

Estimating the rise in expected inflation from higher energy prices*

Paula Patzelt

LSE

Ricardo Reis

LSE

February 2024

Abstract

When the price of electricity increases by 1%, households' average expected inflation increases by 0.9 to 1.3 basis points. But, if those expectations have become unanchored, as happened between the start of 2021 and 2023, then the effect is higher by 0.2 to 0.7 basis points. This paper arrives at these estimates by exploiting variation both in the time series, and especially in the cross section, from newly-available public data on expected inflation by Euro area households across region, gender, education, and income, and on the cost of energy across region and source. The impact of exogenous shocks to energy prices on expected inflation increases for 8 to 12 months, but they can only account for a small share of the rise in expected inflation in 2021-23.

JEL codes: D84, E31, Q43.

Keywords: Great Inflation, Monetary policy, Inattention.

*Contacts: p.h.patzelt@lse.ac.uk and r.a.reis@lse.ac.uk. We are grateful to Julian Callow, Salomé Fofana, Ethan Ilzetzki, Emiliano Rinaldi, Raphael Schoenle, and Gertjan Vlieghe for discussions. This work was supported by the UK Research and Innovation grant number EP/Eo25039/1. First draft: February 2024.

1 Introduction

In May of 2021, annual inflation in the euro area (EA) was precisely at its target of 2%. Twelve months later, inflation had climbed to 8.1%. Two traditional culprits have been put forward: expected inflation and energy prices. After all, during this year they rose by 2.3 percentage points and by 33%, respectively. Yet, one (of many) difficulties with evaluating their contribution to inflation is that each causes the other, and in turn monetary policy responds to both, affecting both as well. This paper makes progress on this challenge by providing empirical estimates to answer two related questions.

The first of these is: *by how much does expected inflation over the next year increase on average when energy prices rise by 1%?* Most of the literature linking energy prices to expected inflation has used time-series variation and has focussed on oil prices. We innovate and improve in several ways.

First, we use the recently-available Consumer Expectations Survey (CES) for the EA, which has between 9,000 and 22,000 respondents. By comparison, the United States (US) Michigan Survey of Consumers (MSC) covers between 500 and 700 households per wave. While the MSC has a longer sample in time, when higher energy prices raise expected inflation, this often comes with an offsetting monetary policy tightening. Time-series estimates will then understate the impact of the energy prices. Likewise, many changes in aggregate demand likely affect both the demand for energy and expectations, creating omitted variables that are difficult to deal with using time-series variation by itself.

Second, we identify the effect by exploiting the large variation in energy prices across 11 countries with sharp changes in a monthly sample between 2020:4 and 2023:12. Variation across regions within a monetary union differences out monetary policy and other aggregate demand factors. While cross-sectional variation alone is not enough, because there are large fixed differences in expected inflation across countries, a short but large panel with plenty of variation in both dimensions can deliver precise estimates.

Third, we focus on the price of electricity paid by households. It is more visible than oil prices since they pay for it frequently. It is also more relevant than the price of gas at the pump and other related energy measures that have been used before, since electricity accounted for 25% of energy consumption in 2021, while oil and petroleum products were a mere 10%. We address the reverse causality from expected inflation potentially driving demand for all goods including energy, and thus affecting electricity prices, by using the features of the electricity market to propose three plausibly exogenous measures of energy

prices. The first relies on a shift-share strategy that exploits cross-region differences in the weight of energy in consumers' baskets, the second uses an exogenous measure of time-series shocks to oil supply, and the third exploits variation in the use of wind to generate electricity across time and region. We complement the regressions with local projections to assess the dynamic effects of the shocks.

All combined, we find that a 1% increase in electricity prices raises expected inflation by 0.96–1.30 basis points (bps). A one-standard deviation exogenous shock to energy prices raises expected inflation by 15–60 bps, with the impact growing until 10 to 12 months after the shock.

The second question is: *by how much more does the 1% rise in energy prices increase inflation expectations when those expectations are less well anchored?* All central banks strive to anchor inflation expectations. If expectations are less responsive to shocks, they will not amplify the initial shock as much, and require a smaller response of policy rates.

While this argument applies to all shocks to inflation, those driven by energy prices play a special role. Because energy shocks tend to be temporary, there is a policy view to “see through them” and not respond by changing policy rates. Yet, if inflation expectations are unanchored, then the energy shock will shift them. Since perceptions tend to persist, they drive persistent movements in inflation that policy should not see through.

Because of the cross-sectional size and richness of the CES data, we provide a first empirical examination of this classic policy channel. Since measures of unanchoring based on disagreement are systematically different not just across country, but also across demographic and socio-economic groups, we can use the cross-sectional variation in the time-series change in disagreement to measure the differential impact of energy shocks. During our short sample, there was significant variability in both energy prices and measures of unanchoring, so the estimates are sharp.

We find that when measures of the expectations anchor drift between their average level during 2021 and during 2023, the 1% increase in electricity prices raises expected inflation by an additional 0.22–0.69bps. Further, the peak of the impulse response to exogenous energy shocks can be more than twice as high when disagreement in expectations is above average versus below average. That significant boost empirically confirms the importance of keeping inflation expectations anchored that theory has suggested.

We apply our estimates to shed light on four related issues. First, we calculate how much of the increase in expected inflation in 2021-22 can be attributed to energy shocks.

They account for a very small part of it. Second, we calculate the impact on expected inflation of a rise in electricity prices at each point in the sample. When expectations were most unanchored, in the first half of 2022, a doubling of electricity prices would raise expected inflation by 80-120bps over the following six months. By the end of our sample, a doubling of electricity prices in the second half of 2023 would raise expected inflation by only 45–60bps by the end of the year. Third, we use US data, with its limited size and cross-sectional variation but a longer time-series sample, to arrive at estimates consistent with the EA ones, but larger, more imprecise, and unstable across specifications. Fourth, we interpret these estimates in light of theories of state-dependent attention, which suggest a link between unanchoring and markup shocks following an energy shock.

The online appendix reviews previous answers to our empirical questions in the literature; here we discuss our marginal contribution.

While the correlation between oil prices and average inflation expectations is often stated as a fact beyond dispute, as with most time-series correlations between two aggregate variables, this one is neither reliable nor meaningful. More carefully, Coibion and Gorodnichenko (2015) and Binder (2018) answer the first question by regressing household expectations on wholesale oil prices and gasoline prices at the pump, and obtain estimates of 1.6bp and 1.0bp, respectively. While they use mostly time-series variation, our estimates use country-group variation within a currency union to deal with the confounding effect of monetary policy and other aggregate demand factors. Kilian and Zhou (2022) uses a structural vector autoregression on aggregate variables identified with zero and sign restrictions, while Känzig (2021) uses local projections on exogenous shocks to expected oil prices, and Miyamoto, Nguyen and Sergeyev (2024) separate periods where nominal interest rates were at the zero lower bound. Six months after the shock, they all find that the impact of higher oil prices on average expected inflation is close to zero, and that oil prices account for a small share of the variation in average expected inflation. We find larger and more persistent effects by using exogenous shocks in the cross-section during a shorter time sample where these were large to achieve identification. Finally, Binder and Makridis (2022) uses state-level variation in real gasoline prices, but for its impact on consumer sentiment, Wehrhofer (2023) uses cross-household variation on the renewal of electricity contracts, but cannot answer our question because it does not have information on how much prices changed, and Hajdini et al. (2024) use cross-county variation in the share of households commuting by car, but without considering the impact of shocks after one week or the role of unanchoring.

Turning to the second question, to our knowledge, this is the first paper to use household micro data to empirically investigate whether the impact of energy prices on expected inflation is different when expectations are unanchored. We follow Kumar et al. (2015) and Bonomo et al. (2024) in measuring anchoring, and relate our findings to the model of state-dependent attention in Flynn and Sastry (2024). Our conclusion that energy prices do not explain the increase in expected inflation during this period contributes to the literature on the 2021-23 inflation disaster (Reis, 2023, Acharya et al., 2023, Gagliardone and Gertler, 2023, Vlieghe, 2024).

The paper is structured as follows. The next section discusses the data and the empirical strategies. Section 3 presents the estimates that answer the two main questions, while section 4 applies them to the related four issues. Section 5 concludes.

2 Data and empirical strategy

We first describe the data and our regression, before explaining how we use cross-sectional variation to achieve identification and build new exogenous energy shocks.

2.1 Variables and data

Let $\pi_{i,c,g,t}^e$ be the answer by household i , who is a resident of region/country c , and is part of a demographic or socio-economic group g , in the month t , to the question: “How much higher (lower) do you think prices in general will be 12 months from now in the country you currently live in?” The data come from the ECB’s Consumer Expectations Survey (CES), where i goes to 9,000 to 22,000 respondents depending on the month, c are eleven countries in the euro area, there are eight demographic groups g from crossing gender (male/female), income bracket (above/below 60th percentile), and education (college/below), and the months t go from April of 2020 to December of 2023 for six countries, and from April 2022 for the remaining five. The data with all the countries and enough cross-sectional richness were first available in February of 2024.

The operator Δ^h refers to the change in a variable relative to its value h months ago. Therefore $\Delta^6 \pi_{i,c,g,t}^e = \pi_{i,c,g,t}^e - \pi_{i,c,g,t-6}^e$ and this variable exists only if a household answered the survey both in month $t - 6$ and again in month t . For expected inflation, choosing $h = 12$ ensures that there is no overlap between the observation frequency and the forecast horizons, while choosing $h = 1$ maximizes the number of observations but

introduces noise. Monthly changes in energy prices are also volatile and transient. We choose $h = 6$ for our baseline, and report estimates with $h = 1, 4, 12$ as well.

Eurostat provides an index for harmonised electricity prices per country paid by households inclusive of taxes and subsidies. Let $e_{c,t}$ denote the log of that index, while e_t is its counterpart for the whole EA. This is a nominal variable, but since we include inflation as a control variable in all regressions, using instead its real equivalent makes little difference. We also consider three alternative measures of energy prices: the consumer price index for energy as opposed to electricity, a measure of wholesale electricity prices, and a regulatory funded public project's measure of electricity prices in capital cities. The appendix describes their sources.

Actual inflation comes from Eurostat, and corresponds to the log change between the harmonised index of consumer prices in date t and 12 months earlier, per country: $\pi_{c,t}$. We denote average year-on-year inflation in the last year by $\bar{\pi}_{c,t} = \sum_{j=1}^{12} \pi_{c,t-j} / 12$.

The degree of anchoring of inflation expectations within a country-group is $a_{c,g,t}$, where a higher $a_{c,g,t}$ stands for more unanchored. The literature has used data on longer-term inflation expectations to measure unanchoring in two ways. One uses higher-order moments of the distribution of inflation expectations, arguing that disagreement among households reflects an unanchoring of expectations. This would be the case in models with incomplete information and dispersed expectations. The other uses the difference between expected inflation and the inflation target, arguing that unanchoring reveals itself as a loss of credibility of the target. Models of learning and reputation support these measures. We use one measure from each of these two classes: the 6-month change in the interquartile range of expected inflation 3-years ahead within country-group, and the 6-month change in the absolute difference between expected inflation 3-years ahead and the ECB's inflation target averaged by country-group.

To build exogenous shocks to energy prices, we use three series: the changes in oil supply expectations from Känzig (2021) that measures the high-frequency change in oil futures prices following OPEC production announcements, k_t ; the share of electricity in household consumption per region in 2019 from the Eurostat HICP, s_c ; and the total energy generated through wind in each region and month from Ember, $w_{c,t}$. Each of these variables is exogenous in the sense that inflation expectations between 2020 and 2023 plausibly did not directly affect them.

2.2 Empirical strategy

Our baseline regression equation on the unbalanced panel is:

$$\Delta^6 \pi_{i,c,g,t}^e = \beta \Delta^6 e_{c,t} + \gamma \Delta^6 e_{c,t} \times \Delta^6 a_{c,g,t} + \alpha_c + \eta_g + \theta \bar{\pi}_{c,t-6} + \varepsilon_{i,c,g,t}, \quad (1)$$

where β and γ answer our first and second questions, respectively, while α_c are country fixed effects, η_g are group fixed effects, θ is the coefficient from controlling for past inflation, and $\varepsilon_{i,c,g,t}$ are residuals.

Starting with the left-hand side, as Fofana, Patzelt and Reis (2024) document, there are large average differences across demographic groups and across countries in average expected inflation. A woman resident of Germany without college that is poorer expects much higher inflation than a richer Italian man with a college degree. Therefore, looking at differences in expected inflation, as opposed to levels, is important when using cross country-group variation.

During this sample period of rising inflation, there was a marked difference in the updating of average expected inflation across countries. This may be due to different levels of trust in monetary policy across countries, or to country-specific characteristics affecting prices. Because of this, we control both for the recent inflation history, as well as for country fixed effects. This also controls for some of the country macro aggregates that affect both variables of interest.

There is no variation in i in any of the right-hand side variables. These are seemingly unrelated regressions, which use the individual variation within country-group to sharpen the estimates of the common coefficients of interest. One may disagree with the implicit assumption that the individual variability has the same structure within country-groups so, as an alternative, we replace the left-hand side variable with $\Delta^6 \pi_{c,g,t}^e = \sum_i \Delta^6 \pi_{i,c,g,t}^e / I_{c,g,t}$, the average expected inflation within a country-group.

On the interpretation of the estimates, we multiply the left-hand side variable by 100, so that β measures the impact on expected inflation in basis points of a 1% increase in oil prices. From a steady state where the anchor remains stable, γ measures by how many basis points more will 1-year ahead expected inflation rise with the increase in electricity prices if unanchoring increased, as measured by a 1 percentage point higher interquartile range of 3-year ahead inflation expectations. Coincidentally, the average disagreement across all households in 2023 was 1.05 percentage points higher than on average in 2021, so γ measures the approximate extra impact of an energy shock between these two years.

2.3 Cross-sectional variation and identification

A concern with estimating the impact of energy prices on expected inflation is the confounding influence of monetary policy. Central banks closely watch both variables and respond to them. In turn, their policies affect inflation and aggregate demand directly, which in turn should move expectations and energy prices. Since our estimates exploit the cross country-group variation within a currency union, they control for the common monetary policy that affected all the people in our survey.

The variation across groups helps to identify the effects of unanchoring. Given the group fixed effects, γ is identified from the change in expected inflation in one country relative to another where electricity prices rose by less and expectations were more anchored relative to the other country-groups. To distinguish between country and group variability, we also consider country variability alone, by aggregating across the groups, thus replacing $a_{c,g,t}$ by $a_{c,t}$. In the other direction, to ascertain whether there is a bias from systematic differences in the way groups within countries changed their expectations during this time, we also include country-group fixed effects.

The time variation across countries and groups is important. We also consider including month fixed effects, so that it is exclusively the cross-group variation that identifies the estimates and, more extremely, a full set of time-country-group joint fixed effects. At one extreme, estimating equation (1) separately for each country gives a set of $\{\beta_c, \gamma_c\}$ estimates that only use the variation in anchoring across groups, and so is close to a time-series regression. It results in very imprecise estimates that are widely different across countries, confirming the importance of the country-group variation.

2.4 Reverse causality and shocks to energy prices

While β and γ answer our two questions, a follow-up question is whether the answer changes if the increase in electricity prices was due to a change in the supply of energy as opposed to the demand for it. Moreover, if the demand shock to electricity prices was itself due to people suddenly expecting higher inflation, then while our estimates would still answer the questions, they would not answer the closely related question of what is the impact on expected inflation of a *shock* to energy prices?

In our sample, it may well turn out that the baseline estimates actually answer this related question. The increase in electricity prices was mostly driven by the invasion of Ukraine. Moreover, each country responded with different measures to the ensuing crisis.

While these differential responses may have been correlated with that country's inflation experience, which recall we control for, they were plausibly not a response to differences in expected inflation.

The cross-sectional variation in the data allows us to explore this further. There are large differences in how much households spend on energy across regions, because of differences in temperature, whether home heating is mostly based on gas, electricity, or solar panels, and the share of electric vehicles, among other factors. In one extreme, energy consumption in Portugal is less than 90 terajoules per person, while Belgium's consumption is more than 200 terajoules per person. The shares of electricity consumption in 2019, before the rise of inflation in our sample, s_c capture the cross-country variability in the impact of higher electricity prices on household budgets, which would proxy for their visibility in forming expectations of inflation. Consider then a shift-share shock series, which multiplies the aggregate time-series variation in energy prices, with the cross-sectional ex ante variation in energy spending: $z_{c,t} = e_t s_c$. Insofar as cross-country differences in expected inflation may drive cross-country spending and electricity prices, but for each country do not affect EA aggregate demand for energy and prices, then this shock series would not be affected by reverse causality.

One concern (implausible to us) is that some households in 2019 foresaw the energy shock that was coming and adjusted their consumption of energy accordingly affecting the s_c . The appendix considers two alternatives for the shares. First, the average expenditure shares over a longer period, 2015 and 2019. Second, to purge from any quantity variation, the network cost of electricity for households in euro per KWh in 2019. This varies across energy markets quite considerably.

One can go even further by exploiting the features of the electricity market in the EA. It is segmented across regions, which do not coincide with countries, but instead with distribution networks. In each of these markets, electricity is produced with renewable sources, which have a large fixed cost but a low marginal cost, and so tend to be infra-marginal. The same applies to nuclear sources, and both renewables and nuclear are hard to expand or contract in the short run. The marginal production of electricity uses oil, natural gas, and solid fossil fuels (like coal), with a competitive market switching between them, and as a result their prices often move together.

Following a cut in the supply of natural gas from Russia, together with environmental and capacity constraints on expanding fossil fuels, oil prices become a proxy for the marginal cost of production of electricity in the EA. The literature has produced changes

in oil prices that are exogenous to demand. Using them leads to a different shift-share shock series: $z_{c,t} = k_t s_c$, where both the shifter and the share are plausibly exogenous.

A final alternative shock series is the production of electricity from wind in each country and month: $z_{c,t} = w_{c,t}$. This is mostly driven by exogenous fluctuations in the weather. When there is more wind, and since the marginal cost of producing electricity for installed turbines is very low, the supply curve for electricity shifts to the right, lowering the price of energy. One concern might be that higher expected inflation could lead to building more wind turbines. Yet, installing this capacity takes time. Moreover, the correlation between our $w_{c,t}$ series and a monthly series for mean wind speed by region is high for most countries, especially for those where wind power is a large share of electricity production. This confirms that most of the variation is indeed exogenous.

Replacing the country-time specific shocks to energy prices $z_{c,t}$ for the energy prices $e_{c,t}$ in equation (1) provides estimates that answer the question posed at the start of this section. As each shock series is in different units, we standardize them, so that β and γ now measure the impact on expected inflation of a one-standard deviation energy shock.

2.5 Dynamics

The regression picks an horizon h over which to measure the impact on expectations. To assess how this may evolve over time, we estimate a local projection in the panel of data for each horizon $h = 1, \dots, 24$:

$$\pi_{c,g,t+h}^e = \beta^h \left(\sum_{p=0}^P z_{c,t-p} \right) + \gamma^h \left(\sum_{p=0}^P z_{c,t-p} \right) A_{c,g,t} + \alpha_c^h + \eta_g^h + \theta^h \bar{\pi}_{c,t-6} + \phi^h + \varepsilon_{c,g,t+h}. \quad (2)$$

This measures the impact on average expected inflation in h months of a cumulative energy shock over the last P months. We set $P = 2$, so the energy shock is over three months, although the results are insensitive to this choice. The dummy variable $A_{c,g,t}$ captures whether unanchoring was above average for that country-group. Therefore, β^h is the impact when expectations unanchor less, while $\beta^h + \gamma^h$ is the impact when they unanchor more. As before, and for the same reasons, we include country fixed effects α_c^h and group fixed effects η_g^h at each horizon, and control for past inflation with coefficient θ^h . Since the left-hand side variable is in levels, we include a horizon intercept ϕ^h .

3 The results

Table 1 reports the results from estimating equation (1).

The first column shows that a 1% increase in electricity prices raises expected inflation by 1.16bp if there is no change in anchoring. However, if disagreement increases by as much as the difference between 2021 and 2023, then the higher electricity prices add an extra 0.67bp effect. Both effects are statistically significant.

The fit of the regression is low, as expected given that no explanatory variable captures the household variation in expected inflation. The second column instead pools observations within country-group. The R^2 dramatically rises, as expected. The effect without unanchoring stays roughly the same, at 0.96bp, while the impact of unanchoring falls to 0.22bp. Both remain significant.

The third column uses as the measure of unanchoring the distance of 3-year expected inflation from target. The effect of higher electricity prices with no change in anchoring is similar, at 0.98bp. Since this measure of unanchoring increased by 76bp between 2021 and 2023, its coefficient now implies that unanchoring contributed to an extra 2.05bp increase in expected inflation following a 1% rise in electricity prices.

The next three columns explore the role of the cross country-group variation in driving the results. Column four shows the results using only country, but no group variation. The impact without unanchoring is slightly larger, while the extra boost from unanchoring is much larger with a point estimate of 1.61bp. With less identifying variation, the confidence bands are wider. Column five includes country-group fixed effects to see if there is a bias from systematic variation in this interaction. It seems not to be the case, as the estimates are very close to the baseline. The final column includes month fixed effects. This soaks much variability leaving only the variation across 11 countries and 8 groups to estimate the coefficients. Both estimates are lower, and marginally statistically significant.

The appendix shows estimates for several alternatives: with three alternative measures of energy prices (table A2), with alternative specifications on the influence of unanchoring (table A3), using a balanced panel of only six countries (also table A3), using median as opposed to mean expected inflation (also table A3), weighting observations by the number of respondents in the country-group (also table A3), separately per country (table A4), with different interactions of fixed effects (table A5), using 1-month, 4-month, and 12-month changes in expected inflation and electricity prices (table A6), and with Huber-White standard errors as well as clustered standard errors per demographic group (table A7). They confirm the baseline results.

Table 1: The impact of electricity prices on expected inflation

Revision of expectation	(1)	(2)	(3)	(4)	(5)	(6)
Change in electricity prices	1.163*** (0.305)	0.961*** (0.107)	0.983*** (0.243)	1.304*** (0.342)	1.154*** (0.304)	0.372** (0.181)
Change in electricity prices × Unanchoring	0.669*** (0.192)	0.220*** (0.063)	2.695*** (0.533)	1.613*** (0.434)	0.692*** (0.194)	0.146 (0.089)
Average past inflation	-0.097*** (0.026)	-0.103*** (0.008)	-0.104*** (0.024)	-0.092*** (0.025)	-0.096*** (0.026)	0.004 (0.079)
Observations	362756	2472	362756	362756	362756	362756
R^2	0.013	0.285	0.015	0.014	0.014	0.032
Country & group fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	No	No	No	No	No	Yes
Country-group fixed effects	No	No	No	No	Yes	No

Note: This table presents estimates of the regression in equation (1): $\Delta^6 \pi_{i,c,g,t}^e = \beta \Delta^6 e_{c,t} + \gamma \Delta^6 e_{c,t} \times \Delta^6 a_{c,g,t} + \alpha_c + \eta_g + \theta \bar{\pi}_{c,t-6} + \varepsilon_{i,c,g,t}$. Column (1) has the baseline estimates, (2) uses the average $\pi_{c,g,t}^e$ as the dependent variable, (3) uses as measure of unanchoring the deviation of long-run expected inflation from target, (4) uses anchoring at the country level only $a_{c,t}$, (5) includes country-group fixed effects, and (6) includes time fixed effects. In parentheses are standard errors clustered by month for the regressions using individual expectations.

Table 2 shows the estimated impact on expected inflation of an energy price shock. The first column still uses electricity prices per country and month. The difference from the first column in table 1 is that the series is now standardized, so that we can compare coefficients across the columns of this new table.

The second column uses instead the shift-share shock series with exogenous energy expenditure shares. The effect of a shock on expected inflation if there is no unanchoring is almost four times larger, while if there is unanchoring, the effect is almost twice larger. This is consistent with the use of exogenous shares dealing with the reverse causality that would be biasing the coefficients downwards in the first column.

The third and fourth column use exogenous time-series variation in oil prices and wind, respectively. In the first case, the impact of the energy shock remains large, but unanchoring no longer plays a role, while in the second case both effects go to zero. Column five explores what might be going on by increasing the horizon to 12 months. The effect of an oil-driven energy shock on expected inflation almost doubles, with the share due to unanchoring now being statistically significant. This suggests that the impacts may accumulate over time, which we inspect next.

Table 2: The impact of energy shocks on expected inflation

Revision of expectation	(1)	(2)	(3)	(4)	(5)
Energy price shock	0.145** (0.057)	0.580*** (0.081)	0.348*** (0.101)	-0.086 (0.100)	0.607** (0.262)
Energy price shock × Unanchoring	0.267*** (0.033)	0.159*** (0.037)	0.006 (0.067)	0.025 (0.079)	0.115** (0.053)
Average past inflation	-0.103*** (0.023)	-0.017 (0.025)	-0.111** (0.041)	-0.132*** (0.029)	-0.041 (0.167)
Observations	362756	362756	305037	362224	197950
R^2	0.017	0.024	0.015	0.010	0.027

Note: This table presents estimates of the regression equation $\Delta^h \pi_{i,c,g,t}^e = \beta \Delta^h z_{c,t} + \gamma \Delta^h z_{c,t} \times \Delta^h a_{c,g,t} + \alpha_c + \eta_g + \theta \bar{\pi}_{c,t-6} + \varepsilon_{i,c,g,t}$ where the first four columns use different measures of $z_{c,t}$. The energy shocks are, in order: the change in HICP electricity prices by country, the h -month change in EA-side HICP electricity times country-specific electricity expenditure weights in 2019, OPEC supply shocks to oil prices cumulated over h months times country-specific expenditure weights in 2019, and the h -month change in wind-source electricity generation. The first four columns set $h = 6$, while the fifth column uses the oil shocks with $h = 12$. In parentheses are standard errors clustered by month.

Figure 1 shows the dynamic effects from the local projections following each of the four energy shocks. In black-bold are pooled estimates that leave out the anchoring interaction term (their confidence bands are in the appendix), while the other two series and their confidence bands show the estimates with below and above average unanchoring. Across all shocks, the impact is negligible in the first four months, but then builds up, reaching between 17bp and 60bp twelve months later. After 12 months, all the estimates approach zero.

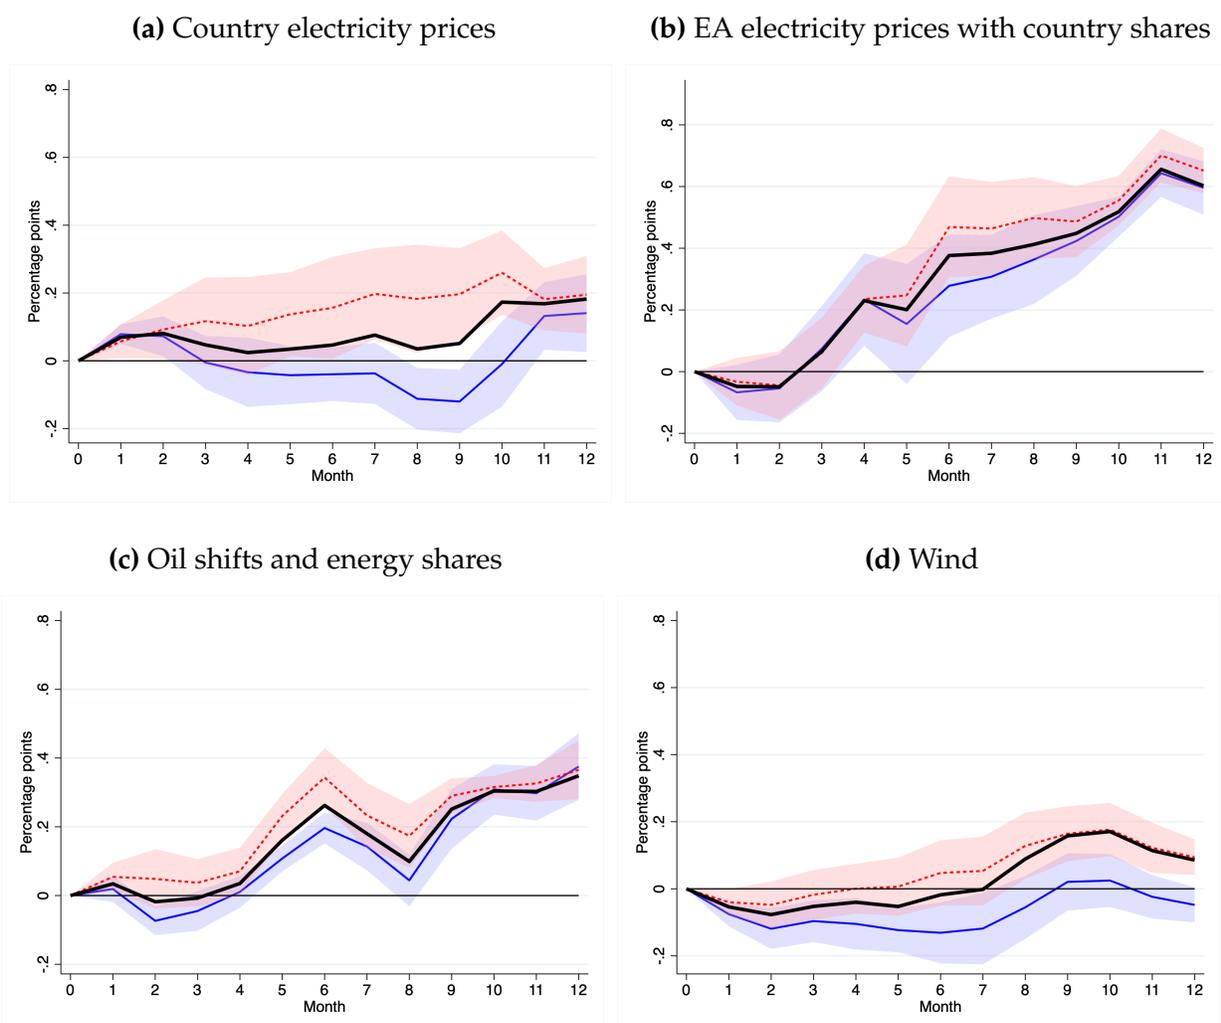
For all the shocks, more unanchored expectations lead to a larger impact of energy prices on expected inflation. Depending on the horizon considered, higher than average unanchoring can as much as double this impact.

The estimates for wind shocks explain the results in table 2. Cumulated over either 6 or 12 months the effects are indeed small and statistically insignificant. The local projection reveals that this is because their impact is only sizeable 8 months after the shock.

4 Four uses of the estimates

Armed with our baseline estimates, we address four related issues.

Figure 1: Impulse response of expected inflation to a shock in energy prices



Note: Local projection of average expected inflation within a region and group on 3-month cumulated energy price shock, controlling for inflation, country and group fixed effects, pooled across states (thick black line), when unanchoring is higher (red dashed line) or lower (blue solid line) in the previous 6 months than average for the country and demographic group. The shocks are scaled by their standard deviation. The shock in panel (a) is the change in electricity price by country and time. The shock in panel (b) is the time-varying EA-wide electricity price times the country-varying expenditure shares. The shock in panel (c) is time-varying oil OPEC supply shocks times the country-varying expenditure shares. The shock in panel (d) is to the country-time contribution of wind to the production of electricity. Standard errors are clustered by country.

4.1 How much of the increase in expected inflation in 2021-22 was due to higher energy prices?

Between May 2021 (when inflation was on target) and one year later, expected inflation on average across all the households, groups and countries, increased by 2.9 percent-

age points. Aggregating the fitted values from our baseline equation (1), which explains expected inflation using past inflation and energy prices, predicted expected inflation would have risen by a meager 0.53 percentage points.

Moreover, most of this increase is explained by the rise in past inflation. The R^2 of a partial regression isolating the contribution of energy prices alone to these predicted values is 0.39. It falls to 0.24 if we focus on the six major countries during the whole sample period, and this is already starting from only 0.53/2.9 explained. The conclusion is that energy prices are important for inflation expectations but by themselves they fall well short of explaining the movements in expected inflation during the inflation disaster.

4.2 How sensitive was expected inflation to electricity prices during the sample?

Figure 2 uses the estimates in table 1 to plot at each date in time, the impact of a doubling of electricity prices over the following 6 months. That is, it plots $(\beta + \gamma \Delta^6 a_t) \log(2)$, where the time variation comes from the unanchoring, averaged across countries and groups. Because of the noise in month-to-month variations in anchoring, we plot a 3rd-order centred moving average.

The estimates show that EA expected inflation was significantly more sensitive to energy prices at the start of 2022 than it was at the start of the sample. The scar of the inflation disaster is noticeable. Reassuringly, the re-anchoring of inflation expectations that came with the tightening of monetary policy and the fall in inflation in 2023 have reduced the impact of energy prices today to their pre-disaster level.

4.3 Estimates using US data

The limitations of the MSC data constrain our empirical strategy. Most importantly, there is no index c for countries. The MSC splits the respondents into only four large US regions, but these cover many states with different energy prices that are difficult to aggregate, and with less regional variability than in the EA. Therefore, there is only cross-group variation over the eight demographic and socio-economic groups. Also, the small sample means that the measures of anchoring are very noisy, as disagreement is calculated over groups that half of the times have fewer than 50 respondents, and even as few as 4.

Energy prices, e_t , now stand for log retail gasoline prices, calculated by the Energy Information Administration. They are not directly comparable to electricity prices, and

Figure 2: The time-varying impact of electricity prices on expected inflation



Note: The figure plots the predicted effect on EA average expected inflation from doubling electricity prices over the following 6 months, calculated as a function of the extent of unanchoring over the same period, using the coefficients estimated in column 1 of table 1. In red are estimates using disagreement about long-run expected inflation as a measure of unanchoring, and in green are those using the absolute difference between expected long-run inflation and target.

are not as important in household budgets. Because the sample now covers a longer period of time, from 1993:4 to 2023:7, the concerns with omitted monetary policy and other aggregate variables are stronger.

With all these caveats in mind, table A8 in the appendix shows estimates of equation (1) and also compares them to Coibion and Gorodnichenko (2015). A 1% rise in US gasoline prices raises expected inflation by 2.72 to 3.92 bp, significantly more than in the Euro area but also with wider confidence bands. The effects of the (poorly measured) unanchoring series are imprecise and unstable across specifications. Energy prices still explain little of the variation in overall expected inflation, with our regression predicting less than one quarter of the observed increase between March of 2021 and 2022, and with partial R^2 's from energy prices ranging between 0.01 and 0.24.

4.4 Interpreting the findings as state-dependent inattention

Households choose how much attention to devote to inflation. It is likely little. But, as long as it is positive, it is to be expected that higher energy prices raise expected inflation. After all, the change in actual inflation is positively correlated with a rise in energy prices. Our finding of a statistically positive but small effect is consistent with households being rational in this weak sense, while also being inattentive and barely noticing that rise.

Understanding the role of unanchoring requires a little more work. Take an agent choosing a variable x with an objective function that depends on other relevant state variables z : $V(x, z)$. She has limited information, which prevents her from observing these variables. Each agent's choices deviate from the average $\bar{x}(z)$ by an individual random error ε that reflects her idiosyncratic noisy signals.

The appendix shows that under some assumptions—in a linear-quadratic approximation, if the costs of attention are the entropy of her decision rule, and with a constant λ marginal cost of an extra bit of attention—then:

$$x = \bar{x}(z) + \sqrt{\left| \frac{\lambda \bar{x}'(z)}{v(z)} \right|} \varepsilon, \quad (3)$$

where ε has a standard normal distribution and $v(z) = -\partial^2 V(\cdot) / \partial x \partial z(\bar{x}(z), z)$. Intuitively, the larger is the cost of attention, λ , the less attention she pays, and so the larger are the errors she makes. In the other direction, the larger is the impact of errors from inattention on her well-being, the more attention she will pay, captured by a lower $\bar{x}'(z) / v(z)$.

Consider then the case where the agent is asked about her expectations of inflation and that one of the relevant state variables is the price of energy. Since for a normal variable the square of the interquartile range is two times the variance, it follows that:

$$\frac{\partial \pi^e}{\partial e} = \left(\frac{v(e)}{2\lambda} \right) a^2(e). \quad (4)$$

More unanchoring is associated with a larger response of inflation expectations to an energy prices, just as we found in the data. The intuition is that when expectations are very sensitive to shocks, then the mistakes in forming those expectations must not be so costly. Therefore, she is less attentive, and so there is more unanchoring.

As part of the energy cycle, anchoring and the sensitivity of expectations will fluctuate. Flynn and Sastry (2024) incorporate this model of attention in a business-cycle

framework, and note that this will lead firms to under and overproduce, depending on whether energy prices are high or low, creating wedges. Energy shocks will then generate endogenous attention wedges that will appear as markup shocks in a Phillips curve.

5 Conclusion

Ever since the 1970s, when large oil price shocks came with a sharp and persistent rise in inflation, economists have been studying the connection between these two variables. An important, but still poorly understood, channel is through inflation expectations. An often-repeated fact is that household expectations of inflation and energy prices are strongly correlated. Sometimes, this is used to assert that this channel is strong, and other times to dismiss expectations data through the same “see through principle” that justifies dismissing energy shocks. Yet, the time-series correlation omits the role of monetary policy in responding to and affecting both variables, it is not causal, it explains little of the variability of expected inflation, and the size of its estimates differ considerably across articles. The 1970s US data also does not support its common use since back then inflation expectations rose well before the oil price shocks (Reis, 2021).

This paper followed in the footsteps of a wave of research in empirical macroeconomics that has used cross-regional variation within a currency union to make progress on identification. Taking advantage of the recently-released household survey of expectations in the EA that has many more respondents identified by country and group, and of the large variability in energy prices in the 2020-23 period, we provided new estimates of the impact of energy prices on expectations. We used the cross-sectional variation to build new exogenous shocks to energy prices, and found that they have a sizeable impact on expected inflation, that is larger when inflation expectations are unanchored. Yet, the energy shocks of 2021-23 explain a small share of the rise in expected inflation. At the end of 2023, the impact of energy price shocks is back to what it was at the start of the sample. These results are consistent with noisier estimates from the US and with theories of state-dependent attention.

References

- Aastveit, Knut Are, Hilde C. Bjørnland, and Jamie L. Cross. 2023. "Inflation Expectations and the Pass-Through of Oil Prices." *Review of Economics and Statistics*, 105(3): 733–743.
- Acharya, Viral, Matteo Crosignani, Tim Eisert, and Christian Eufinger. 2023. "How do Supply Shocks to Inflation Generalize? Evidence from the Pandemic Era in Europe." *CEPR discussion paper 18530*.
- Angeletos, G.-M., and C. Lian. 2016. "Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination." In *Handbook of Macroeconomics*. Vol. 2, , ed. John B. Taylor and Harald Uhlig, Chapter 14, 1065–1240. Elsevier.
- Armantier, Olivier, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar. 2016. "How Do People Revise Their Inflation Expectations?" *Liberty Street economics blog*.
- Arora, Vipin, Pedro Gomis-Porqueras, and Shuping Shi. 2013. "The Divergence Between Core and Headline Inflation: Implications for Consumers' Inflation Expectations." *Journal of Macroeconomics*, 38: 497–504.
- Binder, Carola, and Christos Makridis. 2022. "Stuck in the Seventies: Gas Prices and Consumer Sentiment." *Review of Economics and Statistics*, 104(2): 293–305.
- Binder, Carola Conces. 2018. "Inflation Expectations and the Price at the Pump." *Journal of Macroeconomics*, 58: 1–18.
- Bonomo, Marco, Carlos V. Carvalho, Stefano Eusepi, Marina Perrupato, Daniel Abib, and João Ayres Silvia Matos. 2024. "Abrupt Monetary Policy Change and Unanchoring of Inflation Expectations." *Journal of Monetary Economics*.
- Celasun, Oya, Roxana Mihet, and Lev Ratnovski. 2012. "Inflation Expectations and Commodity Prices in the United States." *IMF Working Paper 89/2012*.
- Coibion, Olivier, and Yuriy Gorodnichenko. 2015. "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation." *American Economic Journal: Macroeconomics*, 7(1): 197–232.
- Conflitti, Cristina, and Riccardo Cristadoro. 2018. "Oil prices and Inflation Expectations." *Bank of Italy Occasional Paper 423*.
- Feldkircher, Martin, and Pierre L. Siklos. 2019. "Global Inflation Dynamics and Inflation Expectations." *International Review of Economics and Finance*, 64: 217–241.
- Flynn, Joel, and Karthik Sastry. 2024. "Attention Cycles." <https://dx.doi.org/10.2139/ssrn.3592107>.
- Fofana, Salome, Paula Patzelt, and Ricardo Reis. 2024. "Facts about Disagreement in Household Inflation Expectations." In *Handbook of Inflation*. , ed. G. Ascari and R. Trezzi, Chapter 16, forthcoming. Edward Elgar Publishing.

- Gagliardone, Luca, and Mark Gertler.** 2023. "Oil Prices, Monetary Policy and Inflation Surges." *NBER working paper 31263*.
- Hajdini, Ina, Edward Knotek II, John Leer, Mathieu Pedemonte, Robert Rich, and Raphael Schoenle.** 2024. "Indirect Consumer Inflation Expectations: Theory and Evidence." *Journal of Monetary Economics*, forthcoming.
- Hensel, Jannik, Giacomo Mangiante, and Luca Moretti.** 2023. "Carbon Pricing and Inflation Expectations: Evidence from France." *CESifo working paper 10552*.
- Känzig, Diego R.** 2021. "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements." *American Economic Review*, 111(4): 1092–1125.
- Kilian, Lutz, and Xiaoqing Zhou.** 2022. "Oil Prices, Gasoline Prices, and Inflation Expectations." *Journal of Applied Econometrics*, 37(5): 867–881.
- Kilian, Lutz, and Xiaoqing Zhou.** 2024. "Oil Price Shocks and Inflation." In *Handbook of Inflation*, ed. G. Ascari and R. Trezzi, Chapter 20, forthcoming. Edward Elgar Publishing.
- Kumar, Saten, Hassan Afrouzi, Olivier Coibion, and Yuriy Gorodnichenko.** 2015. "Inflation Targeting Does Not Anchor Inflation Expectations: Evidence from Firms in New Zealand." *Brookings Papers on Economic Activity*, 151–225.
- Mankiw, N. Gregory, and Ricardo Reis.** 2010. "Imperfect Information and Aggregate Supply." In *Handbook of Monetary Economics*. Vol. 3A, ed. Benjamin Friedman and Michael Woodford, Chapter 5, 183–230. Elsevier: North Holland.
- Miyamoto, Wataru, Thuy Lan Nguyen, and Dmitriy Sergeyev.** 2024. "How Oil Shocks Propagate: Evidence on the Monetary Policy Channel." *CEPR discussion paper 18755*.
- Pfauti, Oliver.** 2023. "The Inflation Attention Threshold and Inflation Surges." *UT Austin manuscript*.
- Reis, Ricardo.** 2021. "Losing the Inflation Anchor." *Brookings Papers on Economic Activity*, 307–361.
- Reis, Ricardo.** 2023. "The Burst of High Inflation in 2021-22: How and Why Did We Get Here?" In *How Monetary Policy Got Behind the Curve—And How to Get it Back*, ed. Michael Bordo, John Cochrane and John Taylor, 203–252. Hoover Institution Press.
- Trehan, Bharat.** 2011. "Household Inflation Expectations and the Price of Oil: It's Déjà Vu All Over Again." *FRBSF Economic Letter 2011-16*.
- Vlieghe, Gertjan.** 2024. "Core Strength: International Evidence on the Impact of Energy Prices on Core Inflation." *LSE - CFM manuscript*.
- Wehrhofer, Nils.** 2023. "Energy Prices and Inflation Expectations: Evidence from Households and Firms." *Deutsche Bundesbank Discussion Paper 28/2023*.
- Wong, Benjamin.** 2015. "Do Inflation Expectations Propagate the Inflationary Impact of Real Oil Price Shocks?: Evidence from the Michigan Survey." *Journal of Money, Credit and Banking*, 47(8): 1673–1689.

Appendix

A Literature Review

The literature has so far not been able to pin down a significant and stable causal link from energy prices to household expected inflation. We survey it here, grouping it by their approach.

Correlations: It is routine in policymaker speeches or as side remarks in academic papers to note the high correlation between households' expected rate of change in consumer prices and the level of oil prices. This is between 0.58 and 0.82 in the US data depending on the decade. However, correlating changes on consumer prices with the level of oil prices is problematic. The units are not comparable, and oil prices have a unit root while inflation does not.

An older literature (Trehan, 2011, Arora, Gomis-Porqueras and Shi, 2013) calculated time-series correlations between households' average expected inflation and core or non-core inflation. These correlations between macroeconomic aggregates are hard to interpret. They are unstable across samples and countries, and are mutually correlated with so many other macroeconomic aggregate time series that controlling for any number of them easily flips the sign of the partial correlation.

From oil prices to expectations: Moving beyond correlations, Coibion and Gorodnichenko (2015) is a classic reference. It used the monthly cross-section of individual inflation expectations over the year ahead from the MSC and regressed it on the rate of change of wholesale oil prices over the past 6 months between 1978 and 2013. The baseline estimate is that a 1% increase in oil prices raises average expected inflation by 1.6bp.

The literature that followed significantly revised this estimate downwards. Binder (2018) replaced wholesale oil prices by gas prices paid at the pump and the estimate fell to 1bp.

Because the oil price is the same for all, and because the micro data is mostly a repeated cross-section with households interviewed only twice, the Coibion and Gorodnichenko (2015) estimate used almost entirely time-series variation. With time fixed effects, there is little variation left, and the estimates become quantitatively and statistically indistinguishable from zero. Armantier et al. (2016) made this observation using the FRB New York Survey of Consumer Expectations.

A complementary literature looked at the impact of changes in oil prices on the ex-

pected inflation over the next 5 years in the MSC (Celasun, Mihet and Ratnovski, 2012, Binder, 2018), which arguably may respond less to other short-term shocks. These measures of long-term expectations move much less than the one-year-ahead ones. With this limited time-series variation, it is even harder to estimate the effect precisely, and results are varied. The few estimates that are statistically significant and different from zero point to a negligible impact of oil prices on expected inflation.

Causal effects that control for monetary policy and aggregate demand: Kilian and Zhou (2022) criticised the Coibion and Gorodnichenko (2015) time-series regression of average expected inflation on the rate of growth of oil prices because inflation expectations will drive spending and so affect the demand for oil. Using a combination of sign and zero restrictions in a vector autoregression linking inflation, inflation expectations and gas prices, it estimated that a 1% increase in gasoline prices raises expected inflation by 3bps on the month of impact, but within 5 months this effect falls to zero.

A related literature uses structural vector autoregression to put an upper bound on the contribution of shocks to the price of oil that are orthogonal to some shocks to inflation or inflation expectations. Kilian and Zhou (2022) found that oil prices account for at most 42% of the variation in inflation expectations, while Wong (2015) found an even weaker connection between oil prices and expected inflation, and Aastveit, Bjørnland and Cross (2023) found an impact response in between the previous two, but one that only more slowly reverts to zero. Results are likewise imprecise with the time-series variation in the expectations from surveys of professionals or firms instead of households. For instance, Feldkircher and Siklos (2019) found a persistent and large impact of oil prices in the consensus survey, Conflitti and Cristadoro (2018) found close to no effect in the Survey of Professional Forecasters, and Hensel, Mangiante and Moretti (2023) found large persistent effects on firms following shocks to carbon pricing.

Känzig (2021) isolated the causal impact of oil prices on multiple macro variables, including median MSC inflation expectations, by constructing high-frequency changes in the oil price expectations reflected in oil price futures around OPEC production announcements. It found that a 1% increase in oil prices raises median expected inflation one month later by 1.8bp. However, within 6 months, the point estimate falls to 0.2bp, and is statistically insignificant from then onwards.

In the same vein, Miyamoto, Nguyen and Sergeyev (2024) locally projected mean expected inflation in Japan on the exogenous oil shocks from Känzig (2021) and also controlling for monetary policy, by separating periods where nominal interest rates were at

the zero lower bound or not. In normal times, they find that oil prices do not have a statistically significant effect on expectations, but at the zero lower bound, a 1% increase in real oil prices raises expected inflation by 3bp after one quarter. The effect falls by more than half after 6 months, and is not statistically significant after that.

Cross-sectional variation: Like us, two recent papers exploited cross-sectional variation. Binder and Makridis (2022) used state-level variation in real gasoline prices, but it measures their impact on indices of consumer sentiment, as opposed to inflation expectations. Wehrhofer (2023) used cross-household variation on when electricity contracts are renewed to find that, in a context of rising energy prices, renewals raise expected inflation by 1.8bp. However, lacking information on how much the electricity price rose with the new contract, it cannot estimate the coefficient of interest for our first question.

Hajdini et al. (2024) is closer by regressing expected inflation from a new survey that asks for cost-of-living inflation on gas prices times the share of households in a US county that use their own car for commuting. Their regressor therefore has the same flavour as the first of our three energy shocks, and they find large effects, close to the upper end of our confidence bands for the US. However, they try to explain levels, as opposed to changes, in expectations. Moreover, they focus on contemporaneous weekly-effects. Our dynamic regressions, and the estimates of the regression with different horizons, both suggest that the effects change considerably over time. Also, they do not investigate the role of unanchoring.

Studies before had used the very limited cross-sectional variation to show that estimates could vary across states and groups. Coibion and Gorodnichenko (2015) had a version of their baseline estimate broken by states and groups. More explicitly, Binder (2018) used the 4 regions in MSC to separately estimate the impact of oil prices on inflation expectations, and found that these line up with expenditure shares, and also correlated expected gas prices and expected headline inflation across the regions.

Anchoring of expectations: Kumar et al. (2015) and Bonomo et al. (2024) discuss a series of measures of anchoring. We further exploit the systematic differences across disagreement documented by Fofana, Patzelt and Reis (2024) to maximize the cross-sectional variation that we extract from the data. Bonomo et al. (2024) discuss an episode in 2021 in Brazil where expectations became quickly unanchored linked to a change in monetary policy. We find a smaller but still significant amount of unanchoring during 2021-22 in the EA linked to oil prices.

Theories of inattention naturally link unanchoring to responsiveness to shocks (An-

geletos and Lian, 2016, Mankiw and Reis, 2010). We use the model of state-dependent attention of Flynn and Sastry (2024) to connect to our estimates and draw implications for attention cycles. Similar conclusions would hold in the model of Pfauti (2023): following the energy shock, inflation surges, the public’s attention to inflation rises, and negative supply shocks become more inflationary. Pfauti (2023) estimates that between a low attention and a high-attention regime the impact of a negative supply shock on inflation expectations doubles, which is consistent with our estimates.

The role of energy prices in the inflation disaster of 2021-23: Finally, a recent literature has tried to explain the 2021-23 inflation disaster with measures of inflation expectations (Reis, 2023), measures of energy prices (Gagliardone and Gertler, 2023), interactions between the two (Acharya et al., 2023), and propagation over time (Vlieghe, 2024). In spite of some brave attempts, quantifying the relative contribution of expectations and energy prices (and other supply shocks) is hard since both affect each other and are related to other major macroeconomic aggregates. We make progress by trying to isolate specific channels that can be pinned.

Kilian and Zhou (2024) use a SVAR with recent data to revise the impact of a 1% gas price shock down to at most 1.2 basis points, still statistically insignificant within 6 months, and explaining only 28% of the variation in the data. Our paper brings some micro data leading to larger estimates, and yet still reaching the conclusion that energy prices do not explain the increase in expected inflation during this period.

B Data: additional information

Our data on inflation expectations comes from the ECB, and was downloaded on 6 February 2024. It ends in December of 2023, and starts in April of 2020 for six countries—Belgium, France, Germany, Italy, the Netherlands, and Spain—and in April of 2022 for another five—Austria, Finland, Greece, Ireland, and Portugal. We censor the individual response if individual point forecasts for inflation exceed 20% in absolute value to ensure robustness to outliers.

Table A1 provides additional information per country-group, with the averages over time of: the number of respondents per group, expected inflation, disagreement, and the average 6-month change in expected inflation. This makes clear that there is significant variation in the cross-section, which our estimates rely on.

Figure A1 plots actual inflation, average expected inflation, and disagreement accord-

ing to the interquartile range, with each of the six major countries in a separate panel, to highlight the time series movements during this period. This figure shows the importance of controlling for country fixed effects as well as for the level of inflation in the regression.

Our data for energy prices from HICP suffers from revisions to the methods to calculate them in the Netherlands in June of 2023. We use their research series to have a consistent series throughout our sample. There was also a change in the methods used for the HICP in Spain from January of 2023 onwards, but there is no research series available.

We also use three alternative series for energy prices. They are: the HICP energy price index that includes all energy prices, not just electricity; wholesale electricity prices from the European Network of Transmission System Operators for Electricity (ENTSO-E) collected by Ember; and the household energy price index (HEPI) from <https://www.energypriceindex.com>, commissioned by VaasaETT and funded by Energie-Control Austria and the Hungarian Energy and Public Utility Regulatory Authority MEKH, copyright 2024 VaasaETT Ltd. Their respective correlations with our HIPC series during this period are 0.60, 0.37, and 0.59, so they are relatively close to each other.

Finally, the data on wind speeds, used to check the correlation with our wind electricity generation series, are sourced from Visual Crossing. We average the daily mean wind speed by EA country and month.

C Alternative specifications

This appendix shows alternative specifications. They complement the baseline results, having the same sign and statistically significant, without overturning them.

Table A2 reproduces the first and second columns of table 1, but replaces the energy estimates with the three alternatives we discussed above. Interestingly, wholesale prices do not seem to be salient in the sense of moving expected inflation as much. The two measures of prices paid by consumers give similar results to our baseline case.

Looking instead to robustness to anchoring, table A3 shows the estimates for different specifications that exclude the anchoring variable, that include it as a separate regressor, that restrict the sample to a balanced panel of 6 countries, and that pool the individual observations via the median as opposed to the average, or weight the country-group averages by their respective number of respondents. The main inferences on β are relatively robust to these different specifications. At the same time, they highlight the importance

of taking disagreement into account when investigating micro data on expected inflation.

Table A4 shows the regression equation estimated separately for each country. The estimates vary significantly across countries showing the importance of exploiting this cross-country variation.

Table A5 adds different interactions of fixed effects to the baseline regression with time fixed effects, using either country-group, country-time, or group-time fixed effects, each for individual and mean expectations. Results are similar to the version in column 6 of table 1 using only country, group and time effects separately, except for electricity prices without unanchoring when adding country-time fixed effects, since electricity prices only vary on this level.

Table A6 compares the results across choices of the revision horizon h . The data is much more noisy so, as expected, with $h = 1$ the R^2 falls and the standard errors rise. At the same time, in size and sign, the estimates remain similar. With $h = 12$, as opposed to 6 months, on the one hand, some of the effect may reverse with the horizon, as the estimate is a little lower at 0.89bp. On the other hand, because now we consider a 12-month change in anchoring as well, the impact of this prolonged unanchoring is larger at 0.88bp.

Table A7 repeats table 1 but now lists three alternatives as the standard errors, the standard ones we produce, together with Huber-White adjustment for heteroskedasticity and clustering per demographic group. The errors rise, but the two key estimates of interest remain statistically significant at conventional significance levels.

Figure A2 shows the impact of exogenous shocks on expected inflation with error bands, omitting the anchoring dummy variable. In the first two rows, effects are, as expected, in between the ones in figure 1, and statistically significant. The figure also shows, in the bottom row, the impact of the oil shift-share, but now using the average expenditure shares between 2015-19, or the network cost of electricity paid by households in 2019. The effects are very similar.

D US estimates

Table A8 shows the results of estimating equation (1) on US data. A 1% rise in gasoline prices raises expected inflation by 3.77bp, significantly more than in the Euro area, while the effects of the (poorly measured) unanchoring are not statistically distinguishable from zero. This last conclusion is not robust though, as slight changes in the specification (like the choice of h) lead to large changes on the coefficient on unanchoring. For instance,

column 2 simply lags the unanchoring variable, and its extra boost rises to a large and statistically significant coefficient of 0.47bp. Column 3 further confirms this by using the alternative measure of anchoring based on the distance from target. The coefficient on anchoring is now quite large, but very imprecise, while the impact of oil prices falls by one third to 2.72bp. Column 4 shows that aggregating across individuals delivers a large R^2 , confirming the role played by time-series variation, and a coefficient of 3.92bp.

Table A9 further investigates the robustness of these results. The first column just repeats the first column of table A8. The second column instead: (i) imposes $\gamma = 0$ so there is no consideration of anchoring, (ii) imposes $\theta = 0$ so there is also no control of actual past inflation, (iii) uses wholesale oil prices as opposed to retail gas prices. This is the specification in Coibion and Gorodnichenko (2015), updated to the more recent data, and censoring observations of expected inflation above 20%, while they censored changes of more than 15%. The estimate of β of 1.98bp is close to their estimates. (Using the same sample and censoring rule as theirs, we replicate their results, with an estimate of 1.68bp.)

The third to fifth column show the impact of each of these changes separately. Switching to consumer gasoline prices, as opposed to wholesale prices, raises β to 3.91, including a control for past inflation lowers it to 1.87, and including the anchoring variable but only at the national level raises β to 2.19.

E Optimal actions with inattention

The result that a linear quadratic approximation of a rational inattention problem leads to normally distributed errors in actions is standard in the literature. We follow Flynn and Sastry (2024)'s formulation, which modifies the cost function to depend on the entropy of the decision function, as opposed to the mutual information between prior and posterior, and so eliminates the influence of the prior on the final solution.

By choosing both how much attention to pay, and how to behave given her imperfect signals, the agent is choosing a stochastic decision rule $p(x|z)$ that maximizes her expected payoff over the possible z 's:

$$\mathbb{E}_z \left[\int (V(x, z) - \lambda \log(p(x|z))) p(x|z) dx \right]. \quad (\text{A1})$$

The second term is the cost of paying attention, which is written here in terms of the entropy of the decision rule and where λ is the marginal cost of an extra bit of attention.

The agent faces the constraint that the function $p(x|z)$ is everywhere non-negative and it integrates to one over all the actions.

Let $\bar{x}(z)$ be the solution to $(\partial V(x, z)/\partial x)(\bar{x}(z), z) = 0$. This need not be the optimal full information solution. After all, $V(x, z)$ may well be an indirect utility function already incorporating other distortions, including distortions to the formation of beliefs. It is simply the solution if $\lambda = 0$ and the agent did not have the rational inattention problem that leads to the dispersion of expectations.

Using the implicit function theorem:

$$\bar{x}'(z) = -\frac{\partial^2 V(x, z)/\partial x \partial z}{\partial^2 V(x, z)/\partial x^2}(\bar{x}(z), z) \quad (\text{A2})$$

It then follows that a quadratic approximation of the objective function around $\bar{x}(z)$ is:

$$\begin{aligned} V(x, z) &\approx V(\bar{x}(z), z) + 0.5 \left[(\partial^2 V(x, z)/\partial x^2)(\bar{x}(z), z) \right] (x - \bar{x}(z))^2 \\ &\propto \left(\frac{v(z)}{\bar{x}'(z)} \right) (x - \bar{x}(z))^2 \end{aligned} \quad (\text{A3})$$

where recall that we defined $v(z) = -\partial^2 V(x, z)/\partial x \partial z(\bar{x}(z), z)$.

Letting z have a density $f(z)$, the optimization problem has the Lagrangean:

$$\begin{aligned} \mathcal{L} &= \int_z \int_x \left((v(z)/\bar{x}'(z)) (x - \bar{x}(z))^2 - \lambda \log(p(x|z) + \kappa(x, z)) \right) p(x|z) dx f(z) dz \\ &+ \int_z \gamma(z) \left(\int_x p(x|z) dx - 1 \right) f(z) dz \end{aligned} \quad (\text{A4})$$

where $\kappa(x, z)$ are the Lagrange multipliers for each choice and state so that their probability is non-negative, and $\gamma(z)$ are the Lagrange multipliers so that at every state, the choice probabilities integrate to 1.

The first-order condition for optimality is:

$$(v(z)/\bar{x}'(z)) (x - \bar{x}(z))^2 + \kappa(x, z) + \gamma(z) = \lambda \log(p(x|z) + \lambda) \quad (\text{A5})$$

Integrating over x and using the constraint that $\int_x p(x|z) dx = 1$, this optimality condition becomes:

$$p(x|z) = \frac{\exp\left(\frac{(x - \bar{x}(z))^2}{|\lambda \bar{x}'(z)/v(z)|}\right)}{\int_x \exp\left(\frac{(x - \bar{x}(z))^2}{|\lambda \bar{x}'(z)/v(z)|}\right) dx} \quad (\text{A6})$$

From this it follows that x follows a normal distribution with mean $\bar{x}(z)$ and with variance $|\lambda\bar{x}'(z)/v(z)|$, just as we wrote in equation (3).

The standard deviation of expectations across agents who each make an idiosyncratic error is $\sqrt{|\lambda\bar{x}'(z)/v(z)|}$. In turn the interquartile range of a standard normal distribution is 1.34898. Therefore:

$$a(z) = 1.34898\sqrt{|\lambda\bar{x}'(z)/v(z)|} \Rightarrow \bar{x}'(z) = \left(\frac{v(z)}{2\lambda}\right) a(z)^2. \quad (\text{A7})$$

Letting $\pi^e = x$ and $e = z$ gives the result in the text.

Table A1: Descriptive statistics by country and group

Country	Group	Number of respondents	Inflation expectation	Dis-agreement	6-month revision	Country	Group	Number of respondents	Inflation expectation	Dis-agreement	6-month revision
AT	1	210	5.66	4.90	-1.07	FR	1	273	2.98	4.45	-0.14
AT	2	143	5.12	4.47	-1.02	FR	2	154	2.83	3.88	0.08
AT	3	75	5.48	3.86	-0.93	FR	3	443	3.25	4.09	0.10
AT	4	79	5.50	2.95	-0.41	FR	4	489	3.30	3.67	0.06
AT	5	268	5.41	5.75	-1.37	FR	5	377	3.13	4.62	-0.09
AT	6	126	5.16	5.00	-1.62	FR	6	117	3.14	4.61	0.02
AT	7	85	5.52	5.69	-1.17	FR	7	540	3.51	4.62	0.07
AT	8	78	4.63	5.35	-0.97	FR	8	373	3.49	4.11	-0.18
BE	1	148	3.59	4.40	-0.03	IE	1	94	4.93	6.31	-1.93
BE	2	74	3.66	4.13	0.05	IE	2	38	4.02	5.94	-1.88
BE	3	131	3.69	3.95	0.03	IE	3	132	5.57	4.54	-1.15
BE	4	157	3.31	2.84	0.06	IE	4	120	5.34	4.80	-1.28
BE	5	182	4.11	5.24	-0.08	IE	5	163	5.36	8.03	-1.45
BE	6	59	4.17	3.89	-0.07	IE	6	49	4.62	6.30	-1.57
BE	7	159	4.00	4.20	0.21	IE	7	219	5.37	6.87	-1.30
BE	8	125	3.95	4.19	0.04	IE	8	171	5.51	5.88	-1.25
DE	1	419	3.44	4.10	0.15	IT	1	488	4.17	5.72	0.03
DE	2	199	2.81	3.63	0.04	IT	2	240	4.23	4.92	0.01
DE	3	348	3.18	4.04	0.31	IT	3	308	3.94	5.26	0.19
DE	4	421	3.16	3.72	0.16	IT	4	324	4.01	4.35	0.07
DE	5	495	3.46	4.52	0.06	IT	5	636	4.57	6.73	-0.04
DE	6	201	3.10	4.01	0.22	IT	6	195	4.87	6.50	0.05
DE	7	306	3.09	3.97	0.22	IT	7	392	4.16	5.33	-0.03
DE	8	266	3.18	3.99	0.11	IT	8	268	4.38	5.26	-0.13
GR	1	137	6.32	10.00	-0.24	NL	1	161	3.63	3.43	-0.03
GR	2	42	7.03	8.83	0.17	NL	2	80	3.84	2.81	0.03
GR	3	163	8.27	11.28	0.07	NL	3	88	3.44	2.82	-0.08
GR	4	170	6.41	9.72	-0.05	NL	4	147	3.58	2.22	0.22
GR	5	148	7.14	10.09	-0.44	NL	5	230	3.99	3.98	-0.03
GR	6	45	6.53	8.49	0.06	NL	6	75	3.84	3.20	0.04
GR	7	204	6.88	9.76	-0.23	NL	7	113	3.55	3.72	0.01
GR	8	124	6.89	9.48	-0.14	NL	8	103	3.76	3.30	-0.01
ES	1	372	3.55	5.23	0.17	PT	1	188	5.22	7.04	-1.01
ES	2	168	3.07	4.28	0.35	PT	2	80	5.45	6.22	-1.26
ES	3	340	3.40	4.34	0.09	PT	3	117	5.04	5.50	-0.79
ES	4	487	3.51	3.40	0.17	PT	4	159	4.74	4.59	-1.27
ES	5	435	3.71	5.90	0.34	PT	5	169	4.92	7.66	-1.36
ES	6	115	4.00	5.71	0.12	PT	6	46	6.00	6.06	-1.19
ES	7	466	3.69	5.52	0.11	PT	7	205	5.24	6.81	-1.02
ES	8	364	3.69	4.69	0.18	PT	8	151	5.49	5.63	-0.88
FI	1	189	4.78	4.92	-1.31						
FI	2	81	4.08	4.53	-1.19						
FI	3	111	4.72	4.80	-1.24						
FI	4	133	4.48	3.77	-1.21						
FI	5	187	4.79	5.44	-1.39						
FI	6	84	5.01	4.76	-1.34						
FI	7	177	5.52	4.78	-1.59						
FI	8	119	4.56	4.33	-1.33						

Note: The table shows average values by country and demographic group across survey waves. Groups are split by: Male (1,2,3,4) or female (5,6,7,8); college education (3,4,7,8) or below (1,2,5,6); and income bracket above 60th percentile (2,4,6,8) or below (1,3,5,7).

Table A2: Alternative measures of energy prices

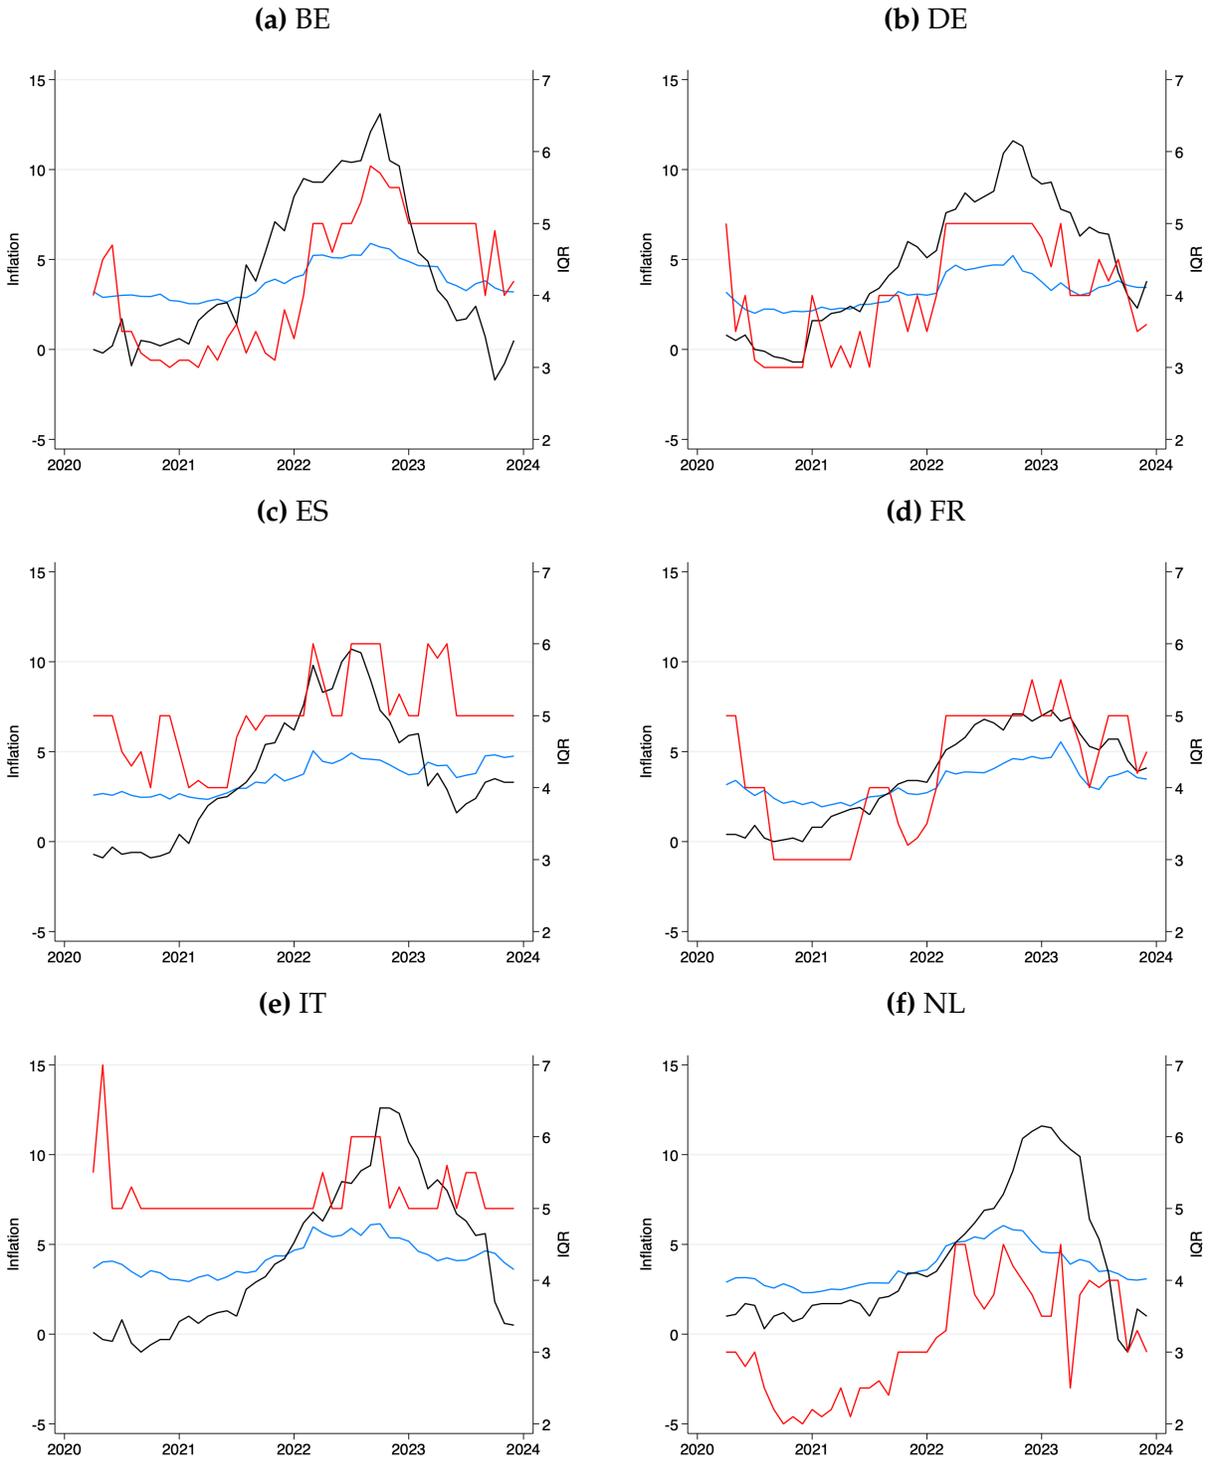
	HICP energy		Wholesale prices		HEPI index	
	(1)	(2)	(3)	(4)	(5)	(6)
Change in energy prices	4.412*** (0.529)	3.860*** (0.144)	0.264 (0.180)	0.450*** (0.044)	2.095*** (0.306)	1.799*** (0.086)
Change in energy prices × Unanchoring	1.288*** (0.228)	0.622*** (0.087)	0.206* (0.118)	0.095*** (0.024)	0.326* (0.173)	0.164*** (0.051)
Average past inflation	-0.025 (0.021)	-0.027*** (0.007)	-0.109** (0.043)	-0.098*** (0.010)	-0.071*** (0.023)	-0.073*** (0.008)
Observations	362756	2472	330729	2112	344597	2296
R ²	0.024	0.438	0.013	0.331	0.019	0.419

Note: This table re-estimates the first two columns of table 1 using alternative measures of energy prices. The first two columns replace the HICP electricity index, with the HICP energy index, the next two with the wholesale electricity price index, and the final two with the household energy price index.

Table A3: Alternative measures of anchoring and of average expectations

	Individual				Median		Wgt. mean
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Change in electricity prices	1.311*** (0.372)	1.101*** (0.255)	0.740*** (0.220)	1.120*** (0.308)	0.603*** (0.118)	0.575*** (0.118)	0.903*** (0.100)
Unanchoring (Disagreement)		0.362*** (0.038)					
Change in electricity prices × Unanchoring (Disagreement)		0.351*** (0.124)		0.724*** (0.194)	0.239*** (0.070)		0.473*** (0.066)
Unanchoring (Target distance)			0.977*** (0.083)				
Change in electricity prices × Unanchoring (Target distance)			1.080*** (0.374)			0.892*** (0.177)	
Average past inflation	-0.099*** (0.027)	-0.078*** (0.021)	-0.078*** (0.021)	-0.099*** (0.026)	-0.155*** (0.009)	-0.157*** (0.008)	-0.088*** (0.006)
Observations	362756	362756	362756	322987	2472	2472	2472
R ²	0.012	0.021	0.025	0.011	0.269	0.273	0.222

Figure A1: Inflation, mean expectations, and the IQR by major EA country



Note: Referring to the left axis, the black line shows actual inflation (HICP yoy, nsa) and the blue line shows mean expected inflation for the next 12 months. Against the right axis, the red line shows the IQR of expected inflation 3-years ahead. The sample is 2020:4 to 2023:12

Table A4: Baseline results by country

	AT	BE	DE	ES	FI	FR	GR	IE	IT	NL	PT
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Change in electricity prices	2.807*** (0.935)	2.384*** (0.432)	-0.028 (1.283)	1.942*** (0.419)	-0.979 (2.276)	-2.656 (5.389)	-7.076*** (2.195)	-6.901* (3.425)	1.678*** (0.549)	-0.714 (0.478)	2.966*** (0.978)
Change in electricity prices × Unanchoring	0.188 (0.364)	0.436* (0.220)	6.485*** (0.884)	0.401* (0.233)	0.161 (0.365)	12.657*** (1.183)	-0.016 (0.242)	0.302 (0.260)	0.477** (0.230)	-0.229 (0.248)	-0.282 (0.371)
Average past inflation	-0.021 (0.066)	-0.077** (0.029)	-0.064** (0.026)	0.016 (0.031)	-0.327 (0.338)	-0.044 (0.079)	0.359*** (0.109)	-0.872* (0.411)	-0.049 (0.046)	-0.335*** (0.041)	-0.073 (0.078)
Observations	9473	25618	68611	66437	10034	71166	4738	6614	66480	24675	8910
R ²	0.010	0.040	0.020	0.010	0.006	0.025	0.006	0.012	0.019	0.050	0.015

Table A5: Baseline results with different fixed effect combinations

	Individual			Mean		
	(1)	(2)	(3)	(4)	(5)	(6)
Average past inflation	0.004 (0.078)	-4.392*** (0.080)	0.002 (0.079)	-0.049** (0.023)	4.321 (4.676)	-0.049** (0.024)
Change in electricity prices	0.368** (0.182)	78.313*** (1.470)	0.386** (0.182)	0.192** (0.094)	-83.676 (96.295)	0.193** (0.096)
Change in electricity prices × Unanchoring	0.145 (0.087)	0.052 (0.072)	0.140 (0.093)	0.001 (0.052)	-0.026 (0.047)	-0.013 (0.054)
Observations	362756	362756	362756	2472	2472	2472
R ²	0.032	0.042	0.034	0.583	0.722	0.603
Country & group fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Country-group fixed effects	Yes	No	No	Yes	No	No
Country-month fixed effects	No	Yes	No	No	Yes	No
Group-month fixed effects	No	No	Yes	No	No	Yes

Table A6: Results for h-month changes in all variables

	1-month changes		4-month changes		12-month changes	
	(1)	(2)	(3)	(4)	(5)	(6)
Average past inflation	-0.013 (0.010)	-0.018*** (0.004)	-0.052** (0.021)	-0.064*** (0.006)	-0.303*** (0.040)	-0.311*** (0.014)
Change in electricity prices	0.963** (0.402)	0.605*** (0.174)	1.132*** (0.280)	0.943*** (0.112)	0.889*** (0.215)	0.560*** (0.117)
Change in electricity prices × Unanchoring	0.255 (0.454)	-0.168 (0.126)	0.581*** (0.171)	0.319*** (0.070)	0.875*** (0.127)	0.566*** (0.055)
Observations	518748	2912	414988	2648	237269	1944
R ²	0.001	0.022	0.006	0.174	0.041	0.568

Note: Columns (1), (3) and (5) show results for individual expectations, $\Delta^h \pi_{i,c,g,t}^e$, while columns (2), (4) and (6) show results for average expectations within country and group, $\Delta^h \pi_{c,g,t}^e$.

Table A7: Results with Huber-White and group-clustered standard errors

	Calendar clustering		Huber-White		Group clustering	
	(1)	(2)	(3)	(4)	(5)	(6)
Average past inflation	-0.097*** (0.026)	-0.104*** (0.024)	-0.097*** (0.004)	-0.104*** (0.004)	-0.097*** (0.014)	-0.104*** (0.014)
Change in electricity prices	1.163*** (0.305)	0.983*** (0.243)	1.163*** (0.062)	0.983*** (0.062)	1.163*** (0.124)	0.983*** (0.079)
Change in electricity prices × Unanchoring (Disagreement)	0.669*** (0.192)		0.669*** (0.046)		0.669*** (0.180)	
Change in electricity prices × Unanchoring (Target distance)		2.695*** (0.533)		2.695*** (0.117)		2.695*** (0.252)
Observations	362756	362756	362756	362756	362756	362756
R ²	0.013	0.015	0.013	0.015	0.013	0.015

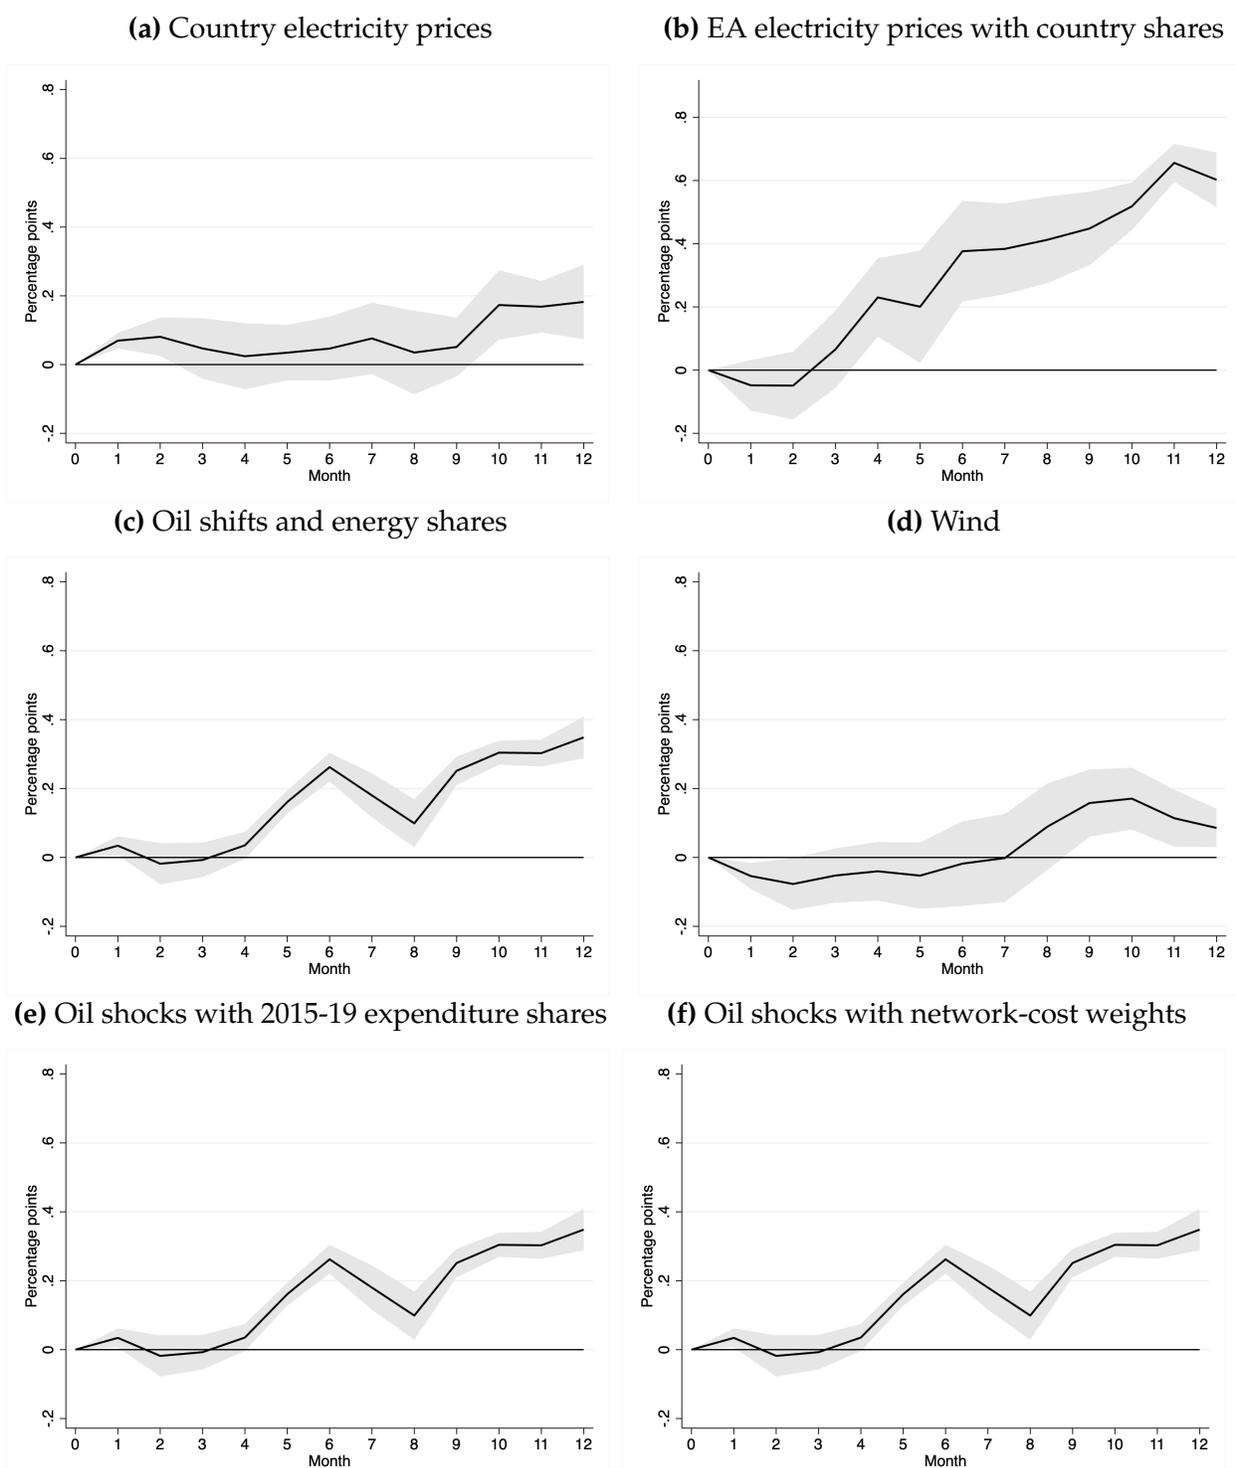
Note: This table re-estimates the first two columns of table 1 using different types of approaches to calculate the standard errors.

Table A8: The impact of electricity prices on expected inflation in the US

Revision of expectation	(1)	(2)	(3)	(4)
Change in gas prices	3.773*** (0.306)	3.657*** (0.276)	2.722*** (0.316)	3.919*** (0.060)
" × Unanchoring	-0.056 (0.149)	0.474*** (0.132)	0.557 (0.787)	0.075** (0.031)
Average past inflation	-0.128*** (0.023)	-0.121*** (0.023)	-0.149*** (0.025)	-0.100*** (0.008)
Observations	59231	59231	24607	2872
R ²	0.024	0.024	0.017	0.624

Note: This table presents estimates of the regression in equation (1) without country-level variation: $\Delta^6 \pi_{i,g,t}^e = \beta \Delta^6 e_t + \gamma \Delta^6 e_t \times \Delta^6 a_{g,t} + \theta \bar{\pi}_{t-6} + \eta_g + \varepsilon_{i,g,t}$. Column (1) has the baseline estimates using gasoline prices, (2) lags anchoring by 6 months to $a_{g,t-6}$, (3) uses as measure of unanchoring the distance of long-run expectations from the inflation target, and (4) uses the average $\pi_{g,t}^e$ as the dependent variable. In parentheses are standard errors clustered by month for the regressions using individual expectations.

Figure A2: Impulse response of expected inflation to a shock to energy prices



Note: Local projection of average expected inflation within a region and group on 3-month cumulated energy price shock, controlling for inflation, country and group fixed effects. The shocks are scaled by their standard deviation. Panels (a) to (d) show the standard errors corresponding to the pooled estimates in figure 1. Panels (e) and (f) investigate the robustness of the oil series by using alternative variables to measure the shares. Standard errors are clustered by country.

Table A9: Decomposition of US estimates

	(1)	(2)	(3)	(4)	(5)
Change in gas prices	3.773*** (0.306)		3.910*** (0.324)		
Change in oil prices		1.980*** (0.230)		1.872*** (0.226)	2.193*** (0.262)
Change in gas prices × Unanchoring	-0.056 (0.149)				
Change in oil prices × Unanchoring					-0.032 (0.293)
Average past inflation	-0.128*** (0.023)			-0.090*** (0.014)	
Observations	59231	89154	59231	89154	73095
R^2	0.024	0.011	0.022	0.013	0.015

Note: The table presents results for individual expectations, $\Delta^6 \pi_{i,t}^e$. Column (1) presents estimates of the regression in equation (1), repeating the first column of table A8. Column (2) presents estimates of the CG specification with wholesale oil prices for our updated sample: $\Delta^6 \pi_{i,t}^e = \beta \Delta^6 e_t + \alpha + \varepsilon_{i,t}$. Column (3) shows the CG estimates when switching from wholesale oil to retail gasoline prices, (4) when controlling for past inflation, $\bar{\pi}_t$, and (5) when including unanchoring at the national level, $\Delta^6 a_t$.