When Do Inflation Perceptions Matter for Monetary Policy? Household Expectations and Energy Prices in Poland

Edyta Wojtyla

ORCID: 0000-0002-5747-1220
Corresponding author: edyta.wojtyla@vp.pl
Running title: Inflation Perceptions and Monetary Policy
Preprint, September 2026

ABSTRACT

Household inflation perceptions often diverge from official inflation, but the policy meaning of that gap depends on what drives it and whether it reaches expectations. Using Polish quarterly data for 2004Q1-2025Q4, the analysis separates retrospective perceived inflation (Q51) from forward-looking expected inflation (Q61) and combines both with source-specific energy prices. Aggregate energy inflation adds little to the perception-gap equation once core and food inflation are controlled for. Fuel, electricity, and gas contain more distinct information, while positive fuel-price movements are associated with Q61 under conventional HAC inference. Small-sample bootstrap evidence is less precise. A semi-structural policy exercise then asks whether perception information improves stabilization when conventional policy coefficients are held fixed. In the baseline activity-gap specification, perception-gap responses of 0.25 and 0.50 reduce model-implied loss in 99.2 and 97.5 percent of accepted parameter draws. The quantitative gain becomes negligible when GDP growth replaces the filtered activity gap. The evidence supports treating household perceptions as supplementary policy information rather than as an additional stabilization target.

Keywords: inflation perceptions; inflation expectations; household expectations; energy prices; monetary policy; central bank information

JEL classifications: E31; E52; E58; D84

1. INTRODUCTION

Central banks routinely observe a gap between official inflation and the inflation households report experiencing. The gap is not itself a policy target. Its relevance depends on what it reveals. It may reflect measurement differences, heterogeneous consumption baskets, or temporary movements in highly visible prices. In other episodes it may signal that recent price changes are beginning to affect beliefs about future inflation. The monetary-policy problem is therefore not whether a large perception gap should trigger an automatic interest-rate response. It is whether the gap contains information that changes the assessment of inflation persistence beyond what conventional macroeconomic indicators already show.

This question matters because household expectations are economically consequential even when they differ from professional forecasts or model-consistent beliefs. Survey expectations display systematic heterogeneity, but they help explain household consumption, saving, borrowing, and responses to policy interventions (Weber et al., 2022; D'Acunto & Weber, 2024). For monetary policy, the relevant issue is consequently not whether household beliefs are statistically unbiased. It is whether they reveal changes in the information environment that may alter spending decisions, wage demands, price expectations, or the persistence of inflation beliefs.

A large literature shows that households do not process all prices symmetrically. Gasoline prices have received particular attention because they are frequent, visible, and volatile. Binder (2018) documents a close relationship between gasoline prices and inflation expectations, while

experimental and quasi-experimental work provides stronger evidence that pump-price changes can shift household expectations (Aidala et al., 2024; Jo & Klopack, 2025). At the same time, the relationship is neither universal nor mechanically proportional to salience. Aidala et al. (2024) find that the effect weakens with the forecast horizon, and Bürgi et al. (2026) report little evidence that households systematically overreact to expected gasoline prices in forming overall inflation expectations. Recent structural evidence also shows that gasoline shocks can affect one-year inflation expectations, but the magnitude and persistence depend on the model and episode (Vatsa et al., 2025). These results make the source of a perception gap potentially informative, while cautioning against treating any single price category as a sufficient statistic for expectations.

Recent monetary-policy research gives a natural setting for this question. Dietrich (2024) shows that category-specific attention can alter the inflation measure relevant for policy in a New Keynesian environment. Hahn and Marenčák (2026) develop a bounded-rationality model in which sectoral prices shape inflation perceptions and, through them, monetary-policy design. Applied studies reach a similar practical issue from different directions. Grajales-Olarte et al. (2026) use survey evidence on credibility to discipline a semi-structural model for Colombia, while Aguilar and Quineche (2026) examine how disaggregated inflation information affects policy design in Peru. The open question is not whether visible prices can influence beliefs. It is whether observed perception data add information that changes a policy assessment.

Poland provides a useful setting for examining this problem. The quarterly sample covers 2004Q1 to 2025Q4 and combines quantitative household survey measures from the European Commission with official consumer-price inflation, real activity, the National Bank of Poland reference rate, and disaggregated household-facing energy prices. The survey design separates Q51, a retrospective measure of perceived inflation, from Q61, a forward-looking measure of expected inflation. This distinction is central to the analysis because a household may report large recent price increases without expecting the same pace of inflation to continue.

The empirical evidence begins with the composition of household-facing energy prices. Once core and food inflation are controlled for, aggregate energy inflation adds little to the perception-gap equation, whereas the source-specific block is jointly informative. Electricity and gas remain associated with the gap, and positive fuel inflation is associated with Q51. The link with Q61 is narrower. Positive fuel inflation is associated with forward-looking expectations and remains positive at subsequent horizons under HAC inference, although block-bootstrap uncertainty is considerably larger. A Brent-price external-shifter exercise gives a strong first stage and estimates with consistent signs, but circular-block bootstrap results rule out a strong causal claim.

Two policy questions are kept separate. The historical reaction-function exercise asks whether survey and source variables help describe the quarterly path of the National Bank of Poland reference rate. Better fit in that exercise says nothing about whether the historical response was desirable. The semi-structural model addresses a different question: whether adding perception information changes stabilization performance within the same estimated stochastic environment.

The policy comparison uses three information sets. Model 0 contains measured inflation and real activity. Model 1 adds the aggregate perception gap. Model 2 further separates a source-

implied component from the residual part of the gap. Since the models are nested, lower optimized loss for a richer rule can arise mechanically and is not treated as evidence. The informative comparison keeps the conventional inflation and activity responses fixed at their Model 0 values and varies only the response to the perception gap. In the baseline activity-gap system, fixed responses of 0.25 and 0.50 lower loss at every tested inflation response. At the baseline point, they lower loss in 99.2 and 97.5 percent of 1,000 accepted parameter draws. This result is not quantitatively stable across activity measures. When GDP growth replaces the filtered activity gap, the admissible response is much smaller and the gain is close to zero. Source differentiation adds only a modest refinement.

The contribution lies in the use of perceptions as policy information, not in proposing a new target. Perception data can help distinguish a visible relative-price movement that remains confined to retrospective beliefs from one that is also associated with forward-looking expectations. Source composition is useful mainly for diagnosis. The estimates do not support assigning independent policy-response coefficients to fuel, electricity, or gas.

The practical implication follows directly from that distinction. A visible price movement that raises Q51 without a corresponding signal in Q61 calls primarily for monitoring and communication. If the same movement is also associated with higher expectations, it becomes more relevant to the assessment of inflation persistence. Household perceptions are therefore neither noise to be discarded nor variables to be stabilized in their own right. Their value depends on the information they add about the outlook for beliefs.

2. Inflation perceptions, household-facing prices, and the policy information problem

Subjective expectations occupy an unusual position in macroeconomics. They are central to models of intertemporal choice and monetary transmission, yet household survey responses are often dispersed, biased relative to realized inflation, and sensitive to the information respondents encounter. The modern survey literature does not treat these properties as reasons to discard household expectations. Rather, they are evidence that households form beliefs with heterogeneous information sets and limited attention. Weber et al. (2022) summarize the measurement, determinants, and behavioral consequences of household and firm inflation expectations, while D'Acunto and Weber (2024) review evidence that survey beliefs help explain consumption, saving, investment, debt choices, and heterogeneous responses to monetary and fiscal policy. For the present paper, this literature provides a measurement premise rather than a rational-expectations assumption: a survey belief can be policy-relevant even if it is not an unbiased forecast.

Retrospective inflation perceptions and forward-looking expectations should nevertheless be kept distinct. The former summarize households' assessment of recent price changes; the latter concern the future inflation rate that enters intertemporal decisions. Bussière et al. (2026), using European Commission consumer-survey data for the euro area, show that perceptions of recent price changes remain important for expectations after controlling for inflation components, household characteristics, and macroeconomic conditions. Their results are heterogeneous across

countries, which is important for a single-country design. They support the idea that a perception measure can contain information for expectations, but they do not imply that every divergence between perceived and measured inflation is an expectations shock.

One reason perceptions can diverge from aggregate inflation is that households encounter prices unevenly. Gasoline is the most studied example. Binder (2018) finds a positive relationship between gasoline prices and household inflation expectations but cautions against a simple overweighting interpretation. Aidala et al. (2024) provide experimental evidence that gasoline-price changes affect inflation expectations, although the effect is small relative to some other goods and weakens at longer horizons. Jo and Klopack (2025) exploit temporary state gas-tax suspensions in the United States and identify a causal link between lower pump prices and lower household inflation expectations when the tax cut passes through to retail prices. Vatsa et al. (2025), using a partially identified Bayesian structural vector autoregression, also find that gasoline-price shocks affect one-year inflation expectations. These studies establish that fuel prices can matter for expectations, so that result cannot be the novelty of the present paper.

Recent evidence also prevents a universal gasoline-salience story. Bürgi et al. (2026) find little evidence in two U.S. household surveys that households overreact to expected gasoline prices when forming overall inflation expectations. The contrast across designs suggests that the relevance of a particular category depends on the price episode, the horizon, the information environment, and the survey question. That is one reason the present analysis retains fuel, electricity, and gas separately rather than imposing a single estimated salience index. It also motivates a distinction between source composition as observed information and salience as a behavioral mechanism proposed in the literature.

Category-level attention has direct implications for monetary policy. Dietrich (2024) shows that non-core inflation forecasts can disproportionately influence headline inflation expectations and embeds that mechanism in a multi-sector New Keynesian model. Under bounded rationality, the inflation measure receiving policy weight depends on how households allocate attention across consumption categories. Hahn and Marenčák (2026) move closer to the perception concept used here. They construct a bounded-rationality New Keynesian model in which sectoral prices disproportionately influence current inflation perceptions and, through those perceptions, affect monetary-policy design. Their latest revision emphasizes the role of food prices. These papers are closest to the conceptual mechanism behind source-sensitive perception data; they also mean that the contribution of the present paper must lie elsewhere.

For a central bank, the difficult part is deciding what a supply-driven price movement means in real time. A relative-price shock may fade without broader consequences, or it may become part of a more persistent inflation process. Beckworth and Horan (2024) emphasize this identification problem: advice to look through temporary supply shocks is useful only if policymakers can distinguish temporary disturbances from changes that threaten the nominal anchor. Household perception data may contribute to that diagnosis when they reveal whether a visible price movement is spreading into beliefs about future inflation, but they cannot determine the policy rate on their own.

Several strands of the policy literature provide the relevant benchmark. De Mendonça (2018) studies credibility and the formation of inflation expectations across emerging inflation-targeting economies, including Poland. Álvarez and Sánchez (2019) treat inflation projections as inputs into monetary-policy decisions rather than as forecasting exercises in isolation. Rodriguez-Rodriguez et al. (2024) use reaction functions to characterize central-bank behavior, while Roy Trivedi and Ghosh (2025) connect central-bank communication with expectation anchoring. Grajales-Olarte et al. (2026) use survey evidence on credibility to discipline a semi-structural policy model for Colombia, and Aguilar and Quineche (2026) show how disaggregated inflation information can affect monetary-policy design in Peru. The common issue is whether an empirical signal changes the policy assessment, not whether the model containing it is more elaborate.

The unresolved question is more specific than whether salient prices or perceptions matter. Existing studies show that visible prices can influence household beliefs, that perceptions can be related to expectations, and that category-specific belief formation may affect monetary-policy design. Less is known about how an observed perception gap should enter the policy information set when both its total size and its source composition are available. The analysis below tests whether a fixed response to the total gap changes stabilization performance with conventional policy coefficients held constant. It then asks whether a limited distinction between source-implied and residual components adds anything further. The first comparison is the main policy test. The second is deliberately secondary because the source component is generated and only weakly validated on its own.

These choices set a clear limit on interpretation. The source regressions test information content rather than structural demand or attention equations. The horizon regressions are predictive exercises rather than impulse responses to identified shocks. Brent is used as an external shifter, but global oil prices can affect Polish expectations through channels other than retail fuel, so the exercise is described as quasi-instrumental rather than structural. The policy comparison is conditional on a semi-structural system and a stabilization criterion, not on a welfare-theoretic Ramsey problem.

3. Data and measurement

The empirical analysis combines quarterly Polish macroeconomic data with respondent-level household survey information and disaggregated household-facing energy prices. The main sample runs from 2004Q1 to 2025Q4, yielding 88 quarterly observations for the survey and principal macroeconomic series. Specifications that require the core- and food-inflation controls use the common actual-data sample ending in 2025Q1; no forecasts or projected observations are used to extend that subsample. The quarterly frequency matches the policy-rate and real-activity blocks and keeps the empirical horizon consistent across the survey and policy exercises.

The household belief measures are constructed from anonymised Polish respondent-level data from the European Commission's harmonised consumer survey. The survey item labelled Q51 elicits a quantitative estimate of the change in consumer prices over the previous 12 months, whereas the item labelled Q61 elicits the percentage change expected over the following 12 months. Q51 is therefore treated as retrospective perceived inflation and Q61 as forward-looking

expected inflation. Baseline quarterly series are survey-weighted means of the respondent-level quantitative answers. The question horizons and quantitative survey structure follow the European Commission methodological user guide (European Commission, 2026a).

The conditional design of the quantitative questions matters for interpretation. Q51 follows a non-zero directional response to the corresponding qualitative question on past price changes, and Q61 follows the analogous directional question about future prices. Across the underlying microdata, about 3,020 respondents contribute to a quarter on average, with roughly 2,318 valid quantitative Q51 responses and 2,085 valid Q61 responses. These counts support construction of the quarterly measures, but inference in the paper is based on the quarterly time series rather than on the pooled respondent count.

The main perception gap is the difference between Q51 perceived inflation and official Polish consumer price index (CPI) inflation over the corresponding 12-month period. CPI inflation is taken from Statistics Poland and remains the benchmark for measured inflation throughout the gap analysis. The gap is an observed survey-official divergence, not a latent structural state by construction. Whether it contains persistent information, reflects the composition of recent price movements, or contributes to monetary-policy assessment is left to the empirical analysis.

A discontinuity issue arises in the Polish consumer-survey history. The European Commission reported in May 2026 that a new provider had been introduced for the Polish consumer survey and that historical data from January 2018 onward had been replaced. The Commission also notes that underlying response shares may display a break at January 2018. Survey-based specifications therefore include an indicator from 2018Q1. The indicator controls for a potential provider-related discontinuity; it is not interpreted as an economic regime shift or as a change in the wording or horizon of Q51 and Q61 (European Commission, 2026b).

The source-composition block is built from Eurostat Harmonised Index of Consumer Prices (HICP) indices for total energy, fuels and lubricants for personal transport, electricity, and gas, accessed as monthly series through Federal Reserve Economic Data (FRED). Monthly non-seasonally adjusted index observations are averaged within each quarter, and year-on-year rates are calculated from the quarterly averages. Monthly observations from 2003 are retained only to construct the first 2004 year-on-year rates. Let f_t denote retail-fuel inflation. Its positive and negative parts are $f_t^+=\max(f_t,0)$ and $f_t^-=\min(f_t,0)$. This decomposition permits a formal asymmetry test; asymmetry is not imposed as an empirical result.

The HICP components are not used to reconstruct national headline CPI inflation or to redefine the perception gap. They enter as harmonised source variables describing household-facing energy-price movements. To reduce the possibility that the source variables merely proxy for broad inflation pressure, the regressions also control for year-on-year core inflation excluding food and energy and for food and non-alcoholic beverage inflation. These controls are Organisation for Economic Co-operation and Development (OECD) series distributed through FRED and are available in the actual-data estimation file through 2025Q1.

Real activity is measured using Eurostat real gross domestic product (GDP). Annualised quarter-on-quarter log GDP growth enters the historical reaction-function specifications. The later

policy exercise uses the real-GDP level to construct a quarterly activity gap and repeats the comparison with GDP growth as an alternative activity measure. The monetary-policy observable is the National Bank of Poland (NBP) reference rate, aggregated to quarterly frequency from the official interest-rate history. These are observed inputs to the empirical models. By contrast, the stabilization losses in Sections 7-8 are counterfactual quantities generated by the estimated semi-structural system and are not observed outcomes.

The external shifter used in the fuel-price identification exercise is the Europe Brent spot price published by the U.S. Energy Information Administration (EIA). Monthly dollar-per-barrel prices are averaged within each quarter. The preferred shifter is the positive part of year-on-year Brent growth; a first-order autoregressive (AR(1)) Brent innovation is retained as a weaker sensitivity measure. Brent is not labelled a structural oil-supply shock, and its exclusion restriction is not treated as automatic. The identification section therefore reports first-stage, reduced-form, two-stage, and block-bootstrap evidence separately.

TABLE 1. DATA, MEASUREMENT, AND SAMPLE CONSTRUCTION

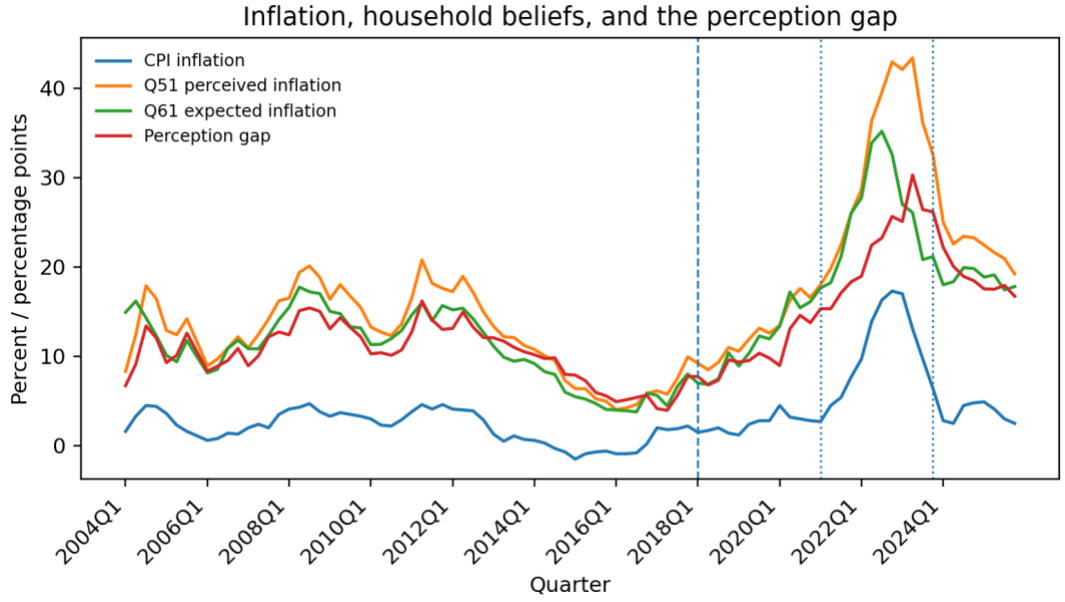
Variable	Meaning	Transformation used	Main coverage	Source	Baseline role	Key measurement caution
Q51	Perceived inflation, past 12 months	Survey-weighted quarterly mean of quantitative respondent answers	2004Q1-2025Q4	European Commission, Directorate-General for Economic and Financial Affairs	Perception measure / perception gap	Retrospective perception, not expectation
Q61	Expected inflation, next 12 months	Survey-weighted quarterly mean of quantitative respondent answers	2004Q1-2025Q4	European Commission, Directorate-General for Economic and Financial Affairs	Forward-looking expectation	Keep distinct from Q51
CPI	Official consumer-price inflation	Year-on-year quarterly CPI	2004Q1-2025Q4	Statistics Poland	Headline comparator / conventional inflation	National CPI remains perception-gap benchmark

Variable	Meaning	Transformation used	Main coverage	Source	Baseline role	Key measurement caution
Real GDP	Real economic activity	Annualised q/q log growth; level used for model activity gap	2004Q1-2025Q4	Eurostat	Activity measure	Model-derived gap is not raw data
NBP rate	Monetary-policy rate	Quarterly aggregation of official reference-rate history	2004Q1-2025Q4	Narodowy Bank Polski	Historical policy observable	Historical fit is not normative evidence
Energy	Aggregate household energy prices	Quarterly average index; y/y rate from quarterly averages	2004Q1-2025Q4	Eurostat HICP via FRED	Aggregate energy benchmark	Not used to reconstruct CPI
Fuel	Fuels/lubricants for personal transport	Quarterly average index; y/y; positive/negative parts retained	2004Q1-2025Q4	Eurostat HICP via FRED	Source-specific price variable	Asymmetry tested, not assumed
Electricity	Household electricity prices	Quarterly average index; y/y rate	2004Q1-2025Q4	Eurostat HICP via FRED	Source-specific price variable	HICP source proxy
Gas	Household gas prices	Quarterly average index; y/y rate	2004Q1-2025Q4	Eurostat HICP via FRED	Source-specific price variable	HICP source proxy
Core inflation	CPI excluding food and energy	Year-on-year quarterly rate	2004Q1-2025Q1	OECD via FRED	Broad-inflation control	Common actual-data sample ends 2025Q1
Food inflation	Food and non-alcoholic beverages	Year-on-year quarterly rate	2004Q1-2025Q1	OECD via FRED	Broad-inflation control	Common actual-data sample ends 2025Q1

Variable	Meaning	Transformation used	Main coverage	Source	Baseline role	Key measurement caution
Brent	Europe Brent spot price	Quarterly average USD/bbl; positive y/y change is preferred shifter	2003 lags; 2004-2025 estimation	U.S. EIA	External fuel-price shifter	Quasi-IV, not structural oil-supply shock

Notes: Q51 and Q61 are European Commission survey item labels. HICP denotes the Harmonised Index of Consumer Prices. The common source-controlled estimation sample ends in 2025Q1 because the core and food controls are not extended with projections.

FIGURE 1. MEASURED INFLATION, HOUSEHOLD PERCEPTIONS AND EXPECTATIONS, AND THE PERCEPTION GAP



Notes: The dashed vertical line marks 2018Q1, the potential provider-related discontinuity in the Polish consumer-survey history. Dotted lines mark the beginning and end of the 2021Q1-2023Q4 high-inflation episode used in robustness analysis.

4. EMPIRICAL STRATEGY AND IDENTIFICATION

The empirical strategy separates information content from causal identification. Contemporaneous regressions ask whether the composition of household-facing prices adds information for perceived or expected inflation after broader inflation conditions are controlled for. Direct horizon regressions ask whether the same source information is associated with survey expectations at later dates. The Brent exercise then supplies an external shifter for retail-fuel

inflation. These layers of evidence are interpreted separately rather than collapsed into a single causal claim.

Equations (1)-(3) distinguish measured year-on-year CPI inflation, Q51 perceived inflation, and Q61 expected inflation. The perception gap is Q51 perceived inflation minus measured CPI inflation, and the source vector contains positive and negative retail-fuel inflation, electricity inflation, and gas inflation. The baseline source equations are:

$$g_t = \alpha_g + \rho_g\, g_{t-1} + \gamma_{gc}\, \pi_t^{core} + \gamma_{gf}\, \pi_t^{food} + \beta_g'\, z_t + \delta_g\, D_t^{2018} + \varepsilon_t^g \quad (1)$$

$$\pi_t^p = \alpha_p + \rho_p\, \pi_{t-1}^p + \gamma_{p\pi}\, \pi_t + \gamma_{pc}\, \pi_t^{core} + \gamma_{pf}\, \pi_t^{food} + \beta_p'\, z_t + \delta_p\, D_t^{2018} + \varepsilon_t^p \quad (2)$$

$$\pi_t^e = \alpha_e + \rho_e\, \pi_{t-1}^e + \gamma_{e\pi}\, \pi_t + \gamma_{ec}\, \pi_t^{core} + \gamma_{ef}\, \pi_t^{food} + \beta_e'\, z_t + \delta_e\, D_t^{2018} + \varepsilon_t^e \quad (3)$$

The 2018 indicator controls for the potential provider-related discontinuity described in Section 3. The perception-gap equation conditions on core and food inflation, whereas the two survey-level equations also condition on measured CPI inflation because the survey levels have not already netted out headline inflation.

Each outcome is first estimated with aggregate energy inflation and then with the source vector. The joint null that all source coefficients are zero is the central test of incremental source information. Fuel-price asymmetry is assessed by the equality restriction on the positive- and negative-fuel coefficients. A significant positive-fuel coefficient combined with an insignificant negative-fuel coefficient is not, by itself, treated as evidence of asymmetry.

Baseline inference uses Newey-West heteroskedasticity- and autocorrelation-consistent (HAC) covariance estimates with four quarterly lags (Newey & West, 1987). The bandwidth allows for approximately one year of serial dependence while keeping the inference rule common across specifications. Because the effective time-series sample is small and several series are persistent, HAC significance is treated as asymptotic evidence rather than as a sufficient robustness criterion. Moving-block resampling is reported where the strength of the empirical claim matters (Künsch, 1989).

For the main forward-looking exercise, the direct horizon regression for h=0,…,4 is:

$$\pi_{t+h}^e = \alpha_h + \rho_h\, \pi_{t-1}^e + \Gamma_h'\, c_t + \beta_h'\, z_t + \delta_h\, D_t^{2018} + u_{t+h} \quad (4)$$

The control vector c_t contains headline, core, and food inflation together with the additional controls used in each robustness specification. The equation follows the direct-projection logic of Jordà (2005) but is used as a predictive validation exercise rather than as a structural impulse-response design. Specifications with a deterministic trend, a post-2020 indicator, a high-inflation indicator, and a combined high-inflation-plus-trend control test whether the source association is concentrated in the recent inflation episode or reflects low-frequency drift. These additions do not convert the horizon coefficients into causal effects (Montiel Olea & Plagborg-Møller, 2021).

A separate robustness specification defines a high-inflation indicator equal to one when CPI inflation is at least 5 percent and interacts it with positive fuel inflation. The threshold is fixed for the diagnostic exercise rather than estimated as a structural regime boundary. The interaction is therefore interpreted as state-dependence robustness only.

The analysis also considered compressing the observed source variables into a generated one-factor source score. Its stand-alone predictive precision is weak under block resampling, so that generated score is not used as a baseline empirical variable. Fuel, electricity, and gas are retained separately in the main empirical specifications. This limitation is also why the source-implied decomposition used later in Model 2 is treated as a secondary model diagnostic rather than as an independently validated policy signal.

The preferred external shifter is the positive part of year-on-year Europe Brent price growth. The first stage is:

$$f_t^+ = \alpha_F + \lambda_B \, b_t^+ + \Lambda_F' \, c_t + u_t^F \qquad (5)$$

The reduced form and second stage for Q61 are:

$$\pi_{t+h}^e = \alpha_h^R + \theta_h \, b_t^+ + \Gamma_h^{R'} \, c_t + u_{t+h}^R \qquad (6)$$

$$\pi_{t+h}^e = \alpha_h^{IV} + \beta_h^{IV} \hat{f}_t^+ + \Gamma_h^{IV'} \, c_t + u_{t+h}^{IV} \qquad (7)$$

The second-stage equation is estimated by two-stage least squares (2SLS). Agreement in sign across the first stage, reduced form, and 2SLS estimates is more informative here than the significance of any one coefficient. The exclusion restriction remains imperfect because Brent prices may affect Polish inflation expectations through imported inflation, exchange-rate developments, macroeconomic news, or other channels that do not operate through retail fuel. Brent is therefore used as an external shifter, and the estimates are interpreted as quasi-instrumental-variable (quasi-IV) evidence rather than as a structural oil-supply experiment.

The main 2SLS estimates use HAC inference, with a circular block bootstrap as a small-sample check (Politis & Romano, 1992). When HAC significance is not supported by block resampling, the coefficient is treated as suggestive rather than as decisive causal evidence. The empirical exercises answer different questions: contemporaneous regressions test incremental information content, horizon regressions examine predictive association, and the Brent shifter provides an additional check on the fuel channel. The counterfactual policy exercise in Sections 7-8 is kept separate from those identification claims.

5. Empirical results: source composition and expectations

Aggregate energy inflation and its individual components convey different information. In the perception-gap equation, adding aggregate energy inflation to core and food inflation produces a coefficient of 0.030 with a HAC p-value of 0.227. Once broader inflation conditions are controlled for, the aggregate measure contributes little independently. Replacing it with fuel, electricity, and gas changes the picture. The four source terms are jointly significant ($\chi^2 = 21.67$, $p < 0.001$), while electricity and gas enter the gap equation individually with coefficients of 0.041 ($p = 0.031$) and 0.026 ($p = 0.031$). The relevant information appears to lie in composition rather than in a generic energy effect.

The Q51 equation shows a similar pattern. Conditional on lagged Q51, headline CPI, core inflation, food inflation, and the 2018 survey indicator, the source block is jointly significant ($\chi^2 = 13.58$, $p = 0.009$). Positive retail-fuel inflation is associated with higher perceived inflation, with a coefficient of 0.058 (HAC standard error = 0.028, $p = 0.040$). Electricity and gas coefficients are 0.056 ($p = 0.022$) and 0.060 ($p = 0.006$), respectively. These coefficients should not be read as

literal CPI-style household weights. They show that individual household-facing energy prices retain information for reported perceptions after conventional inflation controls are included. Supplementary Table S1 reports the descriptive correlation matrix.

The contemporaneous evidence is weaker for forward-looking expectations, which is useful for interpretation. In the Q61 specification, positive fuel inflation enters positively at 0.080 (HAC standard error = 0.038, p = 0.036), whereas the four-variable source block is only marginal as a group ($\chi^2 = 8.18$, p = 0.085). The gas coefficient is also only marginal at conventional levels (0.048, p = 0.100), and electricity is not statistically distinguishable from zero. Thus the data do not support the stronger proposition that all energy components feed contemporaneously into expected inflation. They point instead to a narrower fuel-related signal whose relevance becomes more apparent in the horizon analysis.

The contemporaneous regressions do not establish a fuel-price asymmetry. Equality tests for positive and negative retail-fuel inflation give p-values of 0.659 for the perception gap, 0.527 for Q51, and 0.100 for Q61. Significance of the positive coefficient in some equations, combined with insignificance of the negative coefficient, is not enough to show that households react differently to price increases and declines. Any asymmetry interpretation remains suggestive and is kept outside the main contribution. Supplementary Table S5 reports the formal tests.

The association with Q61 becomes stronger at later horizons. In the baseline projections, the positive-fuel coefficient rises from 0.080 at h = 0 to 0.178 at h = 1, 0.285 at h = 2, 0.308 at h = 3, and 0.310 at h = 4. The corresponding HAC p-values are 0.036, 0.035, 0.014, 0.002, and 0.001. With a high-inflation control and a deterministic trend, the h = 2 to h = 4 coefficients remain positive and statistically distinguishable from zero at 0.275, 0.268, and 0.282. Alternative trend, high-inflation, and post-2020 controls give the same sign pattern, so the positive association is not tied to one baseline specification.

Small-sample uncertainty is much larger than the HAC results suggest. With four-quarter moving blocks, bootstrap p-values for the positive-fuel coefficients in the Q61 regressions are 0.399, 0.277, 0.138, 0.080, and 0.084 across h = 0,...,4; eight-quarter blocks weaken the longer-horizon evidence further. Most confidence intervals include zero. The horizon regressions thus provide a coherent predictive pattern under conventional HAC inference, but not a precisely estimated dynamic causal response in a sample of roughly eighty quarters. Supplementary Table S2 reports the four- and eight-quarter results.

State-dependent specifications offer only limited additional evidence. Fuel-price coefficients are larger in the high-inflation state in several Q51, perception-gap, and Q61 regressions, and the positive-fuel-by-high-inflation interaction is statistically significant in a number of HAC specifications. The threshold is fixed rather than estimated, the high-inflation subsample is short, and the broader block-bootstrap evidence remains fragile. These estimates are used only as a robustness check on possible amplification when inflation is already elevated, not as evidence of a separate nonlinear regime. Supplementary Table S4 reports the state-dependent specifications.

The Brent exercise adds information about the retail-fuel channel but does not solve the identification problem. Positive year-on-year Brent growth strongly predicts positive Polish retail-

fuel inflation: the preferred first-stage coefficient is 0.121 with a HAC standard error of 0.021, implying a robust first-stage statistic of about 33.1. A specification using total Brent and retail-fuel inflation is stronger, whereas an AR(1) Brent innovation is materially weaker. The preferred shifter is chosen because it maps directly into positive fuel inflation, not because it produces the largest second-stage estimate.

Reduced-form Q61 estimates are positive from h = 0 through h = 4 and statistically significant under HAC inference from h = 1 onward. The corresponding 2SLS estimates for instrumented positive fuel inflation are 0.306 at h = 1, 0.545 at h = 2, 0.476 at h = 3, and 0.522 at h = 4, with HAC p-values below 0.05 at each horizon. First-stage, reduced-form, and 2SLS signs are therefore consistent with one another. By contrast, the same exercise is much less informative for the perception gap, where the reduced-form and 2SLS coefficients are not statistically distinguishable from zero at the reported horizons.

Circular-block bootstrap results are less precise. Every 95 percent interval for the Q61 2SLS coefficient includes zero; h = 2 comes closest to conventional significance, with a bootstrap p-value of 0.069. The remaining horizons are weaker. The Brent exercise is therefore best viewed as directionally consistent quasi-IV evidence for a retail-fuel channel into forward-looking expectations, not as identification of a structural causal effect. Supplementary Table S3 reports the bootstrap estimates.

The empirical results narrow the policy question. Aggregate energy inflation is too coarse to summarize the information in household-facing energy prices. Electricity, gas, and positive fuel-price movements contain distinct information for reported perceptions. The connection with forward-looking expectations is more selective. Positive fuel inflation shows the most consistent association with subsequent Q61, and the external-shifter estimates point in the same direction, but finite-sample uncertainty remains substantial. These findings justify carrying perception information into the policy exercise, while stopping short of treating the source coefficients as structural policy parameters.

FIGURE 2. HOUSEHOLD-FACING ENERGY-PRICE COMPONENTS

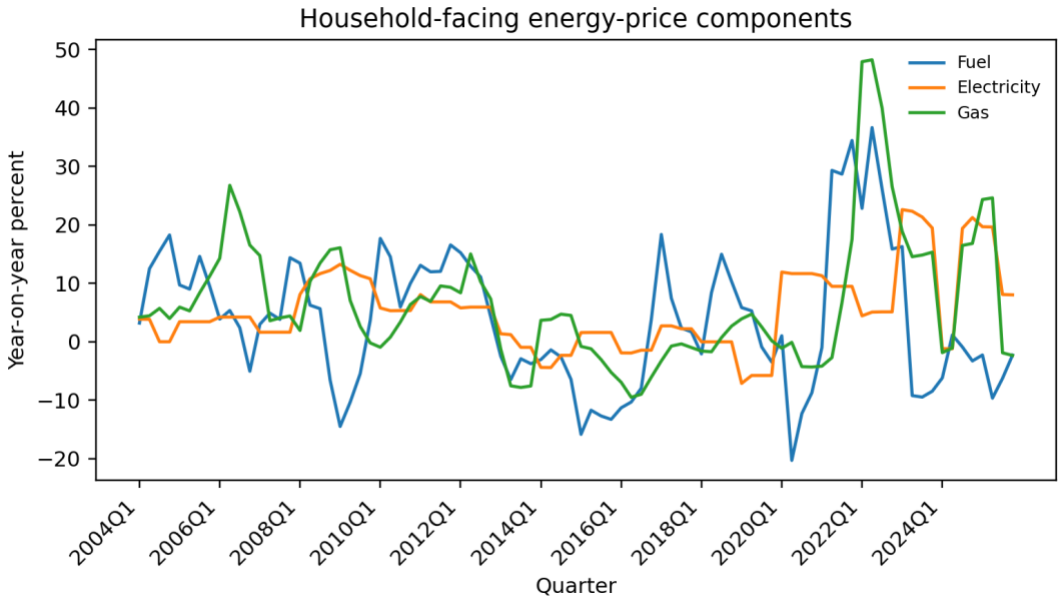


TABLE 2. SOURCE-SPECIFIC INFORMATION FOR PERCEPTIONS AND EXPECTATIONS

Outcome	Source term	Coefficient	HAC SE	p-value	N
Perception gap	Aggregate energy inflation	0.030	0.025	0.227	84
Perception gap	Electricity inflation	0.041	0.019	0.031	84
Perception gap	Gas inflation	0.026	0.012	0.031	84
Q51 perceived inflation	Positive fuel inflation	0.058	0.028	0.040	84
Q51 perceived inflation	Electricity inflation	0.056	0.024	0.022	84
Q51 perceived inflation	Gas inflation	0.060	0.022	0.006	84
Q61 expected inflation	Positive fuel inflation	0.080	0.038	0.036	84
Q61 expected inflation	Gas inflation	0.048	0.029	0.100	84

Notes: SE denotes standard error. All specifications include the relevant lagged outcome, the 2018 survey-history indicator, and the macroeconomic controls described in Section 4. The perception-gap source block is jointly significant (χ^2=21.67, p<0.001); the Q51 block is jointly significant

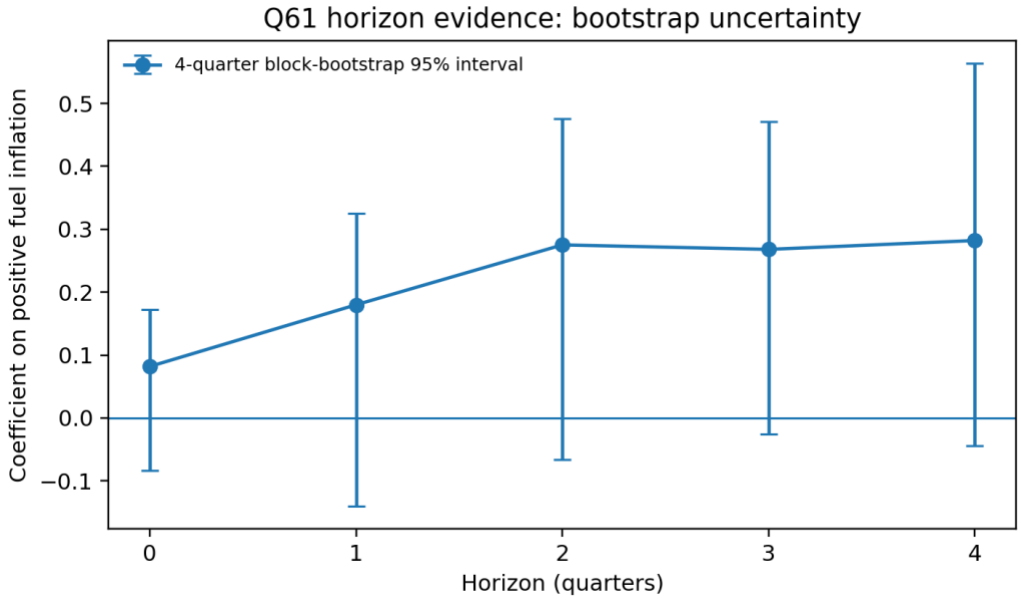
(χ^2=13.58, p=0.009); the Q61 block is marginal (χ^2=8.18, p=0.085). Formal equality tests do not establish fuel-price asymmetry at the 5% level.

TABLE 3. POSITIVE FUEL INFLATION AND FORWARD-LOOKINGEXPECTATIONS

Evidence	h	Coefficient	HAC SE	HAC p	Bootstrap p	N
Q61 robustness	0	0.082	0.045	0.071	0.399	84
Q61 robustness	1	0.180	0.083	0.031	0.277	83
Q61 robustness	2	0.275	0.095	0.004	0.138	82
Q61 robustness	3	0.268	0.075	<0.001	0.080	81
Q61 robustness	4	0.282	0.090	0.002	0.084	80
2SLS / Brent shifter	0	0.126	0.079	0.112	0.388	81
2SLS / Brent shifter	1	0.306	0.130	0.018	0.255	81
2SLS / Brent shifter	2	0.545	0.146	0.000	0.069	81
2SLS / Brent shifter	3	0.476	0.146	0.001	0.143	81
2SLS / Brent shifter	4	0.522	0.179	0.003	0.176	80

Notes: SE denotes standard error. The predictive rows use the high-inflation-plus-trend Q61 horizon specification with four-quarter moving-block-bootstrap p-values; the same matched estimates and intervals are reported in Supplementary Table S2 and Figure 3. The 2SLS rows instrument positive retail-fuel inflation with positive year-on-year Brent growth; their circular-block-bootstrap results are reported in Supplementary Table S3. Bootstrap inference is materially weaker than asymptotic HAC inference, and the 2SLS bootstrap intervals include zero at every reported horizon.

FIGURE 3. Q61 HORIZON EVIDENCE AND BLOCK-BOOTSTRAP UNCERTAINTY



Notes: Point estimates and 95% moving-block-bootstrap intervals use the high-inflation-plus-trend robustness specification with four-quarter blocks. The intervals illustrate the loss of precision relative to conventional HAC inference.

6. HISTORICAL NBP REACTION: DESCRIPTIVE EVIDENCE

The source and survey results do not show that the National Bank of Poland historically used the same information in setting rates. A separate reaction-function exercise examines that narrower descriptive question. All explanatory variables enter with a one-quarter lag, so the specification maps information available before the policy-rate observation into the subsequent quarterly policy setting. The exercise is not an identified monetary-policy rule and is not used to infer the desirable response.

The historical benchmark, Model 0 (M0), is:

$$i_t = \alpha_i + \rho_i\, i_{t-1} + \varphi_\pi\, \pi_{t-1} + \varphi_y\, \Delta y_{t-1} + \varphi_c\, \pi_{t-1}^{core} + \varphi_f\, \pi_{t-1}^{food} + \delta_i\, D_t^{2018} + \varepsilon_t^i \quad (8)$$

In equation (8), the dependent variable is the NBP reference rate and real activity is annualised quarter-on-quarter log real-GDP growth. Model 1 (M1) adds lagged Q61 expected inflation and the lagged perception gap. Model 2 (M2) retains expected inflation but replaces the aggregate gap with lagged source-specific fuel, electricity, and gas inflation. Model 3 (M3) further adds the positive-fuel-by-high-inflation interaction.

M0 fits the quarterly reference rate closely, with an adjusted R^2 of 0.950 and a root mean squared error (RMSE) of 0.417 percentage points. Adding expected inflation and the total perception gap in M1 raises the adjusted R^2 to 0.959 and lowers RMSE to 0.372; the Akaike information criterion (AIC) and Bayesian information criterion (BIC) also decline. Within M1, lagged expected inflation is positively associated with the subsequent policy rate (β=0.076, HAC p=0.010), whereas the lagged perception gap is not (β=0.019, p=0.544).

Replacing the aggregate gap with source-specific prices produces a larger improvement in descriptive fit. In M2 the adjusted R^2 reaches 0.972, RMSE falls to 0.301, AIC declines from 91.3 in M1 to 60.9, and BIC falls from 113.3 to 90.2. The source-price block is jointly significant (χ^2=28.34, p<0.001). Expected inflation remains positively associated with the policy rate, while positive fuel inflation and gas inflation also enter positively. Individual coefficients on headline, core, and food inflation change materially across these highly correlated specifications, so their signs are not assigned a structural interpretation.

Adding the positive-fuel-by-high-inflation interaction in M3 provides little additional evidence of a distinct historical state-dependent response. The interaction has p = 0.268, and the high-inflation indicator is also statistically indistinguishable from zero. AIC falls slightly, but BIC rises relative to M2 because of the additional parameters. The reaction-function exercise is descriptive: it asks whether survey and source variables improve the historical mapping of the policy-rate path. It does not show that the NBP observed the constructed variables in real time, that the coefficients represent structural preferences, or that the historical response minimized an economic loss function.

TABLE 4. HISTORICAL NBP REACTION-FUNCTION FIT

Model	Information set	Adj. R^2	AIC	BIC	RMSE	Key p-value
M0	Conventional lagged macro variables	0.950	106.584	123.683	0.417	CPI: 0.003
M1	M0 + Q61 expected inflation + perception gap	0.959	91.296	113.280	0.372	Q61: 0.010; gap: 0.544
M2	M0 + Q61 expected inflation + source-price vector	0.972	60.936	90.248	0.301	source block: <0.001
M3	M2 + positive fuel inflation × high-inflation indicator	0.973	59.272	93.469	0.291	interaction: 0.268

Notes: The exercise is descriptive and uses lagged information. AIC denotes the Akaike information criterion, BIC the Bayesian information criterion, and RMSE the root mean squared error. Fit improvements are not interpreted as evidence that the historical response was welfare-improving or optimal.

7. SEMI-STRUCTURAL POLICY MODEL

The semi-structural exercise is not a welfare-based dynamic stochastic general equilibrium (DSGE) model. Its purpose is narrower: to ask whether information in household perceptions changes stabilization performance within a common estimated linear environment. The system links measured inflation, real activity, the perception gap, forward-looking survey expectations, the policy rate, and observed source-price variables. Household utility and firm pricing are left outside the model, so the reported losses are decision metrics within the specified rule class rather than Ramsey-optimal welfare outcomes.

The NBP inflation target is 2.5 percent, and the inflation state in the equations is CPI inflation minus that target. Real activity is represented in the baseline by the percentage log deviation of real GDP from its Hodrick-Prescott (HP) trend, using the conventional quarterly smoothing parameter of 1600. The perception state is the Q51-CPI gap and the expectation state is Q61 expected inflation. For simulation, the four source-price variables are demeaned and represented as first-order autoregressive processes. The policy simulations preserve the empirical covariance matrix of the estimated equation innovations.

The estimated semi-structural system is:

$$\tilde{\pi}_t = a_\pi + \rho_\pi \tilde{\pi}_{t-1} + \kappa_x x_{t-1} + \kappa_e \pi_{t-1}^e + \kappa_i i_{t-1} + \psi_\pi' z_t + u_t^\pi \quad (9)$$

$$x_t = a_x + \rho_x x_{t-1} + \sigma_r r_{t-1}^m + \sigma_\pi \tilde{\pi}_{t-1} + u_t^x \quad (10)$$

$$g_t = a_g + \rho_g g_{t-1} + \eta_g \tilde{\pi}_t + \psi_g' z_t + \delta_g D_t^{2018} + u_t^g \quad (11)$$

$$\pi_t^e = a_e + \rho_e \pi_{t-1}^e + \eta_e \tilde{\pi}_t + \lambda_g g_t + \psi_e' z_t + \delta_e D_t^{2018} + u_t^e \quad (12)$$

$$z_{j,t} - \bar{z}_j = \rho_j^z (z_{j,t-1} - \bar{z}_j) + u_{j,t}^z, \; j \in \{f^+, f^-, el, gas\} \quad (13)$$

The real-rate proxy in the activity equation is the lagged NBP reference rate minus lagged Q61 expected inflation. The four main behavioral equations are estimated over 87 quarterly transitions from 2004Q2 to 2025Q4; the source processes use the same quarterly data subject to their lag construction. Coefficient estimates and source-process persistence are reported in Supplementary Tables S10 and S11. These are estimates from actual data; only the subsequent counterfactual policy paths and unconditional losses are simulated quantities.

Three nested policy-information sets are compared. Model 0 (M0) is the conventional inertial state-feedback rule:

$$i_t = \rho_i i_{t-1} + (1-\rho_i)[\varphi_\pi \tilde{\pi}_t + \varphi_x x_t] \quad (14)$$

Model 1 (M1) adds the observed perception gap:

$$i_t = \rho_i i_{t-1} + (1-\rho_i)[\varphi_\pi \tilde{\pi}_t + \varphi_x x_t + \varphi_g g_t] \quad (15)$$

Model 2 (M2) decomposes the observed gap into a source-implied component and a residual component:

$$i_t = \rho_i i_{t-1} + (1-\rho_i)[\varphi_\pi \tilde{\pi}_t + \varphi_x x_t + \varphi_S g_t^S + \varphi_R g_t^R] \quad (16)$$

The source-implied component is obtained from the estimated contribution of the source-price vector to the perception-gap equation, and the residual component is the remainder of the observed gap. The decomposition is generated rather than directly observed or independently validated. In the fixed-differentiation test, the coefficients are parameterized as $\varphi S = \varphi g + \delta$ and $\varphi R = \varphi g$, so $\delta = 0$ reproduces M1 without changing its conventional or aggregate-gap responses. Uncertainty in the Gate-2 source coefficients is propagated jointly using their four-lag Newey-West covariance matrix. Because the generated component has weak stand-alone predictive precision under block resampling, M2 is used only as a secondary diagnostic comparison. The construction of the source component and the treatment of coefficient uncertainty are documented in the Supplementary Material.

Policy inertia is fixed at 0.85 across all three rules. Performance is evaluated using the stabilization criterion:

$$\mathcal{L} = Var(\tilde{\pi}_t) + \lambda_x\, Var(x_t) + \lambda_i\, Var(\Delta i_t) \qquad (17)$$

The baseline loss assigns a weight of 0.5 to activity-gap variance and 0.1 to interest-rate-change variance. Since the activity gap is an empirical proxy and the weights are not derived from household utility, this is a stabilization criterion rather than a welfare loss. Initial searches repeatedly place the conventional inflation-response coefficient at the upper boundary of the admissible range. For that reason, the optimized frontier is reported only as a set of conditional loss envelopes. It is not used to identify information value, because M1 can reproduce M0 when φg=0 and M2 can reproduce M1 when φS=φR. The fixed-rule test instead holds φπ and φx at their M0 frontier values and changes only φg. Unlike the historical reaction function in Section 6, these counterfactual rules use contemporaneous state information by construction. Supplementary Table S6 reports sensitivity to the loss weights.

8. POLICY-INFORMATION RESULTS

The optimized policy frontier must be interpreted in light of model nesting. Because M1 contains M0 when φg=0, and M2 contains M1 when φS=φR, their separately minimized weak ordering is available by construction. Neither that ordering nor its frequency under reoptimization is used as evidence. The informative comparison holds the conventional inflation and activity coefficients at the M0 values reported along the frontier and varies only the perception-gap response.

The fixed-rule comparison provides directional evidence for the aggregate perception gap in the baseline Hodrick-Prescott activity-gap system. A response of φg=0.25 lowers loss at all seven tested values of φπ, with reductions ranging from 4.82 percent when φπ=1.5 to 0.34 percent when φπ=8. A response of φg=0.50 also lowers loss throughout the frontier, with reductions from 7.96 to 0.43 percent. At the baseline point φπ=3 and fixed φx=1.824, the corresponding reductions are 2.21 and 3.61 percent. These comparisons cannot be attributed to reoptimizing the conventional policy coefficients.

Parameter uncertainty supports the direction more strongly than a particular response magnitude. At φπ=3 and the fixed M0 activity response, φg=0.25 lowers loss in 99.2 percent of 1,000 accepted draws; its median reduction is 2.35 percent and its 5th percentile remains positive at 0.70 percent. For φg=0.50, the corresponding share is 97.5 percent, with a median reduction of 3.73 percent and a 5th percentile of 0.50 percent. Larger fixed responses are less secure: the lower tail becomes negative at φg=0.75 and the share with lower loss falls to 87.2 percent at φg=1.00. These are parameter-risk sensitivities with the innovation covariance held fixed, not posterior probabilities. The full uncertainty distribution is reported in Supplementary Table S9B.

The alternative activity measure places the strongest limit on the quantitative result. With demeaned annualised GDP growth, φπ=8 and φx=0 fixed at the M0 values, responses of φg=0.05 and 0.10 lower loss by only 0.0029 and 0.0038 percent. Responses of 0.25 and 0.50 raise loss slightly. The activity-gap estimates therefore do not support a common numerical response across specifications. In the GDP-growth system, only a very small positive response is admissible, and the associated gain is economically negligible. Supplementary Table S7 reports these results.

Excluding the 2021Q1-2023Q4 inflation episode provides an additional magnitude diagnostic. Under separately minimized rule classes, the loss differences fall to about 0.35 percent for M1 and 0.87 percent for M2 relative to M0. Because the models are nested, those weak inequalities are not treated as independent directional validation. The exercise instead shows that the post-pandemic episode contributes materially to the size of the baseline gains. Supplementary Table S8 reports this exclusion exercise.

Source differentiation is assessed without using the optimized ranking of the nested rules. Holding $\varphi\pi$, φx, and φg fixed, the parameterization $\varphi S=\varphi g+\delta$ and $\varphi R=\varphi g$ makes $\delta=0$ the exact M1 benchmark. At $\varphi g=0.50$, setting $\delta=0.25$ reduces loss by 0.62 percent at the point estimate and lowers loss in 98.2 percent of accepted draws after uncertainty in the generated source weights is propagated. The median incremental reduction is 0.57 percent. Other settings give the same direction, but the effect is small and the generated component is not independently validated. M2 is consequently treated as a secondary, exploratory refinement. Supplementary Tables S9A and S9C report the fixed aggregate-gap and source-differentiation comparisons.

The behavior of the optimized coefficients is another reason not to report a unique policy rule. The fixed-rule exercise identifies a limited region in which perception information lowers loss in the baseline system, but it does not identify a universally preferred coefficient. Modest positive responses to the aggregate perception gap can improve stabilization when the conventional rule is held fixed, while both the admissible response and the economic magnitude depend strongly on the activity specification. Source differentiation adds a smaller diagnostic refinement.

The policy interpretation is conditional rather than mechanical. A conspicuous source-price movement that raises current perceptions but does not reach forward-looking expectations gives little basis for a separate interest-rate response. When the perception gap adds information about expectations, monitoring it can improve the policy information set. Source decomposition matters only if it contributes beyond the aggregate gap, and even then the incremental stabilization gain is modest. Monitoring and communication are proportionate when the signal remains confined to retrospective perceptions. Greater policy attention is warranted in the model only when the same information points to broader and more persistent inflationary pressure.

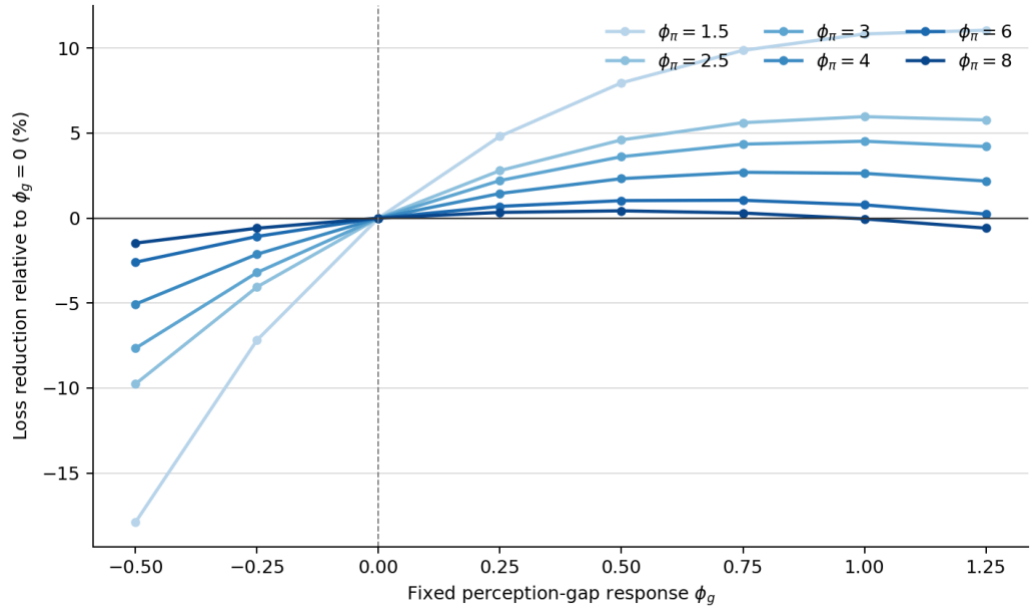
These model results are not a case for adding perceived inflation mechanically to a Taylor rule. The exercise omits a fully microfounded welfare criterion, real-time data revisions, expectations of firms and financial markets, and other channels through which energy shocks may affect the economy. What it does provide is a way to evaluate the decision value of household perception information separately from the question of whether a particular price category causally generated the perception gap.

TABLE 5. FIXED-CONVENTIONAL-RULE VALUE-OF-INFORMATION TEST

Fixed inflation response	Fixed M0 activity response	M0 loss (φg=0)	M1 loss (φg=0.25)	Reduction (%)	M1 loss (φg=0.50)	Reduction (%)
1.500	1.555	6.151	5.855	4.815	5.662	7.959
2.000	1.647	5.558	5.357	3.614	5.226	5.973
2.500	1.737	5.135	4.991	2.795	4.898	4.602
3.000	1.824	4.819	4.713	2.210	4.645	3.614
4.000	1.985	4.390	4.326	1.449	4.288	2.321
6.000	2.256	3.952	3.925	0.689	3.911	1.028
8.000	2.468	3.787	3.774	0.335	3.771	0.429

Notes: For each fixed conventional inflation response, the activity response is held at its M0 loss-minimizing value. M0 is the exact φg=0 benchmark. The M1 columns change only the perception-gap response; conventional coefficients are not reoptimized. Loss reductions are therefore not implied by nesting.

FIGURE 4. FIXED PERCEPTION-GAP RESPONSE AND MODEL-IMPLIED LOSS



Notes: Percentage reduction relative to φg=0 with the conventional inflation and activity responses held at their M0 values. The figure describes the baseline Hodrick-Prescott activity-gap system; the GDP-growth check permits only negligible gains at much smaller responses.

9. INTERPRETING INFLATION PERCEPTIONS FOR MONETARY POLICY

Household inflation perceptions are useful to policy only when they change the interpretation of the inflation outlook. They should not be treated as an additional target beside measured inflation. Their potential value is diagnostic: a perception gap may show that visible price movements are influencing household beliefs in a way that conventional indicators do not fully capture. The operational question for a central bank is whether the gap contains genuinely additional information about forward-looking expectations.

The issue is most relevant when headline inflation is being driven by a small set of visible prices. In the Polish data, aggregate energy inflation does not summarize the information contained in fuel, electricity, and gas equally well. A widening gap between Q51 and measured CPI carries more information when its source can be identified. A gap linked mainly to a temporary movement in one category need not have the same policy meaning as an equally large gap accompanied by broader price pressure and rising Q61 expectations.

The survey distinction between Q51 and Q61 is particularly useful for this purpose. When a visible price movement is reflected in retrospective perceptions but does not appear in forward-looking expectations, the evidence provides little support for a separate interest-rate response to the perception movement itself. The more proportionate response is to monitor whether the effect broadens, while using communication to explain the source and expected persistence of the price change. By contrast, when the same source movement is also associated with higher Q61 at subsequent horizons, it deserves greater attention in the assessment of inflation persistence because the signal is no longer confined to households' description of past price changes. Given the weaker block-bootstrap evidence, this should be read as a warning signal rather than as proof of persistent expectations de-anchoring.

Quantitative household surveys can serve as an additional diagnostic layer within the regular policy information set. They need not replace forecast models or generate a separate target. Q51 and Q61 can be read alongside headline inflation, core inflation, food inflation, wages, and model-based forecasts. Their value is greatest when the standard indicators send conflicting signals, for example when measured inflation is falling while household perceptions remain elevated because a visible category is still rising.

The model also shows that the marginal value of perception information depends on the strength of the conventional policy response. The gain is larger when the response to measured inflation is less aggressive and smaller when that response is stronger. This is not an argument for weakening the inflation response. It means that the contribution of perception data depends on how much of the relevant information is already captured by the existing rule and forecasting system.

Source differentiation calls for more caution. A limited additional response to the source-implied component lowers loss in most accepted parameter draws, but the incremental reduction is small and the component itself is generated. Its practical role is diagnostic. Source information may help explain why household beliefs diverge from measured inflation, but the estimates do not support a standing rule with separate interest-rate coefficients for fuel, electricity, or gas.

For policy communication, the same evidence suggests that explaining the origin of a perception gap can be as important as reporting aggregate inflation itself. When households are

reacting to prices they encounter frequently, communication that focuses only on the headline index may fail to address the experience underlying their beliefs. A central bank can acknowledge the visible source of the divergence while explaining whether it is expected to be temporary, whether it is spreading to broader inflation, and how it enters the policy outlook. This is especially relevant when survey perceptions remain high after measured inflation has begun to fall.

The empirical limitations also constrain policy use. Dynamic fuel estimates are stronger under HAC inference than under block bootstrap, Brent does not eliminate the exclusion-restriction problem, and stabilization gains are sensitive to the activity measure. Perception data should affect how much attention is paid to an inflation signal, not determine the policy decision on their own. They are most informative when they add something beyond standard inflation measures and when the same signal appears in forward-looking expectations.

10. Conclusion

Household inflation perceptions are useful for monetary-policy assessment only to the extent that they add information beyond standard inflation indicators. The Polish evidence shows why the source of the perception gap and its connection with expectations both matter. Gaps of similar size can arise from very different price configurations, so the size of the divergence alone cannot distinguish a visible relative-price movement from a broader change in inflation beliefs.

The Polish evidence points to the importance of composition. After controlling for core and food inflation, aggregate energy inflation contributes little to the perception-gap equation, while individual household-facing energy prices retain information for reported perceptions. The evidence extending to Q61 is narrower. Positive retail-fuel inflation is associated with forward-looking expectations under HAC inference, and the Brent external-shifter estimates have consistent signs. Block-bootstrap inference is less precise, so the pattern is more convincing as predictive and quasi-instrumental evidence than as a structural causal estimate.

The policy exercise reaches a similarly conditional conclusion. The optimized ranking of nested rules is not evidence in itself. With conventional policy coefficients fixed at their Model 0 values, modest responses to the perception gap lower model-implied loss across the baseline activity-gap frontier. At the baseline point, responses of 0.25 and 0.50 lower loss in 99.2 and 97.5 percent of accepted parameter draws. The size of the gain is not robust to the activity measure: the GDP-growth specification admits only negligible improvements at much smaller responses. Limited source differentiation adds a smaller effect. No unique response coefficient is identified.

For monetary-policy practice, the useful distinction is between perceptions that remain retrospective and perceptions that are accompanied by a change in expectations. The former strengthen the case for monitoring and communication. The latter deserve greater attention in the assessment of inflation persistence. Source composition helps interpret why perceptions diverge from measured inflation, but the estimated gains are too small and too specification-sensitive to justify separate reaction coefficients for individual energy prices.

The analysis is based on one country, a quarterly sample, the available survey measures, an imperfect external shifter, and a semi-structural policy environment. The results should not be

generalized mechanically to other inflation-targeting economies or model structures. Comparative work with consistent quantitative perception and expectation measures would provide a natural next test. Within the present evidence, household perceptions are best understood as an input that can sharpen the assessment of forward-looking inflation risks, not as an object of stabilization in their own right.

DATA AVAILABILITY

The public-source macroeconomic and price series used in the analysis are described in Section 3. The European Commission consumer-survey microdata are subject to the provider's access and dissemination terms. The manuscript reports the transformations required to reproduce the quarterly analysis from the underlying sources.

DECLARATION OF COMPETING INTEREST

The author declares no competing interests.

REFERENCES

Aguilar, J., & Quineche, R. (2026). Regional inflation spillovers and monetary policy design. Journal of Policy Modeling 48(2): 468-488. https://doi.org/10.1016/j.jpolmod.2025.10.003.

Aidala, F., Armantier, O., Koşar, G., Somerville, J., Topa, G., & van der Klaauw, W. (2024). Gasoline price changes and consumer inflation expectations: Experimental evidence. Journal of Economic Behavior & Organization 220: 66-80. https://doi.org/10.1016/j.jebo.2024.01.027.

Álvarez, L. J., & Sánchez, I. (2019). Inflation projections for monetary policy decision making. Journal of Policy Modeling 41(4): 568-585. https://doi.org/10.1016/j.jpolmod.2018.09.005.

Beckworth, D., & Horan, P. J. (2024). A two-for-one deal: Targeting nominal GDP to create a supply-shock robust inflation target. Journal of Policy Modeling 46(6): 1071-1089. https://doi.org/10.1016/j.jpolmod.2024.05.014.

Binder, C. C. (2018). Inflation expectations and the price at the pump. Journal of Macroeconomics 58: 1-18. https://doi.org/10.1016/j.jmacro.2018.08.006.

Bürgi, C., Srivastava, P., & Whelan, K. (2026). Expectations for gasoline prices and inflation: Evidence from households. Journal of Money, Credit and Banking 58(2): 619-630. https://doi.org/10.1111/jmcb.13260.

Bussière, M., Gilbert, J., & Grishchenko, O. (2026). The role of inflation perceptions in consumer inflation expectations: Evidence from the euro area. Finance and Economics Discussion Series 2026-038, Board of Governors of the Federal Reserve System. https://doi.org/10.17016/FEDS.2026.038.

D'Acunto, F., & Weber, M. (2024). Why survey-based subjective expectations are meaningful and important. Annual Review of Economics 16: 329-357. https://doi.org/10.1146/annurev-economics-091523-043659.

de Mendonça, H. F. (2018). Credibility and inflation expectations: What we can tell from seven emerging economies? Journal of Policy Modeling 40(6): 1165-1181. https://doi.org/10.1016/j.jpolmod.2018.06.001.

Dietrich, A. M. (2024). Consumption categories, household attention, and inflation expectations: Implications for optimal monetary policy. Journal of Monetary Economics 147: 103594. https://doi.org/10.1016/j.jmoneco.2024.103594.

European Commission. (2026a). The Joint Harmonised EU Programme of Business and Consumer Surveys: Methodological User Guide. Directorate-General for Economic and Financial Affairs, 6 March 2026.

European Commission. (2026b). Methodological concepts: Business and consumer surveys. Directorate-General for Economic and Financial Affairs. May 2026 update on the Polish consumer survey provider and historical data. https://economy-finance.ec.europa.eu/economic-forecast-and-surveys/business-and-consumer surveys/methodology-business-and-consumer-surveys/methodological-concepts_en.

Grajales-Olarte, A., Hamann, F., Naranjo-Saldarriaga, S., & Pulido, J. (2026). Policy implications of losing credibility: Lessons from Colombia's post-pandemic inflation. Journal of Policy Modeling 48(4): 107064. https://doi.org/10.1016/j.jpolmod.2026.107064.

Hahn, V., & Marenčák, M. (2026). Inflation perceptions and monetary policy. SSRN Working Paper, revised 30 July 2026. https://doi.org/10.2139/ssrn.5037474.

Jo, Y. J., & Klopack, B. (2025). Fueling expectations: The causal impact of gas prices on inflation expectations. Journal of Monetary Economics 153: 103797. https://doi.org/10.1016/j.jmoneco.2025.103797.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. American Economic Review 95(1): 161-182. https://doi.org/10.1257/0002828053828518.

Künsch, H. R. (1989). The jackknife and the bootstrap for general stationary observations. The Annals of Statistics 17(3): 1217-1241.

Montiel Olea, J. L., & Plagborg-Møller, M. (2021). Local projection inference is simpler and more robust than you think. Econometrica 89(4): 1789-1823. https://doi.org/10.3982/ECTA18756.

Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55(3): 703-708.

Politis, D. N., & Romano, J. P. (1992). A circular block-resampling procedure for stationary data. In R. LePage & L. Billard (Eds.), Exploring the Limits of Bootstrap (pp. 263-270). New York: Wiley.

Rodriguez-Rodriguez, F. J., Hierro, L. A., & Garzon, A. J. (2024). Fed and ECB reaction functions during quantitative easing: Three phases of monetary policy, both conventional and unconventional. Journal of Policy Modeling 46(5): 928-945. https://doi.org/10.1016/j.jpolmod.2024.03.003.

Roy Trivedi, S., & Ghosh, S. (2025). Do deeds match words? India's monetary policy needs to 'walk the talk' for inflation anchoring. Journal of Policy Modeling 47(3): 662-680. https://doi.org/10.1016/j.jpolmod.2024.06.002.

Vatsa, P., Pino, G., & Clements, A. (2025). Gasoline prices, gasoline price expectations, and inflation expectations in the United States. Energy Economics 146: 108508. https://doi.org/10.1016/j.eneco.2025.108508.

Weber, M., D'Acunto, F., Gorodnichenko, Y., & Coibion, O. (2022). The subjective inflation expectations of households and firms: Measurement, determinants, and implications. Journal of Economic Perspectives 36(3): 157-184. https://doi.org/10.1257/jep.36.3.157.