



BANCA D'ITALIA
EUROSISTEMA

Questioni di Economia e Finanza

(Occasional Papers)

What does the heterogeneity of the inflation expectations of Italian firms tell us?

by Laura Bartiloro, Marco Bottone and Alfonso Rosolia

December 2017

Number

414



BANCA D'ITALIA
EUROSISTEMA

Questioni di Economia e Finanza

(Occasional Papers)

What does the heterogeneity of the inflation expectations
of Italian firms tell us?

by Laura Bartiloro, Marco Bottone and Alfonso Rosolia

Number 414 – December 2017

The series Occasional Papers presents studies and documents on issues pertaining to the institutional tasks of the Bank of Italy and the Eurosystem. The Occasional Papers appear alongside the Working Papers series which are specifically aimed at providing original contributions to economic research.

The Occasional Papers include studies conducted within the Bank of Italy, sometimes in cooperation with the Eurosystem or other institutions. The views expressed in the studies are those of the authors and do not involve the responsibility of the institutions to which they belong.

The series is available online at www.bancaditalia.it.

ISSN 1972-6627 (print)

ISSN 1972-6643 (online)

Printed by the Printing and Publishing Division of the Bank of Italy

WHAT DOES THE HETEROGENEITY OF THE INFLATION EXPECTATIONS OF ITALIAN FIRMS TELL US?

by Laura Bartiloro*, Marco Bottone* and Alfonso Rosolia*

Abstract

Quite a lot. 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 there has no longer been an empirical connection between the dispersion of expectations and the distance from the ECB price stability. These two facts suggest an increased risk of inflation expectations being de-anchored.

JEL Classification: D22, D8, E31.

Keywords: inflation expectations, learning, firms.

Contents

1. Introduction.....	5
2. The Bank of Italy Survey of inflation and growth expectations.....	8
3. The sources of expectations' heterogeneity.....	10
3.1 Information (or lack thereof)	11
3.2 Microeconomic factors	16
3.3 Macroeconomic factors	17
4. Conclusions.....	22
References	25
Tables and figures.....	27

* Bank of Italy - DG Economics, Statistics and Research - Statistical Analysis Directorate. The views expressed are not necessarily shared by the Bank of Italy.

1 Introduction

Close monitoring of inflation expectations is a crucial ingredient for the conduct of monetary policy. The increasing availability of 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 goals¹.

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

¹For 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; Buseti, Delle Monache, Gerali and Locarno (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.

or professional forecasters². When firms are the relevant decision unit, the information available either relates to their own prices or costs or, when about broader inflation measures, it is 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.

Second, we explore the determinants of disagreement among firms 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 policy makers. To our knowledge, only Mankiw, Reis and Wolfers (2003) have conducted a similar empirical analysis for the US, based however on consumers' and professional forecasters' expectations. They show that the dispersion of inflation expectations 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. However, this result relies on expectations collected from *consumers* and from *professional forecasters*, not from *firms*. Our paper thus complements their in providing novel empirical evidence on the determinants and cyclical properties of the dispersion of *firms'* inflation expectations.

Third, 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

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

estimate an upper bound to the delay with which relevant information is incorporated in firms' inflation expectations.

In this paper we use data drawn from the Bank of Italy' Survey of Inflation and Growth Expectations (SIGE). Several studies have used the SIGE or its predecessor to address some of the issues outlined above³. 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 that for periods of economic stability. More recently, Fabiani and Santoro (2012) also concludes against the rationality of firms' inflation expectation. Ropele (2017) is one of the few existing studies that 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 (2017) explore the short-term degree of sensitivity of firms' expectations to monetary policy shocks.

In the next section we briefly describe the SIGE, focussing 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.

³Between 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 launched in collaboration with *il Sole 24 Ore*, the main Italian financial newspaper, the SIGE.

2 The Bank of Italy Survey of inflation and growth expectations

The Bank of Italy started the Survey of businesses' inflation and growth expectations (SIGE) in the fourth quarter of 1999. The SIGE is run quarterly on a sample of currently about 1000 manufacturing and service firms with at least 50 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 quarter⁴. 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 changes⁵; the questionnaire also collects sentiment information on aggregate cyclical developments as well as on own business real and financial conditions⁶.

At the international level, existing surveys focus generally either on consumers' (e.g. Michigan Survey of Consumer Attitudes) or on professional forecasters' (e.g. the Survey of professional forecasters of the ECB and the FED) 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

⁴Specifically, the 1st quarter survey is run in early March, the 2nd quarter survey in early June, the 3rd quarter one in early September and the 4th quarter one in early December.

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

⁶The survey data can be accessed through the Bank of Italy Remote processing system, BIRD; details at www.bancaditalia.it/statistiche/basi-dati/bird/index.html.

collected they typically refer to *own* price or unit costs (e.g. the Atlanta Fed's Business Inflation expectations since 2011) and industry-wide price developments (e.g. the Business surveys run by the Confederation of British Industries; Boneva, Cloyne, Weale and Wieladek (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. An exception is represented by Coibion, Gorodnichenko and Kumar (2015) who, reckoning 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, 6 months, 1 year, 2 years and average between 3 and 5 years). 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 2/3 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:

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 1/3 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] ?

The Survey thus 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-year-ahead Italian consumer price inflation expectations of informed agents. 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 at a level below the ECB’s goal; these developments have partly reflected those of current inflation. 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 was in the 0.4-0.6

percentage points range, around a median inflation expectation gradually converging towards 2 percent. Since 2008, the gap has widened, remaining almost continuously above 0.8 percentage points until mid 2015, when it went back to the initial values but around a much lower median inflation rate; in periods characterised by particularly high median inflation expectations also the 20th-80th percentile gap has been much larger, above 1.5 percentage points.

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 focussing on three factors: information, businesses' observable characteristics, and business cycle developments.

3.1 Information (or lack thereof)

The SIGE allows to assess, though only since 2012:3, the role played by information and firms' awareness in shaping businesses' expectations and their dispersion.

Panel A of figure (2) plots the median one-year-ahead inflation expectation of informed and non informed agents, along with the information on the most recent inflation gauge provided to informed agents over the period for which both questions are asked. The median expected inflation of informed agents closely tracks the most recent official inflation rate provided with the questionnaire, and so does the median expected inflation of non informed agents, although their median expectation fell somewhat less easily below 1 percent.

As concerns their predictive power, the median expectations of informed and non informed respondents both fail to match realized one-year-ahead inflation. Panel B of figure (2) plots

the median forecast error (defined as the difference between median expected and realised inflation) against the quarter in which expectations were collected. It shows that both groups 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 two groups' forecast performances moved in a largely similar way, non informed respondents appear to have taken longer to revise their expectations to the low levels of current inflation gauges, their prediction error being systematically larger than that of informed ones until early 2016; afterwards, their forecast 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 median expectations on current and past inflation rates for the two groups of agents. Columns (1) to (3) show that median 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) to (6) show that non informed respondents' median expectations 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; however, for both groups this unexplained component does not turn out to be in a statistically significant correlation with future inflation realizations.

Looking at this 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} be the prior one-year ahead inflation expectation of firm i at time t and π_t^* the signal

provided; let also σ_{it} and s_t be the variances of the prior and of the signal and assume both are normally distributed. Then i 's posterior 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 et al. (2015), Armantier, Nelson, Topa, van der Klaauw and Zafar (2016), Cavallo, Cruces and Perez-Truglia (2017)). 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) has several interesting implications. First, 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 up to unity, that the constant should be nil and that older signals should have no role in shaping the posterior. Results are presented in table (2), where in columns (1) to (3) we have used the cross-sectional mean expectations in both groups and in columns (4) to (6) the cross-sectional median expectations. In columns (1) and (4) we regress the (mean or median) posterior on the (mean or median)

prior and on the signal, that is the most recent inflation rate. Both factors are given on average essentially the same weight and the data cannot reject that the coefficients sum up to unity and that the constant is zero⁷. In columns (2) and (5) we augment the two regressions with past information, that is with the previous quarter inflation reading. While the coefficients on the prior and the current signal essentially do not change and remain highly statistically significant, older information does not play any role in shaping the posterior. Finally, in columns (3) and (6) we drop from the specification the current inflation reading and show that all its explanatory power is absorbed by the prior expectation rather than by lagged inflation, suggesting that the prior incorporates all the available information. In sum, these results suggest that firms incorporate recent inflation readings in their expectations with a short delay of at most one quarter. While they show that on average the prior and the signal are given a similar weight in forming the posterior, they are not informative, however, 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 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, $\frac{s_t}{s_t + \sigma_t}$. These estimates are displayed in figure (3) along with a 3-terms and 4-terms moving average to smooth out high frequency variability⁸. The figure shows that the information provided to respondents does

⁷The prediction that coefficients sum up to unity refer 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.

⁸To dispel concerns that the ratio of the two cross-sectional standard deviations might reflect the different

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 (2). The figure also suggests that the relative contributions of the new information and of the prior may change over time. Specifically, the weight put on the prior has declined until late 2015 and, correspondingly, the importance of the signal (a persistently declining inflation) has gradually increased. However, this process seems to have reversed over the past two years, when prior expectations have been at their lowest while current inflation was recording somewhat of a recovery (fig. 2). In other words, the uncertainty surrounding firms' priors has first increased (relative to the precision of the signal) when current readings were showing a persistent inflation decline and has then decreased but around a lower expected inflation, a pattern consistent with the gradual entrenchment of low inflation expectations.

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.

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

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.

3.2 Microeconomic factors

The residual dispersion after relevant information has been provided can hardly be traced to firms' characteristics. Although information available for the whole period on firms' characteristics is scant, a regression of expectations on 9 size-class dummies, 4 area dummies, 2 sector dummies and the log of employees for the period 2000:1-2016:4 accounts for less than 1 percent of the total variance. This is not surprising since the regression does not account for the time variation of *average* expectations; however, while adding time dummies to the regression explains about 60 percent of the variance, systematic differences between groups of similar firms still explain only about 1 percent of the total variance; a more flexible specification in which dummies for firms' characteristics are interacted with time dummies so that differences in average expectations between groups of firms can change over time, still explains about 60 percent of total variance; firms' characteristics alone now account between one fifth and one fourth of total variance. Besides, the statistical significance of these estimated differences is often low. In figure (4), each panel displays both the 95 percent confidence interval on the corresponding coefficient estimate obtained under the assumption that it is constant over time and the point estimates and 95 percent confidence intervals from the more flexible model that interacts each characteristic with time dummies. Systematic differences along the size dimension are never statistically significant at customary confidence levels while southern firms and those in the service sector tend to expect, respectively, slightly higher and slightly lower inflation. An assessment of the statistical significance of the differences across groups estimated from the more

flexible model must be more nuanced since they display large swings and changes of sign over time; even if the associated standard errors are large, a formal test of the null that a specific difference is constant over time rejects it in most cases.

These results stand in stark contrast to those usually obtained from studies of the sources of heterogeneity of households' expectations of macro variables. For example, Souleles (2001) finds that in the Michigan Survey of Consumers Attitudes households' inflation forecast errors are correlated their socio-demographic characteristics; Ehrmann et al. (2017) confirm these results and also show that expectations are shaped by the household's financial situation and purchasing habits. Importantly, the fact that differences in expectations across groups of firms are largely unsystematic, displaying even substantial changes from one period to the next, appears to be more consistent with the hypothesis that firms formulate their expectations on the basis of different and changing information sets than with the one that traces their different expectations to different (possibly non rational) methods of formulating expectations from a common information set.

3.3 Macroeconomic factors

In figure (5) we describe the empirical relationship between the cross-sectional dispersion of inflation expectations, as measured by their standard deviation, and specific aspects of actual price dynamics and macroeconomic developments; broken lines visually summarize bivariate relationships. Specifically, panel A 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 in the Euro area here taken to be 2 percent for simplicity⁹. 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 (in red) 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. Panels B and C 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 39 main 3-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

⁹The ECB aim to maintain the year-on-year increase in the Harmonised Index of Consumer Prices (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).

as price setters that monitor average price developments relevant to their business¹⁰. Visual inspection of the two panels suggests an empirical regularity is in place. Expectations appear to be more heterogeneous in quarters characterised by larger absolute changes in the inflation rate and more heterogeneous price developments of the items in the HICP basket. Differently from the evidence in panel A, these relationships appear to have resisted also beyond 2013. Finally, panel D 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 1995:1 and 2016:4. 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, panel D shows that the relationship between the dispersion of inflation expectations and the output gap is at best weak if compared with that shown in panel A with the distance from the ECB’s goal; disagreement 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 (3) 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 expect-

¹⁰Unfortunately, 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.

tations collected in quarter t , $\hat{\pi}_t^2$ is the (squared) deviation from the ECB price stability goal in quarter t , $\Delta\pi_t$ is the change of realised inflation between two consecutive quarters, \hat{y}_{t-1} is the output gap in the previous quarter, 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 selecting the combination that maximizes the model's R^2 when estimated on the full sample 1999:4-2016:4. Each column of the table reports results obtained on a subsample ending in the 4th quarter of the year displayed in the column head.

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 2014:1-2016:4 are able to explain 5 to 10 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). Contemporaneous inflation is only weakly correlated with the dispersion of expectations while the size of the deviation from the price stability target and the magnitude of the short-term change in inflation are positively and strongly correlated with disagreement among firms. The positive correlation of disagreement with the size of the output gap and with the heterogeneity of price developments of HICP items becomes instead 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 between disagreement about future inflation and macroeconomic variables are broadly similar to those detected for the US in the period between the early 1950s and early 2000s by Mankiw et al.

(2003), although the latter focus on households and professional forecasters' expectations rather than on businesses'. They find that current inflation is positively associated to disagreement in all surveys whereas the output gap and the size of the current change in inflation only play a role for disagreement among households. Importantly, they do not specifically assess the correlation 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 disagreement 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 estimate reported in table (3) we obtain the corresponding out-of-sample forecast for disagreement and its 95 percent confidence interval. We plot each confidence interval along with the observed time series of disagreement in the corresponding panel of figure (6). The figure clearly shows that the level of disagreement among businesses recorded since 2014 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 disagreement should have been 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 disagreement 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 2014¹¹.

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 Δ s over j s is therefore the difference between the disagreement at time t predicted by the out-of-sample forecast and the disagreement fitted by the model estimated on the full sample. Thus, for each $t > S$ the quantity Δ_{jt}^S tells how the j th 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 disagreement, especially detected since 2014, is due to the weaker conditional correlation between the deviation of inflation from target and the level of disagreement 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 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

¹¹Specifically, 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.

of Italy, run since 1999 and unique in its collecting 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 consider the most recent available information when formulating their inflation expectations, the delay with which it is finally taken into account is on average rather short, at most one quarter. This updating delay contributes for about a half to 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 US consumers and professional forecasters by Mankiw et al. (2003) 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 since 2014, when disagreement 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. Consistently with evidence based on other methods and sources (e.g. Natoli and Sigalotti (2017)), we read these results as suggestive of a growing risk of de-anchoring of expectations.

Finally, the analysis has shown that combining data on informed and non informed respondents' expectations yield useful indicators to assess the developments of inflation expectations. In particular, it may help to reveal their degree of anchoring, through a high frequency monitoring of both the level of the prior and the weight that the informed respondents put on the prior as opposed to current information.

References

- Armantier, Olivier, Scott Nelson, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar**, “The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment,” *Review of Economic and Statistics*, 2016, 98 (3), 503–523.
- Ball, Laurence, Gregory N. Mankiw, and Ricardo Reis**, “Monetary Policy for Inattentive Economies,” *Journal of Monetary Economics*, 2005, 52 (4), 703 – 725.
- Bank of Italy**, *Methods and Sources: Methodological Notes. Survey of Inflation and Growth Expectations* January 2017.
- Bernanke, Ben S.**, “Inflation Expectations and Inflation Forecasting,” 2007. Speech at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, Massachusetts.
- Boneva, Lena, James Cloyne, Martin Weale, and Tomasz Wieladek**, “The Effect of Unconventional Monetary Policy on Inflation Expectations: Evidence from Firms in the United Kingdom,” *International Journal of Central Banking*, September 2016.
- Bottone, Marco and Alfonso Rosolia**, “The response of firms’ inflation expectations to monetary policy shocks: a regression discontinuity approach,” 2017. Bank of Italy, mimeo.
- Busetti, Fabio, Davide Delle Monache, Andrea Gerali, and Alberto Locarno**, “Trust, but verify. De-anchoring of inflation expectations under learning and heterogeneity,” 2017. ECB Working Paper Series, n. 1994.
- Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia**, “Inflation Expectations, Learning and Supermarket Prices - Evidence from Survey Experiments,” *American Economic Journal: Macroeconomics*, 2017, 9 (3).
- Coibion, Olivier, Yuriy Gorodnichenko, and Rupal Kamdar**, “The Formation of Expectations, Inflation and the Phillips Curve,” 2017. NBER, WP n. 23304.
- , – , and **Saten Kumar**, “How do Firms Form Their Expectations? New Survey Evidence,” 2015. NBER, WP n. 21092.
- Ehrmann, Michael, Damjan Pfajfar, and Emiliano Santoro**, “Consumers’ Attitudes and Their Inflation Expectations,” *International Journal of Central Banking*, June 2017.
- Fabiani, Silvia and Sergio Santoro**, “Rationality of Italian Firms’ Inflation Expectations: Any Change During the Crisis?,” 2012. mimeo, Bank of Italy.
- Gaspar, Vitor, Frank Smets, and David Vestin**, “Adaptive Learning, Persistence, and Optimal Monetary Policy,” *Journal of the European Economic Association*, 2006, 4 (2-3), 376–385.
- Mankiw, Gregory N and Ricardo Reis**, “Sticky Information vs Sticky Prices: a Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, 2002, 117 (4).
- , – , and **Justin Wolfers**, “Disagreement about Inflation Expectations,” 2003. NBER Working Papers 9796.
- Natoli, Filippo and Laura Sigalotti**, “An Indicator of Inflation Expectations Anchoring,” 2017. Bank of Italy Discussion Papers, n. 1103.

- Ropele, Tiziano**, “Inflation expectations and price setting behavior: evidence from business survey data,” 2017. mimeo, Bank of Italy.
- Souleles, Nicholas**, “Expectations, Heterogeneous Forecast Errors, And Consumption: Micro Evidence Form The Michigan Consumer Sentiment Surveys,” *Journal of Money, Credit and Banking*, 2001, 36.
- Visco, Ignazio**, *Price Expectations in Rising Inflation Contributions to Economic Analysis*, North Holland, 1987.

Table 1: Median inflation expectations of informed and non informed respondents.

	(1)	(2)	(3)	(4)	(5)	(6)
	Informed			Non informed		
π_t	0.794 (0.042)	0.706 (0.073)	0.716 (0.081)	0.428 (0.050)	0.202 (0.103)	0.297 (0.092)
π_{t-1}		0.121 (0.065)	0.100 (0.089)		0.365 (0.091)	0.168 (0.101)
π_{t-4}	0.047 (0.032)		0.015 (0.042)	0.194 (0.038)		0.141 (0.048)
Constant	0.359 (0.053)	0.383 (0.046)	0.374 (0.054)	0.753 (0.062)	0.867 (0.064)	0.779 (0.061)
\bar{R}^2	0.968	0.970	0.968	0.919	0.894	0.927

Dependent variable is median one year ahead expected inflation.
Standard errors in parentheses. Sample frame: 2012:3-2017:2.

Table 2: Bayesian learning.

	(1)	(2)	(3)	(4)	(5)	(6)
	Means			Medians		
Prior (β_1)	0.479**	0.436**	0.924**	0.343**	0.354*	0.932*
Signal (β_2)	0.516**	0.507**		0.632**	0.634**	
Lagged signal		0.031	0.192		-0.008	0.187
Constant	0.010	0.049	-0.368	0.086	0.077	-0.368
(p-value)	0.919	0.695	0.187	0.422	0.593	0.242
H0 $\beta_1 + \beta_2 = 1$ (p-value)	0.915	0.621		0.657	0.935	
\bar{R}^2	0.984	0.983	0.906	0.977	0.976	0.869
Observations	20	20	20	20	20	20

Statistical significance: (**) 1%; (*) 5%; (+) 10%.


Note: Sample frame: 2012:3-2017:2. Dependent variable is the cross-sectional mean (cols. 1, 2, 3) or median (cols. 4,5, 6) inflation expectation of informed respondents; Prior is the cross-sectional mean (cols. 1, 2, 3) or median (cols. 4,5, 6) inflation expectation of non informed respondents; Signal is the current (at time of the survey) inflation reading provided to informed agents; Lagged signal is the previous quarter inflation reading.

Table 3: Heterogeneity of inflation expectations and the business cycle.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Estimation sample from 4:1999 to 4:								
	2008	2009	2010	2011	2012	2013	2014	2015	2016
π_t	-0.275+	-0.040	-0.102+	-0.119+	-0.143*	-0.091*	0.050	0.053+	0.054*
$\hat{\pi}_t^2$	0.303**	0.135**	0.157**	0.174**	0.206**	0.170**	0.064**	0.062**	0.061**
$\Delta\pi_t$	-0.041	-0.116*	-0.060	-0.013	-0.015	-0.018	-0.034	-0.034	-0.035
$\Delta\pi_t^2$	0.360**	0.229**	0.239**	0.275**	0.247**	0.258**	0.291**	0.292**	0.291**
\hat{y}_{t-1}	-1.1	-2.5	-1.6	-0.5	0.2	-1.2	-5.2**	-5.4**	-5.4**
\hat{y}_{t-1}^2	-296.2	47.4	-44.9	-126.3	-220.8+	-120.4	173.9*	179.7*	182.6*
Σ_t	2.6	5.0	3.3	6.0	7.4+	7.9*	10.0*	10.1*	9.7*
Obs.	37	41	45	49	53	57	61	65	69
\bar{R}^2	0.59	0.80	0.74	0.75	0.78	0.77	0.68	0.69	0.68

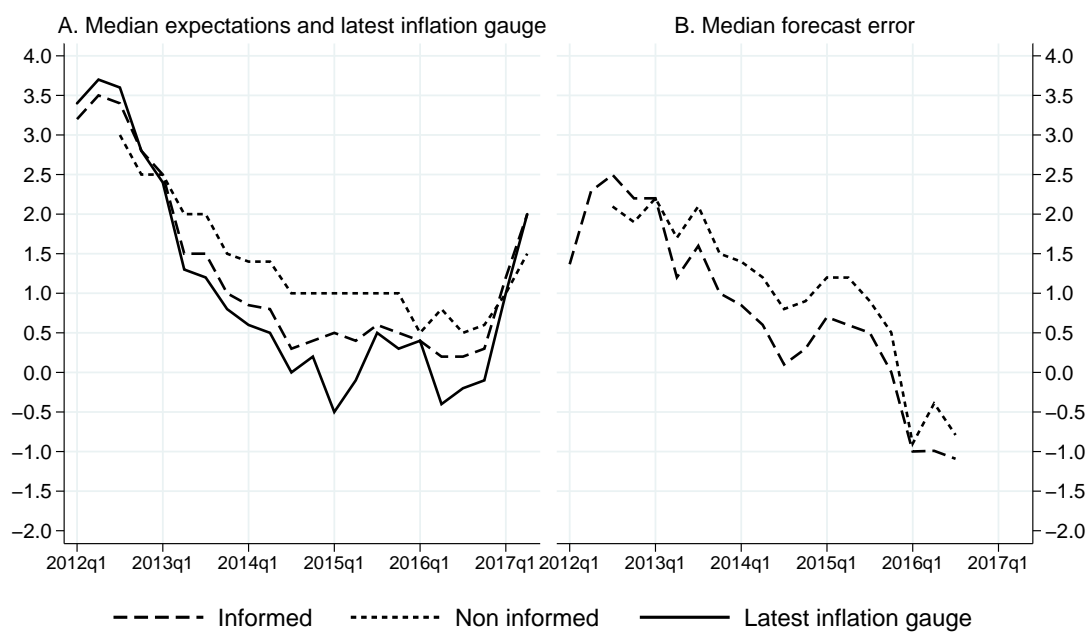
Statistical significance: (**) 1%; (*) 5%; (+) 10%. All regressions also include the logarithm of sample size and a constant.

Figure 1: Businesses' inflation expectations and disagreement.



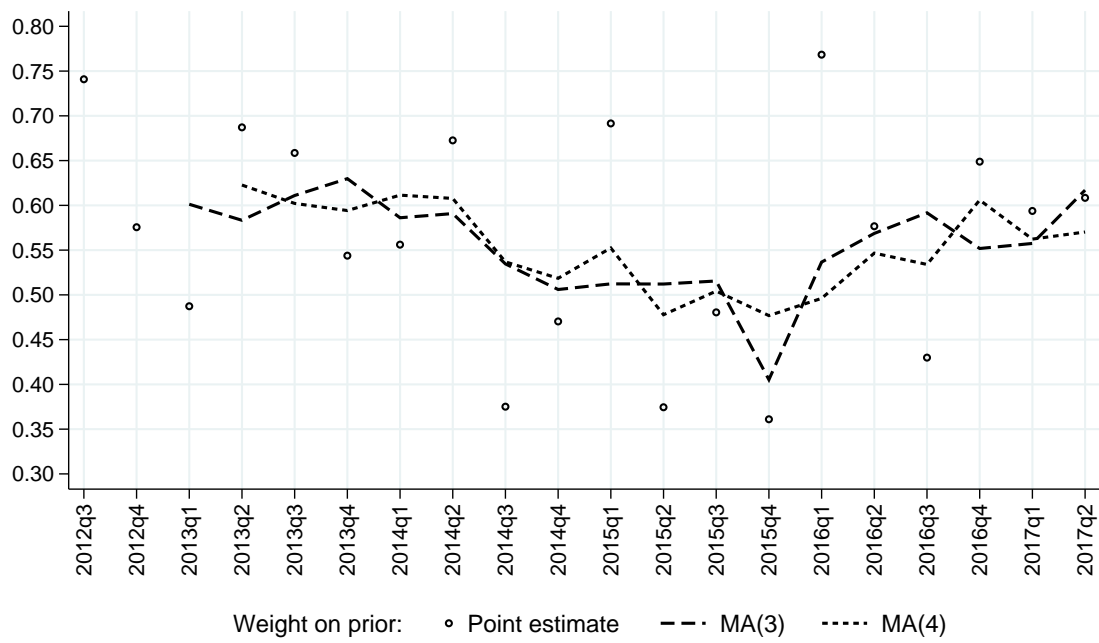
Source: Own elaborations of Bank of Italy Survey of Inflation and Growth Expectations.
 Note: Sample only includes informed agents.

Figure 2: Informed and non informed respondents.



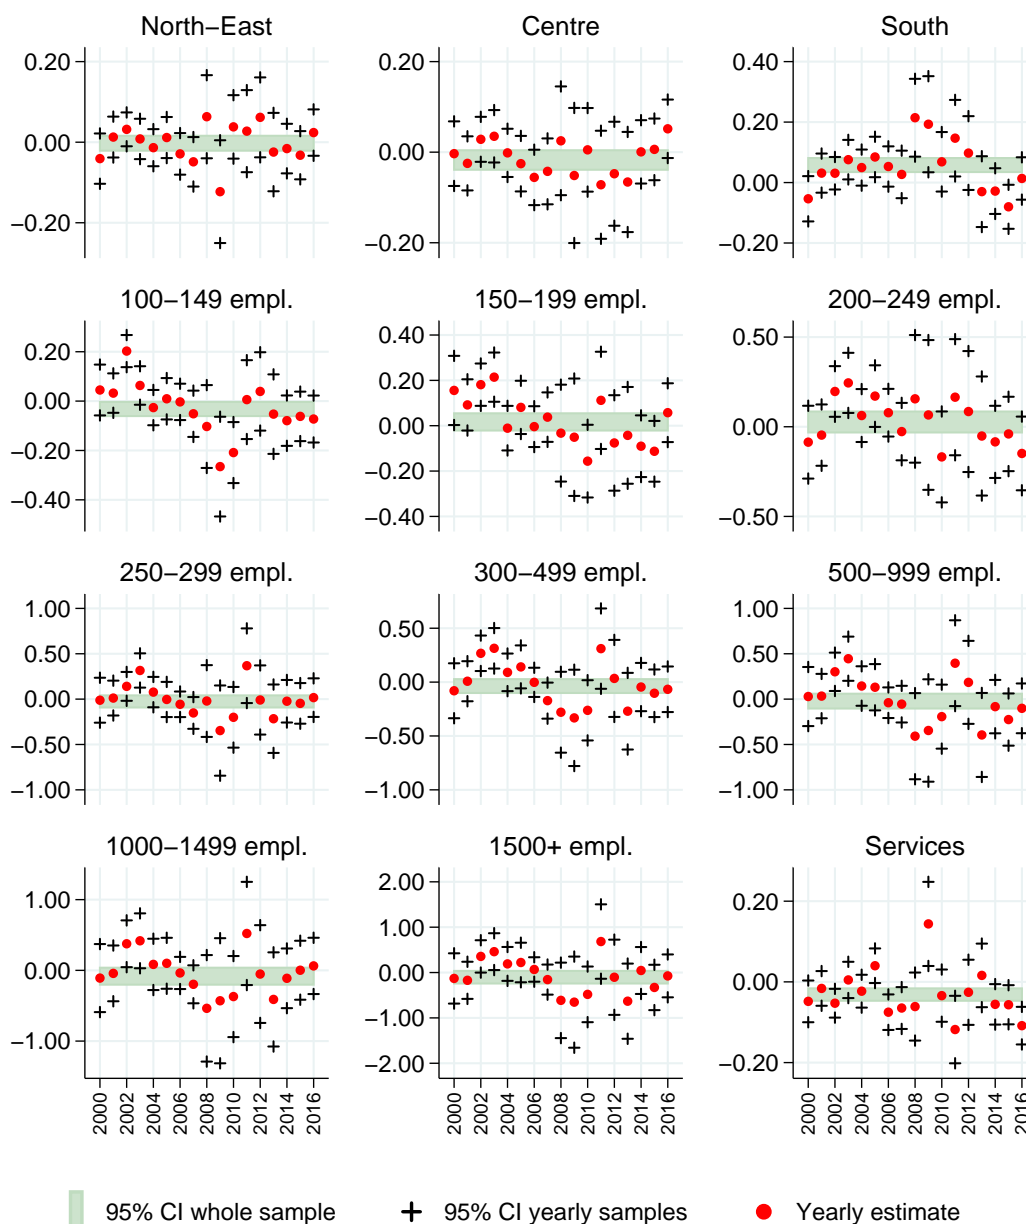
Source: Own elaborations of Bank of Italy Survey of Inflation and Growth Expectations.

Figure 3: Weights on prior expectations.




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

Figure 4: Businesses' characteristics and inflation expectations.



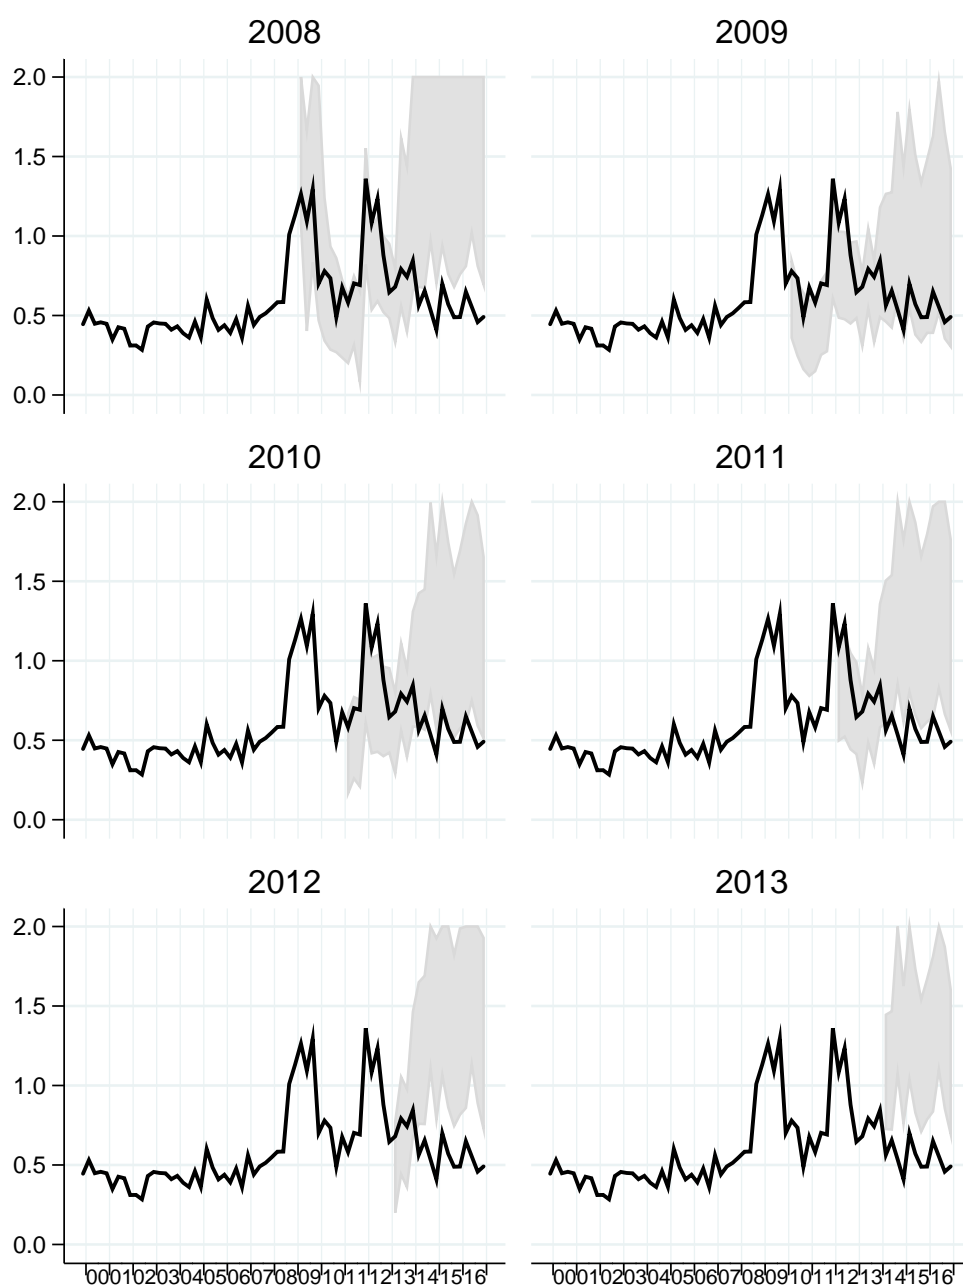
Source: Own elaborations of Bank of Italy Survey of Inflation and Growth Expectations.
 Note: Each panel reports, in gray, the 95% confidence interval of the coefficient estimated on the dummy for the category displayed in the panel head from a regression of inflation expectations of firm characteristics on the entire sample. Red dots are coefficient estimates for the same dummy in cross-sectional regressions based on different yearly subsamples; plus markers are the corresponding bounds of the 95% confidence interval.

Figure 5: Disagreement and macroeconomic developments



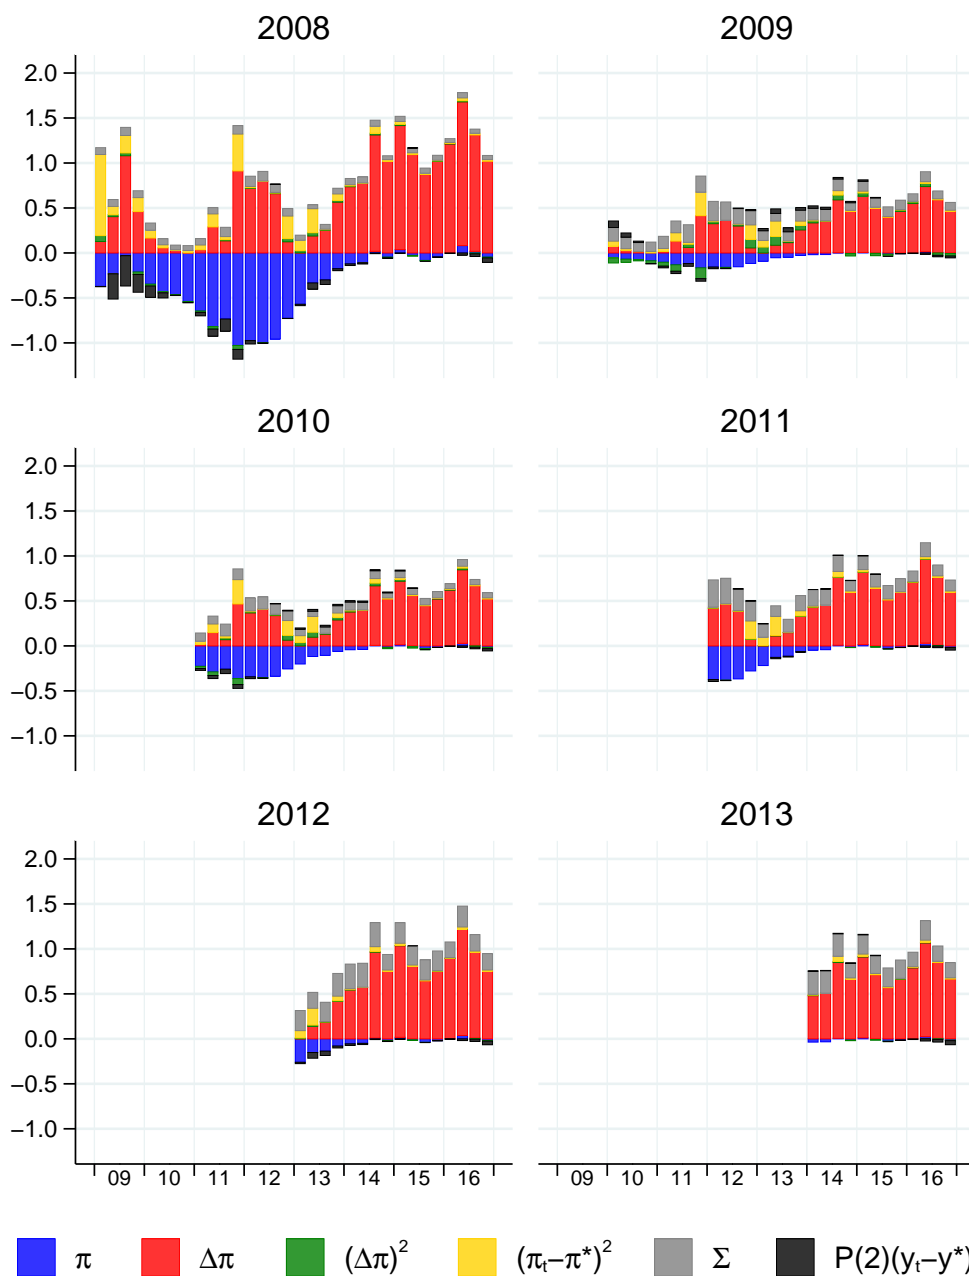
Source: Own elaborations of Bank of Italy Survey of Inflation and Growth Expectations.
 Note: the vertical axis reports the cross-sectional standard deviation of firms' inflation expectation. The horizontal axis (1) the year-on-year HICP inflation rate, (2) the quarter-on-quarter change in inflation rate, (3) the cross-sectional standard deviation of year-on-year percentage changes of the price indexes o 3-digit HICP subitems, (4) the output gap, obtained with a standard Hodrick-Prescott filter applied to chain-linked quarterly log GDP between 1995:1 and 2016:4.

Figure 6: Observed and out-of-sample predictions of disagreement.



Source: Own elaborations of Bank of Italy Survey of Inflation and Growth Expectations.
Note: Each panel reports (gray area) the 95% confidence interval of the out-of-sample forecast of expected inflation dispersion 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 7: The wedge between out-of-sample predictions of and observed disagreement.



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

Note: 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 fig. (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 q-o-q change in inflation; Σ : cross-sectional standard deviation of y-o-y price change of 3-digit HICP items; $P(2)(y_t - y^*)$: second order polynomial of the output gap.