CRESTCENTER FOR RESEARCH IN ECONOMICS AND STATISTICS

WORKING PAPER SERIES

COMPENSATION AGAINST FUEL INFLATION: TEMPORARY TAX REBATES OR TRANSFERS?

Odran Bonnet, Etienne Fize, Tristan Loisel and Lionel Wilner

N°5 /March 2024

CREST
Center for Research in Economics and Statistics
UMR 9194

5 Avenue Henry Le Chatelier
TSA 96642
91764 Palaiseau Cedex
FRANCE

Phone: +33 (0)1 70 26 67 00
Email: communication.crest@ensae.fr
https://crest.science/

Compensation against fuel inflation: Temporary tax rebates or transfers?

Odran Bonnet Étienne Fize Tristan Loisel Lionel Wilner*

January 19, 2024

Abstract

This article exploits both the crude oil price surge consecutive to the invasion of Ukraine and 2022 fuel excise tax rebates in France as quasi-natural experiments to infer the price sensitivity of fuel demand. Based on granular individual bank account data at the transaction level, we properly disentangle anticipation effects from price effects, and estimate an average price elasticity of -0.31. It varies little with respect to income and location but substantially decreases, in absolute, with respect to fuel spending and is higher for retirees. We evaluate financial and distributional effects of the actual tax policy as well as its impact on CO_2 emissions based on counterfactual simulations. We empirically demonstrate that resorting to transfers, be they targeted or not, achieves only imperfect compensation against fuel inflation. However, we show that a policy maker subject to a tight budget constraint and seeking to alleviate excessive losses, relative to income, prefers means-tested transfers to rebates.

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

JEL Classification: C18; C51; D12; H23; H31; L71; Q31; Q35; Q41.

---

*Bonnet: Insee. Fize: Institut des Politiques Publiques, Paris School of Economics. Loisel, Wilner: Insee, Crest. Please address correspondence to lionel.wilner@insee.fr. Address: 88 avenue Verdier, 92 120 Montrouge, France. Tel: (+33)187695917. 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, Erwan Gautier, Daniel Herrera-Araujo, Laurence Jacquet, Etienne Lehmann, Thierry Magnac, Volker Meier, Felix Montag, Stefan Pollinger, Aurélien Saussay, and Ao Wang 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 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

Though global warming is an obvious matter of concern for decision-makers, taxing gasoline at the pump has proven difficult to implement in practice, due to a fierce opposition. In the public debate, a trade-off has emerged between undertaking necessary actions against the 'end-of-the-world' and helping the poor to deal with their 'end-of-the-month', as if policy makers could not do both at the same time. To illustrate, France has experienced numerous demonstrations at the end of 2018: the *Gilets jaunes* (Yellow Vests) movement was firmly opposed to any increase of the carbon tax, and such protests somehow echoed the previous *Bonnets rouges* (Red caps) movement in 2013. In both cases, governments were forced to give up their initial projects: raising the carbon tax and introducing some truck-specific carbon tax, respectively. On the one hand, it is surely desirable for the planet to discourage emissions thanks to the price signal. On the other hand, the temptation might be partly resisted because the corresponding tax burden is mostly borne by low-income individuals: excise taxes on fuel are regressive since the budget share of fuel consumption tends to decrease with income (Douenne, 2020).[1] The 2022 rise in crude oil prices consecutive to the invasion of Ukraine and the following energy crisis only exacerbated this problem: the budget constraint may be binding for households that are unable to adjust their fuel consumption in the absence of any alternate transportation, thinking of those living in rural areas, for instance. In France, the government decided to directly subsidize prices at the pump through rebates.The public intervention was then motivated by the objective of alleviating the burden due to the price surge for individuals devoting an excessive budget share to fuel expenditures. The rebates were removed at the end of 2022, though: at that time, they were accused of (i) providing wrong incentives, from an environmental viewpoint; (ii) costing too much, from a financial viewpoint, and (iii) being unfair in that they benefit more to high-income individuals, from a distributional viewpoint. The government thus decided to replace rebates with a means-tested transfer granted to car-owners in the bottom half of income. Compensation mechanisms such as lump sum or means-tested transfers issued from revenue recycling have long been summoned by economists; determining the optimal compensatory scheme is yet a difficult task, in practice. Sallee (2019) explains that a policy designed to enhance efficiency, like a Pigouvian-based tax-and-transfer scheme, inevitably

[1] Some distinction should yet be made between constrained (typically, daily commuting between home and workplace) and unconstrained consumption (e.g., related to leisure activities).

generates losers, and that an accurate prediction of fuel consumption based on observed characteristics (income, geography, sociodemographics), which is a highly empirical issue, is crucial in order to better tailor transfers, and hereby to reduce the number of losers.

From that viewpoint, recent access to almost real-time, individual transaction-level data like the one issued from bank accounts confers a key role to the econometrician in market design. She may first measure key sufficient statistics in this regard, including the short-run price-elasticity of fuel demand. Moreover, she can provide an almost contemporaneous evaluation of the policy, while comparing the latter with alternative public interventions given distributional, environmental, or financial criteria.

In this paper, we rely on a number of recent quasi-natural experiments provided by the invasion of Ukraine: the 2022 rise in crude oil prices and the tax policy designed to temper that inflation. In several European countries including Germany and France, governments decided to directly subsidize fuel at the pump, which is formally equivalent to excise tax rebates. We exploit those exogenous price variations in order to infer the price sensitivity of fuel demand. Our empirical analysis is based on high-frequency data, namely transaction-level data issued from bank accounts, which includes timestamped operations at the individual level from September 2021 to February 2023. Disposing of daily data enables us to finely disentangle anticipation effects from the pure price effect. We show that credible estimates of the short-run price elasticity require (i) an appropriate source of identifying variability like, e.g., the price shocks mentioned before; (ii) high-frequency data that renders possible to first visualize, then neutralize very short-run (namely, daily) anticipation effects; and (iii) a suited econometric approach that correctly separates anticipation effects from the sole price effect. Lacking of each and any of those ingredients results in identification failure: we estimate that the anticipation bias is about -0.4. From a methodological viewpoint, we build upon a literature devoted to anticipations including, e.g., Coglianese et al. (2017); we also quantify the aggregation bias inherent to less granular datasets.[2] Such an issue has been a long-lasting concern for applied econometricians when

[2] Levin et al. (2017) have already pointed out that low-frequency data suffered from such an aggregation bias due to three distinct reasons: (i) a common price coefficient is imposed while price sensitivities might be heterogeneous, say, at some spatial unit (city) level; (ii) less granular datasets do not allow to include appropriate space or time fixed effects viewed as unobserved components which permit to remove any supply-driven variation induced by fluctuations in price over time, for instance; (iii) aggregation might induce some correlation between average prices and the error term, e.g. when correlation between prices and demand shocks on other days, or in other cities, can cause prices and errors in the aggregated panel data to be correlated.

evaluating tax changes. We believe that the methodology developed in the current paper extends to other settings in public finance, more broadly. For instance, in capital taxation, public announcements as regards wealth taxes are likely to induce strong anticipatory behavior and similar bunching patterns as the ones observed here due to the high mobility of assets. It applies in fact to many other contexts involving any change in commodity taxation, including VAT, or in fiscal incentives, like ecological subsidies, bonus/malus schemes related to green taxation, and other transfers that are provided conditional on buying electrical or on home renovations.

Our contribution is fourfold. First, we measure an average short-run price elasticity[3] of -0.31. Second, taking further advantage of our individual data, we investigate whether the price elasticity displays some heterogeneity across income and location, two possibly relevant dimensions of compensatory tax-and-transfer schemes. It turns out to vary little along those dimensions, which does not preclude yet nominal responses to substantially differ. More strikingly, we find sizeable dispersion of the price sensitivity with respect to fuel spending: the elasticity markedly decreases with that spending, in absolute. Also, among those who consume more fuel, more liquidity-constrained households are more inelastic. In another vein, employees are more inelastic than retirees. These results are helpful to improve the targeting of environment-oriented tax policies since a finer knowledge of fuel spending determinants is useful to reduce the prediction problem mentioned by Sallee (2019), ideally lowering the number of losers in a Pigouvian-based tax-and-transfer scheme.

Third, we evaluate financial, distributional, and environmental impacts of 2022 fuel tax rebates based on counterfactual simulations. Our analysis suggests that the average effect per household amounts to saving €66 (about 0.14% of the average income, or 4.8% of the average annual fuel bill). Alleviating the tax burden has relatively less favored low-income individuals, in nominal terms, since high-income individuals benefited most from those excise tax rebates; when expressed as a fraction of income, that gain nevertheless decreases with income, the effective tax rate being diminished by up to 0.47pp at the lower end of that distribution, against by 0.11pp only at the top. The estimated effect of the policy on CO_2 emissions is +0.36%.

Fourth, from a more normative viewpoint, we determine the optimal policy rule that the decision maker should adopt when seeking to compensate for households' loss consecutive

[3] That elasticity may differ from the long-run price elasticity since it takes time for consumers to adjust their response: for instance, by shifting to another transportation.

to fuel inflation, expressed as a share of income. To that end, we simulate alternate compensating policies: unconditional or means-tested rebates, and transfers that may be lump sum, income-based, conditional on location or fuel spending. We show that means-tested rebates achieve a better job, as regards the government's objective function, at a lower cost; the latter may nevertheless be infeasible due to implementation considerations. Next, we empirically demonstrate that means-tested transfers should be preferred to uniform tax rebates when the government's budget constrain is tight: this is partly due to the fact that she only achieves partial compensation with transfers, which numerous 'losers' may not support. Even endowed with reliable and precise information on households' fuel spending, partial compensation arises due to unobserved heterogeneity along that dimension: we quantify the value of information that rather resides in the possibility of targeting a specific group of consumers, which makes it more affordable to the government. Information per se does not permit to substantially reduce the dispersion of gains and losses within that group. However, a countervailing force is at play: under scarce pecuniary resources for compensation, means-tested transfers are more effective at targeting the right individuals from the policy maker's viewpoint because income acts as a satisfying screening device for the relative loss. When more public funding becomes available, though, rebates achieve an exact compensation at the agent's level, which make them relatively more attractive then.

Previous results help understand why political support for revenue-recycling carbon tax-and-transfer schemes is so difficult to obtain (Young-Brun, 2023), though latter schemes are effective at canceling out regressivity or even at achieving progressivity. To the best of our knowledge, it is the first empirical paper that compares two widely used policy instruments, rebates and transfers, given financial constraints and distributional motives, hereby assessing and quantifying their relative compensating power. By construction, rebates are better tailored to completely offset the impact of inflation. However, they are often criticized for being costly and for favoring low-income households. By contrast, transfers could be targeted towards specific populations that suffer the most from inflation and that the policy maker is willing to help overcome that shock. More generally, we believe that the current analysis provides valuable guidance to decision makers looking for a better design of compensating policy against (fuel) inflation.

Literature A vast empirical literature, surveyed for instance in Dahl and Sterner (1991) and Espey (1998), has been devoted to measuring the short-run price elasticity of fuel demand. Yet those studies have so far relied mostly on disaggregated data, hence facing two fundamental problems.

The first issue is endogeneity arising due to simultaneity, which results in an attenuation bias: in a demand equation, the price is concomitantly determined by the supply side as a function of the quantity purchased. The typical solution consists in finding instruments such as the West Texas Intermediate (WTI) crude oil spot price, the price of the brent, the average price in the market, some tax change. More generally, predictors of prices like cost shifters that are as unrelated as possible with demand usage are potential candidates.

The second empirical issue that complicates demand estimation is anticipatory behavior. When consumers anticipate price changes, they may strategically delay or advance their purchases: intertemporal substitution arises possibly through dips (resp. spikes) when future prices are expected to be lower (resp. higher), which invalidates previous approaches. The current quantity then depends not only on current prices, but also on future prices, which violates usual exclusion restrictions at play in both OLS and IV estimation methods: contemporaneous changes in fuel purchases may be larger than expected, and even IV estimates might overstate the price sensitivity of demand. A typical solution to deal with this problem consists in introducing leads and lags as in Coglianese et al. (2017), or, more recently, in Kilian (2022); yet this approach relies on parametric assumptions as regards the dependence of expectations with respect to future and past prices. Compared to Kilian and Zhou (2023) who dispose of monthly state-level data, our empirical analysis is based on high-frequency data, which enables us to finely disentangle anticipations from price effects; importantly, we rely on the 2022 oil crisis and policy responses as instruments, and we are able to exploit the variation in observed characteristics at the individual level to infer the corresponding heterogeneity of the price elasticity; finally, we focus on a European context.

A few recent studies have resorted to high-frequency data in order to estimate the price sensitivity of fuel demand, including Levin et al. (2017) and Knittel and Tanaka (2021) who disentangle extensive from intensive margins, i.e. driving behavior from travelling distance. We improve upon their methodology by relying on an exogenous source of price variations (the 2022 oil crisis consecutive to the invasion of Ukraine and corresponding policy responses in France, two successive excise tax rebates starting on April 1st and September 1st, at least partly effective until the end of 2022). 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 (about -50% in 6 months). They mostly examine cross-price effects with other spending than fuel, though. Another 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 complement to theirs; on top of that, we explicitly adopt an econometric specification that controls for anticipatory behavior.

Last, our paper contributes to both academic and policy debates about the implementation of carbon taxation, among which Sallee (2019), Douenne (2020), Young-Brun (2023). It builds upon these articles by comparing two policy tools, rebates and transfers, in terms of compensating power against fuel inflation, given financial, distributional and environmental constraints.

The rest of the article is organized as follows. Section 2 presents our data and the institutional background. To illustrate how anticipation effects can be disentangled from price effects, a toy model is exposed 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 of the price elasticity. Section 6 is devoted to counterfactual simulations that quantify financial and distributional impacts of tax rebates as well as their effect on CO_2 emissions; it also compares the actual tax policy to transfers as an alternative compensatory scheme. Section 7 concludes.

2 Data and context

We first present our de-identified bank account data. Our database is issued from the *Crédit Mutuel Alliance Fédérale*, a French group of banks with about 30 million customers, either firms or households. The construction of key variables follows a recent strand of literature exploiting such data including, e.g., Baker (2018), Ganong and Noel (2019) and Andersen et al. (2023). We dispose of transaction-level data on credit and debit card payments,[4] paper checks, cash withdrawals, cash deposits, bank transfers, and direct debits. We observe the amount of each transaction, in euros; we nevertheless base our analysis on a daily aggregation. On top of that, balance sheets are available each month. The statistical unit of observation is a household; the data contains various socio-demographics

[4] In France, the use of credit cards is scarce: it accounts for less than 10% of bank cards.

on households' members like age, sex, *département*,[5] family status, occupation, and the type of location (in three categories: urban, rural, and semi-urban areas).

We define total spending as the sum of outgoing transactions issued by 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 is a sample of about 300,000 households who primarily bank at *Crédit Mutuel-Alliance Fédérale*, this sample being stratified by *départements* of metropolitan France and by 5-year age dummies. To alleviate concerns about representativeness, we proceed to calibration weighting using the method proposed by Deville and Särndal (1992) (see Appendix F for details), and weight our estimating equations using 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 meet previous criteria all over the period, which leaves us with about 180,000 active customers primarily banking at *Crédit Mutuel-Alliance Fédérale* due to attrition since June 2020, namely the time when the sample was drawn from the universe of bank 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 as Andersen et al. (2023) and Gelman et al. (2023) do.[6] Figure 1 displays the distribution of the amount of a transaction, in euros, which seemingly mixes a continuous distribution, the mode of which lies nearly 55 euros,

[5] An administrative division of France 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 the two Corsican *départements*.

[6] We restrict our attention to card payments, excluding cash for instance, which sounds like a mild empirical choice: the distribution of monthly fuel spending with respect to income looks close to the one obtained from the *Budget des familles* representative survey, see Figure F4 in Appendix.

and a discrete distribution over round numbers (typically 30, 20, 50 or 40 euros, as well as other multiples of 5 euros).[7] Figure 2 further shows that the median interpurchase duration, i.e., the time interval between two visits at the pump, is 7 days; the median elapsed time between two visits, measured at the household level, is 16 days. Last, we obtain fuel quantity, in liters, as the ratio of that adjusted fuel spending over a fuel price index: we now explain how the latter is computed.[8]. In our sample, Figure 3 confirms that fuel expenditures increase with income but that the budget share devoted to fuel is decreasing along that dimension; rural households devote a higher share of their budget to fuel. That empirical evidence is consistent with the idea that the burden of fuel consumption is mostly borne by poor and rural households who may depend more on the car, as opposed to other transportation, and who are likely more constrained.

Prices Timestamped and geolocated fuel prices are disclosed at the gas station level by a French governmental website.[9] Such data has already been used by researchers: see, e.g., Montag et al. (2021) or Gautier et al. (2023). It contains information on each and any price change for different kinds of fuel (diesel and gasoline: super unleaded petrol (SP95), super unleaded petrol (SP95-E10), super unleaded petrol (SP98), etc.). In the subsequent analysis, we focus on two types of fuel: diesel and standard gasoline, which we confound with SP95, given that the latter exhibits similar variations over time as both SP95-E10 and SP98 (over the period considered, the correlation is higher than 0.99, which mitigates any concern about substitution between those products). On top of that, the data provides with an identifier and the location of each retailer.

As detailed in Appendix A of Gautier et al. (2023), the first step consists in mapping raw data to a daily panel dataset at the (retailer, type of gasoline) level. Since different price changes may occur within the same day (the typical frequency of price changes being a few

[7]The latter most likely stems from the possibility of prepaying gas in some stations, and we verify below (last raw of Table 3) that this specificity does not dramatically change the picture of our estimated price sensitivity by isolating that dimension, namely purchases corresponding to non-automated fuel dispensers (MCC: 5541).

[8]Cross-border purchases, which result from trade-offs that involve, in particular, the distance to the frontier and the ratio of foreign over 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 some fuel abroad (holidays being kept aside). Foreign transactions occurring during holidays are also removed from the subsequent analysis. For more details on cross-border fuel purchases, see our companion paper Adam et al. (2023).

[9]`https://www.prix-carburants.gouv.fr/rubrique/opendata/`.

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, say, diesel, but inactive for gasoline. We then trim outliers by deleting top and bottom 1% of price observations for each (*département*, type of fuel, day). Admittedly, transaction prices are measured with error since we ignore the exact location of purchase, hence we approximate them with their daily average in the *département*. In that vein, price search is a potential source of measurement error, which we can hardly address with current data.[10] 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 tax rebates 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.[11]

Information on the type of fuel actually purchased, diesel or gasoline, is unavailable in the data, which is yet unimportant provided that both prices similarly covary. Empirically, those prices are very correlated: their Pearson coefficient is 0.8 over the whole period of observation; diesel and gasoline prices sometimes experience different short-run variations due to specific conditions affecting the oil market. We therefore build a fuel price index that weighs diesel and gasoline prices differently within a *département* according to strata based on observed households' characteristics.[12] According to the *Enquête Mobilité* survey,[13] those 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 based on that survey. Again, allowing

[10] Search behavior would lead the econometrician to overestimate price sensitivity, in absolute: transaction prices would be lower, and quantity higher (remember that the latter are obtained from the ratio of expenditures over prices). When we resort to instrumental variables (see below), we obtain a higher elasticity, in absolute, which suggests that we partly tackle that problem.

[11] If 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.

[12] income (in four groups), type of location (urban, rural, semiurban), age group (less than 30, 30-60, more than 60), and 2019 fuel spending (in four groups).

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

for this fuel price index to depend on households' characteristics further mitigates concerns about heterogeneous information among consumers as regards prices, which Montag et al. (2021) have proven to result in heterogeneous pass-through.

Context: invasion of Ukraine, energy crisis, and policy responses (temporary excise tax rebates) Fuel prices have experienced substantial variations from 2019 to 2023, especially in 2022 due to an oil price surge consecutive to the invasion of Ukraine that started on February 24th. In France, the government decided to directly subsidize prices at the pump; since per unit excise taxes represent about 60% of fuel prices, the public intervention in fact consists in offering a tax rebate. On March 12th, Prime Minister Castex made an official announcement to explain that the before-tax gasoline price would be lowered by €0.15 per liter from April 1st onwards (about €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, consecutive to the energy crisis. As announced by Prime Minister Borne at the end of July, a total fuel tax rebate of €0.3 per liter has then been effective on the after-tax price from September 1st onwards, namely an extra €0.12 subsidy for each liter purchased. That second price reduction has prevailed until mid-November 2022 when that rebate has been diminished to €0.1 per liter only, before its complete removal on January 1st, 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. All those reductions disappeared at the end of 2022.

In what follows, we rely on policy-driven fuel price reductions as the primary source of identifying variability in order to causally infer price effects. We view the two public interventions of April 1st and September 1st as quasi-natural experiments, which provide us with exogenous price changes. Those policies were publicly disclosed, hence salient to consumers. Figure 4 suggests that the evolution of fuel demand at the time of announcement is consistent with anticipatory behavior by forward-looking consumers: people strategically refrain from buying and wait for lower future prices once they are aware of lower prices in the future. Figure H3 in Online Appendix further confirms the salience of that intervention: it indicates that consumers mostly adapted to the policy by purchasing less before price reductions (hence adjusting at the intensive margin), rather than by visiting gas stations less often (the number of transactions being a proxy for the extensive margin).

This anticipatory behavior yet renders the identification more subtle, which requires to properly disentangle short-term intertemporal substitution from the true price effect.

3 Theoretical framework

To illustrate how the price effect can be separated from anticipation effects in a dynamic setting, we resort to 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 as regards 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.[14] In our application below, the period is typically a day or a week: it is thus reasonable to assume no depreciation of fuel (when stockpiled), hence a discount factor equal to one along with a zero interest rate.[15] Denoting the instantaneous utility derived from consumption by $u(.)$, fuel purchases by q, fuel prices by p, permanent income by Y, and storage costs by $C(.)$, the program of the agent writes:

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

$$i_t \leq i_{t-1} + q_t - c_t \tag{2}$$

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

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

[14] The model is a simplified version of an inventory model with quasi-linear preferences used, e.g., by Hendel and Nevo (2006).

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

In the optimum, the intertemporal budget constraint also binds, the Euler equation holds, and fuel inventory is ruled by:

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

This equation makes it clear that stockpiling behavior is governed by the expected change in prices. Parametrizing $C(i_t) = \theta i_t^2$ with $\theta > 0$ leads to:

$$i_t = \lambda \frac{p_{t+1} - p_t}{2\theta}, \tag{5}$$

and considering some quadratic utility function of the form $u(c_t) = c_t - \alpha c_t^2$, with $\alpha > 0$, yields a linear demand:[16]

$$c_t = \frac{1 - \lambda p_t}{2\alpha}. \tag{6}$$

The model predicts that observed purchases are given by:

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

hence a specification of the form:

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

where $\beta = -\frac{\lambda}{2\alpha}$, and $\gamma_t = \lambda \frac{(p_{t+1} - p_t) - (p_t - p_{t-1})}{2\theta}$ account for anticipation effects that are non-zero as soon as prices fluctuate, and alter current purchases. They only affect the timing of purchase only and not the quantity consumed c_t; they vanish (resp. are exacerbated) when θ tends to $+\infty$ (resp. 0), i.e. when storage is impossible (resp. not costly) because the tank is full, for instance. By contrast, prices determine the total amount of fuel purchased. This specification provides a microfoundation for the toy econometric model exposed in section 3.2 as well as for our estimating equation (11).[17]

That observed purchases depend on lags and leads of prices on top of current prices, as in the RHS of (7), is reminiscent of Coglianese et al. (2017). Though stylized, this concep-

[16]That function verifies Marshall's second law of demand: its price-elasticity decreases with consumption, in absolute:$-\frac{\partial \log(c_t)}{\partial \log(p_t)} = \frac{\lambda p_t}{1-\lambda p_t} = \frac{1}{2\alpha c_t} - 1$. See also section 5.3 below.

[17]Here 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.

tual framework thus rationalizes any reduced-form approach that involves a regression of purchases on current, past, and future prices when aiming to recover the price-elasticity.[18] It is also immediate to check that 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 be retrieved from the sum of coefficients corresponding respectively to current, past, and future prices.

3.2 Implications for identification

Previous insights shed light as regards identification. We now present a toy econometric approach derived from the model above, which empirically permits to disentangle anticipatory behavior from price-sensitivity. The first specification applies to a single, anticipated price reduction like the one experienced around September 1st, 2022, while the second specification corresponds to an anticipated price surge followed by some expected price reduction of the same magnitude as was the case in March-April 2022.

3.2.1 A single price reduction

A simplified version of the second excise tax rebate is described by Figure 5a. We consider a 4-period model whereby periods are indexed by k and last T_k days with $\sum_{k=1}^{4} T_k = T$. For the sake of simplicity, we ignore any variation either in prices or in purchases within period k: $p_t = p_k, q_t = q_k \quad \forall t \in k$. In the first two periods, prices are thus assumed to be constant and equal to their regular level p. Consecutive to the policy intervention, prices fall to $p - \Delta p$ where $\Delta p > 0$ is the discount. The researcher observes prices and purchases; following previous considerations, she wants to estimate the following linear model:

$$q_t = q_0 + \beta p_t + \gamma_2 \mathbb{1}_{t \in 2} + \gamma_3 \mathbb{1}_{t \in 3} + u_t, \tag{9}$$

based on moments conditions: $\mathbb{E}(u) = 0$, $\mathbb{E}(pu) = 0$ on top of $u_{t\in 2} = u_{t \in 3} = 0$. To make an explicit link with previous subsection, the inventory model would impose supplementary constraints: $\gamma_2 = -\lambda\frac{\Delta p}{2\theta} < 0$ and $\gamma_3 = \lambda\frac{\Delta p}{2\theta} > 0$. Consistently with the econometric specification at stake,[19] $\gamma_2 + \gamma_3 = 0$ (cf. 'constrained estimator' below). Among the four parameters $(q_0, \beta, \gamma_2, \gamma_3)$, the researcher is primarily interested in the marginal effect of

[18] Remember from equation (5) that inventory behavior is the driving force of such an empirical approach.

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

prices. Consistently with observation, she expects $q_2 < q_1 < q_4 < q_3$: (i) $q_1 < q_4$ due to the price effect since $p - \Delta p = p_4 < p_1 = p$; (ii) $q_2 < q_1$ due to strategic delay of purchases in period 2 before the price reduction; (iii) $q_3 > q_4$ since people who refrained from buying when prices are high purchase when prices are lower. Anticipation effects refer to both (ii) and (iii). Those confounders are such that $q_2 + q_3 = q_1 + q_4$: though they affect the timing of purchase, they have no impact on total consumption; they exactly compensate over the *anticipation window* made up of periods 2 and 3.

A 'naive estimator' that would omit to nonparametrically control for anticipation effects in periods 2 and 3, *de facto* imposing that $\gamma_2 = \gamma_3 = 0$. It would yield $\hat{\beta}^n = \frac{T}{T_1+T_2}\frac{\bar{q}-\bar{q}^{34}}{\Delta p}$,[20] which boils down to $\frac{(q_1+q_2)/2-(q_3+q_4)/2}{\Delta p}$ when $T_k = 1, \forall k$ (see Appendix C.1.3 for details). By definition, that estimator does not permit to separate anticipation effects[21] from the pure price effect, which results in spuriously relying on $(q_2,\ q_3)$ to infer the marginal effect of prices. Controlling now for what happens during the *anticipation window*, which is centered around the moment when prices fall, yields the 'unconstrained estimator': $\hat{\beta}^u = \frac{q_1-q_4}{p_1-p_4} = \frac{q_1-q_4}{\Delta p}$ (cf. Appendix C.1.1), which recovers the desired price effect net of any strategic effect. 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 discards any contribution from periods 2 and 3, hence a loss of information. Our preferred estimation procedure consists rather in imposing the constraint that $T_2\gamma_2 + T_3\gamma_3 = 0$, i.e. in positing the zero-sum of anticipation effects over the anticipation window, consistently with theoretical arguments above.[22] Put differently, our identifying assumption states that individuals refrain from buying fuel in period 2 for pure intertemporal substitution motives, because they wait for lower prices, but that they eventually buy an excess quantity in period 3 that exactly corresponds to default quantity from period 2. The 'constrained estimator' $\hat{\beta}^c$ coincides here with the 'unconstrained estimator' (cf. Appendix C.2.2), but 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 (i.e. the case when $T_1 = T_4$ and $T_2 = T_3$). By contrast, the unconstrained estimator $\hat{\beta}^u$ is always independent from (q_2, q_3). Empirically, small price variations during that anticipation window may also be exploited for inference; by definition, not imposing that constraint would render such an inference impossible.

[20]The notation $\bar{q}^{34}$ refers to $\frac{\sum_{k=3}^{4} T_k q_k}{\sum_{k=3}^{4} T_k}$ while $\bar{q}$ stands for $\frac{\sum_{k=1}^{4} T_k q_k}{\sum_{k=1}^{4} T_k}$.

[21]As explained above, those effects are negative when $k = 2$ and positive when $k = 3$.

[22]It is possible to test that constraint: see section 4.2.

Though simplified, this conceptual framework closely resembles the situation that prevails as regards the second excise tax rebate implemented on September 1st. Equipped with the above toy model, observed prices and purchases, and adjusting 2022 data for seasonality based on 2021 observations as in section 4.1 below, we obtain an elasticity of -0.31 in the absence of any constraint on γ_2 and γ_3 (see Appendix C.3). The 'constrained estimator' yields a -0.31 elasticity. The 'naive estimator' amounts to -0.68. Those figures turn out to be close to econometric point estimates (see below), and already give a flavor of the magnitude of the anticipation bias, namely $(-0.68)-(-0.31) \approx -0.37$. Remaining differences with the actual econometric estimation lies in that the model (i) does not account for covariations of price and quantity within each period, and (ii) slightly departs from observation since prices do not behave exactly as in the current theoretical framework.

3.2.2 Price surge + compensating tax rebate

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 rebate 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 resort to a 5-period model as described by Figure 5b: $\forall k \neq 3,\ p_k = p$, 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, consecutive to 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 consecutive to the expected rebate 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:[23]

$$q_t = q_0 + \beta p_t + \gamma_2 \mathbb{1}_{t\in 2} + \gamma_3 \mathbb{1}_{t\in 3} + \gamma_4 \mathbb{1}_{t\in 4} + u_t. \tag{10}$$

[23] Again, 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 $c_3 - \lambda\frac{\Delta p}{\theta} = q_3 < q_1 = q_5 = c_1 < q_2 = q_4 = c_1 + \lambda\frac{\Delta p}{2\theta}$. Consistently 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 $\gamma_2 + \gamma_3 + \gamma_4 = 0$.

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 C.2.1). 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 C.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 C.2.2). 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 C.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.

3.2.3 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 C2, $q_1 \approx 2.373$ 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.[24]

[24] $9.04 \approx T_3(q_1 - q_3) > T_2(q_2 - q_1) + T_4(q_4 - q_5) \approx 7.21$.

4 Empirical analysis

In that section, we explain how observed variations in prices over time constitute a quasi-experimental setting that can be exploited to infer the price sensitivity of fuel demand. In particular, we rely on substantial price changes, including various price increases,[25] combined with two downwards, policy-driven price changes: the €0.18 per liter excise tax rebate from April 1st, 2022, and the extra €0.12 per liter reduction on after-tax prices from September 1st, 2022. These sharp price changes provide us with an identifying variability that enables us to recover the shape of the demand function based on our high-frequency dataset. Our empirical approach consists then in finely disentangling anticipation effects from the aversion to prices. In the very short-run (say, a few days), even 'unexpected' price variations are anticipated: for instance, following the invasion of Ukraine, a peak of purchases can be observed since forward-looking consumers anticipate higher future fuel prices.

4.1 Identification strategy

Our analysis relies on the extra €0.12 per liter excise tax rebate from September 1st, 2022. We view that second intervention as a quasi-natural experiment for the identification of anticipation effects. Indeed, this policy-driven shock acts as an exogenous shock on prices that enables us to retrieve consumer demand. To that end, we adjust prices and purchases for seasonal variations, relying on year 2021 as the baseline. Figure 6 makes it clear that our empirical setting looks close to our theoretical framework: prices are roughly stable both before and after the implementation of the additional rebate. Identification rests on the public intervention being salient to and expected by consumers. The reason why we cannot resort to a similar identification strategy as regards the first public intervention in April 2022 is that we lack of a credible baseline year for seasonal adjustment, due to strict (resp. partial) lockdown occurring in April 2020 (resp. 2021).[26]

Former empirical studies on fuel demand have proposed parametric solutions to deal with anticipation effects: based on monthly-level data, Coglianese et al. (2017) resort to

[25] about +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 1st, 2023 onwards.

[26] We could have relied on April 2019; unfortunately, the information about the MCC has been available in the data since July 2019 only.

price lags and leads. We adopt here a nonparametric approach, which makes full use of our high-frequency dataset. Our inference of the price effect is thus based on the covariation of prices and purchases: we take advantage of the sharp discontinuity around September 1st, net of anticipation effects. Since intertemporal substitution effects act as nuisance factors, or confounders, in the estimation of the price sensitivity of demand, we identify daily anticipation effects during a window that is centered around the rebate, the anticipation window. Consistently with theoretical arguments, we impose that those effects sum up to zero: our identifying assumption is that very short-term (namely daily) variations in fuel purchases around the rebate correspond to pure intertemporal substitution, with no impact on the total quantity purchased within that window.[27]

Figure 7 suggests that the evolution of fuel purchases would have been similar in 2021 and 2022 in the absence of tax rebates; this hypothesis cannot be rejected on the basis of pre-rebate data, at least.[28]

To mitigate any concern about endogeneity, we adopt an IV strategy based on a post-September 1st dummy, which indicates whether the rebate is effective or not, 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, this potential problem arising due to actual fuel purchases being unobserved. Dividing spending by prices could indeed lead to a downward bias of the price coefficient, remember the 'division bias' (Borjas, 1980).

4.2 Econometric specification

To implement our identification strategy, we first aggregate our data into 10,777 cells of individuals with similar *département*, income (in four groups), age (less than 30, 30-60, more than 60), type of location (rural, urban, or semiurban area), and 2019 fuel spending category (in four groups). Our estimations are then weighted according to the sampling

[27]Using a wording borrowed from the bunching literature, there is no excess mass after event once pre-event default mass has been netted out.

[28]Our approach does not require any common trend assumption as in a difference-in-differences or in an event study. Moreover, when checking for the absence of any differential pre-trend, short-run anticipations should be left aside: as already explained, it is largely expected that a policy-induced dip be observed within a one- or two-week anticipation window, followed by some spike once the rebate is effective.

importance of those cells. Proceeding in a such a way substantially alleviates the computational burden inherent to dealing with high-frequency individual data, it does not reduce our identifying power, since diesel and gasoline prices are measured at the *département* $\times$ day level. It is yet worth noting that our fuel price index does vary within a *département* due to the cell-specific fuel mix. Besides, our estimations include cell-specific fixed effects so as to take the heterogeneity of fuel spending across cells into account.

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 the beginning of October[29], both in 2022 and in the baseline year 2021. Let $\underline{t}$ =07-14-2022 (resp. $\bar{t}$ =10-03-2022) designate the beginning (resp. 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 in order to recover the price sensitivity, following insights from section 3:

$$q_{cty} = \beta p_{cty} + \sum_{h=t_2-\Delta}^{t_2+\Delta} \gamma_{hy} \mathbb{1}_{h=t} \mathbb{1}_{y=2022} + \alpha_{cy} + \mu_t + \eta_{cty}, \tag{11}$$

where p_{cty} are prices, α_{cy} is a cell-year fixed effect, and μ_t is a day-of-the-year fixed-effect that accounts for seasonal adjustment, 2021 calendar days being also adjusted so that they coincide with their siblings in 2022. t_2 corresponds to the beginning of the second excise tax rebate, namely September 1st. 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**:
$$\sum_{h=t_2-\Delta}^{t_2+\Delta} \gamma_{h,2022} = 0 \tag{12}$$
- **Unconstrained estimator**: No restriction on γ_{hy}.

Under both constrained and unconstrained estimations, the identification of the price effect β stems from the ratio of the difference in fuel purchases that are adjacent to the

[29]We exclude October 2022 from our sample due to strikes in refineries, leading to shortages in various places.

anticipation window over the difference in corresponding prices (cf. theoretical $\frac{q_1-q_4}{\Delta p}$). Standard errors are computed by two-way clustering at cell and year-day levels.

The choice of the anticipation bandwidth Δ is primarily guided by economical considerations. $\Delta = 14$, namely two weeks, sounds like a reasonable value since the majority of interpurchase durations : 75% of interpurchase duration occur within 14 days. Though consumers may be able to manipulate the timing of their visit, especially when they foresee price changes, they are constrained by their tank capacity.[30] Remembering 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 $\sum_{h=t_2-\Delta}^{t_2+\Delta} \gamma_h = 0$. We hereby verify that the structure imposed by both our theoretical model and our econometric specification on those coefficients is not rejected in the data. Implementing this test on a finite sample yet requires to neutralize the estimation error on those anticipation coefficients: failing to do so would surely lead to rejection (see Appendix D for details).

Moreover, we test for anticipations around the tax rebate. Though we reject the absence of anticipations the two weeks surrounding the tax cuts, we cannot reject the absence of anticipations the weeks further away from the price shock. This statistical argument suggests that most anticipation effects occur the days following and preceding the shock, and supports the idea that our bandwidth is large enough.

On top of that, Figure 8 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 downwards bias results in a -0.76 point estimate. When Δ increases, our estimation method better controls for anticipation effects, and sounds fruitful in disentangling strategic delaying of purchase from contemporaneous response to price change. When $\Delta = 14$, our favorite estimate for the price elasticity becomes -0.31. Reassuringly, estimates obtained with higher values of Δ remain rather stable and not significantly different from -0.31.

To assess the validity of our identification strategy, we perform a falsification or placebo test: we consider a fake rebate on September 1st, 2021, using 2019 as the baseline year. When comparing fuel purchases in 2019 and 2021 (two years without any tax rebate), we

[30] Appendix Figure H3 suggests that households mostly adapted their behavior at the intensive margin.

conclude to the absence of any substantial spike around September 1st (see Figure G2 in Online Appendix) and Table G1 confirms that there is not enough identifying variability in that case.

4.3 External validity

To alleviate any concern about identification being local, we complement previous approach with a similar econometric specification based on the period from September 2021 to February 2023.[31] As regards the period from February 24th to the mid-April 2022, we refer the toy model exposed in section 3.2.2. Doing so allows us to rely on other sources of identifying variability, including the beginning of the invasion of Ukraine on 02-24-2022 (denoted by t_0), the first public intervention on 04-01-2022 (denoted by t_1), the reduction of the temporary tax rebate in mid-November 2022 (denoted by t_3), and the removal of that temporary tax rebate at the beginning of 2023 (denoted by t_4) -except that we do no longer dispose of any relevant baseline year. We therefore adjust fuel spending for seasonal variations thanks 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. Based on that external database, we divide observed fuel spending by the 2019 ratio of daily fuel spending over average fuel spending. We specify:

$$q_{ct} = \beta p_{ct} + \sum_{h=t_0}^{t_1+\Delta} \gamma_h^1 \mathbb{1}_{h=t} + \sum_{k=2}^{4} \sum_{h=t_k-\Delta}^{t_k+\Delta} \gamma_h^k \mathbb{1}_{h=t} + \alpha_c + \mu_t + \eta_{ct}, \tag{13}$$

with $\mu_t \equiv X_t'\beta + \delta t$, where δ captures any linear trend in fuel purchases and X_t account for temporal controls including day-of-the-week fixed effects[32] and holidays.[33] To control for anticipation effects, equation (13) is again estimated under the constraints

$$\sum_{h=t_0}^{t_1+\Delta} \gamma_h^1 = 0, \qquad \sum_{h=t_k-\Delta}^{t_k+\Delta} \gamma_h^k = 0 \quad \forall k = 2, 3, 4. \tag{14}$$

[31] At the exclusion of periods with shortages in October 2022, as already mentioned, and from January 7th to January 27th, 2023.

[32] Daily-level data reveal that fuel purchases exhibit strong within-week variations: tanks are much more often refilled on Fridays and Saturdays. By definition, such a seasonality cannot be observed based on low-frequency data.

[33] Interacted with the day-of-the-week.

5 Results

5.1 Main estimates

Our estimate of the price elasticity is taken at the means from the price coefficient β in equation (11): $\frac{\partial \log(q)}{\partial \log(p)} = \frac{\partial q/\partial p}{q/p} \equiv \beta \frac{p}{q}$. We compute $\hat{\varepsilon} = \hat{\beta}\frac{\overline{p}}{\overline{q}}$, denoting the average of X by $\overline{X}$. Table 1 converts the estimated coefficient $\hat{\beta}$ based on the second rebate from September 1st 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 5%.

Estimations based on the sole period following the invasion of Ukraine, namely March and April 2022, are also in line with these findings (Table G5). The constrained OLS estimate, -0.18 (0.07), is not significantly different at 5% from previous one, -0.19 (0.07). However, the period considered here is short: it mechanically mixes different anticipation effects, viewed as confounders, with the price effect. Hence the estimation sounds more fragile than previous one, based on a single price change, and thus more immune to such confounders.

We next confront our estimates to those issued from equation (13), based on the period ranging from September 2021 to February 2023. The main lesson from Table 2 is that we find an average elasticity comprised between -0.42 and -0.26. This exercise is reassuring from an empirical viewpoint since it somehow assesses the external validity of previous approach. Using sharp price changes as instruments yields very close results comprised between -0.37 and -0.22 (Table G6, Raw 2), which suggests that those shocks provide the main effective source of identifying variation over the period.

Previous estimates fall in 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. Instrumental variables implemented on micro data tend to yield a higher sensitivity, while macro-based time series approaches often point out to a smaller elasticity. In the U.S., on the 1989-2008 period, Coglianese et al. (2017) obtain a -0.37 point estimate, as Knittel and Tanaka (2021) do in Japan. Still in the U.S., and according to Levin et al. (2017), that price elasticity would be comprised between -0.35 and -0.27, while Gelman et al. (2023)'s preferred estimate is -0.2. Based on monthly data at the state level, Kilian (2022) estimates that the price elasticity was -0.31 until 2014, but has amounted to -0.2 since then. To directly compare our results with those of Davis and Kilian (2011) for whom a \$0.1 per gallon tax (namely, a 3.12% price increase) would induce gasoline demand to fall by 1.43%, we estimate that a similar price increase (€0.058 per liter) would depress demand by 0.97%.

The naive approach, which would omit to take anticipations into account by imposing $\gamma_{h,2022} = 0, \;\; \forall h = t_2 - \Delta, \ldots, t_2 + \Delta$, is displayed in Columns I and IV of Table 1: the OLS estimate of the price elasticity, -0.44 (0.07), then suffers from a downward bias, -0.25, and the same prevails with IV, that anticipation bias now being -0.45. When restricted to the sole months of March and April, the naive estimate is -0.73 (0.16), see Column I of Table G5: the magnitude of the anticipation bias tends to be higher, consistently with insights from section 3.2 and Appendix C.2. On the whole, these results concur to an anticipation bias of about -0.4.

When excluding the anticipation window from our estimation sample, indirectly relaxing our identifying assumption that anticipation effects compensate over that time period,[34] the unconstrained OLS estimator amounts to -0.24 (0.06), and the unconstrained IV estimator is -0.29 (0.07), see Columns III and VI of Table 1. Both differences with corresponding constrained estimators are not significant at 5%; remember that constrained and unconstrained estimators should be close. More strikingly, the unconstrained OLS estimator based on the sole months of March and April amounts to -0.22 (0.29): it it is very imprecise as expected from section 3.2. Its theoretical uninformativeness stems from the fact that prices experienced a surge and nearly came back to their initial level after the anticipation window (remember Figure H4). In theory, excluding that anticipation window leads to an infeasible estimator; in practice, the standard error dramatically increases.

Interestingly, the point estimates of naive, constrained and unconstrained estimators (resp. -0.76, -0.31, and -0.29) turn out to be close to the ones derived in the toy econometric specification (-0.68, -0.31, and -0.31, see Appendix C.3), given observed prices and

[34]This is equivalent to nonparametrically control for daily dummies during the anticipation window. Everything happens (almost) as if we discard information provided by that time interval when inferring the price effect.

purchases during corresponding periods acting here as sufficient statistics for the inference of the price elasticity.

The importance of anticipation effects can be assessed by looking at Figure H5 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-run intertemporal substitution effects substantially shape the pattern of demand, on top of the price effect.

To evaluate whether our model is able to accurately predict fuel purchases, Figure H6 provides a comparison of predicted with actual demand. The fit of the model estimated over the period from September 2021 to February 2023 looks quite satisfying in this regard.[35]

5.2 Robustness checks

In that section, we conduct various sensitivity analyses in order to check 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 tackle the issue of possibly heterogeneous anticipations.

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.

As regards our main estimation based on September 1st rebate, 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:

$$q_{cy}^{post} - q_{cy}^{pre} = \beta(p_{cy}^{post} - p_{cy}^{pre}) + \eta_{cy} \tag{15}$$

where q_{cy}^{post} (resp. q_{cy}^{pre}) corresponds to average purchases in cell c on year y after the anticipation window ending on September 14th (resp. 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 G2).

[35]Putting aside what happens on January 2023 when there were substantial threats of fuel rationing: France experienced many refineries blockades then, which were related to the social movement caused by the 2023 retirement reform.

As regards our estimations based on the period from September 2021 to February 2023, we difference out equation (13) 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: we estimate 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.[36] 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:

$$\log(\lambda_{cty}) = \varepsilon \log(p_{cty}) + \sum_{h=t_2-\Delta}^{t_2+\Delta} \gamma_{hy} \mathbb{1}_{h=t} \mathbb{1}_{y=2022} + \alpha_{cy} + \mu_t. \tag{16}$$

The estimation is again subject to the constraint (12), and proceeds from maximum likelihood. We may also consider the period from September 2021 to February 2023: in that case, we posit $q_{ct} \sim \mathcal{P}(\lambda_{ct})$ with

$$\log(\lambda_{ct}) = \varepsilon \log(p_{ct}) + \sum_{h=t_0}^{t_1+\Delta} \gamma_h^1 \mathbb{1}_{h=t} + \sum_{k=2}^{4} \sum_{h=t_k-\Delta}^{t_k+\Delta} \gamma_h^k \mathbb{1}_{h=t} + \alpha_c + \mu_t. \tag{17}$$

The estimation is now subject to the set of constraints (14), 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 G2, which compare well to the ones in the same columns of Table 1. Raw 3 of Table G6 provides the results based on the period from September 2021 to February 2023, which do not differ much from the baseline (Raw 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 sole price of the diesel. Our results are not much affected (Columns II and V of Table G3

[36] In that 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.

for the local estimation as well as Raw 6 of Table G6 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 tax change in diesel prices only alleviates the concern about imperfect information (and possibly heterogeneity of information among consumers).[37]

Fourth, the anticipation window may well be heterogeneous among consumers: for instance, occasional drivers might have longer anticipation windows, as suggested by Figure G3 which depicts the distribution of the interpurchase duration (computed outside anticipation windows) for the four different groups of fuel spending. A possible solution to address that issue is to allow for the anticipation bandwidth to vary with fuel spending (Table G7). 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, in absolute, which is in line with previous findings. On top of that, their estimated price elasticity is not significantly different 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.

5.3 Heterogeneity of the price elasticity

We next wonder whether the price elasticity of fuel demand is homogeneous amongst groups of consumers with similar observed characteristics, or not. We first investigate whether the average price elasticity varies with income Figure 9a[38], location (Figure 9b) or family status (Figure 9c) which turns out not to be the case.[39] Remembering nevertheless that the price

[37] In 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 either, see Table G3 to that end.

[38] Figure H7 is similar but has a higher granularity. Investigating heterogeneity along total card spending, rather than income, yields quite the same picture (Figure H8), which mitigates concerns about the sensitivity of this empirical finding with respect to the measurement of income.

[39] If any, Parisian drivers tend to be more price sensitive (Table G4), due to easier substitution with other transportation. Yet the difference is not statistically significant at 5%. Though rural households use more their car, which could make them more inelastic, they are also older, which tends to make them more elastic according to Figure 9d: on the whole, their average price elasticity does not differ much from the one of urban or semiurban households.

elasticity approximately measures a relative reaction, and that fuel consumption is higher, on average, for wealthier individuals as well as for those living in rural areas. The latter respond more to fuel price changes in nominal terms. Also, employees are more inelastic than retirees (Figure H9), as the rationale would suggest since the former are more likely to commute for professional reasons. By contrast, a dimension along which that average price elasticity exhibits sizeable dispersion is fuel spending as measured in 2019 (Figure 9e): as the intuition suggests, 'dependent households', who rely much 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), while those in the top 25% have an average elasticity of -0.26 (0.08). Figure 9f shows further that 'dependent' and liquidity-constrained individuals are most likely to undergo any rise in fuel prices. Such empirical findings have important policy implications: when designing transfer schemes to compensate 'losers', the policy maker seeks to target such households, remember Sallee (2019). It requires however to dispose of much 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.[40] We resort to the method of causal forests pioneered by Athey et al. (2019): Figure 10 displays sorted group average treatment effects issued from a segmentation of our sample into five groups. We then test for homogeneity of the price elasticity, which we reject: the 20% most price-sensitive households have an elasticity of about -0.9 while the 20% most inelastic have a null price elasticity, and the difference is statistically significant at 5%. This empirical evidence is consistent with fuel being a necessity good for almost every car driver.

Our estimated price elasticity differs quite substantially across households, and it is possible to determine who are the most price-sensitive households in terms of both socioeconomic characteristics and geographic location (Table 3). It turns out that most elastic households have a lower fuel spending, are more often retirees or poorer (in terms of either income or liquidity).

Those results are close to the ones obtained by Kilian (2022): in the U.S., states with lower income, higher unemployment rates, and lower urban shares respond more to price

[40]We resort to a log-log specification here to directly interpret the point estimate as an elasticity, and to avoid the issue of zeroes arising when considering cell-specific differences in purchases over several weeks based on equation (15).

variations. Yet an important difference is that our results are issued 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, namely externality-correcting tax-and-transfer schemes, which arise due to imperfect information and imperfect tagging of individual consumption. Such a market failure limits the planner's control over the final distribution of outcomes; a more accurate prediction helps mitigate that empirical problem (Sallee, 2019), and our findings help enhance the quality of that prediction.

5.4 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 dispose of three ingredients: (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 (remember 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 4, and as confirmed by Figure H2 in Online Appendix, illustrations of these unobserved variations include dips and spikes consecutive to 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,[41] (see Column V of Table G2). When we aggregate our data at the weekly level (Column VI of Table G2), 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.

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

6 Policy implications

6.1 The impact of fuel tax rebates

We now assess financial and distributional impacts of the fuel tax policy as well as its effect on CO_2 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 ($\bar{t}$).[42] 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:

$$\sum_{t=\underline{t}}^{\bar{t}} [q_{ct} - \tilde{q}_{ct}] = \sum_{t=\underline{t}}^{\bar{t}} \hat{\beta}(p_{ct} - \tilde{p}_{ct}) = -\sum_{t=\underline{t}}^{\bar{t}} \hat{\beta}(\Delta p_t) > 0, \tag{18}$$

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

$$\sum_{t=\underline{t}}^{\bar{t}} [p_{ct}q_{ct} - \tilde{p}_{ct}\tilde{q}_{ct}] = -\sum_{t=\underline{t}}^{\bar{t}} (\Delta p_t)\tilde{q}_{ct} - \sum_{t=\underline{t}}^{\bar{t}} \hat{\beta}(\Delta p_t)p_{ct} + \sum_{t=\underline{t}}^{\bar{t}} \hat{\gamma}_t p_{ct}, \tag{19}$$

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*

[42] For each cell of individuals and for each day, we may compute $\tilde{q}_{ct} = \hat{\beta}\tilde{p}_{ct} + \sum_{h=t_0}^{t_a-1} \hat{\gamma}_h^1 \mathbb{1}_{t=h} + \sum_{h=t_a}^{t_a-1+(t_a-t_0)} (-\hat{\gamma}^1_{t_a-h+t_a-1})\mathbb{1}_{t=h} + \hat{\alpha}_c + \hat{\mu}_t + \hat{\eta}_{ct}$ 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 rebate, 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 tax rebate 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, consecutive to reduced prices; (iii) the *anticipatory effect* $\sum_{t=\underline{t}}^{\bar{t}} \hat{\gamma}_t p_{ct}$ is related to the fact that storing (resp. postponing) fuel purchases when prices are low (resp. 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:

$$\sum_{t=\underline{t}}^{\bar{t}} \hat{\gamma}_t p_{ct} = \sum_{t=\underline{t}}^{\bar{t}} \hat{\gamma}_t \tilde{p}_{ct} - \sum_{t=\underline{t}}^{\bar{t}} \hat{\gamma}_t \Delta p_t \approx - \sum_{t=\underline{t}}^{\bar{t}} \hat{\gamma}_t \Delta p_t \tag{20}$$

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 (20) 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.[43]

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 11). 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 latter[44]. Figure 12 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 H10 in Appendix).

[43] Empirically, 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.

[44] The 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 CO_2 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 CO_2.[45] In 2021, the annual carbon footprint of a French household amounted to about 20.3tons: the policy increased that footprint by 0.36% .

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

A delicate issue with carbon taxation, fuel excise taxes, and more generally any fuel price increase, is that these mechanisms are all regressive in the sense that low-income individuals bear a higher burden, relative to their income (Figure 3). To nevertheless render such schemes progressive, or at least to cancel their regressivity, the policy maker may combine them with transfers (Douenne, 2020; Young-Brun, 2023). Sallee (2019) has emphasized how important it is to predict fuel consumption at the household level in order to accurately target such transfers, but this task requires much information.

Seeking here to better tailor public interventions to drivers' needs, we compare actual tax rebates to counterfactual transfers, either lump sum or conditional on income, location, fuel spending. To that end, we specify preferences as regards fuel and the rest of consumption on an annual basis.[46] We consider that agents' preferences can be represented by some quasi-additive utility function $U(f, m) = m + \theta \frac{f^{1+\frac{1}{\varepsilon}} - 1}{1+\frac{1}{\varepsilon}}$ over annual fuel consumption f and the outside good m taken as the *numéraire*, θ indexing the agent's need for driving. We omit here the unnecessary index i though agents may differ in their price-elasticity as well as in their need for driving.[47] Maximizing that utility given the budget constraint: $pf + m \leq Y$, where Y is the annual disposable income and p the price of fuel, leads to an isoelastic fuel demand: $f(p) = \left(\frac{p}{\theta}\right)^{\varepsilon}$, the price elasticity being denoted by $\varepsilon < 0$. Fuel expenditures amount then to $E_f(p) = \theta^{-\varepsilon} p^{1+\varepsilon}$. Following any rise in fuel

[45] In 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.

[46] Our 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 viewed as a sufficient statistics in the subsequent analysis.

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

prices from p to $p+\Delta p$, the policy maker might want to compensate consumers who incur the utility loss L:

$$L = V(p+\Delta p) - V(p) = \frac{E_f(p)}{1+\varepsilon}\frac{p^{1+\varepsilon}-(p+\Delta p)^{1+\varepsilon}}{p^{1+\varepsilon}} = l_\epsilon\left(\frac{\Delta p}{p}\right)E_f(p) < 0, \tag{21}$$

where $V(p) = U(f(p), m(p))$ designates the indirect utility function. Computing the empirical counterpart of (21) requires information on two sufficient statistics for each and any agent: the price elasticity $\hat{\varepsilon}$ and ex ante fuel expenditures $\hat{E}_f(p)$, before any price increase.[48]

In what follows, we focus on the observed price increase in 2022 (+€0.6 per liter, from €1.5 in 2021, 'ex ante' hereafter, to €2.1 in 2022, 'ex post' hereafter) and we consider various compensation schemes that aim at keeping consumer utility constant. By construction, a uniform €0.6 annual rebate would achieve full compensation at the household level. However, this solution is more costly than transfers due to the *behavioral effect*, namely the adjustment of demand to lower prices, which reduces the utility loss, in absolute.[49] Though perfect discrimination transfers would also achieve full compensation, while being cheaper than uniform rebates,[50] implementation considerations often require to restrict our attention to uniform or third degree discrimination transfers; in that case, *aggregation error* adds up to previous effects, which exacerbates imperfect compensation.

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 not directly available to the policy maker.[51] We simulate four alternative policies to excise tax rebates: (i) unconditional or lump sum transfers, (ii) income-based transfers, (iii) location-based transfers and (iv) past-fuel spending based transfers In the first case, the transfer is com-

[48]A first-order approximation for small relative price variations $\mathrm{d}p/p$ leads to $L \approx E_f(p)\frac{\mathrm{d}p}{p}$, i.e., $L \approx E_f(p)\mathrm{d}\log(p)$ as in Astier et al. (2023). Ignoring higher orders of the Taylor expansion requires discarding any behavioral effect.

[49]This is because $g_\epsilon(x) = 1 + (1+\epsilon)x - (1+x)^{1+\epsilon} \equiv (1+\epsilon)(l_\epsilon - l_0)\left(\frac{\Delta p}{p}\right)$ takes positive values when $\epsilon \in (-1, 0)$: $g_\epsilon(0) = 0$ and $g'_\epsilon(x) = (1+\epsilon)[1-(1+x)^\epsilon] > 0$.

[50]Perfect discrimination transfers that exactly compensate each household amount to a share of the fuel bill that is equal to $-l_\epsilon < -l_0 = \frac{\Delta p}{p}$; but $\frac{\Delta p}{p}$ precisely corresponds to the relative increase in fuel bill that is offset by rebates, since ex post after-rebate prices coincide with ex ante prices.

[51]For instance, it may require all households to correctly fill up their income tax file, which includes a proxy for fuel expenditures related to commuting costs, etc. Note also that the recent access to bank account data might be an alternative source of information.

puted based on average ex ante fuel spending as well as on average price elasticity; in the second (resp. third) scenario, the policy maker knows both average ex ante fuel spending in each income (resp. location) group and average price elasticity of the household's income (resp. location) group; in the last scenario, she knows ex ante fuel spending at the household level and the price elasticity of the household's ex ante fuel spending group. Since it is only a matter of interpretation to regard previous fuel inflation as a (carbon) tax, we are effectively looking at tax-and-transfer schemes when focusing on a price increase accompanied with such compensating mechanisms. For each policy, we compute the share of households that receive an exact, positive, or negative compensation; households get a positive compensation if the transfer exceeds their utility loss (remember that utility is quasi-linear with respect to income). Formula (21) makes it clear that full compensation is hard to achieve based on uniform or third degree discrimination transfers due to informational frictions, namely imperfect screening of the price-elasticity or/and 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 rebates provided that she values full compensation at the household level less than €18 per household, i.e. the difference between the cost of the uniform rebate, €454 per HH, and the cost of the second-best lump sum transfer based on average past fuel spending, €436 per HH (Table 4). In that specific case, *forecasting error* is negligible since average fuel spending in 2021 is a good predictor of average fuel expenditures in 2022 in the absence of any price increase. Put differently, lump sum transfers are 4% cheaper, hence more affordable to the government, but 61% of households are 'losers', i.e. would prefer a rebate; the average loss of the 'losers' amounts to €398, while the average gain of the 39% 'winners' is €265. 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: 61% in (i) with an average loss of €379, 61% in (ii) with an average loss of €382 and 49% in (iii) with an average loss of €116. Figure 13 quantifies the value of information on income, location, and ex ante fuel spending for targeting tax-and-transfer schemes. A reliable knowledge of ex ante fuel spending helps a lot to reduce the financial cost of the intervention by tightening the eligibility condition, hence to diminish the heterogeneity of the impact of the public intervention (namely, the dispersion of gains and losses). Note yet that the average loss among households would be high no matter how the transfer is designed (lump sum, income-based, location-based, or based on ex ante fuel spending), i.e. regardless of the information available to the policy maker, due to unobserved heterogeneity in fuel consumption.

As another caveat on top of imperfect compensation, transfers based on ex ante fuel spending[52] do not provide the right incentives to lower emissions, since they are not designed to correct for polluting externalities, contrary to Pigouvian taxation. As regards long-run carbon pricing, a better solution could consist in committing to a gradually rising tax path, while adjusting excise taxes in the short run in the event of a deviation from that path. Such a scheme resembles to a mechanism called the 'floating TICPE' that prevailed in France between 2000 and 2002, whereby fuel excise tax rates could be adjusted over time, depending on circumstances. That tension between short-term desirability and long-term sustainability of the public intervention, which moderates the usage of temporary tax rebates, has been somehow encompassed in the 'end-of-the-month' vs. 'end-of-the-world' trade-off.

6.3 Rebates or transfers?

We last compare policy tools that are likely to compensate for fuel inflation 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 a portion only of the population may be eligible to such transfers, governments facing a tight budget constraint that nonetheless seek to alleviate excessive losses consecutive to some price surge, for specific groups, might favor the latter.

We determine the optimal policy rule as a function of the level of public funding that the government is ready to consent when compensating for fuel inflation. Mimicking here the aggregation of individual utilities into a social welfare function, we posit that the decision maker minimizes a combination of agents' relative utility losses; those losses are expressed, in absolute, as a fraction of income (remember that preferences are quasi-linear). To do so, the government disposes of two policy instruments, rebates r and transfers T, which may be uniform or targeted towards segments of consumers. It also faces some budget

[52] or on location, to a smaller extent.

constraint: the policy cannot cost more than some exogenous C, which may well be issued from revenue recycling (which would approximately amount to €500 per household here). We further assume that the planner weighs more individuals with a high fuel budget share, a particular case being the 'alleviation of excessive (relative) losses', which is reminding of a justice principle called 'alleviating poverty' exposed, e.g., in Saez and Stantcheva (2016).[53] Using similar notations as above, we thus consider the following tool-specific programs with $f_h(p) = \left(\frac{p}{\theta_h}\right)^\epsilon$ referring to household h's fuel consumption, $E_h = \theta_h^{-\epsilon} p^{1+\varepsilon}$ to ex ante fuel spending, and $RL_h = \frac{L_h}{Y_h}$ to the ex post relative loss, namely the utility loss expressed as a share of income:

$$\min_{(r_h)} \sum_{h=1}^{N} \alpha_h(|RL_h^r|)\,|RL_h^r| \quad \text{s.t.} \quad \sum_{h=1}^{N} r_h f_h(p + \Delta p - r_h) \leq C \tag{22}$$

$$\min_{(T_h)} \sum_{i=1}^{N} \alpha_h(|RL_h^T|)\,|RL_h^r| \quad \text{s.t.} \quad \sum_{h=1}^{N} T_h \leq C \tag{23}$$

From previous subsection, $L_h^r = l_\epsilon^r E_h = \frac{E_h}{1+\varepsilon} \frac{p^{1+\varepsilon} - (p+\Delta p - r_h)^{1+\varepsilon}}{p^{1+\varepsilon}}$ and $L_h^T = T_h + l_\epsilon E_h = T_h + \frac{E_h}{1+\varepsilon} \frac{p^{1+\varepsilon} - (p+\Delta p)^{1+\varepsilon}}{p^{1+\varepsilon}}$. Social weights $\alpha_h(\cdot)$ are assumed to be nondecreasing; polar cases correspond to 'Ralwsianism' where $\alpha_h = \mathbb{1}\{|RL_h| = \max_j |RL_j|\}$ and 'pure utilitarianism' where $\alpha_h = 1/N$. A special case is 'alleviating excessive (relative) losses', namely $\alpha_h = \mathbb{1}\{RL_h > \overline{RL}\}$, where $\overline{RL}$ is some threshold that the government considers as excessive. Since the objective function is decreasing in policy instruments, the budget constraint of the government binds in the optima:

$$\sum_{h=1}^{N} E_h \frac{r_h}{p} \left(\frac{p + \Delta p - r_h}{p}\right)^\epsilon = C, \quad \sum_{h=1}^{N} T_h = C \tag{24}$$

At the agent's level, it is always more costly to achieve full compensation with rebates (cf. footnote 49):

$$\forall h, \quad L_h^r = 0 \Longleftrightarrow r_h = \Delta p \Longleftrightarrow \frac{C_h^r}{E_h} = \frac{\Delta p}{p}$$

[53] To 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.

than with transfers:

$$\forall h, \quad L_h^T = 0 \Longleftrightarrow \frac{C_h^T}{E_h} = \frac{1}{1+\varepsilon}\frac{(p+\Delta p)^{1+\varepsilon} - p^{1+\varepsilon}}{p^{1+\varepsilon}} = -l_\epsilon < -l_0 = \frac{\Delta p}{p} = \frac{C_h^r}{E_h}$$

Within that framework, we are able to characterize the optimal policy rule given any level of public funding C. In our application, we consider rebates and transfers that can be either uniform or means-tested, eligible households then corresponding to the bottom half of income. Our simulations indicate that means-tested rebates should be chosen by the government (Figure 14) when its objective function consists in alleviating the top 10% highest losses (in absolute), relative to income; Figure H12 in Online Appendix displays the results with other thresholds. Due to non-discriminatory rules, this first best might be forbidden in practice, though. That solution aside, there exists some funding threshold $\bar{C}$ (about €160 per household) such that means-tested transfers dominate uniform rebates below $\bar{C}$, while the contrary prevails above $\bar{C}$. Put differently, there is a case for means-tested transfers when the budget constraint is tight: despite achieving imperfect compensation at the agent's level, transfers are effective at targeting the right individuals from the planner's viewpoint. This is because income acts as a satisfying screening device for the relative loss. When more money can be devoted to the policy, though, transfers become marginally less attractive than rebates, which are more expensive but achieve exact compensation at the agent's level. This result can be more generally derived under any nondecreasing weights α_h in e_h; the relative location of $\bar{C}$ with respect to the perfect discrimination transfer C^T depends on the social weight on households with high relative losses. Means-tested transfers are preferred to rebates when $\bar{C} < C^T$, hence provided that the government puts enough weight on households with high relative losses. Interestingly, those simulations are conform to the actual choice made by the French government choice when replacing rebates with a means-tested transfer at the beginning of 2023 for both financial and distributional considerations; we therefore believe that they may be of some practical guidance to policy makers.

7 Conclusion

This paper has shown that the researcher who aims at causally inferring the price elasticity of fuel demand should dispose of granular, high-frequency data, on top of relying on exogenous price variations. Daily data permit to disentangle anticipation effects from the price effect; we have leveraged 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, green taxation, etc. Equipped with such ingredients, we estimate an average price elasticity of -0.31. That price elasticity exhibits sizeable dispersion, primarily in the fuel spending dimension: individuals who consume more fuel are more inelastic, including those who commute for professional reasons, and especially when they are also liquidity-constrained. By contrast, income and location are not associated with significantly different average price elasticity. Previous estimates along with anticipation effects have enabled us to evaluate financial, distributional and environmental effects of the current tax policy.

We have also simulated various policy experiments in which counterfactual transfers would have been provided to households, possibly depending on their observed characteristics. We have shown that exact compensation at the household level, as is the case with rebates, costs 4% more than imperfect compensation achieved by lump sum transfers. Assuming further that the government disposes of reliable information on households' ex ante fuel spending would not help much in this regard: partial compensation would still arise due to 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 to substantially reduce the dispersion of gains and losses within that group. Last, a decision maker seeking to alleviate excessive losses, relative to income, would ideally prefer means-tested tax rebates; feasibility might require resorting to means-tested transfers rather than uniform tax rebates when she faces tight budget constraints.

From a market design viewpoint, our results help to better tailor compensation mechanisms to drivers' needs, hereby to improve the targeting of transfer mechanisms; more generally, they apply to other markets exposed to inflation. Further research should therefore concentrate on other institutional settings and other markets in order to investigate whether empirical findings, especially as regards policy rules, extend to different contexts.

References

Adam, M., Bonnet, O., Fize, E., Loisel, T., Rault, M., and Wilner, L. (2023). Is carbon tax truly more salient? evidence from fuel tourism at the France-Germany border. Technical report.

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 hétérogenes 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. *Journal of Political Economy*, 126(4):1504–1557.

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.

Chernozhukov, V., Demirer, M., Duflo, 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.

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.

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.

Hendel, I. and Nevo, A. (2006). Measuring the implications of sales and consumer inventory behavior. *Econometrica*, 74(6):1637–1673.

Kilian, L. (2022). Understanding the estimation of oil demand and oil supply elasticities. *Energy Economics*, 107:105844.

Kilian, L. and Zhou, X. (2023). Heterogeneity in the pass-through from oil to gasoline prices: A new instrument for estimating the price elasticity of gasoline demand.

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.

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.

Montag, F., Sagimuldina, A., and Schnitzer, M. (2021). Does tax policy work when consumers have imperfect price information? Theory and evidence. CESifo Working Paper.

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. (2019). Pigou creates losers: On the implausibility of achieving pareto improvements from efficiency-enhancing policies. *NBER Working Paper 25831*.

Young-Brun, M. (2023). Public acceptability of carbon taxation: a model of political support with income and urban-rural inequality. Working paper.

A Figures

A.1 Descriptive statistics

Figure 1: Distribution of fuel spending

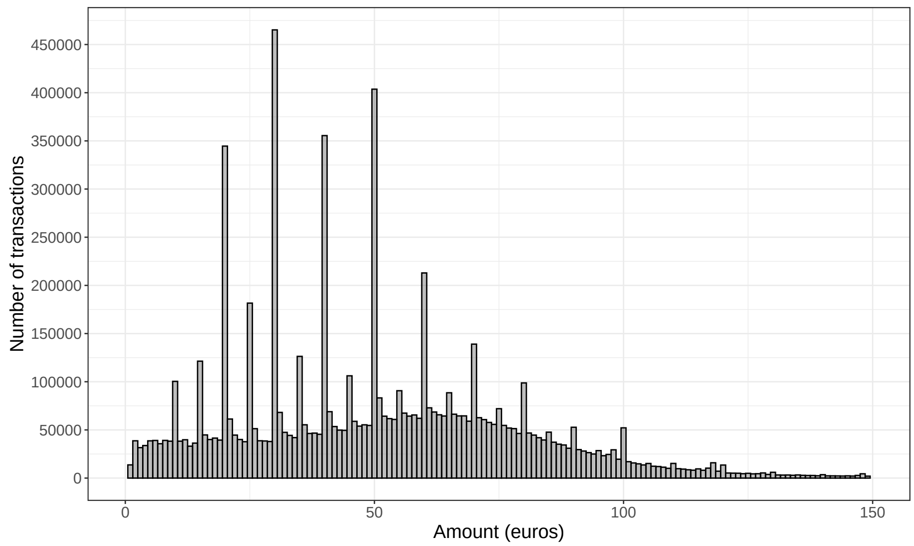

Note. Transaction-level fuel expenditures.

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

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

Figure 2: Interpurchase duration (in days)

(a) Transaction level (b) Household level

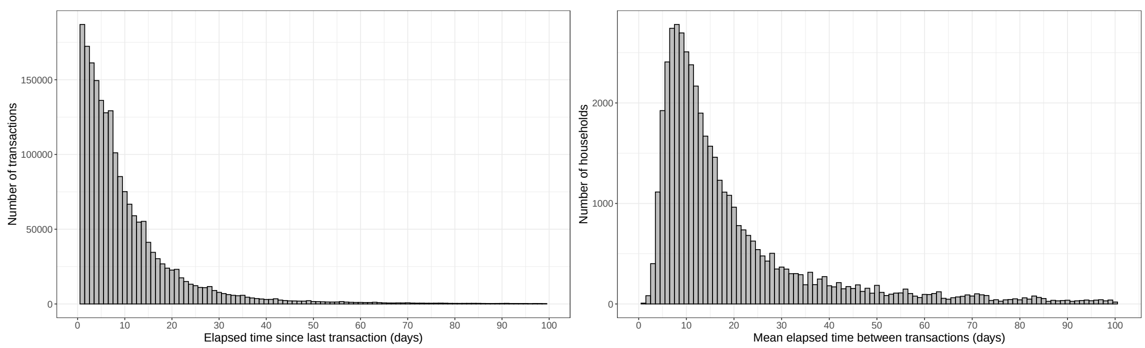

Lecture. 75,000 transactions have occurred 10 days since last purchase. For 2,500 households, the average duration between two transactions is 10 days.
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure 3: Average fuel spending (by income and location)

(a) Income

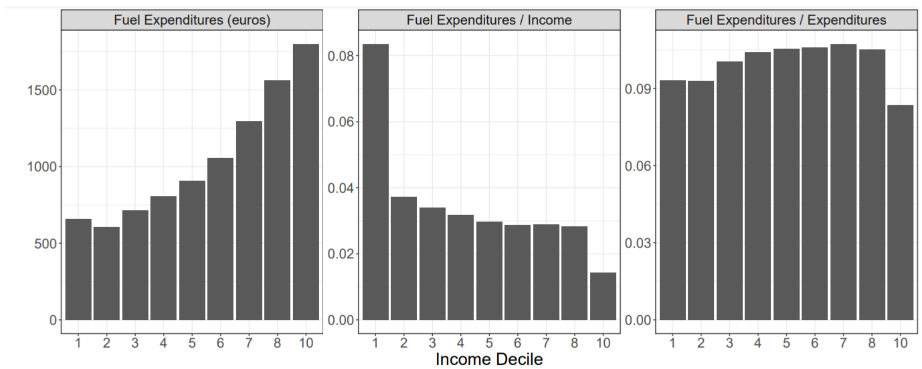

(b) Location

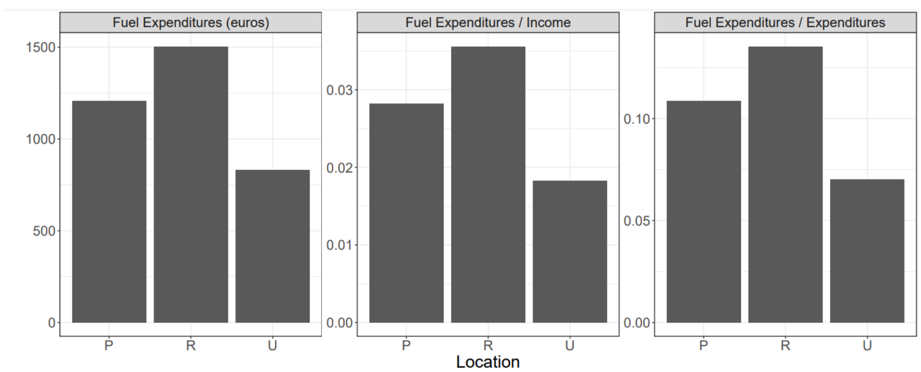

Note. Fuel expenditures: card payments in gas stations. Total expenditures: both card payments and checks. Location: peri-urban (P), rural (R), or urban (U). Fuel spending increases with income. The budget share of fuel decreases with income; it is higher in rural areas.
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale.*

Figure 4: Fuel prices and purchases

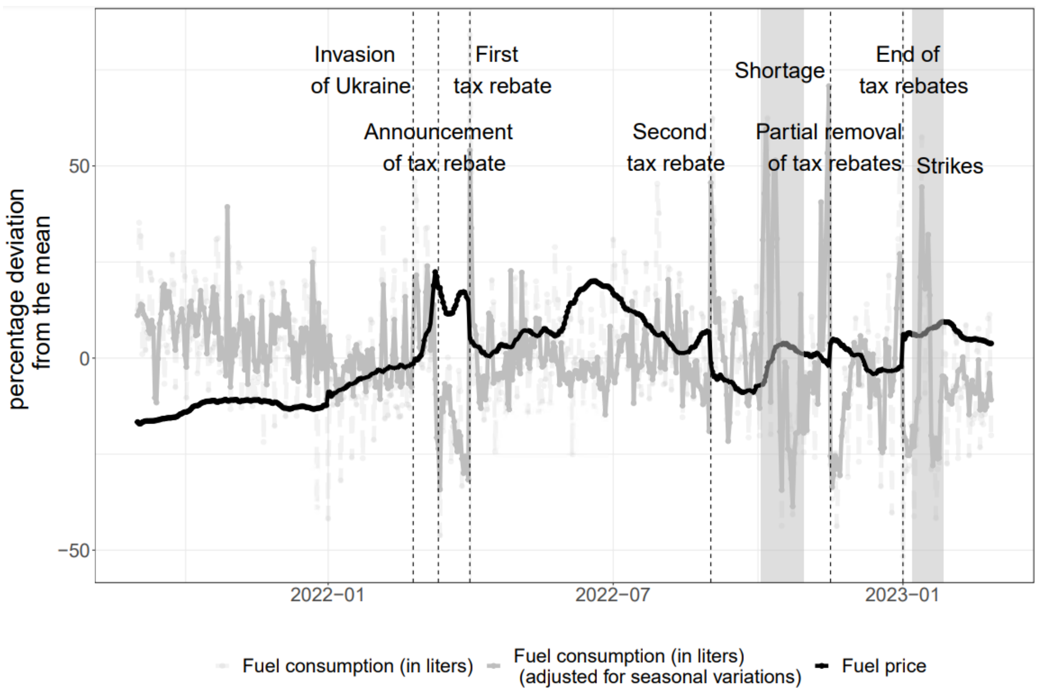

Note. Fuel prices (including taxes) and purchases (in liters). Purchases are adjusted for seasonal variations thanks to GIE-CB data from September 2021 to February 2023. Dashed lines correspond to the invasion of Ukraine and policy interventions. The first intervention on April 1st is a €0.18 per liter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

A.2 Identification

Figure 5: Price versus anticipation effects

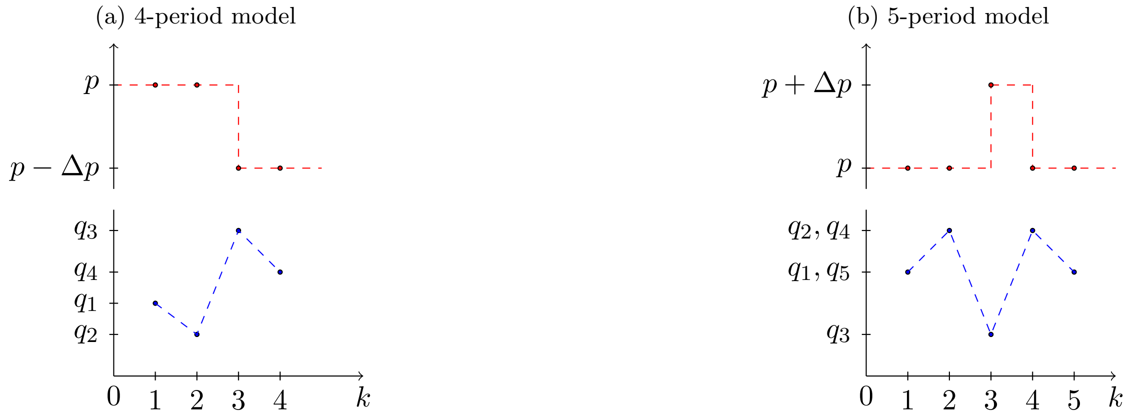

Figure 6: Fuel prices around September 1st

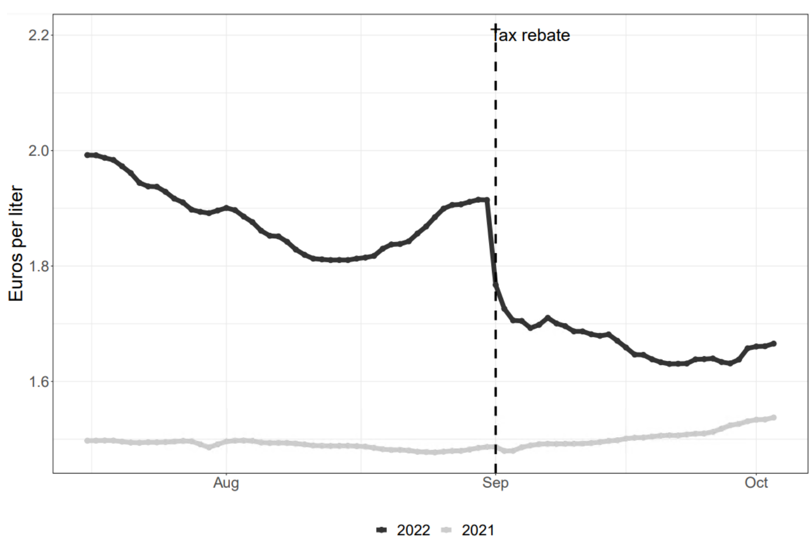

Note. Fuel prices (including taxes). The dashed line corresponds to the extra rebate of €0.12 per liter purchased, implemented from September 1st, 2022.

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

Figure 7: Fuel purchases around September 1st

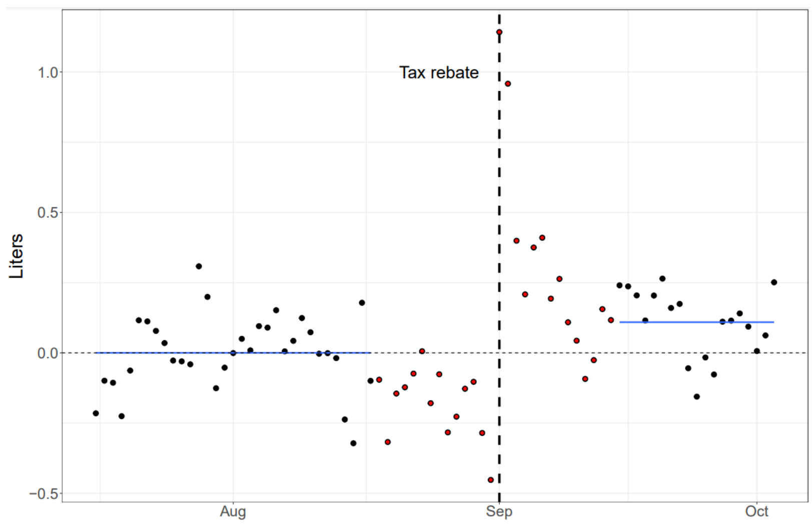

Note. Dots correspond to adjusted daily fuel purchases from July 15th to October 4th, 2022. The adjustment for seasonal variation relies on 2021 as the baseline year. Purchases are normalized so that they sum up to 0 before the anticipation window. Dashed line: Extra rebate of €0.12 per liter starting on September 1st, 2022. Red dots: Anticipation window (7 days before and after the implementation of the rebate). Blue lines: Average purchases before and after the implementation of the rebate, excluding the anticipation window.
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

A.3 Estimation results

Figure 8: Estimated price elasticity (depending on the anticipation bandwidth Δ)

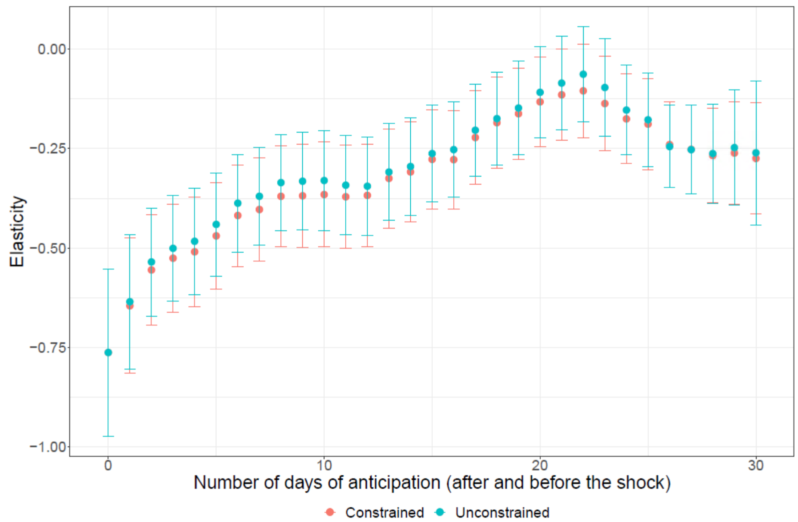

Note. Estimated price elasticity for both constrained and unconstrained specifications (equation (11)), depending on the anticipation bandwidth. Estimation period: from July 15th to October 4th in 2021 and 2022.
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

A.4 Heterogeneity of the price elasticity

Figure 9: Heterogeneity along some observed characteristics

(a) Income

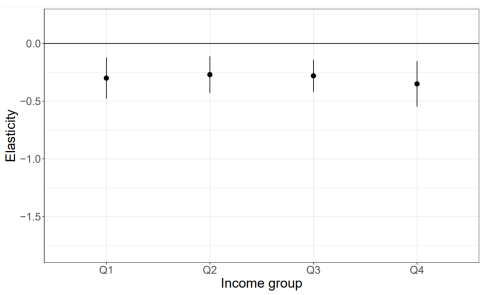

(b) Location

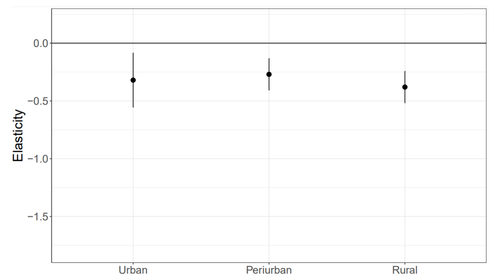

(c) Family status

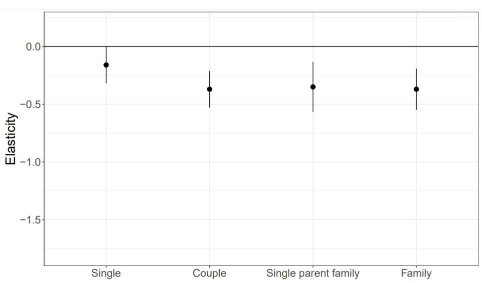

(d) Age

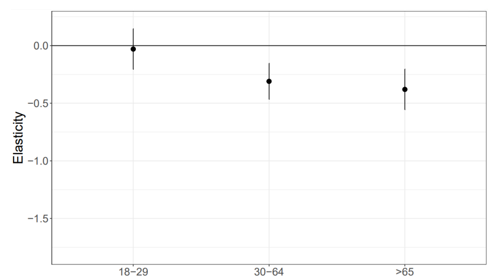

(e) Fuel spending

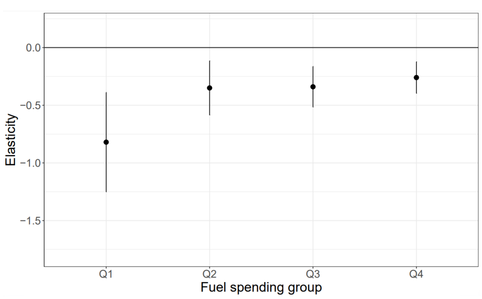

(f) Liquidity × Fuel spending

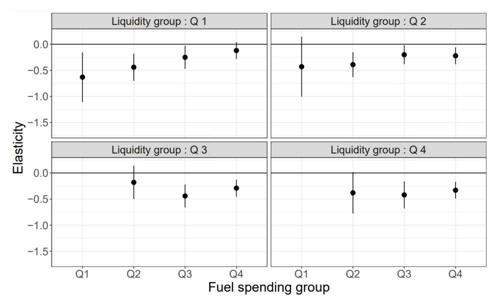

Note. Estimated price elasticity using equation (11). Estimation period: from July 15th to October 4th in 2021 and 2022. Subsamples defined according to corresponding households' characteristics (Continuous variables: Observed before the intervention from January to June 2022, Discrete variables: Observed in June 2022). Two-way clustering of standard errors at cell and year-day levels.
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale.*

Figure 10: Sorted Group Average Treatment Effects

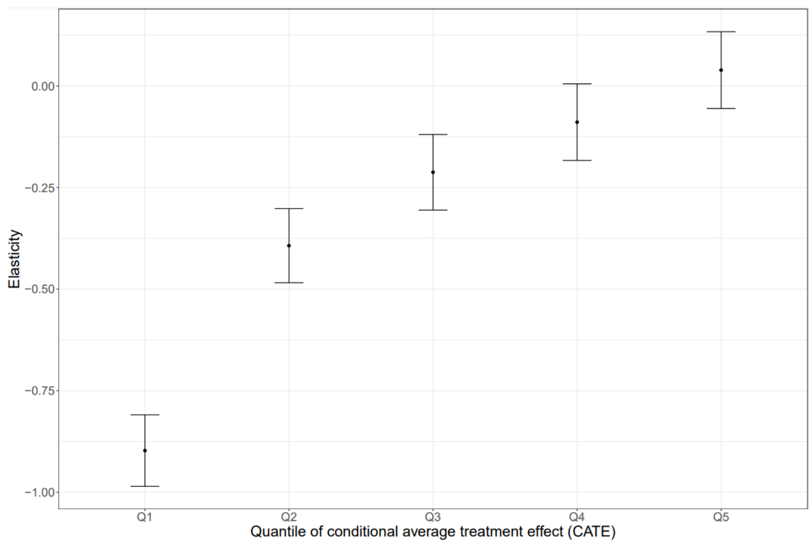

Note. Average elasticity of conditional average treatment effect (CATE) in each quintile. Estimation period: from July 15th to October 4th in 2021 and 2022, excluding the anticipation window, based on corresponding subsamples. Average elasticity estimates and standard errors correspond to medians over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

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

A.5 Counterfactual simulations

Figure 11: Distributional effects of rebates (by income)

(a) in euros (monthly) (b) in % of annual income

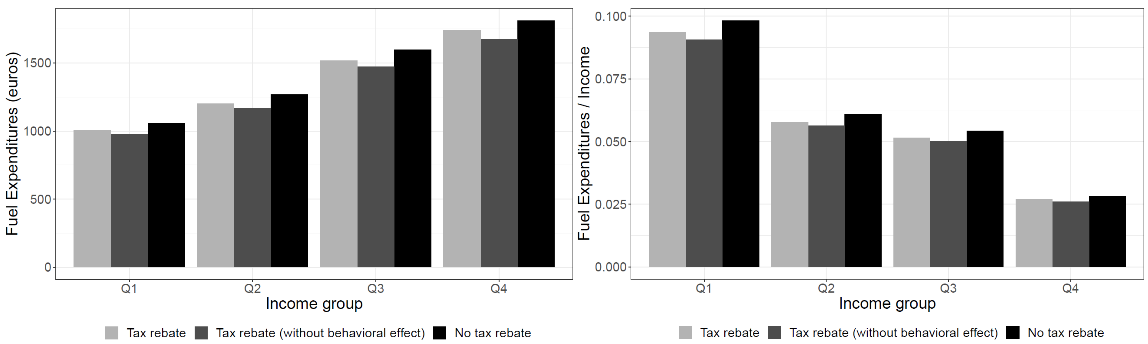

Note. Simulation period: from January 8th 2022 to January 8th 2023.
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure 12: Distributional effects of rebates (by location)

(a) in euros (monthly) (b) in % of annual income

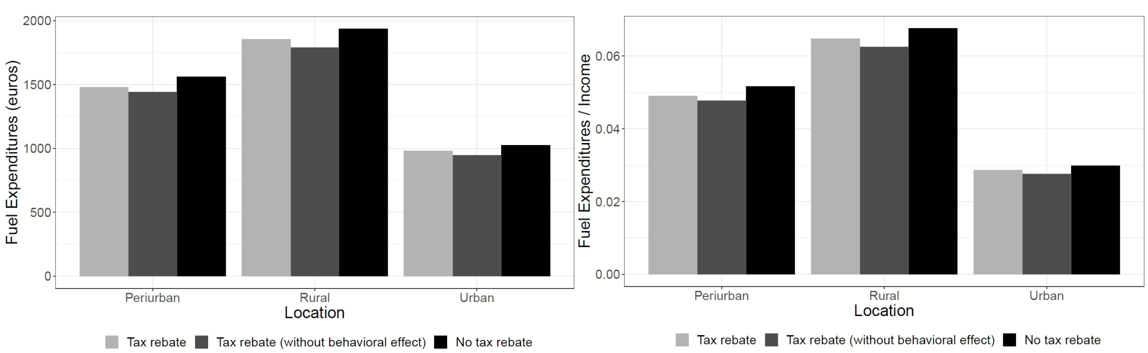

Note. Simulation period: from January 8th 2022 to January 8th 2023. Location is defined by the bank.
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure 13: Gains and losses from 'compensating' tax-and-transfer scheme

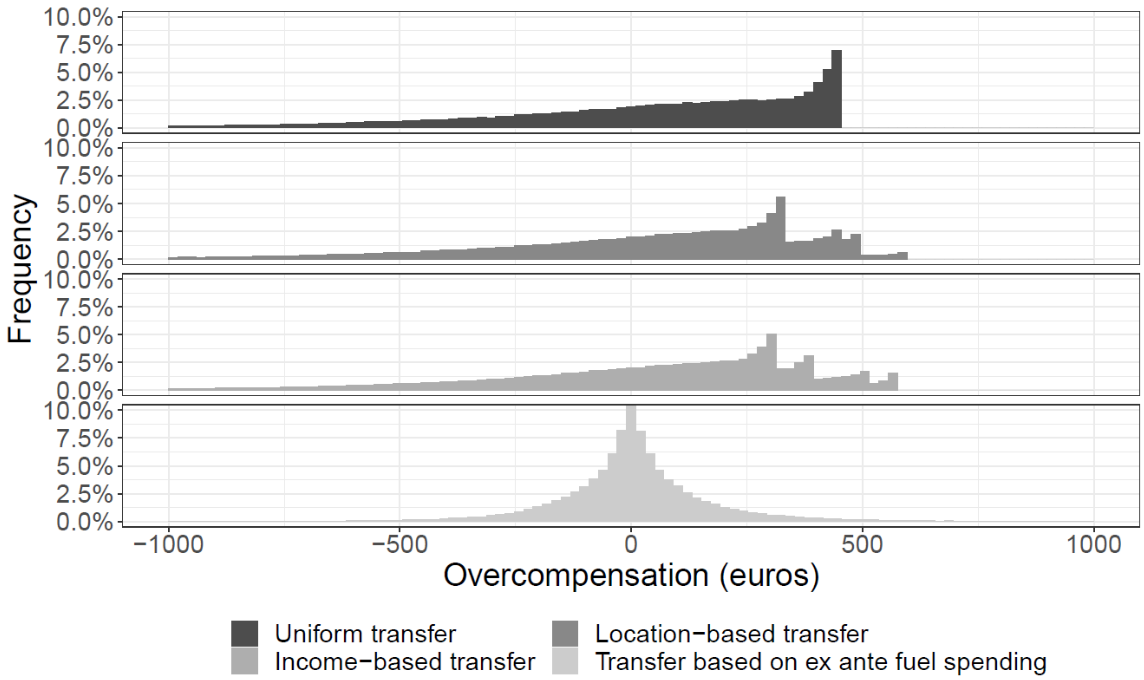

Note. We compare the relative impact of transfers with respect to rebates so as to compensate households for the observed price increase from 2021 to 2022 (+€0.6 per liter, i.e. from €1.5 to €2.1). We compute households' gains and losses, in euros, associated with transfers rather than with rebates.
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure 14: Optimal policy instrument: alleviating top 10% utility loss in nominal terms relative to income

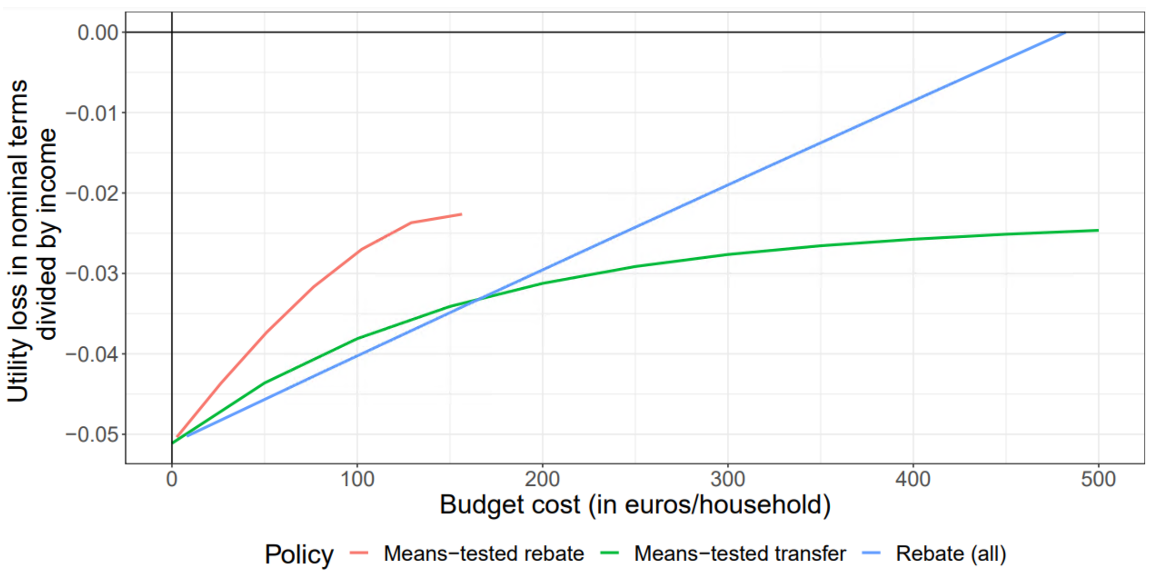

Note. For any given level of public funding (x-axis), we compare (conditional or unconditional) rebates and transfers as policy tools designed to compensate households for the observed price increase from 2021 to 2022 (+€0.6 per liter, i.e. from €1.5 to €2.1).
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale.*

B Tables

Table 1: Estimation based on September 1st tax rebate

	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 (11) 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*.

Table 2: 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 (13) 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*.

Table 3: 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
Age groups						
Share of below 30 years	0.12	0.000	0.06	0.001	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
Location						
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
Familial status						
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
Occupations						
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.*

Table 4: Simulated 'compensating' policies to 2022 fuel inflation (financial and distributional effects)

	Tax rebate	Transfer			
		Uniform	Income-based	Location-based	Based on ex ante fuel spending
Average cost per household (€)	454	436	436	436	436
Share of HH with positive compensation (%)	0	61	61	61	49
Average positive compensation (€)		265	257	258	129
Share of HH with negative compensation (%)	0	39	39	39	48
Average negative compensation (€)		398	379	382	116

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

Online Appendix

C Details on toy econometric specifications

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

C.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:

$$q_0 + \beta \bar{p} = \bar{q} - \sum_{k=2}^{3} \gamma_k \frac{T_k}{T} \tag{25}$$

$$q_0 \bar{p} + \beta \overline{p^2} = \overline{pq} - \sum_{k=2}^{3} \gamma_k p_k \frac{T_k}{T} \tag{26}$$

$$q_k - q_0 - \beta p_k = \gamma_k, \ \forall k = 2, 3 \tag{27}$$

Plugging (27) into (25) 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 (27) into (26) 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

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

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

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

that is,

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

C.1.2 Constrained estimator

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

$$q_0 + \beta\bar{p} = \bar{q} \tag{28}$$

$$q_0\bar{p} + \beta\overline{p^2} = \overline{pq} - \sum_{k=2}^{3}\gamma_k p_k \frac{T_k}{T} \tag{29}$$

$$\gamma_2(T_2 + T_3) = T_3[q_2 - q_3 - \beta(p_2 - p_3)] \tag{30}$$

Plugging the latter into (29) and using (28) combined with the constraint yields

$$\boxed{\hat{\beta}^c = \frac{(T_2 + T_3)[T_1q_1(T_3 + T_4) - T_4q_4(T_1 + T_2)] + (T_2T_4 - T_1T_3)(T_2q_2 + T_3q_3)}{[T_2T_4(T_1 + T_2) + T_1T_3(T_3 + T_4)]\Delta p}}$$

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 boils down to

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

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

$$\hat{\gamma}_2^c = T_3\frac{T_1(T_3 + T_4)(q_2 - q_1) + T_4(T_1 + T_2)(q_4 - q_3)}{T_2T_4(T_1 + T_2) + T_1T_3(T_3 + T_4)} < 0$$

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

C.1.3 Naive estimator

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

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

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

It follows that

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

C.2 Price surge + tax rebate (5-period model)

C.2.1 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:

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

$$q_0 \bar{p} + \beta \overline{p^2} = \overline{pq} - \sum_{k=2}^{4} \gamma_k p_k \frac{T_k}{T} \tag{32}$$

$$q_k - q_0 - \beta p_k = \gamma_k, \ \forall k = 2, \ldots, 4 \tag{33}$$

Plugging (33) into (31) 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 (33) into (32) 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

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

infeasible since $p_1 = p_5$.

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

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

C.2.2 Constrained estimator

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

$$q_0 + \beta \bar{p} = \bar{q} \tag{34}$$

$$q_0 \bar{p} + \beta \overline{p^2} = \overline{pq} - \sum_{k=2}^{4} \gamma_k p_k \frac{T_k}{T} \tag{35}$$

$$\gamma_2 (T_2 + T_3) + \gamma_4 T_4 = T_3 [q_2 - q_3 - \beta (p_2 - p_3)] \tag{36}$$

$$\gamma_2 T_2 + \gamma_4 (T_3 + T_4) = T_3 [q_4 - q_3 - \beta (p_4 - p_3)] \tag{37}$$

It follows from (36) and (37) that

$$\gamma_2 = \frac{[(T_3 + T_4) q_2 - T_3 q_3 - T_4 q_4] - \beta [(T_3 + T_4) p_2 - T_3 p_3 - T_4 p_4]}{T_2 + T_3 + T_4}$$

$$\gamma_4 = \frac{[(T_2 + T_3) q_4 - T_2 q_2 - T_3 q_3] - \beta [(T_2 + T_3) p_4 - T_2 p_2 - T_3 p_3]}{T_2 + T_3 + T_4}$$

Plugging the latter into (35) and using (34) combined with the constraint yields

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

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

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

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

hence

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

and

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

C.2.3 Naive estimator

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

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

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

It follows that

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

C.3 Numerical examples

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

Table C1: Fuel prices and purchases around September 1st

Period	07-15 to 08-17	08-18 to 08-31	09-01 to 09-14	09-15 to 10-04
k	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

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

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

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

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

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

Table C2: 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.

D Testing for the presence of anticipations

We consider the following model:

$$q_{cty} = \beta p_{cty} + \sum_{h=t_2-\Delta}^{t_2+\Delta} \gamma_{hy} \mathbb{1}_{h=t} \mathbb{1}_{y=2022} + \alpha_{cy} + \mu_t + \eta_{cty}. \tag{38}$$

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:

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

The estimation consists then in the regression of q_{cty} on p_{cty}, the set of dummies $\mathbb{1}_{h=t}\mathbb{1}_{y=2022}$ where $h \in [t_2 - \Delta, t_2 + \Delta]$, daily fixed effects μ_t and cell fixed effects α_{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}.

D.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.

D.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 $\sum_h \tilde{\gamma}^2_{hy}$, which we compare to $\sum_t \nu^2_{ty}$. Specifically, we test whether the $\sum_{h'} \tilde{\gamma}^2_{h'y}$ at time h', where we suspect anticipations, is higher than the 90th percentile of the distribution of $\sum_t \nu^2_{ty}$.

Figure D1: Testing the presence of anticipations in weeks surrounding the shock

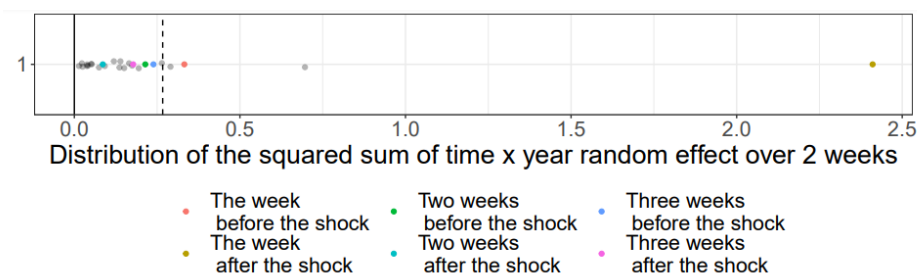

Note. Grey points correspond to placebo estimates

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

From Figure D1 one can reject the absence of anticipations during both the week before and the week after the shock with a 90% level of confidence.

D.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 D2: Testing whether the sum of anticipation is equal to 0 (period from 7 days before to 7 days after the shock)

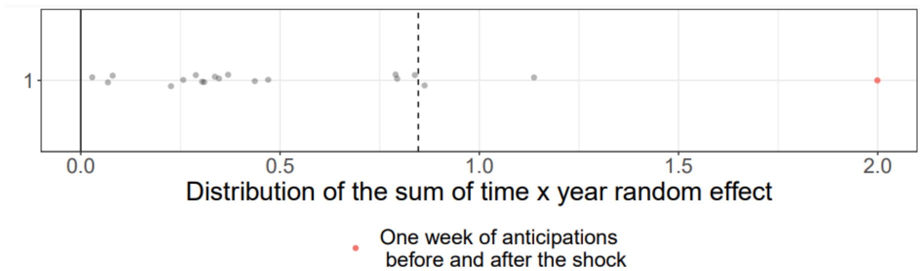

Note. Grey points correspond to placebo estimates

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

Figure D3: Testing whether the sum of anticipation is equal to 0 (period from 14 days before to 14 days after the shock)

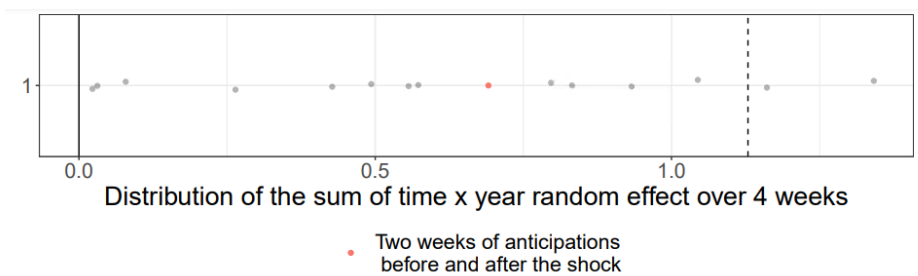

Note. Grey points correspond to placebo estimates

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

From Figure D2 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 D3.

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

D.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):

$$q_{cty} = \beta p_{cty} + \alpha_{cy} + \mu_t + \eta_{cty}. \tag{39}$$

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:

$$q_{cty} = \alpha_{cy} + \mu_t + \eta_{cty} \tag{40}$$

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.

E 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.

F Data details

Two concerns have been raised by the literature as regards the external validity of bank account data (Baker, 2018): representativeness and completeness. We therefore resort to 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 as regards 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 F3). 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 F1 which suggests that Crédit Mutuel customers are wealthier: they dispose of higher income (Figure F1), detain more assets (Figure F2), and spend more than the average (Figure F3). The pregnancy of liquidity constraints can be assessed by looking at the liquid wealth-to-income ratio, about 10, meaning that, on average, households dispose of 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 F4 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 F6) 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 F5).

Second, our measure of income is more volatile (Figure F7) than the one measured by Insee.[54] 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 F8). If anything, Crédit 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.

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

F.1 Data: External validity

Table F1: Summary statistics

	Weighted sample
# of observations	181,527
	Banking variables (sample means)
Monthly Spending	2,721
Fuel (cards)	94
Income	3,622
Financial Assets	
Liquid financial Assets	38,116
Illiquid financial Assets	23,469
Ratio liquid assets/deposit account	3.1
	Household head characteristics (sample means)
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 F2: 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 F3: 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 F4: 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 (resp. who primarily bank at Crédit Mutuel Alliance Fédérale) detain 88.4% (resp. 88.1%) of their financial assets in their main bank.
Source. Histoire de vie et Patrimoine French wealth survey (2017).

Figure F1: Distribution of income (transaction data vs. survey data from *ERFS*, Insee)

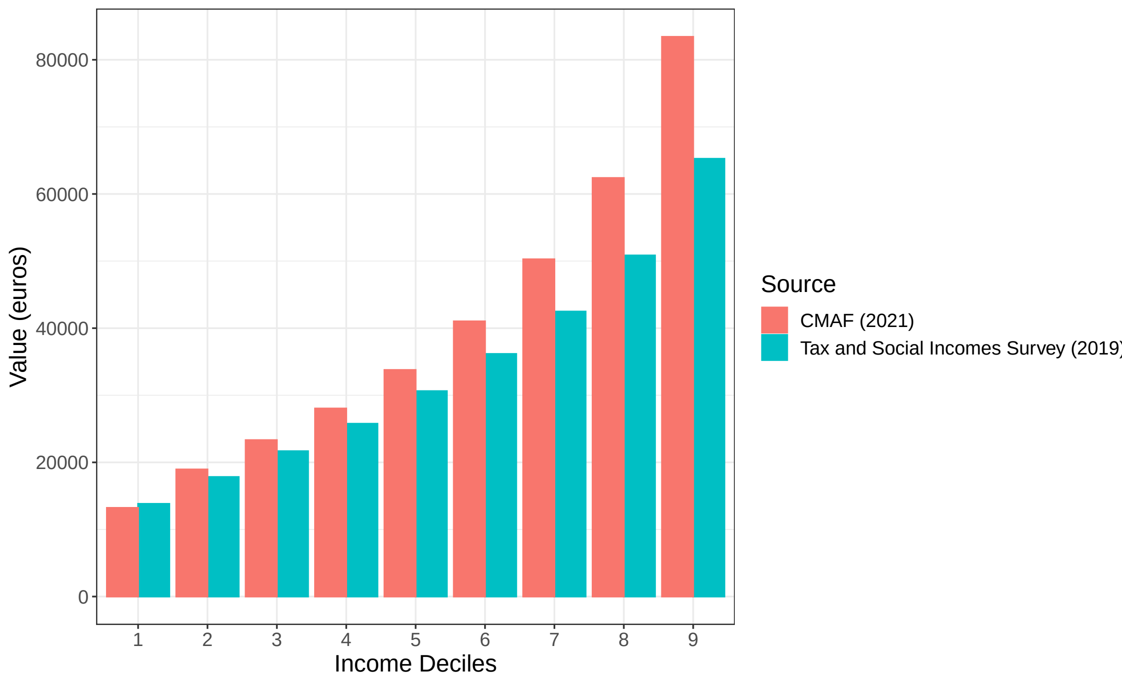

Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*; *Enquête sur les Revenus Fiscaux et Sociaux* (ERFS) survey.

Figure F2: Distribution of household financial wealth by income (transaction data vs. survey data from *Histoire de Vie et Patrimoine*, Insee)

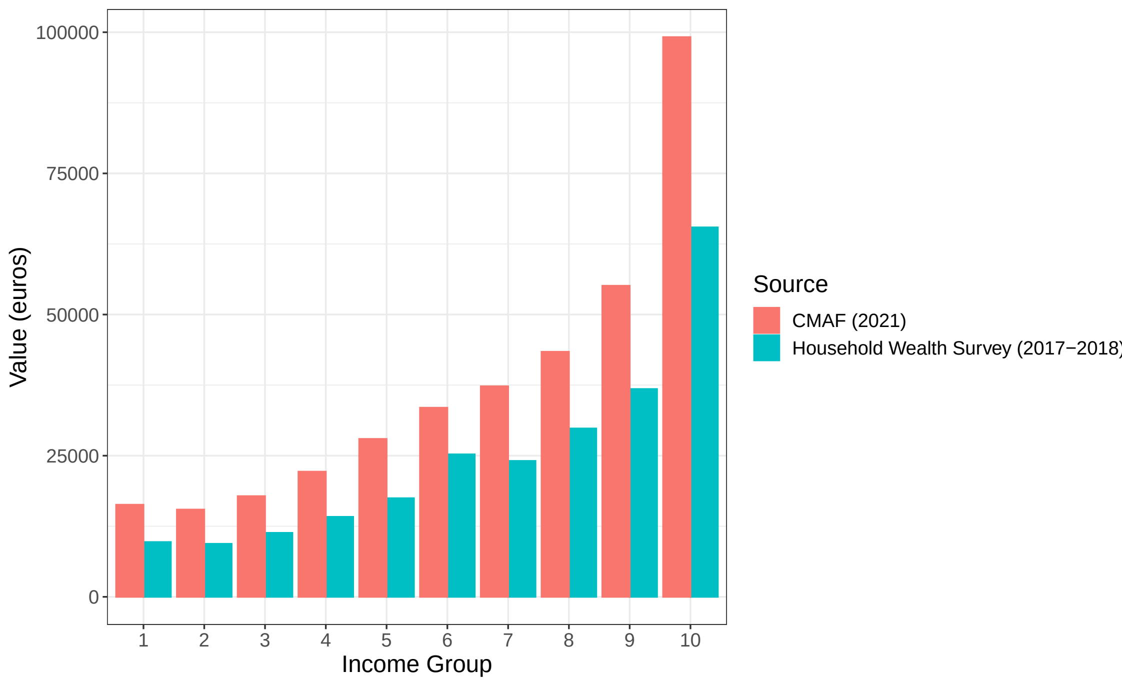

Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*; *Patrimoine* survey.

Figure F3: Distribution of household monthly expenditures by income (transaction data vs. survey data from *Budget des Familles*, Insee)

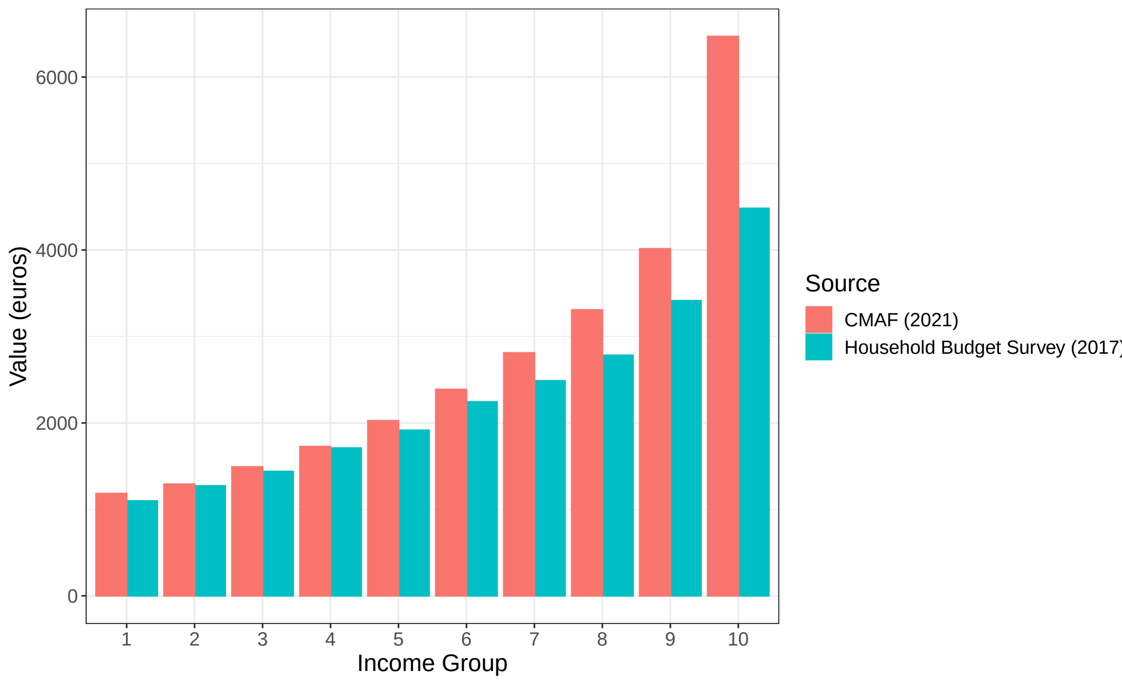

Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*; *Budget des Familles* survey.

Figure F4: Distribution of monthly fuel spending, by income (transaction data vs. survey data from *Budget des Familles*, Insee)

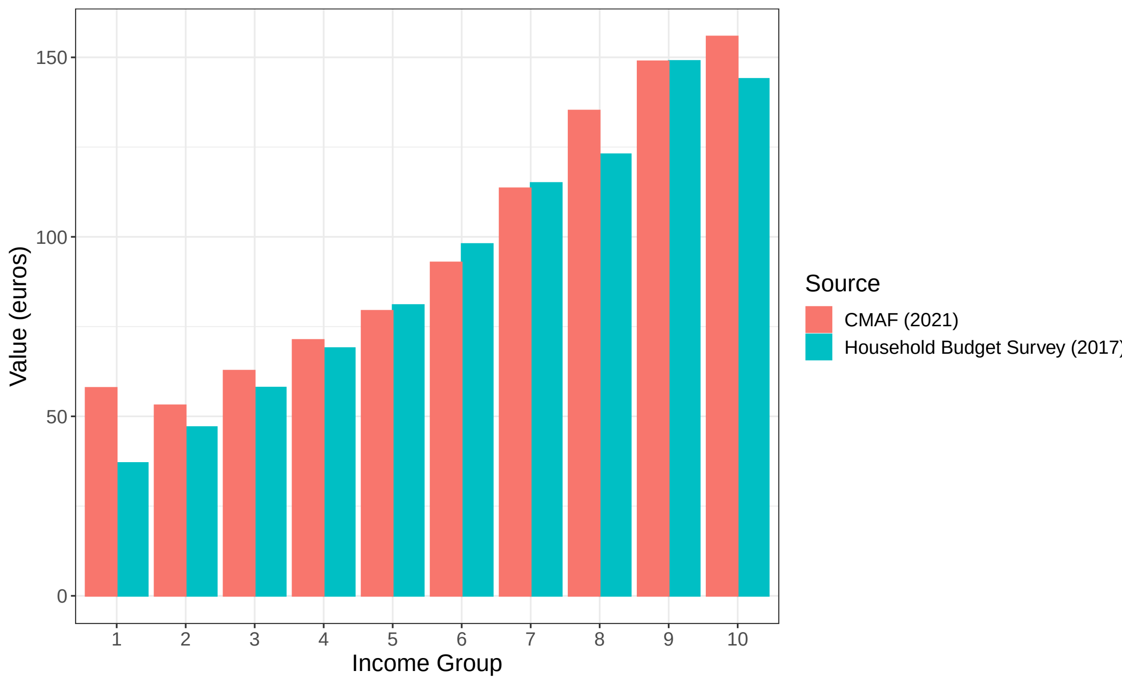

Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*; *Budget des Familles* survey.

Figure F5: Evolution of spending (transaction data vs. aggregate data from the French interbank network)

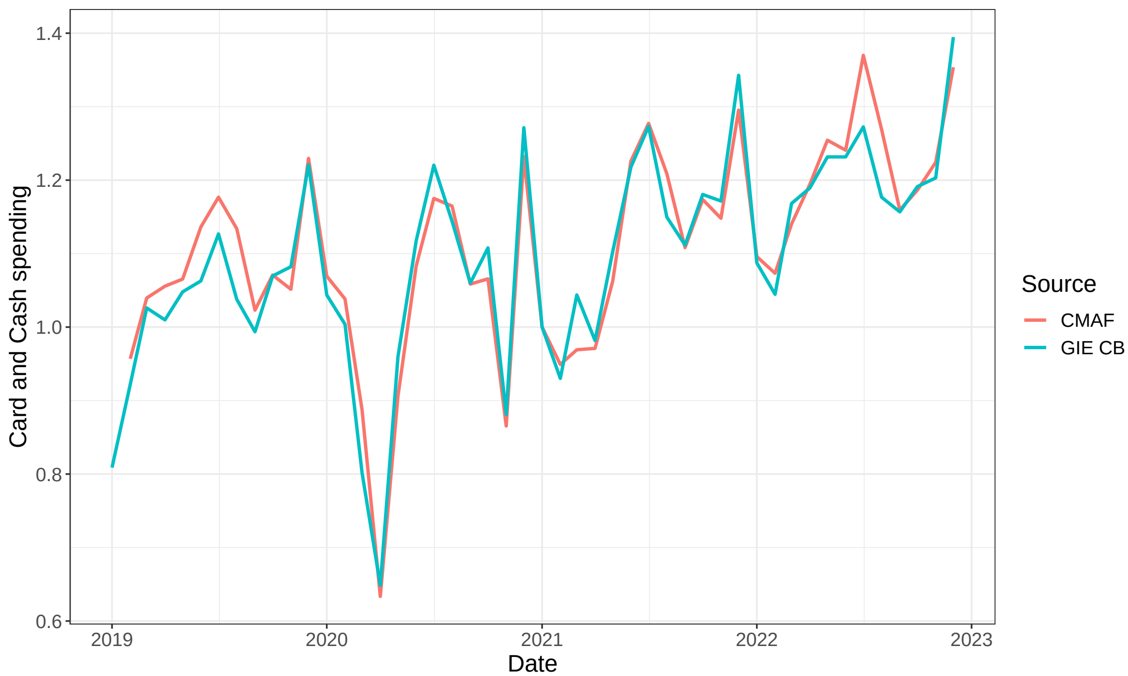

Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*; *GIE-CB* data.

Figure F6: Evolution of fuel spending (transaction data vs. aggregate data from the French interbank network)

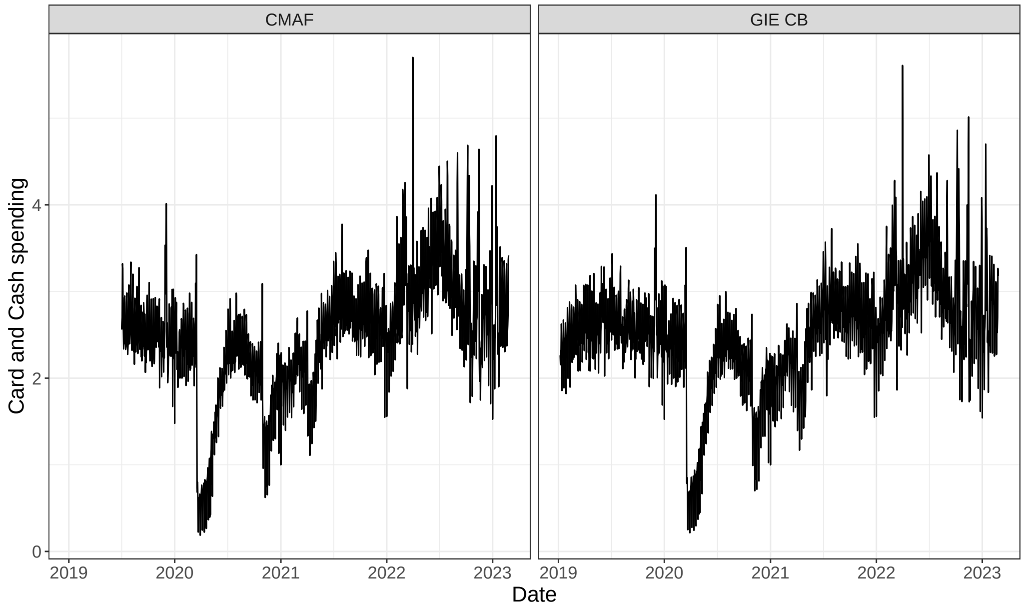

Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*; *GIE-CB* data.

Figure F7: 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 F8: Liquid Assets (transaction data vs. aggregate data from Banque de France)

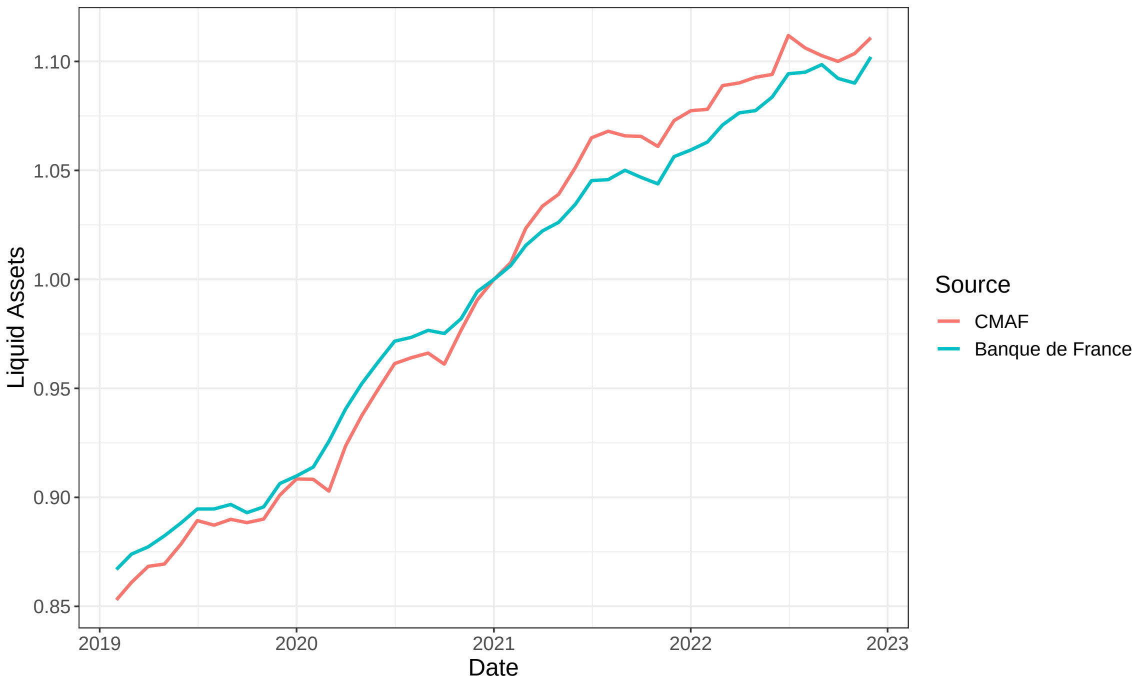

Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*; *Banque de France.*

G Robustness checks

Figure G1: Fake tax rebate on September 1st, 2021 (baseline year: 2019)

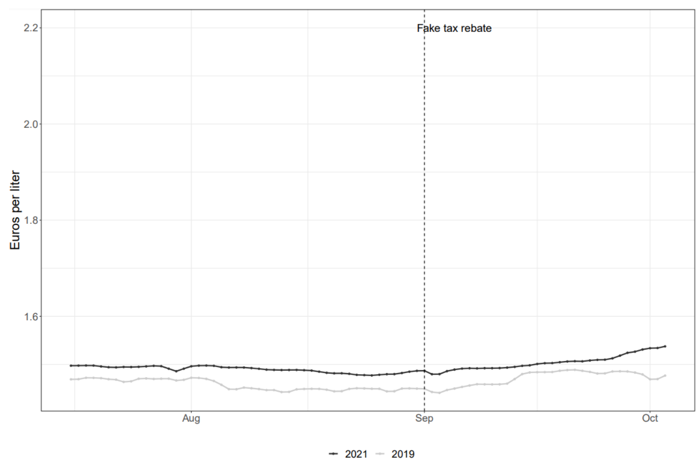

Note. Fuel prices (including taxes). The dashed line corresponds to a fake rebate on September 1st, 2021.
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale.*

Figure G2: Fuel purchases around September 1st, 2021 (baseline year: 2019)

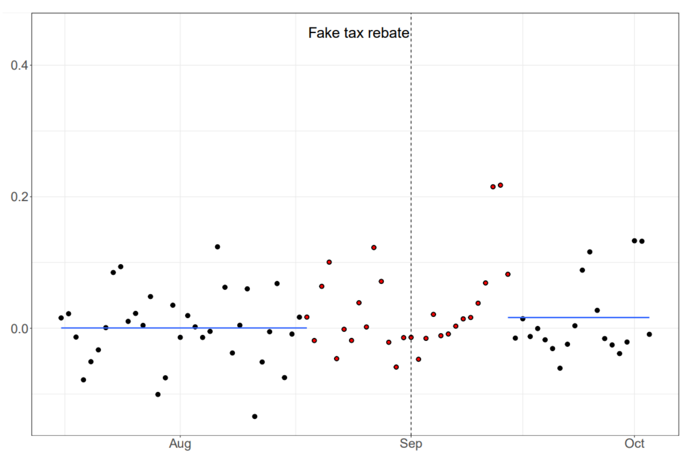

Note. Dots correspond to adjusted daily fuel purchases from July 15th to October 4th, 2021. The adjustment for seasonal variation relies on 2019 as the baseline year. Purchases are normalized so that they sum up to 0 before the anticipation window. Dashed line: (Fake rebate on) September 1st, 2021. Red 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.*

Figure G3: Interpurchase duration (by group of fuel spending, in days)

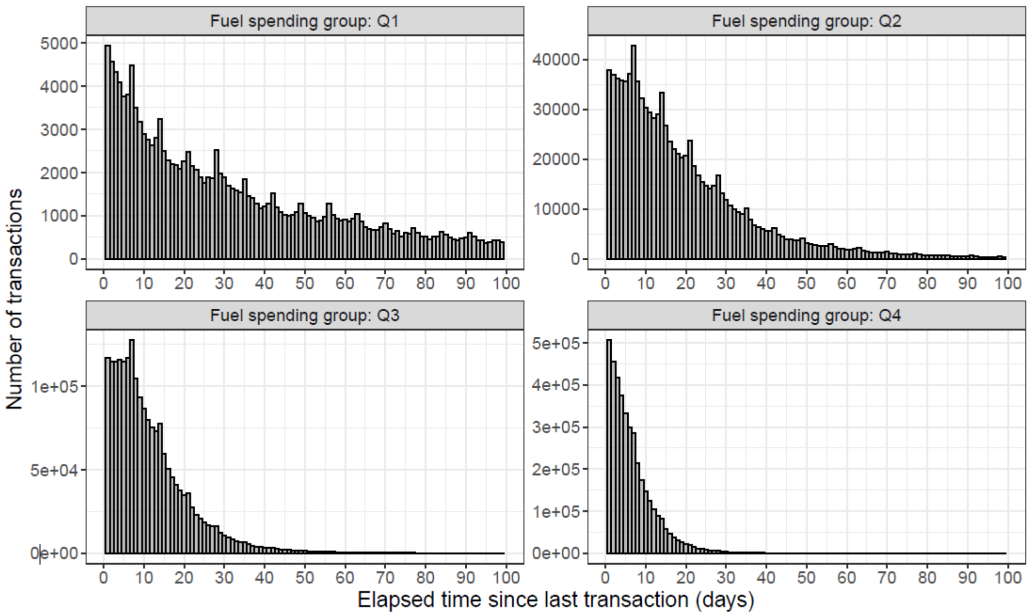

Note. Elapsed time between two transactions from September 2021 to February 2023.
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Table G1: 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*.

Table G2: 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 (11) 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.*

Table G3: Robustness checks with different fuel prices

	I	II	III	IV	V	VI
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 (11) 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 G4: 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 G5: 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 G6: 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)
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 (13) 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 G7: 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 (11) 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.*

H Supplementary figures and tables

Figure H1: Diesel and gasoline prices

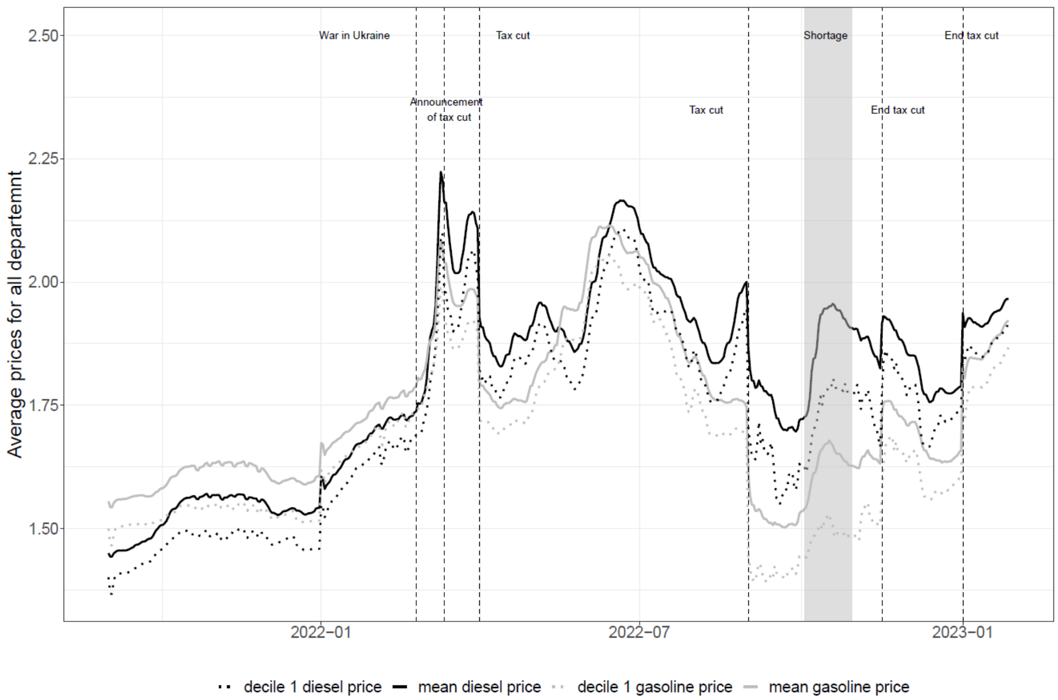

Note. In each *département*, we consider mean diesel price over all stations, mean SP95-E10 price, and the first decile. In the graph we average these quantities at the national level. All prices include taxes. Dashed lines correspond to the invasion of Ukraine and policy interventions. The policy intervention 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 H2: Monthly aggregation of fuel prices and purchases

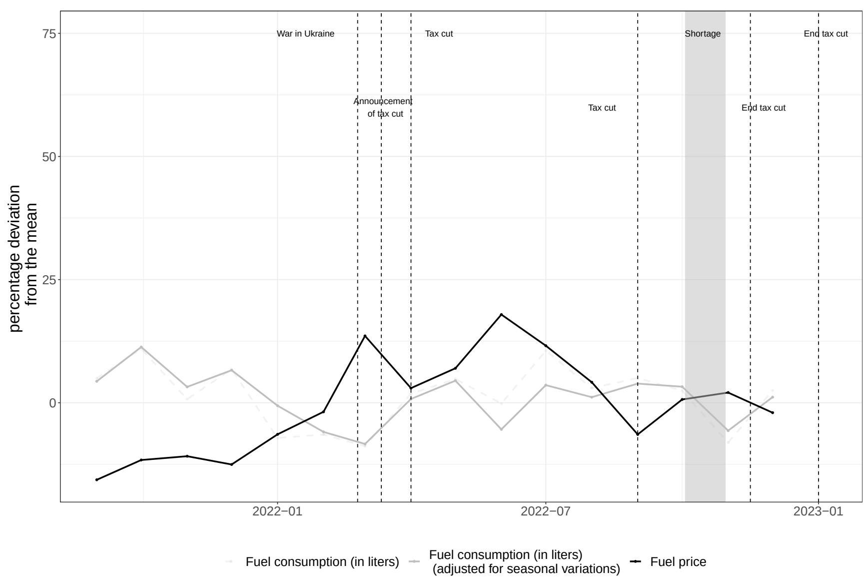

Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Figure H3: Intensive vs extensive margins (quantity per transaction vs # of transactions)

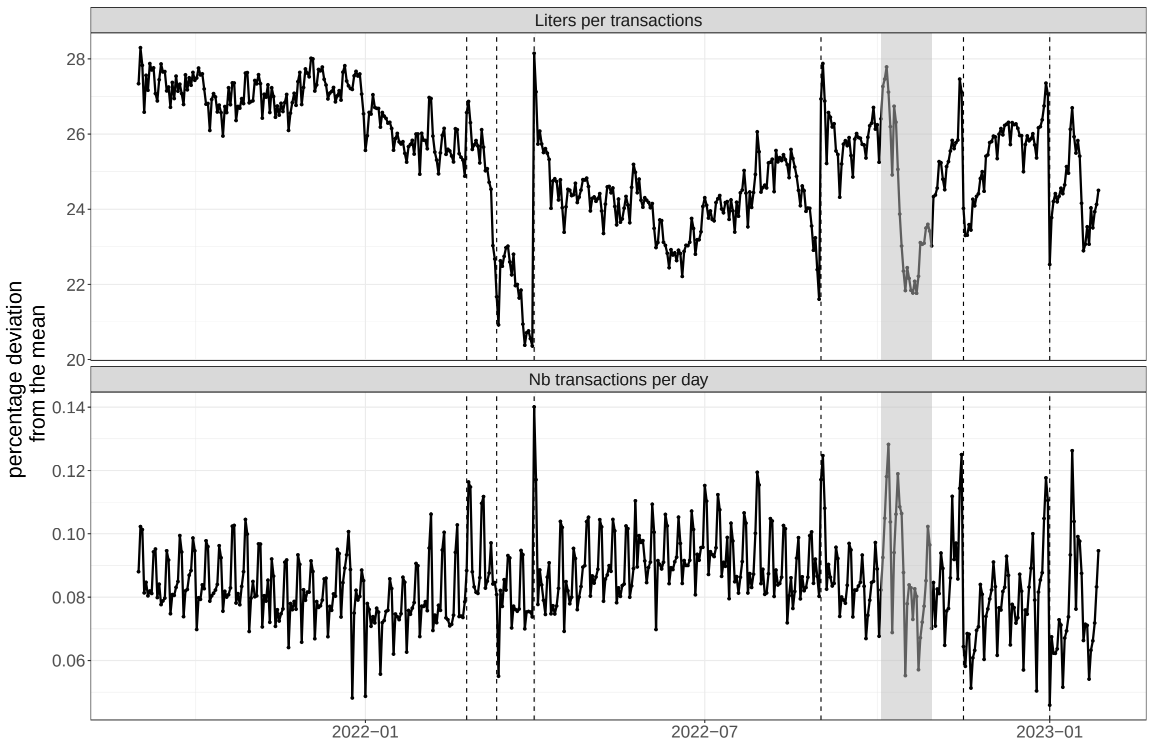

Sources. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale.*

Figure H4: Fuel prices from January to April 2022

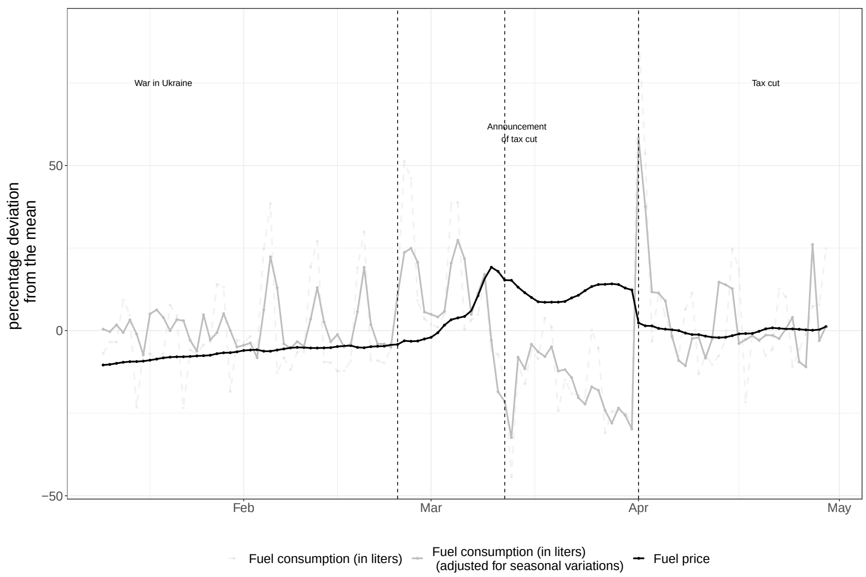

Note. Fuel prices (including taxes) and purchases (in liters). Purchases adjusted for seasonal variations thanks to GIE-CB data from January 8th 2022 to April 30th 2022. The first dashed line corresponds to the invasion of Ukraine, the second dashed line refers to the announcement of the first policy 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 H5: Estimated anticipation parameters (γ coefficients)

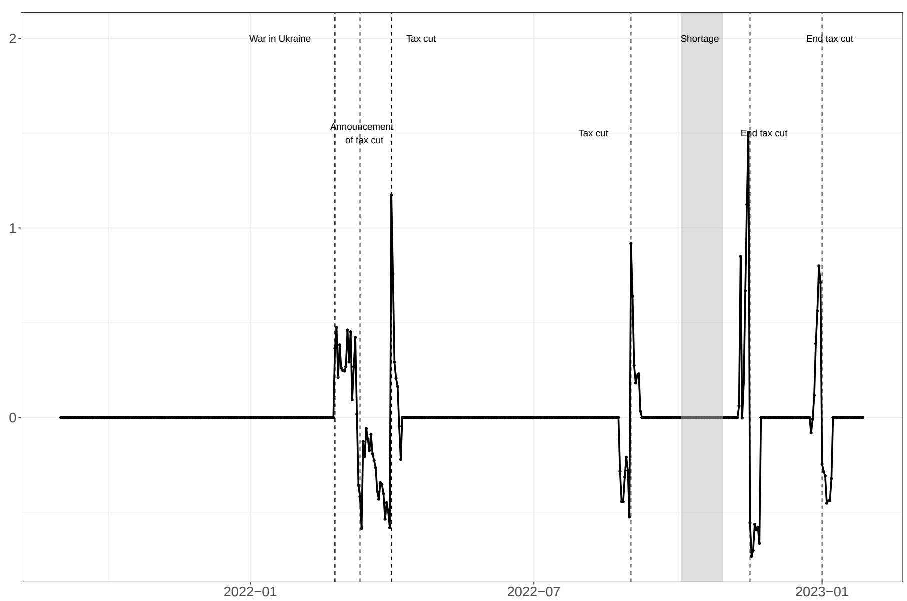

Note. The solid line corresponds to the γ coefficients of equation (13) estimated from September to February. These coefficients capture purchases due to anticipatory behavior, in liters; in each anticipation window, the coefficients sum up to 0. Dashed lines correspond to the invasion of Ukraine and 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 H6: Predicted vs actual demand for fuel

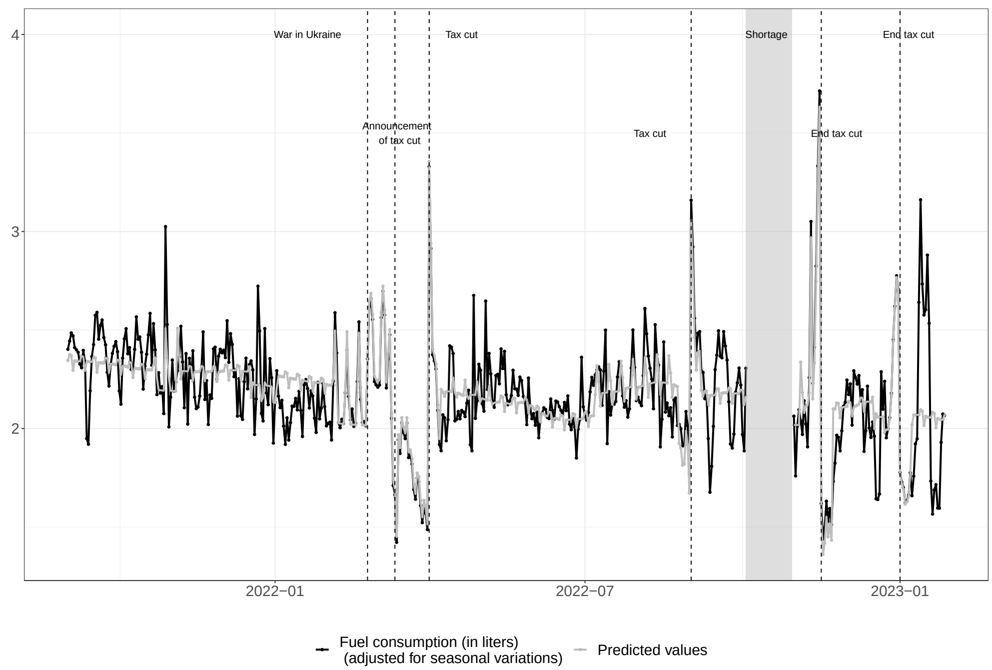

Note. The black line corresponds to observed fuel purchases (in liters) over the period. The grey line corresponds to predicted fuel purchases based on equation (13). Dashed lines correspond to the invasion of Ukraine and policy interventions. The policy intervention of April 1st amounts to a €0.18 per 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 H7: Heterogeneity of the price elasticity with respect to income

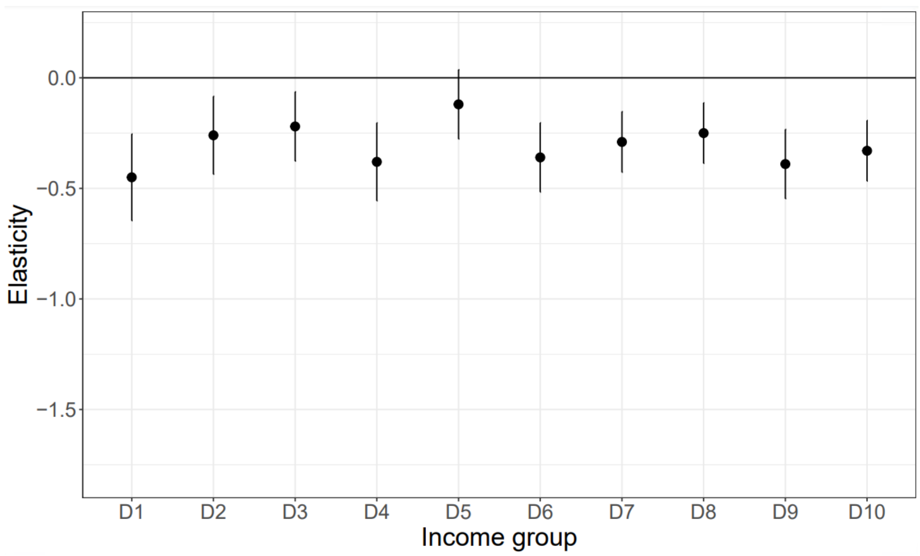

Note. Estimated price elasticity using equation (11). Estimation period: from July 15th to October 4th. Corresponding subsamples depend on households' characteristics (observed before the intervention from January to June, for continuous variables, and in June for discrete variables). Weighted regression 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 H8: Heterogeneity of the price elasticity with respect to total card spending

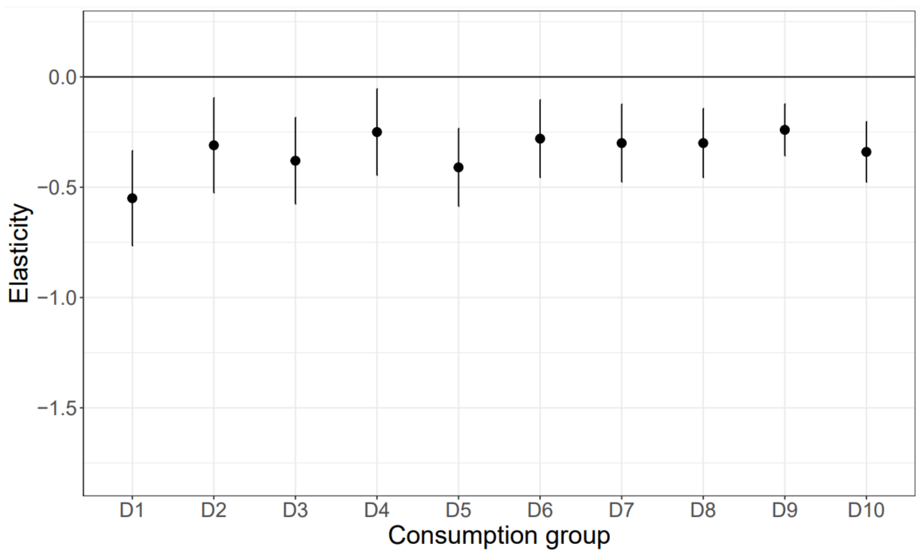

Note. Estimated price elasticity using equation (11). Estimation period: from July 15th to October 4th. Corresponding subsamples depend on households' characteristics (observed before the intervention from January to June, for continuous variables, and in June for discrete variables). Weighted regression 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 H9: Heterogeneity of the price elasticity with respect to occupation

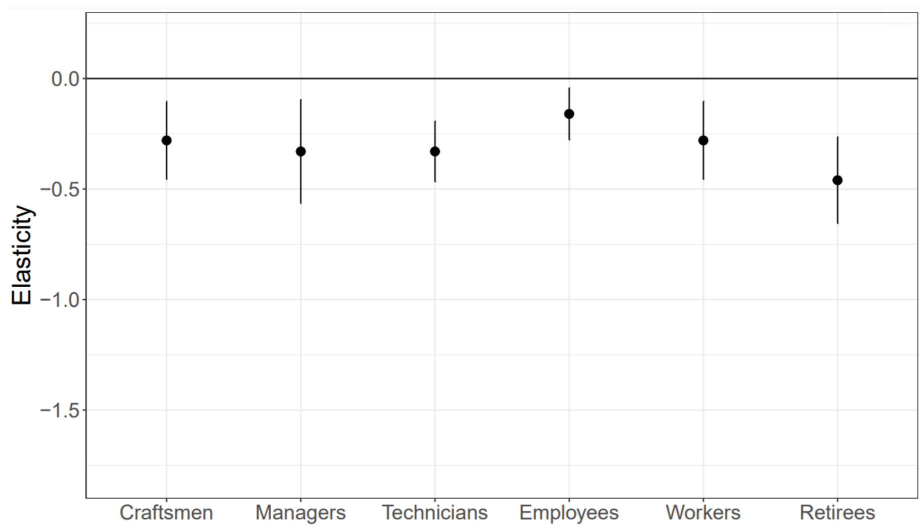

Note. Estimated price elasticity using equation (11). Estimation period: from July 15th to October 4th. Corresponding subsamples depend on households' characteristics (observed before the intervention from January to June, for continuous variables, and in June for discrete variables). Weighted regression 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 H10: Distributional effects of rebates (by fuel spending and income)

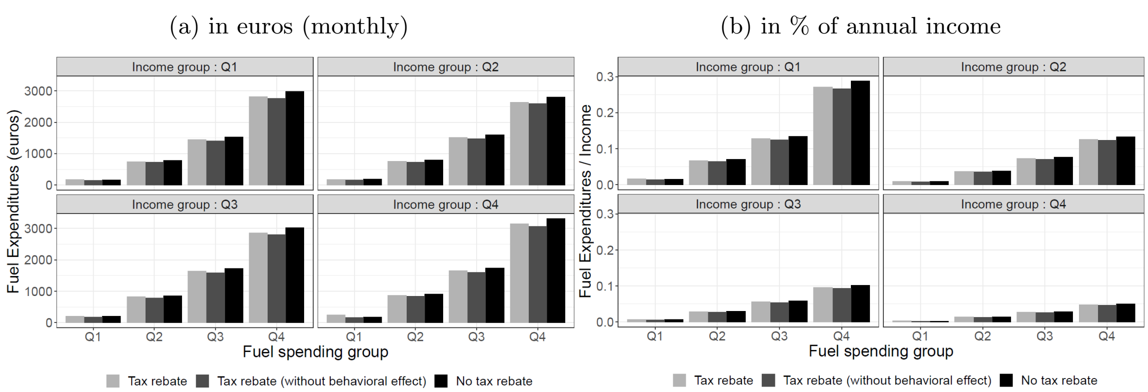

Note. Estimation period: from January 8th 2022 to January 8th 2023.
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale.*

Figure H11: Distributional effects of rebates (by income and location)

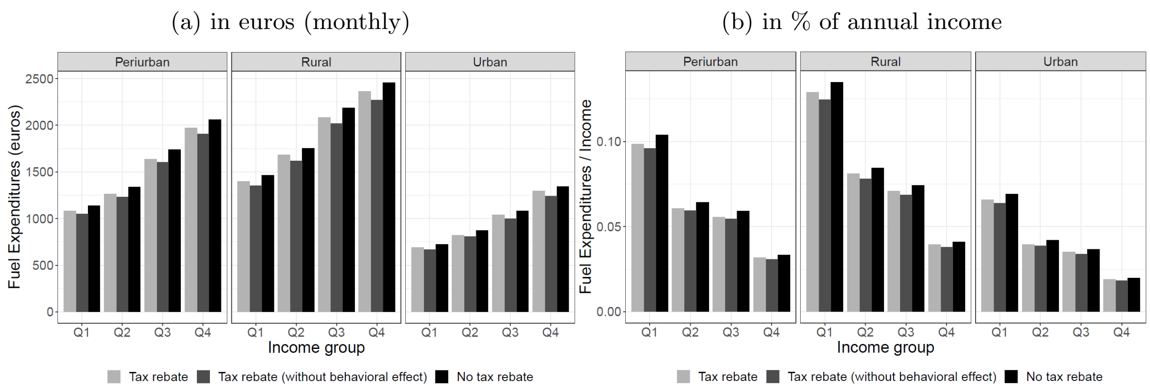

Note. Estimation period: from January 8th 2022 to January 8th 2023.

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

Figure H12: Optimal policy instrument (Objective function: Alleviating excessive relative losses)

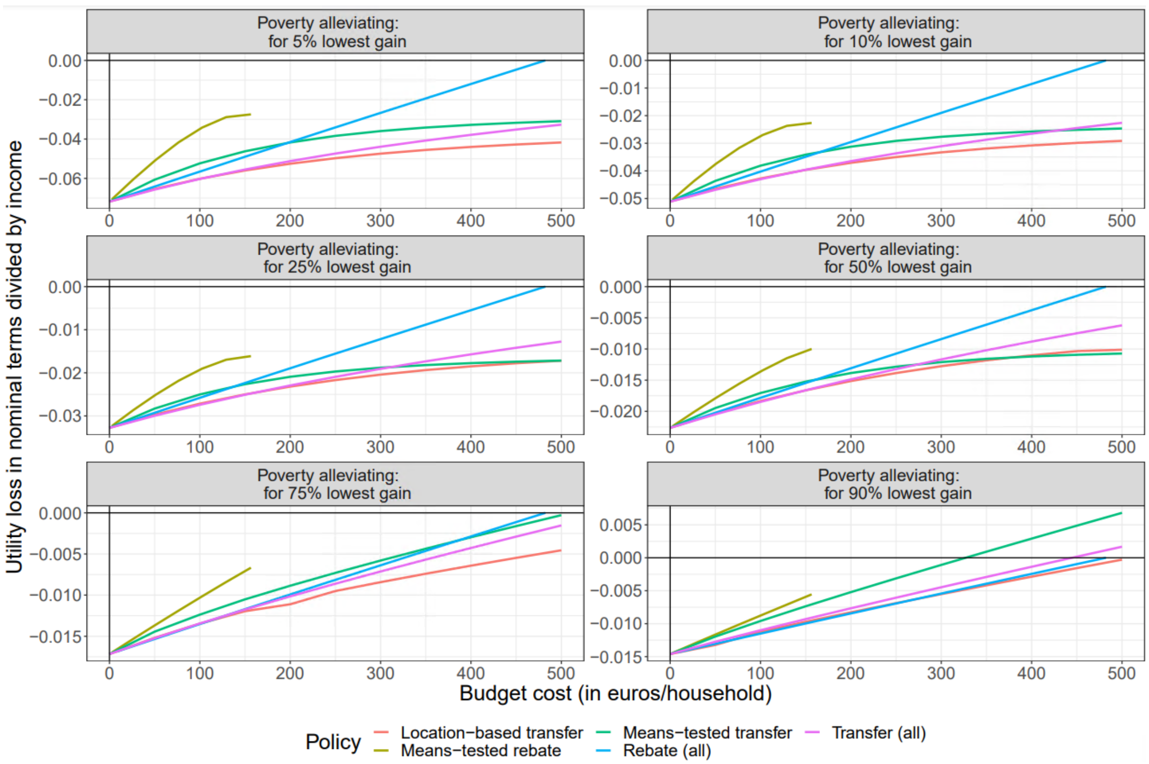

Note. For any given level of public funding (x-axis), we compare (conditional or unconditional) rebates and transfers as policy tools designed to compensate households for the observed price increase from 2021 to 2022 (+€0.6 per liter, i.e. from €1.5 to €2.1).
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale.*

CRESTCENTER FOR RESEARCH
IN ECONOMICS AND STATISTICS

CRESTCENTER FOR RESEARCH
IN ECONOMICS AND STATISTICS

CREST
Center for Research in Economics and Statistics
UMR 9194

5 Avenue Henry Le Chatelier
TSA 96642
91764 Palaiseau Cedex
FRANCE

Phone: +33 (0)1 70 26 67 00
Email: info@crest.science
https://crest.science/

The Center for Research in Economics and Statistics (CREST) is a leading French scientific institution for advanced research on quantitative methods applied to the social sciences.

CREST is a joint interdisciplinary unit of research and faculty members of CNRS, ENSAE Paris, ENSAI and the Economics Department of Ecole Polytechnique. Its activities are located physically in the ENSAE Paris building on the Palaiseau campus of Institut Polytechnique de Paris and secondarily on the Ker-Lann campus of ENSAI Rennes.