The Heterogeneity of the Inflation Expectations of Italian Firms along the Business Cycle*

Laura Bartiloro, Marco Bottone, and Alfonso Rosolia
Bank of Italy

We investigate how the cross-sectional heterogeneity of firms' inflation expectations reflects information availability and awareness of recent macroeconomic developments, observable firm characteristics, and broader macroeconomic developments using the Bank of Italy's survey on businesses' inflation and growth expectations. We find that, on average, about half of the dispersion of expectations is traceable to a lack of information about the most recent price developments; firms incorporate new information into their expectations within a quarter; the dispersion of expectations is related in a statistically significant way to some important aggregate economic variables, and it is greater when current inflation is farther away from the ECB's price stability goal. Since 2015 the weight attributed to prior beliefs of low inflation has steadily increased and the uncertainty surrounding them has decreased. Furthermore, since 2014 the empirical connection between the dispersion of expectations and the distance from the ECB price stability goal has become considerably weaker. These two facts suggest an increased risk of inflation expectations being de-anchored.

JEL Codes: D22, D8, E31.

1. Introduction

Close monitoring of inflation expectations is a crucial ingredient for the conduct of monetary policy. The increasing availability of

*The views expressed are our own and are not necessarily shared by the Bank of Italy. We thank for their comments Roc Armenter, Andrea Brandolini, Giuseppe Ilardi, Lucrezia Reichlin, Luigi Federico Signorini, an anonymous referee, and seminar audiences at the Bank of Italy, the 33rd Annual Congress of the European Economic Association (Cologne, 2018), and the 35th CIRET Conference (Rio de Janeiro, 2018).

high-frequency survey data on expectations of consumers, professional forecasters, and firms has made it possible to complement standard market based measures of expected inflation with those directly reported by economic agents. The availability of microdata on expectations of single decision units has also helped investigate the process by which expectations are defined and underscore the relevance of this knowledge for the correct formulation of macroeconomic models and policy goals.[1]

In this paper we add to this large body of evidence in three respects. First, we provide evidence about *firms*' inflation expectations. Price-setting firms play a central role in shaping aggregate price dynamics; modern New Keynesian macroeconomic models posit that monopolistic firms set prices in a staggered fashion and with an eye to their competitors' current and future price strategies and thus to developments in aggregate prices. However, to use Ben Bernanke's words, "Information on the price expectations of businesses—who are, after all, the price setters in the first instance—as well as information on nominal wage expectations is particularly scarce" (Bernanke 2007). Indeed, all empirical analyses focus on inflation expectations of consumers or professional forecasters.[2]

Second, in the few available sources where firms are the relevant decision unit, the information collected either relates to their own prices or costs or, when about broader inflation measures, it is, with very few exceptions, of a qualitative nature, typically reported in brackets or, more often, in terms of a dichotomous choice between "higher" and "lower" future average price levels. We rely instead on point inflation expectations reported by firms, which we think adds to the existing literature even if using this information requires

[1]For example, Ball, Mankiw, and Reis (2005) show that if price setters only slowly incorporate macroeconomic information in their prices, then the optimal monetary policy should target a price level; Gaspar, Smets, and Vestin (2006) show that optimal monetary policy has to react to cost-push shocks when agents form their expectations through adaptive learning; Coibion, Gorodnichenko, and Kamdar (2017) summarize how alternative mechanisms of expectation formation shape the correct empirical specification of the Phillips curve; Busetti et al. (2017) show that in the face of a sequence of negative shocks, the coexistence of heterogeneous expectation formation mechanisms may bring inflation off target and reinforce a de-anchoring of expectations.

[2]See Ehrmann, Pfajfar, and Santoro (2017) for a recent discussion of available studies.

strong assumptions on the way agents form and report their expectations (Manski 2018).

Third, we explore the determinants of the dispersion among firms of expectations about future price dynamics. While expectations about the same macroeconomic phenomenon are known to be heterogeneous across decision units, little is known about the nature of this heterogeneity, its cyclical properties, and whether it has some informational content of value to policymakers. To our knowledge, there are only a few studies concerned with the cyclical properties of forecast disagreement; however, these studies do not focus on firms. Mankiw, Reis, and Wolfers (2003) show that the dispersion of inflation expectations of U.S. consumers and professional forecasters moves over time with the inflation rate, its changes, and the variability of relative prices. They further argue that, contrary to models with staggered price setting, models in which rational firms set prices constrained by sticky information (e.g., Mankiw and Reis 2002) are able to generate dispersion of (rational) expectations among *firms* with these observed cyclical properties. More recently, Dovern, Fritsche, and Slacalek (2012) explore disagreement about GDP growth, inflation, and interest rates among professional forecasters and show that disagreement about inflation is higher in recessions and when inflation is more volatile, and it is lower under independent central banks; they also document that the patterns of disagreement are quite different when forecasting GDP growth and inflation. Andrade et al. (2016) exploit a survey of professional forecasters to document that disagreement about output growth, inflation, and the federal funds rate can be substantial even for long horizons far in the future but that it is not necessarily larger than that about shorter horizons, and they develop a model of expectations formation that accounts for these regularities. However, these results rely on expectations collected from *consumers* and from *professional forecasters*, not from *firms*. Our paper thus complements theirs in providing novel empirical evidence on the determinants and cyclical properties of the dispersion of *firms*' inflation expectations.

Fourth, we exploit a unique feature of our data to explore how different degrees of awareness about recent macroeconomic developments contribute to the dispersion of expectations and to estimate an upper bound to the delay with which relevant information is incorporated in firms' inflation expectations. We thus contribute to

a growing literature that addresses the process of expectations revision in light of novel information. For example, Lahiri and Sheng (2008) use fixed-event expected GDP growth of professional forecasters to show that disagreement is initially largely due to heterogeneous prior beliefs and, as the forecast horizon becomes shorter, to different interpretations of public information. Patton and Timmermann (2010) also document the role of heterogeneity in priors and learning models in generating disagreement and also show it to increase during recessions.

In this paper we use data drawn from the Bank of Italy's Survey of Inflation and Growth Expectations (SIGE). Several studies have used the SIGE or its predecessor to address some of the issues outlined above.[3] Visco (1987), one of the earliest studies of microdata on inflation expectations, discusses comprehensively the methodological problems involved in extracting aggregate measures of inflation expectations from individual answers and in using them as forecasts of actual inflation; he also explores firms' expectation formation mechanisms and their rationality, a hypothesis he rejects except for periods of economic stability. More recently, Fabiani and Santoro (2012) also conclude against the rationality of firms' inflation expectation. Ropele (2017) is one of the few existing studies empirically verifies the relationship between a firm's own expected pricing decisions and its inflation expectations; he finds a positive and statistically significant correlation which has disappeared with the unfolding of the sovereign debt crisis. Finally, Bottone and Rosolia (2018) show that firms' inflation expectations respond to high-frequency market-based measures of unanticipated monetary policy news in a statistically significant and theoretically consistent way; however, price decisions are largely unaffected by the same shocks.

In the next section we briefly describe the SIGE, focusing especially on the collection of firm-level inflation expectations. We then explore the role of several determinants in shaping the cross-sectional dispersion of expectations, namely information availability, firms' observable characteristics, and macroeconomic developments. We draw our conclusions in the final section.

[3]Between 1952 and 1999 the Mondo Economico survey collected, twice a year, inflation expectations of a panel of experts, including CEOs and entrepreneurs. The survey was discontinued in 1999 and the Bank of Italy, in collaboration with *il Sole 24Ore*, the main Italian financial newspaper, launched the SIGE.

2. The Bank of Italy Survey of Inflation and Growth Expectations

The Bank of Italy started the Survey of Inflation and Growth Expectations (SIGE) in the fourth quarter of 1999. The SIGE is run quarterly on a sample of currently about 1,000 manufacturing and service firms with at least fifty employees; construction firms were added to the sample in 2013. The sample is stratified by sector of activity, firm size, and area (Bank of Italy 2017). The data-collection process lasts at most three weeks, usually in the last month of the reference quarter.[4] To our knowledge this survey is the longest-running survey that systematically collects point expectations of firms about consumer price inflation at several horizons and quantitative information on past and expected own selling price changes;[5] the questionnaire also collects sentiment information on aggregate cyclical developments as well as on own business real and financial conditions.[6]

At the international level, existing surveys focus generally either on consumers' (e.g., Michigan Surveys of Consumers Attitudes) or on professional forecasters' (e.g., the Survey of Professional Forecasters of the European Central Bank (ECB) and the Federal Reserve (Fed); Consensus Forecasts) expectations. When the relevant observation unit is a firm, long-running surveys typically collect *qualitative* information on price developments within the broader context of business climate surveys; alternatively, when point expectations are collected, they typically refer to *own* price or unit costs (e.g., the Federal Reserve Bank of Atlanta's Business Inflation Expectations Survey since 2011) and industry-wide price developments (e.g., the business surveys run by the Confederation of British Industries; Boneva et al. (2016) rather than to market-wide indexes of price developments; measures of expected inflation are generally obtained by subsequent aggregation of firms' own price/cost expected developments.

[4]Specifically, the first-quarter survey is run in early March, the second-quarter survey in early June, the third-quarter one in early September, and the fourth-quarter one in early December.

[5]The survey focuses on consumer price inflation in Italy; until the end of 2004, expectations on consumer price inflation in the euro area were also collected.

[6]The survey data can be accessed through the Bank of Italy remote processing system, BIRD; details are available at www.bancaditalia.it/statistiche/basidati/bird/index.html.

An exception is represented by Coibion, Gorodnichenko, and Kumar (2015) who, acknowledging the lack of suitable information to properly study the formation of inflation expectations of firms, ran a quantitative survey very similar to the Bank of Italy's among New Zealand businesses between the end of 2013 and the beginning of 2015 collecting information on the expected change of prices in the overall economy along with a host of other relevant information.

The Bank of Italy's SIGE collects instead since its beginning firms' *point inflation expectations* at several horizons (currently, six months, one year, two years, and an average between three and five years). Point expectations are collected by means of a single question worded in two different ways and administered to randomly selected subsamples. The first version was used for all respondents until the second quarter of 2012 and is currently administered to about two-thirds of respondents, which we dub "informed" agents. It provides them with information on current inflation developments which, due to dissemination delays, usually refers to two months before the survey. Specifically, the question reads as follows:

> *In [month of most recent inflation release] consumer price inflation, measured by the 12-month change in the HARMONIZED INDEX OF CONSUMER PRICES was [xx] per cent in Italy and [yy] per cent in the euro area. What do you think it will be in Italy . . . [in six months], [in one year], [in 2 years], [on average between 3 and 5 years]?*

The wording used for the remaining one-third of respondents, dubbed "non-informed" agents, since the third quarter of 2012 does not provide any information whatsoever on current developments:

> *What do you think consumer price inflation in Italy, measured by the 12-month change in the HARMONIZED INDEX OF CONSUMER PRICES, will be . . . [in six months], [in one year], [in 2 years], [on average between 3 and 5 years]?*

Unfortunately, respondents are not explicitly requested to report a specific statistic of their (if any) subjective probability distribution. Engelberg, Manski, and Williams (2009) have shown that indeed reporting practices differ across reporting units in the Fed's Survey of Professional Forecasters, although the majority seems to report

values not inconsistent with their subjective means. Lacking additional information, in the following we will make the working assumptions that agents hold probability distributions on price dynamics and that they report the implied mean value. Under these assumptions, the survey offers several angles to look at inflation expectations of firms, their formation, and cyclical properties. To this we turn in the next section.

3. The Sources of Expectations' Heterogeneity

Figure 1 displays the long time series of specific percentiles of the cross-sectional distribution of businesses' one- and two-year-ahead Italian consumer price inflation expectations of informed agents.[7] After a protracted period of substantial stability around a level slightly above the one consistent with the ECB's objective of price stability for the euro area, businesses' inflation expectations have recorded wide swings before settling, for both horizons, at a level below the ECB's goal between 2014 and the beginning of 2017.

The figure shows both the significant degree of heterogeneity of inflation expectations and its variation over time. Until 2007, before the onset of the global financial crisis, the gap between the 80th and 20th percentiles of one-year-ahead expected inflation was in the 0.4–0.6 percentage point range, around a median inflation expectation gradually converging towards 2 percent. Since 2008, the gap has widened, remaining almost continuously above 0.8 percentage point until mid-2015, when it went back to the initial values but around a much lower median inflation rate; in periods characterized by particularly high median inflation expectations also the 20th–80th percentile gap has been much larger, above 1.5 percentage points. Consistent with the evidence in, for example, Lahiri and Sheng (2008) and Patton and Timmermann (2010) and contrary to what was documented for professional U.S. forecasters in Andrade et al. (2016), the dispersion at the longer horizon of twenty-four months is somewhat larger, although the time pattern is broadly similar to that of the shorter-horizon expectations.

[7]Expectations at the six-month horizon are collected only since 2010:Q4 and expectations at the longer three-to-five-year horizon only since 2014:Q1. In the following we focus only on the horizons of twelve and twenty-four months.

Figure 1. Businesses' Inflation Expectations and Dispersions

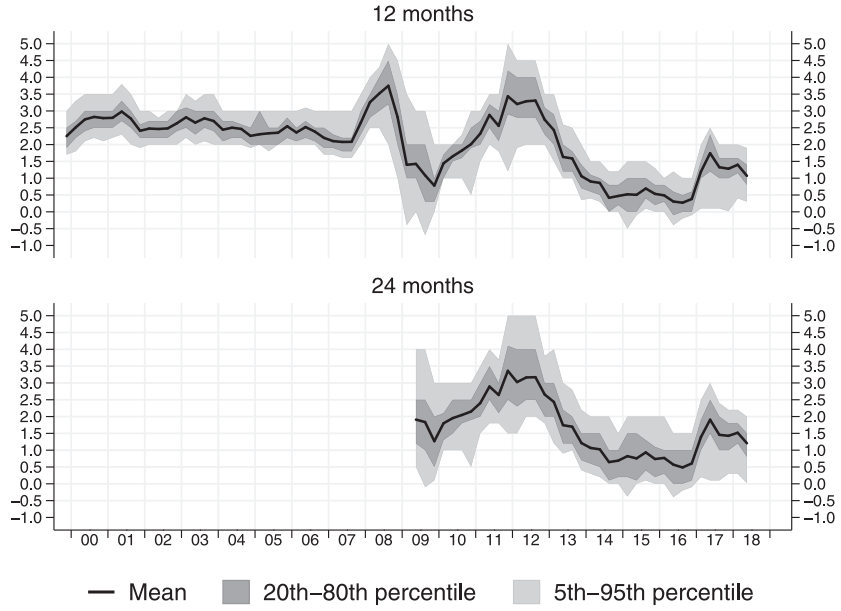

Source: Own elaborations of Bank of Italy Survey of Inflation and Growth Expectations.
Note: Sample only includes informed agents.

The figure thus shows a substantial amount of dispersion of businesses' inflation expectations and the fact that it moves quite a lot over the business cycle. In the rest of the paper we provide evidence on the sources of this heterogeneity, focusing on three factors: information, businesses' observable characteristics, and business cycle developments.

3.1 Information (or Lack Thereof)

The SIGE allows us to assess, though only since 2012:Q3, the role played by information and firms' awareness in shaping businesses' expectations and their dispersion.

The upper panels of figure 2 plot the mean one- and two-year-ahead inflation expectation of informed and non-informed

Figure 2. Expectations of Informed and Non-informed Respondents and Consensus Forecasts

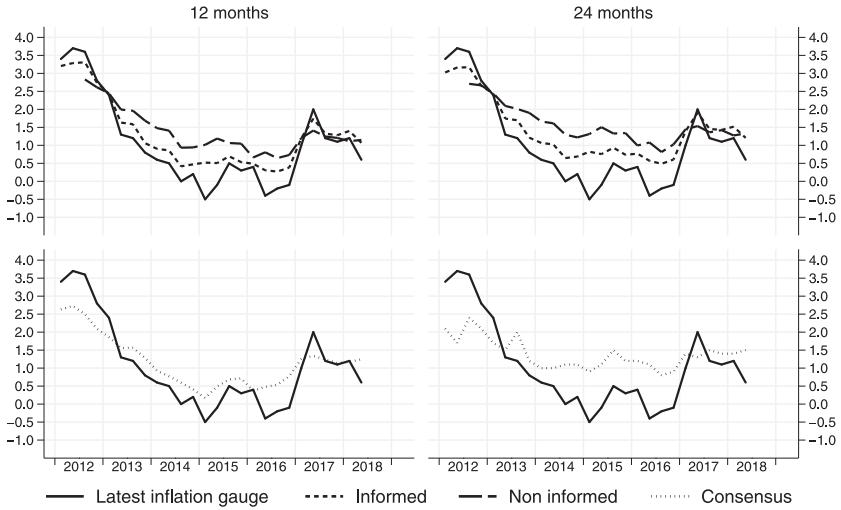

Source: Own elaborations of Bank of Italy Survey of Inflation and Growth Expectations.

respondents, along with the information on the most recent inflation gauge provided to informed agents over the period for which both questions are asked. The mean expected inflation of informed agents closely tracks the most recent official inflation rate provided with the questionnaire, and so does the mean expected inflation of non-informed agents, although the latter fell somewhat less in the aftermath of the sovereign debt crisis. For both groups and at both horizons, expectations did not fall below zero, even in quarters recording negative inflation. Overall, the one-year-ahead inflation expectations of firms behave very much like those of professional forecasters collected by Consensus Forecasts, which also appear to closely track current inflation developments;[8] longer expectations

[8]Monthly Consensus Forecasts are collected with reference to a fixed date (at the end of the current year, at the end of the next year). We transform them into fixed-horizon forecasts (in the next twelve months) following Dovern, Fritsche,

Figure 3. Forecast Errors

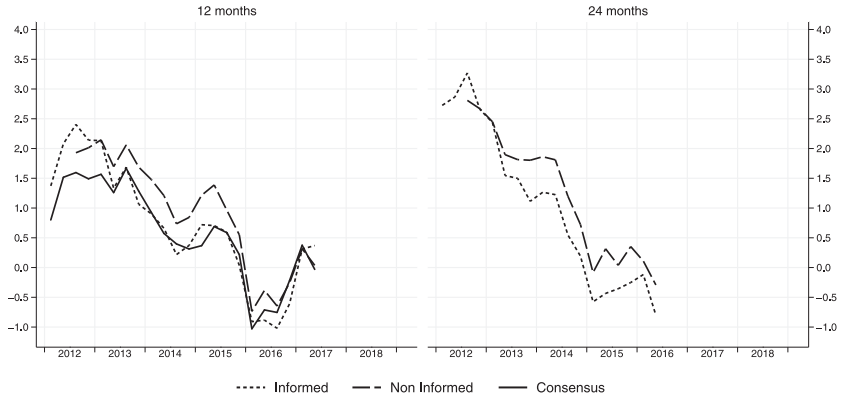

Source: Own elaborations of Bank of Italy Survey of Inflation and Growth Expectations.

expressed by professional forecasters appear instead more stable than those reported by survey respondents.[9]

As concerns their predictive power, the mean expectations of informed and non-informed respondents as well as those of Consensus Forecasters generally fail to match realized inflation. Figure 3 plots the mean forecast error (defined as the difference between mean expected and realized inflation) against the quarter in which expectations were collected. It shows that both groups of firms and Consensus Forecasters have failed to anticipate both the fall in inflation in the aftermath of the euro-area sovereign debt crisis and the more recent rebound in price dynamics. Although the forecast performances of the two groups of firms moved in a largely similar way, non-informed respondents appear to have taken longer to revise

and Slacalek (2012). Specifically, we weight forecasts for the current and following year by the number of months each forecast contributes to the twelve-month horizon.

[9]Consensus Forecasts does not collect information to measure two-year-ahead forecasts. We thus exploit data collected at quarterly frequency on expected yearly consumer price inflation for the following six to seven quarters and approximate the object of interest with the expectations reported quarterly for the farthest quarter available.

their expectations to the low levels of inflation gauges of 2016, their prediction error being systematically larger than that of informed respondents until early 2016; afterwards, their forecasts have been somewhat more precise.

More quantitative descriptive evidence in this sense is presented in table 1. We report results of simple linear regressions of mean expectations at twelve (panel A) and twenty-four (panel B) months of informed (columns 1–3) and non-informed (columns 4–6) respondents and of Consensus Forecasts (columns 7–9) on current and past inflation rates. Columns 1–3 show that mean expectations of informed respondents basically reflect only current inflation developments; past inflation rates correlate at best very weakly and not in a statistically significant way; on the contrary, columns 4–6 show that the mean expectations of non-informed respondents are rather strongly correlated with both current and past inflation developments. Yet, compared with informed respondents, current and past price dynamics explain a lower share of the overall variance of median expectations of non-informed agents. As a comparison, the central tendency of Consensus Forecasts is also correlated in a statistically significant way with both current and past inflation realizations. For all three groups, the unexplained component does not turn out to be in a statistically significant correlation with future inflation realizations.

Table 1 focuses on cross-sectional means. However, as pointed out by Keane and Runkle (1990), these may suffer, among other problems, from serious aggregation biases that may mute or induce some degree of consistency with actual data. Therefore, in table 2 we replicate the same regressions on individual data. In panel A we estimate least-square regressions of individual expectations at twelve (columns 1–8) and twenty-four (columns 8–16) months of informed and non-informed respondents, controlling for a host of individual characteristics with dummies for firm size, sector of activity, and location of headquarters; in panel B we introduce instead firm fixed effects to absorb systematic differences. Point estimates are generally highly consistent with those obtained from aggregate data reported in table 1: the expectations of informed agents are more tightly linked to current developments than those of non-informed respondents. We also show (columns 4, 8, 12, and 16) that including the contemporaneous one-year-ahead Consensus Forecast,

Table 1. Average Inflation Expectations of Informed and Non-informed Respondents and Consensus Forecasts

	Informed			Non-informed			Consensus		
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**
A. Expectations Twelve Months Ahead									
π_t	0.718	0.608	0.633	0.385	0.246	0.337	0.470	0.402	0.449
	(0.032)	(0.060)	(0.058)	(0.040)	(0.113)	(0.077)	(0.038)	(0.082)	(0.073)
π_{t-1}		0.159	0.100		0.268	0.056		0.133	0.026
		(0.052)	(0.057)		(0.098)	(0.077)		(0.072)	(0.073)
π_{t-4}	0.075		0.051	0.196		0.183	0.099		0.093
	(0.023)		(0.026)	(0.029)		(0.035)	(0.028)		(0.033)
$\bar{R}^2$	0.971	0.970	0.974	0.918	0.812	0.917	0.922	0.893	0.919
B. Expectations Twenty-Four Months Ahead									
π_t	0.635	0.553	0.571	0.297	0.179	0.263			
	(0.031)	(0.057)	(0.057)	(0.041)	(0.109)	(0.079)			
π_{t-1}		0.118	0.074		0.235	0.041			
		(0.050)	(0.057)		(0.095)	(0.079)			
π_{t-4}	0.055		0.037	0.177		0.168			
	(0.022)		(0.026)	(0.030)		(0.036)			
$\bar{R}^2$	0.965	0.964	0.966	0.880	0.751	0.875			

Notes: Standard errors are in parentheses. Sample frame: 2012:Q3–2018:Q2. The dependent variable is average expected inflation of informed (columns 1–3) and non-informed (columns 4–6) respondents and Consensus Forecasts (columns 7–9). Consensus Forecasts are adjusted as in Dovern, Fritsche, and Slacalek (2012); not available for the twenty-four-month horizon.

Table 2. Inflation Expectations of Informed and Non-informed Respondents

	Twelve Months Ahead								Twenty-Four Months Ahead							
	Informed				Non-informed				Informed				Non-informed			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
A. Least-Squares Regressions																
π_t	0.719	0.607	0.630	0.553	0.386	0.256	0.337	0.368	0.634	0.549	0.566	0.518	0.302	0.189	0.262	0.341
	(0.007)	(0.012)	(0.013)	(0.021)	(0.018)	(0.033)	(0.033)	(0.056)	(0.008)	(0.014)	(0.014)	(0.024)	(0.019)	(0.034)	(0.035)	(0.060)
π_{t-1}		0.159	0.104	0.100		0.274	0.058	0.058		0.120	0.079	0.077		0.243	0.048	0.047
		(0.011)	(0.013)	(0.013)		(0.028)	(0.033)	(0.033)		(0.012)	(0.014)	(0.014)		(0.030)	(0.035)	(0.035)
π_{t-4}	0.073		0.048	0.033	0.198		0.183	0.190	0.055		0.036	0.026	0.178		0.166	0.183
	(0.005)		(0.006)	(0.007)	(0.013)		(0.015)	(0.018)	(0.006)		(0.007)	(0.008)	(0.013)		(0.016)	(0.019)
CF				0.171				−0.065				0.107				−0.168
				(0.038)				(0.097)				(0.044)				(0.103)
$\bar{R}^2$	0.592	0.592	0.594	0.595	0.216	0.197	0.216	0.216	0.457	0.457	0.458	0.458	0.148	0.133	0.148	0.149
B. Fixed-Effects Regressions																
π_t	0.650	0.575	0.584	0.526	0.389	0.339	0.341	0.281	0.578	0.524	0.531	0.504	0.327	0.280	0.281	0.254
	(0.008)	(0.012)	(0.012)	(0.020)	(0.021)	(0.030)	(0.030)	(0.050)	(0.009)	(0.013)	(0.013)	(0.022)	(0.023)	(0.032)	(0.031)	(0.052)
π_{t-1}		0.157	0.087	0.086		0.229	0.067	0.071		0.117	0.062	0.061		0.190	0.064	0.065
		(0.010)	(0.012)	(0.012)		(0.025)	(0.029)	(0.029)		(0.011)	(0.013)	(0.013)		(0.026)	(0.031)	(0.031)
π_{t-4}	0.096		0.073	0.057	0.183		0.162	0.143	0.074		0.058	0.050	0.145		0.126	0.118
	(0.005)		(0.006)	(0.008)	(0.013)		(0.016)	(0.020)	(0.006)		(0.007)	(0.008)	(0.014)		(0.017)	(0.021)
CF				0.136				0.136				0.063				0.062
				(0.037)				(0.090)				(0.041)				(0.095)
$\bar{R}^2$	0.603	0.600	0.605	0.606	0.451	0.440	0.452	0.452	0.544	0.542	0.545	0.545	0.426	0.420	0.427	0.427
$\bar{R}^2$	0.541	0.542	0.544	0.545	0.224	0.201	0.224	0.224	0.403	0.404	0.405	0.405	0.147	0.130	0.148	0.148

Notes: Standard errors are in parentheses. Sample frame: 2012:Q3–2018:Q2. The dependent variable is average expected inflation of informed (columns 1–4 and 9–12) and non-informed (columns 5–8 and 13–16) respondents twelve months ahead (columns 1–8) and twenty-four months ahead (columns 9–16). Panel A: All regressions include dummy variables for size, sector, and area. Panel B: All regressions include firm fixed effects. Consensus Forecasts (CF) always refer to twelve months ahead. The last row reports adjusted R^2 from a regression of individual data excluding individual controls or fixed effects.

a proxy for other available information, does not modify substantially the patterns of correlation between individual expectations and current and lagged inflation. Importantly, the same specification applied to informed and non-informed individuals can explain a very different share of the overall variance: the specifications in panel A, in which controls for individual observable characteristics are included, explain about half of the variance of informed agents and only roughly one-fifth of that of non-informed respondents; allowing for firm fixed effects, in panel B, does not substantially affect the explanatory power of the regression in the case of informed respondents, while it doubles it among non-informed ones. Finally, the last row of table 2 reports the adjusted R^2 from a regression of individual expectations on only the aggregate variables used in the corresponding column. Interestingly, allowing for observable characteristics or firm fixed effects does not increase the explanatory power of the regressions for informed agents in comparison with this latter simple specification; on the contrary, firm fixed effects explain a larger share of the variance of non-informed respondents while not affecting the correlations of the central tendencies with inflation and its lags.

These results stand in contrast to those usually obtained from studies of the sources of heterogeneity of households' expectations of macro variables. For example, Souleles (2004) finds that in the Michigan Surveys of Consumers Attitudes households' inflation forecast errors are correlated with their observable sociodemographic characteristics; Ehrmann, Pfajfar, and Santoro (2017) confirm these results and also show that expectations are shaped by the household's financial situation and purchasing habits. On the contrary, observable characteristics do not seem to play a major role in driving disagreement among Italian firms, and systematic differences among respondents are significantly dampened by information availability, suggesting they are related to different information sets or priors.

3.2 Updating Expectations

Looking at the data through the lens of Bayesian learning can shed further light on the strength of agents' prior information and on the uncertainty surrounding it. More specifically, let π_{it}^{h} be the mean of the (normally distributed) prior inflation expectation at horizon

$h = \{12, 24\}$ of firm i at time t and σ_{it} its variance; let π_t^* be the (normally distributed) signal and s_t its variance. Then i's posterior mean is

$$\hat{\pi}_{it} = \pi_{it} \frac{s_t}{s_t + \sigma_{it}} + \pi_t^* \frac{\sigma_{it}}{s_t + \sigma_{it}}. \quad (1)$$

The ideal setting to assess how information is used to update one's prior requires that expectations are collected from the same respondent before and after information is provided (e.g., Coibion, Gorodnichenko, and Kumar 2015, Armantier et al. 2016, and Cavallo, Cruces, and Perez-Truglia 2017); alternatively, the updating behavior has been studied using fixed-event forecasts such as those typically collected within high-frequency surveys of professional forecasters that allow us to study how one's expectations evolve as the forecast date approaches (e.g., Lahiri and Sheng 2008, Patton and Timmermann 2010).

However, since in the SIGE information on the most recent official inflation rate is randomly provided to agents, we can adapt equation (1) to our empirical setting. Specifically, we can assume that expectations elicited from informed and non-informed agents are valid estimates of, respectively, posterior and prior expectations of the same population. Therefore, statistics computed on the two samples can be combined to study aspects of the learning patterns of Italian firms as concerns their inflation expectations.

Under this assumption, equation (1) implies that a regression of the mean expectations of informed agents on those of non-informed agents and on the signal gives a sense of the average weight put on both pieces of information in shaping the posterior; note also that equation (1) implies that the two coefficients sum to unity and that the constant should be nil. Moreover, it also provides an indirect test of the underlying assumption that informed and non-informed respondents only differ because of the different degree of awareness about the most recent inflation data. This is important, because the initially (as of 2012:Q3) random grouping of firms changes over time only because of sample attrition and refreshments so that continuing firms always belong to one of the two groups and thus either always or never receive the information within the survey. Clearly, they can, possibly with a delay, autonomously acquire the information so that their information sets are the same by the time the following quarter's survey is administered. If this is the case, in equation (1) older

information (including older signals) should have no independent role in shaping the posterior, since it is already incorporated in the prior.

Results are presented in table 3. Specifically, we augment equation (1) with lagged growth rates of the Industrial Production Index, a lagged quadratic term in the output gap, and lagged quarterly and yearly changes in the inflation rate as well as with the lagged cross-sectional standard deviation of price changes of major items of the Harmonised Index of Consumer Prices (HICP). These variables are meant to capture past available information on the business cycle that should be common to both informed and non-informed respondents. Regression in columns (a) are based on unconditional mean expectations at twelve and twenty-four months; those in columns (b) are instead based on means conditional on observable characteristics.[10]

The weight assigned to the prior is somewhat larger at the shorter horizon and, at both horizons, larger than that put on the signal. For all specifications, the data do not reject the null that the coefficients sum to unity, that the constant is zero, and that also all other coefficients are jointly zero.[11] Importantly, with a p-value of at least 0.74, the explanatory power of the lagged signal is negligible, suggesting that firms incorporate recent inflation readings in their expectations with a short delay of at most one quarter.

Overall, these results show that on average in forming the posterior the prior is given at best a slightly larger weight than the signal. However, the results are not informative as to whether (and how) this weight changes over time. This leads us to a further implication of equation (1).

Under the assumption that the firms face the same degree of uncertainty surrounding their prior (i.e., $\sigma_{it} = \sigma_t$), the ratio of

[10]These latter conditional means have been obtained from linear regression of individual expectations on observable firm characteristics (size, sector, location) and quarterly dummies for informed and for non-informed respondents. In columns (b) the former set of quarterly dummies is the dependent variable and the latter set is the prior.

[11]The prediction that coefficients sum to unity refers to those on the prior and on the signal on future inflation; with a slight abuse, we have assimilated the currently available inflation rate to the signal on future inflation, whereas the true signal should be, in principle, a function of current information.

Table 3. Bayesian Learning

	(1)	(2)	(3)	(4)
	Twelve Months		Twenty-Four Months	
	(a)	(b)	(a)	(b)
Prior (β_1)	0.639	0.642	0.502	0.507
	(0.155)	(0.156)	(0.198)	(0.199)
Signal (β_2)	0.396	0.390	0.418	0.411
	(0.090)	(0.091)	(0.103)	(0.104)
Constant	−0.030	−0.039	0.273	0.243
	(0.261)	(0.259)	(0.369)	(0.361)
Lag 1 of:				
Signal	0.006	0.009	0.025	0.029
	(0.067)	(0.067)	(0.078)	(0.078)
Industrial Production, q/q Change	0.033	0.033	0.028	0.027
	(0.023)	(0.024)	(0.027)	(0.027)
Industrial Production, y/y Change	−0.023	−0.022	−0.021	−0.020
	(0.018)	(0.018)	(0.022)	(0.022)
Output Gap	0.248	0.248	0.181	0.181
	(0.123)	(0.123)	(0.155)	(0.155)
Output Gap, Squared	0.625	0.592	0.376	0.340
	(0.825)	(0.825)	(0.996)	(0.995)
Inflation Rate, Squared q/q Change	−0.004	−0.004	−0.005	−0.005
	(0.009)	(0.009)	(0.011)	(0.011)
Inflation Rate, Squared y/y Change	0.029	0.030	0.077	0.080
	(0.089)	(0.089)	(0.106)	(0.106)
Dispersion in Price Changes	0.830	0.781	−3.652	−3.691
	(9.446)	(9.459)	(11.286)	(11.291)
$\bar{R}^2$	0.98	0.98	0.97	0.97
H0: $\beta_1 + \beta_2 = 1$ (p-value)	0.753	0.770	0.582	0.575
H0: All Coeffs. (Except β_1 and β_2) = 0 (p-value)	0.132	0.124	0.396	0.384

Notes: Standard errors are in parentheses. Sample frame: 2012:Q3–2018:Q2. Column (a): The dependent variable is the cross-sectional mean of the inflation expectations of informed agents at twelve (column 1) and twenty-four (column 3) months; the prior is the corresponding cross-sectional mean for non-informed agents. Column (b): The dependent variable is the set of quarter dummies estimated from an auxiliary least-squares regression of individual inflation expectations on firms' observable characteristics (size dummies, sector dummies, area dummies); "Prior" is obtained from the same regression as the interaction of quarter dummies with a dummy for non-informed agents. "Signal" is the current (at time of the survey) inflation reading provided to informed agents.

Figure 4. Weights on Prior Expectations

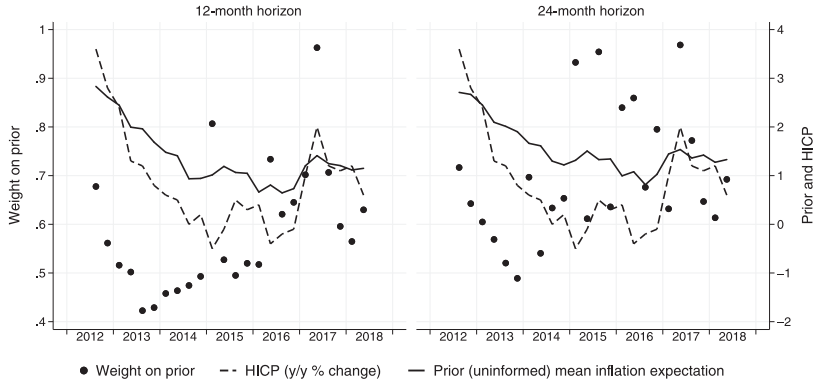

Source: Own elaborations of Bank of Italy Survey of Inflation and Growth Expectations.
Notes: Standard deviations are computed excluding observations in the bottom and top 1 percent of the inflation expectations in each quarter of, respectively, informed and non-informed respondents.

the cross-sectional standard deviations of expectations of informed agents (the posterior) to non-informed ones (the prior) is an alternative estimate of the relative weight agents put on their prior in a given quarter, $\frac{s_t}{s_t+\sigma_t}$. Figure 4 displays, for the twelve- and twenty-four-month horizons, the estimated weight of the prior expectations (left axis) and the expectations of non-informed firms (the prior itself) and current inflation (the incoming information) (both right axis). The figure shows that the information provided to respondents does reduce the dispersion of expectations; on average, the ratio of the two standard deviations is slightly higher than 0.5, consistent with the estimates presented in table 3. The figure also suggests that the relative contributions of the new information and of the prior change over time. As consumer prices quickly slowed down after the sovereign debt crisis, the weight assigned by firms to their prior expectations kept falling, implying a growing role for incoming data (i.e., a persistently declining inflation) to help them form expectations. The downward revisions seem to have stopped when current inflation stabilized at exceptionally low levels, at which point the weight given to the now lower prior inflation expectations rose again. As a consequence, the pickup in inflation recorded since 2017 has not yet been incorporated into firms' expectations, still hovering

around 1 percent at both twelve and twenty-four months, and with a high weight in the updating process. In other words, the uncertainty surrounding firms' priors first increased (relative to the precision of the signal) when current readings were showing a persistent inflation decline and then decreased but around a lower expected inflation, a pattern consistent with the gradual entrenchment of low inflation expectations.[12]

To sum up, the evidence above suggests that firms quickly, although not instantaneously, incorporate new information in their expectations and that over a period of steadily falling inflation the weight given to incoming data has gradually risen until prior expectations have reached a minimum, after which prior expectations have become gradually more relevant. It is also interesting to note that, especially at the twelve-month horizon, when current inflation occasionally fell below zero, the weight put on the (still positive) prior sharply increased. Similarly, when in the second quarter of 2017 inflation hit 2 percent, the weight on the lower prior again increased considerably. Indeed, that inflation figure turned out to be a temporary occurrence. In both cases, faced with uncommon inflation readings, firms were cautious enough not to revise their expectations as suggested by their recent updating strategy.

In light of these results, the lack of (very) recent information explains about half of the dispersion of expectations. To what extent the other half—that is, the dispersion among informed agents—is traceable to macro-determinants is the subject of the next subsection.

3.3 Macroeconomic Factors

In figure 5 we describe the empirical relationship between the cross-sectional dispersion of one-year-ahead inflation expectations

[12] As a robustness check to dispel concerns that the ratio of the two cross-sectional standard deviations might reflect the different sample sizes of informed and non-informed respondents, we bootstrapped 100 subsamples of informed agents of the same size as the non-informed sample and computed the corresponding cross-sectional standard deviations. The resulting distributions were highly concentrated around the observed standard deviations, and the cross-sectional dispersions among non-informed agents were always higher than the maximum bootstrapped standard deviations of informed agents.

Figure 5. Dispersions and Macroeconomic Developments

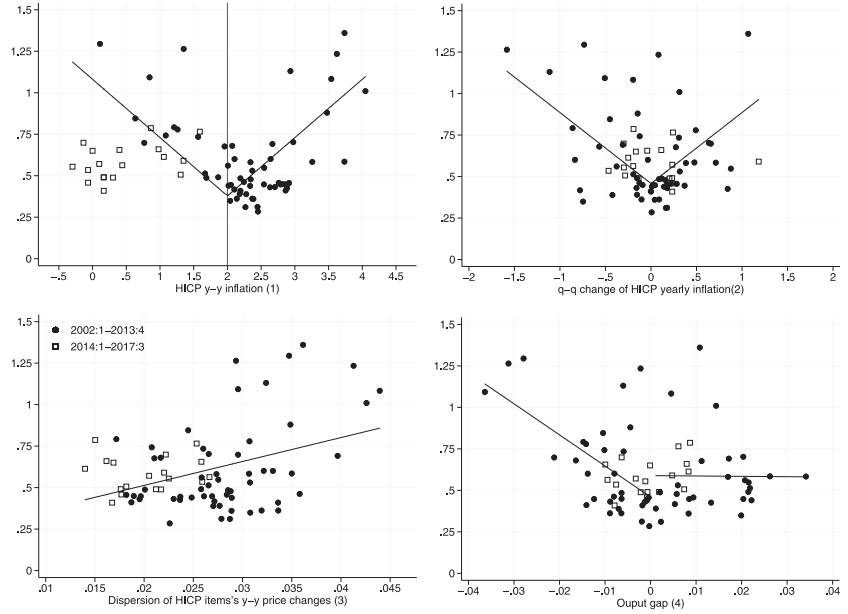

Source: Own elaborations of Bank of Italy Survey of Inflation and Growth Expectations.
Notes: The vertical axis reports the cross-sectional standard deviation of firms' inflation expectation. The horizontal axis reports (i) the year-on-year HICP inflation rate, (ii) the quarter-on-quarter change in inflation rate, (iii) the cross-sectional standard deviation of year-on-year percentage changes of the price indexes of three-digit HICP sub-items, and (iv) the output gap, obtained with a standard Hodrick-Prescott filter applied to chain-linked quarterly log GDP between 1995:Q1 and 2018:Q2.

of informed respondents, as measured by their standard deviation, and specific aspects of actual price dynamics and macroeconomic developments; broken lines visually summarize bivariate linear relationships.[13] Specifically, the upper-left panel of figure 5 plots the dispersion of firms' inflation expectations in a given quarter against the current year-on-year change in the Italian HICP; the vertical line represents the level consistent with the ECB's goal of price stability

[13]We focus only on one-year-ahead expectations to exploit the longer time span covered by the data.

in the euro area, here taken to be 2 percent for simplicity.[14] The figure shows that, for most of the time since the adoption of a common monetary policy, the dispersion of expectations has been larger the farther away current inflation was from the target. This relationship seems, however, to have broken down at the end of 2013 (see hollow squares in the figure) when current inflation was still significantly below the level consistent with the price stability goal but the heterogeneity of expectations was nonetheless low and similar to that observed when current inflation was around the price stability target. The upper-right and lower-left panels of the figure correlate instead the heterogeneity of expectations with proxies for the uncertainty faced by firms. Specifically, we develop two simple indexes. The first is simply the quarter-on-quarter absolute change of the yearly Italian inflation rate measured by the HICP ($\Delta\pi_t$); large sudden changes in inflation may surprise agents or capture their attention, perhaps because they receive more attention in the media, and induce them to revise their expectations; however, even under rational expectations, whether this leads to more similar or more heterogeneous expectations depends on how similar their information sets are. This leads to the second index, which measures the heterogeneity of *observed* price changes for the main items covered by the HICP; we consider the thirty-nine main three-digit groupings and compute the quarter-specific standard deviation of their percentage year-on-year changes, weighting each price change with the weight the item is assigned in the HICP. This index thus captures the potential heterogeneity of the information firms may pay attention to, possibly as price setters that monitor average price developments relevant to their business.[15] Visual inspection of the two panels suggests an empirical regularity is in place. Expectations appear to be more heterogeneous in quarters characterized by larger absolute changes in the inflation rate and more heterogeneous price developments of the items in the HICP basket. Differently from the evidence

[14]The ECB aims to maintain the year-on-year increase in the HICP for the euro area below but close to 2 percent in the medium term. Results do not change if we use a slightly lower value (e.g., 1.9 percent).

[15]Unfortunately, the information on the sector of activity collected by the survey is not sufficient to explore the possibility that firms' inflation expectations are more strongly related to price developments of items closer to the firm's relevant market.

in the upper-left panel, these relationships appear to have resisted also beyond 2013. Finally, the lower-right panel explores the empirical relationship between inflation expectations heterogeneity and the output gap, measured as the cyclical component extracted by a Hodrick-Prescott filter applied to the log of quarterly chain-linked GDP between 1995:Q1 and 2016:Q4. Over a long enough horizon and with well-anchored inflation expectations, the output gap and deviations from the price stability target are clearly related. However, the lower-right panel shows that the relationship between the dispersion of inflation expectations and the output gap is at best weak if compared with that shown in the upper-left panel with the distance from the ECB's goal; dispersion does not seem to increase the farther GDP is from its potential but only when it falls below it, and even this relationship seems to be driven entirely by the first three quarters of 2009, when GDP abruptly fell because of the global financial crisis.

A more formal assessment of these empirical regularities is reported in table 4, where we report results from the estimation of

$$\sigma_t^{\pi} = \alpha + \beta_1 \pi_t + \beta_2 \hat{\pi}_t^2 + \beta_3 \Delta\pi_t + \beta_4 (\Delta\pi_t)^2 + \beta_5 \hat{y}_{t-1} + \beta_6 \hat{y}_{t-1}^2 + \beta_7 \Sigma_t + n_t + \epsilon_t, \quad (2)$$

where the dependent variable σ_t^{π} is the cross-sectional standard deviation of inflation expectations collected in quarter t, $\hat{\pi}_t^2$ is the (squared) deviation from the ECB's price stability goal in quarter t, $\Delta\pi_t$ is the change of realized inflation between two consecutive quarters, $\hat{y}_{t-1}$ is the output gap in the previous quarter, and Σ_t is the weighted standard deviation of annual price changes of HICP items. We also include the (log of) sample size in the quarter to indirectly account for small sample and measurement effects. The choice of lags of the explanatory variables is made by selecting the combination that maximizes the model's adjusted R^2 when estimated on the full sample, 1999:Q4–2018:Q2. Each column of the table reports results obtained on a subsample ending in the fourth quarter of the year displayed in the column head (expect for the last one, ending in the second).

The first observation is that even if the model specification has been selected to maximize the share of explained variance on the entire period, estimates excluding the period since 2014 are able to

Table 4. Heterogeneity of Inflation Expectations and the Business Cycle

	Estimation Sample from Q4:1999 to Q4:									
	2009	**2010**	**2011**	**2012**	**2013**	**2014**	**2015**	**2016**	**2017**	**2018**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**	**(10)**
π_t	−0.058	−0.126*	−0.141*	−0.165**	−0.102*	0.046	0.051+	0.053+	0.046	0.038
$\hat{\pi}_t^2$	0.150**	0.178**	0.191**	0.222**	0.176**	0.060*	0.056**	0.054**	0.047*	0.040*
$\Delta\pi_t$	−0.111*	−0.055	−0.013	−0.015	−0.017	−0.032	−0.033	−0.036	−0.060	−0.067
$(\Delta\pi_t)^2$	0.227**	0.235**	0.266**	0.238**	0.254**	0.293**	0.294**	0.293**	0.255**	0.243**
$\hat{y}_{t-1}$	−2.0	−0.9	0.3	1.2	−0.6	−5.0**	−5.2**	−5.2**	−4.1*	−3.1+
$\hat{y}_{t-1}^2$	−4.7	−114.1	−180.0	−267.5*	−146.1	162.4+	173.4*	180.4*	192.7*	213.6*
Σ_t	4.9	3.3	5.9	6.9+	7.8*	10.2*	10.3*	10.0**	8.5*	6.2
Obs.	41	45	49	53	57	61	65	69	73	75
$\bar{R}^2$	0.80	0.74	0.75	0.79	0.78	0.67	0.67	0.67	0.60	0.56

Notes: Statistical significance: (**) 1 percent; (*) 5 percent; (+) 10 percent. All regressions also include the logarithm of sample size and a constant.

explain about 20 percent more of the overall variance, even accounting for the lower number of data points to be fitted by the model. Second, results generally confirm the empirical associations detected in figure 5. The correlation between contemporaneous inflation and the dispersion of expectations does not follow a discernible pattern and, from slightly significant and negative, becomes negligible when including the post-2014 period. The size of the deviation from the price stability target ($\hat{\pi}_t^2$) and the magnitude of the short-term change in inflation ($(\Delta\pi_t)^2$) are positively and strongly correlated with dispersion among firms. The correlation of dispersion with the size of the output gap ($\hat{y}_{t-1}^2$) and with the heterogeneity of price developments of HICP items (Σ_t) becomes instead positive and statistically significant only since 2014, a period over which the correlation with the size of the deviation of inflation from target becomes weaker. These empirical associations are broadly similar to those detected by Mankiw, Reis, and Wolfers (2003) for U.S. households and professional forecasters in the period between the early 1950s and early 2000s. They find that current inflation is positively associated with dispersion in all surveys, whereas the output gap and the size of the current change in inflation only play a role for dispersion among households. In a similar vein, Dovern, Fritsche, and Slacalek (2012) explore the macroeconomic determinants of disagreement among professional forecasters from major OECD countries over about twenty years between the early 1990s and 2010 and also find that disagreement increases in recessions, with the level of inflation and with the volatility of permanent inflation shocks, and declines under central bank independence. Importantly, in both cases there is no specific assessment of the correlation of dispersion with deviations from a quantitative price stability target, possibly because over such a long period of time substantially different monetary policy regimes were in place.

The above estimates suggest that certain empirical regularities linking dispersion among businesses to macroeconomic developments have broken down over the most recent period of persistently low inflation. To assess how substantial such breakdown is, from each set of estimates reported in table 4 we obtain the corresponding out-of-sample forecast for dispersion and its 95 percent confidence interval. We plot each confidence interval along with the observed time series of dispersion in the corresponding panel of figure 6. The

Figure 6. Observed and Out-of-Sample Predictions of Dispersions

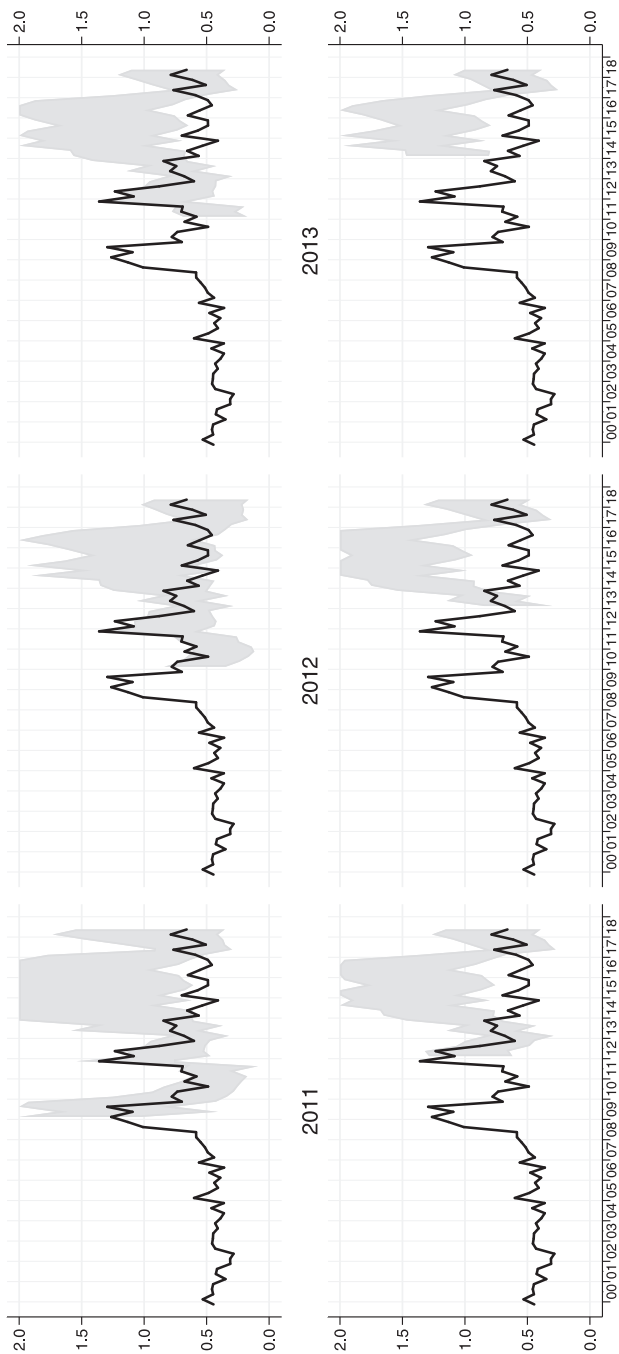

Source: Own elaborations of Bank of Italy Survey of Inflation and Growth Expectations.
Notes: Each panel reports (gray area) the 95 percent confidence interval of the out-of-sample forecast of the dispersion of one-year-ahead inflation expectations based on the estimation of equation (2) in the main text on a sample ending in the year reported on top of the panel and (black line) the actual cross-sectional dispersion.

figure clearly shows that the level of dispersion among businesses recorded between 2014 and 2016 is at odds with the developments foreseeable on the basis of deviations from the price stability target, dispersion and volatility of price changes, and the output gap: given these developments, the level of dispersion should have been much higher than that observed. In particular, figure 5 suggests that the main deviation from the pre-2014 empirical regularity is in the bivariate relationship between dispersion among businesses and the current deviation of consumer price inflation from the price stability goal. Indeed, a statistical test of the stability of the coefficient β_2 loading the deviation from target in equation (2) does reject the null hypothesis of no break and locates it in the first quarter of 2014.[16]

A more formal assessment of the role played by the macroeconomic variables considered in equation (2) is displayed in figure 7. For each right-hand-side variable x_j included in the empirical model, we compute the difference $\Delta_{jt}^S = (\beta_j^S - \beta_j^F)x_{jt}$, where β_j^S and β_j^F are, respectively, the OLS regression coefficients estimated on the subsample ending in year S and on the full sample ending in 2016. The sum of Δs over js is therefore the difference between the dispersion at time t predicted by the out-of-sample forecast and the dispersion fitted by the model estimated on the full sample. Thus, for each $t > S$, the quantity Δ_{jt}^S tells how the jth right-hand-side variable contributes to this difference. A glance at the figure shows that the major driver of the difference between actual and out-of-sample predicted dispersion, especially detected since 2014, is due to the weaker conditional correlation between the deviation of inflation from target and the level of dispersion among agents.

4. Conclusions

This paper is the first to document cross-sectional and time-series properties of consumer price inflation expectations formulated by *firms*. Virtually all existing econometric research on expectations

[16]Specifically, we perform a supremum Wald test to search for a structural break on the relevant coefficient; the null of no break is rejected with a p-value of 0.0001.

Figure 7. The Wedge between Out-of-Sample Predictions and Observed Dispersions

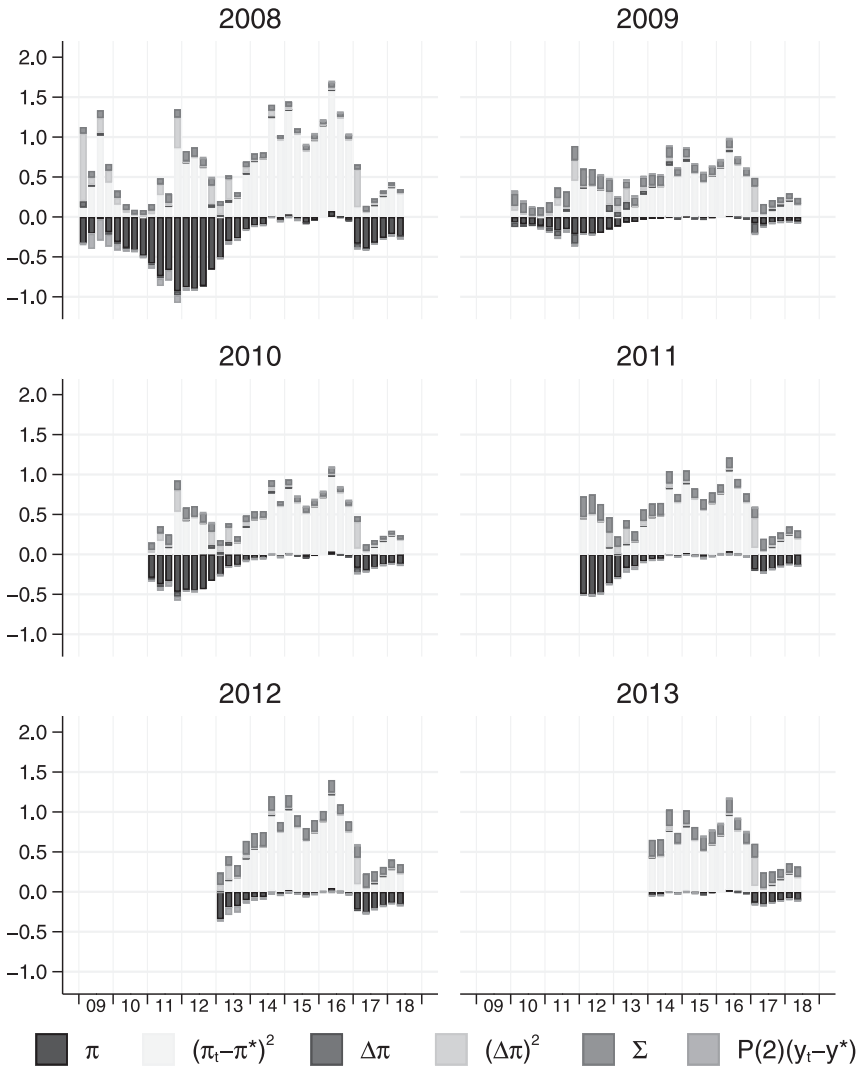

Source: Own elaborations of Bank of Italy Survey of Inflation and Growth Expectations.

Notes: Each panel is based on results of a specific out-of-sample forecast based on a sample ending in the year displayed on top of the panel (see text and figure 6 for details). π: inflation rate at time of the survey; $(\pi_t - \pi^*)^2$: (squared) deviation from inflation target ($\pi^* = 2\%$); $\Delta\pi$: absolute quarter-on-quarter change in inflation; Σ: cross-sectional standard deviation of year-on-year price change of three-digit HICP items; $P(2)$ $(y_t - y^*)$: second-order polynomial of the output gap.

focuses on consumers or professional forecasters; yet, ultimately it is the inflation expectations of price-setting firms that matter for a fuller understanding of price dynamics. The analysis is based on the quarterly Survey of Inflation and Growth Expectations of the Bank of Italy, run since 1999 and unique in that it collects point estimates of firms' consumer price inflation at several horizons along with other information on their price-setting behavior and macroeconomic expectations. We have shown that while firms tend not to be aware of the most recent available information about inflation when formulating their expectations, the delay with which it is finally taken into account is on average rather short, at most one quarter. This updating delay contributes to about half of the average cross-sectional dispersion of expectations. The remaining cross-sectional dispersion is hardly a reflection of firms' heterogeneity itself, a result in contrast with the evidence available for consumers' expectations, in which a large share of heterogeneity is explained by differences in their observable characteristics; it is, however, related in a systematic way to developments in certain economic aggregates. Specifically, our analysis has shown that the dispersion of expectations of informed agents is substantial in all phases of the business cycle but tends to be higher when inflation is farther away from the ECB's price stability goal, when its short-term swings are larger, and when price dynamics of consumption items are more diverse. These features are largely in line with those documented for U.S. consumers and professional forecasters by Mankiw, Reis, and Wolfers (2003) and for OECD professional forecasters by Dovern, Fritsche, and Slacalek (2012) and hardly replicated by standard macroeconomic models.

Importantly, we have shown that the empirical link between the dispersion of expectations and the gap between current inflation and the price stability objective has considerably weakened between 2014 and 2016, when dispersion among informed agents has shrunk even with inflation still short of the ECB's price stability goal. We have also found that, over basically the same period, firms have put a steadily increasing weight on their prior beliefs of low inflation, a reflection of the perceived lower uncertainty surrounding them. Taken together, these results provide further evidence that inflation expectations still suffer from a growing risk of de-anchoring (Buono and Formai 2018, Natoli and Sigalotti 2018).

References

Andrade, P., R. K. Crump, S. Eusepi, and E. Moench. 2016. "Fundamental Disagreement." *Journal of Monetary Economics* 83 (October): 106–28.

Armantier, O., S. Nelson, G. Topa, W. van der Klaauw, and B. Zafar. 2016. "The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment." *Review of Economics and Statistics* 98 (3): 503–23.

Ball, L., G. N. Mankiw, and R. Reis. 2005. "Monetary Policy for Inattentive Economies." *Journal of Monetary Economics* 52 (4): 703–25.

Bank of Italy. 2017. "Methods and Sources: Methodological Notes." Survey of Inflation and Growth Expectations (January 16).

Bernanke, B. S. 2007. "Inflation Expectations and Inflation Forecasting." Speech at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, Massachusetts, July 10.

Boneva, L., J. Cloyne, M. Weale, and T. Wieladek. 2016. "The Effect of Unconventional Monetary Policy on Inflation Expectations: Evidence from Firms in the United Kingdom." *International Journal of Central Banking* 12 (3, September): 161–95.

Bottone, M., and A. Rosolia. 2019. "Monetary Policy, Firms' Inflation Expectations and Prices: Causal Evidence from Firm-level Data." Working Paper No. 1218, Bank of Italy (April).

Buono, I., and S. Formai. 2018. "New Evidence on the Evolution of the Anchoring of Inflation Expectations." *Journal of Macroeconomics* 57 (September): 39–54.

Busetti, F., D. Delle Monache, A. Gerali, and A. Locarno. 2017. "Trust, but Verify. De-anchoring of Inflation Expectations under Learning and Heterogeneity." ECB Working Paper No. 1994.

Cavallo, A., G. Cruces, and R. Perez-Truglia. 2017. "Inflation Expectations, Learning and Supermarket Prices: Evidence from Survey Experiments." *American Economic Journal: Macroeconomics* 9 (3): 1–35.

Coibion, O., Y. Gorodnichenko, and R. Kamdar. 2017. "The Formation of Expectations, Inflation and the Phillips Curve." NBER Working Paper No. 23304.

Coibion, O., Y. Gorodnichenko, and S. Kumar. 2015. "How Do Firms Form Their Expectations? New Survey Evidence." NBER Working Paper No. 21092.

Dovern, J., U. Fritsche, and J. Slacalek. 2012. "Disagreement Among Forecasters in G7 Countries." *Review of Economics and Statistics* 94 (4): 1081–96.

Ehrmann, M., D. Pfajfar, and E. Santoro. 2017. "Consumers' Attitudes and Their Inflation Expectations." *International Journal of Central Banking* 13 (1, February): 225–59.

Engelberg, J., C. F. Manski, and J. Williams. 2009. "Comparing the Point Predictions and Subjective Probability Distributions of Professional Forecasters." *Journal of Business and Economic Statistics* 27 (1): 30–41.

Fabiani, S., and S. Santoro. 2012. "Rationality of Italian Firms' Inflation Expectations: Any Change During the Crisis?" Mimeo, Bank of Italy.

Gaspar, V., F. Smets, and D. Vestin. 2006. "Adaptive Learning, Persistence, and Optimal Monetary Policy." *Journal of the European Economic Association* 4 (2–3): 376–85.

Keane, M. P., and D. E. Runkle. 1990. "Testing the Rationality of Price Forecasts: New Evidence from Panel Data." *American Economic Review* 80 (4): 714–35.

Lahiri, K., and X. Sheng. 2008. "Evolution of Forecast Disagreement in a Bayesian Learning Model." *Journal of Econometrics* 144 (2): 325–40.

Mankiw, G. N., and R. Reis. 2002. "Sticky Information vs Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117 (4): 1295–1328.

Mankiw, G. N., R. Reis, and J. Wolfers. 2003. "Disagreement about Inflation Expectations." NBER Working Paper No. 9796.

Manski, C. F. 2018. "Survey Measurement of Probabilistic Macroeconomic Expectations: Progress and Promise." In *NBER Macroeconomics Annual 2017*, Vol. 32, ed. M. Eichenbaum and J. A. Parker, 411–71 (chapter 5). Cambridge, MA: University of Chicago Press.

Natoli, F., and L. Sigalotti. 2018. "Tail Co-movement in Inflation Expectations as an Indicator of Anchoring." *International Journal of Central Banking* 14 (1, January): 35–71.

Patton, A. J., and A. Timmermann. 2010. "Why Do Forecasters Disagree? Lessons from the Term Structure of Cross-sectional Dispersion." *Journal of Monetary Economics* 57 (7): 803–20.

Ropele, T. 2017. "Inflation Expectations and Price Setting Behavior: Evidence from Business Survey Data." Mimeo, Bank of Italy.

Souleles, N. 2004. "Expectations, Heterogeneous Forecast Errors, and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys." *Journal of Money, Credit and Banking* 36 (February): 39–72.

Visco, I. 1987. *Price Expectations in Rising Inflation Contributions to Economic Analysis*. North Holland.