

How Do Households Update Inflation Expectations?

Evidence from Probabilistic Beliefs and Financial Literacy

May 4, 2026

Abstract

How do households process new information when forming inflation expectations, and what role does financial literacy play? Using the ECB Consumer Expectations Survey—a large panel covering 11 euro area countries from 2020 to 2026—I exploit the unique probabilistic format of the survey to measure both the level and the uncertainty of individual inflation beliefs. I document three novel findings. First, households incorporate about 42% of their forecast errors into updated expectations, indicating substantial but incomplete information processing. Second, financial literacy has opposing effects depending on the information source: literate households update *more* in response to individual-level forecast errors (+2.9 percentage points per literacy point, $t = 7.6$) but *less* in response to exogenous, country-level information (−3.7 pp per point, $t = -14.8$). Third, this sign reversal is consistent with a Bayesian learning model in which literate households have both more precise priors and less noisy signals. The findings imply that financial education may improve expectations anchoring not by increasing responsiveness to news, but by reducing the noise with which households perceive economic signals.

JEL Classification: D83, D84, E31, E52, G53

Keywords: Inflation expectations, Bayesian updating, financial literacy, household beliefs, ECB Consumer Expectations Survey

Contents

1	Introduction	3
2	Theoretical Framework	4
3	Data and Measurement	5
3.1	The ECB Consumer Expectations Survey	5
3.2	Probabilistic Inflation Expectations	6
3.3	Perceived Past Inflation	6
3.4	Forecast Errors	6
3.5	Financial Literacy	7
4	Stylized Facts	7
5	Empirical Strategy	8
6	Main Results	8
6.1	Baseline Updating	9
6.2	Financial Literacy and Individual Forecast Errors	9
6.3	Exogenous Forecast Errors: The Sign Reversal	10
6.4	Interpretation: Bayesian Learning with Heterogeneous Precision	10
7	Heterogeneity Across Inflation Regimes	11
8	Implications for Spending	13
9	Robustness	13
10	Conclusion	14

1 Introduction

Inflation expectations are central to monetary policy transmission. If households expect prices to rise, they may accelerate purchases, push up wages, and set in motion the very inflation they fear. Central banks therefore invest heavily in anchoring expectations—through communication, forward guidance, and credibility. Yet we know surprisingly little about how individual households actually *process* inflation-relevant information. Do they update their beliefs rationally when new data arrives? Does financial sophistication help or hinder this process?

This paper provides new evidence on these questions by exploiting a distinctive feature of the ECB Consumer Expectations Survey (CES): respondents report their inflation beliefs as full probability distributions, allocating 100 probability points across bins of future price changes. From these distributions I compute individual-level means ($E[\pi]$), variances ($\text{Var}[\pi]$), and higher moments—a richness that point-estimate surveys cannot provide. Combined with the CES panel structure (90,000 respondents tracked monthly across 11 euro area countries) and the Lusardi-Mitchell “Big Three” financial literacy questions, this allows me to study how the *entire distribution* of beliefs evolves in response to new information, and how financial literacy moderates the updating process.

The empirical strategy is straightforward. For each respondent observed in consecutive survey waves, I compute the revision in expected inflation, $\Delta E[\pi]_{it} = E[\pi]_{it} - E[\pi]_{i,t-1}$, and relate it to the “forecast error”—the gap between newly perceived actual inflation and the prior expectation. In a simple Bayesian framework, the coefficient on this forecast error equals the Kalman gain: a weight between zero (no updating) and one (full adjustment to new information) that depends on the relative precision of the prior and the signal.

Three main findings emerge.

First, the average updating rate is 0.42: households revise their expectations by about 42 cents for every one-percentage-point forecast error. This is well above zero (they do respond to information) but well below one (they do not fully adjust), placing them in the “sticky information” range documented by [Coibion and Gorodnichenko \(2015\)](#) at the aggregate level.

Second, financial literacy has a strong but nuanced effect. When I use individual-level forecast errors—which combine a true information signal with individual-specific perception noise—literate households update *more*. Each additional correct answer on the Big Three increases the updating coefficient by 2.9 percentage points ($t = 7.6$), so that the most literate households (score 3) update at 45.5% versus 36.8% for the least literate (score 0). However, when I replace the individual forecast error with an exogenous, leave-one-out country-wave mean of perceived inflation, the sign flips: literate households now update *less* (-3.7 pp per point, $t = -14.8$), from 91.7% at score 0 to 80.7% at score 3.

Third, this sign reversal is exactly what a Bayesian model with heterogeneous precision

predicts. Literate households have more precise priors (their $SD[\pi]$ is 2.01 versus 2.70 for the illiterate). When facing their own—noisier—signals, the dominant margin is signal precision: literate households’ signals are cleaner, so they extract more information. When facing an exogenous, common signal, the dominant margin is prior precision: literate households have tighter priors and therefore give less weight to the same external information. The opposing signs are two sides of the same Bayesian coin.

These results have direct policy implications. The conventional wisdom is that “more financially literate households are more responsive to central bank communication.” My findings qualify this: literate households are better at processing information, but their superior priors also make them less sensitive to any single piece of news. Financial education, then, improves expectations anchoring not by amplifying responsiveness but by reducing the noise in the signal-extraction process.

The paper contributes to several literatures. It extends the information rigidity framework of [Coibion and Gorodnichenko \(2015\)](#) from aggregate to individual-level data, showing that the average “stickiness” masks large heterogeneity by financial literacy. It provides a new test of Bayesian updating using individual probability distributions, following the measurement approach pioneered by [Manski \(2004\)](#). It connects to the financial literacy literature ([Lusardi and Mitchell, 2014](#); [Jappelli and Padula, 2013](#)) by providing a structural interpretation of why literacy matters for expectations: not because literate households pay more attention, but because their signals are less noisy. Finally, it adds to the growing body of work using the CES ([D’Acunto et al., 2022](#)) by demonstrating the analytical value of the probabilistic expectation format.

The remainder of the paper is organized as follows. Section 2 presents the theoretical framework. Section 3 describes the data. Section 4 documents stylized facts. Section 5 describes the empirical strategy. Section 6 presents the main results. Section 7 explores heterogeneity across inflation regimes and countries. Section 8 examines implications for spending. Section 9 provides robustness checks. Section 10 concludes.

2 Theoretical Framework

Consider a household i entering period t with a prior belief about inflation that is normally distributed with mean $\mu_{i,t-1} = E[\pi]_{i,t-1}$ and variance $\sigma_{prior,i}^2$. During period t , the household observes a noisy signal s_{it} about actual inflation:

$$s_{it} = \pi_t + \varepsilon_{it}, \quad \varepsilon_{it} \sim N(0, \sigma_{signal,i}^2) \quad (1)$$

where π_t is the true inflation rate and $\sigma_{signal,i}^2$ is the signal noise, which may vary across households.

Bayesian updating yields the posterior mean:

$$E[\pi]_{it} = (1 - \kappa_i) \cdot E[\pi]_{i,t-1} + \kappa_i \cdot s_{it} \quad (2)$$

where $\kappa_i = \sigma_{prior,i}^2 / (\sigma_{prior,i}^2 + \sigma_{signal,i}^2)$ is the Kalman gain. Rearranging:

$$\Delta E[\pi]_{it} = \kappa_i \cdot (s_{it} - E[\pi]_{i,t-1}) = \kappa_i \cdot FE_{it} \quad (3)$$

where $FE_{it} = s_{it} - E[\pi]_{i,t-1}$ is the forecast error.

The Kalman gain κ_i is increasing in prior uncertainty $\sigma_{prior,i}^2$ and decreasing in signal noise $\sigma_{signal,i}^2$. Financial literacy plausibly affects both:

- **More precise priors:** Literate households have better-calibrated beliefs, so $\sigma_{prior,i}^2$ is lower. This *decreases* κ_i .
- **Less noisy signals:** Literate households process economic information more accurately, so $\sigma_{signal,i}^2$ is lower. This *increases* κ_i .

The net effect is ambiguous and depends on which margin dominates. Crucially, the two margins can be separated empirically by varying the signal source:

- When FE_{it} is computed from the individual's own perceived inflation, the signal noise $\sigma_{signal,i}^2$ varies across individuals. Literate households have lower signal noise, so the dominant margin is signal precision: $\partial\kappa/\partial FL > 0$.
- When FE_{it} is computed from an exogenous, common source (e.g., the country-wave average perceived inflation), signal noise is equalized across individuals. The dominant margin is now prior precision: $\partial\kappa/\partial FL < 0$.

This yields the key testable prediction: the sign of the financial literacy interaction should *reverse* when moving from individual to exogenous forecast errors. Observing this reversal would constitute strong evidence for Bayesian learning with heterogeneous precision.

3 Data and Measurement

3.1 The ECB Consumer Expectations Survey

The CES is an online panel survey administered monthly by the European Central Bank since January 2020. It covers 11 euro area countries—Belgium, Germany, Spain, France, Italy, and the Netherlands from wave 1, joined by Austria, Greece, Finland, Ireland, and Portugal from wave 28 (April 2022). Each wave surveys approximately 20,000 respondents, with substantial panel overlap across waves. After restricting to respondents observed in

at least two waves, the estimation sample contains 1,081,658 person-wave observations from 90,144 unique respondents.

3.2 Probabilistic Inflation Expectations

The CES elicits inflation expectations using a probabilistic format following [Manski \(2004\)](#). Respondents allocate 100 probability points across bins of possible future price changes. Two formats are used: an 8-bin version (waves 4–30, bins centered at $\pm 10, \pm 6, \pm 3, \pm 1$) and a 10-bin version (waves 31–75, adding bins at ± 14). For each respondent-wave, I compute:

$$E[\pi]_{it} = \sum_k m_k \cdot p_{ik,t} \quad (4)$$

$$\text{Var}[\pi]_{it} = \sum_k m_k^2 \cdot p_{ik,t} - (E[\pi]_{it})^2 \quad (5)$$

$$\text{SD}[\pi]_{it} = \sqrt{\text{Var}[\pi]_{it}} \quad (6)$$

where m_k is the midpoint of bin k and $p_{ik,t}$ is the probability weight assigned by respondent i to bin k at wave t . Observations where the total probability weight is below 95% are dropped.

3.3 Perceived Past Inflation

The CES also asks respondents to report their quantitative perception of past inflation (variable c1120). I use this as the “signal” s_{it} in the updating framework. Reports exceeding ± 50 percentage points are trimmed as outliers.

3.4 Forecast Errors

Two forecast error measures are constructed:

- **Individual forecast error:** $FE_{it} = \pi_{it}^{perc} - E[\pi]_{i,t-1}$, where π_{it}^{perc} is respondent i ’s perceived past inflation. This combines a true information component with individual perception noise.
- **Exogenous forecast error:** $FE_{it}^X = \bar{\pi}_{-i,ct}^{perc} - E[\pi]_{i,t-1}$, where $\bar{\pi}_{-i,ct}^{perc}$ is the leave-one-out mean of perceived past inflation in country c at wave t (excluding respondent i). By construction, this removes individual perception noise, providing a common signal that varies only at the country-wave level.

3.5 Financial Literacy

Financial literacy is measured using the standard Lusardi-Mitchell “Big Three” questions on compound interest, real interest rates, and risk diversification (Lusardi and Mitchell, 2014). The score $FL_i \in \{0, 1, 2, 3\}$ counts the number of correct answers. In the estimation sample, 50.3% score 3 (all correct), 30.0% score 2, 13.5% score 1, and 5.3% score 0.

4 Stylized Facts

Before turning to regressions, I document three descriptive patterns that motivate the analysis.

Fact 1: Financial literacy predicts both expectation levels and uncertainty.

Table 1 reports summary statistics by financial literacy group. Mean inflation expectations are highest for $FL = 2$ (4.71 pp) and lowest for $FL = 0$ (3.95 pp). More strikingly, expectation uncertainty falls monotonically with literacy: the average $SD[\pi]$ declines from 2.70 for $FL = 0$ to 2.04 for $FL = 3$. This is consistent with literate households having more precise priors.

Fact 2: Literate households have smaller forecast errors. The mean absolute forecast error declines sharply from 6.04 pp ($FL = 0$) to 3.09 pp ($FL = 3$)—a nearly 50% improvement in accuracy. This reflects both better-calibrated priors and more precise perception of actual inflation.

Fact 3: Expectation dynamics vary dramatically across the inflation cycle.

Figure 1 shows that inflation expectations surged across all literacy groups during 2022–2023, peaked around wave 38 (February 2023), and subsequently declined. The expectation *uncertainty* ($SD[\pi]$) also rose during the surge, but less for literate households, suggesting that financial literacy provides anchoring during volatile periods.

Table 1: Summary Statistics by Financial Literacy

	Obs.	Respondents	$\bar{E}[\pi]$	$SD(E[\pi])$	$\overline{SD}[\pi]$	$\overline{ FE }$	\bar{S}
$FL = 0$	57,194	4,635	3.95	5.64	2.70	6.04	0.354
$FL = 1$	145,595	12,186	4.39	5.20	2.42	5.37	0.446
$FL = 2$	324,642	26,702	4.71	4.72	2.21	4.29	0.579
$FL = 3$	544,291	43,746	4.46	4.03	2.04	3.09	0.675
Full sample	1,081,658	90,144	4.50	4.52	2.18	3.92	0.598

Notes: FL denotes the financial literacy score (0–3), counting correct answers to the Lusardi-Mitchell “Big Three” questions. $E[\pi]$ is the probability-weighted mean of the individual’s inflation expectation distribution. $SD[\pi]$ is the corresponding standard deviation. $|FE|$ is the absolute forecast error (perceived past inflation minus prior expected inflation). S is the spending intention index (–2 to +2). Panel respondents with ≥ 2 waves. Sample: CES waves 5–75, 11 euro area countries.




Figure 1: Inflation expectations and uncertainty by financial literacy, 2020–2026

5 Empirical Strategy

The baseline specification regresses the expectation revision on the forecast error with individual and wave fixed effects:

$$\Delta E[\pi]_{it} = \alpha_i + \delta_t + \beta \cdot FE_{it} + \varepsilon_{it} \quad (7)$$

where α_i absorbs all time-invariant individual characteristics (including financial literacy, country, and demographics) and δ_t absorbs common time trends. The coefficient β estimates the average Kalman gain—the fraction of the forecast error incorporated into the updated expectation.

To test for heterogeneity by financial literacy, I interact the forecast error with the literacy score:

$$\Delta E[\pi]_{it} = \alpha_i + \delta_t + \beta_0 \cdot FE_{it} + \beta_1 \cdot FE_{it} \times FL_i + \varepsilon_{it} \quad (8)$$

Even though FL_i is time-invariant and absorbed by α_i , the interaction $FE_{it} \times FL_i$ is identified because FE_{it} varies across waves.

The key identification test uses the exogenous forecast error FE_{it}^X in place of FE_{it} :

$$\Delta E[\pi]_{it} = \alpha_i + \delta_t + \beta_0^X \cdot FE_{it}^X + \beta_1^X \cdot FE_{it}^X \times FL_i + \varepsilon_{it} \quad (9)$$

Under the Bayesian framework in Section 2, we predict $\beta_1 > 0$ (literacy increases updating when signals are noisy) but $\beta_1^X < 0$ (literacy decreases updating when signals are common). All standard errors are clustered at the individual level.

6 Main Results

6.1 Baseline Updating

Table 2 reports the main updating regressions. Column (a) estimates the baseline specification (7): the coefficient on the individual forecast error is $\hat{\beta} = 0.424$ ($t = 125.9$), indicating that households revise their expectations by about 42 cents for every one-percentage-point surprise. The within R^2 of 0.357 means that forecast errors alone explain over one-third of within-individual expectation revisions.

6.2 Financial Literacy and Individual Forecast Errors

Column (b) adds the literacy interaction. The coefficient on $FE \times FL$ is $\hat{\beta}_1 = 0.029$ ($t = 7.6$), strongly positive. The implied updating rates range from 0.368 ($FL = 0$) to 0.455 ($FL = 3$)—a 24% increase from the least to the most literate. Column (c) adds the interaction with prior uncertainty ($SD[\pi]_{i,t-1}$), which enters negatively (-0.034 , $t = -41.6$), confirming that more uncertain households update *less* in response to individual forecast errors. The literacy effect is robust to this control. Column (d) estimates separate updating coefficients by literacy group, confirming the monotonic pattern.

Table 2: Expectation Updating: Individual Forecast Errors

	(a)	(b)	(c)	(d)
Forecast error (FE_{it})	0.4235*** (0.0029)	0.3681*** (0.0072)	0.4727*** (0.0073)	
$FE_{it} \times FL_i$		0.0290*** (0.0031)	0.0218*** (0.0028)	
$FE_{it} \times SD[\pi]_{i,t-1}$			-0.0342*** (0.0008)	
$FE_{it} \times \mathbf{1}[FL = 0]$				0.3977*** (0.0113)
$FE_{it} \times \mathbf{1}[FL = 1]$				0.3808*** (0.0066)
$FE_{it} \times \mathbf{1}[FL = 2]$				0.4204*** (0.0050)
$FE_{it} \times \mathbf{1}[FL = 3]$				0.4620*** (0.0043)
Individual FE	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes
R^2 (within)	0.3509	0.3528	0.3738	0.3495
Observations	1,059,515	1,051,293	1,051,293	1,059,515

Notes: Dependent variable: $\Delta E[\pi]_{it} = E[\pi]_{it} - E[\pi]_{i,t-1}$. Forecast error: $FE_{it} = \pi_{it}^{perc} - E[\pi]_{i,t-1}$, where π_{it}^{perc} is quantitative perceived past inflation. FL is financial literacy (0–3). All specifications include individual and wave fixed effects. Standard errors clustered at the individual level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

6.3 Exogenous Forecast Errors: The Sign Reversal

Table 3 replaces the individual forecast error with the exogenous, leave-one-out country-wave mean. Column (a) shows a much higher baseline updating rate of 0.845 ($t = 141.6$)—households respond nearly one-for-one to the country-level signal, as expected when individual noise is removed. Column (b) reveals the central finding: the literacy interaction is now *negative*, $\hat{\beta}_1^X = -0.037$ ($t = -14.8$). Literate households update *less* in response to the common signal, from 0.917 at FL = 0 to 0.807 at FL = 3.

Figure 2 visualizes this result. The left panel shows the monotonically increasing updating rate with individual forecast errors; the right panel shows the monotonically decreasing rate with exogenous forecast errors. This sign reversal is the paper’s central empirical contribution.

Table 3: Expectation Updating: Exogenous Forecast Errors

	(a)	(b)	(c)	(d)
Exog. forecast error (FE_{it}^X)	0.8453*** (0.0024)	0.9167*** (0.0061)	0.9258*** (0.0062)	
$FE_{it}^X \times FL_i$		-0.0365*** (0.0026)	-0.0358*** (0.0026)	
$FE_{it}^X \times SD[\pi]_{i,t-1}$			-0.0060*** (0.0008)	
$FE_{it}^X \times \mathbf{1}[FL = 0]$				0.9158*** (0.0094)
$FE_{it}^X \times \mathbf{1}[FL = 1]$				0.8753*** (0.0057)
$FE_{it}^X \times \mathbf{1}[FL = 2]$				0.8495*** (0.0041)
$FE_{it}^X \times \mathbf{1}[FL = 3]$				0.8039*** (0.0034)
Individual FE	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes
R^2 (within)	0.4270	0.4276	0.4278	0.4231
Observations	1,066,362	1,058,030	1,058,030	1,066,362

Notes: Dependent variable: $\Delta E[\pi]_{it}$. The exogenous forecast error is $FE_{it}^X = \bar{\pi}_{-i,ct}^{perc} - E[\pi]_{i,t-1}$, where $\bar{\pi}_{-i,ct}^{perc}$ is the leave-one-out country-wave mean of perceived past inflation. This removes the individual’s own perception from the information signal. All specifications include individual and wave fixed effects. Standard errors clustered at the individual level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

6.4 Interpretation: Bayesian Learning with Heterogeneous Precision

The sign reversal is precisely what the Bayesian model predicts. When the forecast error contains individual perception noise ($\sigma_{signal,i}^2$ varies), literacy reduces signal noise and




Figure 2: Updating rates by financial literacy: individual vs. exogenous forecast errors

therefore increases the Kalman gain. When the forecast error is purged of individual noise (the leave-one-out mean), signal precision is equalized across households, and the only remaining heterogeneity is in prior precision: literate households have tighter priors ($SD[\pi] = 2.01$ versus 2.70) and therefore assign less weight to any given signal.

An alternative explanation—that literate households simply pay less attention to country-level news—is hard to reconcile with the positive individual-level coefficient, which shows they are more, not less, responsive to information in general.

7 Heterogeneity Across Inflation Regimes

The sample spans three distinct inflation environments: the pre-surge period (April 2020–January 2022, HICP below 3%), the inflation surge (February 2022–September 2023, HICP reaching 10.6%), and the post-surge period (October 2023–March 2026, HICP returning to target). Table 4 reports separate updating regressions by regime.

The baseline updating coefficient increases over time: 0.338 in the pre-surge, 0.378 during the surge, and 0.451 post-surge. Households learned to update more aggressively as they accumulated experience with large inflation fluctuations. The literacy interaction, conversely, declines: 0.040 pre-surge, 0.026 during the surge, and 0.019 post-surge. When inflation surprises are large and salient, the literacy advantage in signal extraction diminishes—even less literate households can detect a 10-percentage-point inflation shock.

Figure 3 shows forecast accuracy over time. All literacy groups saw accuracy deteriorate during the surge, but the gap between high and low literacy *widened* from about 2 pp before 2022 to nearly 5 pp at the peak. Literate households remained better anchored even when everyone’s forecasts worsened.

Table 4: Expectation Updating by Inflation Regime

	Pre-surge (w4–25)	Surge (w26–45)	Post-surge (w46–75)	Pooled interactions
FE_{it}	0.3377*** (0.0138)	0.3777*** (0.0093)	0.4508*** (0.0132)	0.3169*** (0.0118)
$FE_{it} \times FL_i$	0.0399*** (0.0063)	0.0257*** (0.0041)	0.0191*** (0.0055)	0.0443*** (0.0055)
$FE_{it} \times Surge_t$				0.0284** (0.0117)
$FE_{it} \times Post_t$				0.1184*** (0.0165)
$FE_{it} \times FL_i \times Surge_t$				−0.0194*** (0.0054)
$FE_{it} \times FL_i \times Post_t$				−0.0233*** (0.0073)
Individual FE	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes
R^2 (within)	0.3497	0.3647	0.3924	0.3564
Observations	201,086	327,435	516,490	1,051,293

Notes: Dependent variable: $\Delta E[\pi]_{it}$. Pre-surge: waves 4–25 (Apr 2020 – Jan 2022). Surge: waves 26–45 (Feb 2022 – Sep 2023). Post-surge: waves 46–75 (Oct 2023 – Mar 2026). The pooled specification includes regime interaction terms with the pre-surge period as baseline. Standard errors clustered at the individual level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.




Figure 3: Mean absolute forecast error by financial literacy, 2020–2026

8 Implications for Spending

If expectations matter for economic decisions, households that update differently should also behave differently. Table 5 tests this by regressing spending intentions (-2 to $+2$) on inflation expectations. Column (a) shows a strongly positive relationship: a one-percentage-point increase in $E[\pi]_{it}$ is associated with a 0.019 increase in spending intentions ($t = 48.8$). This is consistent with the intertemporal substitution motive—households expecting higher inflation plan to spend more now.

Column (b) adds expectation variance, which enters negatively: greater uncertainty about future inflation is associated with *lower* spending intentions, consistent with precautionary saving. Column (c) shows that the spending response to expectations is moderated by financial literacy, although the economic magnitude is modest.

Table 5: Inflation Expectations and Spending Intentions

	(a)	(b)	(c)	(d)
$E[\pi]_{it}$	0.0189*** (0.0004)	0.0191*** (0.0004)	0.0063*** (0.0009)	0.0066*** (0.0009)
$\text{Var}[\pi]_{it}$		0.0007*** (0.0001)		0.0007*** (0.0001)
$E[\pi]_{it} \times \text{FL}_i$			0.0064*** (0.0004)	0.0063*** (0.0004)
Individual FE	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes
R^2 (within)	0.0091	0.0093	0.0101	0.0102
Observations	1,066,362	1,066,362	1,058,030	1,058,030

Notes: Dependent variable: spending intention index S_{it} (scaled -2 to $+2$). $E[\pi]_{it}$ is the probability-weighted mean inflation expectation. $\text{Var}[\pi]_{it}$ is the individual-level variance. All specifications include individual and wave fixed effects. Standard errors clustered at the individual level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

9 Robustness

Table 6 reports a battery of robustness checks. Column (a) restricts to the original six countries with complete wave coverage; the literacy interaction is unchanged at 0.027. Column (b) winsorizes forecast errors at the 1st and 99th percentiles to address outlier concerns; results are virtually identical. Column (c) restricts to consecutive survey waves (no gaps between observations) to ensure that updating is measured over a uniform horizon; the coefficient is 0.030. Column (d) excludes Germany, the largest country, to verify that results are not driven by a single economy. Column (e) adds an interaction with top-quintile income, which is small and does not affect the literacy interaction.

Table 6: Robustness Checks

	(a) Orig. 6 countries	(b) Winsorized FE	(c) Consec. waves	(d) Excl. Germany	(e) Q1 vs Q5 income
FE_{it}	0.3840*** (0.0082)		0.3703*** (0.0083)	0.3348*** (0.0076)	0.3284*** (0.0107)
FE_{it} (wins.)		0.4456*** (0.0069)			
$FE_{it} \times FL_i$	0.0270*** (0.0036)		0.0299*** (0.0036)	0.0350*** (0.0033)	0.0135*** (0.0048)
$FE_{it} \times FL_i$ (wins.)		0.0201*** (0.0030)			
$FE_{it} \times \mathbf{1}[Q5]$					0.1388*** (0.0094)
Individual FE	Yes	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes	Yes
Observations	853,959	1,051,293	896,559	879,561	421,603

Notes: Dependent variable: $\Delta E[\pi]_{it}$. Column (a) restricts to original 6 CES countries (BE, DE, ES, FR, IT, NL). Column (b) winsorizes forecast errors at the 1st and 99th percentiles. Column (c) uses only consecutive survey waves (no gaps). Column (d) excludes Germany. Column (e) restricts to bottom and top income quintiles and adds an income interaction. All specifications include individual and wave fixed effects. Standard errors clustered at the individual level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure 4 shows that the updating coefficient varies over time but remains positive in every wave, and Figure 5 displays rolling-window estimates of both the average updating rate and the literacy gap, confirming the time-varying patterns documented in Table 4.

10 Conclusion

This paper uses the probabilistic format of the ECB Consumer Expectations Survey to study how households update inflation beliefs and how financial literacy moderates this process. The central finding is a sign reversal: financial literacy increases updating in response to individual forecast errors (which contain perception noise) but decreases updating in response to exogenous, country-level signals (which do not). This pattern is inconsistent with simple models where literacy uniformly makes households “more responsive” and instead supports a Bayesian framework where literacy affects both prior precision and signal quality.

The findings have two main implications for policy. First, aggregate measures of “information rigidity” mask substantial heterogeneity: the average updating rate of 42% conceals a range from 37% (low literacy) to 46% (high literacy). Monetary policy models that assume a uniform degree of inattention may miss important distributional effects.




Figure 4: Wave-by-wave updating coefficient, 2020–2026




Figure 5: Rolling-window updating rate and literacy gap




Figure 6: Cross-sectional distribution of expectation moments

Second, the case for financial education as a tool for expectations management is more nuanced than often assumed. Literate households are not simply “more attentive to the ECB”—they are better at extracting signal from noise, which makes their expectations both more responsive to genuine news and more resistant to transient shocks.

References

- Coibion, Olivier and Yuriy Gorodnichenko**, “Information rigidity and the expectations formation process: A simple framework and new facts,” *American Economic Review*, 2015, 105 (8), 2644–2678.
- D’Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber**, “The Consumer Expectations Survey: An Overview and First Evaluation,” *ECB Occasional Paper*, 2022, (287).
- Jappelli, Tullio and Mario Padula**, “Financial literacy and consumption,” *Journal of Banking and Finance*, 2013, 37 (8), 2779–2792.
- Lusardi, Annamaria and Olivia S Mitchell**, “The economic importance of financial literacy: Theory and evidence,” *Journal of Economic Literature*, 2014, 52 (1), 5–44.
- Manski, Charles F**, “Measuring expectations,” *Econometrica*, 2004, 72 (5), 1329–1376.