the latter. Ultimately, pass-through may vary with wealth not only through variable demand elasticities but also through correlations of wealth with other determinants of pass-through, such as supply elasticities (Marion and Muehlegger, 2011), market structure (Bajo-Buenestado and Borrella-Mas, 2022), or tax salience (Kroft et al., 2024).

A positive cost shock – such as a tax or input cost increase – elicits a direct change in consumer surplus through two channels: (a) the additional cost of consumption maintained in the face of rising prices; and (b) the utility lost from reduced consumption. Pass-through physically measures the former (per unit consumption), which is the first-order approximation to the consumer surplus impact of a marginal tax change. Total individual welfare is also determined at least by (a) ownership of supply-side capital; (b) externalities; (c) other goods' prices and quantities that are affected in general equilibrium; and (d) the use of government revenues obtained through taxation. In this paper, however, I focus only on the utility derived directly from the purchase and consumption of energy. Sterner (2012) provides a fuller discussion of the various channels through which a tax affects welfare in the context of fuel markets.

[1] Ritz (2024) shows that *cost* convexity can also produce a positive relationship between market power and pass-through, all else equal.

2 Empirical setting and data

The Spanish retail automotive fuel market is oligopolistic, with vertically connected firms.[2] Three companies (Repsol, Cepsa, and BP) own – in my sample window of 2007-2013 and currently – the nine oil refineries producing automotive fuel in Spain,[3]. In my sample window, these companies owned a significant (29.5 percent) share of Compañía Logística de Hidrocarburos ("Hydrocarbon Logistics Company"), the regulated monopoly responsible for distribution of oil products in Spain (Comisión Nacional de la Competencia, 2012). Most importantly for the present study, Spain's refiners are very active in the retail market. In my main analysis sample, 67 percent of retail gas stations offer gasoline bearing the brand of a refiner, and 30 percent of stations are fully vertically integrated with a refiner. These companies have faced significant scrutiny from government and popular media alike, on the grounds of alleged collusion and some of the highest estimated retail margins nationally in all of Europe (see, for example, Noceda and Jimenez, 2015).

2.1 Price and other station attributes

One result of such scrutiny has been closer monitoring of pricing by gas stations. A government mandate which went into effect in January 2007 requires all stations across the country (more than 12,000 today; Comisión Nacional de los Mercados y la Competencia, 2024) to send in their fuel prices to the Ministry of Energy whenever they change, and weekly regardless of any changes. These prices are then posted by the Ministry to a web page - called the *Geoportal* - designed for consumer use.[4] I obtain daily price data for retail diesel at 9,884 Spanish gas stations from January 2007 to June 2013 (diesel has an 80% share of all retail automotive fuel sales in Spain during this time period, according to data from the Spanish entity strategic petroleum reserve [CORES]). The price data come with cross-sectional attributes: location, brand, wholesale contract type, and amenities.[5] Figure 1 maps all Spanish gas stations, excepting those in the Canary Islands and the (African contintental) territories of Ceuta and Melilla, which are each omitted from the main analysis sample because they are exempt from the fuel tax that I study here.

I add two further types of data to the station-level dataset recorded in Spain's Geoportal. First, I obtain station-year sales volumes from 2008 to 2014, also from the Ministry of Energy. Station-level sales volumes are useful because I am primarily interested in estimating population-average pass-through rates, not just unweighted averages across stations. The merge of price and quantity data is imperfect, however, resulting in 89 percent of stations in the price data being successfully matched to annual sales at least once.[6]

[2]For background on the evolution of Spain's oil markets, see Contín-Pilart et al. (2009); Perdiguero and Borrell (2007).

[3]Spain both imports and exports refined petroleum products. In 2011, net imports accounted for 15 percent of refined diesel, but imports have been deemed ineffective at exerting competitive pressure on the refining companies (Comisión Nacional de la Competencia, 2012).

[4]<https://geoportalgasolineras.es/geoportal-instalaciones/Inicio>; there is now a phone application as well. Appendix Figure A1 is a screenshot of the *Geoportal* from 2015.

[5]Station-specific *changes* in brand and contract type over time are thus unobservable. Spain's Comisión Nacional de la Competencia (2012) reports that "the relative positions of the three majors (Repsol, Cepsa, and BP) by number of service stations has remained more or less stable over time." Sales of retail stations over time have overwhelmingly been exchanges among wholesale operators without refining capacity in Spain, and my parameterization of station brand (into three groupings) in analysis is unaffected by such sales.

[6]Five percent of stations in the *sales* dataset (representing three percent of total recorded sales) cannot be matched to the

Figure 1: *Map of Spanish gas stations, 2007-2013*

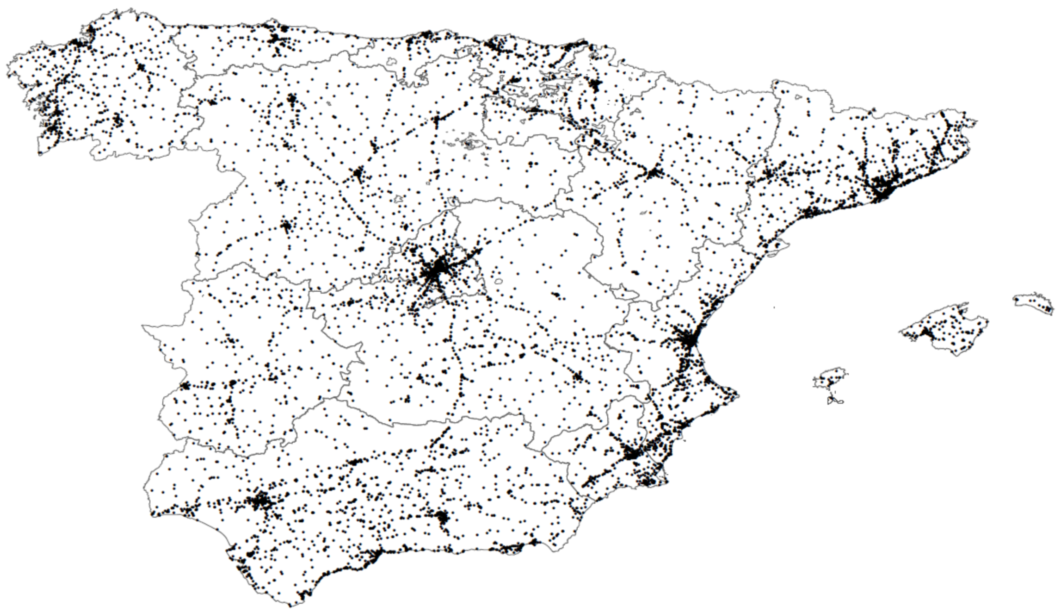

Notes: Data are from Spain's Ministry of Energy. The Canary Islands and Ceuta and Melilla territories are omitted from the map (and main analysis sample).

Second, I calculate two station-specific proxies for market power. One of these is a count of open stations within a 10-minute distance radius, weighted by $\frac{1}{1+d}$, where d is the travel distance (in minutes) between a pair of stations. This count captures degree of spatial isolation from competitors.[7] The second is the proportion of stations within a 10-minute radius that share one's brand; this measure captures the degree of own-brand concentration in local markets. Both of these competition proxies vary cross-sectionally and over time due to entry and exit of stations.

2.2 Taxes

Fuel taxes are large as a proportion of total price in Spain. The unweighted average final retail diesel price in my sample is 1.15 Euros/liter; on average, taxes are 47 percent of final price. There are three taxes applicable to retail diesel in Spain during the sample time period: a national sales tax that rises twice, going from 16 to 18 and then 21 percent; a national excise tax ("Impuesto Especial de Hidrocarburos", or IEH) that changes once, from 27.8 to 30.7 Eurocents per liter (c/l); and a smaller excise tax ("Impuesto sobre las Ventas Minoristas de Determinados Hidrocarburos", or IVMDH) with

price dataset.

[7]A station's relevant competitors are defined in part by typical travel patterns, including commuting (Houde, 2012). Lacking commuting data for the whole of Spain, I use the simpler, commonly used travel time based measure of competition (e.g., Doyle Jr. and Samphantharak, 2008; Miller et al., 2017). Genakos and Pagliero (2022) show, using fuel tax and price data from Greek Islands, that distance-restrictions on local retail fuel market definitions can lead to overestimation of pass-through when the number of nearby stations is small. In my context, however, the logic ends up working the other way, because I find pass-through rising in market power, rather than falling as in Genakos and Pagliero (2022).

both a national component (2.4 c/l) and an optional regional component (Agencia Tributaria, 2024).[8] I rely on the regional component of this last tax for pass-through identification, because it is the only source of panel variation in tax levels. It is colloquially known as the "céntimo sanitario" (" health cent"), because of its modest size – 0 to 4.8 c/l – and the fact that its revenues were earmarked for public health expenditures (Court of Justice of the European Union, 2014).[9]

At the first sub-national level of government, Spain has 17 autonomous communities; these are the "regions" to which I refer throughout the paper. Between 2008 and mid 2013, there are 14 discrete regional tax changes in 11 unique regions (three regions change their level twice), all of them increases.[10] Figure 2 documents the timeline of these tax changes, with average retail and crude oil price trends (Brent benchmark) provided in the background for context. The first tax hike occurs in early 2010, and each month thereafter has 0-4 additional, region-specific tax hikes. The increases are modest; the average size of the 14 tax changes is 2.97 c/L, which is about 2.6 percent of after-tax retail prices.

Figure 2: *Regional tax changes and prices over time*

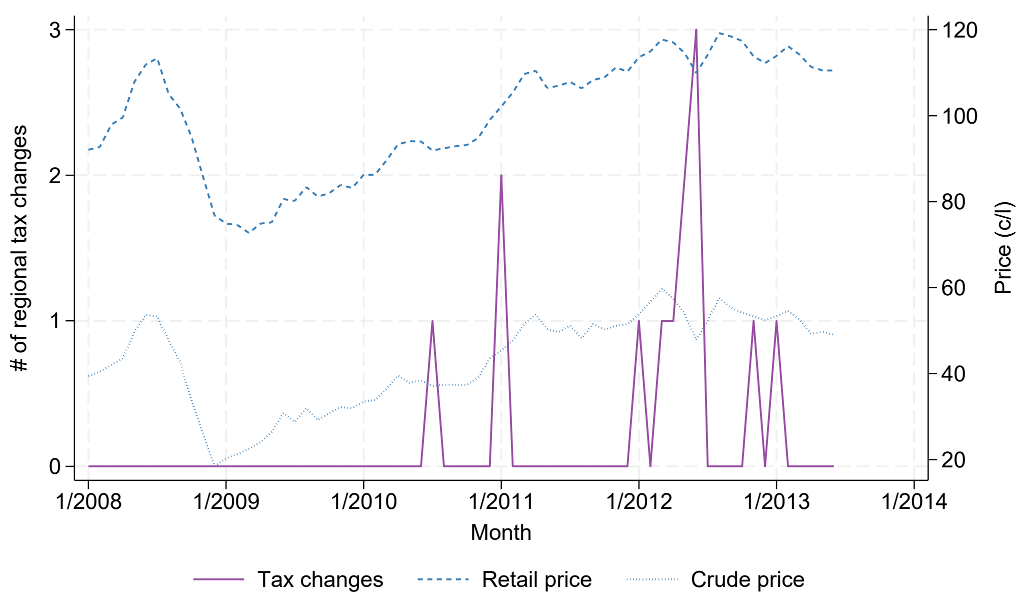

Notes: Retail prices are monthly national averages, taken from daily national averages weighted by station-level quantity sold.

[8]The sales tax is applied on top of the excise taxes. In all analyses, I remove the effect of the sales tax from the final retail price and use this "pre sales tax" price as the outcome of interest.

[9]In 2013, the IVMDH was integrated into the IEH to comply with European Union Law. In 2014, the European Court of Justice ruled that, from 2002 through 2011, the tax was unconstitutional and its revenues must be returned (Court of Justice of the European Union, 2014).

[10]The regional tax changes apply identically to diesel and gasoline, the main alternative auto fuel to diesel.

2.3 Income and population

I combine the station (price) and tax datasets with income data from Spain's Foundation for Applied Economic Studies (FEDEA). I collect per capita income and median income per declarant by municipality-year, from 2008-2013, for all municipalities with population greater than 5,000.[11] I generate municipality means (across years) of the two income variables and assign them to quintiles. I prefer the median income measure in my context because it is representative of *more* households (in a given municipality) experience than the per capita measure (which is influenced by a long right-tail in incomes). I use the latter measure to define quintiles in an Appendix analysis.

Importantly, I retain the large subset of stations in municipalities too small to be included in the income data as a data point in the assessment of pass-through heterogeneity. This subset is, by population, more rural than the five groups based on income quintile. While incomes in this subset are not measured, inspecting the relationship between income and population near the threshold of 5,000 is suggestive. Figure 3 allows for this inspection; I plot a binscatter of municipality median income and population (which is available for towns with more than 1,000 residents from Spain's National Statistical Institute). There is a clear positive relationship, with the appearance of concavity near the threshold, such that income is dropping fast in population nearest the region of interest (<5,000 population). It is certainly possible for the trend to break, but this figure suggests it is unlikely that the "no income" subgroup has incomes *higher* than the bottom income quintile.

2.4 Main analysis sample

We start with 9,984 unique stations observed in the raw Geoportal data. We drop stations in the Canary Islands region and two small territories (Ceuta and Melilla), because these places were not subject to any part of the IVMDH (national or regional).[12] This yields 9,187 stations across 16 regions. My preferred analysis weights station-level prices by station sales in order to estimate population average effects – which restricts the sample to 8,211 stations and an analysis window of 2008-2013.[13] I carry out the main analyses without weights as well, which gives estimates of average pass-through across stations weighted equally regardless of how much they sell.

Table 1 presents summary statistics for the main analysis variables nationally as well as in six income-based groups separately – the five quintiles of municipality median income, and the group of stations without income data. The national average retail price (excluding sales tax, weighted by station sales volume) is 100.23 c/l, and all subgroups have an average price within 0.65 cents of that. 67 percent of stations nationally are "refiner-branded", which is to say, have a contract for branded fuel from a Spanish oil refiner. 13 percent have a contract for branded fuel with a wholesale operator ("wholesaler-branded"). And 20 percent do not have contracts for branded fuel ("independent-branded"). Firm contracts are also distinguished by their degree of vertical integration; 30 percent of stations nationally are both owned and operated by a refiner or wholesaler (often referred to as "company-owned, company-operated").

[11] The FEDEA dataset does not include data from the País Vasco and Navarra regions, because they are not a part of Spain's Common Fiscal Regime Territory. I omit these regions from analysis involving income data.

[12] My findings are robust to the inclusion of these places, with a tax level of zero, in the analysis sample. These results are available upon request.

[13] Appendix Figure A2 shows how the station income distribution differs between the full sample and the sample without.

Figure 3: *Population and median income across municipalities*

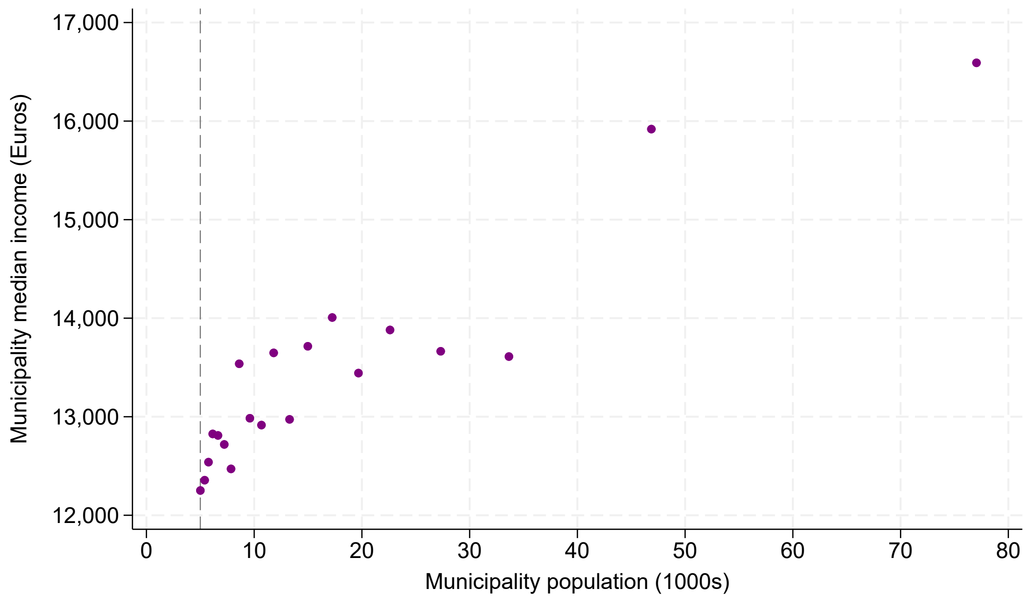

Notes: Data points are from a binscatter of municipal average population and median income; observations are grouped into twenty equal-sized population bins, and then averages of both population and income are calculated within bins. The graph is truncated above at a population of 100,000. The dashed line indicates a population of 5,000, the cutoff for being included in the FEDEA income dataset.

Higher quintile stations tend to have more nearby rivals, more vertically integrated contracts, and (much) larger municipal populations. The group of stations missing income data, whose incomes are likely quite low on the spectrum, has more independent-branded stations, fewer refiner-branded ones, and a higher concentration of its own brand locally than any of the quintile-specific groups. Of course, I am able to control for cross-sectional differences across stations in a difference-in-differences (DD) analysis; it is time *trends* in the treatment group as compared to the control, pre-dating the tax-change treatments, that are the main challenge to DD identification.

3 Methods

3.1 Difference in differences with staggered treatment

Difference-in-differences – comparing "treated" units to untreated ones, before versus after the treatment is applied – is a common method of causal estimation across the social sciences (e.g., Roth et al., 2023; Wing et al., 2024). In the frequent case that the treatment of interest is staggered in time across many units being studied, two way fixed effects (TWFE) regression is a common strategy for specifying the DD comparison. The TWFE model is as follows in my context, which involves a non-binary treatment:

Table 1: *Summary statistics*

	All	No Inc	Q1	Q2	Q3	Q4	Q5
Retail diesel price (c/L)	100.23	99.58	100.07	100.21	100.36	100.49	100.81
	(3.00)	(3.15)	(2.85)	(2.80)	(2.86)	(3.00)	(2.85)
Community-level retail tax (c/L)	1.76	1.36	1.75	1.87	1.80	1.90	2.08
	(0.96)	(1.16)	(0.58)	(0.66)	(0.83)	(0.80)	(0.87)
Refiner-branded (0/1)	0.67	0.63	0.71	0.70	0.67	0.69	0.70
	(0.47)	(0.48)	(0.45)	(0.46)	(0.47)	(0.46)	(0.46)
Wholesaler-branded (0/1)	0.13	0.13	0.11	0.11	0.11	0.12	0.15
	(0.33)	(0.34)	(0.31)	(0.31)	(0.32)	(0.33)	(0.35)
Independent-branded (0/1)	0.20	0.24	0.18	0.19	0.22	0.19	0.16
	(0.40)	(0.43)	(0.38)	(0.40)	(0.41)	(0.40)	(0.37)
# rivals, distance-weighted	1.52	1.01	0.71	1.05	1.40	1.92	2.28
	(1.35)	(1.07)	(0.75)	(1.00)	(1.23)	(1.33)	(1.49)
Proportion of rivals sharing one's brand	0.27	0.33	0.25	0.26	0.24	0.25	0.24
	(0.30)	(0.37)	(0.34)	(0.29)	(0.29)	(0.25)	(0.22)
Vertically integrated (0/1)	0.30	0.31	0.16	0.21	0.25	0.31	0.38
	(0.46)	(0.46)	(0.37)	(0.41)	(0.44)	(0.46)	(0.48)
Vertically independent (0/1)	0.19	0.23	0.19	0.19	0.21	0.18	0.15
	(0.39)	(0.42)	(0.39)	(0.39)	(0.40)	(0.38)	(0.36)
Commission contract (0/1)	0.29	0.28	0.40	0.36	0.28	0.29	0.24
	(0.45)	(0.45)	(0.49)	(0.48)	(0.45)	(0.45)	(0.43)
Firm contract (0/1)	0.19	0.15	0.22	0.21	0.23	0.18	0.20
	(0.39)	(0.36)	(0.42)	(0.41)	(0.42)	(0.38)	(0.40)
Municipality population (1000s)	186.19	13.28	17.59	32.17	32.58	142.56	600.75
	(579.75)	(42.46)	(17.90)	(44.06)	(30.70)	(153.79)	(1,053.78)
Muni. pop. density (1,000s per km^2)	1.29	0.38	0.17	0.31	0.89	1.68	2.94
	(2.55)	(1.26)	(0.23)	(0.67)	(1.73)	(2.91)	(3.44)
Muni. median income (Euros)	15167.71	.	9,385.23	11331.31	13224.41	15478.68	18719.41
	(3,591.57)	(.)	(834.93)	(535.24)	(529.41)	(636.58)	(2,827.13)
Observations	8,211	2,755	615	942	850	1,395	1,654

Notes: Means and standard deviations (in parentheses) are calculated from station-level observations weighted by station sales volume. "All" refers to all stations in the national sample. "No Inc" indicates stations in municipalities too small to have income data reported. *Q*1 through *Q*5 denote (municipal) median income quintiles. All stations are either refiner branded, wholesaler branded, or independent branded. # of rivals is weighted by inverse travel time in minutes between station pairs.

$$P_{it} = \alpha + \beta \text{Tax}_{it} + \theta_t + \phi_i + \epsilon_{it}. \tag{3}$$

P_{it} is the retail price of auto fuel and Tax_{it} the region-specific tax level, both at station i in year-month t. θ_i and ϕ_t are station and year-month fixed effects, and ϵ_{it} is a mean-zero error term. The coefficient β is interpreted as the average pass-through rate over the full post-period after a tax change. Analogously, one might estimate an event study model, where Tax_{it} is replaced by a vector of binary variables (D^j_{it}) equaling one if an observation is j months from a tax change in its region, and relative-month specific coefficients β_j are estimated instead of β:

$$P_{it} = \alpha + \Sigma^b_{j=a} \beta^j D^j_{it} + \theta_t + \phi_i + \epsilon_{it}. \tag{4}$$

Identification of average treatment effects via DD requires that two key assumptions hold; first, that a group's current outcome does not depend on its future treatments ("no anticipation"); and second, that in the absence of treatment, outcomes among treated and untreated groups would have evolved in parallel ("parallel trends") (Roth et al., 2023).[14] The credibility of these assumptions can be examined visually with an event study plot showing pre-treatment data points. There is no uniformly applied statistical criterion for validating the parallel trends assumption, but the individual statistical significance of the pre-treatment coefficients is common. Roth (2022), however, shows how typical statistical tests of the parallel trends assumption tend to be under-powered.

While non-parallel trends (in the treated group versus the control, predating treatment) are widely considered the main threat to credible estimation of treatment effects in the DD setup, a new and rapidly evolving literature has highlighted another significant threat to DD identification using TWFE: heterogeneous treatment effects. Consider that, in the case of binary and staggered treatment, the TWFE estimator is mathematically a weighted average of DD treatment effects across all pairs of units with one unit switching treatment (and the other not), for every observed pre-versus-post time period comparison (Goodman-Bacon, 2021). This includes, for example, comparing a group A that switches from untreated in period 1 to treated in period 2 and a group B that is treated in both periods. If treatment effects are constant, then taking the pre-post difference in group B removes the effect of group B's treatment and makes this DD comparison valid. But if treatment effects evolve over time, then group B's period specific treatment effects do not cancel each other out in the pre-post difference, and one treatment effect enters the TWFE average treatment effect calculation with negative weight. These negative weights result in the TWFE estimator not being a convex combination of the individual DD comparisons and not having the "no sign reversal" property (it is possible for TWFE to produce an estimate of the wrong sign relative to the true value; de Chaisemartin and D'Haultfoeuille, 2023).

A second problematic type of comparison is made when the treatment is additionally non-binary (as is common in published papers using DD; de Chaisemartin and D'Haultfoeuille, 2023): a group A as above, becoming newly treated in period 2, and a group B that is also treated in period 2 but with a smaller treatment *size*. This comparison's validity requires constant treatment effects not just over time, as above, but also across groups. If the latter is not the case, then comparing a more-treated group to a less-treated one conflates treatment size with treatment effect.

Several recent papers – including Callaway and Sant'Anna (2021); Sun and Abraham (2021); Borusyak et al. (2024); de Chaisemartin and D'Haultfoeuille (2024a) – have proposed (and provided code for) alternative DD estimators that are robust to heterogeneous treatment effects under certain conditions, by virtue of avoiding the "forbidden comparisons" described above. I rely here on the methods of de Chaisemartin and D'Haultfoeuille (2024a) to estimate sample-wide and heterogeneous (sub-sample) treatment effects. The treatment in question – a fuel tax increase – is staggered and non-binary; de Chaisemartin and D'Haultfoeuille (2024a) are the only researchers to propose estimators for a non-binary treatment. Their estimators, which I refer to as dCDH (computed in Stata with the -did_multiplegt_dyn- package), compare treated groups with as-yet-untreated groups, and they allow

[14]The DD also requires the stable unit treatment value assumption (SUTVA), which rules out spillover and general equilibrium effects (Roth et al., 2023). In my context, one station's treatment does indeed affect other, competing stations' outcomes. But region-level treatment assignment ensures that stations in competition with each other overwhelmingly receive the same treatment – the exception is with cross (regional) border competition.

for dynamic (that is, lagged) treatment effects.

3.2 Implementation

Following de Chaisemartin and D'Haultfoeuille (2024a), let $Y_{g,t}$ be the retail fuel price at station g in year-month t.[15] F_g is the month in which station g's tax changes. Station g's comparison group is all stations g' having (a) the same initial tax level and (b) not yet been treated as of month F_g. A DD estimate of station g's treatment effect in post-treatment month ℓ is:

$$\mathrm{DD}_{g,\ell} = \left(Y_{g,F_g-1+\ell} - Y_{g,F_g-1}\right) - \sum_{g':F_{g',1}=F_{g,1},F_{g'}>F_g-1+\ell} \left(\frac{W_{g',F_g-1+\ell}}{N^g_{F_g-1+\ell}}\right)\left(Y_{g',F_{g'}-1+\ell} - Y_{g',F_{g'}-1}\right) \quad (5)$$

The baseline, or pre-, period is fixed here to the month immediately preceding group g's tax change; other pre-periods are reserved for "placebo tests" of the parallel trends assumption. $N^g_{F_g-1+\ell}$ is the number of observed groups in g' in a given post-treatment month, and $W_{g',F_g-1+\ell}$ is a weight equal to annual quantity sold at station g' and month $F_g-1+\ell$ (de Chaisemartin et al., 2024). Thus, control group pre-post differences are aggregated into a weighted average before being subtracted from the treatment group pre-post difference. Month-specific event study effects DD_ℓ can then be computed as the average of $DD_{g,\ell}$ across all groups g,

$$\mathrm{DD}_\ell = \sum_{g:F_g-1+\ell\leq T_g} \frac{W_{g,F_g-1+\ell}}{N_\ell} DD_{g,\ell}, \quad (6)$$

where the weight W used is now quantity sold by *treated* station g. Equation 6 is unbiased conditional on the parallel trends and no-anticipation assumptions (de Chaisemartin and D'Haultfoeuille, 2024a).

The -did_multiple_gt- command also estimates an average total effect per unit of treatment, which conveniently translates exactly into an average pass-through rate in my context. Let T_g be the last observed month for which there is an as-yet-untreated station remaining. Furthermore, define $D_{g,\ell}$ as the non-binary, discrete size of the tax faced by station g in post-month ℓ. Then, pass-through $\rho = \frac{\Delta P}{\Delta t}$ can be estimated with

$$\hat{\rho} = \frac{\sum_{\delta=1}^{T_g-F_g+1} w_\ell \mathrm{DD}_\ell}{\sum_{\delta=1}^{T_g-F_g+1} w_\ell \sum_{g:F_g-1+\ell\leq T_g} \frac{W_{g,F_g-1+\ell}}{N_\ell}\left(D_{g,F_g-1+\ell} - D_{g,1}\right)}, \quad (7)$$

where weight w_ℓ equals the proportion of all observed treated station-months that fall in month ℓ. I estimate Equations 5-7 in the full sample of stations, as well as separately for subsamples corresponding to quintile of municipality median income (and the subsample of stations missing income). In most of my analyses, I estimate event study effects for each $\ell \in [1, 12]$ and, correspondingly, average pass-through rates through one year. To assess the appropriateness of the parallel trends and no-anticipation assumptions, I also estimate six pre-treatment placebo effects according to Equations 5 and 6, where ℓ here takes a value in [-6,-1]. I report on the individual statistical significance of these

[15] I also conduct the analysis using municipality-month observations.

effects, as well as their average and the slope of a line fitted by Ordinary Least Squares through them.

As noted earlier, my preferred specification of Equation 5 weights observations by station-year sales volume. However, I include results from unweighted analysis as well. I primarily cluster standard errors by province; while treatment is determined by region, which is one administrative level above province, there are only 17 regions in Spain – 14 of which get used in the main, income-focused analysis – whereas there are 55 provinces. Using the latter as cluster level avoids the "too-few clusters" problem that leads to downward bias in standard errors (Cameron and Gilbert, 2015). I show results from the use of region-level clustering in the Appendix.

4 Results

4.1 Overall fuel tax pass-through

The event study in Figure 4 provides an overall picture of relative price trends among the treatment and control groups defined by dCDH. The point estimates are obtained via Equations 5 and 6. Since Equation 5 uses month 0 (that is, the month before a station's tax changes) as the pre-treatment data point in the DD design, the treatment/control price difference in month 0 is normalized to zero. The levels and trend of the estimates for months -6 through 0 provide evidence on the parallel trends assumption.

Figure 4: *Full-sample average price impacts over time*

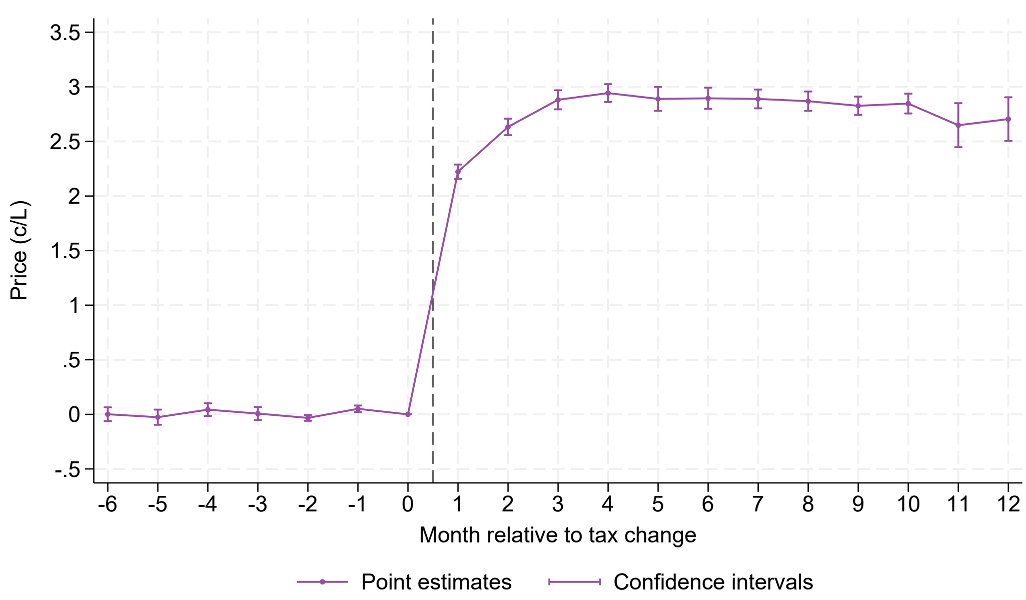

Notes: Point estimates and confidence intervals are for monthly DD coefficients (DD_ℓ) estimated using Equations 5 and 6, where $\ell \in [-6, 12]$. Observations are station-year-months and are weighted by station-year diesel quantity sold. Standard errors are clustered at the the province level.

The pre-treatment trend in Figure 4 appears roughly flat. Estimates jump significantly in the month of the tax change and remain elevated for the duration of the twelve-month analysis window used. The post-treatment price impact hovers between 2.5 and 2.9 cents per liter (c/l), as compared to an average tax change of 2.97.[16] Average standard errors are small – on the order of hundredths of a c/l – and as a result, two of the six pre-treatment placebo effects are significantly different from zero at the five percent level. However, the average pre-treatment point estimate is only 0.006 and is not significantly different from zero (p=0.73). More importantly, a straight line fitted by OLS through point estimates for periods -6 to 0 is also insignificant, with a slope of 0.003 and a p-value of 0.695. Put together, the evidence does not suggest a violation of non-parallel trends in the overall station sample.

Table 2 presents points estimates of and standard errors for the cumulative average pass-through rate in the full sample, out to the end of the first year after a tax change. Pass-through in the preferred, weighted specifications of Panel A is in the range of 91-92 percent on average and significantly different from 100 percent, or fully one-for-one, pass-through. Consistent with Figure 4, columns 1-3 show an insignificant, one percentage point drop in cumulative pass-through rate from month 6 to month 18 (column 2 gives the estimate corresponding exactly to the specification and results of Figure 4). Limiting the sample to those stations used in the income-related analysis (column 4) does not change the estimate meaningfully; nor does using two way fixed effects (column 5). Unweighted pass-through estimates, in Panel B, are smaller than the weighted ones in all dCDH specifications, settling at in the range of 89-90 percent (though 6-month cumulative pass-through, in column 1, is only 86.8 percent). Without weights, the TWFE result deviates further from the dCDH result; the difference is about four percentage points (column 5 minus column 2).

4.2 Fuel tax pass-through by income group

Figure 5 presents event studies identical to that of Figure 4 except now separately for each of six income groups: the five quintiles of municipality median income, and the set of stations whose towns are too small for inclusion in the income data. Because of the number of event studies depicted here, I omit confidence intervals from the figure. Income-group specific pre-trends are visually hard to distinguish from both the zero line and each other. In practice, a quarter of the 36 pre-treatment point estimates are statistically significant, but none of the six income groups have an average pre-treatment estimate different from zero (the lowest p-value is 0.14) and none have a linear trend coefficient different from zero (the lowest p-value is 0.25).

Once treatment occurs, all income groups show a mean shift in prices. The size of the mean shift varies from roughly 2-3.5 c/l; however, the ordering of income groups vertically cannot be interpreted as reflective of relative pass-through rates, because the *sizes* of tax changes experienced are not uniform across income group. The absolute-price impacts shown in Figure 5 must be adjusted to account for tax-change magnitudes; Figure 6 does exactly that. I calculate month-specific pass-through rates analogously to Equation 7 (and following de Chaisemartin et al., 2024), dividing each monthly estimated price impact by a weighted average of tax-change sizes in that period.

Panel A of Figure 6 shows estimated pass-through over time for each of the six income groups.

[16]The month-1 coefficient is mechanically slightly lower than later-month coefficients, because tax changes happen at various points within a month, and so the average tax level in that month is in *between* the pre- and post-treatment tax levels.

Table 2: *Full-sample average tax pass-through*

	dCDH				TWFE
	(1)	(2)	(3)	(4)	(5)
Panel A. Sales-weighted					
Average pass-through	0.923	0.914	0.913	0.924	0.926
	(0.014)	(0.013)	(0.013)	(0.014)	(0.030)
N	29,230	54,903	76,514	52,291	461,976
Post-period	6	12	18	12	12
Sample	Full	Full	Full	Income	Full
Panel B. Unweighted					
Average pass-through	0.868	0.891	0.896	0.902	0.935
	(0.006)	(0.007)	(0.009)	(0.007)	(0.025)
N	31,217	56,890	78,501	54,278	530,082
Post-period	6	12	18	12	12
Sample	Full	Full	Full	Income	Full

Notes: Estimates are average pass-through rates of tax changes into retail prices, $\frac{\Delta P}{\Delta t}$, obtained via Equation 7) in columns 1-4 ('dCDH') and via Equation 3 in column 5 ('TWFE'). An observation is a station-month. 'N' reports the number of treatment-group station-months used in columns 1-4 and the number of total station-month observations used in column 5. The two Panels (A and B) show results with and without analytic weights, respectively. 'Post-period" gives the number of months included in estimation. "Sample" is either "Full" – all stations regardless of income data availability – or "Income" – all stations with matched income data. Standard errors are clustered by province.

Higher income quintiles experienced larger absolute price impacts in part because they experienced larger tax changes, but there remains evidence of pass-through heterogeneity: the range of income group specific rates is consistently 10-20 percentage points wide across months. The lowest income quintile (Q1) and the no-data group (which, I argue, likely has even lower income than Q1), tend to have lower month-specific pass-through estimates than the other four quintiles. Panel B zeros in on quintiles 1 and 5 specifically, to provide a clearer picture of how different pass-through rates are at the bottom and top of the income spectrum. After treatment, Q1 estimates are consistently *lower* than those of Q5 – by as much as ten percentage points. Confidence intervals in Panel B shed some qualitative light on the significance of differences in pass-through rate between Q1 and Q5. Because I cluster standard errors at the province level here rather than the station level, I cannot formally test whether Q1 and Q5 pass-through rates are significantly different without bootstrapping (de Chaisemartin and D'Haultfoeuille, 2024b). However, I *can* do so with a slightly different version of the analysis that uses municipality-month observations and municipality clustering of standard errors. Table 3, which presents full-period (twelve-month) pass-through rates, includes results from

Figure 5: *Income and* absolute *price impacts*

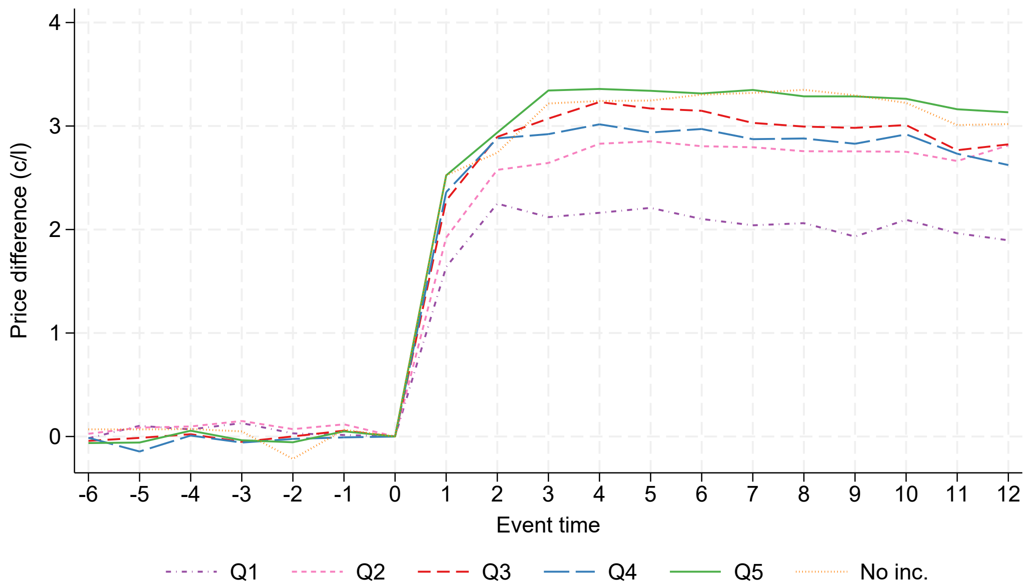

Notes: Lines are constructed from monthly DD point estimates of average treatment effect (DD_ℓ) by income group, estimated using Equation 6, where $\ell \in [-6, 12]$. Observations are station-year-months and are weighted by station-year diesel quantity sold. Q1-Q5 indicate quintiles of the municipality median income distribution; "No Inc" denotes the group of stations in towns too small to be included in the data.

this municipality-level setup.

Estimates in Table 3 are obtained via Equations 5-7, using each income group separately. Panel A uses the primary, station-month observation level, while Panel B uses the municipality-month level that facilitates hypothesis testing about heterogeneity. In both cases, the lowest pass-through rate – 86.5 and 86.8 percent, respectively – is found in the subset of stations in small towns missing income data. The next lowest pass-through rate is in Q1 and is also below 90 percent (89.3 and 87.6 in the two panels). The remaining four income quintiles have pass-through rates between 93 and 96 percent, though the rate does not rise monotonically with quintile. Following the instruction of de Chaisemartin and D'Haultfoeuille (2024b) and using the results of Panel B, I calculate the variance of the difference in two quintile-specific pass-through rates by summing the estimated variances of the two pass-through rates themselves. I am able to reject the null of equal pass-through rates when comparing either of Q1 and the no-income group to any of Q2-Q5 (p-values are less than 0.05).

Table 4 summarizes results from other specifications of pass-through rate estimation by income group. Panel A restricts the dCDH analysis to using a balanced panel of stations that are observed for all twelve post-treatment months; Panel B omits the station sales weights from the analysis; and Panel C uses TWFE.[17] The relative pattern across income groups is consistent across all panels of Table 4 and both panels of Table 3. Pass-through is lowest among the lowest-income groups, jumping from Q1

[17]Appendix Figures A3 and A4 present event studies corresponding to Panels A and B. Appendix Table A1 presents results of weighted and unweighted dCDH analysis with *region*-level clustering of standard errors.

Figure 6: *Income and pass-through*

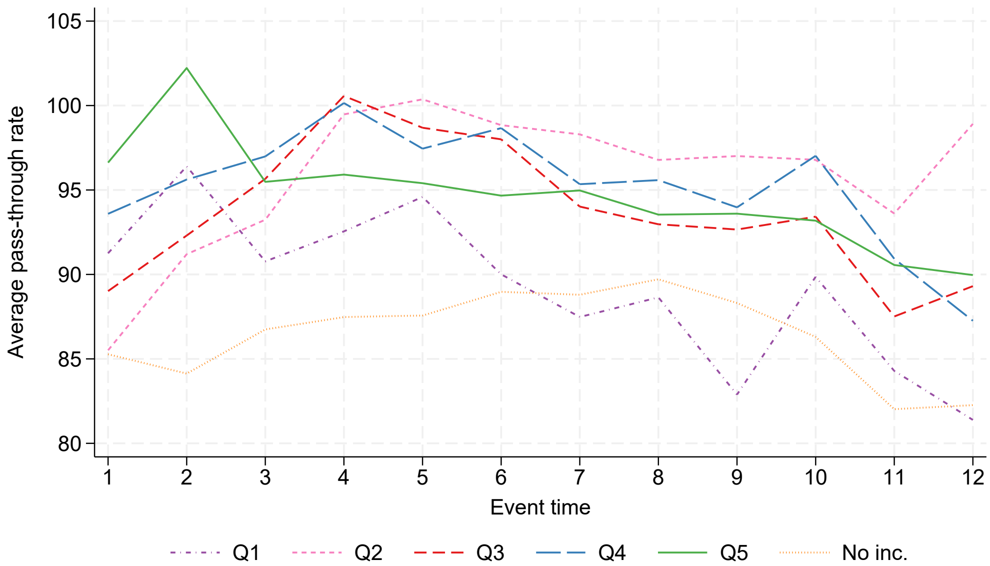

(a) *Avg. pass-through by income quintile*

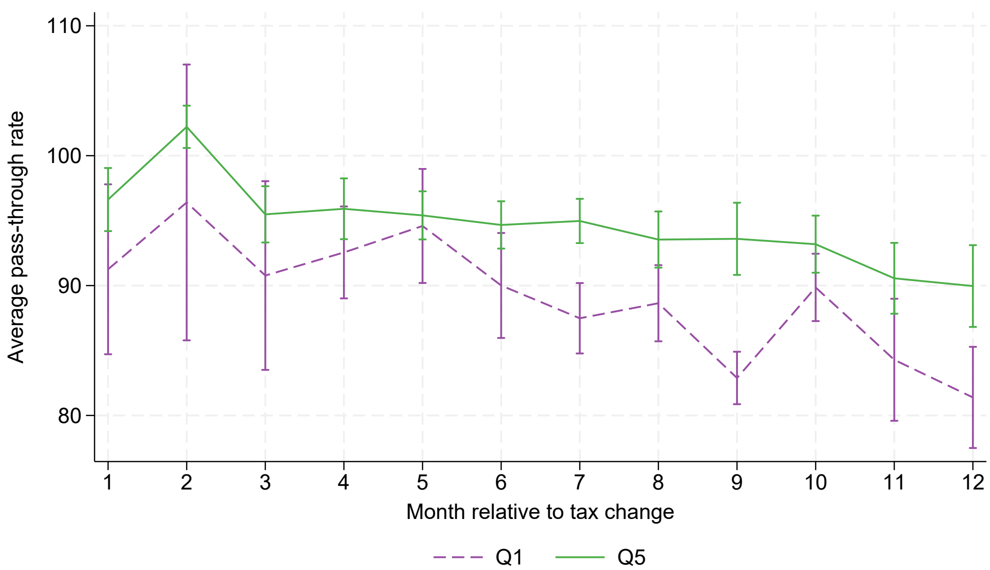

(b) *Avg. pass-through: income quintiles 1 and 5*

Notes: Lines are constructed from monthly DD point estimates of average pass-through rate ($\frac{\Delta P}{\Delta t}$) by income group, estimated using a monthly analog of Equation 7, where $\ell \in [1, 12]$. Panel A depicts all six income groups, while Panel B displays only Q1 and Q5 and includes confidence intervals. Observations are station-year-months and are weighted by station-year diesel quantity sold. Standard errors are clustered by province. Q1-Q5 indicate quintiles of the municipality median income distribution; "No Inc" denotes the group of stations in towns too small to be included in the data.

Table 3: *Average tax pass-through by income group*

	No Inc (1)	Q1 (2)	Q2 (3)	Q3 (4)	Q4 (5)	Q5 (6)
Panel A. Station-month level						
Average pass-through (%)	0.865	0.893	0.960	0.937	0.952	0.953
	(0.049)	(0.014)	(0.016)	(0.011)	(0.018)	(0.006)
N	14,347	5,763	8,647	6,965	10,322	5,905
Panel B. Municipality-month level						
Average pass-through (%)	0.8676	0.8762	0.9378	0.9344	0.9502	0.9459
	(0.049)	(0.015)	(0.022)	(0.016)	(0.018)	(0.011)
N	9,770	1,973	1,809	1,678	1,183	689

Notes: Estimates are average pass-through rates of tax changes into retail prices, $\frac{\Delta P}{\Delta t}$, obtained via Equation 7, for each of six income groups. Panel A displays results with station-month observations; Panel B results are from municipality-month observations. 'N' reports the number of treatment-group station-months used. The post-tax estimation window is 12 months. Observations are weighted by annual station (or municipality) sales. Standard errors are clustered by province. Q1-Q5 indicate quintiles of the municipality median income distribution; "No Inc" denotes the group of stations in towns too small to be included in the data.

to Q2 and generally peaking *before* Q5. Magnitudes vary modestly across dCDH specifications, but, notably, TWFE yields larger estimated differences across groups – 10-17 percentage points between lower and higher pass-through groups, as compared with 7-10 percentage points using dCDH. This implies that the "forbidden comparisons" included in TWFE bias estimates of heterogeneity upwards in this context.[18]

5 Discussion

The previous section catalogs evidence on the answer to the main research question: what is the relationship between income and fuel tax pass-through? Put together, the evidence suggests that average pass-through is relatively lower – in the range of 4-9 percent – in the lowest-income municipalities of spain. In this section, I discuss three extensions of the main research question: (1) How does the pass-through/income relationship affect distributional welfare? (2) What is *causing* pass-through to be higher in higher-income places? (3) Does the pass-through of other elements of retail fuel cost behave the same way with respect to income?

[18]Appendix Table A2 presents results from one further alternative specification, which replaces median income with per capita income in the calculation of income quintile. These results are noteworthy because quintile-1 pass-through estimates are elevated relative to those of the preferred income definition, and not different from those of Q2-Q5. Of 226 municipalities assigned to Q1 by *per capita* income, 55 (24 percent) are in a *higher* quintile when defined by median, and these have a higher average pass-through rate than the remaining 171 municipalities assigned to Q1 under both definitions. At the same time, there are also 55 municipalities assigned to Q1 by *median* income but a higher quintile by per capita; these exhibit relatively *lower* average pass-through.

Table 4: *Average tax pass-through by income group*

	No Inc (1)	Q1 (2)	Q2 (3)	Q3 (4)	Q4 (5)	Q5 (6)
Panel A. Balanced panel						
Average pass-through (%)	0.865	0.892	0.969	0.936	0.952	0.939
	(0.053)	(0.014)	(0.017)	(0.011)	(0.019)	(0.006)
N	13,042	5,635	8,381	6,647	9,856	4,705
Panel B. Unweighted						
Average pass-through (%)	0.887	0.902	0.952	0.951	0.955	0.950
	(0.011)	(0.012)	(0.014)	(0.011)	(0.018)	(0.007)
N	15,025	5,763	8,647	6,965	10,640	5,905
Panel C. Two-way-fixed effects						
Average pass-through (%)	0.833	0.876	0.938	1.001	0.987	0.969
	(0.054)	(0.033)	(0.027)	(0.037)	(0.037)	(0.036)
N	149,548	33,953	53,960	48,808	80,175	95,532

Notes: Estimates are average pass-through rates of tax changes into retail prices, $\frac{\Delta P}{\Delta t}$, obtained via Equation 7, for each of six income groups. The preferred specification, in Table 3, Panel A, uses station-month observations, weighted by station-year diesel sales, with a twelve-month post-period and standard errors clustered by province. In this table, the Panel A specification differs from the base by using only a balanced panel of stations that are observed and have a viable control group for the relevant twelve-month post-period. Panel B shows results from the base specification except without (analytical) weights. Panel C exhibits the results of two way fixed effects regression. Q1-Q5 indicate quintiles of the municipality median income distribution; "No Inc" denotes the group of stations in towns too small to be included in the data. 'N' reports the number of treatment-group station-months used in Panels A and B and the number of total station-month observations used in Panel C.

5.1 Effect on welfare

The fact that tax pass-through is positively correlated with income has significant implications for the distributional welfare impacts of the tax. Mathematically, the first-order approximation of the consumer surplus loss induced by a marginal tax hike ($\frac{dCS}{dt}$) is the additional cost paid to continue consuming the taxed good – in a graph of supply and demand, it is a rectangle with width equal to quantity consumed (Q) and height equal to the pass-through rate ($\frac{dp}{dt}$; Weyl and Fabinger, 2013). Equivalently, and more intuitively, it is likely that the primary burden on a (gas or diesel powered) car owner's mind when a tax is raised is the extra cost of all the fuel that they will continue to purchase. The welfare loss from this extra cost is likely to be significantly larger than the welfare loss due to reduced consumption, because demand for retail automotive fuel tends to be relatively inelastic (though less than previously thought; Kilian and Zhou, 2024).

The progressivity of a marginal tax increase could thus be approximated by estimating, at different points of the wealth (W) spectrum,

$$\frac{Q\frac{dp}{dt}}{W}, \tag{8}$$

This is equivalent to the extra cost of consuming fuel after the tax change, as a proportion of wealth. If the above expression rises with wealth, then the approximation is progressive; if it falls, then the approximation is regressive. It has historically been standard practice to omit pass-through from the calculation (e.g., Poterba, 1991), implicitly assuming that pass-through is uniform across the wealth spectrum. The exception to date is Harju et al. (2022), who estimate fuel tax pass-through rates that *fall* in local incomes and use them in their calculation of $\frac{\Delta CS}{dt}$.

I stop short of calculating pass-through-adjusted changes in consumer surplus here because of data limitations. The Spanish consumer expenditure survey (Encuesta de Presupuestos Familiares) provides household auto fuel consumption, expenditures, and income with survey weights for national representativeness, but location is masked. I can only estimate pass-through rates by quantile of *municipality* median income. In addition, income data are missing for two regions (País Vasco and Navarra) as well as municipalities with fewer than 5,000 residents. Thus, matching pass-through rates to consumption levels requires an unrealistic assumption that quantile of (observed) municipality income equals quantile of household income. Finally, Equation 8 is only an approximation, and a potentially misleading one given evidence that demand elasticity can vary with income. This is because all else equal, more elastic demand causes both lower pass-through *and* a greater reduction in fuel consumption in response a tax change. These effects work in opposite directions; the pass-through effect dominates when demand is inelastic, but the consumption effect offsets it to some degree.[19]

I do, however, present pass-through estimates by municipality income decile alongside proportional fuel consumption by household expenditure decile,[20] to illustrate how heterogeneous pass-through in principle could affect distributional welfare. Figure 7 plots the pass-through estimates and their confidence intervals. It shows a non-monotonic pattern in municipality income decile, with higher pass-through on average in the upper half of the distribution. I include, at income decile "0", the pass-through rate measured in the group of small-town stations with no income data. To the extent that such towns have the lowest incomes, the decile-0 rate contributes further to the progressive pass-through pattern from Q1 to Q10.

Figure 8 shows, using expenditure survey data from 2010 and 2013 (the first and last year in which treated stations are observed), that relatively poorer households consume relatively more fuel per dollar of expenditure. The consumption/wealth relationship is stronger than the pass-through wealth relationship and implies that the consumer surplus impacts of these tax changes are regressive. If the progressive pass-through patterns I identify were to hold across the national household wealth distribution, the effect would be a dampening of the apparent regressivity in Figure 8. In my context,

[19] Harju et al. (2022) get around this issue by assuming that demand is perfecly inelastic, but they, too, acknowledge that their exercise should only be considered an illustration, rather than definitive findings.

[20] Poterba (1991) shows that household expenditure is a more accurate proxy for lifetime wealth than income.

Figure 7: *Average pass-through by income decile*

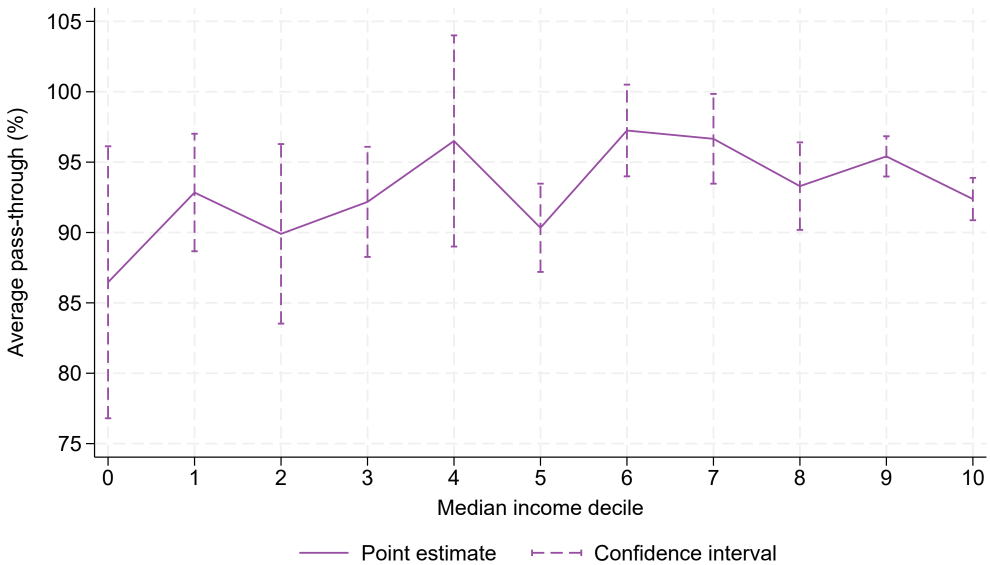

Notes: Data points are twelve-month average pass-through rates of tax changes and associated confidence intervals, obtained via Equation 7 separately for each income decile as well as the group of stations without income data (corresponding to $x = 0$). Observations are station-year-months and are weighted by annual station-level quantity sold. Standard errors are clustered at the province level.

the magnitude of this dampening is likely modest: households in the bottom expenditure decile buy 25-60 percent more fuel per dollar total expenditure, while the maximum decile-specific pass-through rate (Figure 7) is only about 12 percent larger than the minimum.

5.2 Causal mechanisms

My main result is descriptive: pass-through of the *céntimo sanitario* tax is higher in higher-income municipalities. This descriptive fact is what matters for distributional welfare in the Spanish context; but we might wonder what is *causing* this pattern, for instance to inform external validity considerations. Higher pass-through in wealthier areas is consistent with wealthier areas having more inelastic demand. However, many different causal channels are possible and empirically established, from the toughness of competition, to the degree of vertical integration, to the supply elasticity, among many others.

The dCDH method does not accommodate interaction terms with pass-through, so I cannot simultaneously estimate relationships between pass-through and several characteristics like is commonly done with TWFE regression. Instead I focus on jointly considering municipal income and station *brand* as predictors of pass-through. Table 5 breaks pass-through out by brand group – refiner (Panel A) or independent (Panel B) – in addition to income group. Point estimates are obtained from Equation 7 separately in subsamples corresponding to each combination of brand and income groups. Brand group is highly predictive of pass-through overall (column 1) – refiner-brand stations

Figure 8: *Household auto fuel consumption relative to expenditure*

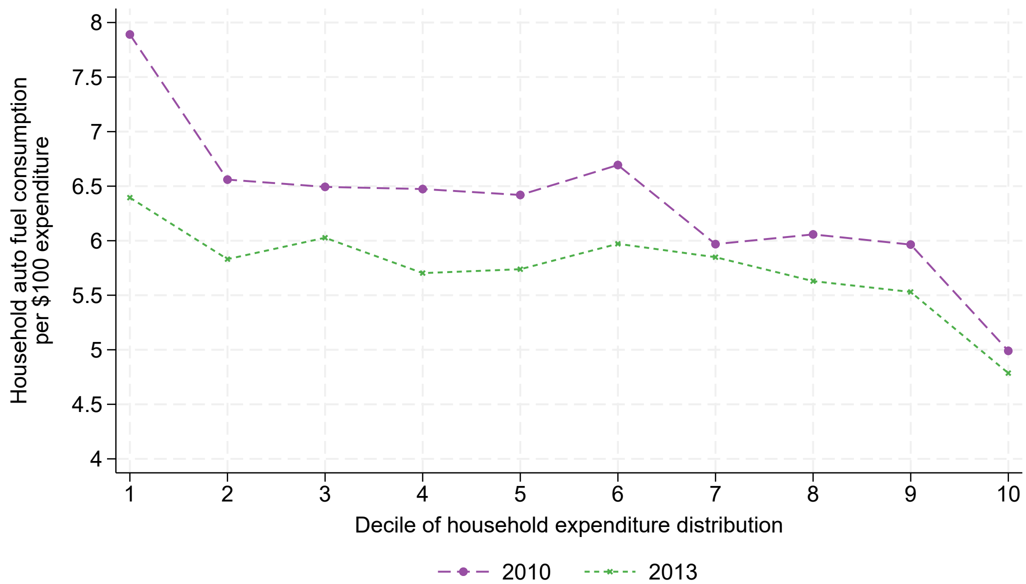

Notes: Data are from the Spanish household expenditure survey (Encuesta de Presupuestos Familares). I apply survey weights for national representativeness.

pass-through 95.4 percent of a tax change on average, while independent ones pass-through 81.7 percent.[21] However, differences in pass-through persist across income groups, even within brand type. Q2-Q5 pass-through rates are higher than the Q1 rate among refiner-branded stations (panel A, columns 3-7), just as in the full national sample. Refiner-branded stations in small towns (column 2) buck this income trend, with a pass-through rate (94.0 percent) closer to those of Q2-Q5.

Pass-through also varies with income among independent-brand stations (Panel B), though the pattern is significantly different from that of refiner-branded stations. Independent stations pass through less of a tax change in quintiles 3 and 5 then in the other three quintiles. More strikingly, pass-through is quite low in small towns missing income data (58.5 percent; column 2). Put together, Table 5 shows that conditioning on brand group does not eliminate all variation in pass-through across income groups. At the same time, the divergent patterns in Panels A and B cast doubt on the notion that variable demand elasticities are the primary cause of the income/pass-through correlation.[22]

[21] Appendix Figure A5 presents event study plots for the two brand groups, and Figure A6 breaks out the event study further into (brand-)income groups.

[22] Appendix Table A3 presents coefficients from four different TWFE "interaction" regressions, successively adding interaction terms between the tax variable and station and local attributes. The income/pass-through relationship remains significant even in the "kitchen-sink" regression including interaction terms between the tax and (a) brand group indicators, (b) proxies for spatial competition, (c) a vertical integration indicator, (d) population density, and (e) four station amenities. I caution against reading too much into TWFE results in my context but include them as the best of flawed options for "horseracing" predictors.

Table 5: *Pass-through and station branding*

	All (1)	No Inc (2)	Q1 (3)	Q2 (4)	Q3 (5)	Q4 (6)	Q5 (7)
Panel A. Refinery-branded stations							
Average pass-through (%)	0.954	0.940	0.841	0.968	0.959	0.970	0.982
	(0.008)	(0.017)	(0.019)	(0.019)	(0.012)	(0.015)	(0.007)
N	32,595	8,465	3,491	5,388	4,555	6,808	3,705
Panel B. Independent-branded stations							
Average pass-through (%)	0.817	0.585	0.925	0.913	0.869	0.942	0.855
	(0.049)	(0.140)	(0.018)	(0.029)	(0.035)	(0.044)	(0.022)
N	12,637	4,338	1,638	2,138	1,360	2,052	797

Notes: Point estimates are average pass-through rates of tax changes into retail prices, estimated via Equation 7. In Panel A, only stations bearing the brand of an oil refiner in Spain are considered; in Panel B, only stations with an *independent* brand are considered. Column 1 displays overall average pass-through rates by brand type, and columns 2-6 display brand-specific rates by income grouping (quintiles $Q1$-$Q5$ and the group of stations without income data). All specifications use station-month observations, weighted by station-year diesel sales, with a twelve-month post-period and standard errors clustered by province. 'N' is the number of treatment-group station-months used.

5.3 Pass-through of other fuel costs

Taxes are nearly half of retail prices in this time period in Spain; the production cost of fuel is the other major source of variable costs of auto fuel supply. There is ongoing interest in whether the fuel elasticity of demand with respect to taxation differs from the elasticity with respect to non-tax costs. A difference is possible because gas taxes are more persistent than (e.g.) the price of crude oil and also sometimes more salient to the consumer (Li et al., 2014), among other potential explanations. The question is policy relevant because it dicates whether one can use either of the estimates to proxy for the other, for instance in a forecast of the impacts of a hypothetical tax change. Many researchers have found that the two elasticites do differ (e.g., Davis and Kilian, 2014; Li et al., 2014), but Kilian and Zhou (2024) is a recent, notable exception in the US context.

I test for heterogeneity in crude oil price pass-through by income group using the same regression specification as Kilian and Zhou (2024):

$$\Delta P_{it} = \alpha + \beta \Delta \text{Crude}_t + \theta_y + \nu_m + \phi_i + \epsilon_{it}, \tag{9}$$

with fixed effects by year θ_y and by month ν_m separately rather than year-month t, to accommodate the time-series of crude oil price as the "treatment" variable. ΔP and ΔCrude are changes in pretax retail price and Brent crude oil price from the previous month to the current one. Table 6 documents the income group specific pass-through estimates. These differ from each only minimally – the highest and lowest rates are separated by only one percentage point in the preferred specification (sales-weighted, in Panel A), and 2.6 percentage points in the case of unweighted analysis (Panel B). While the difference is marginally statistically significant, its magnitude is noticeably smaller than the analogous estimated difference for *tax* pass-through. This suggests that there is, in fact, a difference here between the consumer response to an auto fuel tax versus an oil price shock.

Table 6: *Average crude price pass-through by income group*

	No inc (1)	Q1 (2)	Q2 (3)	Q3 (4)	Q4 (5)	Q5 (6)
Panel A. Sales-weighted						
1-month Δ Brent crude price (c/l)	0.998	0.987	0.990	0.997	0.995	0.998
	(0.003)	(0.003)	(0.003)	(0.004)	(0.003)	(0.002)
N	147,071	33,394	53,103	48,034	78,899	94,002
Panel B. Unweighted						
1-month Δ Brent crude price (c/l)	0.955	0.953	0.967	0.973	0.978	0.979
	(0.004)	(0.003)	(0.004)	(0.005)	(0.004)	(0.003)
N	170,843	38,485	60,398	54,841	90,779	105,547

Notes: Point estimates are average pass-through rates of Brent crude oil price changes into retail prices, estimated via Equation 9, separately for each income quintile ($Q1$-$Q5$) as well as the group of stations without income data ("No Inc""). Panel A displays results with station-month observations; Panel B results are from municipality-month observations. The post-tax estimation window is 12 months. 'N' is the number of station-months used; observations are weighted by annual station (or municipality) sales. Standard errors are clustered by province.

6 Conclusion

Automotive fuel taxes are used in dozens of countries around the world, and energy taxes more generally are a key current and future policy lever for driving reductions in fossil fuel use. The distributional equity of such policies is a fundamental political and moral issue that cannot be assessed by theory alone. In this paper, I contribute to our understanding of the distributional impacts of energy taxation by studying the empirical relationship between auto fuel tax pass-through and income. If pass-through varies with wealth, it has consequences for distributional welfare that are not generally accounted for in welfare analysis.

I find that a regional auto fuel tax in Spain, which was increased 14 times in staggered fashion across 11 regions from 2010-2012, exhibits variable pass-through in income. While my national average pass-through estimate is 91.4 percent, my estimate of pass-through in the bottom quintile of municipal median income is 4-7 percentage points lower than in the upper four quintiles. Furthermore, stations in towns too small to have recorded income data have the lowest estimated pass-through. Given that such towns are likely among the lowest-income municipalities in the country, the evidence suggests that fuel tax pass-through is itself progressive in income. While this stands in direct contrast to the one other existing study of energy tax pass-through and income (Harju et al., 2022), both results are possible because of the many causes of heterogeneous pass-through and the varying empirical contexts being studied. In my context, other observable attributes of stations and markets – such as branding, vertical structure, and spatial competition – do not appear to fully explain the correlation of pass-through with income.

My results illustrate the relevance of heterogeneous pass-through to distributional welfare analysis; a negative pass-through/wealth relationship makes a fuel tax's welfare impacts more regressive than they appear under a uniform pass-through assumption, while a positive relationship makes those welfare impacts *less* regressive. This study also provides an early example of the application of one of the newest difference-in-differences methods, from de Chaisemartin and D'Haultfoeuille (2024a), to facilitate unbiased treatment effect estimation in settings with staggered, non-binary treatments. Future research efforts may profitably focus on improving the wealth measure by which to separately estimate pass-through rates (income is an imperfect proxy for lifetime wealth, and municipal measures are not the same as household ones); determining the causes of the pass-through/wealth correlation; and integrating variable pass-through rates into a fuller welfare analysis that captures the lost welfare from reduced consumption.

References

Agencia Tributaria (2024): "Devolución especial IVMDH: Preguntas Frecuentes," .

Alm, J., E. Sennoga, and M. Skidmore (2009): "Perfect Competition, Urbanization, and Tax Incidence in the Retail Gasoline Market." *Economic Inquiry*, 47, 118–134.

Atkin, D. and D. Donaldson (2015): "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs." Working Paper.

Bajo-Buenestado, R. and M. Borrella-Mas (2022): "The Heterogeneous Tax Pass-through under Different Vertical Relationships," *Economic Journal*, 132, 1684–1708.

Bento, A., L. Goulder, M. Jacobsen, and R. von Haefen (2009): "Distributional and Efficiency Impacts of Increased US Gasoline Taxes." *American Economic Review*, 99, 667–699.

Bonnet, C., P. Dubois, S. Villas Boas, and D. Klapper (2013): "Empirical Evidence on the Role of Nonlinear Wholesale Pricing and Vertical Restraints on Cost Pass-through." *Review of Economics and Statistics*, 95, 500–515.

Borenstein, S., A. Cameron, and R. Gilbert (1997): "Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?" *Quarterly Journal of Economics*, X, 306–339.

Borusyak, K., X. Jaravel, and J. Spiess (2024): "Revisiting Event-Study Designs: Robust and Efficient Estimation." *Review of Economic Studies*, 1–33.

Callaway, B. and P. Sant'Anna (2021): "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics*, 225, 200–230.

Cameron, A. and D. Gilbert (2015): "A Practitioner's Guide to Cluster-Robust Inference." *Journal of Human Resources*, 50, 317–372.

Caspersen, E. and G. Metcalf (1994): "Is a Value Added Tax Regressive? Annual Versus Lifetime Incidence Measures." *National Tax Journal*, 47, 731–746.

Cavallo, A., G. Gopinath, B. Neiman, and J. Tang (2021): "Tariff Pass-through at the Border and at the Store: Evidence from US Trade Policy." *American Economic Review: Insights*, 3, 19–34.

Chetty, R. (2009): "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods." *Annual Review of Economics*, 1, 451–487.

Comisión Nacional de la Competencia (2012): "Report on the Consultation Request Submitted by the Secretary of State for the Economy and Support for Enterprises Regarding the Automotive Fuel Market in Spain." Tech. rep., CNC, España, <https://www.cnmc.es/sites/default/files/1186976_7.pdf>.

Comisión Nacional de los Mercados y la Competencia (2024): "Boletín Informativo de la Distribución de Carburantes en Estaciones de Servicio." Tech. rep., CNC, España, <https://www.cnmc.es/sites/default/files/5218929.pdf>.

Contín-Pilart, I., A. Correljé, and M. Palacios (2009): "Competition, Regulation, and Pricing Behaviour in the Spanish Retail Gasoline Market." *Energy Policy*, 37, 219–228.

Court of Justice of the European Union (2014): "Press Release: The Tax on Retail Sales of Certain Hydrocarbons is Contrary to European Union Law," .

Davis, L. and L. Kilian (2014): "Estimating the Effect of a Gasoline Tax on Carbon Emissions." *Journal of Applied Econometrics*, 62, 1187–1214.

de Chaisemartin, C., D. Ciccia, X. D'Haultfoeuille, F. Knau, M. Malézieux, and D. Sow (2024): "Event Study Estimators of Variance Estimators Computed by the did_multiplegt_dyn Command." .

de Chaisemartin, C. and X. D'Haultfoeuille (2020): "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review*, 110, 2964–2996.

——— (2023): "Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey." *Econometrics Journal*, 26, C1–C30.

——— (2024a): "Difference-in-Differences Estimators of Intertemporal Treatment Effects." *Review of Economics and Statistics*, 1–45.

——— (2024b): "Stata help file: did_multiplegt," .

Doyle Jr., J. and K. Samphantharak (2008): "$2.00 Gas! Studying the Effects of a Gas Tax Moratorium." *Journal of Public Economics*, 92, 869–884.

Fabra, N. and M. Reguant (2014): "Pass-through of Emissions Costs in Electricity Markets." *American Economic Review*, 104, 2872–2899.

Flues, F. and Thomas, A. (2015): "The Distributional Effects of Energy Taxes." Tech. rep., Organization for Economic Cooperation and Development, <https://www.oecd-ilibrary.org/the-distributional-effects-of-energy-taxes_5js1qwkqqrbv.pdf?itemId=%2Fcontent%2Fpaper%2F5js1qwkqqrbv-en&mimeType=pdf>.

Ganapati, S., J. Shapiro, and R. Walker (2020): "Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes." *American Economic Journal: Applied Economics*, 12, 303–342.

Genakos, C. and M. Pagliero (2022): "Competition and Pass-through: Evidence from Isolated Markets." *American Economic Journal: Applied Economics*, 14, 35–57.

Goodman-Bacon, A. (2021): "Difference-in-Differences with Variation in Treatment Timing." *Journal of Econometrics*, 225, 254–277.

Gruber, J. and B. Koszegi (2004): "Tax Incidence When Individuals are Time-Inconsistent: The Case of Cigarette Excise Taxes." *Journal of Public Economics*, 88, 1959–1987.

HARJU, J., T. KOSONEN, M. LAUKKANEN, AND K. PALANNE (2022): "The Heterogeneous Incidence of Fuel Carbon Taxes: Evidence from Station-Level Data." *Journal of Environmental Economics and Management*, 112, 102607.

HOROWITZ, J., J.-A. CRONIN, H. HAWKINS, L. KONDA, AND A. YUSKAVAGE (2017): "Methodology for Analyzing Carbon Taxes." Tech. rep., United States Department of Treasury, Office of Tax Analysis, <https://home.treasury.gov/system/files/131/WP-115.pdf>.

HOUDE, J.-F. (2012): "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline." *American Economic Review*, 102, 2147–2182.

HUGHES, J., C. KNITTEL, AND D. SPERLING (2008): "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand." *Energy Journal*, 29, 113–134.

JAFFE, S. AND E. WEYL (2013): "The First-Order Approach to Merger Analysis." *American Economic Journal: Microeconomics*, 5, 188–218.

JENKIN, H. (1872): "On the Principles that Regulate the Incidence of Taxes." *Proceedings of the Royal Society of Edinburgh*, 7, 618–631.

KAYSER, H. (2000): "Gasoline Demand and Car Choice: Estimating Gasoline Demand Using Household Information," *Energy Economics*, 22, 331–348.

KILIAN, L. AND X. ZHOU (2024): "Heterogeneity in the Pass-through from Oil to Gasoline Prices: A New Instrument for Estimating the Price Elasticity of Gasoline Demand." *Journal of Public Economics*, 232, 105099.

KROFT, K., J.-W. LALIBERTE, R. LEAL-VIZCAINO, AND M. NOTOWIDIGDO (2024): "Salience and Taxation with Imperfect Competition," *Review of Economic Studies*, 91, 403–437.

LADE, G. AND J. BUSHNELL (2022): "Fuel Subsidy Pass-Through and Market Structure: Evidence from the Renewable Fuel Standard." *Journal of the Association of Environmental and Resource Economists*, 6, 563–592.

LI, S., J. LINN, AND E. MUEHLEGGER (2014): "Gasoline Taxes and Consumer Behavior." *American Economic Journal: Economic Policy*, 6, 302–342.

MARION, J. AND E. MUEHLEGGER (2011): "Fuel Tax Incidence and Supply Conditions." *Journal of Public Economics*, 95, 1202–1212.

MARSHALL, A. (1890): *Principles of Economics.*, New York: Macmillan.

MATHUR, A. AND A. MORRIS (2012): "Distributional Effects of a Carbon Tax in Broader U.S. Fiscal Reform." Tech. rep., Brookings Institute.

METCALF, G. (1999): "A Distributional Analysis of Green Tax Reforms," *National Tax Journal*, 52, 655–681.

Miller, N., M. Osborne, and G. Sheu (2017): "Pass-through in a Concentrated Industry: Empirical Evidence and Regulatory Implications." *RAND Journal of Economics*, 48, 69–93.

Muehlegger, E. and D. Rapson (2022): "Subsidizing Low- and Middle-Income Adoption of Electric Vehicles: Quasi-experimental Evidence from California," *Journal of Public Economics*, 216, 104752.

Muehlegger, E. and R. Sweeney (2022): "Pass-through of Own and Rival Cost Shocks: Evidence from the US Fracking Boom," *Review of Economics and Statistics*, 104, 1361–1369.

Noceda, M. and M. Jimenez (2015): "Competition Authorities Fine Five Oil Companies 32.4 Million Euros for Price Fixing." *El País*, <https://elpais.com/economia/2015/02/25/actualidad/1424859682_727919.html>.

Perdiguero, J. and J. Borrell (2007): "La Difícil Conducción de la Competencia por el Sector de las Gasolineras en España." *Economía Industrial*, 365, 113–125.

Pizer, W. and S. Sexton (2020): "The Distributional Impacts of Energy Taxes." *Review of Environmental Economics and Policy*, 13, 104–123.

Pless, J. and A. van Benthem (2019): "Pass-through as a Test for Market Power: An Application to Solar Subsidies." *American Economic Journal: Applied Economics*, 11, 367–401.

Poterba, J. (1991): "Is the Gasoline Tax Regressive?" *Tax Policy and the Economy*, 5, 145–164.

——— (1996): "Retail Price Reactions to Changes in State and Local Sales Taxes." *National Tax Journal*, 49, 165–176.

Rico-Straffon, J., Z. Wang, S. Panlasigui, C. Loucks, J. Swenson, and A. Pfaff (2023): "Forest Concessions and Eco-certifications in the Peruvian Amazon: Deforestation Impacts of Logging Rights and Logging Restrictions." *Journal of Environmental Economics and Management*, 118, 102780.

Roth, J. (2022): "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends." *American Economic Review: Insights*, 4, 305–322.

Roth, J., P. Sant'Anna, A. Bilinski, and J. Poe (2023): "What's Trending in Difference-DIfferences? A Synthesis of the Recent Econometrics Literature." *Journal of Econometrics*, 235, 2218–2244.

Sterner, T. (2012): "Fuel Taxes, Climate and Tax Incidence." *Fuel Taxes and the Poor: The Distributional Effects of Gasoline Taxation and Their Implications for Climate Policy*.

Sun, L. and S. Abraham (2021): "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." *Journal of Econometrics*, 225, 175–199.

West, S. (2004): "Distributional Effects of Alternative Vehicle Pollution Control Policies." *Journal of Public Economics*, 88, 735–757.

Weyl, E. and M. Fabinger (2013): "Pass-through as an Economic Tool: Principles of Incidence under Imperfect Competition." *Journal of Political Economy*, 121, 528–583.

Wing, C., M. Yozwiak, A. Hollingsworth, S. Freedman, and K. Simon (2024): "Designing Difference-in-DIfference Studies with Staggered Treatment Adoption: Key Concepts and Practical Guidelines." *Annual Review of Public Health*, 45, 485–505.

Appendix

Figure A1: Geoportal *screenshot*

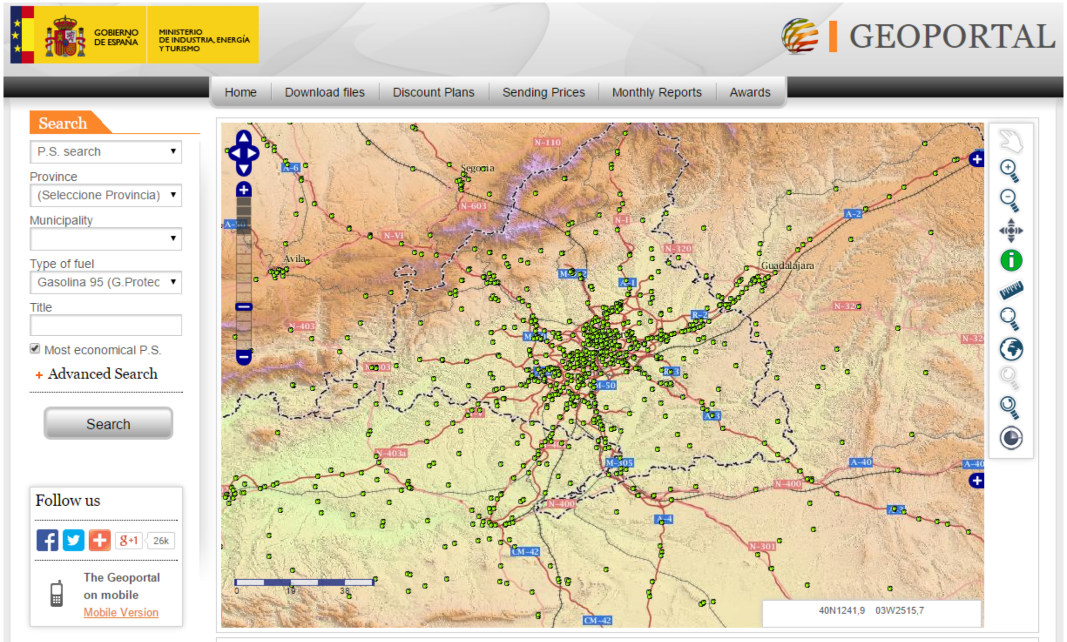

Notes: Green dots are Spanish retail diesel stations. The screenshot shows the Madrid metro area. Source: <https://geoportalgasolineras.es/geoportal-instalaciones/Inicio>, accessed on February 15th, 2015.

Figure A2: *Income distribution of stations, by sales data availability*

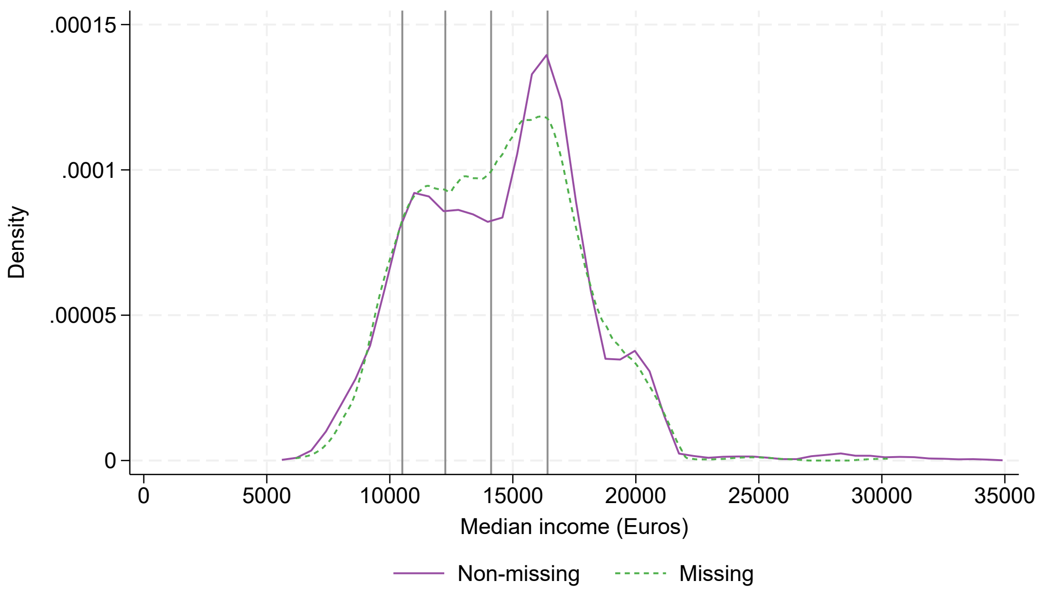

Notes: Plotted curves are kernel densities, using an Epanechnikov kernel, of average municipality median income from 2008-2013 among municipalities missing versus not missing income data.

Figure A3: *Income and pass-through, balanced panel*

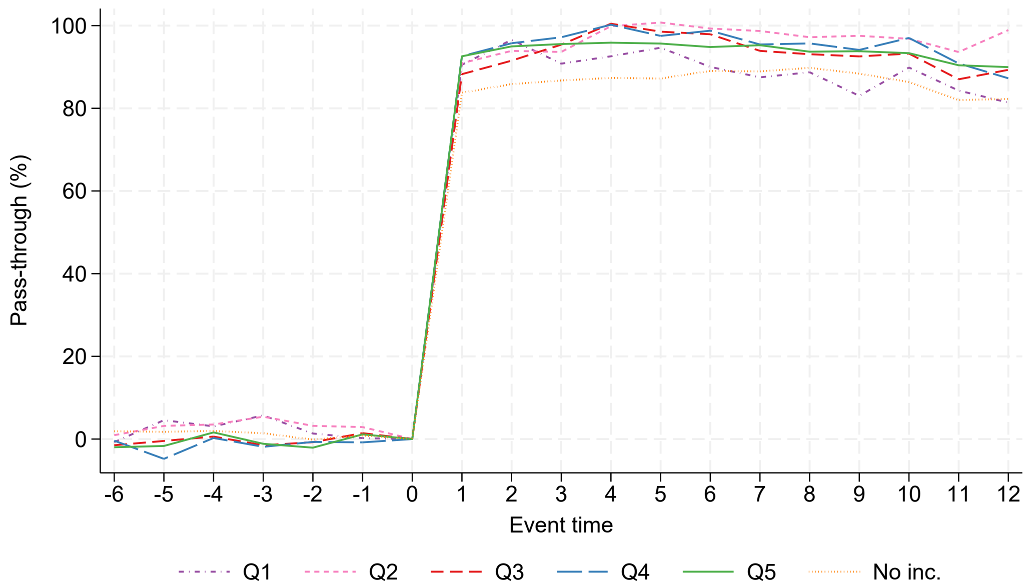

(a) *Average pass-through by income quintile*

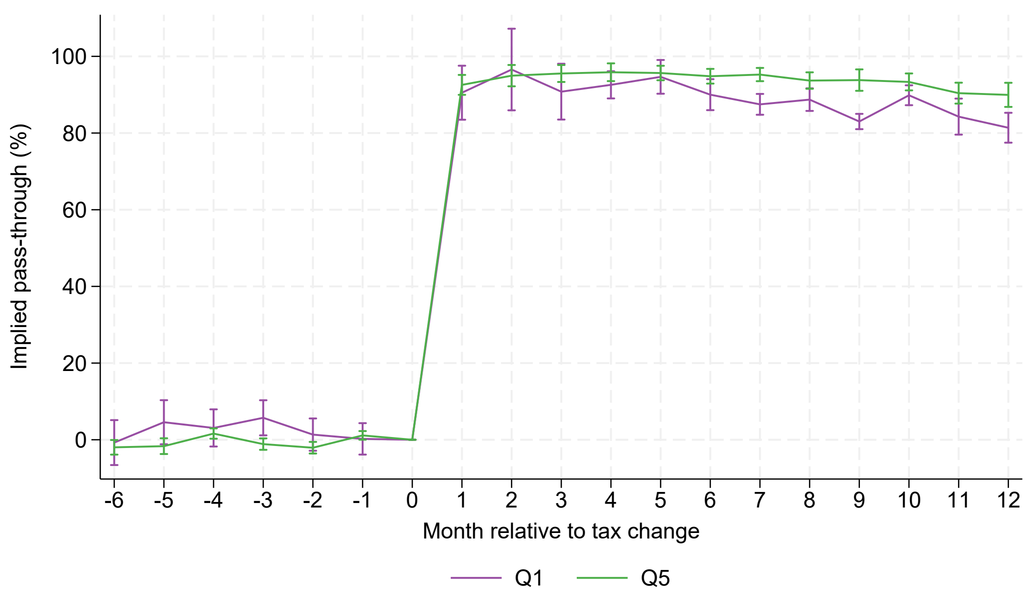

(b) *Average pass-through: income quintiles 1 and 5*

Notes: Lines are constructed from monthly DD point estimates of average pass-through rate ($\frac{\Delta p}{\Delta t}$) by income group, estimated using a monthly analog of Equation 7, where $\ell \in [1, 12]$. The sample is restricted to only using treatment-group stations if they are observed for all twelve post-treatment months. Panel A depicts all six income groups, while Panel B displays only Q1 and Q5 and includes confidence intervals. Observations are station-year-months and are weighted by station-year diesel quantity sold. Standard errors are clustered by province. Q1-Q5 indicate quintiles of the municipality median income distribution; "No Inc" denotes the group of stations in towns too small to be included in the data.

Figure A4: *Income and pass-through, no weights*

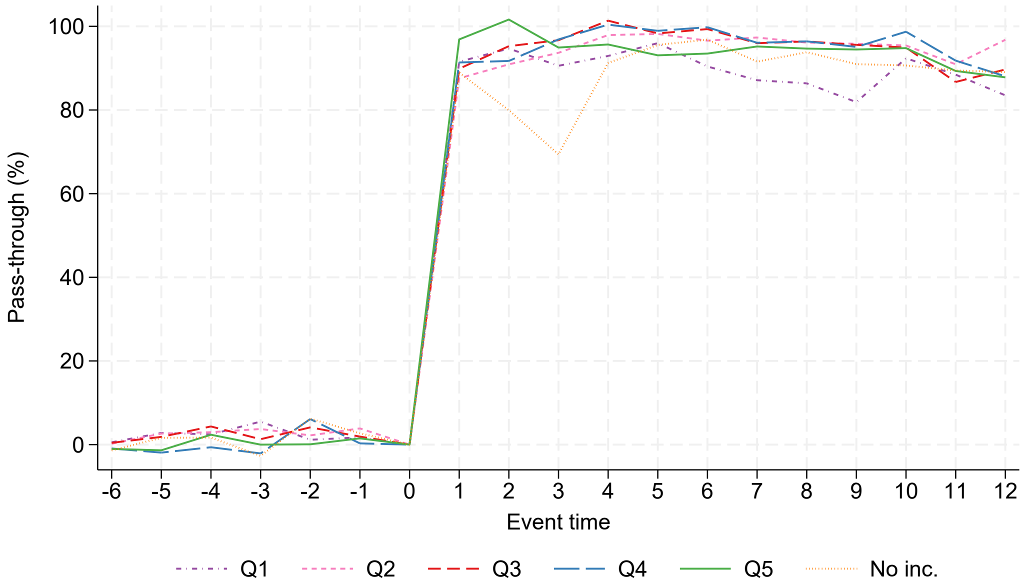

(a) *Average pass-through by income quintile*

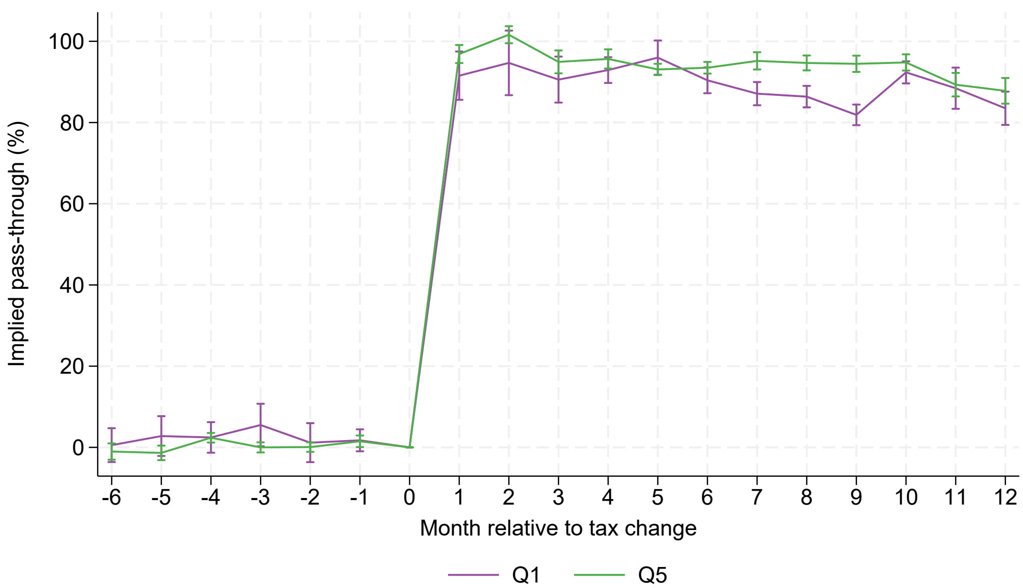

(b) *Average pass-through: income quintiles 1 and 5*

Notes: Lines are constructed from monthly DD point estimates of average pass-through rate ($\frac{\Delta p}{\Delta t}$) by income group, estimated using a monthly analog of Equation 7, where $\ell \in [1, 12]$. Panel A depicts all six income groups, while Panel B displays only Q1 and Q5 and includes confidence intervals. Observations are station-year-months and are unweighted, unlike in the preferred specification. Standard errors are clustered by province. Q1-Q5 indicate quintiles of the municipality median income distribution; "No Inc" denotes the group of stations in towns too small to be included in the data.

Figure A5: *Station brand and pass-through*

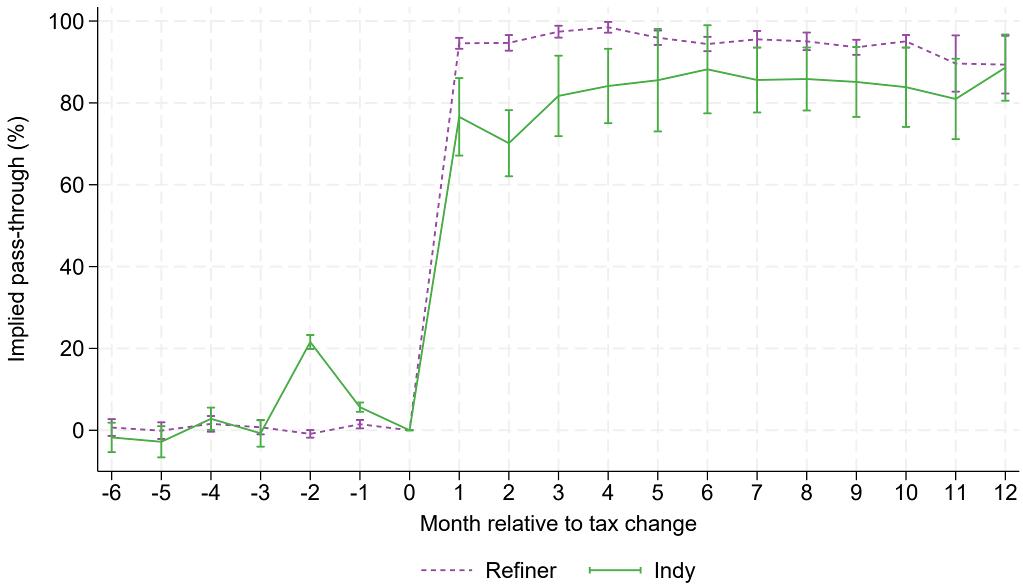

Notes: Lines are constructed from monthly DD point estimates of average pass-through rate ($\frac{\Delta P}{\Delta t}$) by brand group, estimated using a monthly analog of Equation 7, where $\ell \in [1, 12]$. The sample is restricted to only using treatment-group stations if they are observed for all twelve post-treatment months. Observations are station-year-months and are weighted by station-year diesel quantity sold. Standard errors are clustered by province. "Refiner" denotes refiner-branded stations, while "Independent" denotes brands unaffiliated with a refiner or wholesaler.

Figure A6: *Absolute price impacts by brand type and income group*

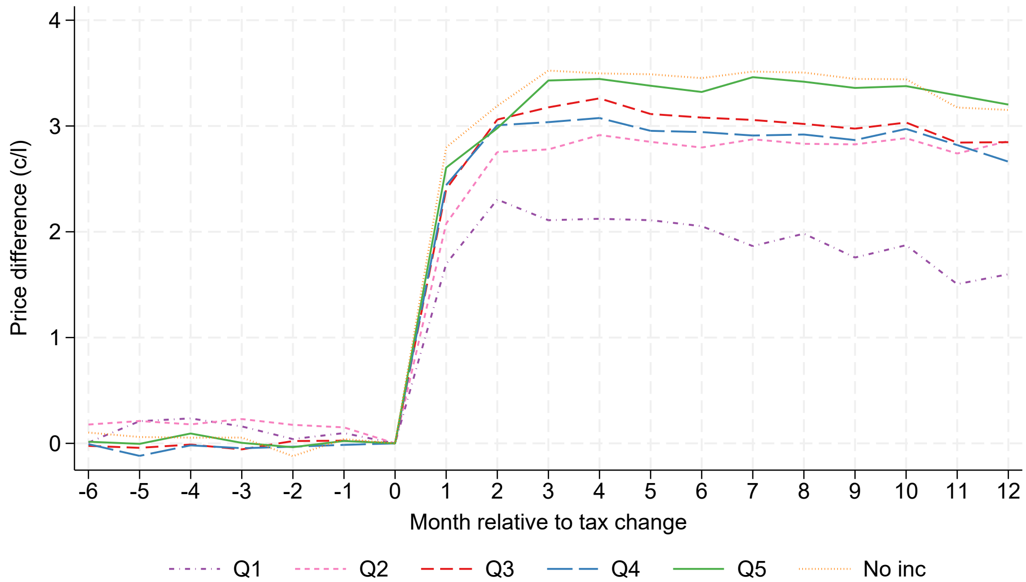

(a) *Refiner-branded stations*

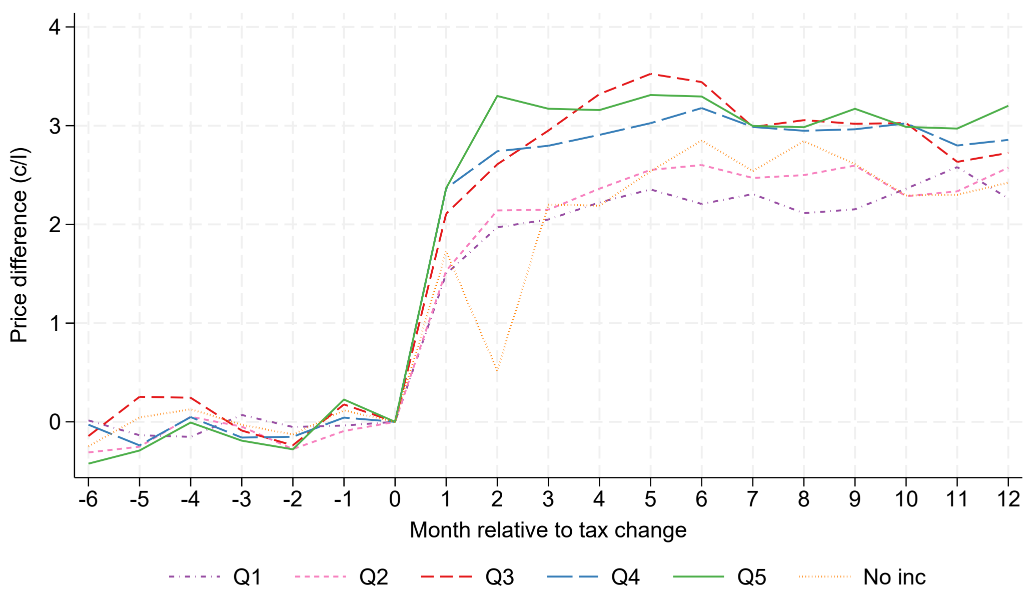

(b) *Independent-branded stations*

Notes: Lines are constructed from monthly DD point estimates of average absolute price impact by income group and brand group, estimated using Equation 6, where $\ell \in [1, 12]$. Observations are station-year-months and are weighted by station-year diesel quantity sold. Standard errors are clustered by province. Q1-Q5 indicate quintiles of the municipality median income distribution; "No Inc" denotes the group of stations in towns too small to be included in the data.

Table A1: *Average tax pass-through by income group – region-level clusters*

	No Inc (1)	Q1 (2)	Q2 (3)	Q3 (4)	Q4 (5)	Q5 (6)
Panel A. Sales-weighted						
Average pass-through	0.865	0.893	0.960	0.937	0.952	0.953
	(0.015)	(0.010)	(0.013)	(0.007)	(0.008)	(0.003)
N	14,347	5,763	8,647	6,965	10,322	5,905
Panel B. Unweighted						
Average pass-through	0.887	0.902	0.952	0.951	0.955	0.950
	(0.006)	(0.006)	(0.012)	(0.009)	(0.008)	(0.005)
N	15,025	5,763	8,647	6,965	10,640	5,905

Notes: Estimates are average pass-through rates of tax changes into retail prices, $\frac{\Delta P}{\Delta t}$, obtained via Equation 7, for each of six income groups. Panel A displays results with sales-weighted observations; Panel B results rely on unweighted observations. The post-tax estimation window is 12 months. Observations are weighted by annual station (or municipality) sales. Standard errors are clustered by region, as opposed to the preferred level of province. Q1-Q5 indicate quintiles of the municipality median income distribution; "No Inc" denotes the group of stations in towns too small to be included in the data. 'N' reports the number of treatment-group station-months used.

Table A2: *Income and tax pass-through, per capita measure*

	No inc	Q1	Q2	Q3	Q4	Q5
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Sales-weighted						
Average pass-through (%)	0.865***	0.942***	0.956***	0.952***	0.953***	0.925***
	(0.049)	(0.010)	(0.016)	(0.013)	(0.012)	(0.010)
N	14,347	6,637	7,935	9,013	8,447	5,127
Panel B. Unweighted						
Average pass-through (%)	0.887***	0.931***	0.957***	0.934***	0.961***	0.944***
	(0.011)	(0.011)	(0.013)	(0.016)	(0.011)	(0.010)
N	15,025	14,347	7,935	9,013	8,447	5,657

Notes: Point estimates are average pass-through rates of tax changes into retail prices, estimated via Equation 7, separately for each income quintile ($Q1$-$Q5$) as well as the group of stations without income data ("No Inc""). Income quintiles are defined according to per capita income instead of median income. Panel A displays results with station-month observations weighted by annual sales; Panel B displays results without weights. The post-tax estimation window is 12 months. 'N' reports the number of treatment-group station-months used. Observations are weighted by annual station sales. Standard errors are clustered by province.

Table A3: *Income versus other local attributes in pass-through prediction*

Tax X...	(1)	(2)	(3)	(4)
1[Income quintile 1]	0.876***	0.761***	0.744***	0.800***
	(0.037)	(0.051)	(0.072)	(0.038)
1[Income quintile 2]	0.061	0.057	0.058	0.046
	(0.035)	(0.031)	(0.031)	(0.032)
1[Income quintile 3]	0.124**	0.124**	0.124**	0.092*
	(0.044)	(0.036)	(0.037)	(0.039)
1[Income quintile 4]	0.111*	0.107**	0.108**	0.096*
	(0.043)	(0.037)	(0.038)	(0.042)
1[Income quintile 5]	0.093*	0.090*	0.091*	0.069
	(0.044)	(0.040)	(0.040)	(0.037)
1[No income data]	-0.044	-0.063	-0.066	0.016
	(0.060)	(0.055)	(0.052)	(0.033)
1[Refiner branded]		0.067	0.069	0.105*
		(0.057)	(0.056)	(0.040)
1[Wholesaler branded]		-0.010	-0.020	0.023
		(0.060)	(0.054)	(0.031)
# of rivals, distance-weighted		-0.012	-0.011	-0.014
		(0.010)	(0.010)	(0.007)
Own-brand share of local stations		0.088**	0.088**	0.047*
		(0.027)	(0.027)	(0.018)
1[Vertically integrated]		0.107*	0.122*	0.057
		(0.051)	(0.058)	(0.039)
Population density (1000s)		0.003	0.003	0.004
		(0.004)	(0.004)	(0.003)
1[Carwash]			-0.030*	-0.005
			(0.012)	(0.008)
1[Tires and fluids]			0.033	0.009
			(0.033)	(0.009)
1[Convenience store]			0.002	-0.028*
			(0.022)	(0.013)
1[Cafeteria]			0.018	0.007
			(0.019)	(0.015)
N	461,976	461,884	461,884	529,990
Weighted	Y	Y	Y	N

Notes: Point estimates are average pass-through rates of tax changes into retail prices, estimated via Equation 3. Rows indicate interaction terms between the regional tax level and the listed variable. 'N' is the number of station-months used; observations are weighted by annual station (or municipality) sales. Standard errors are clustered by province.