How does fuel demand respond to anticipated price changes? Quasi-experimental evidence based on high-frequency data

Odran Bonnet[1] Etienne Fize[2] Tristan Loisel[3] Lionel Wilner[3]

[1] Insee

[2] CAE-IPP

[3] Insee-Crest

EEA-ESEM - August 2023

Motivation

Background and data

Empirical analysis

- Tax cut on September 1st
- External validity: September 2021-September 2022
- Heterogeneity
- Policy evaluation

Outline

Why do we care?

Environmental concerns

- Worldwide policy objective: reducing greenhouse gas (GHG) emissions
- How efficient is carbon taxation to achieve that goal ? The more elastic the demand, the easier to reduce greenhouse gas (GHG) emissions thanks to taxes.

Distributional concerns

- The budget share devoted to fuel consumption decreases with income: a fuel price increase harms more the low income
- How households can adjust fuel consumption when price increases ? Can they smooth the impact on their budget ?
- Highly sensitive... remember the yellow vests movement in 2019 in France

Take-away from the public debate: end-of-the-world vs end-of-the-month trade-off

This paper

Quasi-natural experiments (i.e. exogenous shocks on prices)

- Russian invasion: from 02-24-2022, rise in crude oil prices
- Public interventions in 2022 : tax cuts implemented on April 1st (€0.18 per liter) and a supplementary discount on September 1st (€0.12 per liter). Tax cut was partly removed in mid-November and completely removed by the end of the year.

High-frequency data

- bank account and transaction data : sample of 300,000 customers

Econometric issues

- anticipation effects induce strategic behavior and important variations of demand in the (very) short run

Definition: elasticity is the percentage change in quantity consecutive to a +1% price change

$$\varepsilon = \frac{\partial \log q}{\partial \log p} = \frac{\partial q/q}{\partial p/p}$$

Our contributions

1. Measuring **a short run price elasticity** $\hat{\varepsilon}$
   - we propose a simple method to disentangle direct price effects from indirect anticipation effects
   - Price elasticity lies between -0.39 and -0.23 depending on identification strategies
2. Investigating **the heterogeneity of the price elasticity** along various observable dimensions (income, location, past fuel spending, type of household, age)
   - The average elasticity varies little with income and location, but exhibits sizeable dispersion with respect to fuel spending.
3. Performing a **policy evaluation** exercise, i.e. assessing financial, distributive and environmental impacts of the tax policy

Literature

Huge empirical literature devoted to the estimation of fuel demand's price elasticity: see, e.g., the survey by Dahl and Sterner (1991), Davis and Kilian (2011), etc.

Two main concerns:

- endogeneity (simultaneity). Remedy to upward-biased OLS: IV
  - but tax-based instruments may overestimate ε due to tax aversion
  - same possible bias for persistent price changes
- anticipatory behaviour: downward bias which invalidates both OLS and IV
  - strategic delay of purchases wrt anticipated price reductions
  - more generally, (very) short-term intertemporal substitution

A selection of (rather recent) papers:

- France: Clerc and Marcus (2009): $\hat{\varepsilon} \approx -0.7$ (AIDS), -0.2 (time series); Calvet and Marical (2011): $\hat{\varepsilon} \in [-0.35, -0.25]$ with *Budget des Familles* (AIDS); Douenne (2020): $\hat{\varepsilon} \approx -0.45$ (AIDS); Bureau (2011): -0.22.
- anticipations: Coglianese et al. (2017) introduce leads and lags on monthly-level data
- high-frequency data:
  - Levin et al. (2017): document the aggregation bias; no exogenous price variation; $\hat{\varepsilon} = -0.37$ on US data
  - Knittel and Tanaka (2021): disentangle extensive from intensive margins; no policy-driven price change; $\hat{\varepsilon} \in [-0.35, -0.27]$ on data from Japan
  - Gelman et al. (2022): large, unexpected oil price shocks viewed as an exogenous income shock on other spending than fuel; $\hat{\varepsilon} = -0.2$ (USA)

Short-run vs long-run elasticities (substitution with other transportations)

Outline

Motivation

Background and data

Empirical analysis
- Tax cut on September 1st
- External validity: September 2021-September 2022
- Heterogeneity
- Policy evaluation

High-frequency bank account data

Banking data:

- de-identified panel of 300,000 households who bank at CM from February 2019 to now with
  - transactions: card payments, bank transfers, paper checks, cash withdrawals
  - monthly balance of deposit and savings accounts
  - => spending, income, and savings can be tracked on a near real-time basis
- sociodemographics: age, sex, size of household, proxy for location (*département*), type of location (rural, urban, semiurban), occupation, etc.

Crucial for us: card spending is categorized according to the company that receives the payment

- the 4-digit *Merchant Category Code* (MCC) classification enables us to spot purchase at gas stations
- N.B. MCC available since July 2019 only

Key variables

Fuel spending:

- cards spending with MCC code corresponding to gas station (5541 and 5542).

Prices:

- Governmental data: timestamped and geolocated fuel prices provided at the gas station level.
- Different types of fuel: diesel and gasoline...

Measurement error on transaction prices:

- unknown location of purchase and unknown fuel type (diesel or gasoline)
  - solution: we build a fuel price index depending on the location (*département*) and with weights for diesel and gasoline which depend on the household's characteristics (age group, income group, 2019 fuel spending, type of location)
  - results are not sensitive to the construction of this price index (fuel prices are highly correlated over the period)

Fuel quantity purchased in liters:

- Fuel spending is divided by the fuel price index to get fuel quantity in liters.

Working sample

Initial sample:

- 300,000 clients randomly drawn
- We restrict our attention to active consumers

Working sample:

- 180,000 households calibrated to reproduce exactly known population totals for auxiliary variables (age, sex and *département*)

October 2022 is removed from our sample because strikes occured and led to shortages

Representativeness

Potential concern with bank account data: selection bias (customers from a particular bank only)

- Distribution of income and the distribution of household spending, by income, matches closely the one issued from the representative consumption survey
- same for fuel spending
- the evolution of fuel spending looks quite identical (0.99 correlation) to the one issued from the comprehensive *Groupement d'Intérêt Bancaire-Carte Bleue* (GIE-CB) dataset GIE-CB data

This empirical evidence alleviates legitimate concerns about selection bias

Figure 1: Evolution of fuel prices and purchases (September 2021-January 2023)

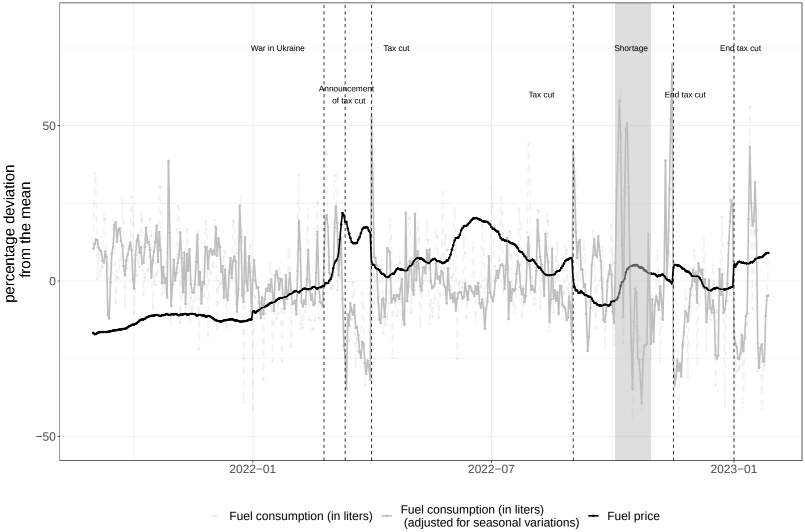

Outline

Motivation

Background and data

Empirical analysis

- Tax cut on September 1st
- External validity: September 2021-September 2022
- Heterogeneity
- Policy evaluation

Identification strategy

Our setting:

- public intervention (2nd/extra €0.12 per liter discount at the pump): quasi-natural experiment
- treatment group: consumers purchasing in year $y = 2022$
- comparison group: consumers purchasing in year $y = 2021$

Figure 2: Fuel prices around September, 1st (in 2021 and 2022)

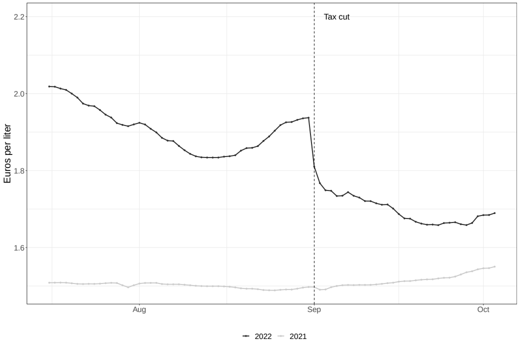

Figure 3: Fuel purchases' response to the September 1st anticipated tax cut : difference between fuel consumption in liters in 2022 and fuel consumption in 2021 (mean difference is normalized to 0 in the pre-period)

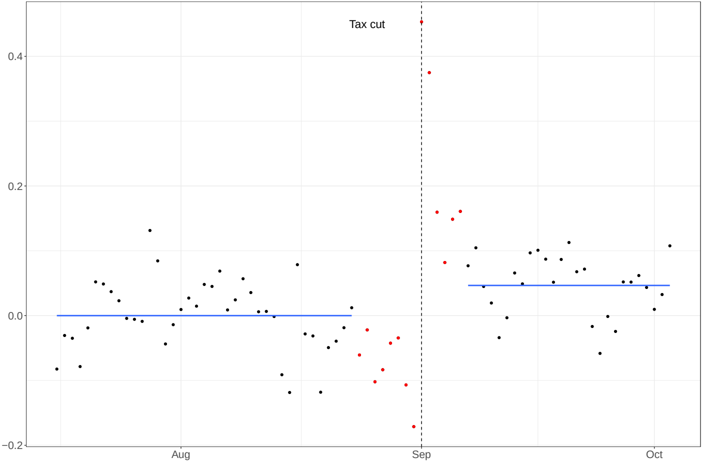

Disentangle price effect and anticipation effect: some intuition

- When price decreases, two distinct effects:
  1. **Price effect**:
     - adjustment of consumption according to prices (short term and medium/long term response)
  2. **Anticipation effect** (red dots period):
     - Strategic delay of purchases to benefit from lower prices (only short term response) : households empty their tank (lower their inventory) before the tax cut and fill their tank (restore their inventory) right after the implementation of the tax cut
- **Implications:**
  - First effect affects aggregate level of consumption over the period
  - Second effect does not impact aggregate level of consumption over the period but just change the timing of purchases

A stylized inventory model of fuel stockpiling behavior rationalizes these ideas (same principles apply for an increase in prices) Theoretical model

Econometric specification

First, for computational reasons, we aggregate our data into 10, 842 cells of individuals who share the same observable characteristics.

Our dependent variable, q_{cty}, is the fuel quantity, in liters, purchased by individuals belonging to cell c on day t of year y, adjusted for seasonal variations.

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

where p_{ct} are prices, α_{cy} is a cell-year fixed effect, μ_t is a day-of-the-year fixed-effect. t_2 corresponds to the beginning of the second excise tax cut, namely September 1st.

N.B. Characteristics for the cells are : département, income group (in four intervals), age group (less than 30, 30-60, more than 60), type of location (rural, urban, or semiurban), and 2019 fuel spending category (in four intervals).

Three estimators

$$q_{cty} = \beta p_{cty} + \sum_{h=t_2-\Delta}^{t_2+\Delta} \gamma_{hy} \mathbb{1}_{h=t} \mathbb{1}_{y=2022} + \alpha_{cy} + \mu_t + \eta_{cty}, \qquad (2)$$

Three estimators:

- **naive estimator**: $\forall h\ y\ \gamma_{hy} = 0$
  - Do not take into account anticipatory behavior => overestimate reaction to prices.
- **constrained estimator**: $\sum_{h=t_2-\Delta}^{t_2+\Delta} \gamma_{h,2022} = 0$
  - Anticipation effects compensate over the anticipation window
- **unconstrained estimator**: no constraints on the γ (act as fixed effects)
  - "equivalent" to delete the period of anticipations (the confounding period) for estimation

The bandwidth Δ of the anticipation window around event is set to 7 days (estimated elasticities are stable above this level) Choosing the right bandwidth

Main estimates

Table 1: DinD (September 1st)

	I	II	III	IV	V	VI
price elasticity	-0.45 (0.07)	-0.23 (0.07)	-0.30 (0.06)	-0.76 (0.14)	-0.39 (0.08)	-0.36 (0.07)
IV (Instrument: post- 9/1 dummy)				✓	✓	✓
Anticipation dummies		✓			✓	
Excluding anticipation window			✓			✓
Cell FE	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓
# of cells	10,842	10,842	10,842	10,842	10,842	10,842

Note. Estimation of equation (2) with a sample of 10,842 cells of customers observed from July, 15th to October, 4th. Treatment group: 2022, Comparison group: 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*.

Main estimates

As predicted by the model:

- Naive estimator over estimate the reaction: implied elasticity of -0.45 (anticipation bias)
- Constrained and unconstrained estimator are close: implied elasticity of -0.23 and -0.30
- anticipation bias is ≈ 0.2

Main estimates

We also made IV estimations (using the tax cut as an IV)

- there the sole source of identifying variability is the price reduction on September 1st
- (it also address possible concerns about measurement error)
- constrained and unconstrained estimators point out to a -0.39 price elasticity and -0.36
- naive estimator is -0.76.

$=>$ **Test of the constraint:** we can not reject the model (hypothesis that the sum of anticipations effect average to 0).

External validity

Admittedly, previous identification strategy is local

To assess its external validity, we also rely on the whole observation period from September 2021 to January 2023.

No valid comparison group here: lockdowns occurring the year(s) before + no data on fuel consumption before July 2019

Seasonal adjustment, though, based on 2019 aggregate fuel consumption by cards data (source, GIE-CB).

$$q_{ct} = \beta p_{ct} + \sum_{h=t_0}^{t_1+\Delta} \gamma_h^1 \mathbb{1}_{h=t} + \sum_{k=2}^{4} \sum_{h=t_k-\Delta}^{t_k+\Delta} \gamma_h^k \mathbb{1}_{h=t} + \alpha_c + \mu_t + \eta_{ct}, \quad (3)$$

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

Estimation based on the whole period

Table 2: Estimations based on the whole period (September 2021-February 2023)

	I	II	III	IV	V	VI
price elasticity	-0.44 (0.05)	-0.49 (0.05)	-0.45 (0.07)	-0.32 (0.04)	-0.39 (0.04)	-0.26 (0.03)
Anticipation dummies				✓	✓	✓
Seasonality controls		✓	✓		✓	✓
Linear trend			✓			✓
Cell FE	✓	✓	✓	✓	✓	✓
# of cells	7,000	7,000	7,000	7,000	7,000	7,000

Note. Estimation of equation (3) with a sample of 7,000 cells of customers observed from September 2021 to February 2023. Two-way clustering of standard errors at cell and day levels.
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

Main lessons:

- previous estimates are close to the ones obtained over a longer period (-0.26 and -0.39 depending on the set of controls)

Heterogeneity of price elasticities

First we assess heterogeneity by estimating our september estimation framework on sub-populations:

- Average price elasticity does not vary much with income or location.
- By contrast, a dimension along which that average price elasticity exhibits sizable dispersion is (past) fuel consumption:
  - households with high level of fuel consumption are less elastic
  - it is particularly the case for low liquidity households => low liquid households with high fuel consumption are particularly harmed by a price increase

Second, we use a causal forest's approach to make an agnostic search on the observable dimensions of heterogeneity:

- Elasticity range between -1 for the 20% most price-sensitive individuals and 0 for the 20% most inelastic (the difference being statistically significant at 5%) Sorted Group Average Treatment Effect .
- => on the short run, elasticity is not very high and not very heterogeneous along observable characteristics

Counterfactuals and policy evaluation: financial effects

Method:

- We simulate demand at counterfactual prices (without tax cuts) based on previous estimations.
- We then evaluate the impact of the policy on consumption, in liters, by computing the difference between observed (ex post) and simulated (ex ante) consumption:

Results:

- Overall financial impact of the policy is to reduce fuel spending by about €71 per household, on average, i.e. €8 monthly, which represents about 0.3% of income.
- Omitting to take into account adjustment to prices would result in a €13 monthly gain, on average (65% higher).

Counterfactuals and policy evaluation: distributive effects

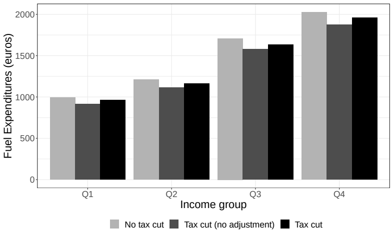

(a) in euros (monthly)

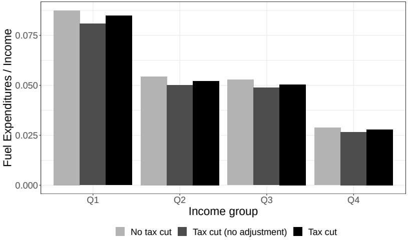

(b) in % of income

Figure 4: Distributive effects of excise tax cuts on fuel

Counterfactuals and policy evaluation: environmental effects

- Last, the impact of the policy on CO_2 emissions has been rather limited (thanks to low elasticity of demand).
- It represents nearly 56 kilograms of CO_2, hence 0.3% of a French annual carbon footprint (20.3 tons in 2021).

Future works on policy evaluation

- We would like to compare the discount at the pump to alternative policies like transfers to all households or to a targeted sub-population.
- Who are the winners and losers when we compare discount and transfers ?
- What is the best policy ? => depends on the social planner objective function

Appendix

Remerciements

Ce travail repose sur l'exploitation de données de comptes bancaires auxquelles le Crédit Mutuel Alliance Fédérale a permis l'accès.

Nous remercions ce partenaire pour nous avoir facilité cet accès technique mais aussi pour sa disponibilité lors des échanges.

Données du Crédit Mutuel Alliance Fédérale

Première banque à adopter le statut 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.

Outline

Theoretical model

Comparisons with other sources

Heterogeneities

How to disentangle price effect and anticipation effects ?

A stylized inventory model of fuel stockpiling behavior
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 agent solves:

$$\max_{(c,i)} \sum_{t=0}^{T} [u(c_t) - C(i_t)] \quad \text{s.t.} \quad \sum_{t=0}^{T} p_t q_t \leq Y \tag{4}$$

$$i_t \leq i_{t-1} + q_t - c_t \tag{5}$$

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

$$C'(i_t) = \lambda(p_{t+1} - p_t) \tag{6}$$

How to disentangle price effect and anticipation effects ?

Solving the model 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:

$$i_t = \lambda \frac{p_{t+1} - p_t}{2\theta}, \tag{7}$$

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

$$c_t = \frac{1 - \lambda p_t}{2\alpha}. \tag{8}$$

The model therefore predicts that observed purchases are given by a specification such that:

$$q_t = q_0 + \beta p_t + \gamma_t \max(\mathbb{1}_{p_{t-1} \neq p_t}, \mathbb{1}_{p_{t+1} \neq p_t}) \tag{9}$$

where anticipation effects $\gamma_t = \lambda \frac{(p_{t+1} - p_t) - (p_t - p_{t-1})}{2\theta}$, which are non-zero as soon as prices fluctuate, alter current consumption.

How to disentangle price effect and anticipation effects ?

Take away:

- when prices are stable, $i_t = 0$ and $q_t = c_t$
- when prices increase/decrease, consumers fill $(q_t > c_t)$ / empty$(q_t < c_t)$ their tank $i_t = \lambda \frac{p_{t+1} - p_t}{2\theta}$ and go back to their usual level of inventory just after: $i_t = 0$
- $=>$ sum of the anticipation effect goes to 0.

return

An illustration with a price reduction

Consider the following 4-period toy model:

- 2 periods at price p
- 2 periods with a price reduction Δp announced in period 2

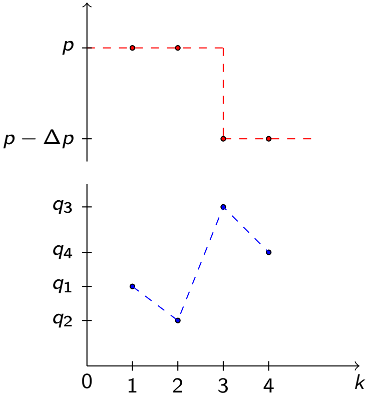

The choice of the bandwidth

Δ is the bandwidth of the anticipation window around event starting at $t_2 =$09-01-2022.

Empirically, $\Delta = 7$:

- economical considerations: timing of purchase can be manipulated up to tank capacity
  - typical interpurchase duration: one week
- Empirical considerations choice of bandwidth:
  - For small values of Δ, the estimation overestimate the reaction to price changes by incorrectly attributing short-run intertemporal substitution motives to price sensitivity
  - Starting at $\Delta = 7$, elasticity stabilizes at -0.24

return

Outline

Theoretical model

Comparisons with other sources

Heterogeneities

Table 3: 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	0.37
Urban areas	0.19

Note. Estimation period: 2021 for transactions (spending, income), January 2021 for assets and socio-demographics. Pecuniary amounts in €. The oldest member of the household is the head of the household.
Source. Sample of households who primarily bank at *Crédit Mutuel Alliance Fédérale*.

return

Figure 5: Distribution of income (transaction data vs. survey data from *ERFS*, Insee)

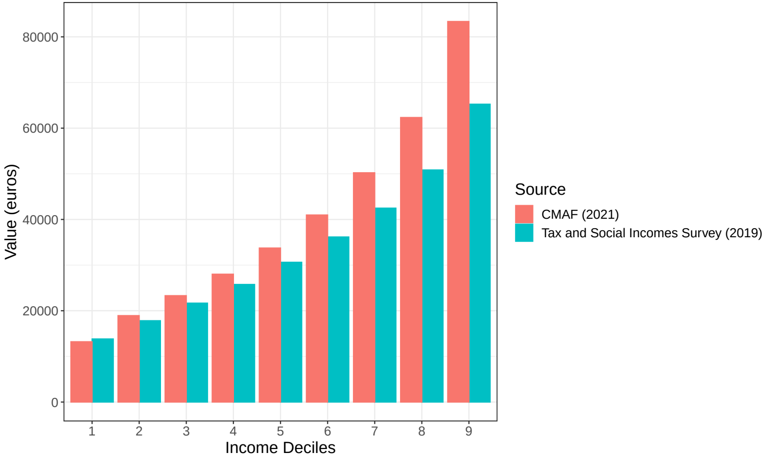

return

Figure 6: Household expenditures by income (transaction data vs. survey data from *Budget des Familles*, Insee)

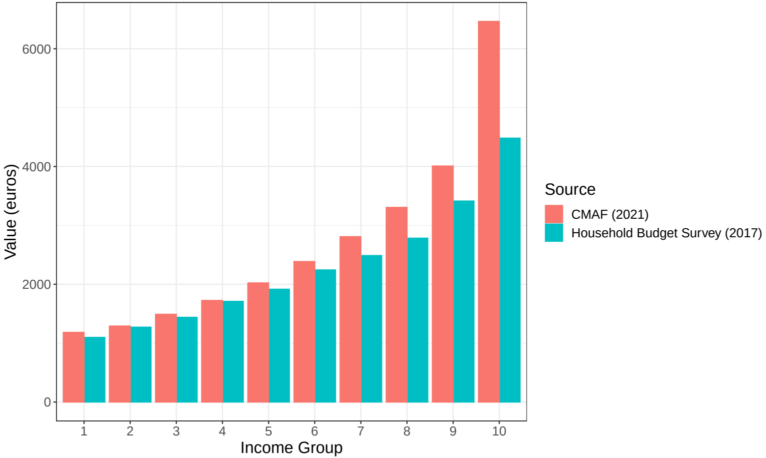

return

Figure 7: Distribution of fuel spending, by income (transaction data vs. survey data from *Budget des Familles*, Insee)

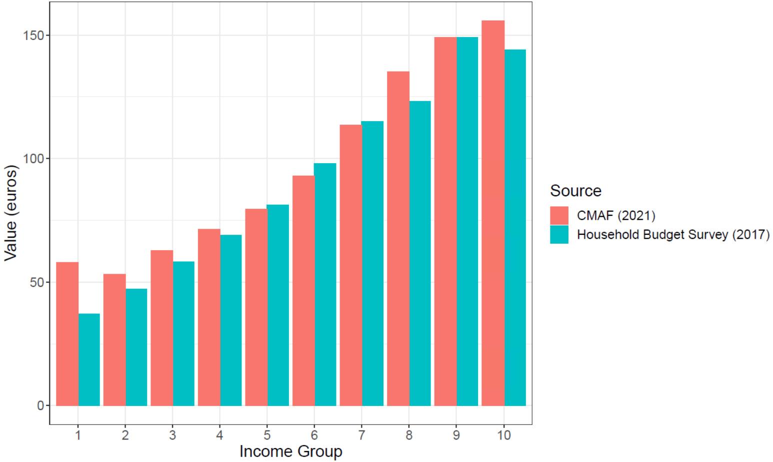

return

Figure 8: Evolution of fuel spending (transaction data vs. aggregate data from the French interbank network)

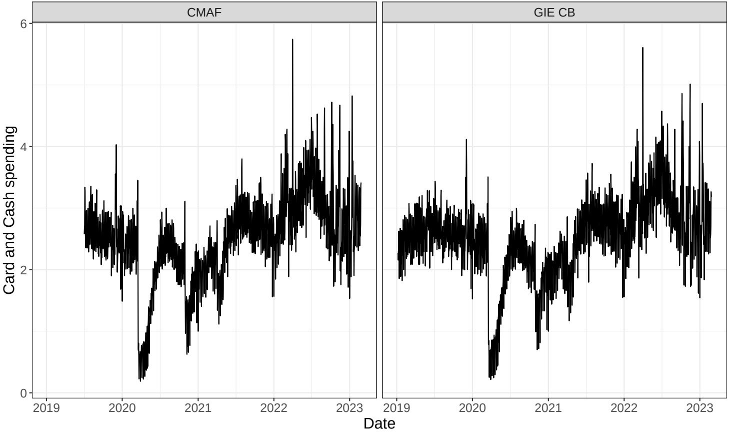

Correlation between both time series: 0.99

return

A statistical criterion for Δ

Figure 9: The BIC as a function of the anticipation bandwidth Δ

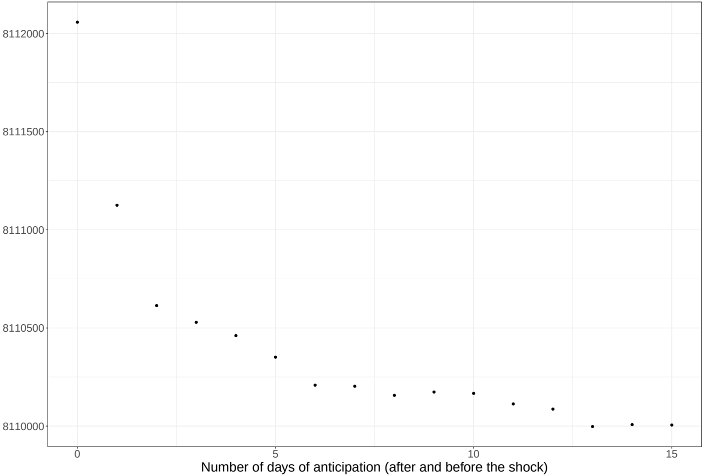

return

Distribution of interpurchase durations

Figure 10: Interpurchase duration

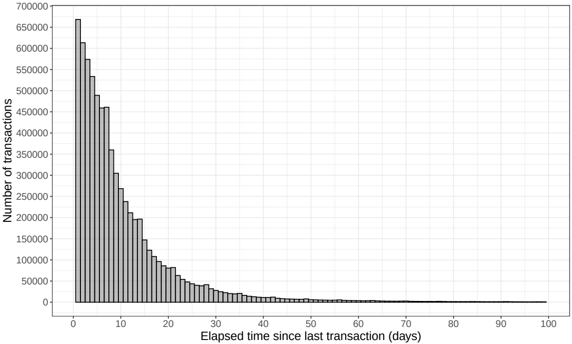

return

Price elasticity as a function of Δ

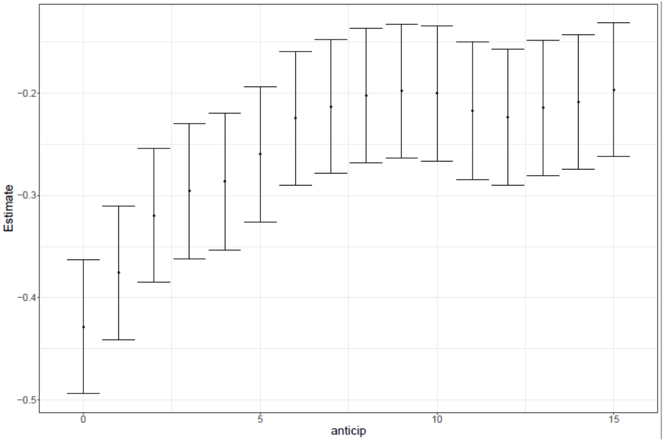

return

Outline

Figure 11: Heterogeneity along income

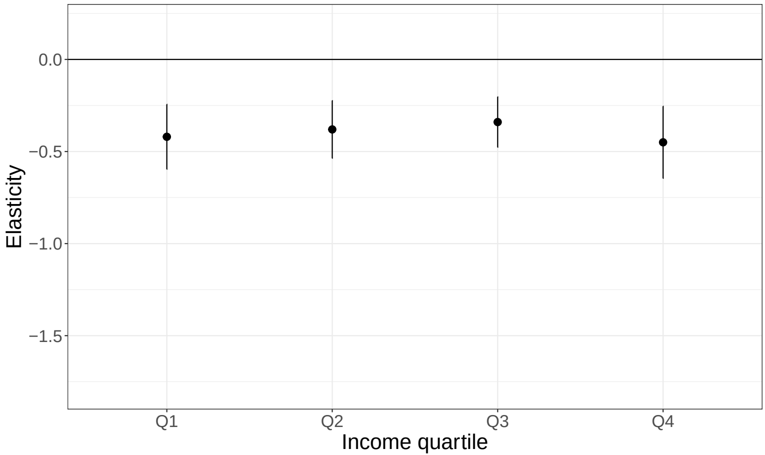

return

Figure 12: Heterogeneity along localisation

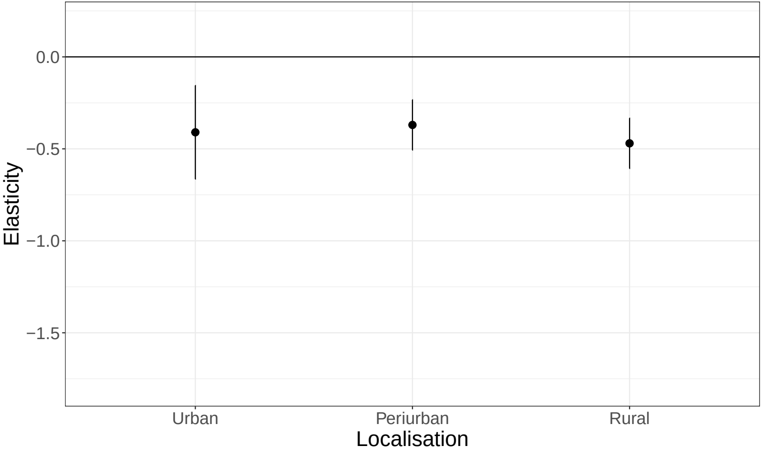

return

Figure 13: Heterogeneity along past fuel consumption

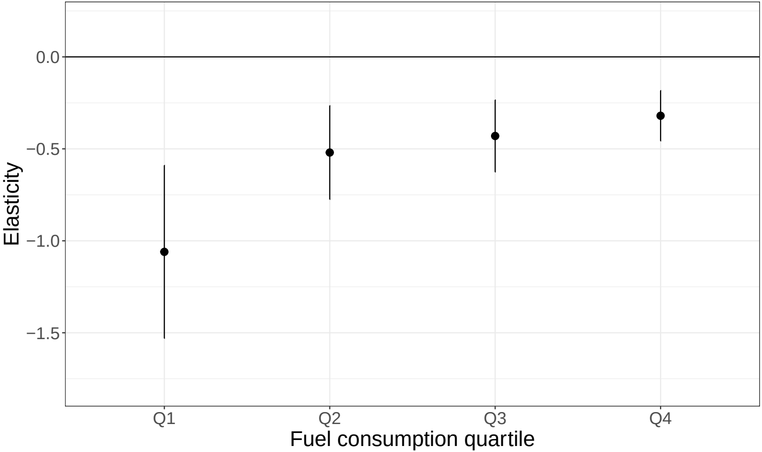

return

Figure 14: Heterogeneity along liquidity assets and past fuel consumption

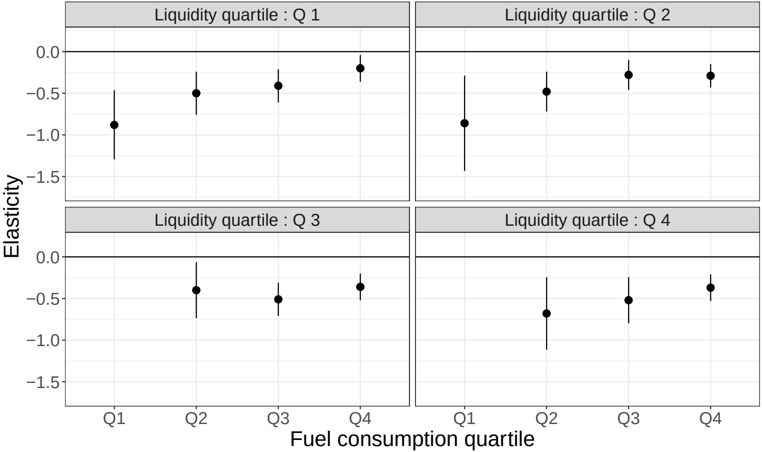

return

Figure 15: Sorted Group Average Treatment Effects

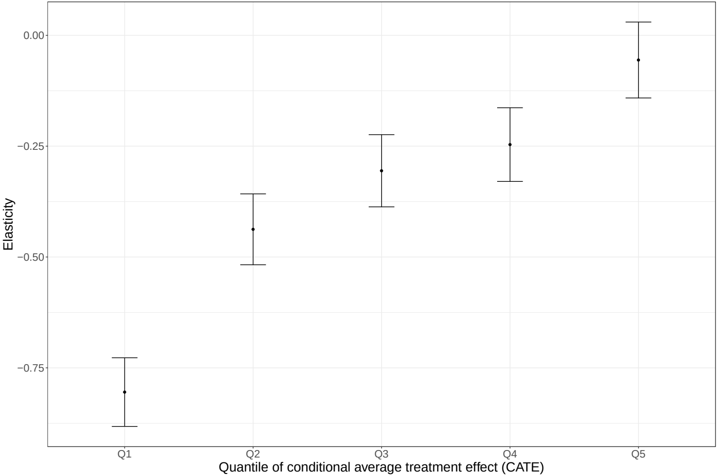

return