Commodity Terms of Trade and the Informativeness of Inflation Expectations in Small Open Economies

Fernando Letelier
Rutgers University

August 2026

Abstract

What inflation expectations are most informative for central banks in open economies? Using a quarterly panel of twelve open economies from 2000 to 2025, we compare one year ahead inflation expectations from non-expert agents and expert forecasters. Unlike recent U.S. evidence, experts are more accurate in both demand and supply driven inflation episodes. The paper's main result is that this expert advantage declines systematically with commodity terms of trade exposure during global supply episodes. The narrowing reflects lower non-expert forecast errors rather than deteriorating expert performance. Fair–Shiller encompassing regressions further show that the non-expert/expert forecast gap becomes increasingly informative about subsequent inflation as exposure rises, with the incremental information most clearly reflected in future services inflation. Contemporaneous exchange rate movements do not account for the exposure gradient. A parsimonious small open economy New Keynesian model shows that this informational channel, rather than commodity transmission alone, narrows the difference between optimal policy responses across demand and supply regimes.

1 Introduction

Inflation expectations are central to monetary policy analysis, yet central banks observe several competing measures: household and business surveys, professional forecasts, and market based indicators. A growing literature argues that their relative usefulness is state dependent. In the United States, the demand/supply decomposition of Shapiro (2026) provides an operational measure of the source of inflation, and recent evidence for the United States, which we document in Letelier and Luna (2026), suggests that non-expert expectations can become relatively more informative when inflation is demand driven. This matters for policy because demand and supply disturbances generate different trade offs between inflation and activity Hofmann et al. (2024).

This paper asks whether that logic generalizes to open economies. The key difference is that a common global disturbance can have very different domestic incidence depending on economic structure. A global commodity disturbance may raise imported costs everywhere, while in commodity exposed economies it can also be associated with movements in export income, exchange rates, fiscal revenues, activity, wages, and spending. The relevant state may therefore be not only whether global inflation is demand or supply driven, but also how commodity terms of trade exposure transforms that disturbance when it reaches the domestic economy.

We assemble a quarterly panel of twelve open economies over 2000–2025: Albania, Australia, Brazil, Chile, Colombia, Iceland, New Zealand, Norway, Peru, Serbia, South Africa, and the United Kingdom. For each country we match one year ahead inflation expectations from non-expert agents (households, unions, or businesses), with forecasts from professional economists or financial analysts. We combine these series with realized inflation, macroeconomic controls, and the IMF's country specific commodity terms of trade indices (Gruss and Kebhaj, 2019).

We classify inflation regimes in three ways. The first is a global measure based on the U.S. decomposition of Shapiro (2026). The other two are country specific classifications obtained from structural VARs in real GDP growth and CPI inflation, identified with sign restrictions and Blanchard–Quah long-run restrictions. The two domestic classifications agree with each other in roughly 80 percent of country-quarters, while each agrees with the global classification in only about 40 percent. This distinction is itself informative: the global shock mix and its domestic incidence are not the same object in an open economy.

The first result is a *no-reversal* benchmark. Experts are more accurate than non-experts in both demand and supply episodes under all three regime classifications. Under the global classification, the expert advantage in mean absolute error is 0.33 percentage points in demand episodes and 0.44 percentage points in supply episodes, and the difference between the two is not statistically significant. The U.S. style ranking reversal therefore does not generalize mechanically to the open economy panel.

The central result is instead that the size of the expert advantage varies systematically with commodity terms of trade exposure. During global supply episodes, a one standard deviation increase in exposure reduces the expert accuracy advantage by approximately 0.13–0.19 percentage points across the main specifications. The baseline estimate is -0.179. It survives quarter fixed effects, controls for emerging market status and inflation volatility, and leave one country out re-estimation, remaining significant in eleven of twelve cases. The gradient is already negative outside the 2020–2023 inflation surge, at -0.107, and becomes substantially larger during the surge, reaching an implied -0.374. We interpret this as episode specific amplification, not as a mechanical relationship with the size of commodity price movements.

The narrowing also has a clear decomposition. The supply exposure gradient in non-expert absolute forecast errors is -0.200 ($p = 0.016$), whereas the corresponding expert gradient is -0.021 ($p = 0.64$). Thus

$$-0.200 - (-0.021) = -0.179,$$

which reproduces the baseline relative loss coefficient. Greater exposure is associated with a reduction in the non-expert forecast disadvantage, not with a systematic deterioration in professional forecast accuracy.

Standalone accuracy and incremental information are different objects. Following the logic of Fair–Shiller forecast encompassing regressions (Fair and Shiller, 1990), we ask whether deviations of non-expert expectations from expert forecasts contain information about subsequent inflation conditional on the expert forecast. The pooled coefficient on the interaction of the forecast gap, the global supply regime, and commodity exposure is 0.215, significant at the 1 percent level. A more flexible specification with quarter fixed effects and the full set of non collinear lower order interactions yields a similar point estimate, 0.199, although with lower precision.

Component level regressions provide evidence on where this incremental information subsequently appears. The exposure dependent forecast gap coefficient is 0.164 for standardized services inflation and remains 0.176 after controlling for lagged services inflation, both significant at the 5 percent level. The corresponding standardized coefficients for food and energy are small and imprecise. Because the surveys forecast aggregate inflation, this does not imply that non-experts directly forecast services better. It shows where the predictive information contained in the aggregate forecast gap is most visible.

A natural alternative explanation is exchange rate pass through. Controlling for contemporaneous NEER growth leaves the baseline relative accuracy coefficient essentially unchanged, from -0.179 to -0.178, while controlling for REER growth reduces it to -0.148, still statistically significant. These are not causal mediation tests, but they show that contemporaneous currency movements do not mechanically summarize the cross country exposure gradient.

We embed the evidence in a parsimonious small open economy New Keynesian model with heterogeneous inflation expectations. The model separates physical commodity transmission from an informational channel disciplined by the forecast encompassing evidence. Physical transmission alone does not make optimal demand and supply regime policy responses converge with exposure. Under the conservative informational mapping, the ratio of the optimal demand to supply regime inflation coefficients falls from 1.16 at low exposure to 1.11 at the high exposure anchor. The associated policy loss effects are modest, so we interpret the model as illustrating the mechanism rather than as delivering a welfare value of non-expert information.

The policy implication is deliberately narrow. Professional forecasts remain the stronger standalone benchmark. Non-expert surveys matter because their incremental information can vary systematically with the external state. For highly commodity exposed economies, disagreement between non-experts and professionals can therefore provide a useful cross check during global supply episodes, especially when domestic propagation is strong.

The paper contributes to three literatures. First, it adds to work on heterogeneous expectations and information rigidity (Coibion and Gorodnichenko, 2015; Cornand and Hubert, 2022; Weber et al., 2022) by showing that relative and incremental forecast informativeness varies with an observable open economy characteristic. Second, it contributes to work on demand/supply inflation decompositions and state dependent policy (Davig and Leeper, 2007; Hofmann et al., 2024; Shapiro, 2026) by distinguishing the global shock mix from its domestic incidence. Third, it adds an informational margin to the literature on commodity and terms of trade transmission (Drechsel and Tenreyro, 2018; Fernández et al., 2017; Gali and Monacelli, 2005; Gruss and Kebhaj, 2019).

2 Related Literature

The paper is closest to work comparing how households, firms, and professional forecasters form expectations and use information. Professional forecasts typically outperform household expectations in standalone prediction (Verbrugge and Zaman, 2021), while information rigidities and salient personal experiences generate systematic heterogeneity across agents (Coibion and Gorodnichenko, 2015; Weber et al., 2022). Our distinction between standalone accuracy and conditional information follows the forecast encompassing logic of Fair and Shiller (1990).

Micro evidence makes the local information mechanism plausible. Using Chilean firm level data, Albagli et al. (2022) show that firms use input price changes observed along their supply chains when forming aggregate inflation expectations. Wehrhöfer (2026) finds that households exposed to electricity price increases revise inflation expectations upward and move away from professional forecasts. These studies show how locally visible prices and costs can generate disagreement across agents. Our contribution is to ask whether that disagreement contains predictive information, and whether its content varies with an open economy exposure measure.

Our component results also relate to evidence on domestic services price dynamics and sector specific expectations. Bobeica et al. (2024) show that selling price expectations of euro area services firms contain leading information about services inflation. Our exercise differs because we observe aggregate inflation expectations rather than services price expectations: the result is that the part of aggregate non-expert expectations not summarized by expert forecasts becomes more informative about future services inflation as commodity exposure rises during global supply episodes.

Finally, the paper connects the heterogeneous expectations literature to open economy commodity transmission. Commodity price and terms of trade shocks are important drivers of output, inflation, exchange rates, and external balances in small open economies (Drechsel and Tenreyro, 2018; Fernández et al., 2017; Gruss and Kebhaj, 2019). We add that economic structure can also affect which agents' expectations contain useful information about the domestic propagation of those shocks.

3 Empirical Framework

Let NE denote non-expert agents and EX expert forecasters. Forecast errors are origin dated:

$$FE^{a}_{i,t} = \pi_{i,t+4} - E^{a}_{i,t}\pi_{i,t+4}, \qquad a \in \{NE, EX\}. \tag{1}$$

We summarize relative standalone accuracy with

$$DL_{i,t} = |FE^{NE}_{i,t}| - |FE^{EX}_{i,t}|, \tag{2}$$

so $DL_{i,t} > 0$ denotes an expert accuracy advantage.

Let $S_t = 1$ denote a global supply driven episode and let χ_i be standardized commodity terms of trade exposure. The central relative accuracy specification is

$$DL_{i,t} = \alpha_i + \delta_t + \beta(S_t \times \chi_i) + X'_{i,t}\Gamma + \varepsilon_{i,t}. \tag{3}$$

Country fixed effects absorb the level of exposure and other permanent country differences. Quarter fixed effects absorb shocks common to all countries, including the level of S_t. Identification therefore comes from whether the cross country exposure gradient in relative forecast performance changes specifically during global supply episodes. The open economy hypothesis predicts $\beta < 0$.

To identify which forecast accounts for changes in relative loss, we estimate the same specification separately for $L^{a}_{i,t} = |FE^{a}_{i,t}|$. With the same sample and regressors,

$$\beta = \beta_{NE} - \beta_{EX}. \tag{4}$$

Relative accuracy does not establish whether the less accurate forecast contains information absent from the more accurate forecast. Define

$$Gap_{i,t} = E^{NE}_{i,t}\pi_{i,t+4} - E^{EX}_{i,t}\pi_{i,t+4}. \tag{5}$$

A conventional Fair–Shiller regression using both forecasts is exactly equivalent to a regression containing the expert forecast and $Gap_{i,t}$; the coefficient on the gap equals the coefficient on the non-expert forecast in the conventional parameterization. A nonzero gap coefficient therefore indicates incremental information conditional on the expert forecast. Its sign need not be positive: a negative coefficient can represent contrarian information.

Our pooled test allows the conditional content of both the expert forecast and the forecast gap to vary with the global regime and exposure:

$$\pi_{i,t+4} = \alpha_i + \delta_t + \theta(S_t\chi_i) + \sum_{j=1}^{4} \beta_j Z_{j,i,t}^{EX} + \sum_{j=1}^{4} \gamma_j Z_{j,i,t}^{Gap} + \varepsilon_{i,t+4}, \tag{6}$$

where

$$Z_{j,i,t}^{EX} \in \left\{ E_{i,t}^{EX}\pi_{i,t+4}, E_{i,t}^{EX}\pi_{i,t+4}S_t, E_{i,t}^{EX}\pi_{i,t+4}\chi_i, E_{i,t}^{EX}\pi_{i,t+4}S_t\chi_i \right\},$$

and $Z_{j,i,t}^{Gap}$ is defined analogously using $Gap_{i,t}$. The coefficient of interest is γ_4, on $Gap_{i,t} \times S_t \times \chi_i$. The exposure gradient in gap informativeness is γ_3 in demand episodes and $\gamma_3 + \gamma_4$ in supply episodes. Under the preferred local information interpretation, the directional prediction is

$$\gamma_4 > 0. \tag{7}$$

This prediction concerns the change in the exposure gradient; it does not require the gap coefficient to be positive at every exposure level.

For inflation component k, we estimate the same fully interacted specification with $\pi_{i,t+4}^k$ as the dependent variable. The corresponding triple interaction coefficient, γ_{4k}, identifies where the exposure dependent incremental information in aggregate expectations is subsequently most visible. It does not identify the microeconomic signal originally observed by non-experts.

4 Data and Regime Measurement

Our matched panel contains twelve open economies at the quarterly frequency over 2000–2025. Non-expert expectations come from households, unions, or businesses depending on country level availability; expert expectations come from professional economists or financial analysts. All direct comparisons are restricted to country-quarters in which both forecasts are observed, yielding 935 observations under the global regime measure and 851 under the domestic measures.

We combine the surveys with headline CPI inflation, unemployment, monetary policy rates, nominal and real effective exchange rates, and the IMF country specific commodity terms of trade indices of Gruss and Kebhaj (2019). We also use harmonized inflation components from the Central Bank of Chile's Global Inflation Dashboard of Bajraj et al. (2023), including goods, services, food, energy, and the available tradable/non-tradable services decomposition.

We use three regime measures. First, d_t^{US} applies the Shapiro (2026) classification of U.S.

inflation to all countries as an indicator of the global shock mix. Second, for each country we estimate a VAR(4) in quarterly real GDP growth and CPI inflation with seasonal and 2020 outlier controls. Sign restrictions identify demand shocks as innovations that move output and inflation in the same direction and supply shocks as innovations that move them in opposite directions, using the median target rotation of Fry and Pagan (2011). Third, a Blanchard–Quah identification imposes that demand shocks have no permanent effect on the level of output (Blanchard and Quah, 1988). A domestic country-quarter is classified as demand driven when demand shocks account for the majority of trailing four quarter inflation. Appendix A provides the construction details.

Table 1: Data coverage, regime frequencies, and agreement between regime measures

Country	N	First year	EM	Exposure χ_i	UNCTAD dep. (%)	Share of demand quarters d^{US}	d^{sign}	d^{BQ}
Norway	90	2002	No	2.16	88.2	0.34	0.53	0.86
Iceland	52	2012	No	1.79	84.5	0.38	0.83	0.85
Chile	79	2005	Yes	1.25	82.8	0.35	0.71	0.72
Serbia	64	2009	Yes	0.70	33.2	0.41	0.66	0.92
Peru	92	2002	Yes	0.62	91.9	0.34	0.72	0.76
Colombia	99	2000	Yes	0.54	80.1	0.32	0.65	0.65
Australia	86	2002	No	0.50	91.6	0.33	0.72	0.81
New Zealand	100	2000	No	0.34	81.0	0.33	0.52	0.65
Albania	57	2010	Yes	0.23	32.1	0.42	0.67	0.63
South Africa	98	2000	Yes	0.23	62.7	0.32	0.56	0.72
Brazil	41	2014	Yes	0.23	75.9	0.39	0.63	0.73
UK	77	2004	No	0.04	35.2	0.36	0.75	0.71

Agreement between measures, matched sample:
d^{US} vs d^{sign}: 0.407 d^{US} vs d^{BQ}: 0.405 d^{sign} vs d^{BQ}: 0.773

Notes: Matched sample. Exposure χ_i: s.d. of net-export commodity terms-of-trade growth. UNCTAD dep.: commodities as % of merchandise exports, 2022–2024. Shares of demand quarters under the global (d^{US}) and domestic (d^{sign}, d^{BQ}) regime measures.

The two domestic classifications agree with one another in approximately 80 percent of country-quarters; each agrees with the global classification in only about 40 percent. We therefore interpret the global shock mix and its domestic incidence as related but distinct objects.

Our baseline exposure measure is

$$\chi_i = sd(\Delta \log CTOT_{i,t}) \,, \tag{8}$$

standardized across countries in the regression analysis. We separately construct time varying commodity terms of trade innovations $z_{i,t}$ as residuals from country specific autoregressive processes with seasonal controls. Because these processes are estimated using the available sample, we interpret the residuals as empirical measures of unexpected terms of trade movements rather than as literal real time signals observed by survey respondents.

Inference uses Driscoll–Kraay standard errors with seven lags, allowing for serial correlation from the overlapping four quarter forecast horizon and general cross sectional dependence. Key specifications are also evaluated with alternative bandwidths, two way clustering, quarter fixed effects, and leave one country out re-estimation.

5 Empirical Results

5.1 Experts remain more accurate across regimes

Table 2 establishes the first result. Experts are more accurate than non-experts in every demand/supply cell under all three regime classifications. Under the global measure, the mean absolute error advantage is 0.33 percentage points in demand quarters and 0.44 percentage points in supply quarters. Non-experts also tend to overpredict inflation, while expert forecast errors are smaller and not systematically different from zero.

Table 2: Forecast accuracy by agent and inflation regime

Regime	N	RMSE		MAE		Bias	
		Non-Experts	Experts	Non-Experts	Experts	Non-Experts	Experts
Panel A: global measure (d^{US})							
Demand	328	2.84	2.54	1.98	1.65	-0.24	0.52
Supply	607	2.64	2.15	1.90	1.46	-0.77	0.12
Panel B: domestic, sign restrictions (d^{sign})							
Demand	555	2.65	2.37	1.90	1.56	-0.37	0.44
Supply	296	2.78	2.32	1.98	1.56	-0.75	0.12
Panel C: domestic, Blanchard–Quah (d^{BQ})							
Demand	645	2.64	2.30	1.87	1.53	-0.56	0.25
Supply	206	2.87	2.52	2.11	1.66	-0.32	0.56

Notes: RMSE, MAE, and bias, defined as the mean of $\pi_{t+4} - E_t\pi_{t+4}$, for non-experts and expert one-year-ahead forecasts. The sample is restricted to country-quarters with both non-experts and expert forecasts available.

Country fixed effect regressions using the loss differential in equation (2) produce the same conclusion: the demand minus supply contrast in relative forecast loss is small and statistically insignificant under all three classifications. The no reversal result survives alternative Driscoll–Kraay lags, exclusion of 2020–2023, alternative timing conventions, stricter regime thresholds, squared loss, two way clustering, and leave one country out estimation. Selected robustness results are reported in Appendix B.1.

5.2 Commodity exposure narrows the expert advantage

Before turning to forecast performance, we first document that the domestic inflationary incidence of commodity terms of trade shocks differs sharply between commodity exporters and non-exporters. We construct country specific innovations from the IMF net export commodity terms of trade index. The monthly index is first aggregated to the quarterly frequency, quarterly growth is computed as $100\Delta\log(CTOT_{i,t})$, and its predictable component is removed using a country specific AR(4) with seasonal controls. The residual, $z_{i,t}$, is our commodity terms of trade innovation. Hence, a one unit positive innovation corresponds approximately to an unexpected 1 percent improvement in the commodity terms of trade.

Column (1) of Table 3 shows that, for non-commodity exporters, a favorable innovation of this magnitude is associated with a 1.651 percentage point decline in annual inflation four quarters later. The economic magnitude of a one percent innovation, however, differs substantially across countries because their commodity terms of trade have very different volatilities. For example, the standard deviation of quarterly CToT growth is only about 0.23 percent in Albania, a non-exporter, so a one percent innovation represents a very large movement relative to its usual CToT fluctuations. By contrast, the corresponding standard deviation is about 1.25 percent in Chile, a commodity exporter, making a one percent innovation roughly comparable to a typical quarterly CToT movement.

For commodity exporters, the estimated differential response is +1.670, implying a net inflation response of only

$$-1.651 + 1.670 = 0.020,$$

which is statistically indistinguishable from zero. Thus, using Albania and Chile as illustrative examples, the pooled estimates imply that the same one percent favorable CToT innovation is associated with a substantial decline in inflation for a non-exporter such as Albania, but essentially no change in inflation for an exporter such as Chile. These are responses implied by the pooled exporter status regression rather than country specific estimates.

This contrast is consistent with two offsetting forces in commodity exporting economies. A favorable commodity terms of trade shock can generate a disinflationary imported cost effect, for example through an appreciation of the domestic currency that lowers the local currency price of imported goods and inputs. For commodity exporters, however, the same shock also raises export income and can stimulate domestic activity, spending, wages, and price pressures. The imported-cost and export income channels therefore work in opposite directions. We do not interpret the regression as a causal decomposition of these mechanisms. Rather, it establishes the empirical fact relevant for our analysis: the domestic inflationary incidence of the same favorable external commodity disturbance differs sharply with a country's export structure.

This motivates our main question: does economic structure also affect the relative informativeness of different agents' expectations? Our hypothesis is that the relative forecasting performance associated with commodity exposure changes specifically during global supply episodes.

Figure 1 provides the raw cross country evidence. The vertical axis is the supply minus demand difference in the non-expert forecast disadvantage. Its correlation with commodity terms of trade exposure is -0.53: countries with greater exposure tend to experience a larger narrowing of the expert advantage during global supply episodes.

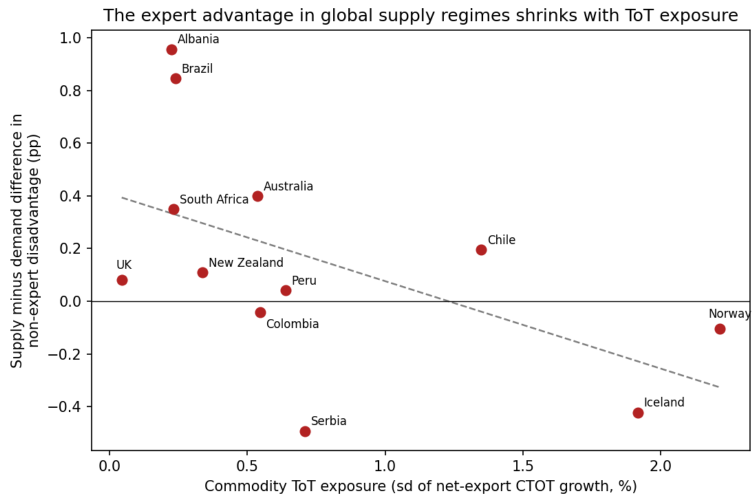

Figure 1: Expert advantage in global supply regimes and commodity exposure.
The vertical axis is the non-expert forecast disadvantage (MAE of non-experts minus MAE of experts) in global supply regimes minus the corresponding disadvantage in global demand regimes. The correlation with commodity terms of trade exposure is -0.53.

Table 3 formalizes this relationship. Across the baseline, quarter fixed effect, and horse race specifications, a one standard deviation increase in exposure reduces the expert accuracy advantage by approximately 0.13–0.19 percentage points during global supply episodes. The baseline estimate is -0.179. With quarter fixed effects, identification comes from cross country differences in exposure within the same global quarter. The estimate is also stable to leaving out each country in turn, ranging from -0.141 to -0.268 and remaining statistically significant in eleven of twelve re-estimations.

Replacing the price based exposure measure with UNCTAD commodity dependence produces a smaller and statistically insignificant coefficient. This suggests that the result is more closely associated with realized commodities terms of trade exposure than with export concentration alone. Allowing the exposure gradient to differ during the 2020–2023 inflation surge yields an estimate of -0.107 outside the surge and an additional -0.267 during it, implying a gradient of -0.374 in 2020–2023. The difference between the two periods is statistically significant ($p = 0.025$). We interpret this as episode specific amplification rather than as evidence that unusually large commodity shocks are required for the mechanism.

The decomposition in equation (4) clarifies what drives the narrowing in relative accuracy. For non-experts,

$$\widehat{\beta}_{NE} = -0.200 \qquad (\text{s.e.} = 0.082,\ p = 0.016),$$

whereas for experts,

$$\widehat{\beta}_{EX} = -0.021 \qquad (\text{s.e.} = 0.045,\ p = 0.64).$$

Table 3: Commodity exposure and the supply-regime expert advantage

	(1) Inflation response	(2) Baseline	(3) Time FE	(4) Horse race	(5) UNCTAD dep.	(6) Episode size
ToT shock	-1.651**					
	(0.759)					
ToT shock × commodity exporter	1.670**					
	(0.788)					
Supply × exposure χ		-0.179**	-0.187**	-0.134***		-0.107*
		(0.071)	(0.063)	(0.046)		(0.054)
Supply × emerging market				0.312		
				(0.282)		
Supply × inflation volatility				-0.247		
				(0.155)		
Supply × UNCTAD dependence					-0.040	
					(0.095)	
Supply × exposure × 2020–23						-0.267**
						(0.118)
Observations	935	935	934	935	935	935
Exporter net effect	0.020					
s.e.	0.105					
p-value: gradient outside 2020–23 = 0						0.051
Gradient during 2020–23						-0.374
s.e.						0.127
p-value: equal gradients						0.025

Notes: In columns (2)–(6), the dependent variable is $|FE_{i,t}^{NE}| - |FE_{i,t}^{E}|$, so positive values indicate larger non-experts forecast errors and a greater expert advantage. Column (1) estimates the response of realized inflation to terms-of-trade shocks and their interaction with commodity-exporter status. Columns (2)–(6) estimate how commodity exposure affects the non-experts–expert loss differential in supply-driven episodes, where supply is defined as $1 - d_t^{US}$. Exposure χ_i is standardized, so coefficients are per one standard deviation. All specifications include country fixed effects; column (3) also includes quarter fixed effects. Standard errors are reported in parentheses. Because exposure is time invariant, country fixed effects absorb its level; the interaction identifies how the exposure gradient differs in supply relative to demand episodes.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Thus,

$$-0.200 - (-0.021) = -0.179,$$

which reproduces the baseline relative loss coefficient up to rounding. Greater exposure therefore narrows the expert advantage because non-expert forecast errors decline, not because professional forecasts systematically deteriorate.

Taken together, these results establish two distinct forms of heterogeneous domestic incidence. Commodity exporting economies exhibit a different inflation response to favorable terms of trade innovations, and greater commodity exposure is associated with a smaller expert forecasting advantage during global supply episodes. The next question is whether this narrowing reflects only relative accuracy or whether non-expert expectations also contain information about future inflation that is absent from the expert forecast.

5.3 Incremental information in the forecast gap

Table 4 separates standalone accuracy from incremental information. In the descriptive cells, the gap coefficient is statistically indistinguishable from zero in global demand episodes. In global supply episodes it is negative at low exposure (-0.292) and positive but imprecise at high exposure (0.282). This is consistent with the fact that incremental information can enter with either sign.

The pooled interaction is the sharper test. The coefficient on

$$Gap_{i,t} \times S_t \times \chi_i$$

is 0.215, significant at the 1 percent level, while the corresponding expert forecast interaction is small and statistically insignificant. The positive sign matches the directional prediction in equation (7): during global supply episodes, the conditional predictive content of the non-expert/expert forecast gap becomes increasingly positive as commodity exposure rises.

Put differently, conditional on the expert forecast, a larger positive gap between non-expert and expert expectations becomes increasingly associated with higher subsequent inflation in more commodity exposed economies. This does not imply that non-experts become more accurate than experts on average. Rather, their disagreement with experts contains additional state and exposure dependent information.

A fully flexible specification with quarter fixed effects and all non-collinear lower order interactions yields a similar point estimate of 0.199, significant at the 10 percent level. When the gradient is allowed to differ during 2020–2023, the estimate outside the surge is 0.051 and imprecise, while the implied gradient during the surge is 0.345 (s.e. = 0.055).

Table 4: Forecast Encompassing and Incremental Information

	Cell estimates				Pooled Fair–Shiller		
	(1) Demand-Lo	(2) Demand-Hi	(3) Supply-Lo	(4) Supply-Hi	(5) Baseline FS	(6) Flexible FS	(7) 2020–23
Expert forecast	0.751**	0.542***	0.735***	0.579***	0.623***	0.623***	0.632***
	(0.334)	(0.124)	(0.148)	(0.099)	(0.150)	(0.109)	(0.121)
Forecast gap (NE–EX)	-0.024	-0.370	-0.292**	0.282	-0.135	-0.112	-0.091
	(0.107)	(0.336)	(0.144)	(0.227)	(0.106)	(0.097)	(0.082)
Gap × Supply					0.084	-0.020	-0.196
					(0.145)	(0.155)	(0.132)
Gap × Supply × exposure					0.215***	0.199*	0.051
					(0.067)	(0.107)	(0.126)
Expert × Supply					0.005	-0.124	-0.162
					(0.134)	(0.097)	(0.116)
Expert × Supply × exposure					-0.019	0.054	-0.093
					(0.034)	(0.152)	(0.061)
Gap × Supply × exposure × 2020–23							0.294**
							(0.145)
Expert × Supply × exposure × 2020–23							0.029
							(0.060)
Implied gap gradient, 2020–23							0.345***
							(0.055)
Observations	160	168	299	308	935	935	935
Country fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Quarter fixed effects	No	No	No	No	No	Yes	No
Full lower-order interactions	No	No	No	No	No	Yes	No

Notes: The dependent variable is realized headline inflation, $\pi_{i,t+4}$. Define $Gap_{i,t} = NE_{i,t} - EX_{i,t}$, where NE and EX denote non-expert and expert one-year-ahead inflation forecasts. Conditional on the expert forecast, a nonzero coefficient on the forecast gap indicates that non-expert expectations contain incremental information about subsequent inflation. The sign of the coefficient indicates the direction of that conditional information. Columns (1)–(4) estimate the Fair–Shiller encompassing regression separately by global inflation regime and a median split of commodity terms-of-trade exposure. Column (5) is the exact EX–Gap reparameterization of the conventional pooled Fair–Shiller specification used in the baseline analysis. Consequently, the coefficient on Gap × Supply × exposure is numerically identical to the coefficient on NE × Supply × exposure in the equivalent EX–NE representation. Column (6) includes quarter fixed effects and the full set of lower-order interactions required to allow both the expert forecast and the forecast gap to vary flexibly with the global supply regime and standardized commodity exposure. The gap gradient remains similar in magnitude (0.199), while the corresponding expert-forecast gradient is small and statistically insignificant (0.054). Column (7) allows the exposure gradient in the baseline specification to differ during the 2020–2023 inflation surge. The four-way coefficients report the additional surge effects, while the implied gap gradient during the surge sums the three-way and four-way terms. Once this episode-specific heterogeneity is allowed for, the estimated gap gradient outside 2020–2023 is 0.051 and statistically indistinguishable from zero, whereas the implied gradient during the surge is 0.345 and highly significant. All specifications include country fixed effects. Standard errors are Driscoll–Kraay with seven lags and are reported in parentheses. The coefficients are forecast-encompassing coefficients, not constrained forecast-combination weights. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

We therefore interpret 2020–2023 as a period in which the informational mechanism was especially pronounced, not as evidence that the mechanism is necessarily absent outside that episode.

5.4 Where does the incremental information appear?

Table 5 repeats the encompassing exercise using future inflation components as dependent variables. The specification is the flexible one of Table 4, with country and quarter fixed effects and both forecasts fully interacted with the supply indicator and exposure, so identification again comes from cross country differences in exposure within the same global quarter. Because components differ in volatility, the comparison across rows uses the standardized column, in which coefficients measure standard deviations of future component inflation per unit of forecast disagreement.

The clearest estimate is for services: 0.164, significant at the 5 percent level. The corresponding coefficients are essentially zero for food (0.001) and small and insignificant for energy (0.050) and goods (0.066). The non tradable services point estimate exceeds the tradable services estimate, but their difference is too imprecise to support a formal distinction. Since these components are estimated on a smaller harmonized sample, the headline row is included as an internal benchmark rather than as an estimate of the aggregate gradient, which is the one reported in Table 4.

Two features of this pattern are worth noting. First, it is not what a simple salience account would predict. If non-experts merely observed commodity linked prices earlier than professional forecasters, the incremental information should appear in food and energy inflation, where the estimates are closest to zero. Second, it does not appear to be extrapolation of recent inflation in the same component: controlling for the one quarter lag of each component leaves the services estimate intact, rising slightly to 0.176 and remaining significant at the 5 percent level.

The exercise identifies the *destination* of the incremental information, not its source. The surveys forecast aggregate inflation, so the results do not imply that non-experts directly forecast services better than experts. The original signal could concern salient commodity prices, input costs, wages, income, activity, or spending; its predictive content can subsequently appear in a different part of the inflation basket. Identifying which signal is at work would require expectations disaggregated by sector or region, which the cross country panel does not provide.

Table 5: Incremental Information in Non-Expert Expectations by Inflation Component

	Raw	Standardized	Standardized + $L1$	N	N ($L1$)
Component	(1)	(2)	(3)		
Headline	0.179	0.103	0.116	849	838
	(0.133)	(0.071)	(0.072)		
Core	0.129	0.087	0.085	849	838
	(0.110)	(0.066)	(0.069)		
Services	0.213*	0.164**	0.176**	849	838
	(0.109)	(0.077)	(0.087)		
Non-tradable services	0.208**	0.119*	0.114	849	838
	(0.101)	(0.067)	(0.082)		
Tradable services	0.269	0.065	-0.066	793	778
	(0.474)	(0.117)	(0.095)		
Goods	0.150	0.066	0.063	849	838
	(0.170)	(0.070)	(0.068)		
Food	0.107	0.001	-0.003	845	834
	(0.322)	(0.077)	(0.083)		
Energy	0.885	0.050	0.075	849	838
	(0.668)	(0.068)	(0.077)		

Notes: Each entry reports the coefficient on $Gap_{i,t} \times Supply_t \times \chi_i$, where $Gap_{i,t} = NE_{i,t} - EX_{i,t}$ is the difference between non-expert and expert one year ahead forecasts of *headline* inflation and χ_i is standardized commodity terms of trade exposure. The dependent variable is four quarter ahead inflation in the indicated component. All specifications include country and quarter fixed effects and fully interact both forecasts with the supply indicator and exposure, as in the flexible specification of Table 4. Driscoll–Kraay standard errors with seven lags in parentheses.

Column (1) is in percentage points and is not comparable across rows, since components differ in volatility. Column (2) standardizes the dependent variable within country, making rows comparable; the paper's component comparisons use this column. Column (3) adds the one quarter lag of the same component, interacted with the supply indicator and exposure. Components are from the Central Bank of Chile's Global Inflation Dashboard with harmonized weights, a smaller sample than Table 4; the headline row is an internal benchmark, not the paper's estimate of the aggregate gradient.

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

5.5 Exchange rate robustness

A natural alternative is that the exposure gradient simply reflects exchange rate movements. Appendix B.2 reports the corresponding regressions. Controlling for contemporaneous NEER growth changes the main coefficient from -0.179 to -0.178. Controlling for REER growth reduces it to -0.148, but the interaction remains statistically significant. We interpret this as evidence that the currency is one margin of open economy transmission, but not the sole explanation for the exposure gradient.

6 Expectation Formation and Diagnostics

The main empirical results do not rely on a structural model of expectation formation. We nevertheless use forecast rationality and expectation formation regressions to characterize differences between the two groups and to discipline the heterogeneous expectations block of the model.Standard forecast error diagnostics indicate larger departures from forecast efficiency among non-experts than among experts.

The expectation formation regressions play a more direct role in the quantitative model. We regress one year ahead inflation expectations on eight lags of inflation, current and lagged unemployment, and current and lagged policy rates, separately by agent type and inflation regime. The sum of the eight lagged inflation coefficients ranges from approximately 0.52 to 0.65 for non-experts and from approximately 0.30 to 0.36 for experts.

These results show that non-experts are substantially more backward-looking than professional forecasters, while differences across demand and supply regimes are comparatively small. We therefore use

$$\omega^{NE} = 0.60, \qquad \omega^{EX} = 0.35$$

as representative calibration values in the model. These values are not structural estimates of the model's expectation parameters; they are a calibration mapping from the empirical one year ahead expectation regressions into the one period ahead expectation process used in the model.

The underlying regressions and coefficient sums are reported in Appendix C.1.

7 A Small Open Economy Model with Heterogeneous Expectations

The model asks whether physical commodity transmission and exposure dependent information have the same implication for optimal policy. They do not. The exercise is deliberately parsimonious and is intended to organize the empirical mechanism rather than provide a complete structural model of commodity transmission.

7.1 Environment and expectations

Each economy is described by an IS curve, a New Keynesian Phillips curve, and a Taylor rule. Non-expert and expert inflation expectations combine lagged inflation with the model consistent forecast:

$$\pi_{i,t}^{e,NE} = \omega^{NE}\pi_{i,t-1} + (1-\omega^{NE})E_t\pi_{i,t+1}, \tag{9}$$
$$\pi_{i,t}^{e,EX} = \omega^{EX}\pi_{i,t-1} + (1-\omega^{EX})E_t\pi_{i,t+1}. \tag{10}$$

We set

$$\omega^{NE} = 0.60, \qquad \omega^{EX} = 0.35,$$

as a calibration mapping from the expectation formation evidence above. These are not structurally estimated model parameters.

The expectation relevant for price setting is

$$\pi_{i,t}^{e} = \mu_{i,t}\pi_{i,t}^{e,NE} + (1-\mu_{i,t})\pi_{i,t}^{e,EX}, \tag{11}$$

where $\mu_{i,t}$ governs the effective informational relevance of non-expert beliefs. The encompassing coefficient does not structurally identify this weight. We preserve only its comparative static:

$$\mu_D = 0, \qquad \mu_S(\chi_i) = \min\Big\{1,\ \max\big\{0,\ \alpha_\mu + \lambda b\widetilde{\chi}_i\big\}\Big\}, \qquad b = 0.215, \tag{12}$$

where the mapping is linear in exposure and truncated to $[0,1]$, the truncation enforcing the domain restriction in equation (11). The intercept $\alpha_\mu = -\lambda b\widetilde{\chi}^{NE}$ sets $\mu_S = 0$ at the exposure of the average non-exporter ($\chi^{NE} = 0.326$), so it depends on λ. The more flexible empirical specification yields 0.199, but we use the more precisely estimated baseline pooled coefficient as the calibration target. The scalar λ maps the reduced form gradient into the model; $\lambda = 0.5$ is the conservative baseline and $\lambda = 1$ a sensitivity. Appendix C.3 reports the implied weights at each exposure anchor.

7.2 Physical commodity transmission and policy

A favorable commodity movement affects exporters through an income channel and all economies through imported costs:

$$y_{i,t} = E_t y_{i,t+1} - \frac{1}{\sigma}\left(i_{i,t} - \pi^e_{i,t} - r^{nat}_{i,t}\right), \qquad r^{nat}_{i,t} = \psi_g g_t + \psi_y X_i \chi_i z_t, \tag{13}$$

$$\pi_{i,t} = \beta \pi^e_{i,t} + \kappa y_{i,t} + u_{i,t} - \psi_m m_i z_t. \tag{14}$$

Here X_i is an exporter indicator and m_i import cost exposure. The physical parameters are calibrated before introducing informational state dependence. The non-exporter inflation response disciplines $\psi_m = 1.18$, while an exporter activity response, translated using an Okun coefficient of 2, disciplines $\psi_y = 0.86$. The corresponding model moments closely match their targets. The net exporter inflation response is left untargeted: the data estimate is 0.020 (s.e. 0.105) while the model generates -1.19. We leave this validation gap explicit rather than add another free parameter. Appendix C reports the compact calibration and robustness details.

The central bank follows

$$i_{i,t} = \phi^{S_t}_{\pi,i} \pi_{i,t} + \phi_y y_{i,t}, \tag{15}$$

and minimizes the policy loss

$$\mathcal{L}_i = \frac{1}{2}\left[\lambda_\pi \mathrm{Var}(\pi_{i,t}) + \lambda_y \mathrm{Var}(y_{i,t}) + \lambda_i \mathrm{Var}(i_{i,t})\right]. \tag{16}$$

We call this a policy loss rather than literal welfare because it is not derived from micro founded small open economy preferences.

7.3 Policy experiment

We evaluate the model at three observed exposure anchors: the United Kingdom ($\chi = 0.04$), Chile ($\chi = 1.25$), and Norway ($\chi = 2.16$). For each economy we solve for the optimal demand and supply regime inflation coefficients and compare the resulting loss with the optimal state invariant rule.

Table 6: Optimal state dependent policy: transmission versus information

	UK	Chile	Norway
Panel A: Experts only benchmark			
ϕ_π^{D*}	1.88	1.88	1.82
ϕ_π^{S*}	1.62	1.60	1.54
$\phi_\pi^{D*}/\phi_\pi^{S*}$	1.16	1.17	1.18
Policy-loss reduction (%)	0.59	0.68	0.64
Panel B: Information channel, baseline $\lambda = 0.5$			
$\omega_S(\chi)$	0.350	0.387	0.424
ϕ_π^{D*}	1.88	1.88	1.82
ϕ_π^{S*}	1.62	1.64	1.64
$\phi_\pi^{D*}/\phi_\pi^{S*}$	1.16	1.15	1.11
Policy-loss reduction (%)	0.59	0.42	0.23
Panel C: Information channel, sensitivity $\lambda = 1$			
$\omega_S(\chi)$	0.350	0.424	0.497
ϕ_π^{D*}	1.88	1.88	1.82
ϕ_π^{S*}	1.62	1.70	1.72
$\phi_\pi^{D*}/\phi_\pi^{S*}$	1.16	1.11	1.06
Policy-loss reduction (%)	0.59	0.24	0.08

Notes: All panels hold the physical commodity transmission block fixed. Panel A switches off informational state dependence. Panels B and C activate the exposure dependent mapping with $\lambda = 0.5$ and $\lambda = 1$, respectively. Loss is evaluated using equation (16).

Panel A shows that physical transmission alone does not generate convergence. The ratio $\phi_\pi^{D*}/\phi_\pi^{S*}$ rises from 1.16 in the United Kingdom to 1.18 in Norway. Once the informational channel is activated, the ratio falls with exposure. Under the conservative baseline it declines from 1.16 to 1.15 to 1.11; under $\lambda = 1$ it reaches 1.06 in Norway. The mechanism is that non-expert expectations are more backward-looking, so a rising informational weight in supply episodes raises the effective backward-looking component entering the Phillips curve and makes the optimal supply regime inflation response more aggressive.

The policy loss effects are modest. This does not mean that non-expert information lowers welfare. Rather, the informational channel makes the two regime specific optimal coefficients more similar, leaving less to gain from a coarse binary regime rule. We therefore do not interpret these numbers as a welfare value of information.

8 Policy Implications

Three objects should remain distinct:

standalone accuracy, incremental information, structural policy weight.

Professional forecasts dominate on standalone accuracy. The narrowing expert advantage at high exposure does not arise because experts become worse: expert absolute errors have essentially no exposure gradient. The practical case for using non-expert surveys is therefore as a complement, not a replacement.

The encompassing results provide the relevant complementarity. In exposed economies, disagreement between non-experts and professionals can contain information about subsequent inflation that is not summarized by the expert forecast alone. This potential value is especially visible during global supply episodes. For relative accuracy, the gradient remains negative outside 2020–2023 but becomes much larger in the surge. For incremental information, the surge specification attributes most of the estimated gradient to 2020–2023, while the pooled full sample regression remains positive and significant. We therefore view the surge as a period in which the mechanism was particularly pronounced rather than as proof that it exists only in extreme episodes.

The results also caution against mechanical interpretations of survey spikes. Non-expert expectations are more backward-looking, so sharp increases during commodity driven inflation episodes can partly reflect recently observed inflation and salient local conditions rather than durable de-anchoring. That does not make them uninformative: the same movements can contain incremental information about future inflation in sufficiently exposed economies. Central banks should therefore interpret survey levels together with professional forecasts, forecast disagreement, recent forecast errors and the external environment.

The encompassing coefficient is not a structural policy weight and does not deliver a stable real time forecast combination rule. A leave one country out forecast combination exercise does not improve out of sample RMSE relative to the expert forecast alone. The appropriate policy role of non expert surveys is thus state dependent cross checking and diagnostic information, not a fixed weight in the policy process. Appendix D reports these checks.

9 Conclusion

The U.S. style demand regime ranking reversal does not generalize to our twelve country open economy panel: experts remain more accurate in both demand and supply episodes. The main result is instead that the expert advantage shrinks with commodity terms of trade exposure during global supply episodes, and this narrowing reflects lower non-expert forecast errors rather than deteriorating expert performance.

Fair–Shiller regressions show that the same exposure also makes the non-expert/expert forecast gap increasingly informative about subsequent inflation. The incremental information is most clearly visible in future services inflation and is not accounted for by contemporaneous exchange rate movements. The regressions do not identify the underlying local signal. Doing so would require more granular survey data linking individual, regional, or sectoral expectations to observed prices, costs, wages, income or activity.

The model provides a qualitative interpretation. Physical commodity transmission changes domestic incidence but does not by itself narrow optimal demand and supply regime policy responses. The convergence appears only when the model incorporates the exposure dependent informational channel. The result is robust, but the policy loss magnitudes are modest and the mapping from encompassing coefficients to structural informational relevance remains a calibration assumption.

More broadly, the demand/supply character of a global inflation episode is not sufficient to determine which expectations are informative in an open economy. The same external disturbance can have different domestic incidence depending on trade structure, and the informational content of different agents' expectations can change with that incidence. The relevant question for small open economies is therefore not only where global inflation comes from, but how the disturbance is transmitted domestically and which agents' expectations contain useful information once it arrives.

References

Albagli, E., Grigoli, F., and Luttini, E. (2022). Inflation expectations and the supply chain. Technical report, International Monetary Fund Washington, DC.

Bajraj, G., Carlomagno, G., and Wlasiuk, J. M. (2023). Where is the inflation?: The diverging patterns of prices of goods and services. Technical report, Banco Central de Chile.

Blanchard, O. J. and Quah, D. (1988). The dynamic effects of aggregate demand and supply disturbances.

Bobeica, E., Paredes, J., Renault, T., and Rousseau, F. (2024). Selling price expectations for services: what do they tell us about consumer price pressures? *Economic Bulletin Boxes*, 5.

Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–2678.

Cornand, C. and Hubert, P. (2022). Information frictions across various types of inflation expectations. *European Economic Review*, 146:104175.

Davig, T. and Leeper, E. M. (2007). Generalizing the taylor principle. *American Economic Review*, 97(3):607–635.

Drechsel, T. and Tenreyro, S. (2018). Commodity booms and busts in emerging economies. *Journal of International Economics*, 112:200–218.

Fair, R. C. and Shiller, R. J. (1990). Comparing information in forecasts from econometric models. *The American Economic Review*, pages 375–389.

Fernández, A., Schmitt-Grohé, S., and Uribe, M. (2017). World shocks, world prices, and business cycles: An empirical investigation. *Journal of International Economics*, 108:S2–S14.

Fry, R. and Pagan, A. (2011). Sign restrictions in structural vector autoregressions: A critical review. *Journal of Economic Literature*, 49(4):938–960.

Gali, J. and Monacelli, T. (2005). Monetary policy and exchange rate volatility in a small open economy. *The Review of Economic Studies*, 72(3):707–734.

Gruss, B. and Kebhaj, S. (2019). *Commodity terms of trade: A new database*. International Monetary Fund.

Hofmann, B., Manea, C., and Mojon, B. (2024). Targeted taylor rules: monetary policy responses to demand-and supply-driven inflation. *BIS Quarterly Review*, 241.

Letelier, F. and Luna, F. (2026). Heterogeneous expectations across inflation regimes: Evidence and implications for monetary policy. Working paper.

Shapiro, A. H. (2026). Decomposing supply-and demand-driven inflation. *Journal of Money, Credit and Banking*, 58(2):365–388.

Verbrugge, R. J. and Zaman, S. (2021). Whose inflation expectations best predict inflation? *Economic Commentary*, (2021-19).

Weber, M., d'Acunto, F., Gorodnichenko, Y., and Coibion, O. (2022). The subjective inflation expectations of households and firms: Measurement, determinants, and implications. *Journal of Economic Perspectives*, 36(3):157–184.

Wehrhöfer, N. (2026). Energy prices and inflation expectations: Evidence from households and firms. *Journal of Political Economy Macroeconomics*, 4(1):135–163.

A Domestic Regime and Commodity Shock Construction

For each country, we estimate a VAR(4) in quarterly real GDP growth and CPI inflation,

$$Y_{i,t} = A_{i,1}Y_{i,t-1} + \cdots + A_{i,4}Y_{i,t-4} + C_i D_t + u_{i,t}, \qquad Y_{i,t} = \begin{pmatrix} \Delta y_{i,t} \\ \pi_{i,t} \end{pmatrix},$$

where D_t contains seasonal controls and 2020 outlier dummies.

Under sign restrictions, a positive demand shock moves output growth and inflation in the same direction, while a positive supply shock moves them in opposite directions. Identification uses the median target rotation. As an alternative, we impose the long run restriction that demand shocks have no permanent effect on the level of output. Historical decompositions classify a country-quarter as demand driven when demand shocks account for more than one half of trailing four quarter inflation.

The two domestic classifications agree in approximately 80 percent of country-quarters, while each agrees with the U.S. based global classification in only about 40 percent. This motivates distinguishing the global shock mix from its domestic incidence.

Commodity terms of trade innovations are constructed from quarterly log changes in the IMF net export commodity terms of trade index. We remove predictable variation using country specific autoregressive processes with seasonal controls and denote the residual by $z_{i,t}$. Because these processes are estimated over the available sample, the residual is treated as an empirical innovation rather than as a literal real time signal observed by survey respondents.

B Selected Empirical Robustness

B.1 No reversal result

Table 7 shows that the absence of a demand/supply ranking reversal is robust to alternative Driscoll–Kraay bandwidths, exclusion of 2020–2023, alternative forecast timing, stricter regime thresholds, squared rather than absolute loss, two way clustered standard errors, and leave one country out estimation. Across these exercises, experts remain more accurate on average in both regimes.

Table 7: Robustness of the no-reversal result

Specification	Contrast (D–S)	s.e.
Baseline (d^{sign}, DK(7))	-0.026	(0.119)
DK(4) lags	-0.026	(0.116)
DK(10) lags	-0.026	(0.122)
Excluding 2020–2023	-0.097	(0.126)
Target-quarter convention	0.096	(0.157)
Strict regimes (share ≥ 0.6 / ≤ 0.4)	0.037	(0.127)
Squared loss	-0.685	(0.869)
Two-way clustered SEs	-0.026	(0.142)
Leave-one-country-out range	[-0.109, 0.042]	

Notes: Each row reports the demand-minus-supply contrast in the loss differential, where the dependent variable is $|FE^{NE}| - |FE^{E}|$. The baseline uses d^{sign}, country fixed effects, and Driscoll–Kraay standard errors with 7 lags.

B.2 Exposure gradient and exchange rates

The main supply exposure coefficient ranges from -0.141 to -0.268 when countries are excluded one at a time and remains statistically significant in eleven of twelve re-estimations.

Table 8 examines whether contemporaneous exchange rate movements account for the exposure gradient in relative forecast performance. Controlling for NEER growth changes the baseline coefficient from -0.179 to -0.178, while controlling for REER growth reduces it to -0.148, with the interaction remaining statistically significant. Because exchange rates respond endogenously to terms of trade disturbances, these specifications are robustness checks rather than causal mediation tests.

Table 8: Exchange rate channel of the terms of trade mechanism

	(1) NEER response	(2) REER response	(3) Mechanism + NEER	(4) Mechanism + REER	(5) NEER-based channel
ToT shock	1.454	6.713			
	(1.019)	(6.951)			
ToT shock × commodity exporter	-0.480	-5.699			
	(0.987)	(6.865)			
Supply × exposure χ_i			-0.178*	-0.148**	
			(0.070)	(0.051)	
NEER growth, QoQ %			0.018		
			(0.019)		
REER growth, QoQ %				0.050*	
				(0.020)	
Supply × NEER growth					0.019
					(0.023)
Observations	934	721	934	721	934
Exporter net effect	0.974	1.014			
s.e.	0.198	0.194			
p-value: exporter net effect = 0	0.000	0.000			
Country fixed effects	Yes	Yes	Yes	Yes	Yes

Notes: NEER and REER are nominal and real effective exchange-rate indices from the IMF, normalized to 2010 = 100. NEER and REER growth are quarter-on-quarter log changes; positive values denote appreciation. Columns (1)–(2) estimate the exchange-rate response to a terms-of-trade shock and its interaction with commodity-exporter status. The exporter net effect is the sum of the two reported coefficients, and the p-value tests whether that sum is zero. Columns (3)–(4) augment the baseline terms-of-trade mechanism regression with contemporaneous exchange-rate growth. Column (5) replaces terms-of-trade exposure with realized NEER growth interacted with the global supply-regime indicator. The dependent variable in columns (3)–(5) is $|FE_{i,t}^{NE}| - |FE_{i,t}^{EX}|$, so positive values indicate a larger non-expert forecast disadvantage, or equivalently a larger expert advantage. The global supply indicator is based on the Shapiro U.S. demand–supply classification. NEER data are available for all 12 countries in the matched sample. REER data are unavailable for Peru, Serbia, and Albania, so columns (2) and (4) are estimated on a reduced 9-country sample. All specifications include country fixed effects. Standard errors are Driscoll–Kraay with seven lags and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

C Expectation Formation and Model Calibration

C.1 Expectation formation

The heterogeneous expectations block is disciplined by regressions of one year ahead inflation expectations on eight lags of realized inflation, current and lagged unemployment, and current and lagged monetary policy rates, estimated separately by agent type and regime.

For agent type $a \in \{NE, EX\}$, we summarize backward-lookingness by

$$B_r^a = \sum_{j=1}^{8} \beta_{j,r}^a,$$

where r denotes the inflation regime.

For non-experts, the cumulative lagged inflation coefficient ranges from approximately 0.52 to 0.65 across the reported specifications. The strict adaptive benchmark

$$H_0 : \sum_{j=1}^{8} \beta_j = 1$$

is rejected, so the estimates indicate substantial backward lookingness without implying a literally adaptive expectations process.

For experts, the corresponding cumulative coefficient is materially smaller, ranging from approximately 0.30 to 0.36. Variation across agent types is therefore much larger than variation across demand and supply regimes within each group. We use these ranges to motivate the representative calibration

$$\omega^{NE} = 0.60, \qquad \omega^{EX} = 0.35.$$

These values are calibration mappings from the empirical one year ahead regressions into the model's one period ahead expectation process, not structural estimates.

Table 9: Adaptive expectations tests: Non-Experts

	(1) US-Demand	(2) US-Supply	(3) Sign-Demand	(4) Sign-Supply	(5) BQ-Demand	(6) BQ-Supply
Unemployment t	-0.019	-0.024	0.019	0.079*	0.038	0.044
	(0.038)	(0.049)	(0.020)	(0.046)	(0.026)	(0.043)
Unemployment $t-1$	-0.014	-0.038	-0.016	-0.016	-0.031	-0.020
	(0.061)	(0.041)	(0.026)	(0.042)	(0.023)	(0.050)
Policy rate t	0.534***	0.415***	0.538***	0.251***	0.457***	0.403***
	(0.119)	(0.036)	(0.063)	(0.068)	(0.057)	(0.082)
Policy rate $t-1$	-0.403***	-0.345***	-0.451***	-0.060	-0.337***	-0.290***
	(0.114)	(0.046)	(0.065)	(0.067)	(0.067)	(0.076)
Observations	311	560	531	289	617	203
Sum of inflation-lag coefficients	0.616	0.636	0.555	0.522	0.541	0.650
s.e.	0.063	0.042	0.040	0.055	0.046	0.056
p-value: sum of inflation lags = 1	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001
p-value: macro variables = 0	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001

Notes: The dependent variable is the one-year-ahead non-experts inflation expectation. Each column estimates the expectation-formation equation separately by inflation regime. The regressors are eight lags of inflation, current and lagged unemployment, and current and lagged policy rates. The table reports the sum of the inflation-lag coefficients rather than the individual lag coefficients. All specifications include country fixed effects. Standard errors are Driscoll–Kraay with 7 lags and are reported in parentheses. The test $\sum_j \beta_j = 1$ evaluates the strict adaptive-expectations benchmark. The joint test of macro variables evaluates whether unemployment and policy rates add explanatory power beyond lagged inflation. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 10: Adaptive expectations tests: experts

	(1) US-Demand	(2) US-Supply	(3) Sign-Demand	(4) Sign-Supply	(5) BQ-Demand	(6) BQ-Supply
Unemployment t	0.024	-0.002	-0.020	0.102**	0.008	0.049
	(0.030)	(0.025)	(0.024)	(0.050)	(0.022)	(0.040)
Unemployment $t-1$	-0.004	-0.002	0.027	-0.073	0.006	-0.050
	(0.039)	(0.022)	(0.022)	(0.054)	(0.024)	(0.050)
Policy rate t	0.480***	0.463***	0.463***	0.481***	0.441***	0.597***
	(0.102)	(0.059)	(0.064)	(0.066)	(0.069)	(0.083)
Policy rate $t-1$	-0.377***	-0.421***	-0.428***	-0.360***	-0.385***	-0.542***
	(0.093)	(0.060)	(0.065)	(0.072)	(0.072)	(0.097)
Observations	311	560	531	289	617	203
Sum of inflation-lag coefficients	0.327	0.364	0.347	0.298	0.345	0.363
s.e.	0.048	0.036	0.030	0.043	0.033	0.051
p-value: sum of inflation lags = 1	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001
p-value: macro variables = 0	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001

Notes: The dependent variable is the one-year-ahead expert inflation expectation. Each column estimates the expectation-formation equation separately by inflation regime. The regressors are eight lags of inflation, current and lagged unemployment, and current and lagged policy rates. The table reports the sum of the inflation-lag coefficients rather than the individual lag coefficients. All specifications include country fixed effects. Standard errors are Driscoll–Kraay with 7 lags and are reported in parentheses. The test $\sum_j \beta_j = 1$ evaluates the strict adaptive-expectations benchmark. The joint test of macro variables evaluates whether unemployment and policy rates add explanatory power beyond lagged inflation. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

C.2 Calibration

Table 11 summarizes the remaining calibration.

Table 11: Model calibration

Parameter	Value	Empirical discipline / role
β, σ, κ	$0.99, 1.00, 0.10$	Standard quarterly NK values
ω^{NE}, ω^{EX}	$0.60, 0.35$	Tables 9–10
b	0.215	Baseline pooled encompassing gradient
λ	0.5 (baseline), 1 (sensitivity)	Reduced form to model mapping
ψ_m	1.18	Non exporter inflation response
ψ_y	0.86	Exporter activity response / Okun mapping
ρ_g, ρ_u, ρ_z	$0.80, 0.50, 0.22$	VAR / panel persistence
$(\sigma_g, \sigma_u)_D$	$(0.784, 0.387)$	Shapiro demand regime moments
$(\sigma_g, \sigma_u)_S$	$(0.683, 0.606)$	Shapiro supply regime moments
$(\sigma_{z,D}, \sigma_{z,S})$	$(0.782, 1.066)$	CTOT factor by regime
ϕ_y	0.125	Policy rule benchmark
$(\lambda_\pi, \lambda_y, \lambda_i)$	$(1, 0.5, 0.5)$	Policy loss weights

The physical transmission block is calibrated under the experts only benchmark before activating informational state dependence. The non-exporter inflation response is -1.651 (s.e. 0.759) in the data and -1.640 in the model, disciplining ψ_m. The exporter activity target is 0.128, compared with 0.129 in the model, disciplining ψ_y.

The exporter net inflation response is deliberately untargeted. Its empirical estimate is 0.020 (s.e. 0.105), while the model generates -1.19. We leave this validation gap explicit rather than introduce an additional free parameter.

C.3 Informational mapping and Policy robustness

The informational channel is disciplined by the pooled encompassing coefficient

$$b = 0.215.$$

This coefficient measures the exposure gradient in the conditional predictive content of the non-expert/expert forecast gap; it does not structurally identify the model's informational weight. The model maps this empirical gradient into

$$\mu_D = 0, \qquad \mu_S(\chi_i) = \min\Big\{1,\ \max\big\{0,\ \alpha_\mu + \lambda b \widetilde{\chi}_i\big\}\Big\}, \qquad \alpha_\mu = -\lambda b \widetilde{\chi}^{NE},$$

using $\lambda = 0.5$ as the baseline and $\lambda = 1$ as a sensitivity. The intercept sets $\mu_S = 0$ at the average non-exporter's exposure ($\chi^{NE} = 0.326$); the linear form carries no structural content and is used because it transmits the estimated gradient without introducing curvature the data do not identify. The resulting weights are

			$\lambda = 0.5$	$\lambda = 1.0$
Anchor	χ_i	$\widetilde{\chi}_i$	ω_S	ω_S
United Kingdom	0.04	-1.015	0.350	0.350
Chile	1.25	0.794	0.387	0.424
Norway	2.16	2.154	0.424	0.497

where $\omega_S(\chi_i) = 0.35 + 0.25\,\mu_S(\chi_i)$. At the United Kingdom anchor the truncation binds, so Panels B and C of Table 6 coincide with the experts-only benchmark by construction.

We evaluate the policy convergence result across alternative policy loss weights, Phillips curve slopes, Taylor rule output responses, cost channel exposure specifications, Okun coefficients, and expert/non-expert backward-lookingness gaps. In all forty-eight specification–λ combinations, activating the informational channel moves the optimal demand and supply regime policy coefficients closer together when convergence is measured by

$$\left| \log\left(\frac{\phi_\pi^{D*}}{\phi_\pi^{S*}} \right) \right|.$$

One extreme case slightly overshoots parity but still represents convergence under this symmetric metric.

For scale, imposing the sharper U.S. style regime switch

$$\omega^D = 0.30, \qquad \omega^S = 0.05,$$

with no exposure dependent informational channel generates a policy loss reduction of 7.12%, substantially larger than the baseline open economy magnitudes.

D Additional Policy Relevant Forecasting Checks

The encompassing regressions establish incremental information in the panel, but they do not imply a stable real time forecast combination rule. In a leave one country out exercise, estimating the encompassing relationship on the remaining countries does not improve out of sample RMSE relative to the expert forecast alone. Incremental information and out of sample forecast improvement are therefore distinct objects.

Standalone accuracy also remains distinct from incremental information during the 2020–2023 inflation surge. Among commodity exporters, the mean non-expert forecast disadvantage declines from approximately 0.41 percentage points outside the surge to 0.18 during it. The latter is statistically indistinguishable from zero but remains positive. Among non-exporters during the same period, the corresponding expert advantage is approximately 0.91 percentage points.

Thus, even in the episode in which the exposure dependent informational gradient is strongest, the evidence is one of a narrowing expert advantage and greater incremental information rather than a systematic reversal in standalone forecast performance.