



HAL
open science

Compensating against fuel price inflation: Price subsidies or transfers?

Odran Bonnet, Étienne Fize, Tristan Loisel, Lionel Wilner

► **To cite this version:**

Odran Bonnet, Étienne Fize, Tristan Loisel, Lionel Wilner. Compensating against fuel price inflation: Price subsidies or transfers?. *Journal of Environmental Economics and Management*, In press. hal-04799412

HAL Id: hal-04799412

<https://hal.science/hal-04799412v1>

Submitted on 22 Nov 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Compensating against fuel price inflation: Price subsidies or transfers?

Odran Bonnet Étienne Fize Tristan Loisel Lionel Wilner*

November 22, 2024

Abstract

Compensating agents against substantial and sudden shocks requires both targeting tax policies and taking behavioral responses into account. Based on transaction-level data from France, this article exploits quasi-experimental variation provided by 2022 fuel price inflation and excise tax cuts. After disentangling anticipation from price effects, we estimate a price elasticity of fuel demand of -0.31, on average, which varies little with respect to income and location but substantially decreases with fuel spending, in absolute value. Using targeted transfers only achieves imperfect compensation, yet a budget-constrained policy-maker seeking to alleviate excessive losses relative to income prefers income-based transfers to price subsidies.

Keywords: Commodity taxation; Excise fuel tax; Tax-and-transfer schemes; Gasoline price elasticity; Anticipatory behavior; Transaction-level data.

JEL Classification: D12; H23; H31.

*Bonnet, Loisel: Insee. Fize: Institut des Politiques Publiques, Paris School of Economics. Wilner: CREST, ENSAE Paris, Institut Polytechnique de Paris. Please address correspondence to lionel.wilner@ensae.fr. Address: 5 avenue Henry Le Châtelier, 91120 Palaiseau, France. Tel: (+33)170266947. We thank participants of AFSE (Paris 2023), ESEM (Barcelona 2023), IAAE (Oslo 2023), IPDC (Amsterdam 2023), JMA (Strasbourg 2023), and SEHO (Copenhagen 2023) conferences as well as audiences at CRED Paris 2 (GT Fiscalité seminar), Franco-German Fiscal Policy Seminar, ifo Lunchtime Seminar (Munich), Insee, Université Paris Dauphine-PSL (LEDa seminar). We are also grateful to Pierre Boyer, Mathias Dolls, Xavier D'Haultfoeuille, Clemens Fuest, Romain Gaté, Erwan Gautier, Daniel Herrera-Araujo, Laurence Jacquet, Xavier Jaravel, Etienne Lehmann, Thierry Magnac, Volker Meier, Felix Montag, Stefan Pollinger, Aurélien Saussay, Ao Wang, and four anonymous referees for useful suggestions, as well as to the *Groupement des Carte Bancaires* (GIE-CB), Corinne Darmaillac, Thomas Laurent and Pierre Leblanc at Insee for providing with daily-level transaction data used to adjust for seasonality. We thank Marine Adam and Marion Rault for outstanding research assistance. We are extremely grateful to Crédit Mutuel-Alliance Fédérale for sharing the data with us, and in particular to key employees for their precious help. All individual data used in this analysis have been anonymized and no single customer can be traced in the data. All data processing has been conducted following the bank's strict data privacy guidelines.

1 Introduction

“These tax cuts are a mistake.
Cutting fuel taxes is a bad idea.
Governments should support household incomes instead. ”
The Economist, Mar 26th 2022.

In 2022, explicit fuel subsidies amounted to \$1.3 trillion, according to the IMF.¹ Following the fuel price inflation precipitated by the invasion of Ukraine, approximately +30% within a few days, policy responses worldwide have mostly consisted of helping people buy cheaper petrol. Examples include EU countries (Germany, France, Italy, and Sweden, among others) as well as the UK, Japan, West Africa, Cuba, and several US states. In France, the government offered excise tax rebates, which is equivalent to subsidizing prices at the pump. This public intervention was intended to provide financial relief to motorists. Price subsidies were nonetheless removed at the end of the year owing to concerns about their costs, regressivity, and poorly aligned environmental incentives. The government decided to replace them with income-based transfers granted to car owners in the bottom half of the income distribution.

This article addresses the challenge that governments face when seeking to alleviate welfare losses that accompany substantial and sudden fuel price increases. We empirically compare the efficiency and equity of two competing policy approaches, price subsidies (rebates) and direct transfers to motorists.² Many environmental and public economists think about rebates versus transfers as trading off targeting against externality gains, which can be thought of as a trade-off between distributive and environmental concerns. Here, we study a distinct trade-off, which exists without considering externalities: the one that arises between financial and distributive concerns for a policy-maker that seeks to alleviate losses incurred by households with a high fuel budget share. On the one hand, price subsidies may completely offset the impact of fuel price inflation by providing full compensation to households, while imperfectly targeting households with a high fuel budget share. On the other hand, transfers achieve only partial compensation, but may be more effective in

¹Explicit subsidies refer to undercharging for supply costs. Implicit subsidies refer to undercharging for environmental costs and forgone consumption taxes, after accounting for preexisting fuel taxes and carbon pricing. Including implicit fuel subsidies yields \$7 trillion (<https://www.imf.org/en/Blogs/Articles/2023/08/24/fossil-fuel-subsidies-surged-to-record-7-trillion>).

²For a similar analysis in the UK with regard to the energy crisis, see [Levell et al. \(2024\)](#).

targeting those households. Both instruments are equivalent when fuel demand is perfectly inelastic or when price shocks are small (the burden of fuel price inflation being simply the product of price change and initial fuel purchases). However, they differ as soon as behavioral responses are taken into account and in the presence of non-marginal price changes: the envelope theorem does not apply. Transfers allow agents to reoptimize their purchase behavior with regard to not only fuel but also goods and services other than fuel; price subsidies do not. Price subsidies have been criticized for being costly and favoring high-income individuals.³ In contrast, transfers can be targeted toward specific groups that suffer the most from fuel price increases and that the policy-maker is willing to help overcome that shock. As noted by [Sallee \(2022\)](#), transfers are all the less attractive when (i) it is difficult to make an accurate prediction of fuel demand based on consumers' observable characteristics, (ii) the behavioral response is small, (iii) environmental externalities are small, and (iv) distributive concerns are absent.

We investigate the situation observed in France, where fuel prices rose by 33% between 2021 and 2022. We first estimate short-term household-level price elasticities to serve as sufficient statistics for welfare loss, which depend on the shape of Hicksian demand and, hence, on compensated price elasticity. If that elasticity is high, consumers can substitute away from fuel, thereby alleviating their welfare loss. Transfers are then more attractive and cheaper than price subsidies. Our approach consists of estimating the uncompensated price elasticity and deriving the welfare loss on the basis of some parameterized utility function. We next simulate various compensatory mechanisms and count how many 'losers' still incur some welfare loss after compensation. Using high-frequency data, this article illustrates how the theoretical framework pioneered by [Sallee \(2022\)](#) can be brought into an observational context of compensation offered to households following sudden fuel price increases. In particular, our work empirically assesses the condition under which compensatory and targeted transfers may lead to a Pareto improvement following fuel price inflation that

³A quantified assessment of that statement can be found in [Appendix J](#), which provides a policy evaluation exercise.

harms households. Indeed, condition 2 in [Sallee \(2022\)](#)⁴ states that if targeting errors, which refer to the gap between welfare loss⁵ and compensatory transfer, are large relative to efficiency gains, compensatory transfers cannot result in any Pareto improvement.

For this purpose, we characterize consumer behavior in response to exogenous and policy-driven price changes. Our empirical analysis is based on a quasinatural experiment involving a tax policy designed to temper fuel price inflation following the invasion of Ukraine. We use high-frequency, transaction-level data issued from bank accounts, which includes timestamped operations from September 2021 to February 2023. Daily data enable us to disentangle anticipation effects properly from the pure price effect. We show that credible estimates of short-term price elasticity require (i) an appropriate source for identifying variability, (ii) high-frequency data that permit first visualization and then neutralization of daily anticipation effects, and (iii) a suitable econometric approach. The lack of one of these criteria results in identification failure. The method adopted here to separate anticipation effects from price effects easily extends to other public finance contexts. In capital taxation, for example, public announcements with regard to wealth taxes likely induce strong anticipatory behavior, and bunching patterns similar to those observed here are expected due to the high mobility of assets. It in fact applies to many other settings (see, e.g., [Mertens and Ravn, 2012](#)) involving any change in commodity taxation, including VAT, or in fiscal incentives, such as ecological subsidies, bonus/malus schemes related to green taxation, and other transfers that may be provided when buying electrical equipment or undertaking home renovations.

We obtain an average elasticity of -0.31⁶ and estimate an anticipation bias of approximately -0.4. Using machine learning methods, we document the heterogeneity of the estimated elasticities. Surprisingly, price sensitivity varies little across income and location; however, it exhibits sizeable dispersion with respect to past fuel spending. Motorists who

⁴This condition is: 'Let c_i be the initial costs from a policy assuming no behavioral responses, N be the number of agents, X_i be a vector of exogenous covariates observed by the planner, $T(\mathbf{X}_i)$ be a transfer scheme, $\bar{\Delta}$ the average funding gap, \bar{g} be the average efficiency gain, and \bar{b} be the average private welfare gain from behavioral responses. If the average absolute targeting error exceeds twice the sum of the average efficiency gain and behavioral adjustment gain minus the average funding gap; i.e., $\frac{1}{N} \sum_i |c_i - T(\mathbf{X}_i)| > 2(g + \bar{b}) - \bar{\Delta}$ then there is no distribution of g with $g_i \geq 0 \ \forall i$ for which the policy and transfers create a Pareto improvement'.

⁵The welfare loss before any transfer but also before any behavioral response has been taken into account corresponds to the increase in fuel spending given the original level of consumption, i.e., to the above burden.

⁶That elasticity may differ from the long-run price elasticity since it takes time for consumers to adjust their response, for example, by shifting to other types of vehicles or transportation.

drive very frequently are more inelastic, especially the most liquidity-constrained among them. Retirees are more elastic than employees are. In addition to being useful to the policy-maker to take agents' behavior into account when designing public interventions, these results help improve the targeting of fuel excise taxes and solve the “prediction problem” discussed by [Sallee \(2022\)](#). Specifically, an accurate prediction of fuel consumption on the basis of observed characteristics is required to better tailor transfers to reduce the number of losers that inevitably accompanies any policy designed to increase efficiency. Observing the distributions of fuel spending across households with similar observable characteristics enables us to go a step further and empirically quantify the magnitude of that problem.

Finally, we compare price subsidies and transfers. We first adopt a positive approach in which we compute the number of losers associated with each policy instrument. Targeting transfers yields a relatively high fraction of losers even when detailed high-frequency information on consumer purchases is available; this is because much unobserved heterogeneity in fuel spending prevails even after conditioning on observables such as income or location. Although the magnitude of the behavioral response is moderate (in the short run, consumers are not able to substitute away from fuel), it matters in the decision-making process; on the basis of compensating variations, transfers are slightly cheaper (5%) than price subsidies.

In a normative approach, we next compare the efficiency and equity of policy tools when the decision-maker seeks to alleviate consumer welfare loss, expressed as a share of income (relative loss hereafter). We simulate various counterfactual compensatory policies, namely, unconditional or income-based price subsidies, as opposed to transfers that may be lump sum, income-based, conditional on location or fuel spending. Income-based price subsidies achieve a better job with respect to distributive objectives and at a lower cost. Such targeted subsidies may nevertheless be infeasible due to implementation considerations. Transfers achieve partial compensation, which hurts numerous losers.⁷ However, a countervailing force to the smaller compensatory power of transfers is at play. Under scarce pecuniary resources for compensation, income-based transfers are more effective at targeting the right households given distributive objectives. Income acts as a screening device

⁷This helps explain why political support for revenue-recycling tax-and-transfer schemes, i.e., compensatory mechanisms that have long been summoned by economists, is so difficult to obtain ([Sallee, 2022](#); [Douenne, 2020](#); [Young-Brun, 2023](#)), although latter schemes are effective at canceling out regressivity.

for relative loss; specifically, many losers are high-income households for whom the relative loss is quite small. When more public funding becomes available, price subsidies provide full compensation, which makes them more desirable but still benefits most of those who drive more, essentially high-income households.

Literature A vast amount of empirical literature, surveyed, for instance, in [Dahl and Sterner \(1991\)](#), [Espey \(1998\)](#), and more recently [Kilian \(2022\)](#), has been devoted to measuring the short-run price elasticity of fuel demand. However, those studies have thus far mostly relied on disaggregated data and, hence, faced two fundamental problems: simultaneity and anticipatory behavior. The typical solution for simultaneity (see, e.g., [Borenstein et al., 1997](#); [Hughes et al., 2008](#); [Kilian and Zhou, 2024](#)) consists of finding instruments, typically predictors of prices, such as cost shifters that are as unrelated as possible with demand usage, such as the West Texas Intermediate (WTI) crude oil spot price, the price of the Brent, the average price in the market, and tax changes. With respect to anticipations, demand estimation has to take into account that when consumers foresee price changes, they may strategically alter their purchase behavior; specifically, intertemporal substitution possibly arises through dips (spikes) when future prices are expected to be lower (higher), which invalidates previous approaches. The current quantity then depends not only on current prices but also on future prices, which violates the usual exclusion restrictions at play in both OLS and IV estimation methods. A typical solution to address this problem consists of introducing leads and lags, as in [Coglianese et al. \(2017\)](#) and [Kilian \(2022\)](#). Such approaches rely on parametric assumptions as regards the dependence of expectations with respect to future and past prices.

Compared to [Kilian and Zhou \(2024\)](#), who avail of monthly state-level data, our empirical analysis is based on high-frequency individual data.⁸ We also build upon that literature

⁸We also rely on the 2022 oil crisis and policy responses as instruments, and we are able to exploit the variation in observed characteristics at the household level to infer the heterogeneity of price elasticity. Another difference stems from our identification of variation resulting from uniform tax changes over the French territory, whereas their empirical strategy leverages state-level price shocks.

by quantifying the aggregation bias inherent to less granular datasets.⁹ This issue has long been a concern for applied econometricians when evaluating any tax change.

A few recent studies have used high-frequency data to estimate the price sensitivity of fuel demand, including [Levin et al. \(2017\)](#) and [Knittel and Tanaka \(2021\)](#), who disentangle driving behavior (extensive margin) from traveling distance (intensive margin). We improve upon their methodology by relying on an exogenous source of price variation. From that viewpoint, the closest paper to ours is [Gelman et al. \(2023\)](#), in which the authors exploit a dataset issued from banking accounts while relying on large, unexpected shocks. They mostly examine cross-price effects with spending other than fuel.¹⁰

Finally, our paper contributes to both academic and policy debates about fuel taxation, following the works of [Douenne \(2020\)](#), [Sallee \(2022\)](#), and [Young-Brun \(2023\)](#). To the best of our knowledge, ours is the first empirical paper that compares price subsidies and transfers as policy tools in terms of compensating power against fuel price inflation, based on financial and distributional objectives. Quantifying the comparative advantage of transfers is made possible by our granular individual data, which enables us to go further than annual panels in the estimation of short-term response but also to accurately recover the distribution of annual fuel consumption across households. In contrast, consumption surveys rely on data collected over a few weeks each year; corresponding estimates would spuriously suggest that the share of households purchasing no fuel is excessively high, for example.

The rest of the article is organized as follows. Section 2 presents our data and the institutional background. A simple conceptual framework that illustrates how anticipation effects can be disentangled from price effects is presented in Section 3. The empirical analysis from Section 4 includes our identification strategy as well as our econometric specification. Section 5 contains our results, including an investigation of the heterogeneity

⁹[Levin et al. \(2017\)](#) noted that low-frequency data suffer from such aggregation bias for three distinct reasons: (i) a common price coefficient is imposed while price sensitivities may be heterogeneous, e.g., at some spatial unit (city) level; (ii) less granular datasets do not allow the inclusion of appropriate space or time fixed effects viewed as unobserved components, which permits the removal of any supply-driven variation induced by fluctuations in the price over time; and (iii) aggregation may induce some correlation between average prices and the error term, e.g., when the correlation between prices and demand shocks on other days, or in other cities, causes prices and errors in the aggregated panel data to be correlated.

¹⁰Another empirical difference lies in their analysis being based on large but continuous price changes; in contrast, we leverage sharp, sudden variations arising at publicly known dates. We therefore view our identification strategy as a complement to theirs; moreover, we explicitly adopt an econometric specification that controls for anticipatory behavior.

of price elasticity. Section 6 is devoted to counterfactual simulations that compare actual price subsidies to alternate policies, including transfers. Section 7 provides a conclusion.

2 Data and context

We first present our anonymous bank account data. Our database is issued from the *Crédit Mutuel Alliance Fédérale*, a French group of banks with more than 30 million customers, firms or households. The construction of key variables follows a recent strand of literature exploiting such data, including Baker (2018), Ganong and Noel (2019), Andersen et al. (2023), Boehm et al. (2023) and Bounie et al. (2020), which use the same data source. We use transaction-level data on credit and debit card payments,¹¹ paper checks, cash withdrawals, cash deposits, bank transfers, and direct debits. We base our analysis on a daily aggregation of each transaction, measured in euros. Balance sheets are also available each month. The unit of observation is a household. The data contain various sociodemographics on household members, such as age, sex, *département*,¹² family status, occupation, and type of location (in three categories: urban, rural, and semiurban areas).

We define total spending as the sum of outgoing transactions issued by a card. We measure disposable income as the sum of monthly incoming transfers, up to a €40,000 threshold. Liquid assets are proxied by the sum of balances on different bank accounts (deposit account and savings accounts) and provide us with a measure of liquid wealth. Illiquid assets are equal to the sum of balances on life insurance, stocks, bonds, mutual funds and certificates of deposits. In France, banks are not in charge of retirement savings plans.

Working sample Our estimation period runs from September 2021 to February 2023. Our initial raw data are a sample of approximately 300,000 households; this sample is stratified by *départements* of metropolitan France and by 5-year age dummies. To alleviate concerns about representativeness, we proceed with calibration weighting via the

¹¹In France, the use of credit cards is scarce, accounting for less than 10% of bank cards.

¹²*Département* is an administrative division of France, one of the three levels of government under the national level between the administrative regions and the municipalities. It also corresponds to the EU NUTS-3 classification and is roughly similar to a county in the U.S. Mainland France (namely, France at the exclusion of Corsica and overseas) is divided into 94 *départements*. Metropolitan France includes two Corsican *départements*.

method proposed by [Deville and Särndal \(1992\)](#) (see [Appendix B](#) for details) and weight our estimating equations via calibration weights. We further restrict our attention to households with the same number of adults (aged at least 18) over the period. We focus on customers who spend at least €150 during three rolling months, either by card or in cash. Moreover, we impose that customers be present and purchase fuel in France at least once over the observation period, in addition to meeting previous criteria throughout the period. We remove households that consume most of their fuel abroad since we are unable to track prices and quantities for these consumers. Factoring in attrition since June 2020, this selection leaves us with approximately 180,000 active customers.

Fuel spending Our bank account data provide the Merchant Category Code (MCC) classification. Based on that taxonomy, we consider that spending categorized with codes 5541 and 5542 corresponds to fuel spending, in accordance with [Andersen et al. \(2023\)](#) and [Gelman et al. \(2023\)](#).¹³ [Figure I1](#) displays the distribution of the amount of a transaction, in euros, which seemingly mixes a continuous distribution, the mode of which is nearly 55 euros, and a discrete distribution over round numbers (typically 30, 20, 50 or 40 euros, as well as other multiples of 5 euros).¹⁴ [Figure I2](#) further shows that the median interpurchase duration, i.e., the time interval between two visits at the pump, is 7 days.

Finally, we obtain fuel quantity, in liters, as the ratio of that adjusted fuel spending over a fuel price index.¹⁵ In our sample, [Figure I3](#) confirms that fuel expenditures increase with income but that the budget share devoted to fuel decreases along that dimension. Rural households devote a greater share of their budget to fuel ([Figure I4](#)), which is consistent with the idea that the burden of fuel consumption is mostly borne by poor and rural households, who may depend more on cars than other types of transportation and are likely more constrained. In addition, annual fuel spending, expressed as a share of income,

¹³We restrict our attention to card payments, excluding cash, for instance, which sounds like a benign empirical choice, as the distribution of monthly fuel spending with respect to income looks close to that obtained from the *Budget des familles* representative survey; see [Figure B4](#) in the Appendix.

¹⁴The latter most likely stems from the possibility of prepaying gas at some stations, and we verify below (last row of [Table 4](#)) that this specificity does not dramatically change the picture of our estimated price sensitivity by isolating that dimension, namely, purchases corresponding to nonautomated fuel dispensers (MCC: 5541).

¹⁵Cross-border purchases, which result from trade-offs that involve, in particular, the distance to the frontier and the ratio of foreign prices to domestic prices, are excluded from the current analysis. More precisely, we exclude individuals living close to a border, which we identify in the data as soon as they purchase fuel abroad (excluding holidays). Foreign transactions occurring during holidays are also removed from the subsequent analysis.

considerably varies within income and location groups (Figure I5). For instance, the fuel budget share of many urban households exceeds that of rural households, and the fuel budget share of many high-income households exceeds that of low-income households. Such an heterogeneity makes effective targeting difficult.

Prices Timestamped and geolocated fuel prices are disclosed at the gas station level via a French governmental website.¹⁶ Such data have already been used by researchers, e.g., Montag et al. (2021) or Gautier et al. (2023). The website contains information on each and any price change for different types of fuel (e.g., diesel and gasoline: super unleaded petrol (SP95), super unleaded petrol (SP95–E10), and super unleaded petrol (SP98)). In the subsequent analysis, we focus on two types of fuel, diesel and standard gasoline, namely, SP95–E10, given that the latter exhibits similar variations over time as SP95 and SP98 (over the period considered, the correlation is greater than 0.99, which mitigates any concern about substitution between those products). We assume that individuals are unlikely to change the type of fuel consumed in the short run because such an adjustment would presumably imply a more substantial change, that of a vehicle. The data also provide an identifier and the location of each retailer. More details on how we recover prices can be found in Appendix B.2.

Information on the type of fuel actually purchased, diesel or gasoline, is unavailable in the dataset, which is not important when both prices similarly covary. Empirically, those prices are highly correlated; their Pearson coefficient is 0.8 throughout the entire observation period (see Figure I6). Diesel and gasoline prices sometimes experience different short-run variations because specific conditions affect the oil market. We therefore build a fuel price index that weighs diesel and gasoline prices differently within a *département* according to strata on the basis of observed household characteristics.¹⁷ According to the *Enquête Mobilité* survey,¹⁸ these characteristics are good predictors of the type of fuel purchased; we thus attribute a weight to diesel in the fuel mix at the stratum level on the basis of that survey (0.66 on average).

¹⁶<https://www.prix-carburants.gouv.fr/rubrique/opendata/>.

¹⁷The characteristics that define strata are as follows: income (in four groups), type of location (urban, rural, or semiurban), age group (less than 30, 30–60, or more than 60), and 2019 fuel spending (in four groups).

¹⁸Detailed results of this survey are available online at <https://www.statistiques.developpementdurable.gouv.fr/resultats-detailles-de-lenquete-mobilite-despersonnes-de-2019>.

Admittedly, transaction prices are measured with error since we ignore the exact location of purchase; hence, our price index is computed from daily average values in the *département*. We perform sensitivity analyses with respect to the chosen price index and consider various alternate measures of prices, including the price of diesel, the most commonly used type of fuel in France, or the first decile (instead of the average) of the price distribution in each *département*. Moreover, we use instrumental variables in our empirical analysis, which partly addresses that issue.

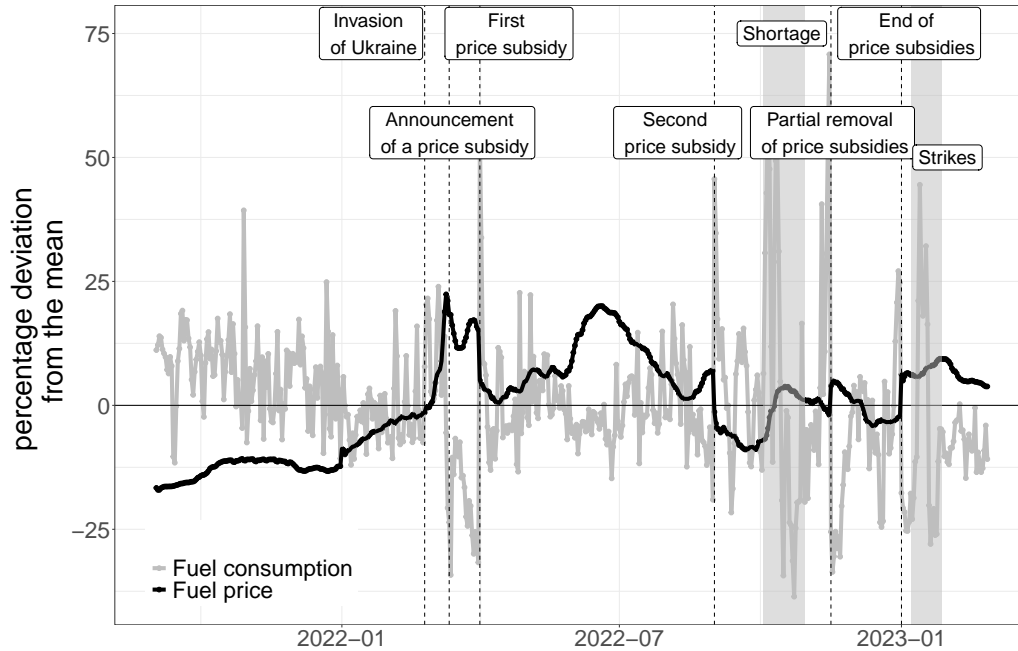
Context: invasion of Ukraine, energy crisis, and policy responses Fuel prices have experienced substantial variations from 2019 to 2023, especially in 2022, due to an oil price surge following the invasion of Ukraine that started on February 24. In France, the government decided to directly subsidize prices at the pump. Since per unit excise taxes represent approximately 60% of fuel prices, the public intervention in fact consists of offering a tax rebate.

On March 12, Prime Minister Castex made an official announcement that the before-tax gasoline price would be lowered by €0.15 per liter from April 1 onward (approximately €0.18 per liter including VAT, with some minor geographic variations due to *département*-specific VAT rates). While this first public intervention was bound to last until the end of Summer 2022, the Parliament decided to extend it to the beginning of October in response to the energy crisis. As announced by Prime Minister Borne at the end of July, a total fuel price subsidy of €0.3 per liter was then effective on the after-tax price from September 1 onward, resulting in an extra €0.12 subsidy for each liter purchased. The second price reduction prevailed until mid-November 2022, when the price subsidy was reduced to €0.1 per liter, before its complete removal on January 1, 2023. Note also that before the implementation of those substantial price cuts, prices were already increasing at a high pace, even before the invasion of Ukraine.

According to the government, price subsidies were indeed granted to fuel suppliers.¹⁹ They were then passed through to gas stations, and ultimately reflected in the price paid by consumers. Upon arriving at a gas station, consumers saw fuel prices already discounted. As a result, they paid the reduced fuel price. Given the high level of competition in

¹⁹See official press release on March 26, 2022: https://presse.economie.gouv.fr/26-03-2022-modalites-de-mise-en-oeuvre-de-la-remise-de-15-centimes-deuro-par-liter-pour-lacquisition-#_ftn1.

Figure 1: Fuel prices and purchases



Note. Fuel prices (including taxes) and purchases (in liters). Purchases are adjusted for seasonal variations via GIE-CB data from September 2021 to February 2023. The dashed lines correspond to the invasion of Ukraine and policy interventions. The first intervention on April 1 is a €0.18 per liter tax rebate, including VAT. The second intervention on September 1 is an extra €0.12 per liter subsidy, which prevailed until mid-November 2022. The residual price subsidies were removed on January 1, 2023. *Source.* Sample of households that use banks primarily at *Crédit Mutuel Alliance Fédérale*.

the market, most players were already encouraged to pass through the price subsidy to avoid losing market shares. Additionally, retailers and suppliers have publicly signed a commitment to pass through the price subsidy. The observed evolution of prices suggests that this agreement was adhered to. The effective pass-through is close to 100% and quite homogeneous (see Appendix D for details).

In the following, we rely on policy-driven fuel price reductions as the primary source of variability to infer price effects causally. We view the two public interventions of April 1 and September 1 as quasinatural experiments. Figure 1 suggests that the evolution of fuel demand at the time of announcement is consistent with anticipatory behavior; people

strategically refrain from buying once they are aware of lower prices in the future.²⁰

3 Theoretical framework

To illustrate how the price effect can be separated from anticipation effects in a dynamic setting, we leverage a simple conceptual framework, namely, a stylized inventory model of fuel stockpiling behavior. We then explain how this setting can shed light on the empirical analysis, especially with respect to identification.

3.1 A stylized inventory model of fuel stockpiling behavior

Let a single agent maximize her intertemporal utility with respect to her fuel consumption c and her outside good m , the *numéraire*, given her intertemporal budget constraint and the law of motion of fuel inventory i .²¹ The period here is typically a day or a week. We thus assume that fuel, when stockpiled, does not depreciate and consider a discount factor equal to one along with a zero interest rate.²² Denoting the instantaneous utility derived from consumption by $u(\cdot)$, the benefit of income by $b(\cdot)$, fuel purchases by q , fuel prices by p , permanent income by Y , and storage costs by $C(\cdot)$, the agent's maximization is as follows:

$$\max_{(c,i,m)} \sum_{t=0}^T [u(c_t) + b(m_t) - C(i_t)] \quad \text{s.t.} \quad \sum_{t=0}^T [p_t q_t + m_t] \leq Y \quad (1)$$

$$i_t \leq i_{t-1} + q_t - c_t \quad (2)$$

Under the assumption of no fuel waste, the law of motion of the fuel inventory binds at all periods:

$$\max_{(c,i,m)} \sum_{t=0}^T [u(c_t) + b(m_t) - C(i_t)] \quad \text{s.t.} \quad \sum_{t=0}^T [p_t c_t + p_t (i_t - i_{t-1}) + m_t] \leq Y \quad (3)$$

²⁰We exclude October 2022 from our sample due to strikes in refineries, leading to shortages in various places. We also exclude periods with shortages from January 7 to January 27, 2023; France experienced many refinery blockades at that time, which were related to the social movement caused by the 2023 retirement reform.

²¹The model is a simplified version of an inventory model used, e.g., by [Hendel and Nevo \(2006\)](#).

²²Taking the tank's capacity constraint into account would not dramatically alter the conclusions of the model.

Let us denote by λ the Lagrange multiplier of the intertemporal budget constraint. In the optimum, this constraint is binding, and λ is equal to the marginal benefit of income ($\lambda = 1$ when preferences are quasilinear). First-order conditions yield the Euler equation and the optimal fuel inventory:

$$C'(i_t) = \lambda(p_{t+1} - p_t). \quad (4)$$

Stockpiling behavior is governed by expected price changes. Introducing the parameter $\theta > 0$, we assume that storage costs $C(i_t) = \theta i_t^2$ are a quadratic function of inventory, which leads to:

$$i_t = \lambda \frac{p_{t+1} - p_t}{2\theta}. \quad (5)$$

We also consider a quadratic utility function of the form $u(c_t) = c_t - \alpha c_t^2$, with $\alpha > 0$, which yields a linear demand:

$$c_t = \frac{1 - \lambda p_t}{2\alpha}. \quad (6)$$

Note that we allow for risk aversion with regard to consumption, but not with regard to prices -an interesting extension. From equations (5), (6), and the binding law of motion for fuel inventory, the model predicts that the observed purchases are as follows:

$$q_t = c_t + \lambda \frac{p_{t+1} + p_{t-1} - 2p_t}{2\theta} = \frac{1}{2\alpha} - \lambda \left(\frac{1}{2\alpha} + \frac{1}{\theta} \right) p_t + \frac{\lambda}{2\theta} p_{t-1} + \frac{\lambda}{2\theta} p_{t+1}, \quad (7)$$

hence a specification of the form:

$$q_t = q_0 + \beta p_t + \gamma_t \max(\mathbb{1}_{p_{t-1} \neq p_t}, \mathbb{1}_{p_{t+1} \neq p_t}) \quad (8)$$

where $\beta = -\frac{\lambda}{2\alpha}$. The coefficients $\gamma_t = \lambda \frac{(p_{t+1} - p_t) - (p_t - p_{t-1})}{2\theta}$ account for anticipation effects that are nonzero as soon as prices fluctuate and that alter current purchases. Anticipation effects affect the timing of purchase but not the quantity consumed c_t . They vanish (are exacerbated) when θ tends to $+\infty$ (0), i.e., when storage is impossible (not costly) because the tank is full, for example. Since they are equal to the change in inventory over time ($\gamma_t = i_t - i_{t-1}$), they sum to $i_T - i_0$, i.e., the difference between the final and initial inventory levels. From equation (5), this difference is null when prices are constant both at the beginning and at the end of the period considered. Our identifying assumption can thus


be stated as follows. *When a shock occurs within that period, anticipation effects sum to zero provided that prices remain stable away from that shock.* In contrast, prices determine the total amount of fuel purchased. This specification provides a microfoundation for the econometric model presented in section 3.2 as well as for our estimating equation (10).²³

That observed purchases depend on lags and leads of prices on top of current prices, as on the right-hand side of equation (7), is reminiscent of Coglianesse et al. (2017). Although stylized, this conceptual framework thus rationalizes any reduced-form approach that involves a regression of purchases on current, past, and future prices when aiming at recovering price elasticity.²⁴ The price effect $\beta = -\frac{\lambda}{2\alpha} = -\lambda \left(\frac{1}{2\alpha} + \frac{1}{\theta}\right) + \frac{\lambda}{2\theta} + \frac{\lambda}{2\theta}$ can also be retrieved from the sum of coefficients corresponding to current, past, and future prices.

3.2 Implications for identification

We now present an econometric approach derived from the model above, which empirically permits disentangling anticipatory behavior from price sensitivity. The specification applies to a single, anticipated price reduction such as the one experienced around September 1, 2022.²⁵

Figure 2: Price versus anticipation effects



²³Here, the anticipation effects γ_t are shaped by parametric assumptions made in the inventory model; in the agnostic approach below, these coefficients are neither derived from nor subject to such restrictions.

²⁴As equation (5) makes clear, inventory behavior is the driving force of such an empirical approach.

²⁵In Appendix E.2, we provide a second specification that corresponds to an anticipated price surge followed by some expected price reduction of the same magnitude, as was the case in March and April 2022.

Table 1: Fuel prices and purchases around September 1, 2022

	07-15 to 08-17	08-18 to 08-31	09-01 to 09-14	09-15 to 10-04
Period	1	2	3	4
Liters per day (2021)	2.818	2.828	2.667	2.609
Liters per day (2022)	2.529	2.363	2.682	2.430
Price in € (2021)	1.493	1.481	1.490	1.514
Price in € (2022)	1.889	1.873	1.701	1.645
Length of period (days)	33	14	14	19

A simplified version of the second excise tax rebate is described in Figure 2. We consider a 4-period model whereby periods are indexed by t .²⁶ In the first two periods, prices are equal to their regular level p . Denoting the price subsidy, or rebate, by $\Delta p > 0$, prices fall to $p - \Delta p$ following the policy intervention between periods 2 and 3. Our identifying assumption is that *prices remain stable between periods 1 and 2, as well as between periods 3 and 4*. Put differently, prices do not vary much away from the policy intervention. Equipped with observed prices and purchases, the researcher considers the linear model:

$$q_t = q_0 + \beta p_t + \gamma_2 \mathbb{1}_{t=2} + \gamma_3 \mathbb{1}_{t=3} + u_t. \quad (9)$$

To make an explicit link with equation (8), the inventory model imposes supplementary constraints: $\gamma_2 = -\lambda \frac{\Delta p}{2\theta} < 0$ and $\gamma_3 = \lambda \frac{\Delta p}{2\theta} > 0$.²⁷ The anticipation effects perfectly compensate over the *anticipation window* composed of periods 2 and 3: $\gamma_2 + \gamma_3 = 0$. Regardless of the structure imposed by previous behavioral assumptions, but consistent with observations (Table 1), we expect $q_2 < q_1 < q_4 < q_3$. First, $q_1 < q_4$ due to the price effect since $(p - \Delta p) < p$. Second, $q_2 < q_1$ due to the strategic delay of purchases in period 2. Third, $q_3 > q_4$ because people who refrained from buying in period 2 purchased at lower prices in period 3. Anticipation effects explain both $q_2 < q_1$ and $q_3 > q_4$. Moreover, those confounders affect the timing of purchase but have no impact on total purchases: $q_2 + q_3 = q_1 + q_4$.

The ‘naive estimator’ $\hat{\beta}^n = \frac{q_1 + q_2 - q_3 + q_4}{2\Delta p}$ omits anticipation effects, *de facto* imposing that $\gamma_2 = \gamma_3 = 0$ (see Appendix E.1.3 for details). By definition, that estimator does not permit separating anticipation effects from the price effect: it relies on (q_2, q_3) to infer a

²⁶In Appendix E.1, we allow for periods with different lengths.

²⁷The model predicts that $i_1 = i_3 = i_4 = 0, i_2 = -\lambda \frac{\Delta p}{2\theta}$; hence, $q_2 = c_1 - \lambda \frac{\Delta p}{2\theta} < q_1 = c_1 < q_4 = c_4 < q_3 = c_4 + \lambda \frac{\Delta p}{2\theta}$. It follows that $\gamma_2 = q_2 - c_2 = -\lambda \frac{\Delta p}{2\theta} = q_2 - q_1, \gamma_3 = q_3 - c_3 = \lambda \frac{\Delta p}{2\theta} = q_3 - q_4$.

spurious price effect.

The ‘unconstrained estimator’ controls for what happens during the *anticipation window*: $\hat{\beta}^u = \frac{q_1 - q_4}{p_1 - p_4} = \frac{q_1 - q_4}{\Delta p}$ (cf. Appendix E.1.1), and therefore recovers the price effect net of any strategic effect. The anticipation effects are $\hat{\gamma}_2^u = q_2 - q_1 < 0$ and $\hat{\gamma}_3^u = q_3 - q_4 > 0$. By construction, this procedure entails a loss of information by discarding any contribution from periods 2 and 3.

The ‘constrained estimator’ $\hat{\beta}^c$, our preferred estimator, is obtained when the constraint $\gamma_2 + \gamma_3 = 0$ is imposed, consistent with previous arguments. We provide an empirical test of this constraint in section 4.2. In that stylized case, the ‘constrained estimator’ coincides with the ‘unconstrained estimator’ (see Appendix E.2.3).²⁸

Although simplified, this conceptual framework closely resembles the situation that prevails with regard to the second price subsidy. Equipped with the above model and observed prices and purchases (Table 1), as well as adjusting 2022 data for seasonality on the basis of 2021 observations as in section 4.1 below, we obtain an elasticity of -0.31 for both ‘constrained’ and ‘unconstrained’ estimators (see Appendix E.3 for numerical computations). The ‘naive estimator’ amounts to -0.68 . These figures are close to the econometric point estimates (section 5), and suggest that the anticipation bias is approximately -0.37 .

4 Empirical analysis

This section explains how observed variations in prices over time constitute a quasiexperimental setting that can be exploited to infer the price sensitivity of fuel demand. We rely on substantial price changes, some upward and unintentional²⁹ but also two downward, policy-driven price changes (the €0.18 per liter excise tax rebate from April 1, 2022, and the extra €0.12 per liter reduction in after-tax prices from September 1, 2022). In our empirical approach, we disentangle anticipation effects from the aversion to prices.

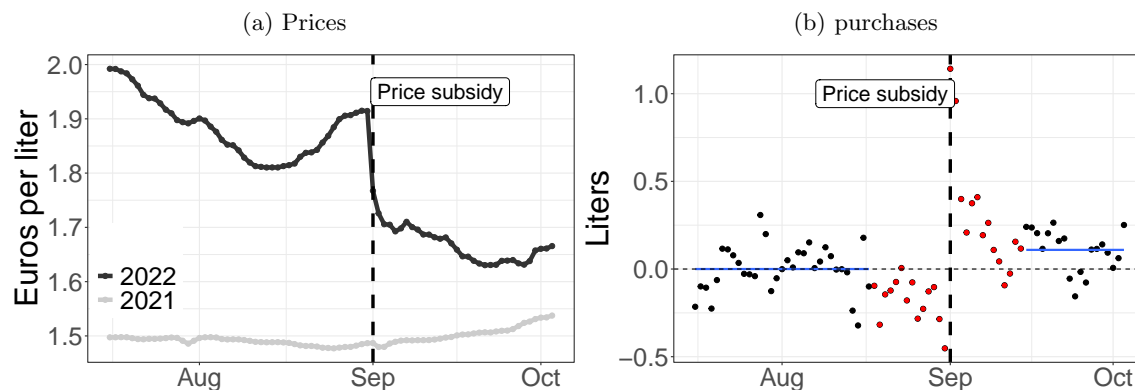
²⁸However, its independence from (q_2, q_3) results from both the peculiar price process considered here and the symmetry of the episode with respect to the moment when prices fall. In contrast, the unconstrained estimator $\hat{\beta}^u$ is always independent of (q_2, q_3) .

²⁹Those price changes are as high as +50% from September 2021 to the end of February 2022, +30% during the two weeks following the declaration of Ukrainian war, +20% in May-June 2022, +€0.2 per liter from mid-November 2022, and +€0.1 per liter from January 1, 2023 onward.

4.1 Identification strategy

Our main analysis relies on the extra €0.12 per liter price subsidy from September 1, 2022. This second public intervention enables us to identify anticipation effects after adjusting prices and purchases for seasonal variations, thereby relying on the year 2021 as the baseline. Our empirical setting (Figure 3a) is aligned with the previous theoretical framework. Prices are roughly stable both before and after the implementation of this additional price subsidy, which supports our identifying assumption. The reason why we cannot use a similar identification strategy with respect to the first public intervention in April 2022 is that we lack a credible baseline year for seasonal adjustment because lockdowns occurred in April 2020 and 2021.³⁰

Figure 3: Fuel prices and purchases around September 1



Note (left panel). Fuel prices (including taxes). The dashed line corresponds to the extra price subsidy of €0.12 per liter purchased, implemented from September 1, 2022.

Note (right panel). Dots correspond to adjusted daily fuel purchases from July 15 to October 4, 2022. The adjustment for seasonal variation relies on 2021 as the baseline year. The purchases are normalized so that they sum to 0 before the anticipation window. Dashed line: Extra price subsidy of €0.12 per liter starting on September 1, 2022. Red dots: Anticipation window (7 days before and after the implementation of the price subsidy). Blue lines: Average purchases before and after the implementation of the price subsidy, excluding the anticipation window.

Source. Sample of households that primarily bank at *Crédit Mutuel Alliance Fédérale*.

³⁰One could have relied on April 2019. Unfortunately, information about the MCC has been available in the data since July 2019 only. In France, COVID-19 policy-related measures included a first strict lockdown from March 17, 2020 to May 5, 2020, a second partial lockdown from October 29, 2020 to December 15, 2020, and a third partial lockdown from April 3, 2021 to May 3, 2021.

Previous empirical studies on fuel demand have proposed parametric solutions to address anticipation effects. Based on monthly-level data, [Coglianese et al. \(2017\)](#) use price lags and leads. We adopt a nonparametric approach, which makes full use of our high-frequency dataset. We identify daily anticipation effects during the anticipation window centered on the price subsidy. Consistent with theoretical arguments, we impose that those effects sum to zero.

Figure [3b](#) suggests that the evolution of fuel purchases would have been similar in 2021 and 2022 in the absence of the price subsidy; this hypothesis cannot be rejected on the basis of the prerebate period.^{[31](#)}

To mitigate any concern about endogeneity, we adopt an IV strategy based on a post-September 1 dummy, which indicates whether the price subsidy is effective as an instrument for prices. This approach enables us to compare what one would obtain when relying on sharp tax-based price changes as the sole source of variability, as opposed to other smaller price fluctuations. It also addresses the issue of measurement error on the dependent variable, which arises because actual fuel purchases are unobserved. Dividing spending by prices could indeed lead to a downward bias in the price coefficient ([Borjas, 1980](#)).

4.2 Econometric specification

To implement our identification strategy, we first aggregate our data into 10,777 cells of individuals (those defined by footnote [17](#)). Our estimations are then weighted according to the sampling importance of those cells. Proceeding in such a way substantially alleviates the computational burden inherent to dealing with high-frequency individual data but does not reduce our identifying power since diesel and gasoline prices are measured at the *département* \times day level. Notably, our fuel price index does vary within a *département* because of the cell-specific fuel mix. Our estimations include cell-specific fixed effects to consider the heterogeneity of fuel spending across cells.

Our estimating equation distinguishes calendar time t , measured at the daily level, from year $y \in \{2021, 2022\}$. We restrict our sample to observations ranging from mid-July to

³¹Our approach does not require any common trend assumption as in a difference-in-differences or in an event study. Moreover, when the absence of any differential pretrend is checked, short-term anticipations should be left aside. As already explained, it is largely expected that a policy-induced dip in purchases be observed within a one- or two-week anticipation window, followed by some spike once the price subsidy is effective.

the beginning of October, both in 2022 and in the baseline year 2021. Let \underline{t} = July 14, 2022 (\bar{t} = October 3, 2022) designate the beginning (end) of that subperiod. Our dependent variable q_{cty} is the fuel quantity, in liters, purchased by individuals belonging to cell c on day t of year y .

We estimate the following equation to recover price sensitivity, following insights from section 3:

$$q_{cty} = \beta p_{cty} + \sum_{h=t_2-\Delta}^{t_2+\Delta} \gamma_{hy} \mathbb{1}_{h=t} \mathbb{1}_{y=2022} + \xi_{cy} + \mu_t + \eta_{cty}, \quad (10)$$

where p_{cty} are prices, ξ_{cy} is a cell-year fixed effect, and μ_t is a day-of-the-year fixed effect that accounts for seasonal adjustment, with 2021 calendar days being adjusted so that they coincide with their siblings in 2022. t_2 corresponds to the beginning of the second price subsidy, namely, September 1. The time interval $[t_2 - \Delta, t_2 + \Delta]$ accounts for the anticipation window around that date: Δ is a bandwidth parameter set by the researcher.

As explained in section 3.2, we consider three estimators:

- **Naive estimator:** $\forall h, y \quad \gamma_{hy} = 0$

- **Constrained estimator:**

$$\sum_{h=t_2-\Delta}^{t_2+\Delta} \gamma_{h,2022} = 0 \quad (11)$$

- **Unconstrained estimator:** No restriction on γ_{hy} .

Standard errors are computed via two-way clustering at the cell and year-day levels.

The choice of the anticipation bandwidth Δ is guided primarily by economic considerations. We set $\Delta = 14$, namely, two weeks, which seems a reasonable value since 75% of interpurchase durations occur within 14 days. Although consumers may be able to manipulate the timing of their visit, especially when they foresee price changes, they are constrained by their tank capacity.³²

We test for anticipations around the price subsidy. We reject the absence of anticipations for the two weeks surrounding the price subsidy, which constitutes empirical evidence against the naive estimator. Moreover, we cannot reject the absence of anticipations the weeks further away, namely, during both the week before and the week after the price

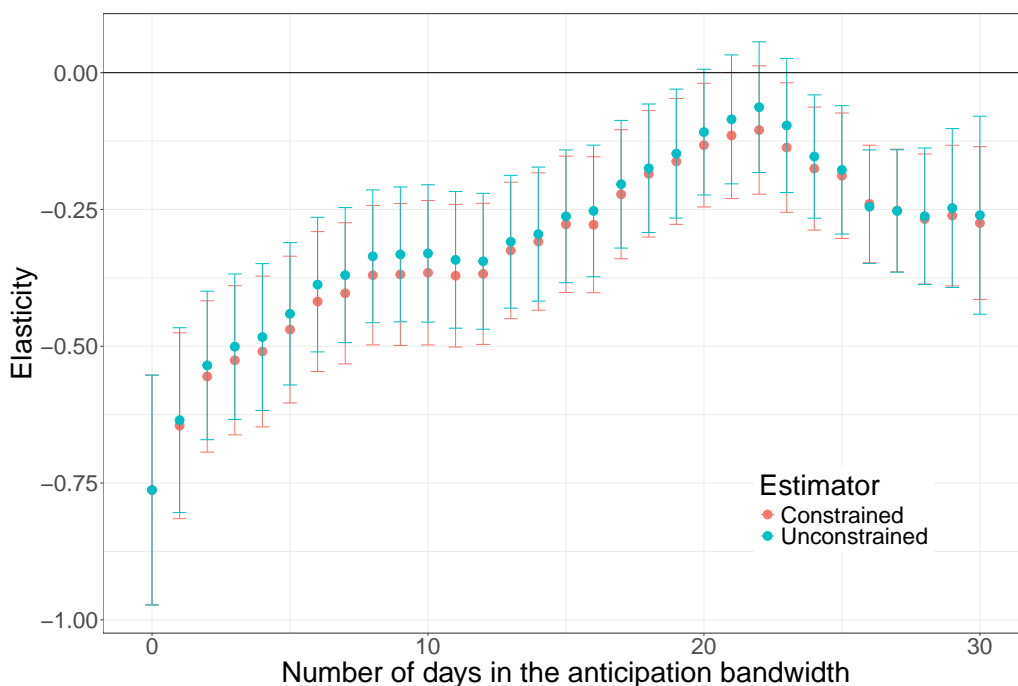
³²Figure 18 suggests that households mostly adapt their behavior at the intensive margin.

shock, with a 90% level of confidence. This statistical argument suggests that most anticipation effects occur the days following and preceding the shock. It also supports the idea that our bandwidth is large enough.

Considering that our constrained estimation procedure rests on the assumption that anticipation effects exactly compensate over the anticipation window, $\Delta = 14$ should pass the statistical test of the constraint $\sum_{h=t_2-\Delta}^{t_2+\Delta} \gamma_{h,2022} = 0$.

Implementing this test on a finite sample requires neutralizing the estimation error on those anticipation coefficients. Failing to do so would surely lead to rejection (see Appendix F for details). Thus, passing that test can be interpreted as empirical evidence in favor of our constrained estimator, as opposed to the unconstrained estimator. Overall, previous tests suggest considering the constrained estimator as our preferred estimator.

Figure 4: Estimated price elasticity (depending on the anticipated bandwidth Δ)



Note. Estimated price elasticity for both constrained and unconstrained specifications (equation (10)), depending on the anticipated bandwidth. Estimation period: from July 15 to October 4 in 2021 and 2022. *Source.* Sample of households that primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure 4 displays how the estimated price elasticity varies with the bandwidth Δ . For

small values of Δ , the estimation does not properly control for anticipations, mechanically overestimating the reaction to price changes by incorrectly attributing intertemporal substitution motives to price sensitivity. This downward bias results in a -0.76 point estimate. When Δ increases, our estimation method becomes fruitful in disentangling strategic delays in purchases from contemporaneous responses to price changes. When $\Delta = 14$, our preferred estimate for price elasticity becomes -0.31. Reassuringly, estimates obtained with higher values of Δ remain rather stable and not significantly different from -0.31.

To further assess the validity of our identification strategy, we perform a falsification or placebo test. We consider a fake price subsidy on September 1, 2021, using 2019 as the baseline year. When comparing fuel purchases in 2019 and 2021, two years without price subsidies, we conclude that there was no substantial spike around September 1 (Figure H2 in the Online Appendix). Table H1 confirms that there is not enough identifying variability in that case.

4.3 External validity

To alleviate concerns about identification being local, we complement the previous approach with a similar econometric specification based on the period from September 2021 to February 2023. With respect to the period from February 24 to mid-April 2022, we refer to the model presented in Appendix E.2. Doing so allows us to rely on other sources for identifying variability, including the beginning of the invasion of Ukraine on February 24, 2022 (denoted by t_0), the first public intervention on April 1, 2022 (denoted by t_1), the reduction in price subsidies in mid-November 2022 (denoted by t_3), and the removal of price subsidies at the beginning of 2023 (denoted by t_4). However, we no longer have any relevant baseline years. We therefore adjust fuel spending for seasonal variations owing to card transaction data provided at the daily level by the *Groupement des Cartes Bancaires* (GIE-CB), the French interbank network that is in charge of centralizing the data. This dataset is almost exhaustive for the universe of French credit card spending. On the basis of that external database, we divide observed fuel spending by the 2019 ratio of daily fuel spending to average fuel spending. We specify:

$$q_{ct} = \beta p_{ct} + \sum_{h=t_0}^{t_1+\Delta} \gamma_h^1 \mathbb{1}_{h=t} + \sum_{k=2}^4 \sum_{h=t_k-\Delta}^{t_k+\Delta} \gamma_h^k \mathbb{1}_{h=t} + \xi_c + \tilde{\mu}_t + \eta_{ct}, \quad (12)$$

with $\tilde{\mu}_t \equiv X'_t\beta + \delta t$, where δ captures any linear trend in fuel purchases and X_t accounts for temporal controls, including day-of-the-week fixed effects³³ and interactions with holidays. To control for anticipation effects, equation (12) is again estimated under the following constraints:

$$\sum_{h=t_0}^{t_1+\Delta} \gamma_h^1 = 0, \quad \sum_{h=t_k-\Delta}^{t_k+\Delta} \gamma_h^k = 0 \quad \forall k = 2, 3, 4. \quad (13)$$

5 Results

5.1 Main estimates

Our estimate of the price elasticity is taken at the means from the price coefficient β in equation (10): $\frac{\partial \log(q)}{\partial \log(p)} = \frac{\partial q / \partial p}{q/p} \equiv \beta \frac{p}{q}$. We compute $\hat{\varepsilon} = \hat{\beta} \frac{\bar{p}}{\bar{q}}$, denoting the average of X by \bar{X} . Table 2 converts the estimated coefficient $\hat{\beta}$ based on the second price subsidy from September 1 into a -0.19 (0.07) price elasticity, obtained with the constrained OLS estimator (Column II). Our preferred estimate of the price elasticity is obtained with the corresponding IV estimator (Column V). It is slightly higher, in absolute: -0.31 (0.08), which is likely explained by measurement error (attenuation bias), simultaneity (upward bias), or by the fact that consumers expect more persistent price changes from a tax shock than from a before-tax price shock. The IV estimate is not much more imprecise than the OLS, and the difference between both estimates is statistically significant at the 5% level.

Table 2: Estimation based on September 1, 2022 price subsidy

	I	II	III	IV	V	VI
price coefficient	-0.62 (0.10)	-0.26 (0.10)	-0.34 (0.09)	-1.06 (0.18)	-0.43 (0.11)	-0.41 (0.10)
price elasticity	-0.44 (0.07)	-0.19 (0.07)	-0.24 (0.06)	-0.76 (0.13)	-0.31 (0.08)	-0.29 (0.07)
IV (Instrument: post- 9/1 dummy)				✓	✓	✓
Anticipation dummies		✓			✓	
Excluding anticipation window			✓			✓
Cell FE	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓
# of cells	10,777	10,777	10,777	10,777	10,777	10,777

Note. Estimation of equation (10) with a sample of 10,777 cells of customers. Estimation period: from July, 15th to October, 4th. Weighted regressions according to the size of the sample within each cell. Baseline year: 2021. Two-way clustering of standard errors at cell and year-day levels.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

³³Daily level data reveal that fuel purchases exhibit strong within-week variations; tanks are much more often refilled on Fridays and Saturdays. By definition, such seasonality cannot be observed on the basis of low-frequency data.

Estimations based on the period following the invasion of Ukraine, namely, March and April 2022, are also in accordance with these findings (Table H5). The constrained OLS estimate, -0.18 (0.07), is not significantly different at the 5% level from the previous estimate, -0.19 (0.07). However, since the period considered here is short, it mechanically mixes different anticipation effects, viewed as confounders, with the price effect. Hence, the estimation is more fragile than the previous estimation, is based on a single price change, and is thus more immune to such confounders.

We next compare our estimates with those issued from equation (12), which is based on the period from September 2021 to February 2023. The main lesson from Table 3 is that we find an average elasticity spanning -0.42 to -0.26. This exercise is reassuring from an empirical viewpoint since it provides a gauge of the external validity of the previous approach. Using sharp price changes (tax changes and the invasion of Ukraine) as instruments yields very close results that range between -0.37 and -0.22 (Table H6, Row 2), which suggests that those shocks provide the main effective source of identifying variation over the period. Previous estimates fall within the range of existing results in the literature, from -0.46 to -0.1 according to Davis and Kilian (2011), depending on the identification strategy.

The naive approach, which would preclude anticipations by imposing $\gamma_{h,2022} = 0, \forall h = t_2 - \Delta, \dots, t_2 + \Delta$, is displayed in Columns I and IV of Table 2. The OLS estimate of the price elasticity, -0.44 (0.07), suffers from a downward bias of -0.25. The same applies to the IV, with anticipation bias now being -0.45. When restricted to the months of March

Table 3: Estimations based on the period from September 2021 to February 2023

	I	II	III	IV	V	VI
price coefficient	-0.56 (0.06)	-0.63 (0.06)	-0.54 (0.08)	-0.42 (0.05)	-0.51 (0.05)	-0.31 (0.04)
price elasticity	-0.46 (0.05)	-0.52 (0.05)	-0.45 (0.07)	-0.35 (0.04)	-0.42 (0.04)	-0.26 (0.07)
Anticipation dummies				✓	✓	✓
Seasonality controls		✓	✓		✓	✓
Linear trend			✓			✓
Cell FE	✓	✓	✓	✓	✓	✓
# of cells	11,031	11,031	11,031	11,031	11,031	11,031

Note. Estimation of equation (12) based on a sample of 11,031 cells of customers (at rate 1/5 for computational issues). Results reported here correspond to median estimates over 100 replications. Estimation period: from September 2021 to February 2023. Regressions are weighted according to the sample size within each cell. Two-way clustering of standard errors at cell and day levels.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

and April, the naive estimate is -0.73 (0.16); see Column I of Table H5. As expected from section 3.2 and Appendix E.2, the anticipation bias is greater in that case. Overall, these results suggest an anticipation bias of approximately -0.4.

When we exclude the anticipation window from our estimation sample, we indirectly relax the assumption that anticipation effects compensate over that time period.³⁴ The unconstrained OLS estimator is now -0.24 (0.06), and the unconstrained IV estimator amounts to -0.29 (0.07), see Columns III and VI of Table 2. Both differences with corresponding constrained estimators are not significant at the 5% level, which is in accordance with our theoretical predictions according to which constrained and unconstrained estimators should be close. This empirical finding further suggests that the naive estimator is biased, and it is all the more likely that the direction of the bias is as expected. More strikingly, the unconstrained OLS estimator based on the months of March and April amounts to -0.22 (0.29) and is very imprecise, as expected from section 3.2. Its theoretical uninformativeness stems from the fact that prices experienced a surge and nearly returned to their initial level after the anticipation window (see Figure I9). In theory, excluding that anticipation window leads to an infeasible estimator. In practice, the standard error dramatically increases.

Interestingly, the point estimates of the naive, constrained and unconstrained estimators (-0.76, -0.31, and -0.29) turn out to be close to those derived in the econometric specification (-0.68, -0.31, and -0.31, see Appendix E.3).

The importance of anticipation effects can be assessed by looking at Figure I10, which depicts the $\hat{\gamma}$ coefficients recovered over the period from September 2021 to February 2023. It is confirmed that ignoring such effects in demand estimation is highly misleading since those short-term intertemporal substitution effects substantially shape the pattern of demand, in addition to the price effect.

To evaluate whether our model is able to accurately predict fuel purchases, Figure I11 provides a comparison of predictions with actual demand. The fit of the model estimated over the period from September 2021 to February 2023 looks quite satisfactory in this regard.

³⁴This is equivalent to nonparametrically controlling for daily dummies during the anticipation window.

5.2 Robustness checks

In this section, we conduct various sensitivity analyses to assess the robustness of previous evidence with respect to methodological choices: (i) we estimate a first-difference version of our model, (ii) we consider an alternative parametric specification, (iii) we address possible concerns related to the measurement of prices, and (iv) we address the issue of possibly heterogeneous anticipations. These robustness checks confirm our baseline results. The use of a first-difference or a pseudo-Poisson estimation yields very similar elasticities. The estimation is also robust to various measures of the price index, such as the sole price of diesel or the first decile of prices within the *département*. Finally, considering alternate anticipation windows does not substantially alter our estimates. Details can be found in Appendix H.

5.3 Heterogeneity of price elasticity

We next explore whether the price elasticity of fuel demand is homogeneous among groups of consumers with similar observed characteristics. We first investigate whether the average price elasticity varies with income (Figure 5a).³⁵ location (Figure 5b) or family status (Figure 5c), which turns out not to be the case.³⁶ However, the price elasticity approximately measures a relative reaction, and fuel consumption is higher, on average, for wealthier individuals as well as for those living in rural areas; hence, these individuals respond more to fuel price changes in nominal terms. Additionally, employees are more inelastic than retirees are (Figure I14), as the rationale would suggest since the former are more likely to commute for professional reasons. In contrast, a dimension along which average price elasticity exhibited sizeable dispersion was fuel spending in 2019 (Figure 5e). As intuition suggests, 'dependent households', who rely on the car as their primary transportation, are less elastic. Individuals in the bottom 25% of fuel spending have an average elasticity of -0.82 (0.22), whereas those in the top 25% have an average elasticity of -0.26 (0.08).

³⁵Figure I12 is similar but has greater granularity. Investigating heterogeneity in total card spending, rather than income, yields the same result (Figure I13), which mitigates concerns about the sensitivity of this empirical finding with respect to the measurement of income.

³⁶If any, Parisian drivers tend to be more price sensitive (Table H4) due to easier substitution with other transportation. However, the difference is not statistically significant at the 5% level. Although rural households use cars more, which could make them more inelastic, they are also older, which tends to make them more elastic, according to Figure 5d. Overall, their average price elasticity does not differ much from that of urban or semiurban households.

Figure 5f shows that 'dependent' and liquidity-constrained individuals are most likely to undergo any rise in fuel prices. These empirical findings have important policy implications. When designing transfer schemes to compensate losers, the policy-maker seeks to target such households (Sallee, 2022), which requires availing information about spending and liquidity.

We next perform a more systematic search of the relevant dimensions of heterogeneity, whereby we allow for the elasticity to directly depend on observed characteristics.³⁷ We use the method of causal forests pioneered by Athey et al. (2019). Figure 6 displays the sorted group average treatment effects issued from a segmentation of our sample into five groups. We test for homogeneity of price elasticity, which we reject. The 20% most price-sensitive households have an elasticity of -0.9 , whereas the 20% most inelastic households have null price elasticity, and the difference is statistically significant at the 5% level. This empirical evidence is consistent with fuel being a necessity for almost every car driver.

Our estimated price elasticity differs substantially across households, and it is possible to determine who the most price-sensitive households are in terms of both socioeconomic characteristics and geographic location (Table 4). Most elastic households have lower fuel spending and are more likely to retire or be poorer (in terms of either income or liquidity).³⁸

These results are consistent with those of Frondel et al. (2012) from Germany and Goetzke and Vance (2021) from the US. Using quantile regressions, both studies reveal a similar result as that derived here, namely, that people who drive a lot have a low elasticity of demand for fuel. The same holds for Gillingham et al. (2015) in Pennsylvania, who find substantial heterogeneity in price responsiveness depending on fuel economy, with the lowest number of fuel economy vehicles corresponding to greater responsiveness. They are also in accordance with Kilian (2022), who found that U.S. states with lower income, higher unemployment rates, and lower urban shares respond more to price variations. However, an important difference is that our results are derived from micro data. As a matter of relevance for public policy, previous findings should be viewed as a contribution to the optimal market design of second-best policies like tax-and-transfer schemes, which arise due to imperfect information and imperfect tagging of individual consumption. Such market failure limits the planner's control over the final distribution of outcomes. A more

³⁷We adopt a log-log specification here to directly interpret the point estimate as an elasticity.

³⁸Those estimates, along with anticipation effects, may serve as sufficient statistics for a policy evaluation of the environmental, financial, and distributional impacts of actual price subsidies effective in 2022, as shown in Appendix J.

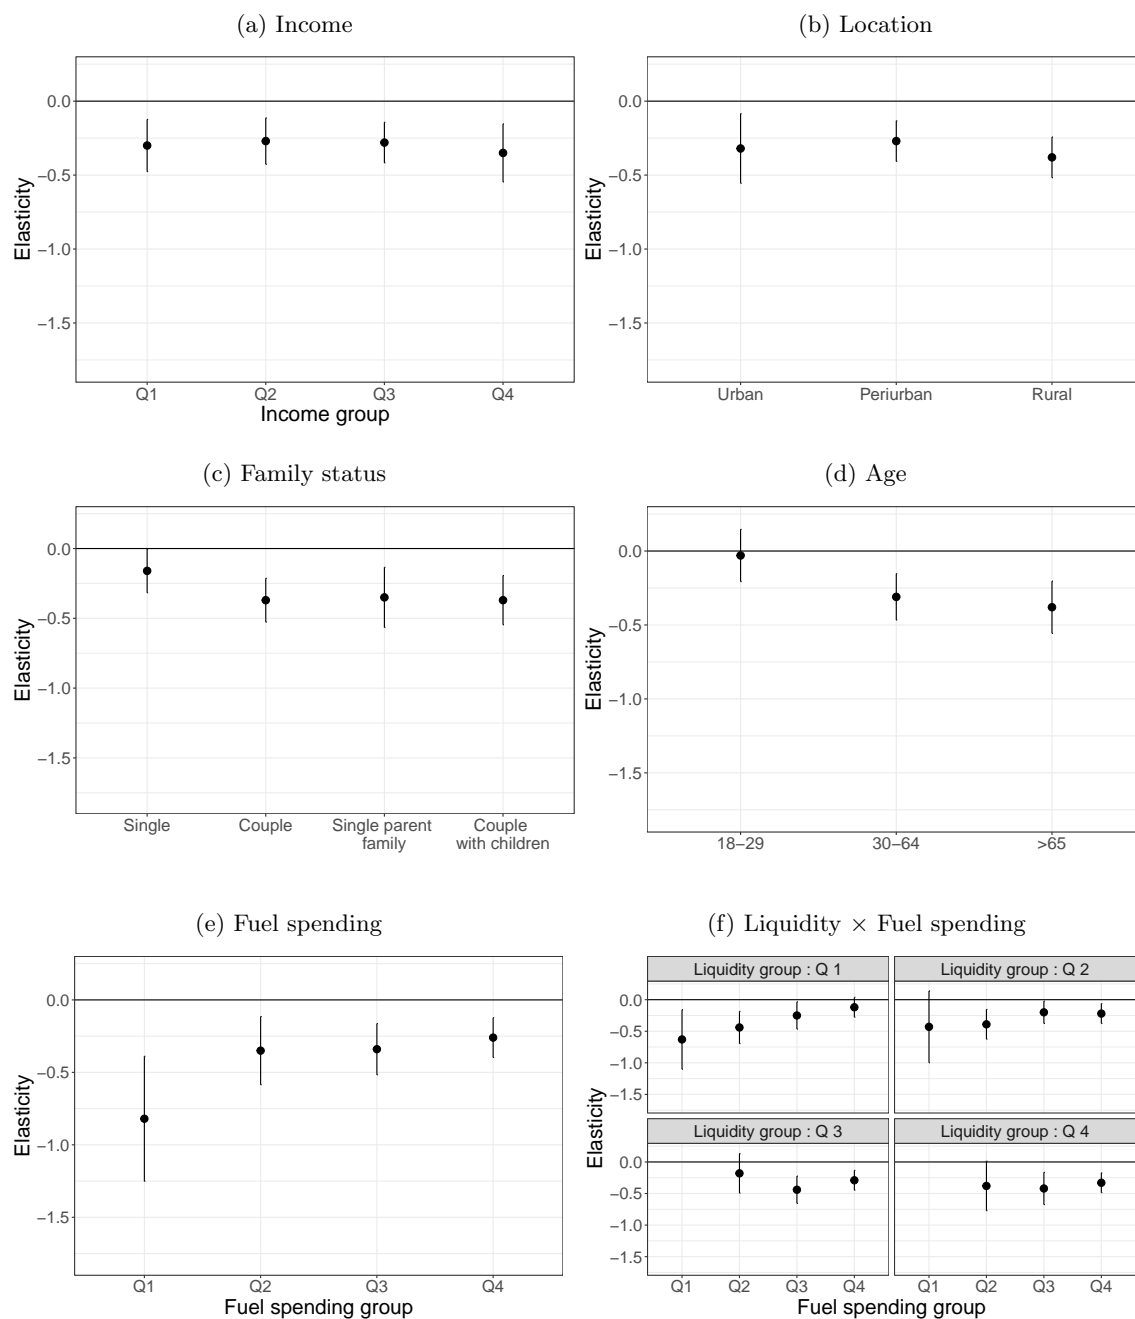
Table 4: Socio-economic characteristics (by elasticity group)

	All		20 % most elastic		20 % least elastic	
	Avg.	Sd.	Avg.	Sd.	Avg.	Sd.
Average age	51	0.04	55	0.08	45	0.08
<i>Age groups</i>						
Share of below 30 years	0.12	0.000	0.06	0.001	0.24	0.002
Share of between 30 and 64 years	0.63	0.001	0.60	0.02	0.61	0.002
Share of above 65	0.25	0.00	0.33	0.002	0.16	0.002
Monthly fuel spending (euros)	99	0.16	58	0.34	124	0.35
Monthly income (euros)	2,758	4.52	2,377	9.79	3,029	10.18
Liquid wealth (euros)	42,077	223	36,063	486	44,973	505
Fuel spending-to-income ratio	0.05	0.000	0.04	0.00	0.07	0.00
<i>Location</i>						
Share of periurban	0.45	0	0.41	0	0.46	0
Share of rural	0.24	0	0.25	0	0.20	0
Share of urban	0.31	0	0.34	0	0.34	0
<i>Familial status</i>						
Single parents	0.04	0	0.05	0	0.03	0
Couples without child	0.34	0	0.38	0	0.27	0
Couples with children	0.26	0	0.28	0	0.21	0
<i>Occupations</i>						
Craftsmen, merchants and business owners	0.08	0	0.08	0	0.08	0
Managerial and professional occupations	0.14	0	0.13	0	0.15	0
Technicians and associate professionals	0.15	0	0.12	0	0.17	0
Employees	0.19	0	0.18	0	0.20	0
Workers	0.14	0	0.13	0	0.15	0
Retirees	0.19	0	0.26	0	0.12	0
# of consumption units	1.49	0	1.57	0	1.37	0
Transactions with mostly round amounts	0.07	0	0.09	0	0.06	0

Note. Characteristics in each quintile of conditional average treatment effect (CATE). Estimation period: from July 15th to October 4th at the exclusion of the anticipation window. Average means and standard errors correspond to medians over 50 estimations on the sample following [Chernozhukov et al. \(2018\)](#). Each estimation is performed on a randomly drawn partition of 5 folds. *Source.* Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

accurate prediction helps mitigate that empirical problem (Sallee, 2022), and our findings help enhance the quality of that prediction.


Figure 5: Heterogeneity in some observed characteristics



Note. Estimated price elasticity via equation (10). Estimation period: from July 15 to October 4 in 2021 and 2022. Subsamples are defined according to the corresponding household characteristics (continuous variables: observed before the intervention from January to June 2022; discrete variables: observed in June 2022). Two-way clustering of standard errors at the cell and year-day levels.

Source. Sample of households that primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure 6: Sorted Group Average Treatment Effects



Note. Average elasticity of the conditional average treatment effect (CATE) in each quintile. The estimation period is from July 15 to October 4 in 2021 and 2022, excluding the anticipation window, which is based on corresponding subsamples. The average elasticity estimates and standard errors correspond to medians based on over 50 estimations on the sample following [Chernozhukov et al. \(2018\)](#). Each estimation is composed of a randomly drawn partition of 5 folds.

Source. Sample of households that primarily bank at *Crédit Mutuel Alliance Fédérale*.

6 Policy implications

In what follows, we focus on the counterfactual price increase in 2022 that would have prevailed in the absence of any price subsidy, namely a substantial price shock of +33% with respect to 2021 (from €1.5 per liter, 'ex ante' hereafter, to €2 per liter, 'ex post' hereafter). To characterize which policy governments should enforce to compensate households exposed to fuel price inflation, we proceed in two steps. First, we adopt a positive approach based on counterfactual policies, either price subsidies or transfers, and evaluate the share of losers under imperfect targeting. Second, we take a normative stand and compare the efficiency and equity of two policy instruments, price subsidies and transfers.

Following any price increase from p to $p' > p$, the *behavioral response* refrains agents from buying fuel, which attenuates their welfare loss with respect to the one that they incur due to the sole *mechanical response*. In particular, consumers may reallocate their expenditures away from fuel. Compensating for the behavioral response through transfers is more affordable than compensating for the mechanical response through price subsidies. Denoting the total expenditure function by E , the compensating variation is written as $CV = E(p', U_0) - E(p, U_0)$, where $U_0 = V(p, Y)$ corresponds to the initial indirect utility level, i.e., before any price increase, of an agent endowed with annual disposable income Y . From Shephard's lemma, $h(p, U_0) = \frac{\partial E(p, U_0)}{\partial p}$; hence, $CV = \int_p^{p'} h(x, U_0) dx$. From the law of demand, it follows that $CV \leq \int_p^{p'} h(p, U_0) dx = (p' - p)h(p, U_0) = (p' - p)h(p, V(p, Y))$. Finally, $CV \leq (p' - p)f(p) = CV_0$ where $f(p)$ corresponds to the quantity of fuel demanded at the initial price p , and CV_0 denotes the compensating variation under inelastic demand. Since $f(p) = h(p, V(p, Y))$, CV_0 is the cost of a price subsidy that exactly compensates for fuel price inflation in the absence of any behavioral adjustment (when $\varepsilon = 0$) and when the price change is marginal (when $(p' - p) \ll p$).

Transfers are equivalent to price subsidies when equality holds, either when price changes are marginal or when the behavioral response can be ignored, which is, for instance, the case if Hicksian demand is inelastic. In contrast, as soon as price changes are substantial, which corresponds to the realistic policy case where there is a point in easing consumers' burden, equality no longer holds. Moreover, the magnitude of the difference between CV , the ex post burden (i.e., once the *behavioral* response has been taken into account), and CV_0 , the *mechanical* or ex ante burden, depends on the shape of the compensated demand. Our empirical approach below consists of measuring the compensated

price elasticity of fuel demand, based on (i) our estimated uncompensated elasticity, (ii) the Slutsky decomposition for elasticities, and (iii) a parametric specification of preferences.

Transfers allow agents to reoptimize their purchase behavior with regard to not only fuel but also goods and services other than fuel. Price subsidies do not, but they provide full compensation for the burden CV_0 . However, in the presence of a non-marginal behavioral response, agents may partly reallocate their consumption such that transfers become a strictly cheaper option ($CV < CV_0$).³⁹ It may come from regressivity concerns, too. Transfers turn out to be more appealing since they can be targeted to specific groups, typically low-income households incurring high losses relative to income (Figures I3 and I5). The policy-maker may thus use them as part of a tax-and-transfer scheme that cancels the regressivity (Douenne, 2020; Young-Brun, 2023). The choice between price subsidies and transfers depends on the decision maker’s preferences, which we emphasize in section 6.2. Choosing a policy instrument is, hence, an empirical issue that requires both taking agents’ behavior into account and targeting the desired segment of consumers given distributive objectives.

6.1 Tax-and-transfer schemes: The limits of targeting

Although it is tempting to use transfers for distributive motives, predicting fuel consumption at the household level is essential to accurately target such transfers. This task requires much information, including the distribution of annual fuel spending across households. To illustrate, we compare observed price subsidies with counterfactual transfers, either lump sum or conditional on income, location, or fuel spending. To that end, we specify preferences with regard to fuel and the rest of consumption on an annual basis.⁴⁰ We consider that agents’ preferences can be represented by some CES utility function $U(f, m) = (m^\rho + Af^\rho)^{\frac{1}{\rho}}$, $\rho < 1$, over annual fuel consumption f and the outside good m

³⁹The tension between both policy tools may also come from the desire to correct for externalities, as is the case in the environmental literature. Since the impact in terms of emissions is small in our setting (the policy increased the carbon footprint by 0.36%, see Appendix J), externalities play a limited role here.

⁴⁰Our inventory model above was designed to model anticipations at a daily or weekly frequency, and it now provides us with an estimate of the short-run price elasticity coefficient viewed as sufficient statistics in the subsequent analysis.

taken as the *numéraire*, with A indexing household h 's preference for fuel.⁴¹ We omit the unnecessary index h , although households may differ in their price elasticity as well as in their preference for fuel. Maximizing utility given the budget constraint: $pf + m \leq Y$, leads to the fuel demand: $f(p, Y) = Y \frac{(A/p)^\sigma}{1 + A^\sigma p^{1-\sigma}}$ where $\sigma = \frac{1}{1-\rho} > 0$. Fuel expenditures amount to $E_f(p, Y) = Y \frac{A^\sigma p^{1-\sigma}}{1 + A^\sigma p^{1-\sigma}}$. Following any rise in fuel price from p to p' , the policy-maker may seek to alleviate consumers' welfare losses on the basis of direct transfers amounting to the following:

$$CV = Y \left[\left(\frac{1 + A^\sigma (p')^{1-\sigma}}{1 + A^\sigma p^{1-\sigma}} \right)^{\frac{1}{1-\sigma}} - 1 \right]. \quad (14)$$

Computing the empirical counterpart of (14) requires information on two parameters viewed as sufficient statistics for each household: σ and A . These parameters are calibrated to match the uncompensated price elasticity at the observed price $\varepsilon^u(p) = -\frac{\sigma + A^\sigma p^{1-\sigma}}{1 + A^\sigma p^{1-\sigma}}$, and the observed fuel budget share $s_f = \frac{E_f}{Y} = \frac{A^\sigma p^{1-\sigma}}{1 + A^\sigma p^{1-\sigma}}$ at the household level.⁴² The compensated elasticity that governs the shape of the Hicksian demand and, thus, the difference between CV and CV_0 is obtained from Slutsky's decomposition for elasticities: $\varepsilon^u = \varepsilon^c - \varepsilon_R s_f = \varepsilon^c - s_f$ since $\varepsilon_R = 1$ under CES preferences, hence $\varepsilon^c(p) = \varepsilon^u(p) + \frac{A^\sigma p^{1-\sigma}}{1 + A^\sigma p^{1-\sigma}} = -\frac{\sigma}{1 + A^\sigma p^{1-\sigma}}$.

We next quantify how both instruments perform in terms of compensatory power. We assume that the policy-maker has access to information on income, location, and ex ante fuel spending. The latter assumption is more demanding since it is usually not available to the policy-maker.⁴³ In addition to a uniform €0.5 per liter annual price subsidy, we simulate (i) unconditional or lump sum, (ii) income-based, (iii) location-based, and (iv) past-fuel spending-based transfers. In the first case, the transfer is computed on the basis of average ex ante fuel spending as well as average price elasticity; in the second (or third) scenario, the policy-maker knows both the average ex ante fuel spending in each income (or location) group and the average price elasticity of the household's income (or location)

⁴¹In Appendix K we also consider quasilinear preferences with iso-elastic demand as a robustness check and reach similar conclusions. Note also that our reduced-form elasticity is only assumed to hold around the equilibrium price/quantity pair, an assumption that would be unnecessary based on an Almost Ideal Demand System (AIDS) à la Deaton and Muellbauer (1980), for instance. We nevertheless believe that estimating an entire demand system is beyond the scope of the current article.

⁴²If those parameters were calibrated to match the average moments, one would obtain $\sigma \approx 0.277$ and $A \approx 6 \times 10^{-6}$.

⁴³Note that recent access to bank account data may be an alternative source of information.

group; in the last scenario, s/he knows ex ante fuel spending at the household level and the price elasticity of the household’s ex ante fuel spending group. For each policy, we compute the share of households that receive exact, positive, or negative compensation. Households receive positive compensation if the transfer exceeds their welfare loss. Formula (14) shows that full compensation is difficult to achieve on the basis of uniform or third-degree discrimination transfers due to informational frictions, namely, imperfect screening of price elasticity and/or ex ante fuel spending at the household level. As a result, *forecasting error* arises in the planner’s prediction of ex post fuel spending.

We estimate that the policy-maker would opt for lump sum transfers rather than price subsidies provided that s/he values full compensation at the household level less than €22 per household level, i.e., the difference between the cost CV_0 of the uniform price subsidy, €404 per HH, and the cost of the second-best lump sum transfer CV on the basis of average past fuel spending, €382 per HH (Table 5). *Forecasting error* is negligible here because average fuel spending in 2021 is a good predictor of average fuel expenditures in 2022 in the absence of any price increase. Lump sum transfers are thus 5% cheaper.⁴⁴ However, 41% of households are losers in the sense that they would prefer a price subsidy. The average loss of the losers amounts to €342, whereas the average gain among the 59% of ‘winners’ is €219.

Table 5: Simulated compensatory policies to 2022 fuel price inflation

	Price subsidy	Transfer			
		Uniform	Income-based	Location-based	Based on ex ante fuel spending
Average cost per household (€)	404	382	382	382	378
Share of HH with positive compensation (%)	0	59	57	58	44
Average positive compensation (€)		219	204	212	107
Share of HH with negative compensation (%)	0	41	43	42	56
Average negative compensation (€)		342	298	327	108

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

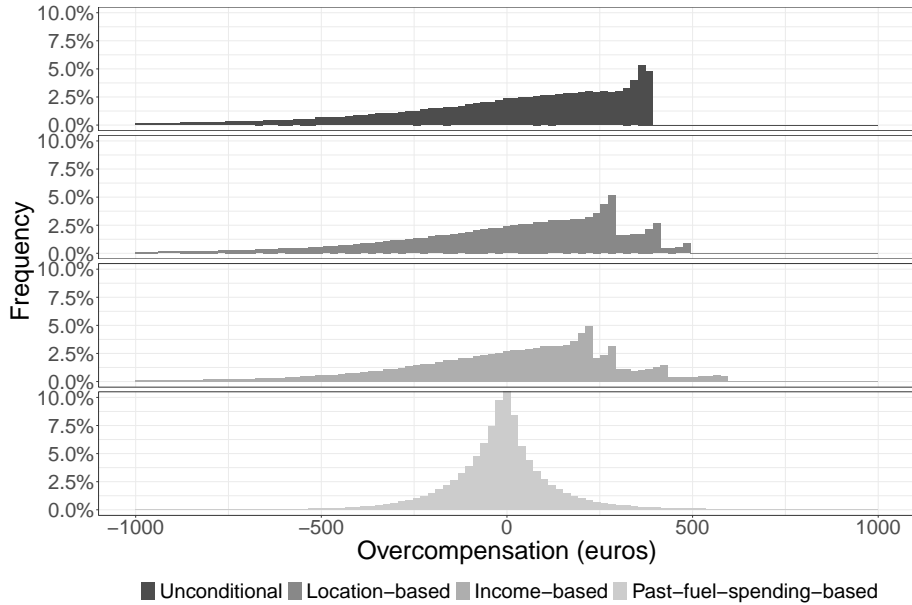
The policy-maker could try to reduce the number of losers by conditioning the transfer on income, location or ex ante fuel spending, as in the above policies (ii), (iii) and (iv). However, the number of losers remains fairly high: 43% in (ii), with an average loss of €298, 42% in (iii), with an average loss of €327, and even 56% in (iv) but with a substantially

⁴⁴Transfers based on past fuel spending are 6% cheaper, hence availing more information translates into an additional (limited) pecuniary gain for the government. Under quasi-linear preferences and iso-elastic demand (Appendix K), lump sum transfers would be 4% cheaper.

smaller loss of €108, on average.

Figure 7 quantifies the value of information on income, location, and ex ante fuel spending for targeting tax-and-transfer schemes. Reliable knowledge of that fuel spending helps reduce the financial cost of the intervention by tightening the eligibility condition. It is also fruitful in reducing the dispersion of gains and losses, hence the heterogeneity of the impact of public intervention. The average loss among eligible households would be high regardless of how the transfer is designed, and regardless of the information available to the policy-maker, due to unobserved heterogeneity in fuel consumption. This empirical result concurs with the finding of Douenne (2020) that designing policies to cope with horizontal distributive effects is a difficult task.

Figure 7: Gains and losses from compensatory tax-and-transfer schemes



Note. We compare the relative impact of transfers with respect to price subsidies to compensate households for the observed price increase from 2021–2022 (+€0.5 per liter, i.e., from €1.5– €2). We compute households’ gains and losses, in euros, associated with transfers rather than with price subsidies. *Source.* Sample of households that primarily bank at *Crédit Mutuel Alliance Fédérale*.

As another caveat beyond partial compensation, transfers based on ex ante fuel spending (or on location, to a smaller extent) do not provide the right incentives to lower emissions, since they are not designed to correct for polluting externalities.

Finally, it should be acknowledged that the previous analysis neglects general equilib-

rium effects, by which fuel usage may increase when the decision-maker opts for transfers because transfers represent an additional source of income. In theory, this income effect may generate additional inflation. In practice, we believe that the magnitude of that effect is sufficiently small. In our example, where income elasticity is equal to one because fuel demand is linear in income, the concerned transfers (€382 per HH) account for less than 1% of the average annual income (€43,464 in our sample). This income effect would then lead to an increase in fuel demand of 1%, at most. Even in the extreme case of perfectly inelastic supply, that increase in demand would translate into a supplementary price increase of 1%. Without mentioning that the supply may be elastic or that the income elasticity may be an order of magnitude lower than 1, it seems quite a reasonable approximation to neglect this channel in a setting where the price change is as high as 33%.

6.2 Price subsidies or transfers?

Finally, we compare price subsidies and transfers by emphasizing the trade-off that arises between financial and distributional motives. We have empirically demonstrated that transfers generate a certain number of losers, even when much targeting information is available. However, since only a portion of the population may be eligible for such transfers, governments may favor transfers when they seek to alleviate excessive losses for specific groups while providing a limited budget for their compensatory policies.

We compare the efficiency and equity of two competing policy approaches depending on the level of public funding that the government is ready to consent when compensating for fuel price inflation. Mimicking here the aggregation of individual utilities into a social welfare function, we posit that the decision-maker minimizes a combination of agents' welfare losses. Those losses are expressed, in absolute terms, as a fraction of income. The government has two available policy instruments: price subsidies or rebates r and transfers T , which may be uniform or targeted toward segments of consumers. It also faces a budget constraint; specifically, the policy cannot cost more than some exogenous C , which may well be issued from revenue recycling (a plausible value of C would typically amount to €404 per household given the previous analysis). We further assume that the planner weighs more individuals with a high fuel budget share. A particular case is the 'alleviation of excessive relative losses', which is reminiscent of the justice principle called

'alleviating poverty' by [Saez and Stantcheva \(2016\)](#).⁴⁵ Using similar notations as above, we thus consider the following tool-specific programs with $f_h(p, Y_h) = Y_h \frac{(A_h/p)^\sigma}{1+A_h^\sigma p^{1-\sigma}}$ referring to household h 's fuel consumption, $E_h(p, Y_h) = Y_h \frac{A_h^\sigma p^{1-\sigma}}{1+A_h^\sigma p^{1-\sigma}}$ referring to ex ante fuel spending, and $RL_h = \frac{-CV_h}{Y_h} < 0$ referring to the ex post relative loss:

$$\min_{(r_h)} \sum_{h=1}^N \omega_h(|RL_h^r|) |RL_h^r| \quad \text{s.t.} \quad \sum_{h=1}^N r_h f_h(p' - r_h) \leq C \quad (15)$$

$$\min_{(T_h)} \sum_{h=1}^N \omega_h(|RL_h^T|) |RL_h^T| \quad \text{s.t.} \quad \sum_{h=1}^N T_h \leq C \quad (16)$$

From the previous analysis, one obtains the following:

$$RL_h^r = 1 - \left(\frac{1 + A_h^\sigma (p' - r_h)^{1-\sigma}}{1 + A_h^\sigma p^{1-\sigma}} \right)^{\frac{1}{1-\sigma}} \quad (17)$$

and

$$RL_h^T = \frac{T_h}{Y_h} + 1 - \left(\frac{1 + A_h^\sigma (p')^{1-\sigma}}{1 + A_h^\sigma p^{1-\sigma}} \right)^{\frac{1}{1-\sigma}} \quad (18)$$

Social weights $\omega_h(\cdot)$ are assumed to be nondecreasing. The polar cases correspond to $\omega_h = \mathbb{1}\{|RL_h| = \max_j |RL_j|\}$ and $\omega_h = 1/N$. 'Alleviating excessive (relative) losses' is a special case where $\omega_h = \mathbb{1}\{RL_h > \overline{RL}\}$, with \overline{RL} being some threshold that the government considers excessive. Since the objective function is decreasing in policy instruments, the budget constraint of the government binds in the optima:

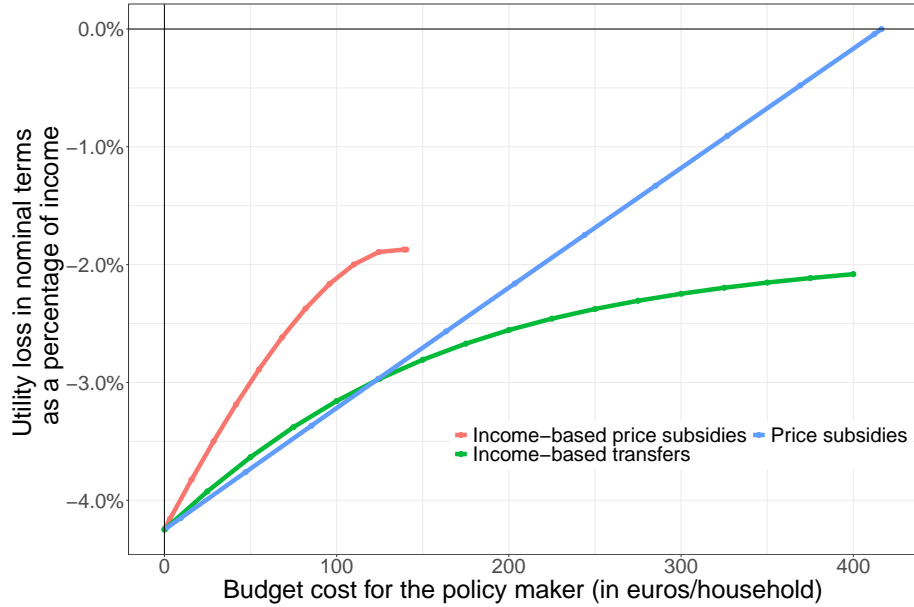
$$\sum_{h=1}^N E_h \frac{r_h}{p} \left(\frac{p}{p' - r_h} \right)^\sigma \frac{1 + A_h^\sigma p^{1-\sigma}}{1 + A_h^\sigma (p' - r_h)^{1-\sigma}} = C, \quad \sum_{h=1}^N T_h = C \quad (19)$$

In our application, we consider price subsidies and transfers that can be either uniform or income-based (eligible households then correspond to the bottom half of the income distribution). Our simulations indicate that income-based price subsidies are preferred (Figure 8) when the government's objective function consists of alleviating the top 10%

⁴⁵To the extent that the objective function considered here is a reduced-form, as opposed to a social welfare function, this setting strongly resembles [Saez \(2002\)](#)'s framework. In particular, a given tax change is desirable if the sum of the mechanical effect, the behavioral effect, and the welfare (here, the objective function) effect is positive.

highest relative losses. A sensitivity analysis of alternative thresholds is provided in Figure I15. Due to nondiscriminatory rules in addition to technical issues, this first-best solution may be hardly feasible, if not forbidden in practice.

Figure 8: A comparison of policy instruments (Policy objective: Alleviating the top 10% welfare loss, in nominal terms, relative to income)



Note. For any given level of public funding (x-axis), we compare (conditional or unconditional) price subsidies and transfers as policy tools designed to compensate households for the observed price increase from 2021–2022 (+€0.5 per liter, i.e., from €1.5 to €2).

Source. Sample of households that primarily bank at *Crédit Mutuel Alliance Fédérale*.

That option aside, there exists some funding threshold \bar{C} , approximately €160 per household, such that income-based transfers dominate uniform price subsidies below \bar{C} , whereas the contrary prevails above \bar{C} . In other words, there is a case for income-based transfers when the budget constraint is tight. Despite achieving partial compensation at the household level, transfers are effective at targeting the right households from the planner’s viewpoint. This is because income acts as a satisfying screening device for relative loss. More than 80% of the households in the top 10% of relative losses are classified as low-income (see Figure I16). Consequently, targeted transfers to this group effectively mitigate losses among those most severely hit. When more money can be devoted to the policy, however, the share of low-income households within the top 10% of losers decreases

because of diminishing marginal returns of transfers. Transfers become marginally less attractive than price subsidies, which are more expensive but achieve full compensation at the household level. Put differently, transfers are less effective at high levels of financial commitment.

This result can be more generally derived under any nondecreasing weights ω_h . The relative location of \bar{C} with respect to perfect discrimination transfer⁴⁶ C^T depends on the social weight of households with high relative losses. Income-based transfers are preferable to price subsidies when $\bar{C} < C^T$, hence provided that the government places enough weight on households with high relative losses. Despite their lower compensatory power, transfers may be chosen by governments with limited funding, as noted by [Sallee \(2022\)](#). These findings shed new insights into policy rules that may provide practical guidance to decision-makers. Interestingly, previous simulations are consistent with the actual choice made by the French government when price subsidies were replaced with an income-based transfer at the beginning of 2023, for both financial and distributional considerations.

7 Conclusion

This paper has shown that researchers who aim at causally inferring the price elasticity of fuel demand should ideally have access to granular, high-frequency data, in addition to relying on exogenous price variations. Daily data permit disentangling anticipation effects from the price effect. We leverage an econometric specification that is less parametric than other approaches in the literature, which may also be useful in other public finance contexts, including capital taxation, commodity taxation, and green taxation. Equipped with such ingredients, we estimate an average price elasticity of -0.31. This price elasticity exhibits sizeable dispersion, primarily in the fuel spending dimension. Motorists who drive more, including those who commute for professional reasons, are more inelastic, especially when they are also liquidity constrained. In contrast, income and location are not associated with significantly different average price elasticities.

We have also simulated various policy experiments in which counterfactual transfers would have been provided to households, possibly depending on their observed charac-

⁴⁶This perfect discrimination transfer achieves full compensation while being cheaper than uniform price subsidies. Each household receives an amount, in proportion to the fuel bill, that is slightly less than $\frac{\Delta p}{p}$, i.e., the relative increase that is offset by price subsidies (CV_0/Y). However, implementation considerations make it infeasible in practice.

teristics. We have shown that full compensation at the household level, as is the case with price subsidies, costs 5% more than imperfect compensation achieved by lump sum transfers. Assuming that the government has reliable information on households' ex ante fuel spending would not help much in this regard. Partial compensation would still arise because of unobserved heterogeneity in fuel consumption. The value of information rather resides in the targeting of a specific group of consumers, which makes it more affordable to the government, but information per se does not permit substantially reducing the dispersion of gains and losses within that group. Finally, a decision-maker seeking to alleviate excessive losses, relative to income, would prefer income-based price subsidies. When s/he faces a tight budget constraint, feasibility may require the use of income-based transfers rather than uniform price subsidies. In this case, transfers may indeed be chosen despite their lower compensatory power.

From a market design viewpoint, our results help better tailor compensatory mechanisms to drivers' needs. More generally, the methodology developed here applies to any other market exposed to inflation. Further research should therefore concentrate on other institutional settings to investigate whether previous empirical findings, especially with respect to policy rules, extend to different contexts.

References

- Andersen, A. L., Jensen, A. S., Johannesen, N., Kreiner, C. T., Leth-Petersen, S., and Sheridan, A. (2023). How do households respond to job loss? Lessons from multiple high-frequency datasets. *American Economic Journal: Applied Economics*, 15(4):1–29.
- Astier, J., Jaravel, X., and Peron, M. (2023). Mesurer les effets hétérogènes de l'inflation sur les ménages. CAE Focus n°99.
- Athey, S., Tibshirani, J., and Wager, S. (2019). Generalized random forests. *The Annals of Statistics*, 47(2):1148–1178.
- Baker, S. R. (2018). Debt and the response to household income shocks: Validation and application of linked financial account data. *The Journal of Political Economy*, 126(4):1504–1557.

- Boehm, J., Fize, E., and Jaravel, X. (2023). Five facts about mpccs: Evidence from a randomized experiment. Working paper.
- Borenstein, S., Cameron, A. C., and Gilbert, R. (1997). Do gasoline prices respond asymmetrically to crude oil price changes? *The Quarterly journal of economics*, 112(1):305–339.
- Borjas, G. J. (1980). The relationship between wages and weekly hours of work: The role of division bias. *The Journal of Human Resources*, 15(3):409–423.
- Bounie, D., Camara, Y., Fize, E., Galbraith, J., Landais, C., Lavest, C., Pazem, T., Savatier, B., et al. (2020). Consumption dynamics in the covid crisis: real time insights from french transaction bank data. *Covid Economics*, 59:1–39.
- Chernozhukov, V., Demirer, M., Dufflo, E., and Fernandez-Val, I. (2018). Generic machine learning inference on heterogeneous treatment effects in randomized experiments, with an application to immunization in India. Technical report, National Bureau of Economic Research.
- Coglianese, J., Davis, L. W., Kilian, L., and Stock, J. H. (2017). Anticipation, tax avoidance, and the price elasticity of gasoline demand. *Journal of Applied Econometrics*, 32(1):1–15.
- Dahl, C. and Sterner, T. (1991). Analysing gasoline demand elasticities: a survey. *Energy Economics*, 13(3):203–210.
- Davis, L. W. and Kilian, L. (2011). Estimating the effect of a gasoline tax on carbon emissions. *Journal of Applied Econometrics*, 26(7):1187–1214.
- Deaton, A. and Muellbauer, J. (1980). An almost ideal demand system. *The American economic review*, 70(3):312–326.
- Deville, J.-C. and Särndal, C.-E. (1992). Calibration estimators in survey sampling. *Journal of the American Statistical Association*, 87(418):376–382.
- Douenne, T. (2020). The vertical and horizontal distributive effects of energy taxes: A case study of a French policy. *The Energy Journal*, 41(3).

- Espey, M. (1998). Gasoline demand revisited: an international meta-analysis of elasticities. *Energy Economics*, 20(3):273–295.
- Fagereng, A., Holm, M. B., Moll, B., and Natvik, G. (2019). Saving behavior across the wealth distribution: The importance of capital gains. NBER Working paper 26588.
- Frondel, M., Ritter, N., and Vance, C. (2012). Heterogeneity in the rebound effect: Further evidence for Germany. *Energy Economics*, 34(2):461–467.
- Ganong, P. and Noel, P. (2019). Consumer spending during unemployment: Positive and normative implications. *The American Economic Review*, 109(7):2383–2424.
- Gautier, E., Marx, M., and Vertier, P. (2023). How do gasoline prices respond to a cost shock? *Journal of Political Economy Macroeconomics*, 1(4):707–741.
- Gelman, M., Gorodnichenko, Y., Kariv, S., Koustas, D., Shapiro, M. D., Silverman, D., and Tadelis, S. (2023). The response of consumer spending to changes in gasoline prices. *American Economic Journal: Macroeconomics*, 15(2):129–160.
- Gillingham, K., Jenn, A., and Azevedo, I. M. (2015). Heterogeneity in the response to gasoline prices: Evidence from Pennsylvania and implications for the rebound effect. *Energy Economics*, 52:S41–S52.
- Goetzke, F. and Vance, C. (2021). An increasing gasoline price elasticity in the United States? *Energy Economics*, 95:104982.
- Hendel, I. and Nevo, A. (2006). Measuring the implications of sales and consumer inventory behavior. *Econometrica*, 74(6):1637–1673.
- Hughes, J. E., Knittel, C. R., and Sperling, D. (2008). Evidence of a shift in the short-run price elasticity of gasoline demand. *The Energy Journal*, 29(1):113–134.
- Kilian, L. (2022). Understanding the estimation of oil demand and oil supply elasticities. *Energy Economics*, 107:105844.
- Kilian, L. and Zhou, X. (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.

- Knittel, C. R. and Tanaka, S. (2021). Fuel economy and the price of gasoline: Evidence from fueling-level micro data. *Journal of Public Economics*, 202:104496.
- Levell, P., O'Connell, M., and Smith, K. (2024). Distributional effects of the European energy crisis. Working paper.
- Levin, L., Lewis, M. S., and Wolak, F. A. (2017). High frequency evidence on the demand for gasoline. *American Economic Journal: Economic Policy*, 9(3):314–47.
- Mertens, K. and Ravn, M. O. (2012). Empirical evidence on the aggregate effects of anticipated and unanticipated us tax policy shocks. *American Economic Journal: Economic Policy*, 4(2):145–181.
- Montag, F., Sagimuldina, A., and Schnitzer, M. (2021). Does tax policy work when consumers have imperfect price information? Theory and evidence. CESifo Working Paper 9138.
- Saez, E. (2002). The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes. *Journal of Public Economics*, 83(2):217–230.
- Saez, E. and Stantcheva, S. (2016). Generalized social marginal welfare weights for optimal tax theory. *American Economic Review*, 106(01):24–45.
- Sallee, J. M. (2022). Pigou creates losers: On the implausibility of achieving Pareto improvements from efficiency-enhancing policies. Working paper.
- Young-Brun, M. (2023). Public acceptability of carbon taxation: a model of political support with income and urban-rural inequality. Working paper.

Online Appendix

A Data-related acknowledgements (in French)

Data from *Crédit Mutuel Alliance Fédérale*:

Première banque à adopter la qualité d'entreprise à mission, Crédit Mutuel Alliance Fédérale a contribué à cette étude par la fourniture de données de comptes bancaires sur la base de deux échantillons : un échantillon d'entreprises et un échantillon de ménages par tirage aléatoire et construit de telle sorte qu'on ne puisse pas identifier les entreprises (exclusion de sous populations de petite taille) ou les ménages. Toutes les analyses réalisées dans le cadre de cette étude ont été effectuées sur des données strictement anonymisées sur les seuls systèmes d'information sécurisés du Crédit Mutuel en France. Pour Crédit Mutuel Alliance Fédérale, cette démarche s'inscrit dans le cadre des missions qu'il s'est fixées :

- contribuer au bien commun en oeuvrant pour une société plus juste et plus durable : en participant à l'information économique, Crédit Mutuel Alliance Fédérale réaffirme sa volonté de contribuer au débat démocratique ;
- protéger l'intimité numérique et la vie privée de chacun : Crédit Mutuel Alliance Fédérale veille à la protection absolue des données de ses clients.

B More details on data and institutional setting

Two concerns have been raised by the literature with regard to the external validity of bank account data (Baker, 2018): representativeness and completeness. We therefore use several external sources to assess both representativeness and completeness of our databases.

Representativeness To alleviate concerns about representativeness, and to build upon previous works afore mentioned, we proceed to calibration weighting using the method proposed by Deville and Särndal (1992). We compute weights that exactly replicate exogenous targets for auxiliary variables, attached to the whole population, while ensuring that these calibrated weights are as close as possible to original sampling weights. By construction, the weighted sample has the same distribution with regard to the corresponding variables as the whole population. We consider the following dimensions, called margins: age, sex and *département*, for that auxiliary information.

The distribution of household expenditures with respect to their position in the standard of living distribution obtained in transaction data matches closely the one issued from the representative consumption survey *Budget des Familles* (Figure B3). In particular, putting aside both ends of the income distribution, spending-to-income ratios look remarkably similar and decreasing from 1 to 0.75, which mitigates previous concerns related to measurement error on income. If anything, our data overestimate spending, probably because Crédit Mutuel customers tend to be richer. This is confirmed by Table B1 which suggests that Crédit Mutuel customers are wealthier: they have a higher income (Figure B1), detain more assets (Figure B2), and spend more than the average (Figure B3). The pregnancy of liquidity constraints can be assessed by looking at the liquid wealth-to-income ratio, about 10, meaning that, on average, households have liquidity equivalent to 10 months of income. It decomposes into a 3.5 ratio of liquid assets over end-of-month balances on deposit accounts (this number compares well with the one documented in the U.S. by Baker (2018)), and another 3.5 ratio of end-of-month balances on deposit accounts over monthly income. Finally, these customers are younger, on average, and tend to live in more peripheral areas. Figure B4 focuses on the sole fuel category: it can be verified that our sample spends systematically a bit more, probably because it is composed of richer customers. Reassuringly, the evolution of fuel spending looks yet quite identical (Figure B6) to the one issued from the comprehensive *Groupement des Cartes Bancaires*

(GIE-CB) dataset, with a 0.99 correlation. On top of supporting external validity, this empirical evidence provides some grounds for a seasonal adjustment based on the data issued from that French interbank network. More generally, we believe that it alleviates legitimate concerns about selection bias.

Completeness First, our measure of spending exhibits quite the same evolution as the one issued from the *Groupement des Cartes Bancaires CB*, the French national interbank network (Figure B5).

Second, our measure of income is more volatile (Figure B7) than the one measured by Insee.⁴⁷ This higher dispersion is rather expected: it is intrinsically related to the fact that we do not observe income directly, but rather all incoming transfers. Yet it is reassuring to see that the magnitude of possible measurement error is limited.

Third, our measure of liquid assets is slightly more dynamic than the one reported by *Banque de France* that centralizes information from all other bank networks (Figure B8). If anything, Cr dit Mutuel customers likely enjoy higher capital gains (Fagereng et al., 2019) but that composition effect looks again rather limited.

On the whole, these comparisons with external sources suggest (i) that representativeness is not too much of a concern, (ii) that the calibration weighting contributes to alleviate this problem, and (iii) that the remaining differences on earnings and assets are mostly due to differences in concepts, rather than to incompleteness.

⁴⁷Namely, the gross standard of living as the ratio of gross disposable income over the number of consumption units.

B.1 Data: External validity

Table B1: Summary statistics

	Weighted sample
# of observations	181,527
	<i>Banking variables (sample means)</i>
Monthly Spending	2,721
<i>Fuel (cards)</i>	94
Income	3,622
Financial Assets	
<i>Liquid financial Assets</i>	38,116
<i>Illiquid financial Assets</i>	23,469
Ratio liquid assets/deposit account	3.1
	<i>Household head characteristics (sample means)</i>
Age	53
Female	0.41
Craftsmen, merchants and business owners	0.08
Managerial and professional occupations	0.13
Technicians and associate professionals	0.12
Employees	0.17
Workers	0.11
Periphery areas	0.41
Rural areas	0.19
Urban areas	0.37

Note. Statistics computed in 2021 for transactions (spending, income), January 2021 for assets and socio-demographics. Pecuniary amounts in €. The head of the household if the oldest member of that household.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Table B2: Deciles of disposable income

	All	CMAF
D1	14,530 (210)	15,580 (1063)
D2	18,590 (222)	19,620 (1208)
D3	22,540 (221)	23,650 (1625)
D4	26,610 (339)	30,250 (2138)
D5	31,670 (434)	34,980 (2148)
D6	37,440 (370)	43,480 (1983)
D7	43,880 (430)	49,920 (2043)
D8	52,440 (474)	57,870 (1508)
D9	66,420 (856)	69,570 (3859)

Note. The 9th decile of disposable income is 66,420 euros for all households and 69,570 for households who primarily bank at Crédit Mutuel Alliance Fédérale.

Source. *Histoire de vie et Patrimoine* French wealth survey (2017).

Table B3: Deciles of financial wealth

	All	CMAF
D1	350 (34)	300 (241)
D2	1,051 (53)	1,600 (523)
D3	2,712 (153)	3,975 (934)
D4	5,750 (255)	8,951 (2,175)
D5	11,000 (399)	15,500 (2,261)
D6	19,206 (801)	24,761 (4,099)
D7	32,000 (951)	39,590 (5,861)
D8	56,410 (1,750)	63,334 (7,228)
D9	117,000 (3,729)	114,162 (27,385)

Note. The 9th decile of financial wealth is 117,000 euros for all households and 114,162 for households who primarily bank at Crédit Mutuel Alliance Fédérale.

Source. *Histoire de vie et Patrimoine* French wealth survey (2017).


Table B4: Share of financial assets in their main bank

	Income deciles
D1	0.92 (0.01)
D2	-0.00 (0.01)
D3	-0.02 (0.01)
D4	-0.04 (0.01)
D5	-0.04 (0.01)
D6	-0.04 (0.01)
D7	-0.05 (0.01)
D8	-0.07 (0.01)
D9	-0.06 (0.01)
D10	-0.10 (0.01)

Note. Households in the bottom 10% of income detain 92% of their financial assets in their main bank, against 82% for those in the top 10%. On average, households (who primarily bank at Crédit Mutuel Alliance Fédérale) detain 88.4% (88.1%) of their financial assets in their main bank.


Source. *Histoire de vie et Patrimoine* French wealth survey (2017).

Figure B1: Distribution of income (transaction data vs. survey data from *ERFS*, Insee)




Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*; *Enquête sur les Revenus Fiscaux et Sociaux* (ERFS) survey.

Figure B2: Distribution of household financial wealth by income (transaction data vs. survey data from *Histoire de Vie et Patrimoine*, Insee)




Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*; *Patrimoine* survey.

Figure B3: Distribution of household monthly expenditures by income (transaction data vs. survey data from *Budget des Familles*, Insee)




Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*; *Budget des Familles* survey.

Figure B4: Distribution of monthly fuel spending, by income (transaction data vs. survey data from *Budget des Familles*, Insee)




Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*; *Budget des Familles* survey.

Figure B5: Evolution of spending (transaction data vs. aggregate data from the French interbank network)



Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*; *GIE-CB* data.

Figure B6: Evolution of fuel spending (transaction data vs. aggregate data from the French interbank network)




Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*; *GIE-CB* data.

Figure B7: Income (transaction data vs. aggregate data from national accounts, Insee)



Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*; French National Accounts.

Figure B8: Liquid Assets (transaction data vs. aggregate data from Banque de France)



Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*; *Banque de France*.

B.2 Prices

As detailed in Appendix A of [Gautier et al. \(2023\)](#), the first step consists of mapping raw data to a daily panel dataset at the retailer and gas-type level. Since different price changes may occur within the same day (the typical frequency of price changes being a few days), we consider the price that prevails at 5pm, as [Montag et al. \(2021\)](#) do. In a second step, we remove inactive stations, which we define as stations that have not experienced any price change since at least 30 days, following [Gautier et al. \(2023\)](#). Note that a station may be active for diesel, but inactive for gasoline: in that case, the sole diesel-type station remains in the sample. We then trim outliers by deleting the top and bottom 1% of price observations for each *département*, type of fuel, and day.

Possible concerns: Search Price search is a potential source of measurement error, which we can hardly address with current data. Search behavior would lead to overestimate price sensitivity: transaction prices would be lower, and quantity higher since they are derived from the ratio of expenditures over prices. We obtain a higher elasticity, in absolute, when using instrumental variables, which suggests that instrumenting helps mitigate this issue of possible measurement error. Besides, there is suggestive evidence that consumers did not substantially alter their purchase behavior by switching from a gas station to another over the considered period. From comprehensive monthly-level VAT data, which contains information on stations' revenue, the rank-rank and log sales correlation between the beginning of the year and other month of the year is close to 1. It hardly varies over time, especially between January 2022 and September 2022 despite substantial price fluctuations that might stimulate search behavior: the only period when that correlation is significantly lower is April 2020, which corresponds to the lockdown. The year 2022 is very similar to the previous years. Further details are provided in Appendix C.

Possible concerns: Heterogeneous pass-through Consumers may also be imperfectly informed about prices, or there might be some heterogeneity in their degree of information in this regard, which may result in both heterogeneous pass-through ([Montag et al., 2021](#)) and measurement error on prices. Yet the price subsidies we focus on were publicly announced, hence quite salient. In our empirical analysis, we do not exclude that possibility, but implicitly assume that this error is constant over time. In France, this

concern may be less of an issue than in Germany due to the quasi-absence of within-day price variation (Fronzel et al., 2012).⁴⁸ We nevertheless investigate that issue further in Appendix D and conclude that this pass-through is very high, close to 100%. It is not surprising due to the price setting mechanism explained in section 2. This pass-through is also quite homogeneous across gas stations and does not significantly impact stations' ranking in terms of gas average prices.

⁴⁸If imperfect information is a concern, our IV strategy then suffers from imperfect compliance: our price coefficient shall then be interpreted as an average reaction to prices by heterogeneous consumers with respect to their level of information, yet this is still a policy-relevant parameter. Another parameter, which we do not estimate here, is the price-sensitivity of fully informed consumers.

C Sales ranking

The volatility of fuel prices was very high in 2022 when compared to 2021. Figure 1 indeed shows that prices markedly increased, especially after the invasion of Ukraine, until the first price subsidy in April, a second tax cut occurred in September. The tax cuts were then repealed in November and December. In front of such shocks, households may have adapted their purchase behavior by searching and switching to other service stations. In principle, this could endanger our estimation of fuel price sensitivity, which may suffer from a downward bias.

To address this issue, we use comprehensive administrative VAT file at the monthly level, which every French firm has to fill out so as to report its sales. We perform two exercises. First, we rank gas stations based on their monthly sales in January of every year from 2018 to 2022. We then looked at the Spearman rank correlation between the rank in January and the rank in any other month of the year. We focus on April and May where prices were the highest, as well as on June and September, i.e. after the two price subsidies. Figure C1 depicts the rank correlation for every year between January and September from 2018 to 2022. We have two main findings. First, this correlation is always very high regardless of the year or the month considered (except during Covid-19 lockdown in April 2020). Second, it is neither higher nor lower in 2022, which suggests that consumers did not massively change gas stations on that year. Indeed, if any substantial change in search behavior were at stake, we would expect this correlation to vary over time (cheapest stations would go up in the ranking and the most expensive ones would go down, which is not the case). As an additional exercise, we allow for *département*-specific rankings, and reach similar conclusions.

Second, for every year we look at the correlation (instead of the rank) of total sales, in logarithm, between January and September. We also find a strong and positive relationship, with a coefficient close to 1 and that does not differ in 2022 from what it amounts to the other years. Figure C2 depicts that relationship using bins of equal size according to January's sales. This empirical evidence does not point out to substantial change occurring in 2022 and affecting the relationship between sales in January and sales in September.

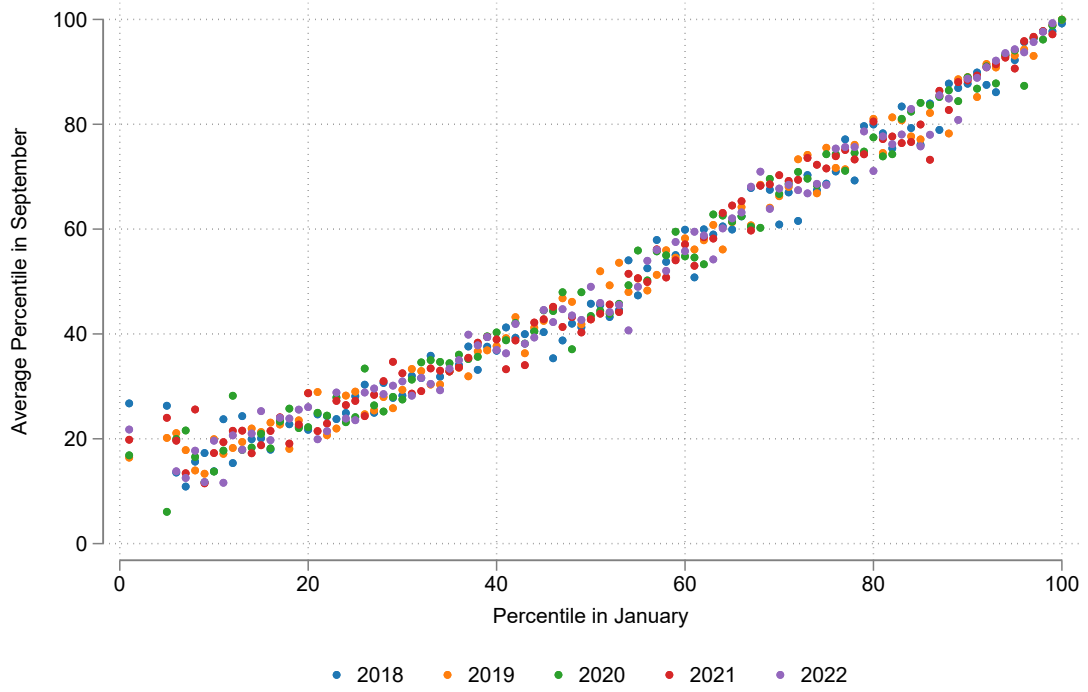


Figure C1: Rank correlation of gas stations' sales between January and September



Figure C2: Correlation of the logarithm of gas stations' sales between January and September

D Pass-Through

To empirically document both the level and the heterogeneity of the pass-through of the tax rebate, we estimate its actual impact on prices around September 1st, 2022.

To first elicit the level of the pass-through, we rely on average weekly diesel prices in neighboring European countries⁴⁹ as a control group for French prices (Figure D1) in an event study analysis. We find empirical evidence in favor of common trends before the public intervention, followed by a significant and substantial price decrease in France after the intervention. The result is suggestive of a high pass-through, close to 100%.

We then investigate how homogeneous that pass-through was. To that end, we leverage French gas station-level data and rank the stations in five groups according to their pre-rebate average daily diesel prices. We then examine the pattern of prices around the September intervention (Figure D2). We observe a sharp and persistent decrease in prices regardless of the group considered, which suggests a quite homogeneous pass-through. If anything, the most expensive stations were also the ones that reduced their prices the most, but the ranking of stations in terms of prices remained the same before and after the intervention.

⁴⁹We exclude Germany from our analysis because price subsidies in this country were removed at the same time on September 1st.




Figure D1: Weekly diesel price evolution by countries

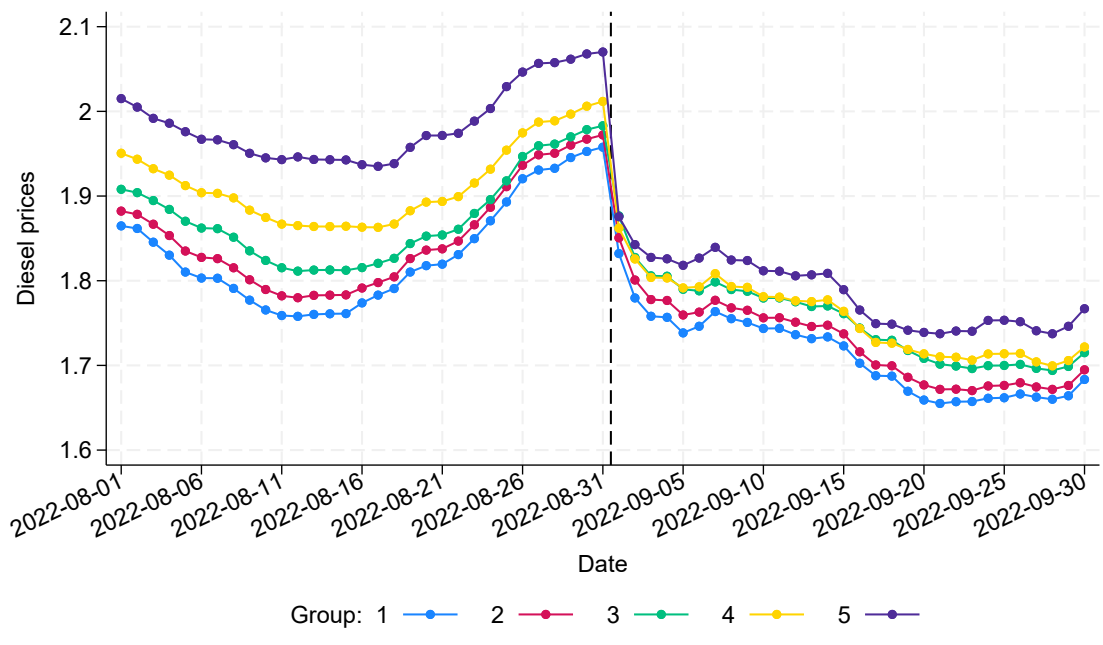


Figure D2: Daily price evolution by groups of gas station

E Details on econometric specifications

E.1 A single price reduction (4-period model)

E.1.1 Unconstrained estimator

Program:

$$\min_{(q_0, \beta, \gamma_2, \gamma_3)} \sum_{k=1}^4 T_k \left(q_k - q_0 - \beta p_k - \sum_{k=2}^3 \gamma_k \right)^2$$

FOC:

$$q_0 + \beta \bar{p} = \bar{q} - \sum_{k=2}^3 \gamma_k \frac{T_k}{T} \quad (20)$$

$$q_0 \bar{p} + \beta \bar{p}^2 = \bar{p} \bar{q} - \sum_{k=2}^3 \gamma_k p_k \frac{T_k}{T} \quad (21)$$

$$q_k - q_0 - \beta p_k = \gamma_k, \quad \forall k = 2, 3 \quad (22)$$

Plugging (22) into (20) yields $q_0(T_1 + T_4) + \beta(T_1 p_1 + T_4 p_4) = T_1 q_1 + T_4 q_4$

Plugging (22) into (21) yields $q_0(T_1 p_1 + T_4 p_4) + \beta(T_1 p_1^2 + T_4 p_4^2) = T_1 p_1 q_1 + T_4 p_4 q_4$

Thus

$$\hat{\beta}^u = \frac{q_1 - q_4}{p_1 - p_4} < 0$$

$$\hat{q}_0^u = \frac{p_1 q_4 - p_4 q_1}{p_1 - p_4}$$

$$\hat{\gamma}_k^u = q_k - q_1 - (q_4 - q_1) \frac{p}{\Delta p} - p_k \frac{q_1 - q_k}{\Delta p} \quad \forall k = 2, 3$$

that is,

$$\hat{\gamma}_2^u = q_2 - q_1 < 0, \hat{\gamma}_3^u = q_3 - q_4 > 0$$

E.1.2 Constrained estimator

Under the constraint $\sum_{k=2}^3 T_k \gamma_k = 0$, FOC now write:

$$q_0 + \beta \bar{p} = \bar{q} \quad (23)$$

$$q_0 \bar{p} + \beta \bar{p}^2 = \bar{p} \bar{q} - \sum_{k=2}^3 \gamma_k p_k \frac{T_k}{T} \quad (24)$$

$$\gamma_2(T_2 + T_3) = T_3[q_2 - q_3 - \beta(p_2 - p_3)] \quad (25)$$

Plugging the latter into (24) and using (23) combined with the constraint yields

$$\hat{\beta}^c = \frac{(T_2 + T_3)[T_1 q_1 (T_3 + T_4) - T_4 q_4 (T_1 + T_2)] + (T_2 T_4 - T_1 T_3)(T_2 q_2 + T_3 q_3)}{[T_2 T_4 (T_1 + T_2) + T_1 T_3 (T_3 + T_4)] \Delta p}$$

When the episode is symmetric with respect to the moment when prices fall, i.e. $T_1 = T_4$ and $T_2 = T_3$, the latter formula simplifies to

$$\hat{\beta}^c = \frac{q_1 - q_4}{\Delta p}$$

In the absence of symmetry, the information about purchases during the anticipation window may be used to infer β .

$$\hat{\gamma}_2^c = T_3 \frac{T_1(T_3 + T_4)(q_2 - q_1) + T_4(T_1 + T_2)(q_4 - q_3)}{T_2 T_4 (T_1 + T_2) + T_1 T_3 (T_3 + T_4)} < 0$$

Under symmetry, the latter expression simplifies to $\hat{\gamma}_2^c = (q_2 - q_1)/2 + (q_4 - q_3)/2$.

E.1.3 Naive estimator

Under $\gamma_k = 0, \forall k = 2, 3$, FOC now write:

$$q_0 + \beta \bar{p} = \bar{q}$$


$$q_0 \bar{p} + \beta \bar{p}^2 = \bar{p} \bar{q}$$

It follows that

$$\hat{\beta}^n = \frac{\bar{p}q - \bar{p}\bar{q}}{p^2 - \bar{p}^2} = \frac{T}{T_1 + T_2} \frac{\bar{q} - \bar{q}^{34}}{\Delta p} < 0$$

E.2 Price surge + compensating price subsidy (5-period model)

Figure E1: Price versus anticipation effects (5-period model)



The main advantage of our estimation procedure (the fact of imposing a zero-sum for the γ coefficients over the anticipation window) is even more striking when we consider a price surge followed by a price subsidy that brings prices back to their initial level. Though simplified, this framework once again resembles the situation that prevailed at the beginning of the invasion of Ukraine, followed a few weeks later by the first excise tax rebate announced on March 11th and implemented on April 1st. We then leverage a 5-period model as described by Figure E1: $\forall k \neq 3$, $p_k = p$, and $p_3 = p + \Delta p$. Here the researcher expects that $q_3 < q_1 = q_5 < q_2 = q_4$ due to positive anticipation effects in period 2, following the expected price surge in period 3 as well as to negative anticipation effects in period 3 (as a consequence of the latter, but also following the expected price subsidy in period 4), on top of the sole price effect ($p_3 > p$). In the same vein as before, we consider the following linear specification:⁵⁰

⁵⁰Again, the inventory model would impose further constraints: $\gamma_2 = \lambda \frac{\Delta p}{2\theta} = \gamma_4 > 0$ and $\gamma_3 = -\lambda \frac{\Delta p}{\theta} < 0$. The model would then predict that $i_1 = i_4 = i_5 = 0$, $i_2 = \lambda \frac{\Delta p}{2\theta}$, $i_3 = -\lambda \frac{\Delta p}{2\theta}$, and that $c_3 - \lambda \frac{\Delta p}{\theta} = q_3 < q_1 = q_5 = c_1 < q_2 = q_4 = c_1 + \lambda \frac{\Delta p}{2\theta}$. Consistent with the econometric specification at stake, $\gamma_2 = q_2 - c_2 = \lambda \frac{\Delta p}{2\theta} = q_2 - q_1 = \gamma_4 = q_4 - c_4 = q_4 - q_5$, $\gamma_3 = q_3 - c_3 = -\lambda \frac{\Delta p}{\theta} = (q_3 - q_1) - \left(\frac{q_2 + q_3 + q_4}{3} - \frac{q_1 + q_2 + q_3 + q_4 + q_5}{5} \right) \frac{15}{2}$, such that $\gamma_2 + \gamma_3 + \gamma_4 = 0$.

$$q_t = q_0 + \beta p_t + \gamma_2 \mathbb{1}_{t \in 2} + \gamma_3 \mathbb{1}_{t \in 3} + \gamma_4 \mathbb{1}_{t \in 4} + u_t. \quad (26)$$

Not imposing that anticipation effects exactly compensate over the anticipation window $T_2\gamma_2 + T_3\gamma_3 + T_4\gamma_4 = 0$ would be equivalent to discard the whole episode when inferring price sensitivity: since the price is identical in periods 1 and 5, the unconstrained estimator is now infeasible (cf. Appendix E.2.2). In practice, the imprecision, namely the standard error, should dramatically increase. Under the naive approach, $\hat{\beta}^n = \frac{T}{T-T_3} \frac{q_3 - \bar{q}}{\Delta p} < 0$ (cf. Appendix E.2): this estimator mostly relies on the sole time period when the price effectively varies, and compares the demand in that period with the average demand over the whole episode. By definition, such an approach does not account for any short-term intertemporal substitution. Under the assumption that anticipation effects exactly compensate over the anticipation window, $\hat{\beta}^c = \frac{T}{T_3} \frac{T_2 + T_3 + T_4}{T_1 + T_5} \frac{\bar{q}^{234} - \bar{q}}{\Delta p}$ (see Appendix E.2.3). Interestingly, the ‘constrained estimator’ exploits the information contained in the anticipation window to infer the price effect, while adjusting for anticipatory behavior. A numerical example based on observed prices and purchases during that episode suggests that the anticipation bias would be even more pronounced here. The naive estimated elasticity would reach 1.55, in absolute (see Appendix E.3), yet the constrained estimation would amount to -0.32 only, still in absolute. On the whole, this example also suggests that the estimation procedure is perhaps more fragile when relying on that sole episode, compared with the one based on the single price reduction.

To sum up, the main insights of this exercise are the following: (i) anticipations bias the naive estimator downwards; (ii) the constrained estimator should not differ much from the unconstrained estimator, when the latter is feasible; (iii) the former estimator is more precise, which matters when the latter is empirically uninformative.

E.2.1 Testing the model

Though it is not possible to test the model in the sense that $T_2\gamma_2 + T_3\gamma_3 + T_4\gamma_4 = 0$ holds by construction, it is yet possible to test whether $q_1 = q_5$ and $q_2 = q_4$ in the data. From Table E1, $q_1 \approx 2.373$ liters per day and $q_5 \approx 2.360$ liters per day, hence a tiny 0.5% difference; similarly, $q_2 \approx 2.694$ liters per day and $q_4 \approx 2.554$ liters per day, a 5.5% gap. Besides, the model could be rejected if the condition $q_1 - q_3 > (q_2 - q_1) + (q_4 - q_5)$ was not met; once those periods have been appropriately weighted according to their duration, it

cannot be rejected, though.⁵¹

E.2.2 Unconstrained estimator

Program:

$$\min_{(q_0, \beta, \gamma_2, \gamma_3, \gamma_4)} \sum_{k=1}^5 T_k \left(q_k - q_0 - \beta p_k - \sum_{k=2}^4 \gamma_k \right)^2$$

FOC:

$$q_0 + \beta \bar{p} = \bar{q} - \sum_{k=2}^4 \gamma_k \frac{T_k}{T} \quad (27)$$

$$q_0 \bar{p} + \beta \bar{p}^2 = \bar{p} \bar{q} - \sum_{k=2}^4 \gamma_k p_k \frac{T_k}{T} \quad (28)$$

$$q_k - q_0 - \beta p_k = \gamma_k, \quad \forall k = 2, \dots, 4 \quad (29)$$

Plugging (29) into (27) yields $q_0(T_1 + T_5) + \beta(T_1 p_1 + T_5 p_5) = T_1 q_1 + T_5 q_5$

Plugging (29) into (28) yields $q_0(T_1 p_1 + T_5 p_5) + \beta(T_1 p_1^2 + T_5 p_5^2) = T_1 p_1 q_1 + T_5 p_5 q_5$

Thus

$$\hat{\beta}^u = \frac{q_1 - q_5}{p_1 - p_5}$$

infeasible since $p_1 = p_5$.

$$\hat{q}_0^u = \frac{p_1 q_5 - p_5 q_1}{p_1 - p_5}$$

$$\hat{\gamma}_k^u = \frac{[q_k(p_1 - p_5) - p_k(q_1 - q_5)] - (p_1 q_5 - p_5 q_1)}{p_1 - p_5} \quad \forall k = 2, \dots, 4$$

E.2.3 Constrained estimator

Under the constraint $\sum_{k=2}^4 T_k \gamma_k = 0$, FOC now write:

$$q_0 + \beta \bar{p} = \bar{q} \quad (30)$$

⁵¹9.04 $\approx T_3(q_1 - q_3) > T_2(q_2 - q_1) + T_4(q_4 - q_5) \approx 7.21$.

$$q_0\bar{p} + \beta\bar{p}^2 = \bar{p}\bar{q} - \sum_{k=2}^4 \gamma_k p_k \frac{T_k}{T} \quad (31)$$

$$\gamma_2(T_2 + T_3) + \gamma_4 T_4 = T_3[q_2 - q_3 - \beta(p_2 - p_3)] \quad (32)$$

$$\gamma_2 T_2 + \gamma_4(T_3 + T_4) = T_3[q_4 - q_3 - \beta(p_4 - p_3)] \quad (33)$$

It follows from (32) and (33) that

$$\gamma_2 = \frac{[(T_3 + T_4)q_2 - T_3q_3 - T_4q_4] - \beta[(T_3 + T_4)p_2 - T_3p_3 - T_4p_4]}{T_2 + T_3 + T_4}$$

$$\gamma_4 = \frac{[(T_2 + T_3)q_4 - T_2q_2 - T_3q_3] - \beta[(T_2 + T_3)p_4 - T_2p_2 - T_3p_3]}{T_2 + T_3 + T_4}$$

Plugging the latter into (31) and using (30) combined with the constraint yields

$$\hat{\beta}^c = \frac{T}{T_3} \frac{T_2 + T_3 + T_4}{T_1 + T_5} \frac{\bar{q}^{234} - \bar{q}}{\Delta p} < 0$$

N.B. $\hat{\beta}^c < 0$ because $\bar{q}^{234} < \bar{q}$ due to $\bar{p}^{234} > \bar{p}$.

$$\hat{q}_0^c = \bar{q}^{15} - (\bar{q}^{234} - \bar{q}) \frac{p}{\Delta p} \frac{T}{T_3} \frac{T_2 + T_3 + T_4}{T_1 + T_5}$$

$$\hat{\gamma}_k^c = q_k - \bar{q}^{15} + (\bar{q}^{234} - \bar{q}) \frac{p - p_k}{\Delta p} \frac{T}{T_3} \frac{T_2 + T_3 + T_4}{T_1 + T_5} \quad \forall k = 2, \dots, 4$$

hence

$$\hat{\gamma}_k^c = q_k - \bar{q}^{15} > 0 \quad \forall k = 2, 4$$

and

$$\hat{\gamma}_3^c = q_3 - \bar{q}^{15} - (\bar{q}^{234} - \bar{q}) \frac{T}{T_3} \frac{T_2 + T_3 + T_4}{T_1 + T_5} < 0$$

E.2.4 Naive estimator

Under $\gamma_k = 0$, $\forall k = 2, \dots, 4$, FOC now write:

$$q_0 + \beta \bar{p} = \bar{q}$$

$$q_0 \bar{p} + \beta \overline{p^2} = \overline{p\bar{q}}$$

It follows that

$$\hat{\beta}^n = \frac{\overline{p\bar{q}} - \bar{p}\bar{q}}{\overline{p^2} - \bar{p}^2} = \frac{T}{T - T_3} \frac{q_3 - \bar{q}}{\Delta p} < 0$$

E.3 Numerical examples

4-period model Cf. seasonal adjustment wrt baseline year 2021.

Unconstrained estimator: $\hat{\varepsilon}^u \approx -0.31$.

Constrained estimator: $\hat{\varepsilon}^c \approx -0.31$.

Naive estimator: $\hat{\varepsilon}^n \approx -0.68$.

Table E1: Fuel prices and purchases following the invasion of Ukraine

Period	01-10 to 02-24	02-25 to 03-09	03-10 to 03-31	04-01 to 04-14	04-15 to 04-30
k	1	2	3	4	5
Liters per day	2.283	2.702	1.955	2.522	2.393
Liters per day (adjusted)	2.373	2.694	1.962	2.554	2.360
Price in €	1.712	1.870	2.063	1.830	1.837
Length of period (days)	44	14	22	14	15

Adjustment with respect to 2019 purchases based on GIE-CB data.

5-period model Constrained estimator: $\hat{\varepsilon}^c \approx -0.32$.

Naive estimator: $\hat{\varepsilon}^n \approx -1.55$.

F Testing for the presence of anticipations

We consider the following model:

$$q_{cty} = \beta p_{cty} + \sum_{h=t_2-\Delta}^{t_2+\Delta} \gamma_{hy} \mathbf{1}_{h=t} \mathbf{1}_{y=2022} + \xi_{cy} + \mu_t + \eta_{cty}. \quad (34)$$

We assume that the error term η_{cty} is made of two components: $\eta_{cty} = \nu_{ty} + \varepsilon_{cty}$. ν_{ty} corresponds to a shock on purchases that is common to all individuals or cells on day t in year y , while ε_{cty} is an idiosyncratic shock. We normalize $\nu_{ty=2021} = 0$ without loss of generality since μ_t is a daily fixed effect. The reason why we want to control further for ν_{ty} is that we observe sizeable differences in purchases in 2022, relative to 2021, which cannot be explained by price variations (even in periods with no anticipations, i.e., in periods far from any price shock). That model can be rewritten as:

$$\begin{aligned} q_{cty} &= \beta p_{cty} + \sum_{h=t_2-\Delta}^{t_2+\Delta} \gamma_{hy} \mathbf{1}_{h=t} \mathbf{1}_y + \xi_{cy} + \mu_t + \nu_{ty} + \varepsilon_{cty} \\ &= \beta p_{cty} + \sum_{h=t_2-\Delta}^{t_2+\Delta} (\gamma_{hy} + \nu_{hy}) \mathbf{1}_{h=t} \mathbf{1}_{y=2022} + \xi_{cy} + \mu_t \\ &\quad + \sum_{h \notin [t_2-\Delta, t_2+\Delta]} \nu_{hy} \mathbf{1}_{h=t} \mathbf{1}_{y=2022} + \varepsilon_{cty} \\ &= \beta p_{cty} + \sum_{h=t_2-\Delta}^{t_2+\Delta} (\tilde{\gamma}_{hy}) \mathbf{1}_{h=t} \mathbf{1}_{y=2022} + \xi_{cy} + \mu_t + \varepsilon_{cty} \end{aligned}$$

The estimation consists then of regressing q_{cty} on p_{cty} , the set of dummies $\mathbf{1}_{h=t} \mathbf{1}_{y=2022}$ where $h \in [t_2 - \Delta, t_2 + \Delta]$, daily fixed effects μ_t and cell fixed effects ξ_{cy} . Note that the coefficients $\tilde{\gamma}_{hy}$ correspond to the sum of the anticipation effects γ_{hy} and the 2022-day specific shock ν_{ty} .

F.1 Testing the model

First, we want to test for the presence of anticipations: can we reject the null that the coefficients γ_{hy} are all equal to 0? Second, we want to test that those anticipation effects

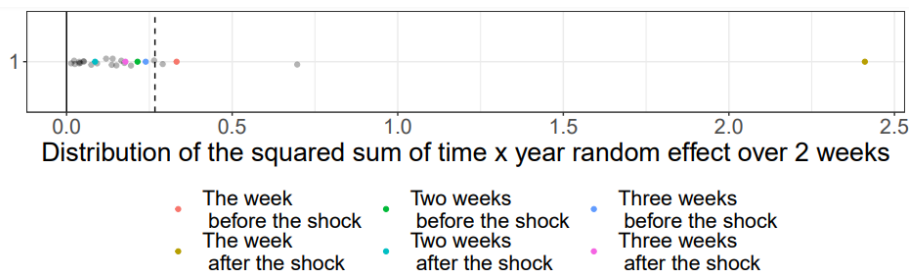
sum up to 0 around a price shock.

F.1.1 Testing the presence of anticipations

Testing for the presence of anticipations requires in fact slightly more than a naive test of γ_{hy} being equal to 0 since those coefficients are not directly estimated. Indeed, the econometrician recovers $\tilde{\gamma}_{hy}$, the sum of potential anticipation effects and day \times year shocks; even without any anticipation, those terms simplify to ν_{ty} , and are generally not equal to 0.

To nevertheless implement the test, we rely on placebo comparisons. We test for the presence of anticipations during weeks surrounding the shock. To do so, we consider the distribution of $\sum_h \tilde{\gamma}_{hy}^2$, which we compare to $\sum_t \nu_{ty}^2$. Specifically, we test whether the $\sum_{h'} \tilde{\gamma}_{h'y}^2$ at time h' , where we suspect anticipations, is higher than the 90th percentile of the distribution of $\sum_t \nu_{ty}^2$.

Figure F1: Testing the presence of anticipations in weeks surrounding the shock



Note. Grey points correspond to placebo estimates

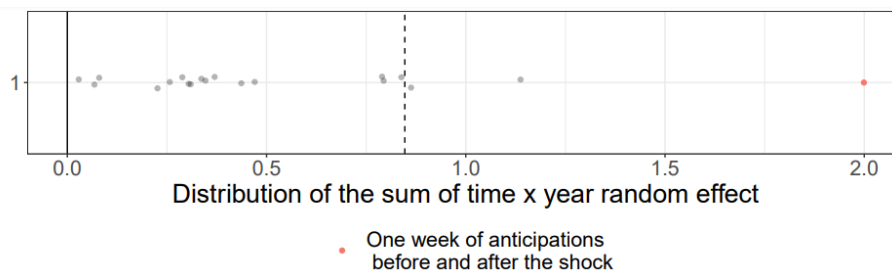
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

From Figure F1 one can reject the absence of anticipations during both the week before and the week after the shock with a 90% level of confidence.

F.1.2 Testing that anticipation effects sum up to 0

In the same vein, to test that the sum of anticipation is equal to 0, we compare the estimates of $\sum_{h=t_2-\Delta}^{t_2+\Delta} \tilde{\gamma}_{hy}$ to the distribution of $\sum_{t=t'-\Delta}^{t'+\Delta} \nu_{ty}$ (estimated in periods with no price shock, see below). If the sum is higher than the 95th centile or lower than the 5th centile of the distribution of $\sum_{t=t'-\Delta}^{t'+\Delta} \nu_{hy}$ in periods without anticipations, we reject the null that the sum is equal to 0.

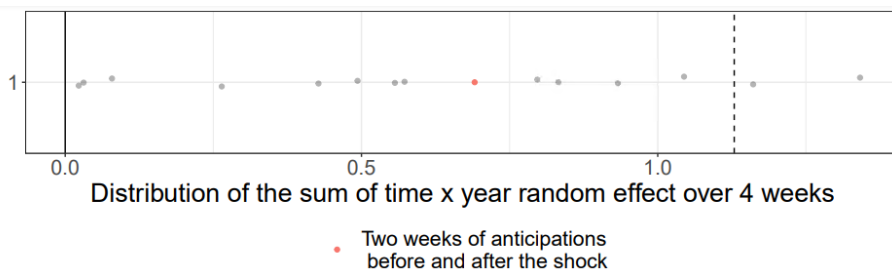
Figure F2: Testing whether the sum of anticipation is equal to 0 (period from 7 days before to 7 days after the shock)



Note. Grey points correspond to placebo estimates

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure F3: Testing whether the sum of anticipation is equal to 0 (period from 14 days before to 14 days after the shock)



Note. Grey points correspond to placebo estimates

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

From Figure F2 one can reject that the sum of anticipation effects is equal to 0 with a 90% level of confidence when we consider 2 weeks surrounding the shock (one week after and one week before); however, one cannot reject that this sum is equal to 0 when considering 4 weeks F3.

=> **In our main specification, we consider a bandwidth $\Delta = 14$ days with regard to the anticipation window.**

F.2 Placebo estimation

To estimate the distribution of ν_{ty} , we rely on periods with no price shock, hence with no anticipations. A first solution could be to estimate the following equation in the period ranging from mid-July to the beginning of October in 2019, 2020 and 2021 (excluding

therefore 2022 due to the price shock and the corresponding anticipations):

$$q_{cty} = \beta p_{cty} + \xi_{cy} + \mu_t + \eta_{cty}. \quad (35)$$

Unfortunately, β cannot be accurately estimated based on those periods, precisely because there is no price shock, hence no identifying variation. We thus remove that price part of the equation which is not predictive of purchases during that period (an alternative could be to set β at our estimated price sensitivity, yet our results are not sensitive to this empirical choice), and we estimate:

$$q_{cty} = \xi_{cy} + \mu_t + \eta_{cty} \quad (36)$$

We then collect the residuals $\hat{\eta}_{cty}$ and compute the average η_{cty} over cells c for each day t and year y , normalizing those terms to 0 in 2021.

G The aggregation bias: disposing of high-frequency data matters

In order to correctly infer the short-run price elasticity of fuel demand, the researcher should rely on: (i) exogenous price variations, (ii) high-frequency data, and (iii) a suited econometric method to control for consumer anticipations. It is perhaps pointless to illustrate how essential the first ingredient is to identification: removing exogenous price changes results in an identification failure (see our falsification test, for instance). We have also explained why taking anticipations into account is crucial so as not to confound them with the price effect (cf. naive estimations).

We thus aggregate our data at the monthly level and show that this aggregation is misleading. Though widely used due to the lack of more granular datasets, monthly data miss short-term variations in fuel purchases. As made clear by Figure 1, and as confirmed by Figure I7 in Online Appendix, illustrations of these unobserved variations include dips and spikes following anticipated tax changes. As a result, it is impossible for the econometrician to isolate the price effect. To quantify the magnitude of the aggregation bias, we replicate the identification strategy developed in section 4.1 based on monthly data (though without IV and without anticipation, by construction). We then obtain a higher price elasticity of demand, in absolute, namely -0.65 ,⁵² (see Column V of Table H2). When we aggregate our data at the weekly level (Column VI of Table H2), we obtain a point estimate of -0.09 , which is smaller, in absolute, than the one obtained at the daily level. In our view, this empirical evidence supports the claim that daily data are truly necessary to properly control for anticipations.

⁵²It is not possible to cluster standard errors in the time dimension in that case, hence we do not comment precision here.

H Robustness checks

First, we estimate a first-difference (FD) version of our model. Choosing the length of the FD operation is tricky: a daily FD would lack practical significance given that households need time to adjust their behavior.

With regard to our main estimation based on the September 1st price subsidy, we consider a FD corresponding to the time of an anticipation window, that is, the difference between what happens after and what happens before the price shock. We thus estimate:

$$q_{cy}^{post} - q_{cy}^{pre} = \beta(p_{cy}^{post} - p_{cy}^{pre}) + \eta_{cy} \quad (37)$$

where q_{cy}^{post} (q_{cy}^{pre}) corresponds to average purchases in cell c on year y after the anticipation window ending on September 14th (before the anticipation window starting on August 19th). We find an elasticity of -0.27, close to the one obtained with our preferred IV estimator (Table H2).

With regard to our estimations based on the period from September 2021 to February 2023, we difference out equation (12) at a monthly frequency. Reassuringly, this operation also yields a close estimate of -0.35.

Second, we document the sensitivity of our results with respect to the functional form by estimating a quasi-Poisson regression instead of a linear model. Such a parametric assumption is motivated by our dependent variable taking either null or positive values, on the one hand, and by the ease of interpretation of the price coefficient as a price elasticity, on the other hand.⁵³ In the same vein as our local identification strategy, we consider a model where $q_{cty} \sim \mathcal{P}(\lambda_{cty})$ in which we specify:

$$\log(\lambda_{cty}) = \varepsilon \log(p_{cty}) + \sum_{h=t_2-\Delta}^{t_2+\Delta} \gamma_{hy} \mathbb{1}_{h=t} \mathbb{1}_{y=2022} + \xi_{cy} + \mu_t. \quad (38)$$

The estimation is again subject to the constraint (11), and proceeds from maximum likelihood. We may also consider the period from September 2021 to February 2023. In that

⁵³In the quasi-Poisson regression, $\frac{\partial \log(\mathbb{E}q)}{\partial \log(p)} = \frac{\partial \log(\lambda)}{\partial \log(p)} \equiv \varepsilon$, which refers to the price-elasticity of the average demand.

case, we posit $q_{ct} \sim \mathcal{P}(\lambda_{ct})$ with

$$\log(\lambda_{ct}) = \varepsilon \log(p_{ct}) + \sum_{h=t_0}^{t_1+\Delta} \gamma_h^1 \mathbb{1}_{h=t} + \sum_{k=2}^4 \sum_{h=t_k-\Delta}^{t_k+\Delta} \gamma_h^k \mathbb{1}_{h=t} + \xi_c + \tilde{\mu}_t. \quad (39)$$

The estimation is now subject to the set of constraints (13), and still proceeds from maximum likelihood. In both cases, standard errors are computed by two-way clustering at cell and day levels.

Empirically, the choice of the functional form sounds rather innocuous: when opting for the quasi-Poisson regression instead of a linear model, we obtain estimates in Columns I to III of Table H2, which compare well to the ones in the same columns of Table 2. Row 3 of Table H6 provides the results based on the period from September 2021 to February 2023, which do not differ much from the baseline (Row 1). Moreover, the (quite high) quality of the prediction does not depend much on the parametric specification adopted.


Third, to alleviate any concern about our fuel price index, we replace it with the price of the diesel. Our results are not much affected (Columns II and V of Table H3 for the local estimation as well as Row 6 of Table H6 for the estimation from September 2021 to February 2023). Montag et al. (2021) find that diesel drivers are better informed about prices, hence perfect compliance is more likely for them. The fact that our results remain unaltered when we instrument prices by the tax change in diesel prices alleviates the concern about imperfect information (and possibly heterogeneity of information among consumers).⁵⁴

Fourth, the anticipation window may well be heterogeneous among consumers. For instance, occasional drivers might have longer anticipation windows, as suggested by Figure H3, which depicts the distribution of the interpurchase duration (computed outside anticipation windows) for the four different groups of fuel spending. A possible solution to this issue is to allow for the anticipation bandwidth to vary with fuel spending (Table H7). Even in the case where drivers in the bottom 25% of fuel spending have an anticipation window as large as one month, their estimated price sensitivity is higher, which is in line with previous findings. Moreover, their estimated price elasticity is not significantly dif-

⁵⁴In the same vein, since heterogeneity in information (and pass-through) arises along the price dimension itself, one could think of another instrument like the change in the minimum or average price in the *département*. To avoid outliers, we use the price located at the 10th percentile of the distribution rather than the minimum price. This operation does not dramatically affect our results (Table H3).

ferent from -0.82, the one obtained with a one-week anticipation bandwidth. Finally, the same pattern of heterogeneity is obtained regardless of the choice of the bandwidth.

Figure H1: Fake price subsidy on September 1st, 2021 (baseline year: 2019)



Note. Fuel prices (including taxes). The dashed line corresponds to a fake price subsidy on September 1st, 2021.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

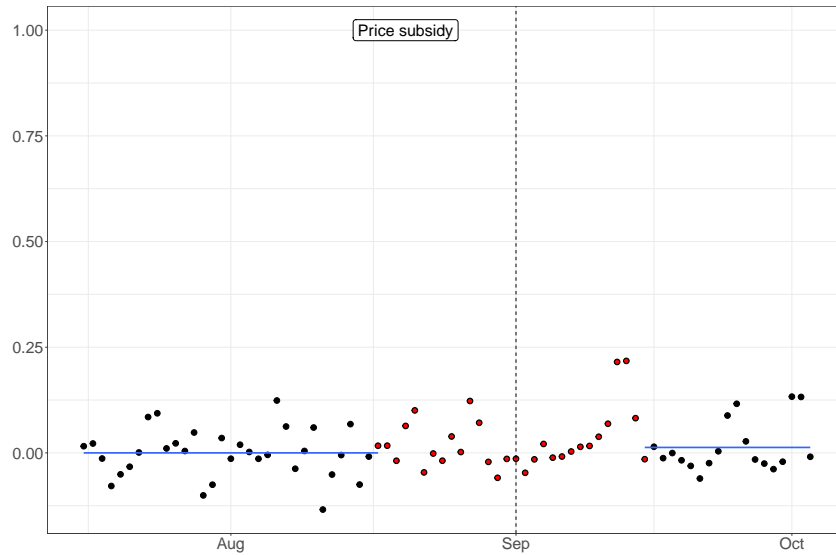
Table H1: Falsification test (Fake rebate on September 1st, 2021)

	I	II	III
price coefficient	-0.97 (1.23)	-1.58 (1.22)	-1.49 (1.41)
price elasticity	-0.53 (0.67)	-0.86 (0.67)	-0.81 (0.77)
Cell FE	✓	✓	✓

Note. Estimation sample: 10,767 cells of customers. Estimation period: from July, 16th to October, 3th 2021. Baseline year: 2019. OLS estimates are reported here. Weighted regressions according to the size of the sample within each cell. Two-way clustering of standard errors at cell and day levels.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure H2: Fuel purchases around September 1st, 2021 (baseline year: 2019)



Note. Dots correspond to adjusted daily fuel purchases from July 15th to October 4th, 2021. The adjustment for seasonal variation relies on 2019 as the baseline year. Purchases are normalized so that they sum up to 0 before the anticipation window. Dashed line: (Fake price subsidy on) September 1st, 2021. Red dots: Anticipation window (7 days before and after September 1st, 2021). Blue lines: Average purchases before and after September 1st, 2021, excluding the anticipation window.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

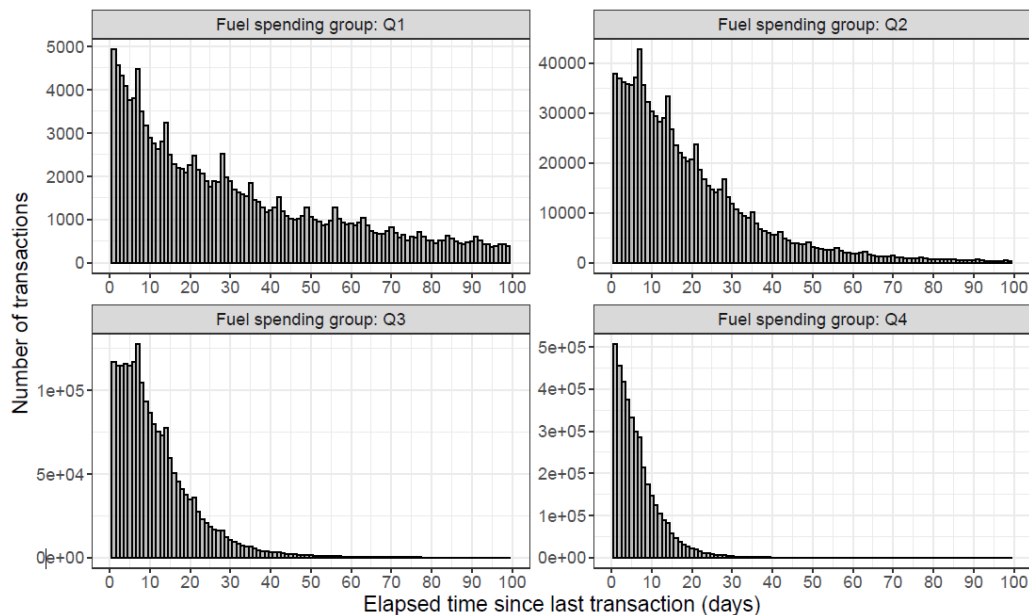
Table H2: Robustness checks - Estimation based on September 1st tax rebate

	I	II	III	IV	V	VI	VII
price elasticity	-0.41 (0.07)	-0.16 (0.07)	-0.19 (0.07)	0.96 (0.38)	-0.65 (.)	-0.09 (.)	-0.27 (.)
Quasi-Poisson regression	✓	✓	✓				
Linear model				✓	✓	✓	
First difference							✓
Anticipation dummies		✓			✓		
Excluding anticipation window			✓			✓	✓
Monthly aggregation					✓		
Weekly aggregation						✓	
Cell FE	✓	✓	✓	✓	✓	✓	
Day FE	✓	✓	✓	✓			
Week FE						✓	
Month FE					✓		
# of cells	10,777	10,777	10,777	10,777	10,777	10,777	10,777

Note. Estimation of equation (10) with a sample of 10,777 cells. Estimation period: customers observed from July, 15th to October, 4th. Baseline year: 2021. Weighted regressions according to the size of the sample within each cell. Two-way clustering of standard errors at cell and year-day levels.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure H3: Interpurchase duration (by group of fuel spending, in days)



Note. Elapsed time between two transactions from September 2021 to February 2023.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Table H3: Robustness checks with different fuel prices

	I	II	III	IV	V	VI
	Index derived from average prices	Average diesel price	Index derived from 1st decile of prices	Index derived from average prices	Average diesel price	Index derived from 1st decile of prices
price coefficient	-0.29 (0.10)	-0.28 (0.10)	-0.32 (0.08)	-0.43 (0.11)	-0.52 (0.14)	-0.38 (0.09)
price elasticity	-0.19 (0.07)	-0.20 (0.07)	-0.23 (0.06)	-0.31 (0.08)	-0.37 (0.10)	-0.27 (0.06)
IV (Instrument: post- 9/1 dummy)				✓	✓	✓
Anticipation dummies	✓	✓	✓	✓	✓	✓
Excluding anticipation window						
Cell FE	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓
# of cells	10,777	10,777	10,777	10,777	10,777	10,777

Note. Estimation of equation (10) with a sample of 10,777 cells of customers. Estimation period: from July, 15th to October, 4th. Baseline year: 2021. Weighted regressions according to the size of the sample within each cell. Two-way clustering of standard errors at cell and year-day levels. Columns I and IV correspond to the regression of quantities on a fuel price index based on mean prices by *département*; Columns II and V: based on diesel prices; Columns III and VI: based on the first decile of prices within the *département*.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Table H4: Further heterogeneity by location

	Urban	Periurban	Rural	Paris	All
price elasticity	-0.32 (0.12)	-0.27 (0.07)	-0.38 (0.07)	-0.67 (0.23)	-0.31 (0.08)
Cell FE	✓	✓	✓	✓	✓

Note. Estimation sample: 10,777 cells of customers. Estimation period: from July, 15th to October, 4th 2022. Baseline year: 2021. Constrained IV estimates are reported here. Weighted regressions according to the size of the sample within each cell. Two-way clustering of standard errors at cell and day levels.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Table H5: Estimation based on the period following the invasion of Ukraine

	I	II	III
price elasticity	-0.73 (0.16)	-0.18 (0.07)	-0.22 (0.29)
Anticipation dummies		✓	
Excluding anticipation window			✓
Cell FE	✓	✓	✓
# of cells	10,754	10,754	10,754

Note. Estimation sample: 10,826 cells of customers. Estimation period: from January, 10th to April, 30th 2022. Weighted regression according to the size of the sample within each cell. Two-way clustering of standard errors at cell and day levels.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Table H6: Robustness checks (Estimation based on the period from September 2021 to February 2023)

	I	II	III	IV	V	VI
OLS estimates	-0.46 (0.05)	-0.52 (0.05)	-0.45 (0.07)	-0.35 (0.04)	-0.42 (0.04)	-0.26 (0.04)
IV estimates	-0.38 (0.06)	-0.42 (0.06)	-0.41 (0.08)	-0.29 (0.04)	-0.37 (0.04)	-0.22 (0.04)
Quasi-Poisson regression	-0.45 (0.04)	-0.52 (0.05)	-0.45 (0.07)	-0.37 (0.03)	-0.46 (0.03)	-0.27 (0.04)
Before war (2022-02-24)	-0.80 (0.11)	-1.04 (0.11)	-0.33 (0.26)	-0.80 (0.11)	-1.04 (0.11)	-0.33 (0.26)
Sole diesel price	-0.50 (0.05)	-0.55 (0.05)	-0.42 (0.07)	-0.41 (0.04)	-0.47 (0.04)	-0.22 (0.04)
Monthly estimates	-0.51 (0.47)					
Anticipation dummies				✓	✓	✓
Seasonality controls		✓	✓		✓	✓
Linear trend			✓			✓
Cell FE	✓	✓	✓	✓	✓	✓

Note. Estimation of equation (12) with a sample of 11,031 cells of customers. Estimation were made on one fifth of the sample for computational issues. Results correspond to median estimates over 25 replications. Estimation period: from September 2021 to February 2023. Weighted regression according to the size of the sample within each cell. Two-way clustering of standard errors at cell and day levels.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Table H7: Heterogeneous anticipation bandwidths with respect to fuel spending


	Q1 fuel spending	Q2 fuel spending	Q3 fuel spending	Q4 fuel spending	All
no anticipation	-1.39 (0.24)	-0.95 (0.17)	-0.85 (0.16)	-0.65 (0.12)	-0.76 (0.13)
7 days of anticipation	-1.06 (0.24)	-0.52 (0.13)	-0.43 (0.10)	-0.32 (0.07)	-0.40 (0.08)
14 days anticipation	-0.82 (0.22)	-0.35 (0.12)	-0.34 (0.09)	-0.26 (0.08)	-0.31 (0.08)
21 days anticipation	-0.52 (0.19)	-0.17 (0.11)	-0.14 (0.09)	-0.06 (0.07)	-0.11 (0.07)
28 days anticipation	-0.71 (0.16)	-0.32 (0.09)	-0.21 (0.10)	-0.26 (0.07)	-0.27 (0.08)
Cell FE	✓	✓	✓	✓	✓

Note. Estimation of equation (10) with a sample of 10,777 cells of customers. Estimation period: from July, 15th to October, 4th. Constrained IV estimates are reported here for various bandwidths of the anticipation window. Weighted regressions according to the size of the sample within each cell. Two-way clustering of standard errors at cell and day levels.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

I Supplementary figures and tables

Figure II: Distribution of fuel spending

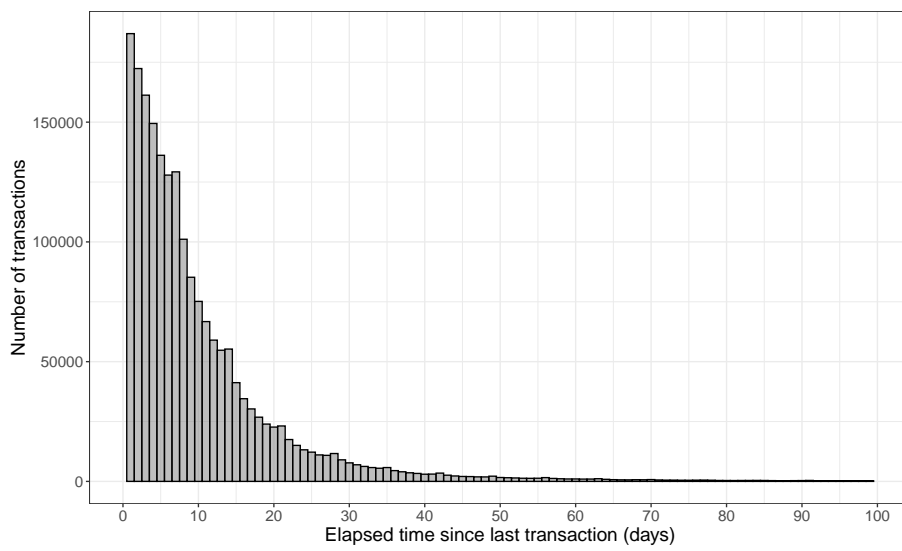


Note. Transaction-level fuel expenditures.

Lecture. 400,000 transactions amount to between 50 and 51 euros.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

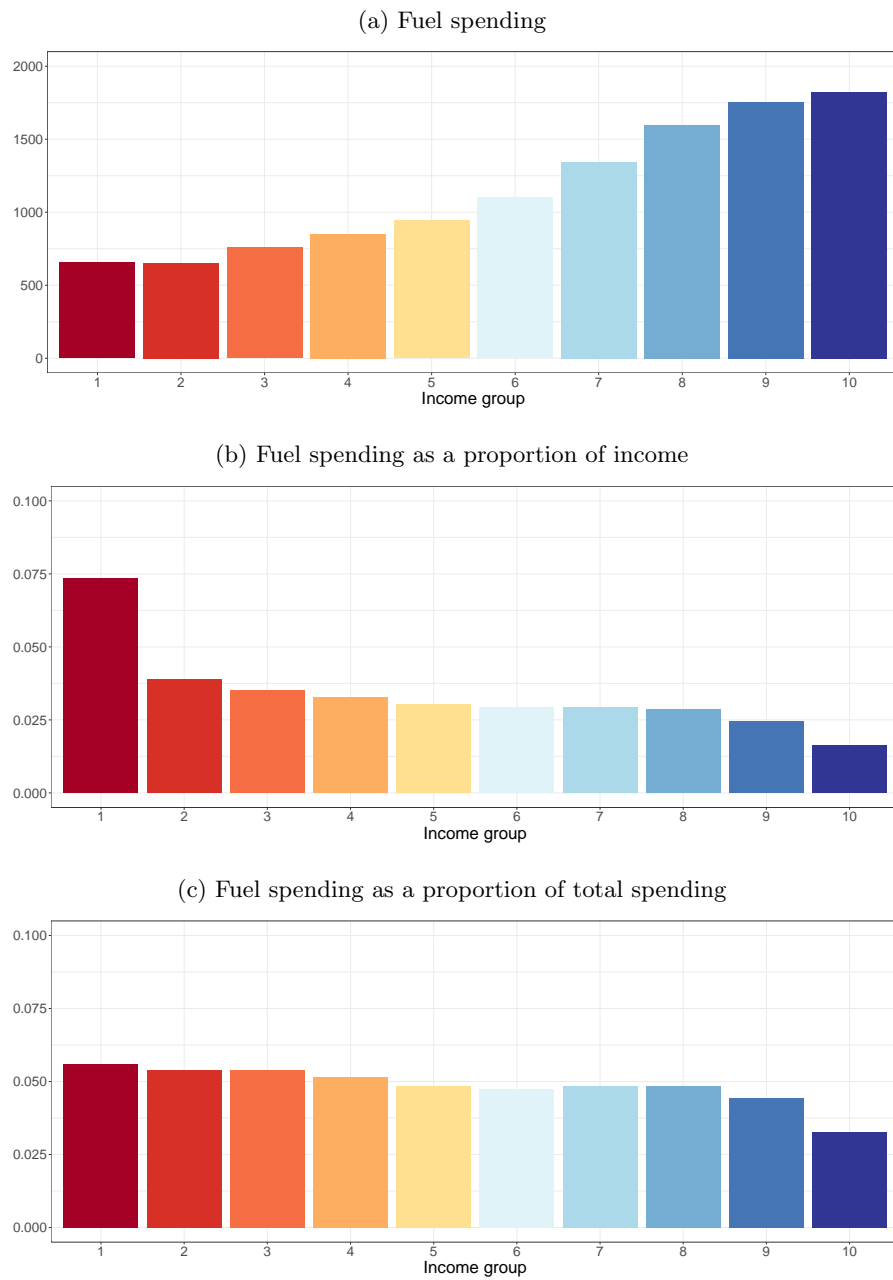
Figure I2: Interpurchase duration (in days)



Lecture. 75,000 transactions have occurred 10 days since last purchase. For 2,500 households, the average duration between two transactions is 10 days.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure I3: Average fuel spending of motorists by income group (deciles)

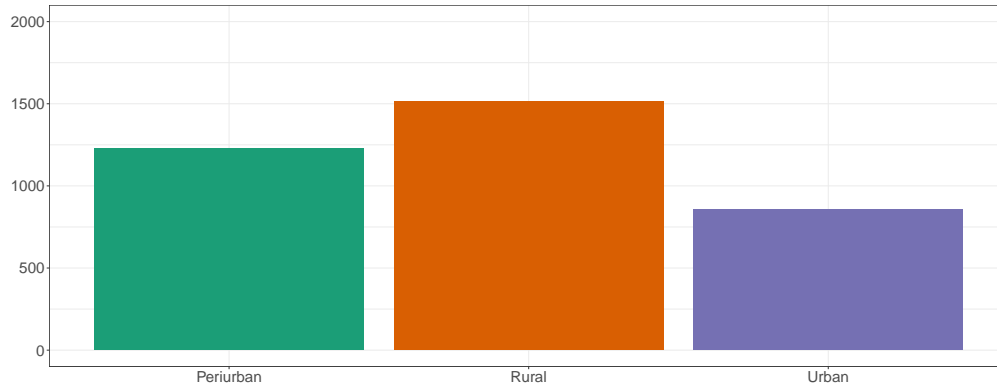


Note. Fuel expenditures: card payments in gas stations. Total expenditures: both card payments and checks. Location: peri-urban (P), rural (R), or urban (U). Fuel spending increases with income. The budget share of fuel decreases with income; it is higher in rural areas. The sample is restricted to motorists defined as households who buy fuel at least once.


Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure I4: Average fuel spending of motorists by location


(a) Fuel spending



(b) Fuel spending as a proportion of income



(c) Fuel spending as a proportion of total spending

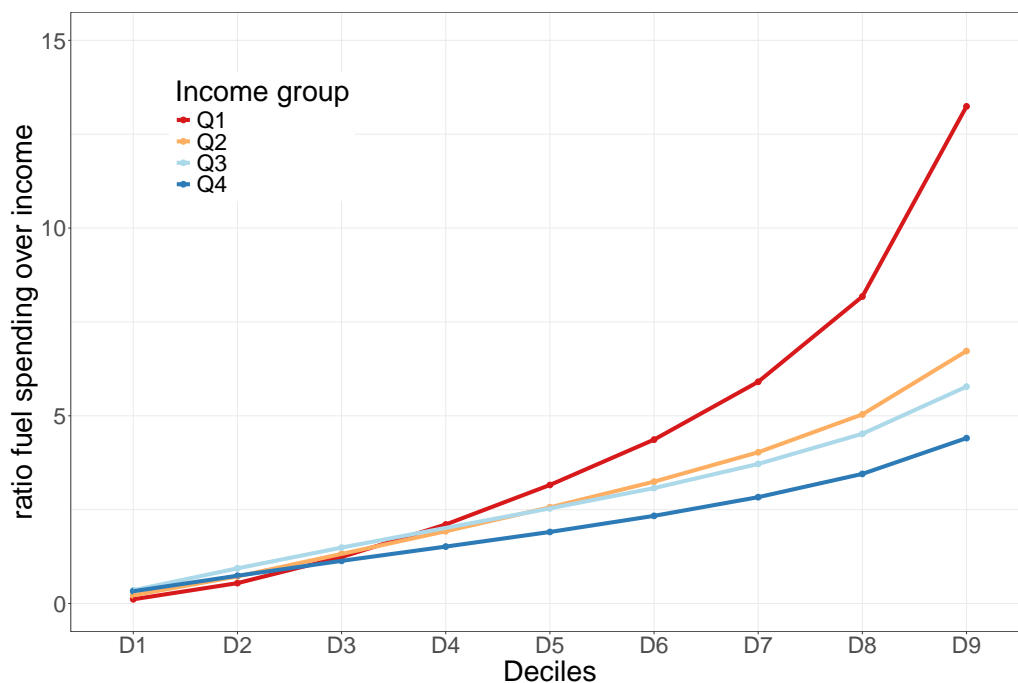


Note. Fuel expenditures: card payments in gas stations. Total expenditures: both card payments and checks. Location: peri-urban (P), rural (R), or urban (U). The sample is restricted to motorists defined as households who buy fuel at least once.

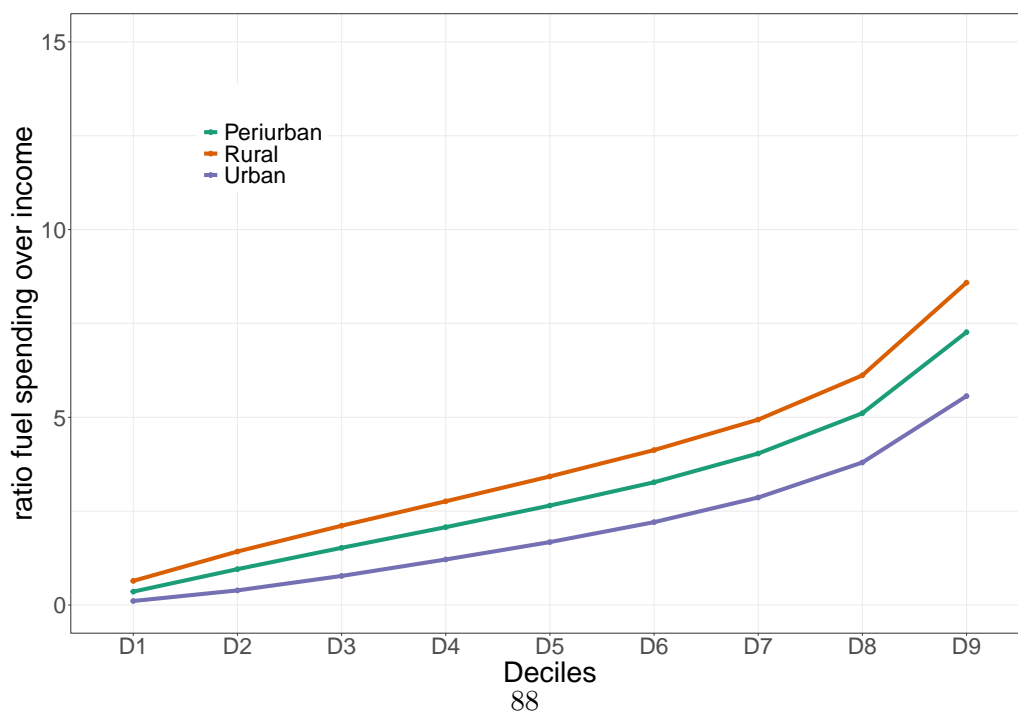
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure I5: Deciles of fuel budget shares

(a) By income group (quarters)



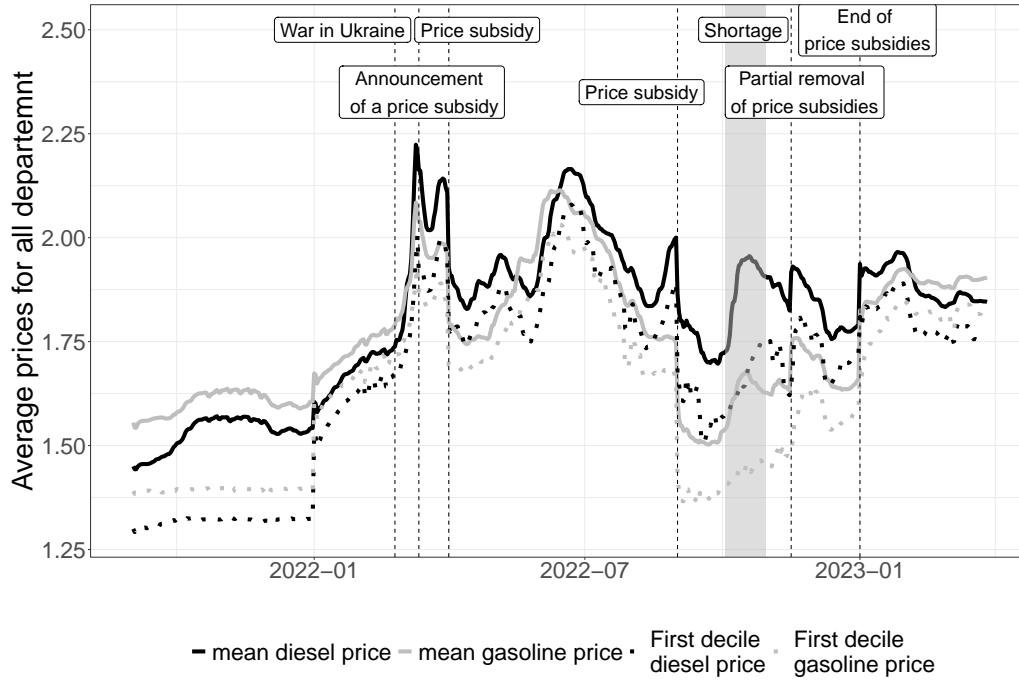
(b) By location



Note. Fuel expenditures: card payments in gas stations. Budget shares are equal to fuel spending divided by income 2021. Location: peri-urban (P), rural (R), or urban (U). The sample is restricted to households who buy fuel at least once.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

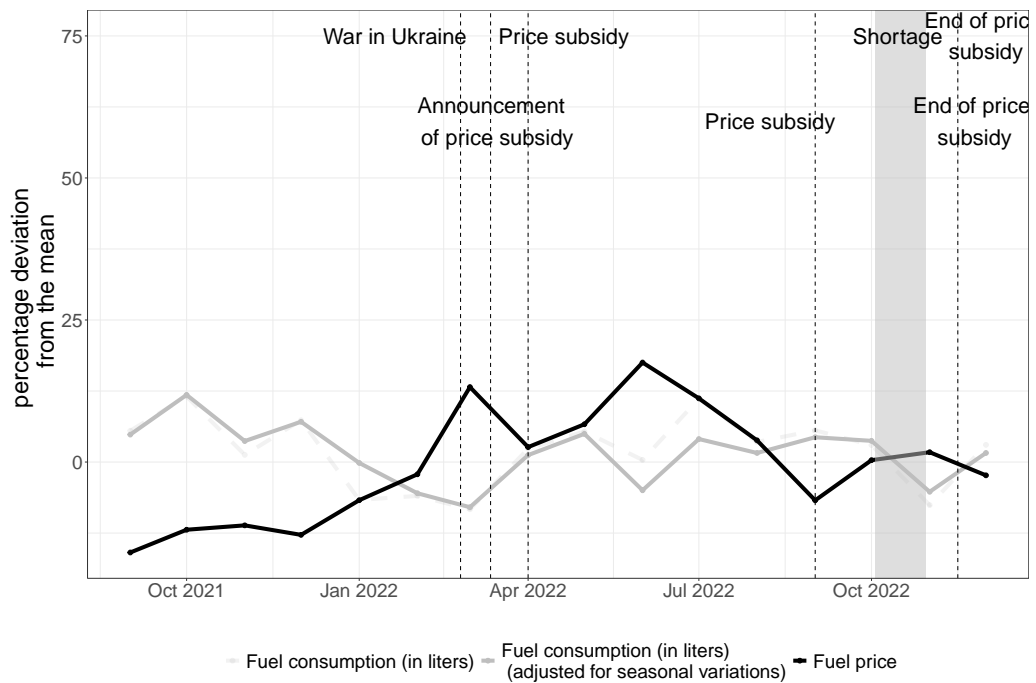
Figure I6: Diesel and gasoline prices



Note. In each *département*, we consider mean diesel price over all stations, mean SP95-E10 price, and the first decile. In the graph we average these quantities at the national level. All prices include taxes. Dashed lines correspond to the invasion of Ukraine and policy interventions. The policy intervention of April 1st amounts to a €0.18 per liter tax rebate (including VAT). The policy intervention of September 1st amounts to an extra €0.12 per liter subsidy, which has prevailed until mid-November 2022. The remaining subsidy was removed on January 1st 2023.

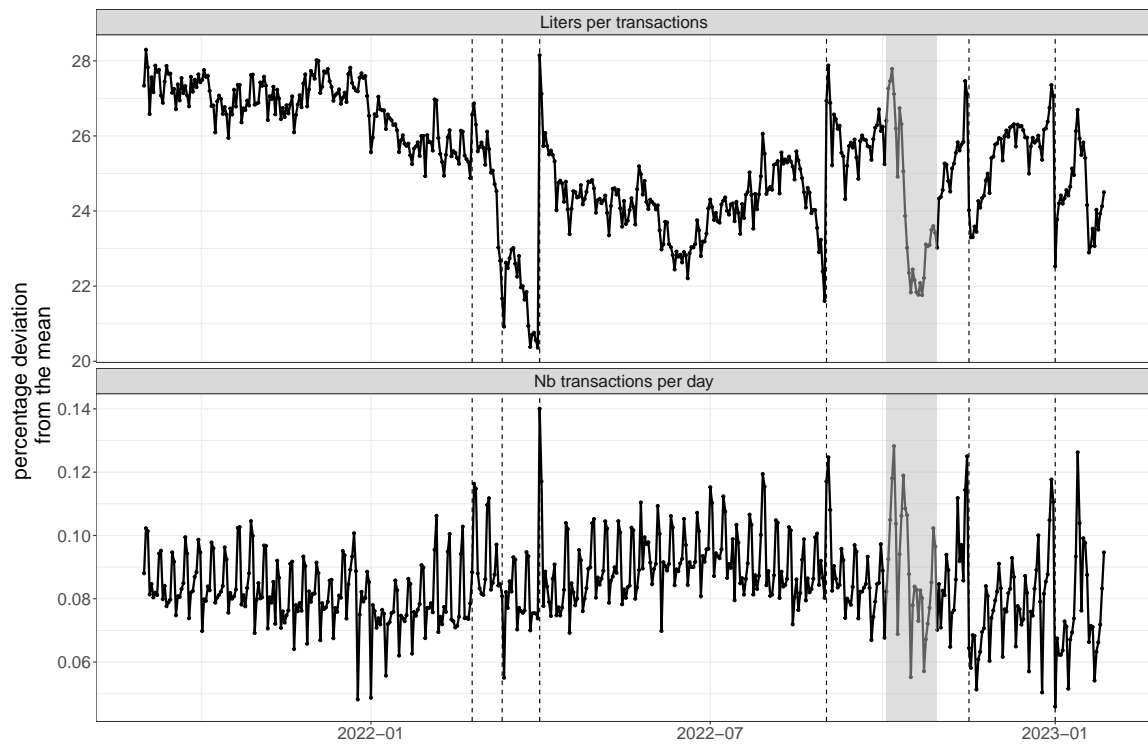
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure I7: Monthly aggregation of fuel prices and purchases




Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure I8: Intensive vs extensive margins (quantity per transaction vs # of transactions)



Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.


Figure I9: Fuel prices from January to April 2022



Note. Fuel prices (including taxes) and purchases (in liters). Purchases adjusted for seasonal variations thanks to GIE-CB data from January 8th 2022 to April 30th 2022. The first dashed line corresponds to the invasion of Ukraine, the second dashed line refers to the announcement of the first policy intervention, a subsidy of €0.18 per liter (including VAT), and the last dashed line indicates the effective implementation of the intervention.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.


Figure I10: Estimated anticipation parameters (γ coefficients)



Note. The solid line corresponds to the γ coefficients of equation (12) estimated from September to February. These coefficients capture purchases due to anticipatory behavior, in liters; in each anticipation window, the coefficients sum up to 0. Dashed lines correspond to the invasion of Ukraine and policy interventions. The policy intervention of April 1st amounts to a €0.18 per liter tax rebate (including VAT). The policy intervention of September 1st amounts to an extra €0.12 per liter subsidy, which has prevailed until mid-November 2022. The remaining subsidy was removed on January 1st 2023.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.


Figure I11: Predicted vs actual demand for fuel



Note. The black line corresponds to observed fuel purchases (in liters) over the period. The grey line corresponds to predicted fuel purchases based on equation (12). Dashed lines correspond to the invasion of Ukraine and policy interventions. The policy intervention of April 1st amounts to a €0.18 per liter tax rebate (including VAT). The policy intervention of September 1st amounts to an extra €0.12 per liter subsidy, which has prevailed until mid-November 2022. The remaining subsidy was removed on January 1st 2023.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.


Figure I12: Heterogeneity of the price elasticity with respect to income



Note. Estimated price elasticity using equation (10). Estimation period: from July 15th to October 4th. Corresponding subsamples depend on households' characteristics (observed before the intervention from January to June, for continuous variables, and in June for discrete variables). Weighted regression according to the size of the sample within each cell. Two-way clustering of standard errors at cell and year-day levels.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

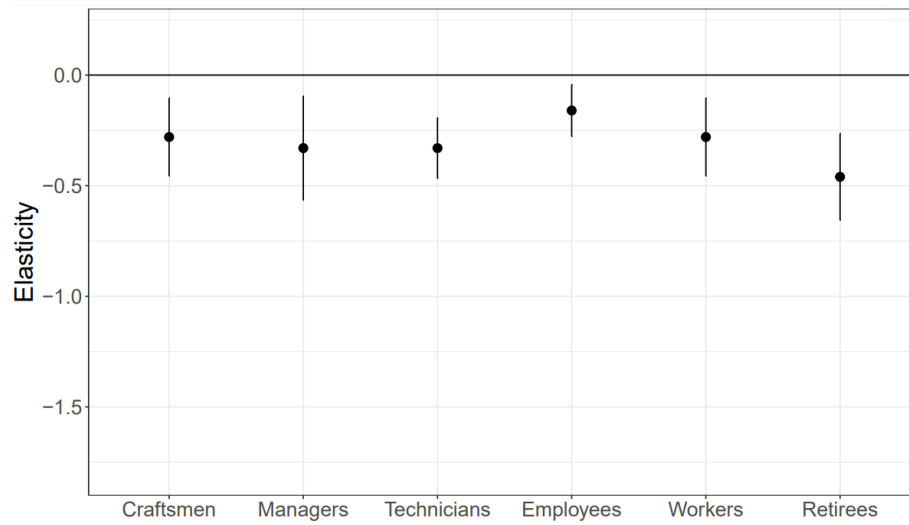
Figure I13: Heterogeneity of the price elasticity with respect to total card spending



Note. Estimated price elasticity using equation (10). Estimation period: from July 15th to October 4th. Corresponding subsamples depend on households' characteristics (observed before the intervention from January to June, for continuous variables, and in June for discrete variables). Weighted regression according to the size of the sample within each cell. Two-way clustering of standard errors at cell and year-day levels.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.


Figure I14: Heterogeneity of the price elasticity with respect to occupation



Note. Estimated price elasticity using equation (10). Estimation period: from July 15th to October 4th. Corresponding subsamples depend on households' characteristics (observed before the intervention from January to June, for continuous variables, and in June for discrete variables). Weighted regression according to the size of the sample within each cell. Two-way clustering of standard errors at cell and year-day levels.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

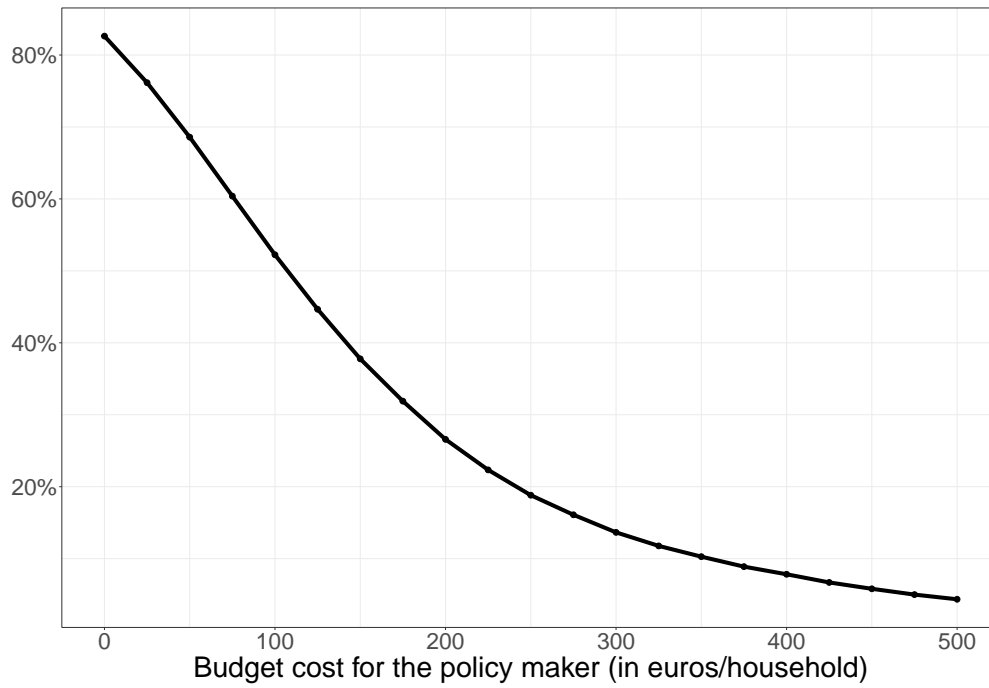
Figure I15: Comparing policy instruments (Objective function: Alleviating excessive relative losses)



Note. For any given level of public funding (x-axis), we compare (conditional or unconditional) price subsidies and transfers as policy tools designed to compensate households for the observed price increase from 2021 to 2022 (+€0.5 per liter, i.e. from €1.5 to €2).

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure I16: Share of low-income households within the top 10% relative losses



Note. The share of low-income households (i.e., whose income is below the median) within the top 10% relative losses is computed for any level of public funding (x-axis).

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

J An evaluation of actual fuel price subsidies: Financial, distributive and environmental aspects

We assess here financial and distributional impacts of the fuel tax policy as well as its effect on CO₂ emissions. To that end, we first simulate a counterfactual that would have prevailed in the absence of excise tax rebates. To evaluate the impact of the sole public interventions, we assume full pass-through of tax changes to consumers. We predict fuel spending \tilde{q}_{ct} at prices $\tilde{p}_{ct} = p_{ct} + \Delta p_t$ from January 8th, 2022 (\underline{t}) to January 8th, 2023 (\bar{t}).⁵⁵ The after-tax price differential Δp_t is equal to zero until the end of March 2022, then amounts to +€0.18 per liter from April 1st to the end of August, and up to +€0.30 per liter only from September 1st onwards; it is then reduced to +€0.10 per liter from November, 16th until the end of 2022 when it vanishes.

We then evaluate the impact of the policy on fuel purchases, in liters, by computing the difference between observed and simulated demand:

$$\sum_{t=\underline{t}}^{\bar{t}} [q_{ct} - \tilde{q}_{ct}] = \sum_{t=\underline{t}}^{\bar{t}} \hat{\beta}(p_{ct} - \tilde{p}_{ct}) = - \sum_{t=\underline{t}}^{\bar{t}} \hat{\beta}(\Delta p_t) > 0, \quad (40)$$

noting that anticipation effects cancel out over each anticipation window but the first one. The change in fuel spending is computed as follows:

$$\sum_{t=\underline{t}}^{\bar{t}} [p_{ct}q_{ct} - \tilde{p}_{ct}\tilde{q}_{ct}] = - \sum_{t=\underline{t}}^{\bar{t}} (\Delta p_t)\tilde{q}_{ct} - \sum_{t=\underline{t}}^{\bar{t}} \hat{\beta}(\Delta p_t)p_{ct} + \sum_{t=\underline{t}}^{\bar{t}} \hat{\gamma}_t p_{ct}, \quad (41)$$

which makes clear that β and $\gamma = (\gamma^1, \gamma^2, \gamma^3, \gamma^4)$ are sufficient statistics for this evaluation exercise. Three effects are at stake: (i) $-\sum_{t=\underline{t}}^{\bar{t}} [(\Delta p_t)\tilde{q}_{ct}] < 0$ corresponds to the *mechanical*

⁵⁵For each cell of individuals and for each day, we may compute $\tilde{q}_{ct} = \hat{\beta}\tilde{p}_{ct} + \sum_{h=t_0}^{t_a-1} \hat{\gamma}_h^1 \mathbb{1}_{t=h} + \sum_{h=t_a}^{t_a-1+(t_a-t_0)} (-\hat{\gamma}_{t_a-h+t_a-1}^1) \mathbb{1}_{t=h} + \hat{\xi}_c + \hat{\mu}_t + \hat{\eta}_{ct}$ from previous estimates. In the absence of any sharp, policy-driven price change as is the case for the latter three anticipation windows and the second part of the first anticipation window, anticipation effects should be neutralized. During the first anticipation window, we assume that (stored) fuel purchases observed during the period from t_0 to March 10th, the day before the announcement of the first price subsidy, denoted by $(t_a - 1)$, would have been exactly compensated the days after, from t_a to $t_a - 1 + (t_a - t_0)$ according to some opposite and symmetric scheme. Note that the latter assumption is unimportant to our policy evaluation exercise: it is only required that default purchases during the rest of that window, from t_a to $t_a - 1 + (t_a - t_0)$, exactly compensate excess purchases from t_0 to $t_a - 1$, corresponding to storage.

effect, namely the direct effect of the price subsidy on fuel spending, consumption being fixed; (ii) the *behavioral effect* $-\sum_{t=\underline{t}}^{\bar{t}}[(\hat{\beta}\Delta p_t)p_{ct}] > 0$ corresponds to the impact of the increase in consumption on spending, following reduced prices; (iii) the *anticipatory effect* $\sum_{t=\underline{t}}^{\bar{t}}\hat{\gamma}_t p_{ct}$ is related to the fact that storing (postponing) fuel purchases when prices are low (high) does not alter total consumption, but may increase or decrease spending depending on how prices evolve over time. Since there is no anticipation at the exclusion of anticipation windows, the latter term is negligible in an annual policy evaluation. Formally, the latter term rewrites:

$$\sum_{t=\underline{t}}^{\bar{t}}\hat{\gamma}_t p_{ct} = \sum_{t=\underline{t}}^{\bar{t}}\hat{\gamma}_t \tilde{p}_{ct} - \sum_{t=\underline{t}}^{\bar{t}}\hat{\gamma}_t \Delta p_t \approx -\sum_{t=\underline{t}}^{\bar{t}}\hat{\gamma}_t \Delta p_t \quad (42)$$

because counterfactual prices \tilde{p}_{ct} do not vary much during anticipation windows, contrary to observed prices p_{ct} ; the first of the two terms in the decomposition (42) is almost equal to the average counterfactual price during each anticipation window, multiplied by the sum of anticipation effects over that period, i.e., zero, hence it can be neglected.⁵⁶

Based on the marginal price effect $\hat{\beta} \approx -0.43$ corresponding to the average elasticity $\hat{\varepsilon} \approx -0.31$, we estimate that the financial impact of the policy has been to reduce fuel spending by €66 per household, on average, in 2022: this economy represents 0.14% of the average income and 4.8% of the annual fuel bill. The *mechanical effect* amounts to a €109 reduction in fuel spending, while the *behavioral effect* is estimated to a €43 increase: this countervailing response has therefore attenuated the *mechanical effect* by about 39%.

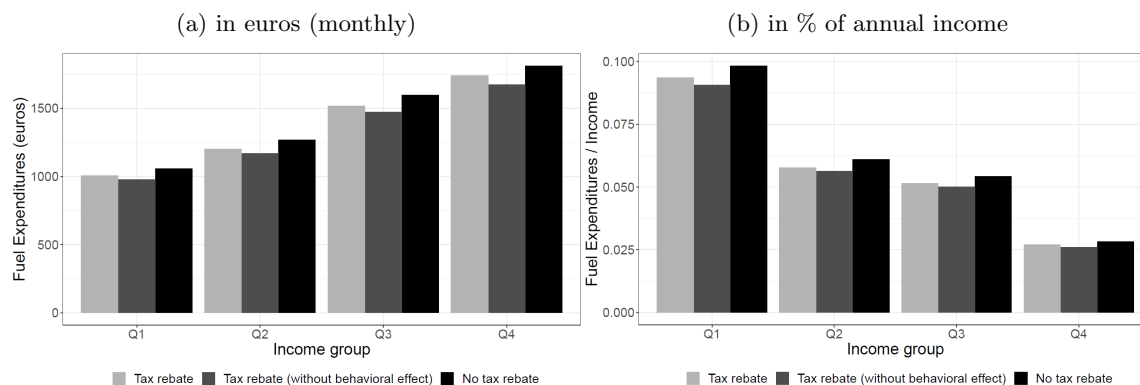
To quantify distributional effects at play, we further allow for $\hat{\beta}$ to vary depending on the same observed characteristics as in section 5.3 (Figure J1). The impact of the policy ranges from €51 saved by the bottom 25% of income to €71 saved by the top 25% of income. Those figures represent 0.47% of income for the former and 0.11% of income for the latter⁵⁷. Figure J2 confirms that households living in rural areas, whose share of income devoted to fuel expenditures is higher, benefited more from the policy in nominal terms (especially low-income households as shown by Figure J3).

⁵⁶Empirically, this effect is of same magnitude as the price effect, but during anticipation windows only: as a result, it does not matter much in that annual evaluation exercise, despite the importance of anticipation effects $\hat{\gamma}$ in the estimation.

⁵⁷The former devote 9.35% of their budget to fuel expenditures, and that share would have increased to 9.82% in the absence of any intervention. For the latter, the corresponding figures are 2.71% and 2.82%, respectively.

Last, the impact of the policy on CO₂ emissions has been rather limited. The extra fuel consumption amounts to 24 liters per household, an increase of +3.3%. This effect displays substantial heterogeneity, though: it amounted to 36 liters for the top 25% of fuel consumption but to 16 liters only for the bottom 25%. Based on the observed fuel mix between gasoline and diesel, we estimate that this supplementary consumption represents 73 kilograms of CO₂.⁵⁸ In 2021, the annual carbon footprint of a French household amounted to about 20.3tons: the policy increased that footprint by 0.36%.

Figure J1: Distributional effects of price subsidies (by income)

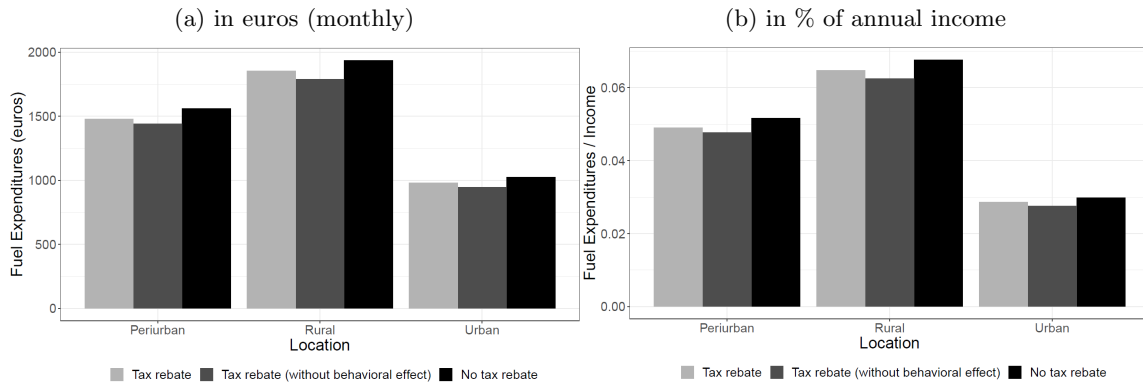


Note. Simulation period: from January 8th 2022 to January 8th 2023.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

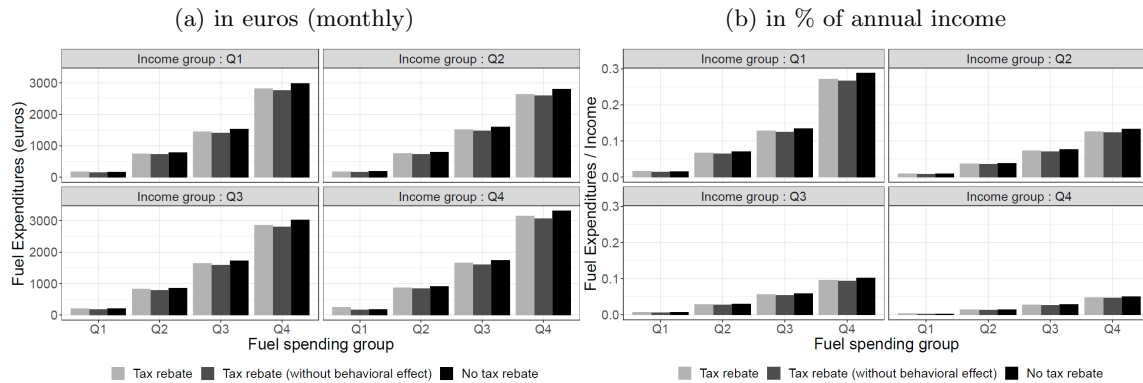
⁵⁸In the polar case of pure diesel, corresponding estimates would amount to 76 kilograms. In the polar case of pure gasoline, they would amount to 67 kilograms.

Figure J2: Distributional effects of price subsidies (by location)



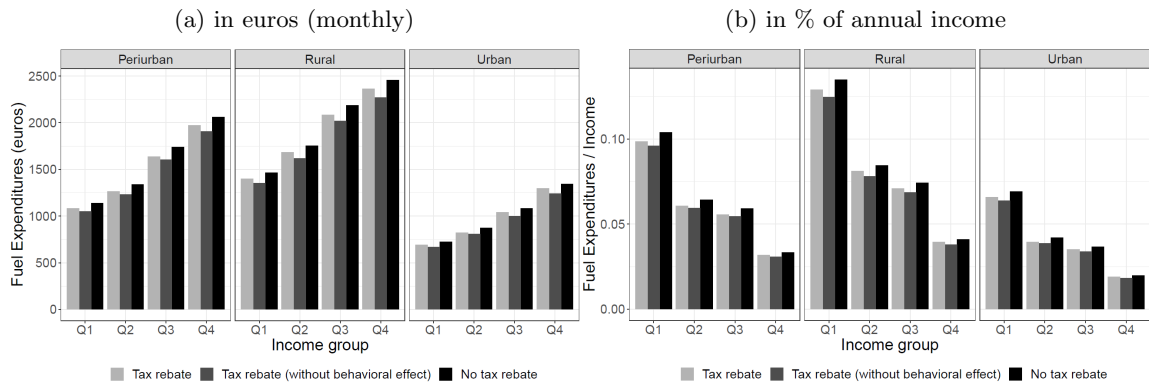
Note. Simulation period: from January 8th 2022 to January 8th 2023. Location is defined by the bank.
 Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure J3: Distributional effects of price subsidies (by fuel spending and income)



Note. Estimation period: from January 8th 2022 to January 8th 2023.
 Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure J4: Distributional effects of price subsidies (by income and location)



Note. Estimation period: from January 8th 2022 to January 8th 2023.

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

K Policy analysis under quasi-linear preferences

We consider here that households' preferences can be represented by some quasi-additive utility function $U(f, m) = m + B \frac{f^{1+\frac{1}{\varepsilon}} - 1}{1+\frac{1}{\varepsilon}}$ over annual fuel consumption f and the outside good m taken as the *numéraire*, with parameters B indexing the household h 's need for driving and $\varepsilon < 0$. We omit the unnecessary index h , though households may differ in their price-elasticity as well as in their need for driving.⁵⁹ Given annual disposable income Y and the price of fuel p , agents maximize their utility subject to the budget constraint: $pf + m \leq Y$, which leads to an iso-elastic fuel demand: $f(p) = \left(\frac{p}{B}\right)^\varepsilon$ with price elasticity ε . Fuel expenditures amount to $E_f(p) = B^{-\varepsilon} p^{1+\varepsilon}$. Following any rise in fuel prices from p to $p' > p$, the policy maker might want to compensate consumers based on the corresponding compensating variation, which coincides with the change in consumer surplus CS because preferences are quasi-linear. Denoting by $l_\varepsilon(x) = \frac{1-(1+x)^{1+\varepsilon}}{1+\varepsilon} < 0$ consumer's utility loss, by $V(p) = U(f(p), m(p))$ her indirect utility function, and from $V(p', Y + CV) = V(p, Y) = Y + CS(p)$,

$$CV = CS(p) - CS(p') = \frac{E_f(p) p^{1+\varepsilon} - (p')^{1+\varepsilon}}{1 + \varepsilon} = -l_\varepsilon\left(\frac{p' - p}{p}\right) E_f(p) > 0. \quad (43)$$

Computing the empirical counterpart of (43) requires information on two sufficient statistics for each agent: the price elasticity $\hat{\varepsilon}$ and ex ante fuel expenditures $\hat{E}_f(p)$, before any price increase.⁶⁰

Importantly, denoting by $l_0(x) = \lim_{\varepsilon \rightarrow 0} l_\varepsilon(x) = -x$, we have $g_\varepsilon(x) = 1 + (1 + \varepsilon)x - (1 + x)^{1+\varepsilon} = (1 + \varepsilon)(l_\varepsilon - l_0)(x)$, which takes positive values when $\varepsilon \in (-1, 0)$ because $g_\varepsilon(0) = 0$ and $g'_\varepsilon(x) = (1 + \varepsilon)[1 - (1 + x)^\varepsilon] > 0$. As a result, $CV \leq CV_0 = \frac{\Delta p}{p} E_f(p)$: price subsidies are more costly than transfers due to the behavioral response, namely the adjustment of demand to lower prices, since that adjustment reduces consumer welfare loss, in absolute. Also, the gap between both instruments shrinks either when $\varepsilon \rightarrow 0$, i.e. when the behavioral effect is small, and when $(p' - p) \ll p$, i.e. when the price change is marginal. However, the equivalence between both policy instruments does no longer hold in other cases.

⁵⁹Since preferences are weakly separable, the parameter B should indeed be heterogeneous so that there be a case for commodity taxation on top of nonlinear taxation of income, cf. [Saez \(2002\)](#).

⁶⁰When price shocks are marginal, i.e. for small $\frac{p' - p}{p}$, a first-order approximation leads to $CV \approx -E_f(p) \frac{p' - p}{p}$ as in [Astier et al. \(2023\)](#). Ignoring higher orders of the Taylor expansion requires discarding any behavioral effect.