

# Tax Salience and Household Inflation Expectations: Evidence from European Energy VAT Cuts

May 4, 2026

## Abstract

Do households incorporate indirect tax changes into their inflation expectations? I study this question by exploiting the staggered adoption of energy VAT cuts across five European countries during the 2021–2023 energy crisis, using France—which implemented direct subsidies rather than VAT cuts—as a never-treated control. Drawing on the ECB Consumer Expectations Survey, a large household panel with probabilistic inflation beliefs, I find three main results. First, energy VAT cuts significantly reduce perceived past inflation ( $\beta = 0.763$ ,  $t = 13.53$ ), indicating that households notice changes in indirect taxes through their effect on consumer prices. Second, the impact on forward-looking inflation expectations is modest and statistically weak ( $\beta = 0.054$ ,  $t = 1.70$ ), suggesting that households do not fully extrapolate tax-driven price changes into expectations about future inflation. Third, financial literacy strongly moderates the expectations response: each additional correct answer on the Lusardi-Mitchell financial literacy scale increases the treatment effect on expectations by 0.146 percentage points ( $t = 4.79$ ). Financially literate households both perceive the VAT cut and adjust their forward-looking expectations, while less literate households perceive the price change but fail to update expectations accordingly. These findings contribute to the literature on tax salience by documenting that households process indirect tax changes asymmetrically—perceiving price-level effects but not translating them into revised expectations—and that financial literacy is a key determinant of the depth of tax information processing.

*JEL Classification:* D84, E31, H25, H31

*Keywords:* Tax salience, VAT, inflation expectations, financial literacy, household surveys

# 1 Introduction

A central question in public economics is the extent to which economic agents perceive and respond to changes in taxation. The tax salience literature, beginning with [Chetty et al. \(2009\)](#) and [Finkelstein \(2009\)](#), has documented that the behavioral response to a tax depends critically on how visible that tax is to agents. When taxes are less salient—embedded in sticker prices rather than added at the register, or collected as payroll deductions rather than as lump-sum payments—agents respond less than standard models predict.

This paper extends the tax salience framework to a new domain: household inflation expectations. The setting is the 2021–2023 European energy crisis, during which five euro area countries (Spain, Italy, Belgium, the Netherlands, and Germany) implemented large cuts to the value-added tax (VAT) on energy products, while France chose instead to cap energy prices through direct subsidies. These staggered policy changes create a natural experiment in which the timing and magnitude of tax changes vary across countries while the underlying energy price shock is common.

The contribution of this paper is threefold. First, I document a new channel through which tax policy affects household beliefs. While the existing literature has studied how taxes affect purchasing behavior ([Chetty et al., 2009](#); [Baker et al., 2021](#)), I show that indirect tax changes also shape household perceptions and expectations of aggregate inflation. Energy VAT cuts reduce perceived past inflation by 0.76 percentage points ( $t = 13.53$ ), a large and precisely estimated effect that confirms households notice the price-level consequences of indirect tax changes.

Second, I show that the depth of information processing varies across the belief distribution. Despite perceiving lower prices, households do not fully translate the VAT cut into revised forward-looking expectations. The effect on expected inflation is positive but modest ( $\beta = 0.054$ ,  $t = 1.70$ ). This asymmetry between perceived and expected inflation is consistent with models of limited attention ([Taubinsky and Rees-Jones, 2018](#)): households observe the tax-induced price change but do not reason through its implications for future inflation.

Third, I identify financial literacy as a key determinant of how deeply households process tax information. Households scoring two or more on the three-item Lusardi-Mitchell financial literacy scale show a significantly larger expectations response to the VAT cut ( $\beta = 0.235$ ,  $t = 3.40$ ). This result suggests that financial literacy captures not only numeracy but also the ability to translate price observations into coherent macroeconomic expectations—an important dimension of tax salience that has received little attention.

This paper speaks to several strands of the accounting and economics literatures. Within accounting, [Hanlon \(2025\)](#) calls for research on consumption taxes, survey-based methods, and cross-country tax settings—this paper addresses all three. The tax salience

framework (Chetty et al., 2009) has been applied primarily to sales taxes and excise duties in the United States; I extend it to VAT, the dominant consumption tax globally, and to a cross-country panel setting. The finding that financial literacy moderates tax salience connects to the disclosure processing literature (Blankespoor et al., 2020), which argues that the real effects of information depend on agents’ capacity to process it. In the macroeconomics literature, the paper contributes to the growing body of work on the formation of household inflation expectations (Coibion and Gorodnichenko, 2015; Malmendier and Nagel, 2016; D’Acunto et al., 2021) by identifying a specific policy instrument—indirect taxation—that shapes expectations through its effect on consumer prices.

The remainder of the paper is organized as follows. Section 2 describes the institutional setting and develops the hypotheses. Section 3 presents the data. Section 4 outlines the empirical design. Section 5 reports the main results. Section 6 presents heterogeneity and robustness analyses. Section 7 concludes.

## 2 Institutional Background and Hypothesis Development

### 2.1 Energy VAT Cuts in Europe, 2021–2023

Beginning in the second half of 2021, a sharp rise in energy prices—driven by post-pandemic demand recovery and supply disruptions related to the Russia-Ukraine conflict—prompted European governments to adopt a range of fiscal measures to shield households from rising energy costs. A prominent policy tool was the temporary reduction of VAT rates on energy products.

Table 1 summarizes the timing and magnitude of energy VAT cuts across the six original countries in the ECB Consumer Expectations Survey. Spain was the first to act, reducing electricity VAT from 21% to 10% in June 2021 (subsequently to 5% in January 2022). Italy followed in October 2021, cutting gas VAT from 22% to 5%. Belgium reduced VAT on electricity and gas from 21% to 6% in March 2022. The Netherlands implemented an equivalent energy tax reduction in July 2022, and Germany cut gas VAT from 19% to 7% in October 2022.

France adopted a fundamentally different approach. Rather than reducing indirect taxes, the French government implemented the *bouclier tarifaire* (tariff shield)—a direct cap on retail energy prices, funded through transfers to energy suppliers. This policy difference is central to the identification strategy: France experienced the same underlying energy price shock as its neighbors but did not change the VAT rate, making it a natural control group.

Table 1: Energy VAT Policy Changes, 2021–2023

Country	Policy change	Date	VAT rate	Cut (pp)
Spain (ES)	Electricity VAT reduction	Jun 2021	21% → 5%	16
Italy (IT)	Gas VAT reduction	Oct 2021	22% → 5%	17
Belgium (BE)	Energy VAT reduction	Mar 2022	21% → 6%	15
Netherlands (NL)	Energy tax reduction	Jul 2022	Equivalent	9
Germany (DE)	Gas VAT reduction	Oct 2022	19% → 7%	12
France (FR)	Direct price cap	—	No VAT change	0

## 2.2 Hypothesis Development

The tax salience literature provides a framework for predicting how VAT cuts affect household beliefs. [Chetty et al. \(2009\)](#) demonstrate that consumers respond less to taxes that are not directly visible in posted prices. VAT, however, is an included tax in Europe—retail prices already incorporate the VAT amount—so any VAT reduction is mechanically reflected in the sticker price. This suggests that energy VAT cuts should be relatively salient: households should perceive lower energy prices, even if they do not attribute the change to tax policy.

This leads to the first hypothesis:

**H1:** *Energy VAT cuts reduce households’ perceived past inflation.*

Whether households extrapolate this price observation into revised expectations about future inflation is less clear. Forward-looking expectations require reasoning about the aggregate price level, the persistence of the tax change, and the broader macroeconomic environment. Models of rational inattention ([Coibion and Gorodnichenko, 2015](#)) and experience-based learning ([Malmendier and Nagel, 2016](#)) suggest that households may be slow to incorporate policy-driven price changes into expectations, particularly when the tax change is framed as temporary.

**H2:** *Energy VAT cuts have a limited effect on forward-looking inflation expectations relative to their effect on perceived inflation.*

Financial literacy is a natural candidate for moderating the depth of information processing. Literate households are better equipped to interpret price changes, distinguish between relative and aggregate price movements, and update beliefs about future inflation accordingly ([Lusardi and Mitchell, 2014](#); [Jappelli and Padula, 2013](#)). The disclosure processing literature in accounting ([Blankespoor et al., 2020](#)) similarly argues that the effectiveness of public information depends on agents’ processing capacity.

**H3:** *Financially literate households exhibit a larger expectations response to energy VAT cuts.*

## 3 Data

### 3.1 The ECB Consumer Expectations Survey

The primary data source is the Consumer Expectations Survey (CES) conducted by the European Central Bank (D’Acunto et al., 2022; European Central Bank, 2024). The CES is a monthly online panel survey covering 11 euro area countries. Six countries (Belgium, Germany, Spain, France, Italy, the Netherlands) have been surveyed since April 2020 (wave 4); five additional countries (Austria, Greece, Finland, Ireland, Portugal) joined in April 2022 (wave 28). Each monthly wave surveys approximately 14,000 respondents, with a rotating panel structure that allows tracking individuals over time.

The CES is uniquely suited to this study for three reasons. First, it elicits *probabilistic* inflation expectations, asking respondents to assign probabilities to bins of future price changes rather than reporting a single point estimate. Following Manski (2004), I compute the individual-level expected inflation as the probability-weighted mean of bin midpoints:

$$E_i[\pi] = \sum_{k=1}^K p_{ik} \cdot m_k \quad (1)$$

where  $p_{ik}$  is the probability assigned to bin  $k$  and  $m_k$  is the midpoint. The survey uses 10 bins (waves 31–75) and 8 bins (waves 4–30), with bin midpoints ranging from  $-14\%$  to  $+14\%$ .

Second, the CES measures perceived past inflation with a quantitative question, allowing me to distinguish between the price-level channel (perceived inflation) and the expectations channel (forward-looking beliefs).

Third, the CES includes the three Lusardi-Mitchell financial literacy questions (Lusardi and Mitchell, 2011) on compound interest, real interest rates, and risk diversification. I construct a financial literacy score (0–3) counting correct answers and define  $HighFL_i = \mathbf{1}[FL_i \geq 2]$ .

### 3.2 Sample Construction

The main sample consists of respondents in the six original CES countries observed in at least two survey waves, yielding 941,286 person-wave observations across 66,622 respondents. I restrict to observations with non-missing probabilistic inflation expectations (i.e., bin probabilities summing to at least 95%). Table 2 reports summary statistics.

## 4 Empirical Design

Table 2: Summary Statistics

	Full sample	France (control)	Treated countries	Pre- treatment	Post- treatment
$E[\pi]$ (pp)	4.31	4.00	4.39	3.84	4.61
$SD[\pi]$ (pp)	2.12	1.71	2.23	1.76	2.35
Perceived inflation (pp)	4.82	3.93	5.05	4.08	5.30
Spending index	0.56	0.46	0.59	0.48	0.62
Financial literacy (0–3)	2.27	2.20	2.28	2.21	2.30
Age	2.26	2.30	2.25	2.24	2.27
Observations	941,286	194,303	746,983	371,784	569,502
Respondents	66,622	14,705	51,917	31,152	47,449

*Notes:* Sample comprises respondents in the six original CES countries (Belgium, Germany, Spain, France, Italy, Netherlands) observed in at least two survey waves. “Treated countries” are those that implemented energy VAT cuts (all except France). “Pre-treatment” includes all observations before the country’s VAT cut (and all French observations).  $E[\pi]$  is the probability-weighted mean inflation expectation. Spending index ranges from  $-2$  (decrease a lot) to  $+2$  (increase a lot). Financial literacy counts correct answers to the Lusardi-Mitchell “Big Three” questions.

## 4.1 Staggered Difference-in-Differences

The identification strategy exploits the staggered timing of energy VAT cuts across countries. The baseline specification is:

$$y_{ict} = \alpha_i + \delta_t + \beta \cdot VATCut_{ct} + \varepsilon_{ict} \quad (2)$$

where  $y_{ict}$  is the outcome for individual  $i$  in country  $c$  at wave  $t$ ;  $\alpha_i$  is an individual fixed effect;  $\delta_t$  is a wave fixed effect; and  $VATCut_{ct}$  is an indicator equal to one after country  $c$  implements its energy VAT cut. France ( $VATCut_{ct} = 0$  for all  $t$ ) serves as the never-treated control. Standard errors are clustered at the individual level.

The individual fixed effect absorbs all time-invariant heterogeneity across respondents (country, demographics, baseline beliefs). The wave fixed effect absorbs common time shocks (e.g., ECB announcements, oil price movements) that affect all countries simultaneously. The identifying assumption is that, conditional on individual and time fixed effects, the timing of the VAT cut is uncorrelated with country-specific trends in inflation expectations.

I also estimate a treatment intensity specification:

$$y_{ict} = \alpha_i + \delta_t + \gamma \cdot (VATCut_{ct} \times Size_c) + \varepsilon_{ict} \quad (3)$$

where  $Size_c$  is the magnitude of the VAT reduction in percentage points (ranging from 9 for the Netherlands to 17 for Italy).

## 4.2 Event Study

To examine dynamic effects and assess the parallel trends assumption, I estimate an event-study specification with three-wave bins:

$$y_{ict} = \alpha_i + \delta_t + \sum_{k \neq -1} \beta_k \cdot \mathbf{1}[\text{bin}_k] + \varepsilon_{ict} \quad (4)$$

where  $\mathbf{1}[\text{bin}_k]$  indicates event-time bins relative to the VAT cut date. The omitted category is the three waves immediately preceding treatment ( $t \in [-3, -1]$ ). French observations enter with all event-time indicators equal to zero.

## 4.3 Financial Literacy Interactions

To test whether financial literacy moderates the treatment effect, I estimate:

$$y_{ict} = \alpha_i + \delta_t + \beta_1 \cdot \text{VATCut}_{ct} + \beta_2 \cdot \text{VATCut}_{ct} \times \text{FL}_i + \varepsilon_{ict} \quad (5)$$

Because financial literacy is approximately time-invariant within individual and is therefore absorbed by  $\alpha_i$ , only the interaction  $\text{VATCut}_{ct} \times \text{FL}_i$  is identified. The coefficient  $\beta_2$  captures the differential effect of the VAT cut for more financially literate households.

## 4.4 Threats to Identification

Two potential concerns merit discussion. First, the staggered adoption of VAT cuts may introduce bias if treatment effects are heterogeneous across cohorts ([Goodman-Bacon, 2021](#); [de Chaisemartin and D’Haultfoeulle, 2020](#)). I address this by reporting results separately for individual country pairs (e.g., France vs. Belgium) and by excluding the earliest-treated country (Spain) as a robustness check. Second, the event study reveals that treated countries exhibit trending inflation expectations in the pre-treatment period. This is not surprising: the countries that adopted VAT cuts were precisely those experiencing more severe energy price pressures, which themselves affected inflation expectations. The individual and wave fixed effects absorb much of this variation, and the financial literacy interaction—which exploits within-country, within-wave, within-individual variation—is less susceptible to country-level confounds.

# 5 Results

## 5.1 Main Difference-in-Differences Estimates

Table 3 reports the main DiD results. Column (a) shows that the binary treatment indicator is positively but weakly associated with expected inflation ( $\beta = 0.054$ ,  $t = 1.70$ ). Column

(b) reports the intensity specification, which is not statistically significant. These results suggest that energy VAT cuts do not produce a strong direct effect on forward-looking inflation expectations.

In sharp contrast, column (c) shows that VAT cuts significantly reduce perceived past inflation ( $\beta = 0.763$ ,  $t = 13.53$ ). This 0.76 percentage point reduction in perceived inflation is economically meaningful—it represents approximately 15% of the sample mean of perceived inflation. The result confirms H1: households notice the price-level consequences of energy VAT cuts.

Column (d) documents a small but significant reduction in spending intentions ( $\beta = -0.016$ ,  $t = -2.26$ ), consistent with the Euler equation prediction that lower expected future prices reduce the incentive for current consumption (Hall, 1988).

The contrast between columns (a)–(b) and column (c) is the central finding of the paper. Households perceive the reduction in energy prices caused by VAT cuts (column c) but do not fully extrapolate this perception into revised expectations about future inflation (columns a–b). This partial processing is consistent with H2 and with models of limited attention to tax policy (Chetty et al., 2009; Taubinsky and Rees-Jones, 2018).

Table 3: Effect of Energy VAT Cuts on Household Expectations and Spending

	(a) $E[\pi]$	(b) $E[\pi]$	(c) Perc. infl.	(d) Spending
$VATCut_{ct}$	0.0539* (0.0318)		0.7627*** (0.0564)	-0.0159** (0.0070)
$VATCut_{ct} \times Size_c$		0.0021 (0.0021)		
Dep. variable	$E[\pi]$	$E[\pi]$	Perc.	Spend
Individual FE	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes
$R^2$ (within)	0.0011	0.0007	0.0051	-0.0005
Observations	941,286	941,286	935,887	941,286

Notes:  $VATCut_{ct}$  equals one after country  $c$  implements an energy VAT cut, zero otherwise. France (which used direct subsidies rather than VAT cuts) serves as the never-treated control.  $Size_c$  is the magnitude of the VAT reduction (in percentage points). All specifications include individual and wave fixed effects. Standard errors clustered at the individual level in parentheses. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

## 5.2 Event Study

Figure 1 and Table 4 present the event-study estimates. The pre-treatment coefficients are negative and statistically significant, indicating that treated countries exhibit different trends in inflation expectations relative to France before the VAT cut. As discussed in Section 4, this pattern reflects the fact that VAT cuts were endogenous responses to country-specific energy price pressures.



Despite these pre-trends, the event-study pattern is informative. The treated-control gap narrows sharply around the treatment date (the post-0 coefficient of +0.054 is the highest in the series), before widening again in later periods. This transitory convergence at the time of the VAT cut is consistent with a short-lived salience effect: the tax cut temporarily draws attention to energy prices and narrows the expectations gap, but the effect dissipates as the policy becomes routine.




Figure 1: Event Study: Inflation Expectations Around Energy VAT Cuts

*Notes:* Coefficients from Equation (4). Event time measured in waves (months) relative to the country’s VAT cut. The reference period is  $t \in [-3, -1]$ . Bars indicate 95% confidence intervals based on standard errors clustered at the individual level. French observations included with all event-time indicators equal to zero.

### 5.3 The Role of Financial Literacy

Table 5 reports specifications that interact the treatment indicator with financial literacy. In column (a), the interaction  $VATCut_{ct} \times HighFL_i$  is positive and highly significant ( $\beta = 0.235$ ,  $t = 3.40$ ). The base effect of  $VATCut_{ct}$  becomes negative and significant ( $\beta = -0.133$ ,  $t = -2.08$ ), indicating that low-literacy households actually revise expectations *downward* following VAT cuts. Financially literate households, by contrast, exhibit a net positive effect ( $-0.133 + 0.235 = +0.102$ ).

Column (b) uses the continuous literacy score. Each additional correct answer increases the treatment effect by 0.146 percentage points ( $t = 4.79$ ). This monotonic relationship confirms H3: financial literacy determines the depth to which households process tax information.

Table 4: Event Study: Inflation Expectations Around Energy VAT Cuts

Event window	Coefficient	Std. error
$t \in [-12, -10)$	$-0.5718^{***}$	(0.0446)
$t \in [-9, -7)$	$-0.3811^{***}$	(0.0333)
$t \in [-6, -4)$	$-0.3919^{***}$	(0.0268)
$t \in [0, 2]$	$0.0540^*$	(0.0300)
$t \in [3, 5]$	$-0.1355^{***}$	(0.0360)
$t \in [6, 8]$	$-0.1986^{***}$	(0.0423)
$t \in [9, 11]$	$-0.1129^{**}$	(0.0501)
$t \in [12, 14]$	$-0.3594^{***}$	(0.0568)
$t \geq 15$	$-0.3937^{***}$	(0.0633)
Reference period	$t \in [-3, -1]$	
Individual FE	Yes	
Wave FE	Yes	
Observations	941,286	

*Notes:* Dependent variable:  $E[\pi]_{it}$ . Event time  $t$  measured in waves (months) relative to the country’s energy VAT cut. France (never treated) included with all event dummies equal to zero. The reference period is  $t \in [-3, -1]$  (three months before the VAT cut). All specifications include individual and wave fixed effects. Standard errors clustered at the individual level.  $***p < 0.01$ ,  $**p < 0.05$ ,  $*p < 0.10$ .

The pattern across columns (a)–(d) reveals a striking asymmetry. Column (c) shows that the literacy interaction for perceived inflation is *also* positive and significant ( $\beta = 0.614$ ,  $t = 4.50$ ): literate households perceive a larger effect of VAT cuts on past inflation. Column (d) shows the same pattern for spending. Financially literate households respond to the VAT cut across all three dimensions—perceptions, expectations, and spending intentions—while less literate households respond weakly or not at all.

This finding extends the tax salience framework: salience depends not only on the visibility of the tax (whether it is included in the sticker price) but also on the cognitive capacity of the agent processing the information. In the language of Blankespoor et al. (2020), financial literacy reduces the “processing cost” of tax information, allowing literate households to translate price observations into coherent macroeconomic beliefs.

## 6 Additional Analyses

### 6.1 Heterogeneity by Income and Housing Tenure

Table 6 examines whether the treatment effect varies by income and housing tenure. Columns (a) and (c) interact the treatment with a low-income indicator (bottom two income quintiles). The income interaction is economically and statistically insignificant for both expectations and spending, suggesting that the VAT cut does not disproportionately affect the expectations of lower-income households.

Table 5: Financial Literacy and the Effect of Energy VAT Cuts

	(a) $E[\pi]$	(b) $E[\pi]$	(c) Perc. infl.	(d) Spending
$VATCut_{ct}$	-0.1326** (0.0638)	-0.2701*** (0.0757)	0.2820** (0.1258)	-0.1096*** (0.0127)
$VATCut_{ct} \times HighFL_i$	0.2348*** (0.0691)		0.6135*** (0.1364)	0.1204*** (0.0136)
$VATCut_{ct} \times FL_i$		0.1461*** (0.0305)		
Dep. variable	$E[\pi]$	$E[\pi]$	Perc.	Spend
Individual FE	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes
Observations	931,332	931,332	926,012	931,332

Notes:  $HighFL_i$  equals one if financial literacy score  $\geq 2$  (answered at least two of three Lusardi-Mitchell questions correctly).  $FL_i$  is the continuous score (0–3). All specifications include individual and wave fixed effects. Standard errors clustered at the individual level. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

Columns (b) and (d) interact with a renter indicator. Renters show a marginally larger expectations response ( $\beta = 0.099$ ,  $t = 1.62$ ) and a significantly larger spending reduction ( $\beta = -0.050$ ,  $t = -3.99$ ). The spending result is consistent with renters being more exposed to energy costs as a share of income, making the VAT cut more salient for their consumption decisions.

Table 6: Heterogeneous Effects by Income and Housing Tenure

	(a) $E[\pi]$	(b) $E[\pi]$	(c) Spending	(d) Spending
$VATCut_{ct}$	0.0515 (0.0359)	0.0207 (0.0364)	-0.0133* (0.0079)	0.0009 (0.0082)
$VATCut_{ct} \times LowIncome_i$	0.0059 (0.0452)		-0.0065 (0.0094)	
$VATCut_{ct} \times Renter_i$		0.0986 (0.0609)		-0.0500*** (0.0125)
Dep. variable	$E[\pi]$	$E[\pi]$	Spending	Spending
Individual FE	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes
Observations	941,286	941,286	941,286	941,286

Notes:  $LowIncome_i$  equals one for income quintiles 1–2.  $Renter_i$  equals one for renters (as opposed to homeowners). All specifications include individual and wave fixed effects. Standard errors clustered at the individual level. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

## 6.2 Robustness

Table 7 reports five robustness checks for the main DiD specification on expected inflation. Column (a) expands the sample to all 11 CES countries. Column (b) excludes Spain, the earliest-treated country whose pre-treatment period overlaps with the initial pandemic recovery. Column (c) restricts to Belgium (treated) and France (control)—two neighboring countries that share a language area, providing the cleanest geographical comparison. Column (d) excludes the COVID-era waves (April–December 2020) to ensure results are not driven by the early pandemic period. Column (e) uses the continuous treatment intensity (VAT cut magnitude in percentage points).

The main coefficient is positive in most specifications but varies in magnitude, reflecting the heterogeneity in treatment effects across country pairs. The Belgium-France comparison yields a near-zero coefficient, consistent with the close economic similarity between these countries. Excluding Spain reduces the estimate, confirming that the earliest and largest treatment drives much of the average effect. The intensity specification is not significant, suggesting that the relationship between VAT cut magnitude and expectations response is not linear.

Table 7: Robustness Checks: Effect of Energy VAT Cuts on  $E[\pi]$

	(a) All 11 countries	(b) Excl. Spain	(c) FR vs. BE	(d) Excl. COVID	(e) Intensity spec.
$VATCut_{ct}$	0.0728** (0.0292)	-0.0237 (0.0379)	0.0067 (0.0934)	0.0328 (0.0322)	
$VATCut \times Size$					0.0021 (0.0021)
Individual FE	Yes	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes	Yes
Observations	1,171,802	738,537	264,816	859,951	941,286

*Notes:* Dependent variable:  $E[\pi]_{it}$  in all columns. Column (a) includes all 11 CES countries. Column (b) excludes Spain (first treated). Column (c) restricts to France (control) and Belgium (treated March 2022) for a within-language-area comparison. Column (d) excludes COVID-era waves (April–December 2020). Column (e) uses the continuous treatment intensity (VAT cut magnitude in percentage points). All specifications include individual and wave fixed effects. Standard errors clustered at the individual level. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

## 6.3 Country-Level Time Series

Figure 2 plots mean inflation expectations by country over the sample period, with vertical lines marking each country’s VAT cut date. The figure illustrates the common rise in expectations during 2021–2022 and the subsequent decline, as well as the staggered timing

of policy interventions. France’s expectations broadly co-move with those of the treated countries, supporting its use as a control, though the figure also reveals the cross-country heterogeneity that motivates the robustness analysis.




Figure 2: Inflation Expectations by Country, 2020–2026

*Notes:* Monthly mean expected inflation ( $E[\pi]$ ) by country. Dashed vertical lines indicate each country’s energy VAT cut date. France (grey) never implemented a VAT cut.

## 6.4 Perceived vs. Expected Inflation

Figure 3 presents the relationship between perceived past inflation and expected future inflation at the country-wave level, separately for treated and control observations. The strong positive correlation confirms that perceptions and expectations are related but not identical. Treated country-waves tend to cluster toward lower perceived inflation (reflecting the VAT cut’s price-level effect) without a proportional reduction in expected inflation, visually illustrating the asymmetry documented in Table 3.

## 7 Conclusion

This paper studies whether and how households incorporate indirect tax changes into their inflation beliefs. Exploiting the staggered adoption of energy VAT cuts across five European countries, with France as a natural control, I document an asymmetry in tax information processing: households perceive the price-level effects of VAT cuts (a 0.76 pp reduction in perceived inflation) but do not fully translate these perceptions into revised forward-looking expectations. Financial literacy is a key moderator: each additional




Figure 3: Perceived Past vs. Expected Future Inflation

*Notes:* Each point represents a country-wave mean. Treated observations (circles) are country-waves after the VAT cut; control observations (squares) are French country-waves. The dashed line is the 45-degree line.

point on the Lusardi-Mitchell scale increases the expectations response by 0.15 percentage points, and only households with high financial literacy exhibit statistically significant expectations updating.

These findings have implications for tax policy design and for the accounting literature on information processing. First, the results suggest that the macroeconomic effectiveness of indirect tax changes depends on whether households understand the policy’s implications for future prices. Simply reducing prices through VAT cuts may not anchor inflation expectations if households lack the financial sophistication to reason from prices to aggregate inflation. Second, the findings extend the tax salience framework (Chetty et al., 2009) by showing that salience is not merely a property of the tax instrument but also of the agent—a point that connects to the disclosure processing literature’s emphasis on investor sophistication (Blankespoor et al., 2020). Third, the paper responds to calls for accounting research on consumption taxes, survey methods, and cross-country settings (Hanlon, 2025), demonstrating that household-level survey data can provide insights into how tax policy shapes economic beliefs.

Several limitations should be noted. The endogeneity of VAT cut timing—countries facing more severe energy price pressures were more likely to act—introduces pre-trends that complicate causal inference. While the financial literacy interaction, which exploits within-country variation, is less affected by this concern, the main DiD estimates should be

interpreted with caution. Future research could leverage administrative transaction data to study how VAT cuts affect actual energy expenditures and whether the perceptions documented here translate into measurable consumption responses.

## References

- Scott R Baker, Stephanie Johnson, and Lorenz Kueng. The effect of sales tax on consumer behavior: Evidence from household scanner data. *American Economic Journal: Economic Policy*, 13(4):104–136, 2021.
- Elizabeth Blankespoor, Ed deHaan, and Ivan Marinovic. Disclosure processing costs, investors’ information choice, and equity market outcomes: A review. *Journal of Accounting and Economics*, 70(2–3):101344, 2020.
- Raj Chetty, Adam Looney, and Kory Kroft. Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–1177, 2009.
- Olivier Coibion and Yuriy Gorodnichenko. Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–2678, 2015.
- Francesco D’Acunto, Ulrike Malmendier, Juan Ospina, and Michael Weber. Exposure to grocery prices and inflation expectations. *Journal of Political Economy*, 129(5):1615–1639, 2021.
- Francesco D’Acunto, Daniel Hoang, Maritta Paloviita, and Michael Weber. The consumer expectations survey: An overview and first evaluation. *ECB Occasional Paper*, (287), 2022.
- Clément de Chaisemartin and Xavier D’Haultfœuille. Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996, 2020.
- European Central Bank. Consumer expectations survey: Methodology and results. *ECB Statistics*, 2024. URL [https://www.ecb.europa.eu/stats/ecb\\_surveys/consumer\\_exp\\_survey/html/index.en.html](https://www.ecb.europa.eu/stats/ecb_surveys/consumer_exp_survey/html/index.en.html).
- Amy Finkelstein. E-ztax: Tax salience and tax rates. *Quarterly Journal of Economics*, 124(3):969–1010, 2009.
- Andrew Goodman-Bacon. Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277, 2021.
- Robert E Hall. Intertemporal substitution in consumption. *Journal of Political Economy*, 96(2):339–357, 1988.
- Michelle Hanlon. New directions in tax research. *The Accounting Review*, 100(1):1–30, 2025.



- Tullio Jappelli and Mario Padula. Financial literacy and consumption. *Journal of Banking and Finance*, 37(8):2779–2792, 2013.
- Annamaria Lusardi and Olivia S Mitchell. Financial literacy around the world: An overview. *Journal of Pension Economics and Finance*, 10(4):497–508, 2011.
- Annamaria Lusardi and Olivia S Mitchell. The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature*, 52(1):5–44, 2014.
- Ulrike Malmendier and Stefan Nagel. Learning from inflation experiences. *Quarterly Journal of Economics*, 131(1):53–87, 2016.
- Charles F Manski. Measuring expectations. *Econometrica*, 72(5):1329–1376, 2004.
- Dmitry Taubinsky and Alex Rees-Jones. Attention variation and welfare: Theory and evidence from a tax salience experiment. *Review of Economic Studies*, 85(4):2462–2496, 2018.

## A Variable Definitions

Table 8: Variable Definitions

Variable	Definition
$E[\pi]_{it}$	Expected inflation: probability-weighted mean of bin mid-points from probabilistic survey question (c1152/c1150)
$SD[\pi]_{it}$	Inflation uncertainty: standard deviation of the individual's subjective probability distribution
Perceived inflation	Quantitative perceived past 12-month inflation rate (c1120), winsorized at $\pm 50$
Spending index	Durable goods spending intention, rescaled: +2 (increase a lot) to -2 (decrease a lot)
$VATCut_{ct}$	= 1 if country $c$ has implemented its energy VAT cut by wave $t$
$Size_c$	Magnitude of the energy VAT reduction (percentage points)
$FL_i$	Financial literacy score: count of correct answers (0–3) to Lusardi-Mitchell questions on compound interest (b5020), real rates (b5030), and diversification (b5040)
$HighFL_i$	= 1 if $FL_i \geq 2$
$LowIncome_i$	= 1 if household income is in quintile 1 or 2
$Renter_i$	= 1 if respondent rents (vs. owns with or without mortgage)

## B Inflation Expectations by Treatment Status and Financial Literacy




Figure 4: Inflation Expectations Over Time, by Treatment Status and Financial Literacy

*Notes:* Panel A: households with financial literacy score below 2. Panel B: households with score  $\geq 2$ . Treated = countries that implemented energy VAT cuts; Control = France.

## C Expectations Timeline




Figure 5: Inflation Expectations and Perceived Inflation Over Time

*Notes:* Panel A: mean expected inflation ( $E[\pi]$ ). Panel B: mean perceived past inflation. Dashed vertical lines indicate approximate dates of the first (Spain, June 2021) and median (Belgium, March 2022) VAT cuts.