LUISS

Cattedra

RELATORE

CORRELATORE

CANDIDATO

Anno Accademico

Learning from Prices:
Input Costs and Inflation Expectations

Francesco Fabris [1]

Rome Masters in Economics
EIEF and LUISS

Abstract

We study the inflation expectation-formation process of firms. We provide empirical support to the channel of exposure to posted prices in order to explain the cross-sectional heterogeneity of inflation expectations. Using a randomized control trial (RCT) and survey-based micro-data on reported inflation expectations and input cost inflation, we show that firms rely differentially on their input costs inflation when forming inflation expectations, as a result of exogenous changes in their information sets. We characterize how this channel changes when transitioning from low- to high-inflation settings. Lastly, we show how informed firms use the new information provided by the treatment, vis-à-vis their individual input inflation, to update their expectations about future inflation.

[1] I am indebted to my supervisor, Luigi Paciello, for his guidance, and to Mounu Prem and all the EIEF faculty for their suggestions. I also thank Laura Pittalis and Francesco Boninsegna for their support. All errors are my own.

1 Introduction

Expectations are ubiquitous in economic agents' decisions that contain an inter-temporal dimension: contemporaneous choices depend on agents' perceptions about future economic outcomes. This prominent role of expectations has long been recognised by macroeconomists. Starting at least from the 1960s, economists have been continuously trying to incorporate expectations into their models. While different hypotheses have been formulated over the years, most modern business cycle models are built on the FIRE assumption: the idea that agents have full information about the economy and form their expectations rationally. Rational expectations were first theorised in Muth (1961) and later gained a central position in macroeconomics, as a result of the so-called rational expectations revolution of the 1970s (Lucas, 1972, 1973, 1976; Lucas and Sargent, 1979).

Among economic expectations, inflation expectations are the most closely monitored. The increasing availability of surveys of households and firms has proved the importance of inflation expectations in the determination of economic outcomes at the individual level. Inflation expectations enter firms' investment, pricing and employment plans (Coibion et al., 2019), and households' consumption, investment and labour market decisions (Bachmann et al., 2015; Coibion et al., 2023; Pilossoph and Ryngaert, 2023).

Beyond their influence at the individual level, inflation expectations are also central in the determination of aggregate outcomes. For example, the standard micro-founded New Keynesian models use NKPCs which incorporate inflation expectations in the determination of aggregate inflation realisations (Galí, 2008). Moreover, when the developed economies hit the effective lower bound, central bankers started managing inflation expectations as an instrument of monetary policy (Draghi, 2015; Yellen, 2016). Price stability also fundamentally depends on the anchoring of long-term inflation expectations. Given such centrality, describing correctly the process of expectations-formation and incorporating it into the models is a first-order issue for the conduct of policy.

FIRE predicts that all agents have the same expectations about the future. The evidence from micro-level survey-based data for households, firms and professional forecasters, however, deviates sensibly from the FIRE assumption: individual expectations are biased, volatile across time and differ across individuals (D'Acunto et al., 2023). Relaxing the FIRE assumption entails allowing for models of imperfect information. Imperfect information is not sufficient per se to explain the heterogeneity in expectations. When agents are uninformed about the aggregate state they receive noisy signals about the aggregate. They optimally decide to respond weakly to these signals. Yet they are very attentive about the price signals they receive at low cost, such as the prices they observe in their day-to-day activity. Whether agents use such private signals to infer about current inflation, given the high heterogeneity consumption bundles (or input mix for firms) and the heterogeneity in price developments across products, can be a relevant factor to explain the cross-sectional heterogeneity of agents' expectations.

One of the main drivers of the heterogeneity in inflation expectations for households is the personal experience of posted prices in day-to-day shopping: variations in the prices of frequently purchased goods significantly shape consumers' expectations regarding aggregate inflation. The weights consumers assign to price changes in their grocery bundle in their expectations-formation process depend on the frequency of purchase (D'Acunto et al., 2021). These findings buttress the assumptions to the class of Islands' models, first theorised by Lucas (1972, 1973), whereby agents learn about the aggregate state from the prices they observe in their local markets. While the literature on households' expectation formation has established the causality between inflation expectations and posted prices, there is still relatively little evidence on whether an analogous channel is operating through input prices for firms' expectations.

If input inflation were heterogeneous across firms and if firms used their input cost inflation to infer aggregate inflation, this would be a potential explanation for the disagreement in firms' inflation expectations documented in the literature. Figure 1 documents the heterogeneity in input cost inflation and inflation expectations using survey data from Italian firms. In particular, panels (a) and (b) plot the average input inflation and inflation expectations by sector of activity. Input inflation at the sectoral level was similar across sectors, even though more pronounced during the inflation spike for inputs to the industrial and construction sectors. Inflation expectations, on the other hand, are almost undistinguishable by sector. Panels (c) and (d) document the cross-sectional dispersion of input inflation and inflation expectations at the aggregate level and disaggregated by sector of activity. Construction and Industrial inputs, again, show the greatest dispersion during the inflation outburst but sectoral dispersions follow similar paths. The dispersion of inflation expectations, again, does not seem to be significantly different across sectors. First, these panels show that there exists sizeable input inflation heterogeneity at the firm level, which follows similar paths across sectors and is positively correlated with the level of inflation realisations. Second, input inflation dispersion positively correlates with the dispersion of inflation expectations at the individual level.

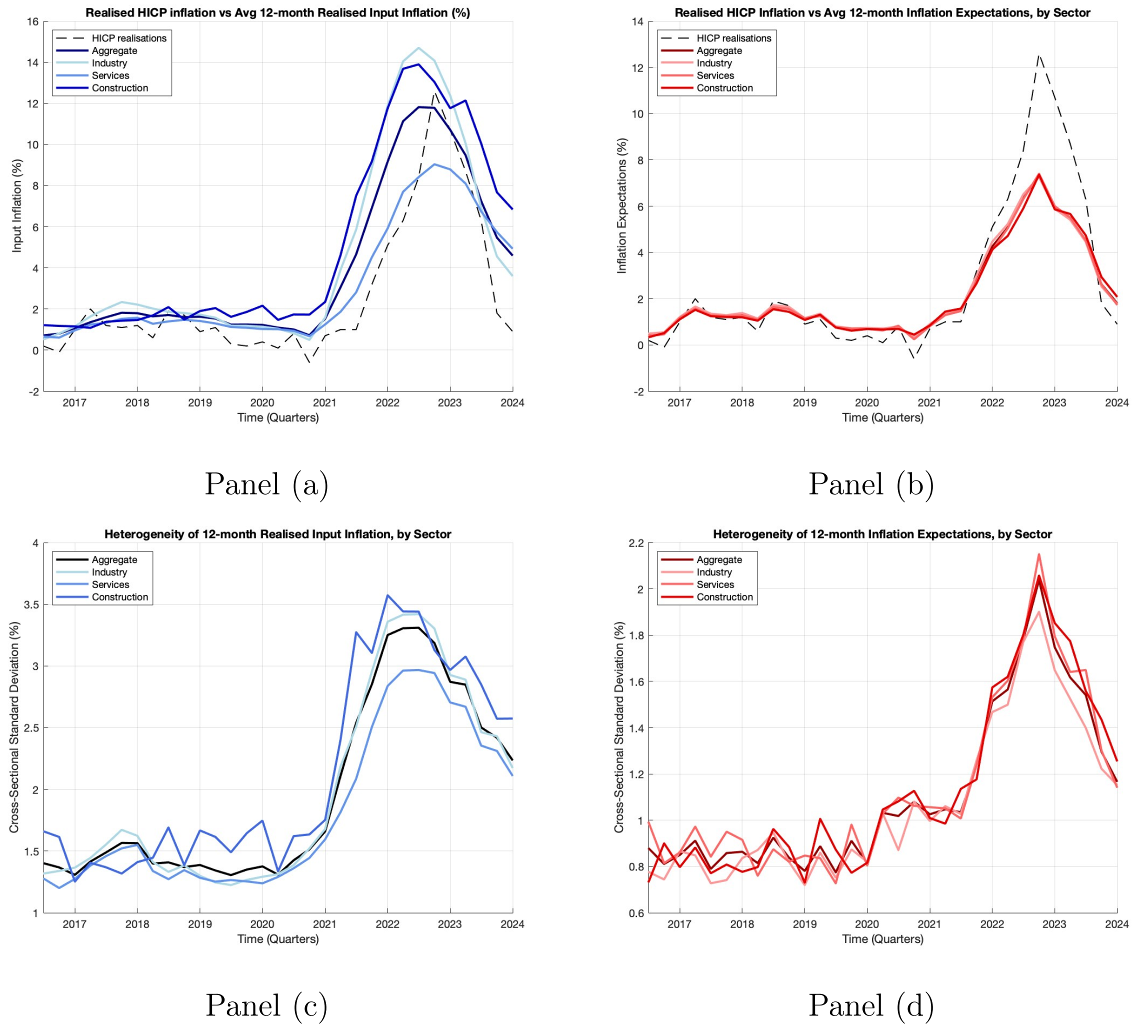

Figure 1: Heterogeneity in Input Costs Inflation and Inflation Expectations. SIGE data

Using the firm-level data of expectations and reported input cost inflation from the Survey on Inflation and Growth Expectations (SIGE), this work aims to test the assumption underlying Islands models with Italian data. First, we exploit the design of SIGE, which features a randomized informational treatment, to estimate the effect of a change in the information set on the reliance on input costs of firms. Our results support the hypothesis that firms learn from prices: being provided with additional information about aggregate inflation, firms levy less on the signal from input inflation. Second, because of the variation in aggregate inflation realization within the sample period, we characterise to which extent learning from price changes across inflation regimes. We show that, in line with the existing literature whereby agents are more attentive to aggregate conditions when inflation is high (Weber et al., 2023), treatment effects are attenuated in periods of high inflation. The coefficient estimating a differential elasticity to input prices when transitioning to different inflation regimes is, however, not significant. Third, we document how treated firms use the new information provided by the treatment vis-à-vis their private prior, constituted by input inflation, to update their expectations about future inflation. Informed firms perceive low inflation on their input side to be predictive of lower aggregate future inflation.

First, this work relates to the relatively unpopulated empirical literature of learning from prices for firms. Andrade et al. (2022) uses a survey of French manufacturing firms to document how firms' expectations by both aggregate and industry-specific conditions. In response to industry-level shocks orthogonal to the aggregate state, firms' aggregate expectations respond persistently. In the same spirit, Albagli et al. (2022) finds that Chilean firms extrapolate an aggregate value for future inflation from a local signal obtained from their business purchasing prices orthogonal to current inflation.

Second, the thesis connects to previous work exploiting SIGE data. The most closely related papers are Coibion et al. (2019), henceforth CGR, and Ropele et al. (2023), which use similar empirical strategies to provide causal evidence between inflation expectations heterogeneity and firm-level economic outcomes and aggregate TFP losses, respectively. Other papers used SIGE to show the response of inflation expectations to monetary policy shocks (Bottone and Rosolia, 2019), and the correlation between inflation expectations and expected own-price changes (Ropele, 2019).

This work is organized as follows. Section 2 provides the basic conceptual framework of Islands' models. Section 3 describes the data used from the SIGE. Section 4 provides the empirical strategy and the estimation for average treatment effects. Section 5 estimates how treated firms use their private signal and the aggregate signal to form their expectations. Section 6 concludes.

2 Conceptual Framework

This section is devoted to describing the conceptual framework typical to Islands' models, with particular reference to Lucas (1973). These models sought to explain the short-term pro-cyclicality of output and inflation while maintaining the assumption of long-run money neutrality. The main ingredients of these models are competitive markets, imperfect information and rational expectations.

In this context, firms have to form expectations about the current aggregate price level as a function of the new information from the private signal they receive from their current input prices and their expectation about current-period inflation formulated one period before. They solve a signal extraction problem which optimally gives weight to both the signal from input shocks and to the expectations of the current aggregate price level conditional on the previous period information set.

The economy is constituted by N islands, indexed with $j = 1, ...N$. In each island j, there exists a continuum of identical final-good producers. The supply curve for producers in island j is an increasing function of the relative price of good j, $p_t(j)$, with respect to the economy-wide price level p_t:

$$y_t(j) = \gamma(p_t(j) - p_t)$$

Because producers do not have full information about input prices in the whole economy, each final producer has to form expectations about the current price level of the economy. Let $I_t(j)$ be the information set available to a producer in market j at period t. By assumption, $I_t(j)$ comprises the current local input price $p_t(j)$ as previous period price level p_{t-1}. The incomplete information supply curve is then:

$$y_t(j) = \gamma(p_t(j) - \mathbb{E}[p_t|I_t(j)])$$

Rational expectations are used to characterise the imperfect information supply curve. In particular, it must hold that $p_t = \mathbb{E}(p_t|I_{t-1}) + \varepsilon_t$, where the forecast error ε_t cannot be predicted at time $t-1$, with $\mathbb{E}(\varepsilon) = 0$ and $\mathbb{E}(\varepsilon^2) = \sigma^2$. The term ε is the aggregate shock, common to all islands.

Furthermore, input prices are set by local input producers in each market j. Input producers are price-setters and are subject to both the aggregate shock ε and an island-specific shock η_{jt} (e.g. local productivity shock). In the simplest setting, one can assume random exogenous shocks across islands $p_t(j) = p_t + \eta_{jt}$: zero-mean $\mathbb{E}(\eta_{jt}) = 0$ for all j, with common variance $\mathbb{E}(\eta_{jt}) = \tau^2$ and with uncorrelated shocks at all leads and lags $\mathbb{E}(\eta_{jt}, \eta_{i,t-k}) = 0$ for all i,j,k. Because of the information rigidity present in the economy, input prices carry a composite signal about the shocks occurring in the economy at period t, $\eta_{jt} + \varepsilon_t$, whose individual factors cannot be isolated by final producers.

In deciding how much output to produce, final producers need to solve a signal extraction problem. They need to find the optimal weights to assign to island-specific and aggregate shocks when observing the composite signal $\eta_j + \varepsilon$. By assumptions of the model, uncertainty concerns only one period: at time t, firms know the realisations of $\eta_{j,t-1}$ and ε_{t-1} and therefore the aggregate price level p_{t-1}. To estimate which fraction of the composite signal is due to island-specific shocks, the solution to the signal extraction problem yields:

$$\eta_{jt} = \frac{\tau^2}{\tau^2 + \sigma^2} \left(\eta_{jt} + \varepsilon_t\right) + u_t$$

With $\frac{\tau^2}{\tau^2+\sigma^2} \equiv \theta$. By assumption of random exogenous shocks, we can write the expected price level conditional on the information set of the firm in island j:

$$\begin{aligned} \mathbb{E}(p_t|I_{t-1}(j), p_j(t)) &= p_t(j) - \mathbb{E}(\eta_{jt}|I_{t-1}(j), p_t(j)) \\ &= p_t(j) - \theta(p_t(j) - \mathbb{E}(p_t|I_{t-1})) \\ &= (1-\theta)p_t(j) + \theta\,\mathbb{E}(p_t|I_{t-1}) \end{aligned}$$

The coefficient $1 - \theta$ describes "learning from prices", ie. the extent to which firms optimally rely on their input prices to make inferences about the current aggregate price level. It is a function of the relative volatility of idiosyncratic shocks to aggregate shocks. The higher the volatility of idiosyncratic shocks η_j, the less informative input prices are to determine the aggregate price level, and the lower the optimal weight attached to them. Such a coefficient is a result of the information rigidity present in the economy. If firms could observe aggregate shocks, or equivalently if they could observe all input prices and all production functions across islands, there would be no scope for a signal extraction problem and no weight on input costs would be attached.

This framework can be adopted when thinking about firms' expectations about future inflation realisation. When facing the problem of predicting future period inflation, firms would need to use their input cost inflation to predict the current inflation level. Furthermore, they could partly levy on their input inflation as predictive of future inflation, via production networks and input-output linkages. Firms would then use their input to both predict current and future inflation realisations at the aggregate level. Therefore, this class of models predicts that changing the degree of informedness of firms about the current aggregate state should change the extent of firms learning from their input prices. Firms which know with certainty the present inflation level should use information coming from their inputs insofar as they are informative about future inflation. The following sections will test this prediction and provide evidence in favour of imperfect information and "learning from prices".

3 Data

The data source is the Survey of Inflation and Growth Expectations (SIGE), a quarterly business survey run since December 1999 by the Bank of Italy. The reference universe consists of firms operating in the industrial sector (ATECO codes B-E), constructions and non-financial private services, with administrative headquarters in Italy and employing 50 or more workers. The sample is stratified by sector of economic activity (industry, non-financial private services, and construction), geographical area (northwest, north-east, centre, and south and islands), and number of employees (50–199, 200–999, 1,000 and over). In recent years, each wave has about 1,000 firms (400 in industry excluding construction, 400 in nonfinancial private services, and 200 in construction). The list of firms used to construct the representative sample is drawn from the INPS and InfoCamere databases and is updated every five years on average.

The purpose of the survey is to obtain information on firms' expectations concerning inflation, the general economy, own-product prices and demand, investment, and employment. While most of the data is qualitative, firms are required to give quantitative estimates about own-product price changes (past and expected), own-input price changes, inflation expectations, and current number of employees. The survey is conducted by a specialist firm that distributes the questionnaire to company managers, who are best informed about the topics covered in the survey. About 90 per cent of the data is collected through computer-assisted web interviews in the form of an online questionnaire featuring a purpose-designed interface, while the remaining 10 per cent are collected through computer-assisted telephone interviews. Data are collected in the first three weeks of March, June, September and December. The response rate is on average 45 per cent.

Starting 2012Q2, SIGE features a randomised informational treatment, whereby two-thirds of the surveyed firms are being told the latest Italian inflation reading before declaring their inflation expectations, whereas the other third of the sample does not receive any informational treatment. Assignment into groups is random and lasts on average every 5 years. New assignment into treatment and control is contemporaneous with a new resampling of SIGE. The list of firms in the representative sample is updated from the population of firms present in the INPS and Infocamere databases. As noted in CGR, a non-negligible share of treated firms ends up in the control group and vice versa: during the reassignment that occurred in 2019Q2 nearly 60% of firms from the control group moved into the original treatment group, and nearly 20% of firms in the treatment group moved to the control group.

The question about inflation expectations comes at the beginning of the questionnaire, before eliciting a quantitative measure of inflation expectations by firms. Firms in the treatment group are subject to the following informational treatment when asked about their expectations: "In [previous month], consumer price inflation measured by the 12-month change in the Harmonized Index of Consumer Prices was [X.X]% in Italy and [Y.Y]% in the Euro area. What do you think it will be in Italy ..." Firms in the control group are asked: "What do you think consumer price inflation in Italy, measured by the 12-month change in the Harmonized Index of Consumer Prices, will be ...". Individual inflation expectations are elicited also at a 6 month, 2 years and 4 years horizon. All other questions in the survey are identical.

Since 2016Q3, all firms have also been asked to report the upstream inflation rate they experienced on their inputs: they are asked to report what is "the average per cent change in your firm's prices of goods and services bought in Italy and abroad in the last 12 months". We will therefore use all waves in SIGE in the period 2016Q3-2023Q4.

4 Empirical Strategy

4.1 Baseline Specification

The randomised informational treatment featuring the SIGE sets up the ideal conditions for a Randomised Control Trial (RCT). With random assignment[2], treated and control units are identical on average, conditional on treatment status. Selection bias is resolved and the estimated coefficients have a causal interpretation.

The first hypothesis we want to test is whether firms use the signals present in their input cost inflation to form expectations about current inflation realisations. While CGR uses SIGE to show that firms are not fully informed, we want to test whether informed firms rely differentially on signals coming from input cost inflation to predict aggregate inflation. Within the Islands' framework, if all firms received an undistorted signal about current

[2] CGR provides empirical evidence of random assignment by regressing treatment status on observables and verifying that none of the regressors predict assignment into neither treatment or control groups.

aggregate inflation in the economy, their reliance on input cost inflation to predict future inflation should not change on whether the firm received the informational treatment of the latest inflation reading. The baseline specification then reads:

$$\begin{aligned}\mathbb{E}_{it}[\pi_{t+k}] =& \alpha_1 Treat_{it} + \alpha_2 Treat_{it} \times S_t + \beta_1 \Delta c_{it} + \beta_2 \Delta c_{it} \times Treat_{it} \\ & + \gamma_i + \gamma_{SEC \times t} + \gamma_{AREA \times t} + \epsilon_{it}\end{aligned} \tag{1}$$

Where $\mathbb{E}_{i,t}[\pi_{t+k}]$ is the inflation expectation at time t of unit i for k quarters ahead. We will consider inflation expectations at a 1-year horizon in the rest of the analysis. S_t is the realisation of HICP inflation in Italy at time t, Δc_{it} is the average per cent change in input prices in the 12 months before period t faced by firm i as reported in the survey; $Treat_{it}$ is an indicator variable that takes value 1 if the unit was treated at period t, 0 otherwise. The regression contains individual fixed effects γ_i, sectoral-time and area-time fixed effects, $\gamma_{SEC \times i}$ and $\gamma_{AREA \times t}$, to net out potential confounding factors not attributable to the treatment itself. Following CGR, we allow for serial cross-sectional and time correlation in the errors using Driscoll and Kraay (1998) standard errors.

We claim that all coefficients are properly identified and have a causal interpretation. Contrary to CGR, because the composition of treatment and control groups changes during our sample period and a non-negligible share of both groups transitions into the other in the subsequent periods, we can properly identify the treatment effect associated with mere inclusion or not into the treatment. The specification allows us to estimate a direct treatment effect affine in the level of aggregate inflation realisations. The coefficient α_1 estimates the average difference in inflation expectations between treated and controls at a zero-inflation level. The coefficient α_2 estimates how this difference changes linearly as a function of current inflation realisations. The coefficient of interest is β_2, the indirect treatment effect mediated by input costs. It identifies the differential elasticity of inflation expectations to input costs for treated firms vis-à-vis control firms. Conditional on input inflation Δc and aggregate inflation realisation S, the Average Treatment Effect (ATE) on inflation expectations is:

$$\mathbb{E}[\pi|\Delta c, S, T] - \mathbb{E}[\pi|\Delta c, S, NT] = \alpha_1 + \alpha_2 S + \beta_1 \Delta c$$

Table (1) reports the sequential estimation leading to the results of column 9, corresponding to the preferred specification (1). Column 2 directly speaks to CGR as to their baseline specification. The qualitative results point in the same direction: a positive coefficient α_2 means that treated firms have differential inflation expectations with respect to controls and such are increasing less than one-to-one in the current level of inflation. While not precisely identified, the sign of coefficient α_1 is suggestive evidence that treated firms have relatively lower inflation expectations at low levels of inflation realisations. Both facts still hold when correctly identifying α_1 by including individual fixed effects and using sector-time and area-time fixed effects, instead of sector-quarter fixed effects alone in column 8. The coefficient associated with input inflation β_1 is significant and positive across all specifications. The coefficient of interest β_2 turns statistically insignificant when including fixed effects at the individual level. Across all specifications, however, the sign of the estimated coefficient is consistently negative. This might suggest that indeed the prediction from Islands' models is verified.

We restrict the estimation of specification (1) into time subsamples. In particular, we split the sample into three sub-periods: 2016Q3-2019Q4, which was a period of relative calm for the Italian economy; 2020Q1-2021Q4, which was the period affected by Pandemic-related restrictions and 2022Q1-2023Q4, characterised by the inflation outburst and the subsequent descent. Table (2) reports in each column the estimation results depending on the reference period. The sign of the estimated coefficients for α_1 and α_2 does not change across columns, being respectively negative and positive in all samples in which they are identified. We note that α_1 is not estimated in sample 2022Q1-2023Q4 because assignment into treatment was not redrawn in the period. This results in the treatment indicator $Treat_{it}$ and individual fixed effects being collinear. As for the main coefficient of interest β_2, the estimation from the three subsamples shows that the non-significant coefficient estimated across the whole sample period is a composition of heterogeneous effects in the subsamples. During 2016Q3-2019Q4 treated firms indeed have a lower elasticity of inflation expectations to input inflation. This coefficient turns insignificant but negative during the Pandemic period and insignificant but positive during the inflation outburst sample.

To formally test the prediction implied by the framework, we can carry out a one-sided hypothesis test. With the research hypothesis being whether treated firms rely relatively less on their input costs when predicting inflation, this amounts to an F-test $H_0 : \beta_2 \geq 0$ versus $H_1 : \beta_1 < 0$. Table 3 shows that one can reject the null of weakly positive differential elasticity to input costs between 2016Q3-2019Q4 and 2020Q1-2021Q4 at, respectively, 1% and 10% significance levels. The significance level for the whole sample period is slightly above the 10% level. These results suggest, at different significance levels for the two subsamples before 2022, that firms used their input inflation as signals to infer aggregate inflation when forming their expectations about the future.

Furthermore, one can test whether the signal of the current inflation realisation is a sufficient statistic for informed firms to predict future inflation realisations. Again, we can use a linear hypothesis test, with the research hypothesis being that the elasticity of inflation expectations to input costs for fully informed firms is non-negative. For all subsamples considered, the results in Table (4) show that treated firms maintain a positive elasticity of inflation expectations to input costs. One way to defend this result within Rational Expectations is that fully informed firms keep using the component from their input cost inflation that is predictive of future realisations of aggregate inflation to inform their expectations.

Table 1: BASELINE REGRESSION

Regression 1 Dependent Variable: WHOLE SAMPLE	(1) $E\pi_{t+12}$	(2) $E\pi_{t+12}$	(3) $E\pi_{t+12}$	(4) $E\pi_{t+12}$	(5) $E\pi_{t+12}$	(6) $E\pi_{t+12}$	(7) $E\pi_{t+12}$	(8) $E\pi_{t+12}$
$Treat_{it}$ (α_1)	.2622	-0.3719***	-0.2034***	-0.1713***	-0.1106***	-0.1074	-0.1094	-0.1596***
	(0.212)	(0.054)	(0.0592)	(0.0565)	(0.074)	(0.0667)	(0.066)	(0.0525)
S_t		0.4240***	0.3934***	0.3739***	0.2240***			
		(0.0281)	(0.0241)	(0.0256)	(0.023)			
$Treat_{it} \times S_t$ (α_2)		0.2184***	0.1986***	0.2243***	0.1371***	0.1331***	0.1319***	0.1799***
		(0.0172)	(0.0158)	(0.0128)	(0.020)	(0.0192)	(0.0189)	(0.0178)
Δc_{it} (β_1)			0.0410***	0.0616***	0.0401***	0.0285***	0.027***	0.0258***
			(0.0023)	(0.0092)	(0.0064)	(0.0075)	(0.0075)	(0.0093)
$Treat_{it} \times \Delta c_{it}$ (β_2)				-0.0265**	-0.0151*	-0.0165**	-0.0145*	-0.0119
				(0.0102)	(0.0082)	(0.007)	(0.0079)	(0.0091)
L.$E\pi_{t+12}$					0.3982***	0.464***	0.4625***	0.2615***
					(0.0306)	(0.0290)	(0.0293)	(0.0327)
Sector × Quarter FE	Yes	Yes	Yes	Yes	Yes			
Sector × Time FE						Yes	Yes	Yes
Area × Time FE							Yes	Yes
Firm FE								Yes
Constant	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,737	32,737	32,737	32,737	24,916	24,916	24,916	24,541

Source: Bank of Italy, Survey on inflation and growth expectations, 2016Q3-2023Q4
Driscoll-Kray Standard Errors in Parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 2: BASELINE REGRESSION SUBSAMPLES

Regression 1	(1)	(2)	(3)	(4)
Dependent Variable:	$E\pi_{t+12}$	$E\pi_{t+12}$	$E\pi_{t+12}$	$E\pi_{t+12}$
Sample:	WHOLE SAMPLE	2016Q3-2019Q4	2020Q1-2021Q4	2022Q1-2023Q4
$Treat_{it}$ (α_1)	-0.1596***	-0.2037***	-2.049***	
	(0.0458)	(0.0445)	(0.0935)	
$Treat_{it} \times S_t$ (α_2)	0.1799***	0.382***	0.413***	0.2603***
	(0.022)	(0.038)	(0.0429)	(0.0265)
Δc_{it} (β_1)	0.0258***	0.998***	0.0356***	-0.0044
	(0.010)	(0.029)	(0.004)	(0.0280)
$Treat_{it} \times \Delta c_{it}$ (β_2)	-0.0119	-0.0902***	-0.0157	0.0208
	(0.011)	(0.0265)	(0.0104)	(0.0250)
L.$E\pi_{t+12}$	0.2615***	0.0474	0.267***	0.0087
	(0.0225)	(0.0385)	(0.0796)	(0.0228)
Constant	1.413***	0.939***	2.265***	2.801***
Sector × Time FE	Yes	Yes	Yes	Yes
Area × Time FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Observations	24,916	9,654	6,149	7,062

Source: Bank of Italy, Survey on inflation and growth expectations, 2016Q3-2023Q4
Driscoll-Kraay Standard Errors in parentheses
Significance Levels: *** p<0.01, ** p<0.05, * p<0.1

Table 3: Test $H_0 : \beta_2 \geq 0$ vs. $H_1 < 0$

SAMPLE PERIOD	Point Estimate	**p-value**
WHOLE SAMPLE	-0.0119	0.105
2016Q3-2019Q4	-0.0902	0.0025
2020Q1-2021Q4	-0.0158	0.090
2022Q1-2023Q4	0.0207	0.781

Table 4: Test $H_0 : \beta_1 + \beta_2 \leq 0$ vs. $H_1 > 0$

SAMPLE PERIOD	Point Estimate	**p-value**
WHOLE SAMPLE	0.0139	0.0000
2016Q3-2019Q4	0.0123	0.0132
2020Q1-2021Q4	0.0198	0.0018
2022Q1-2023Q4	0.0164	0.0025

4.2 Differential Treatment Effects

The results from the baseline specification point to heterogeneous treatment effects across periods. Given the heterogeneity in HICP inflation realisations across the whole sample, this section is devoted to testing whether the extent of learning from prices changes when transitioning from periods of low to high inflation, and vice-versa. To answer this question, the baseline regression is complemented with an indicator variable $High_t$ which takes value 1 when HICP inflation realisations in Italy were above the 2% threshold set by the ECB's mandate. The regression model allowing for differential treatment effects depending on the inflation regime then reads:

$$\begin{aligned}
\mathbb{E}_{it}[\pi_{t+k}] = \; & \alpha_1 Treat_{it} + \alpha_2 \, Treat_{it} \times High_t + \alpha_3 \, Treat_{it} \times S_t + \alpha_4 \, Treat_{it} \times S_t \times High_t \\
& + \beta_1 \, \Delta c_{it} + \beta_2 \, \Delta c_{it} \times Treat_{it} \\
& + \beta_3 \, \Delta c_{it} \times High_t + \beta_4 \, \Delta c_{it} \times Treat_{it} \times High_t \\
& + \gamma_i + \gamma_{SEC \times t} + \gamma_{AREA \times t} + \epsilon_{i,t,k}
\end{aligned} \tag{2}$$

This specification allows us to estimate the average treatment effects in low- and high-inflation periods separately. Let the former be indexed with L, the latter with H. The average treatment effects, conditional on input cost inflation Δc and aggregate inflation realisation S, will be:

$$\mathbb{E}[\pi|\Delta c, S, T, L] - \mathbb{E}[\pi|\Delta c, S, NT, L] = \alpha_1 + \alpha_3 S + \beta_2 \Delta c$$
$$\mathbb{E}[\pi|\Delta c, S, T, H] - \mathbb{E}[\pi|\Delta c, S, NT, H] = (\alpha_1 + \alpha_2) + (\alpha_3 + \alpha_4) \, S + (\beta_2 + \beta_4) \Delta c$$

The differential treatment effect between low- and high-inflation, conditional on S and

Δc, is estimated by the difference-in-differences estimator (DiD):

$$\text{DiD} = \alpha_2 + \alpha_4 S + \beta_4 \Delta c$$

Table 5 reports the estimation results. The coefficient α_1 shows that, on average, treated firms at zero inflation have lower inflation expectations than their controls. The coefficient α_3 shows that an additional 1 percentage point increase in realised inflation increases inflation expectations differentially for treated firms. The slope of the average treatment effect on aggregate inflation realisations S is higher for firms in low inflation settings as α_4 is negative: treated firms react more strongly to a 1 percentage point increase in HICP with respect to their controls inflation during low inflation periods. In other words, the treatment effect of an additional 1 percentage point inflation is lower in periods of high inflation. These estimation results square with the literature of learning in low- and high-inflation settings: as the economic environment switches to a high-inflation regime, households and firms become more attentive and informed about inflation, leading to respond less to exogenous information about inflation and monetary policy (Weber et al., 2023).

Conditional on firms being more informed about the current state, our framework would predict that both groups would be setting a lower weight on input costs when transitioning to high inflation periods as they have a higher level of informedness about the aggregate. This would amount to estimating β_3 and $\beta_3 + \beta_4$ as negative coefficients. This prediction is not verified in the data as none of the coefficients is significant and they have opposite signs. One possible explanation would be that firms understood the inflationary shock as coming from their inputs and therefore acknowledged a greater predictive power of inputs during high inflation[3]. This greater reliance on input costs would cancel the effect of firms having a better signal about aggregate inflation, therefore yielding to null effects.

The non-significance of β_4, in particular, implies that treated firms do not have a differential elasticity to input costs versus their controls when transitioning from low- to high-inflation periods. The negative point estimate, however, goes in the same direction of attenuated treatment effects during high inflation. As for low inflation periods, the estimate for β_2 is significant and negative, in line estimation of specification 1.

In the same fashion as in the previous section, we can test the hypotheses formulated by our conceptual framework with a series of linear hypotheses tests. Table 6 reports the

[3]Evidence of the prominence of supply shocks in inflation realisations is found in ECB (2022)

Table 5: DIFFERENTIAL TREATMENT EFFECTS

Regression 2 Dependent Variable:	(1) $E\pi_{t+12}$
$Treat_{it}$ (α_1)	-0.187***
	(0.0431)
$Treat_{it} \times High_t$ (α_2)	0.6988**
	(0.1973)
$Treat_{it} \times S_t$ (α_3)	0.232**
	(0.035)
$Treat_{it} \times High_t \times S_t$ (α_4)	-0.1237***
	(0.0442)
Δc_{it} (β_1)	0.0372***
	(0.0161)
$Treat_{it} \times \Delta c_{it}$ (β_2)	-0.0311**
	(0.015)
$High_t \times \Delta c_{it}$ (β_3)	-0.0100
	(0.0146)
$Treat_{it} \times \Delta c_{it} \times High_t$ (β_4)	0.0179
	(0.0164)
L.$E\pi_{t+12}$	Yes
Sector × Time FE	Yes
Area × Time FE	Yes
Firm FE	Yes
Observations	24,541

Source: Bank of Italy, SIGE, 2016Q3-2023Q4
Driscoll-Kraay Standard Errors in parentheses
Significance Levels: *** p<0.01, ** p<0.05, * p<0.1

results from such tests. First, average treatment effects on elasticities are negative and significant at least at 10% level during both regimes. Thus, even in the context of high inflation, informed firms levy relatively less on their input inflation than their controls. Second, we cannot reject the null that treated firms have the same treatment effects for elasticity to input costs as compared to the controls, in both inflation regimes. The point estimates are such that during high inflation periods, treated units have relatively higher elasticity to input costs versus their controls than with respect to low inflation periods, but the p-value for the one-sided test prevents us from conclusive statements in such regards. Again, this is suggestive of agents being more attentive during high inflation settings and having lower informational treatment effects.

Table 6: TREATMENT EFFECTS BY SETTING

H_0 (vs. H_1)	Point Estimate	**p-value**
$\beta_2 \geq 0 \quad$ (vs. < 0)	-0.0311	0.027
$\beta_2 + \beta_4 \geq 0$ (vs. < 0)	-0.0130	0.0897
$\beta_4 \leq 0$ (vs. > 0)	0.0179	0.146

Table 7: INPUT ELASTICITIES FOR THE TREATED

H_0 (vs. H_1)	Point Estimate	**p-value**
$\beta_1 + \beta_2 \leq 0$ (vs. > 0)	0.0064	0.1014
$\sum_i^4 \beta_i \leq 0$ (vs. > 0)	0.0141	0.000

Lastly, the results in table (7) indicate whether treated firms during both inflation regimes use their inputs to predict future inflation realisations: the signal about current HICP inflation is not a sufficient statistic to form inflation expectations during high inflation periods. The same test during low inflation periods is suggestive of the same pattern, even though the p-value for the test is slightly above 10%. Again assuming the FIRE hypothesis, this evidence points towards individual input inflation being predictive of future aggregate inflation mostly during the inflation outburst between 2021Q3 and 2023Q3.

4.3 Persistence of the Treatment

Analogously to CGR, one can also use the SIGE data to characterise the persistence of the informational treatment on inflation expectations. In particular, we can estimate whether exposition to the treatment leads treated firms to rely differentially on input inflation at leads further ahead in time than the contemporaneous one. To evaluate the persistence of the informational treatment when forming expectations at a horizon h while sitting j periods after the treatment. This specification builds on the baseline model 1 introducing lagged terms for the independent variables and keeping the same fixed effects.

$$
\begin{aligned}
\ln \mathbb{E}_{it}[\pi_{t+h}] = &+ \sum_{j=0}^{q} \gamma_{h,j} \Delta c_{i,t-j} \\
&+ \sum_{j=0}^{q} \delta_j \; Treat_{i,t-j} + \sum_{j=0}^{q} \theta_j \; Treat_{i,t-j} \times S_{t-j} \\
&+ \sum_{j=0}^{q} \beta_j \; \Delta c_{i,t-j} \times Treat_{i,t-j} \\
&+ \gamma_i + \gamma_{SEC \times t} + \gamma_{AREA \times t} + \epsilon_{itk}
\end{aligned} \tag{3}
$$

The average treatment effect of being treated j periods before, conditional on Δc_{t-j} and S_{t-j} is: $ATE(S_{t-j}, \Delta c_{t-j}) = \delta_j + \theta_j S_{t-j} + \beta_j \Delta c_{t-j}$. The parameter of interest is $\beta_{h,j}$, which estimates the average differential elasticity of inflation expectations to changes in input inflation that occur at period $t - j$ of treated firms with respect to control firms when forming expectations at period t.

We report the estimation from this specification for subsample 2016Q3-2019Q4, which is the only subsample that reported significant treatment effects for the elasticity term from the estimation of the baseline specification 1. We report the estimation of regression model 3 in table 8, allowing the maximum lag q to run from 0 (ie. the contemporaneous effect only) to four quarters before the treatment. The number of observations at each column shrinks because we are restricting the estimation to the subgroups of firms present in the survey for at least $q + 1$ consecutive periods. Consistently with the persistence of treatment effects found in CGR for the direct treatment effect, the estimation shows no persistence in average treatment effect as far as the elasticity to input costs is concerned.

5 Within the Treated

The previous subsections aimed at estimating the average treatment effects resulting from informing a share of firms about contemporaneous aggregate inflation. We have shown that treated firms continue to use their part of the signal coming from input cost inflation. We now restrict the sample to the panel comprising treated firms. We aim at characterising how treated firms use the new information provided by the treatment vis-à-vis their private prior, constituted by input inflation, to update their expectations of future inflation.

Table 8: PERSISTENCE: 2016Q3-2019Q4

Regression 3 Dependent Variable:	(1) $E\pi_{t+12}$	(2) $E\pi_{t+12}$	(3) $E\pi_{t+12}$	(4) $E\pi_{t+12}$
$\Delta c_{it} \times Treat_{it}$ (β_0)	-0.069***	-0.0982**	-0.121***	-0.106*
	(0.021)	(0.0153)	(0.021)	(0.0259)
$L.\Delta c_{it} \times L.Treat_{it}$ (β_1)		-0.0061	-0.035	-0.0262
		(0.023)	(0.029)	(0.030)
$L2.\Delta c_{it} \times L2.Treat_{it}$ (β_2)			0.037	0.044
			(0.028)	(0.043)
$L3.\Delta c_{it} \times L3.Treat_{it}$ (β_3)				-0.0145
				(0.036)
Firm FE	Yes	Yes	Yes	Yes
Area x Time FE	Yes	Yes	Yes	Yes
Sector x Time FE	Yes	Yes	Yes	Yes
Observations	24,541	8,876	6,726	5,286

Source: Bank of Italy, SIGE, 2016Q3-2023Q4
Driscoll-Kraay Standard Errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Let us define I_{it} the informational shock firm i receives at period t as the individual-specific time-varying variable that describes the difference between inflation in input costs and realised inflation in the same period: $I_{it} = S_t - \Delta c_{it}$. Therefore, when input inflation is lower than aggregate inflation, the firm will be subject to an informational shock with a positive sign. We are interested in estimating how the sign of the informational shock I_{it} affects an average firm i's inflation expectations. The first specification within the treated group is therefore:

$$\mathbb{E}_{it}[\pi_{t+k}] = \beta_1 POS_{it} + \gamma_i + \gamma_{SEC \times t} + \gamma_{AREA \times t} + \epsilon_{it} \tag{I}$$

Where the variable POS_{it} is an indicator for the sign of the informational shock:

$$POS_{it} = \begin{cases} 1 & \text{if } S_t - \Delta c_{it} \geq 0 \\ 0 & \text{if } S_t - \Delta c_{it} < 0 \end{cases}$$

The coefficient of interest β_1 is negative and significant across all subsamples in Table 9. Firms whose private signal of input inflation is lower than the aggregate inflation level today will tend to predict lower inflation in the next 12 months. One rationalisation of this finding is that firms predict lower inflation on their cost side to transmit to the overall economy in later periods, leading to lower aggregate inflation. This piece of evidence squares with the fact that treated firms continue to use their input costs to form expectations and suggests that firms perceive input-output linkages in the economy.

Table 9: TREATED GROUP: SIGN OF THE SHOCK

Regression I	(1)	(2)	(3)	(4)
Dependent Variable:	$E\pi_{t+12}$	$E\pi_{t+12}$	$E\pi_{t+12}$	$E\pi_{t+12}$
	WHOLE SAMPLE	2016Q3-2019Q4	2020Q1-2021Q4	2022Q1-2023Q4
$POS_{it}(\beta_1)$	-0.1273***	-0.08172***	-0.1858***	-0.2730***
	(0.0282)	(0.0201)	(0.0484)	(0.0731)
L.$E\pi_{t+12}$	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Area x Time FE	Yes	Yes	Yes	Yes
Sector x Time FE	Yes	Yes	Yes	Yes
Observations	16,674	6,047	4,290	4,994

Source: Bank of Italy, Survey on inflation and growth expectations, 2016Q3-2023Q4
Driscoll-Kraay Standard Errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Given how treated firms use the sign of the informational shock to update their inflation expectations, we want to test whether firms use their input cost inflation beyond the information contained in the sign of the informational shock and whether such elasticity is differential for negative and for positive informational shocks. This second specification reads:

$$\mathbb{E}_{it}[\pi_{t+k}] = \beta_1 POS_{it} + \beta_2\, \Delta c_{it} + \beta_3\, \Delta c_{it} \times POS_{it} + \gamma_i + \gamma_{SEC\times t} + \gamma_{AREA\times t} + \epsilon_{it} \tag{II}$$

Table 10 reports the estimation for specification II. When controlling for the level of input inflation, the sign of the informational shock is only relevant in subsample 2016Q3-2019Q4. What is more, the level of input inflation per se is relevant beyond the information already contained in the sign of the informational shock in the whole sample. Lastly, the elasticity to input costs is not asymmetric, as the coefficient β_2 is never significant in all subsamples: treated firms that receive a positive informational shock do not levy on

Table 10: TREATED GROUP: SIGN AND ELASTICITY INTERACTION

Regression II	(1)	(2)	(3)	(4)
Dependent Variable:	$E\pi_{t+12}$	$E\pi_{t+12}$	$E\pi_{t+12}$	$E\pi_{t+12}$
	WHOLE SAMPLE	2016Q3-2019Q4	2020Q1-2021Q4	2022Q1-2023Q4
POS_{it} (β_1)	-0.0786**	-0.0738**	-0.1169	-0.2624
	(0.0298)	(0.0116)	(0.086)	(0.1705)
Δc_{it} (β_2)	0.00935*	0.0012	0.01386	0.0052
	(0.0055)	(0.0274)	(0.0104)	(0.009)
$\Delta c_{it} \times POS_{it}$ (β_3)	0.0047	-0.0174	-0.0031	0.0074
	(0.0142)	(0.0274)	(0.0279)	(0.0227)
L.$E\pi_{t+12}$	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	Yes
Area x Time FE	Yes	Yes	Yes	Yes
Sector x Time FE	Yes	Yes	Yes	Yes
Observations	16,674	6,047	4,290	4,994

Source: Bank of Italy, Survey on inflation and growth expectations, 2016Q3-2023Q4
Driscoll-Kraay standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

their input inflation differentially with respect to firms with negative informational shocks.

6 Conclusion

While research on firms' expectations provided evidence rejecting the FIRE hypothesis (Coibion et al., 2019), deviations from Full Information are not sufficient to explain the heterogeneity in expectations. Firms devoting little attention to aggregate conditions receive noisy signals about the aggregate. They optimally decide to respond weakly to these signals. Yet they are very attentive to their own business signals. Whether firms use such private signals to infer about current inflation, given the high variance of input inflation across firms, can be a relevant factor in explaining the cross-sectional heterogeneity of firms' expectations. We argue that this is indeed the case.

Using micro-data on reported inflation expectations and input cost inflation, we show that firms rely differentially on their input cost inflation to predict future inflation realisations, as a function of exogenous changes in their information sets to form expectations. Specifically, firms informed the about current inflation realisation levy less heavily on signals from their inputs inflation. Furthermore, we show that these treatment effects are attenuated in high-inflation settings. Lastly, the distance between individual input inflation and aggregate inflation is a relevant factor in forming inflation expectations within the cross-section of treated firms. More generally, these results provide a potential explanation for the disagreement among firms about aggregate economic conditions.

This work provides empirical support for deviations from Full Information through the channel of learning from local prices by firms, adding to Andrade et al. (2022) and Albagli et al. (2022). Firms, analogously to households, use part of the signal coming from their local prices to infer about the aggregate. These results provide the direct micro-level empirical evidence supporting the mechanism of "island" models pioneered by Lucas (1972), in which firms observe signals which are combinations of idiosyncratic or industry shocks and aggregate shocks.

This evidence has important macroeconomic implications. Within the setting of Islands models, recent work has theorised a new role for monetary policy and its communication strategies. Central banks can affect firms' pricing behaviour and equilibrium prices by effectively communicating the current inflation developments. It has been shown that the level of equilibrium markups in the economy is increasing in the extent to which agents learn from prices (Janssen and Shelegia, 2015; Gaballo and Paciello, 2023).

More broadly, our results speak to the growing gap between macroeconomic models assuming the FIRE hypothesis and the recent empirical evidence of the presence of non-negligible information rigidities on the part of economic agents. In standard NK models, expectations about the future adjust immediately to shocks and can provide a powerful propagation mechanism even for small aggregate shocks into the decisions of forward-looking agents. On the contrary, evidence of imperfect information implies that this channel of expectations might be much weaker when it comes to macroeconomic shocks. This friction is important as it attenuates the power of stabilisation policies relying on macroeconomic expectations.

Lastly, consistent with the previous literature on learning in different inflation settings, we find that the treatment effect is attenuated in periods of high inflation, as it is easier for firms to be aware of the current inflation state. We find, however, that both treated and untreated groups keep respectively the same reliance on inputs when transitioning to a higher inflation regime. We speculate that firms understood the nature of the inflation shock in the sample period. Then two balancing forces would be at play. First, the pervasiveness of the inflation read could reduce the reliance on input costs as everybody is more informed about the inflation level. Second, the supply-shock nature of the inflation shock would increase the reliance on input costs as such signals are now more predictive of future inflation. Understanding how firms use their input prices depending on how they interpret inflationary shocks, might constitute a future path of research.

References

E. Albagli, F. Grigoli, and E. Luttini. Inflation expectations and the supply chain. Working paper 22/161, International Monetary Fund, Aug. 2022.

P. Andrade, O. Coibion, E. Gautier, and Y. Gorodnichenko. No firm is an island? how industry conditions shape firms' expectations. *Journal of Monetary Economics*, 125:40–56, 2022. ISSN 0304-3932. doi: https://doi.org/10.1016/j.jmoneco.2021.05.006. URL https://www.sciencedirect.com/science/article/pii/S0304393221000568.

R. Bachmann, T. O. Berg, and E. R. Sims. Inflation expectations and readiness to spend: Cross-sectional evidence. *American Economic Journal: Economic Policy*, 7(1): 1–35, February 2015. doi: 10.1257/pol.20130292. URL https://www.aeaweb.org/articles?id=10.1257/pol.20130292.

M. Bottone and A. Rosolia. Monetary policy, firms' inflation expectations and prices: causal evidence from firm-level data. Temi di discussione (Economic working papers) 1218, Bank of Italy, Economic Research and International Relations Area, 2019. URL https://EconPapers.repec.org/RePEc:bdi:wptemi:td_1218_19.

O. Coibion, Y. Gorodnichenko, and T. Ropele. Inflation Expectations and Firm Decisions: New Causal Evidence*. *The Quarterly Journal of Economics*, 135(1):165–219, 09 2019. ISSN 0033-5533. doi: 10.1093/qje/qjz029. URL https://doi.org/10.1093/qje/qjz029.

O. Coibion, D. Georgarakos, Y. Gorodnichenko, and M. van Rooij. How does consumption respond to news about inflation? field evidence from a randomized control trial. *American Economic Journal: Macroeconomics*, 15(3):109–52, July 2023. doi: 10.1257/mac.20200445. URL https://www.aeaweb.org/articles?id=10.1257/mac.20200445.

F. D'Acunto, U. Malmendier, and M. Weber. What do the data tell us about inflation expectations? In *Handbook of economic expectations*, pages 133–161. Elsevier, 2023.

M. Draghi. Press conference, October 22 2015. URL https://www.ecb.europa.eu/press/pressconf/2015/html/is151022.en.html. Accessed: 2024-05-30.

J. Driscoll and A. Kraay. Consistent covariance matrix estimation with spatially dependent panel data. *The Review of Economics and Statistics*, 80(4):549–560, 1998. URL https://EconPapers.repec.org/RePEc:tpr:restat:v:80:y:1998:i:4:p:549-560.

F. D'Acunto, U. Malmendier, J. Ospina, and M. Weber. Exposure to grocery prices and inflation expectations. *Journal of Political Economy*, 129(5):1615–1639, 2021.

ECB. Economic bulletin, issue 7. Technical report, by Tobias Schuler and Hannah-Maria Hildenbrand and Martina di Sano, 2022. URL https://www.ecb.europa.eu/press/economic-bulletin/focus/2022/html/ecb.ebbox202207_01~cd275e145d.en.html.

G. Gaballo and L. Paciello. Spending allocation under nominal uncertainty: A model of effective price rigidity*, February 2023.

J. Galí. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press, 2008.

M. Janssen and S. Shelegia. Consumer search and double marginalization. *American Economic Review*, 105(6):1683–1710, June 2015. doi: 10.1257/aer.20121317. URL https://www.aeaweb.org/articles?id=10.1257/aer.20121317.

R. Lucas. Expectations and the neutrality of money. *Journal of Economic Theory*, 4 (2):103–124, 1972. URL https://EconPapers.repec.org/RePEc:eee:jetheo:v:4:y:1972:i:2:p:103-124.

R. Lucas. Some international evidence on output-inflation tradeoffs. *American Economic Review*, 63(3):326–34, 1973. URL https://EconPapers.repec.org/RePEc:aea:aecrev:v:63:y:1973:i:3:p:326-34.

R. Lucas and T. Sargent. After keynesian macroeconomics. *Quarterly Review*, 3(Spr), 1979. URL https://EconPapers.repec.org/RePEc:fip:fedmqr:y:1979:i:spr:n:v.3no.2.

R. J. Lucas. Econometric policy evaluation: A Critique. *Carnegie-Rochester Conference Series on Public Policy*, 1(1):19–46, January 1976. URL https://ideas.repec.org/a/eee/crcspp/v1y1976ip19-46.html.

J. F. Muth. Rational expectations and the theory of price movements. *Econometrica*, 29 (3):315–335, 1961. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1909635.

L. Pilossoph and J. M. Ryngaert. Job search, wages, and inflation, September 2023. JEL Codes: E31, J3, J6.

T. Ropele. Inflation expectations and price setting behavior: evidence from business survey data. Banca d'Italia Occasional Papers 1218, Bank of Italy, Economic Research and International Relations Area, 2019.

T. Ropele, Y. Gorodnichenko, and O. Coibion. Inflation expectations and misallocation of resources: evidence from italy. Temi di discussione (Economic working papers) 1437, Bank of Italy, Economic Research and International Relations Area, 2023. URL https://EconPapers.repec.org/RePEc:bdi:wptemi:td_1437_23.

M. Weber, B. Candia, T. Ropele, R. Lluberas, S. Frache, B. Meyer, S. Kumar, Y. Gorodnichenko, D. Georgarakos, O. Coibion, G. Kenny, and J. Ponce. Tell me something i don't already know: Learning in low and high-inflation settings. Technical report, IZA - Institute of Labor Economics, 2023. URL http://www.jstor.org/stable/resrep57509.

J. L. Yellen. Macroeconomic Research After the Crisis : a speech at The Elusive 'Great' Recovery: Causes and Implications for Future Business Cycle Dynamics. Speech 915, Board of Governors of the Federal Reserve System (U.S.), Oct. 2016. URL https://ideas.repec.org/p/fip/fedgsq/915.html.