

WP 2116 – September 2021

Information frictions in inflation expectations among five types of economic agents

Camille Cornand, Paul Hubert

Abstract:

We compare disagreement in expectations and the frequency of forecast revisions among five categories of agents: households, firms, professional forecasters, policymakers and participants to laboratory experiments. We provide evidence of disagreement among all categories of agents. There is however a strong heterogeneity across categories: while policymakers and professional forecasters exhibit low disagreement, firms and households show strong disagreement. This translates into a heterogeneous frequency of forecast revision across categories of agents, with policymakers revising more frequently their forecasts than firms and professional forecasters. Households last revise less frequently. We are also able to explore the external validity of experimental expectations.

Keywords:

inflation expectations, information frictions, disagreement, forecast revisions, experimental forecasts, survey forecasts, central bank forecasts

JEL codes:

E3, E5, E7

Information frictions in inflation expectations among five types of economic agents*

Camille Cornand*

Paul Hubert*

27 August 2021

We compare disagreement in expectations and the frequency of forecast revisions among five categories of agents: households, firms, professional forecasters, policymakers and participants to laboratory experiments. We provide evidence of disagreement among all categories of agents. There is however a strong heterogeneity across categories: while policymakers and professional forecasters exhibit low disagreement, firms and households show strong disagreement. This translates into a heterogeneous frequency of forecast revision across categories of agents, with policymakers revising more frequently their forecasts than firms and professional forecasters. Households last revise less frequently. We are also able to explore the external validity of experimental expectations.

Keywords: inflation expectations, information frictions, disagreement, forecast revisions, experimental forecasts, survey forecasts, central bank forecasts.

JEL classification: E3, E5, E7.

* We would like to thank Erwan Gautier, Cars Hommes and Johannes Wohlfart for helpful comments. We are also grateful to Cars Hommes, Domenico Massaro, Luba Petersen, Damjan Pfajfar, Matthias Weber and Blaz Žakelj for sharing their experimental data. This research was performed within the framework of the LABEX CORTEX (ANR-11-LABX-0042) of Université de Lyon, within the program “Investissements d’Avenir” (ANR-11-IDEX-007) operated by the French National Research Agency (ANR). This paper contains the views of the authors and not necessarily those of Banque de France.

* Univ Lyon, CNRS, GATE. Address: L-SE UMR 5824, F-69130 Ecully, France. Email: cornand@gate.cnrs.fr.

* Banque de France & Sciences Po – OFCE. Address: 31 rue Croix des Petits Champs, 75001 Paris, France. Email: paul.hubert@sciencespo.fr.

1. Introduction

Information frictions play a key role in various theories of expectations formation in macroeconomics. Models of sticky information (Mankiw and Reis, 2002), dispersed and heterogeneous information (Angeletos and Lian, 2019), and rational inattention (Sims, 2003; Mackowiak and Wiederholt, 2009) exhibit disagreement in expectations among individuals. This disagreement is found to be key to macroeconomic dynamics (Mankiw et al., 2004). It may reflect heterogeneity in the rate at which agents update their information sets or differences in these information sets.¹ Because the cost of collecting and processing information may be different for various categories of agents, the strength of information frictions within and across various categories may vary dramatically. As macroeconomic dynamics strongly depend on expectation processes, understanding how the degree of information frictions varies among economic agents is of utmost importance.

While the recent empirical literature focuses on information frictions among households and firms (see e.g. Coibion and Gorodnichenko, 2018, or Link et al., 2021), in this paper, we compare disagreement in inflation expectations and the frequency of forecast revisions among five categories of agents: households, firms, professional forecasters, policymakers, and participants to laboratory experiments. Disagreement is at the core of models with information frictions, and these models are compatible with agents revising their forecast infrequently. Predictions from theoretical contributions could in principle be applied to any category of economic agents who form expectations. However, these five categories of agents exhibit different individual characteristics, are subject to different objectives and incentives and may collect and process information differently. These differences may impact their degree of information frictions. By considering these five categories of agents, our paper takes a broader view on information frictions and escapes from the magnifying glass effect associated with the opposition between firms and households. We emphasize the fact that differences in information frictions observed between households and firms should be considered relative to those with policymakers, professional forecasters and participants to experiments.

Our methodology follows in part Andrade and Le Bihan (2013) who quantify the disagreement among professional forecasters and the frequency at which they revise their forecasts. In contrast, we compare the degree of information frictions in inflation expectations *within* and *across* various categories of economic agents in the vein of Coibion and Gorodnichenko (2015). Another novelty relates to the inclusion of experimental data. We proceed in two steps. First, we document disagreement in expectations within each category of agents and across these categories and investigate how disagreement is affected by inflation shocks. To tackle the latter issue, we estimate the conditional variance of inflation from a GARCH model. Second, for each category of expectations, we measure the frequency with which economic agents revise their forecasts and document whether there is some heterogeneity in this frequency within and across the five categories of economic agents. We also assess whether this frequency is affected by inflation shocks. To do so, we use the Michigan and Livingston surveys, respectively for households and firms, the Survey of Professional Forecasters for professional forecasters, the FOMC forecasts for policymakers, and experimental data for participants to Learning-to-

¹ Note that disagreement could also be driven by heterogeneity in beliefs about the underlying structure of the economy (Angeletos et al., 2021). Such heterogeneity generates disagreement in expectations even if all agents have the same information about previous realizations of macroeconomic variables and of shocks. See Andrade et al. (2016) and Andre et al. (2021) for an empirical analysis of heterogeneity in subjective models.

Forecast Experiments (LtFEs) from a series of papers (Petersen, 2014, Pfajfar and Žakelj, 2018; Cornand and M'baye, 2018a, b; and Hommes et al, 2019).²

In tackling these issues, our contribution is also methodological. We are able to explore the external validity of experimental inflation expectations in terms of disagreement and frequency of forecast revisions relative to four categories of field expectations. Laboratory experiments are increasingly used to test the predictions of macroeconomic models or their assumptions (Duffy, 2008, 2016). Establishing the external validity of experimental inflation forecasts is essential if laboratory experiments are to be used as decision-making tools for monetary policy.³ Conclusions that can be drawn from experiments would only be valid if the experimental expectations present similarities with those observed in field data – in particular regarding information frictions.⁴

Our main results can be summarized as follows. First, we provide evidence of a strong heterogeneity in disagreement among the different categories of agents: while policymakers, professional forecasters and participants to experiments exhibit low disagreement, firms and households show strong disagreement. Second, we find that an inflation shock has a positive and significant effect on disagreement for all categories of agents, except for policymakers. Third, we observe a heterogeneous frequency of forecast revisions across categories of agents, with policymakers revising more frequently than participants to experiments, firms and professional forecasters, who themselves revise more frequently than households. Fourth, an inflation shock has a positive and significant effect on the frequency of forecast revisions for all categories of agents. In terms of magnitude, the frequency of forecast revision increases the least for households, firms, and professional forecasters. By contrast, participants to experiments revise more and policymakers even more.

Overall, our results qualify the view according to which firms and households have to be opposed in terms of information frictions. Especially when it comes to disagreement, there is more difference between two groups (one composed of policymakers, professional forecasters and participants to experiments and the other of households and firms) than within each of these two groups. Regarding the frequency of revision, policymakers and participants to experiments are the categories that stands in stark contrast to the others.

Regarding the external validity of experimental inflation expectations, while participants are supposed to play the role of firms in the considered set of experiments, our analysis suggests that the behavior of participants to experiments does not mimic much that of firms. It reproduces more closely the behavior of policymakers (and, to some extent, of professional forecasters), which might be due to the salience of information and to the New-Keynesian (NK) data generating process enacted in LtFEs.

² In LtFEs, participants' task is to provide their expectations about an economic variable (say, inflation). Their payoffs depend negatively on their forecast error. The expectations that are formed by participants are aggregated (using the mean or median) and this summary statistic is introduced into the theoretical model as the aggregate expectation of agents. Most recent experiments have used variants of the standard 3-equation New-Keynesian model: IS curve, Phillips curve, and policy rule. This model is directly implemented via a computer program, except for the expectations determined by participants. The computer program then derives the current values of variables conditional on the model parameters. See Hommes (2011) for a survey.

³ A recent growing macro-experimental literature (see Duffy, 2016 and Hommes, 2021, section 3) has considered inflation and/or output expectation formation in the laboratory. Laboratory experiments – particularly LtFEs – are used to validate expectation hypotheses and learning models and also serve as important tools for central bankers by providing a test bed for competing policy actions (Cornand and Heinemann, 2014, 2019).

⁴ In addition, the comparison between experimental and field data can be valuable for experimenters to improve the design of their macro-experiments in order to mimic real world situations.

Our paper contributes to the empirical literature that documents disagreement in inflation expectations obtained from survey data, within and across different categories of economic agents. The closest paper to ours is Andrade and Le Bihan (2013). They show that forecasters do not systematically update their forecasts even when new information is released and that forecasters who update also disagree on their forecasts.⁵ The main difference however is that they analyze one agent only (the ECB' Survey of Professional Forecasters). We apply and extend their methodology to five categories of agents. In this respect, we are close to Carroll (2003) who compares professional forecasters to consumers and Coibion and Gorodnichenko (2015) who consider forecasts from professional forecasters, firms, households and policymakers. Coibion and Gorodnichenko (2012) apply a different test of the expectation formation process to U.S. and international survey data from professional forecasters. Mankiw et al. (2004) document the extent of disagreement and show that it varies over time and with other aggregate variables. Finally, we complement this literature by including experimental data, allowing us to compare it to the four categories of field expectations and test the external validity of experimental inflation expectations as in Cornand and Hubert (2020).⁶ While Coibion and Gorodnichenko (2012, 2015) and Cornand and Hubert (2020) consider aggregate data, our dataset exploits individual data to document the degree of information frictions.

The literature that compares experimental data to field data in macroeconomics is scarce. There is a recent literature on information-provision experiments in surveys.⁷ This literature combines large-scale survey data of households and firms with Randomized Controlled Trial (RCT) experiments. RCT experiments provide a random subset of survey respondents with a piece of information and measure the corresponding effects on expectations. In particular, Link et al. (2021) study how information frictions (the dispersion of expectations and the learning rate from information, which can be comparable, to some extent, to the frequency of forecast revisions) vary between households and firms. They show that firms' expectations are less dispersed than those of households and more closely aligned with expert forecasts. Our research question is closely related to theirs, but our methodology is different. While RCT experiments allow for representative data of real-world expectations and actual decisions, most of them only allow collecting cross-sectional data. This does not allow to study how information frictions vary with the state of the economy. By contrast, we conduct a panel data analysis (allowing for a proper measurement of the frequency of forecast revisions⁸) and control for inflation dynamics.

In terms of policy implications, our findings may inform central banks about the public they should target to improve their communication strategy in order to cope with information frictions, both within and across categories of economic agents. In particular, acknowledging the size of disagreement within and across each category of agents (implying that the information released by the central bank may not reach all categories of agents and also all agents within each category in the same manner), targeted communication towards each

⁵ Clements (2020) complements this analysis by documenting whether inefficiencies in the use of information can explain the accuracy of forecasts and disagreement between forecasters. He provides evidence that the inefficient use of information is responsible for persistent differences in accuracy across forecasters.

⁶ Compared to Cornand and Hubert (2020), the experimental data of Adam (2007) is excluded (for which only average expectations per group – and not individual ones – were available) because the focus here is on individual rather than aggregate data. For the same reason, we also exclude financial market expectations and Greenbook data from our field sample. Moreover, compared to the experimental dataset used by Cornand and Hubert (2020), we include an additional experimental paper (Petersen, 2014).

⁷ Armantier et al. (2015), Armantier et al. (2016), Cavallo et al. (2017), Coibion et al. (2018a) and Coibion et al. (2018b) use this method to study inflation expectations.

⁸ By contrast, in Link et al. (2021), the learning rate is based on an exogenous allocation of information.

category and toward specific groups of agents (presenting the same characteristics) within each category might represent a useful tool.

The paper is structured as follows. Section 2 presents the data. Section 3 depicts testable hypotheses. Sections 4 and 5 respectively describe the empirical results in terms of disagreement and frequency of forecast revisions. Section 6 concludes the paper.

2. Data

We collect inflation expectation data from three types of measures (survey and policymaker data as well as experimental data), corresponding to five categories of agents (households, industry, professional forecasters, policymakers and participants in experiments).

2.1. Survey data

Households (*Michigan*)

The Michigan Survey of Consumer Attitudes and Behavior surveys a cross-section of the population about their expectations over the next year. Most papers using the Michigan survey cover only the period since 1978, during which these data have been collected monthly and on a quantitative basis: respondents were asked to state their precise quantitative inflation expectations. Before then, the Michigan survey was qualitative. It has been conducted quarterly since 1946, although for the first 20 years, the respondents were asked only whether they expected prices to rise, fall, or stay the same. Each month, a sample of approximately 500 households is interviewed, in which the sample is chosen to statistically represent households in the US, excluding Alaska and Hawaii. Survey respondents are questioned twice on average, sometimes thrice. The monthly phone call survey focuses on respondents' perceptions and expectations regarding personal finances, business conditions and news regarding the economy in general, as well as macroeconomic aggregates, such as unemployment, interest rates and inflation. Furthermore, the survey collects individual and household socioeconomic characteristics.⁹

Firms (*Livingston*)

The Livingston Survey was started in 1946 by the late columnist Joseph Livingston. It is the oldest continuous survey of firms' expectations. It summarizes the forecasts of analysts and economists working in the industry sector in the US. The Federal Reserve Bank of Philadelphia took responsibility for the survey in 1990. It is conducted twice per year, in June and December, so it has a semiannual frequency. It provides twelve-month Consumer Price Index (CPI) inflation forecasts from approximately 50 survey respondents. We consider that expectations collected via the Livingston survey represent firms' expectations. But, as these are expectations of firms' economists, we acknowledge that they could share the properties of that of professional forecasters. The subsequent results provided in Sections 4 and 5 suggest that Livingston expectations differ from those of professional forecasters in various respects.

Professional forecasters (*Survey of Professional Forecasters*)

The Survey of Professional Forecasters (SPF) is collected and published by the Federal Reserve Bank of Philadelphia. It focuses on professional forecasters mostly in the banking sector in the US. Surveys are sent to approximately 40 panelists at the end of the first month of the quarter,

⁹ We acknowledge that the Michigan survey includes questions formulated in a very broad manner rather than targeted on inflation, which could induce a bias toward more dispersion in inflation expectations.

the deadline for submission is the second week of the second month of the quarter, and forecasts are published between the middle and end of February, May, August, and November. GDP price index forecasts (available since 1968) are fixed-horizon forecasts for the current and the next four quarters. They are provided as annualized quarter-over-quarter growth rates. We also perform our analysis with CPI forecasts provided since 1981. We consider the median of individual responses, rather than the mean, which could be affected by potential outliers.

2.2. Policymakers: Federal Open Market Committee (FOMC)

The FOMC has published forecasts for inflation and real GDP growth twice per year in the Monetary Policy Report to the Congress since 1979. Since October 2007, their publication has been quarterly. We consider forecasts of the Consumer Price Index until 1999 and then the Personal Consumption Expenditures (PCE) measure of inflation following the focus of the FOMC. These forecasts are fourth quarter-over-fourth quarter growth rates for the current and next calendar years. Until 2005, the forecast for the next year was published only per year. Individual members' FOMC forecast are made public since 1992, but only summary statistics or anonymous individual data are published in real-time, with an embargo of 10 years.¹⁰

2.3. Laboratory experiment data

We collect a sample of macro-experimental data on inflation expectation from five published papers.¹¹ The Learning-to-Forecast design, based on the NK reduced-form model or some variant, offers the incentives to form accurate inflation forecasts. Four out of five considered experimental papers implement variants of the standard NK three equation model, with the IS curve, Phillips curve, and policy rule:

$$\begin{aligned} y_t &= E_t y_{t+1} - \varphi(i_t - E_t \pi_{t+1}) + g_t \\ \pi_t &= \lambda y_t + \rho E_t \pi_{t+1} + u_t \\ i_t &= \bar{\pi} + \phi_\pi(\pi_t - \bar{\pi}) + \phi_y(y_t - \bar{y}), \end{aligned}$$

where π_t and y_t are the inflation rate and output gap in period t , $\bar{\pi}$ and \bar{y} are their steady state values, i_t is the nominal interest rate, g_t and u_t are exogenous disturbances, $E_t \pi_{t+1}$ is the average expected inflation, $E_t y_{t+1}$ is the average expected output gap, φ , λ , ρ , ϕ_π , and ϕ_y are positive parameters. The economy is qualitatively described to participants in experiments. Instructions include an explanation of the mechanisms that govern model equations (in particular, monetary policy is described as central bank intervention reducing interest rates to increase inflation and conversely).

Pfajfar and Žakelj (2018) (henceforth PZ) presented an LtFE conducted at the University Pompeu Fabra in Spain and the University of Tilburg in the Netherlands, based on the above-presented reduced form of the NK model. They asked subjects to form a prediction of the $t+1$ period inflation and the 95% confidence interval of their inflation prediction. The computer program feeds the model with naïve output gap expectations: $E_t y_{t+1} = y_{t-1}$ in the above model specification. The parameters values are standard: $\rho = 0.99$, $\lambda = 0.3$, $\varphi = 0.164$, and $\bar{\pi} = 3$. Since they investigated the targeting rule that best stabilizes the economy, they considered four different treatments, corresponding to four different policy rules: inflation forecast

¹⁰ We do not consider Greenbook data since our analysis of information frictions requires individual data.

¹¹ For a systematic comparison of the characteristics of experimental designs used in the five considered papers, see Table A1 in the Appendix.

targeting, with three different degrees of monetary policy aggressiveness: $\phi_y = 0$ and $\phi_\pi = 1.5$ or 1.35 or 4; and contemporaneous inflation targeting, with an intermediate degree of monetary policy aggressiveness: π_t is replaced by $E_t\pi_{t+1}$ in the monetary policy rule with $\phi_\pi = 1.5\%$. The disturbances g_t and u_t follow an AR-1 process. Subjects observe the history of macroeconomic variables: at each period t , participants observe inflation, the output gap and the interest rate up to period $t-1$. There are 70 periods, and each period corresponds to one quarter. The number of observations amounts to 24 independent groups.

Cornand and M'baye (2018a) (henceforth CMa) focused on a design very close to that of PZ. Their experiment relied on the same model, but the considered parameter values were slightly different: $\rho = 0.99$, $\lambda = 0.3$, $\varphi = 1$, $\bar{\pi} = 5$. As in PZ, CMa asked subjects to state inflation expectations but not output expectations, for which they assumed again $E_t y_{t+1} = y_{t-1}$. CMa studied the role of the central bank's Inflation Target (IT) communication by comparing treatments in which the central bank explicitly announces its IT to treatments in which the central bank does not announce the IT. They considered four treatments differing by the type of IT procedure.¹² In contrast to PZ, g_t and u_t are independently and normally distributed shocks with a standard deviation of 0.1. For each treatment, CMa has four sessions, and each session lasted 50 periods.

Cornand and M'baye (2018b) (henceforth CMb) is similar to CMa in terms of design. The parameter values are the same, and they focus on the case in which the central bank stabilizes both inflation and the output gap ($\phi_\pi = 1.5$, $\phi_y = 0.5$). They considered four different treatments differing with respect to whether the central bank implements a band or point IT and also by the size of shocks (g_t and u_t are independently and normally distributed shocks with standard deviations of 0.08 or 0.4 depending on the treatment).¹³ The authors had four sessions for each treatment, and each session lasted 60 periods. Both experiments by Cornand and M'baye were conducted at the GATE-Lab of the University of Lyon in France.

Hommes et al. (2019) (henceforth HMW) presented an LtFE conducted at the CREED lab at the University of Amsterdam in the Netherlands, based again on the simple standard version of the NK model considered above. The parameter values are the same as in CMa, except for the IT, which is lower: $\bar{\pi} = 3.5$. A main difference is that subjects' task consists of forming *both* inflation and output gap expectations in period t for period $t+1$. They considered two different treatments, corresponding to the implementation of two different policy rules by the central bank: one in which the central bank reacts to inflation only ($\phi_\pi = 1.5$, $\phi_y = 0$); and one in which the central bank additionally reacts to the output gap ($\phi_\pi = 1.5$, $\phi_y = 0.5$). As in PZ, CMa and CMb, subjects observe the history of macroeconomic variables up to period $t-1$. Shocks g_t and u_t are independently and normally distributed with a standard deviation of 0.1. There are 50 periods; the number of observations amounts to 43 independent groups.

¹² More precisely, under strict IT, the sole objective of the central bank is to stabilize inflation ($\phi_\pi = 1.5$, $\phi_y = 0$). Under the explicit strict IT treatment, the central bank announces its 5% target, while under the implicit strict IT treatment, there is no announcement about the target value. Under a flexible IT, the central bank has both an inflation objective and an output gap stabilization objective ($\phi_\pi = 1.5$, $\phi_y = 0.5$). Again, depending here on whether the flexible IT is made explicit or not, the central bank communicates its IT or not.

¹³ More precisely, the considered treatments were the following. In the band IT with small shock treatment, the central bank simply announces a band IT (interval [4% - 6%]) to the public, in a context in which shocks have low variance. In the point IT with small shock treatment, the central bank explicitly communicates its 5% numerical IT with a tolerance band of +/-1% around its target in a context in which the variance of shocks is low. In the band IT with large shock treatment, the central bank simply announces the band IT ([4% - 6%]) to the public but in a context in which the variance of shocks is relatively high. In the point IT with large shock treatment, the central bank explicitly communicates its 5% numerical target with a tolerance band of +/-1% around its target but in a context in which the variance of shocks is relatively high.

Petersen (2014) presents an LtFE conducted in Montreal, Quebec (with both students and non-students), based on a slightly modified four equation version of the above NK economy where households and firms make optimal decisions given their expectations:

$$\begin{aligned} y_t &= E_t y_{t+1} - \varphi(i_t - E_t \pi_{t+1} - r_t^n) \\ \pi_t &= \lambda y_t + \rho E_t \pi_{t+1} \\ i_t &= \bar{\pi} + \phi_\pi (E_{t-1} \pi_t - \bar{\pi}) + \phi_y (E_{t-1} y_t - \bar{y}) \\ r_t^n &= \phi r_{t-1}^n + \epsilon_t \end{aligned}$$

where r_t^n is the natural rate of interest (ϵ_t is assumed to be drawn from a normal distribution with mean zero and standard deviation 1.13) and chosen parameter values ($\rho = 0.989$, $\lambda = 0.13$, $\varphi = 1$, $\bar{\pi} = 0$, $\bar{y} = 0$, $\phi_\pi = 1.5$, $\phi_y = 0.5$, $\phi = 0.57$) are intended to mimic the Canadian economy. Each experimental session consisted of two stationary repetitions, with 9 subjects interacting within a group, consisting of approximately 50 periods each. Each period, subjects were provided information about the current period's interest rate, shock to the natural rate of interest, and the expected shock size in the following period. While two treatments are considered in this experiment, the data we managed to collect refer to a single information treatment, in which subjects are also provided with forecast error information on their screen. They were asked to provide forecasts for next period's inflation and output gap in basis points. The current period's inflation and output and the next period's nominal interest rate are then computed using the median forecasts for inflation and output.

2.4. Macroeconomic data

Regarding experimental data, inflation is generated by a computer program that implements a model of the economy, conditional on the parameters and on the expectations that participants to the experiment are asked for (inflation expectations for all experiments considered in this paper as well as output gap expectations in Hommes et al. (2019) and Petersen (2014)). For the observed macroeconomic data, we use the monthly Consumer Price Index for All Urban Consumers (FRED mnemonic: CPI-AUCSL).

In order to measure inflation shocks, we estimate a GARCH(1,1) model of the inflation rate. This specification has various advantages for our approach. First, a GARCH model provides a parsimonious identification of the conditional variance combined with an agnostic statistical representation of the inflation process. Second, the GARCH model fits very well data-generating processes in which the volatility of a series varies over time. So it captures this time-dependent volatility as a function of observed prior volatility. Third, the question of whether disagreement and the frequency of forecast revisions – two variables defined in the positive space – are correlated to inflation shocks requires these inflation shocks to be defined in the positive space as well. Using as a proxy for inflation shocks, the residuals from an AR(1) estimation of the inflation process or from a Phillips curve would answer a different question as these residuals would be defined in the positive and negative spaces. Although we could consider the absolute value of these residuals, estimating a GARCH model has the benefit of focusing our analysis on the conditional variance of inflation, which is the key dimension we want to investigate for the question of information frictions.

The GARCH model is estimated with maximum likelihood and based on these two equations:

$$\pi_t = \beta_0 + \beta_1 \pi_{t-1} + \epsilon_t, \quad \epsilon_t \sim (0, \sigma_t^2) \quad (1)$$

$$\sigma_t^2 = \gamma_0 + \gamma_1 \epsilon_{t-1}^2 + \gamma_2 \sigma_{t-1}^2 \quad (2)$$

where π_t is the inflation rate regressed on its own lags in the mean equation. The number of lags in the variance equation for both the error term and its variance is set to one. We estimate these equations for both the observed macroeconomic data and the experimental data.

2.5. Summary of inflation expectations across the five categories of agents

Table 1 presents the source, frequency and sample of inflation expectations for our five categories of agents. We acknowledge the heterogeneity of the different datasets with respect to their frequency and the sample period considered. While frequency may differ from one set to the other, it is worth emphasizing that it corresponds to the frequency of usual economic decisions for each category of agents. For experimental forecasts, frequency is abstract. Regarding the sample period for field data, we complement, as robustness tests, our empirical analysis with tests performed on a common sample period, from 1992 to 2009, for comparability purposes between types of agents, as well as macroeconomic and structural environments.

Table 1 – Characteristics of inflation expectations data

	Source	Frequency	Sample
Professional F.	SPF	Quarterly	1981q2 - 2020q2
Policymakers	FOMC	6 months - Quarterly	1992m1 - 2009m11
Firms (Industry)	Livingston	6 months	1948h1 - 2020h1
Households	Michigan	6 months	1978m1 - 2020m5
Experiments	Experiments	Quarterly equivalent	NA

The various datasets that we consider provide two different forms of inflation forecasts: fixed-event or fixed-horizon forecasts. Fixed-horizon forecasts are preferable for the analysis of disagreement since they are not influenced by decreasing forecasting horizon. Following Dovern et al. (2012), we construct fixed-horizon forecasts (at the 1-year horizon) as a weighted average of fixed-event forecasts (using current-year and next-year forecasts as well as the number of months forecasted in each year). We are therefore able to compare all forecasts on the same ground with a similar fixed-horizon (1-year) scheme. We provide alternative measures of disagreement and the frequency of forecast revisions using the raw fixed-event data in the Appendix to ensure that our results are not driven by this transformation.

3. Predictions within and across the different categories of agents

In an economy without frictions, there should be no disagreement and forecasts should be revised continuously. By contrast, according to models with information frictions, economic agents do not revise their forecasts continuously and disagree in their forecasts. In sticky information models (Mankiw and Reis, 2002), agents update their information set infrequently, with a given constant probability. Consequently, at each date, only a fraction of the population has access to macroeconomic news and revise their expectations. Therefore, disagreement among agents rises after any economic shock, but the probability to revise their forecast is set exogenously. In models with dispersed and heterogenous information (Angeletos and Lian, 2019), agents update their information but the news they get is imperfect, so that they partly integrate it in their forecast. In models of rational inattention (Sims, 2003), agents may rationally not account for information and use different information sets. These latter two classes of models can be grouped into what the literature calls noisy information

models. Under noisy information models, the amount of disagreement is independent of shocks, but the frequency of forecast revisions rises after an economic shock.¹⁴

The theoretical results associated with models of information frictions could in principle be applied to any category of economic agents who form inflation forecasts. Forecasting inflation accurately is obviously crucial for monetary policymakers. It is also central for firms to predict their future demand or their future financial constraints. As argued by Link et al. (2021), households have an interest in holding accurate expectations about changes in the cost of living. Professional forecasters also have incentives, especially in terms of credibility and participants to experiments have direct monetary incentives. In LtFEs, the tasks of participants to experiments are incentivized. Therefore, all our categories of agents may perceive high stakes in being informed about future inflation.

However, some categories of agents may face lower costs of acquiring and processing information than others. Carroll (2003) shows that the opinion of professional forecasters spreads to households. It thus influences their expectations and decisions. As Andrade and Le Bihan (2013, p. 968), we expect “*the extent of attention to news among professional forecasters [to be] an upper bound for other agents’ attention to aggregate conditions*”. Another upper bound is naturally that of policymakers’ forecast revisions. However, in contrast to FOMC and SPF expectations, households’ ones are expected to be less regularly revised. Indeed, policymakers and professional forecasters have in principle more information about future shocks. By contrast, we expect households to less revise their expectations. While the expectations formation process in an LtFE depends on the design features, participants to LtFEs generally play the role of price-setting firms.¹⁵ It is thus natural to postulate that laboratory subjects behave like firms. Regarding disagreement, we expect policymakers and professional forecasters to form more consensual forecasts than other categories of agents (firms, households and participants to experiments) since they might update more frequently their information set.

We can summarize our predictions around the following four items:

H1 - Disagreement within each category of agents. Under the assumption of no information friction, we expect no disagreement. By contrast, in the case of information frictions, we expect disagreement to react either positively to an inflation shock (sticky information) or not (noisy information).

H2 - Disagreement across categories of agents. In the case of information frictions, we expect policymakers and professional forecasters to disagree less than firms and households. We conjecture experimental data to mimic firms’ behavior.

¹⁴ The empirical literature shows that there is a correlation of disagreement and forecast revisions with macroeconomic variables and in particular inflation dynamics. Carlson and Valev (2003), Carroll (2003), Cukierman and Wachtel (1979), Mankiw et al. (2004) and Souleles (2004) provide evidence of a positive relationships between the cross-sectional dispersion in inflation expectations and the level of the inflation rate, Cukierman and Wachtel (1979) and Mankiw et al. (2004) document the link between the variance of measured inflation and the cross-sectional dispersion. Consistent with noisy information models, Coibion and Gorodnichencko (2015) cannot reject the hypothesis of no response of disagreement to shocks but can reject the hypothesis that disagreement responds in the manner predicted under sticky information. They conclude that noisy information models better render account for expectations of professional forecasters, consumers, firms and central bankers.

¹⁵ Pfajfar and Zakelj (2018) argue that they do not assign the participants a particular role in the economy they consider and that participants rather play the role of professional forecasters who provide firms with their inflation forecasts. However, their NK model is directly fed with participants’ inflation forecasts, in the place of standard expectations of firms.

H3 - Frequency of forecast revisions within each category of agents. Under the assumption of no information friction, we expect a probability of forecast revision equal to 1 in each period. So an inflation shock does not affect the frequency of forecast revisions. By contrast, in the case of information frictions, we expect a probability of forecast revision strictly below 1 in each period, that is either not affected (sticky information) or positively affected (noisy information) by an inflation shock.

H4 - Frequency of forecast revisions across categories of agents. In the case of information frictions, we expect policymakers and professional forecasters to revise their forecast more frequently than firms. We expect households to be the category of agents that revises the least frequently. Again, we conjecture participants to experiments to mimic firms.

4. Disagreement

We measure the disagreement in expectations within our five categories of economic agents and proceed to comparisons across categories of agents.

4.1. Descriptive statistics

Following Andrade and Le Bihan (2013), a natural measure of disagreement among economic agents about inflation is the cross-section standard deviation of inflation expectations for each category of agents:

$$SD^{f(\pi)} = \sqrt{\frac{1}{n} \sum_{i=1}^n (f_i^\pi - f^\pi)^2} \quad (3)$$

where f_i^π denotes individual i 's fixed-horizon forecasts for inflation π on the whole sample (of the considered category of agents), one relevant period ahead (4 quarters for the considered categories), $f^\pi = \frac{1}{n} \sum_{i=1}^n f_i^\pi$ is the average (or consensus) inflation forecast, with n the number of respondents to the survey. The variable π denotes the year-on-year inflation rate.

We also consider two other additional measures of disagreement, the interquartile range:

$$IQR^{f(\pi)} = p_{75}^{f(\pi)} - p_{25}^{f(\pi)} \quad (4)$$

and the interdecile range:

$$IDR^{f(\pi)} = p_{90}^{f(\pi)} - p_{10}^{f(\pi)} \quad (5)$$

where $p_x^{f(\pi)}$ denotes the percentile x of inflation forecasts.

Table 2 presents the mean and median forecasts, the standard-deviation of inflation forecasts as calculated by (3), the interquartile range as calculated by (4), the interdecile range as

calculated by (5), as well as the normalized values (over the mean of the whole sample) of the interquartile range (IQR_N).¹⁶ Figure 1 depicts the distribution of expectations.¹⁷

All our categories of agents exhibit disagreement about inflation.¹⁸ However, there is a strong heterogeneity among the different categories of agents. While policymakers, professional forecasters and participants to experiments exhibit low disagreement (with respective IQR of 1.05, 1.30, and 0.70 in Table 2), firms and households are characterized by a stronger disagreement (with respective IQR of 2.94 and 4). As shown on Figure 1, the distribution of inflation expectations is flatter for these two categories of agents, which is also consistent with the larger spread between IQR and IDR for these two categories. These comparisons across categories are robust to the normalization of our measures of disagreement, implying that this disagreement does not mechanically depend on the level of inflation.

Table 2 - Descriptive statistics about disagreement - Eq. (3), (4), (5)

	(1)	(2)	(3)	(4)	(5)	(6)
	IQR	IDR	SD	Mean	Median	IQR_N
SPF	1.30	3.02	1.40	2.92	2.60	0.45
FOMC	1.05	1.93	0.72	2.15	2.13	0.49
Livingston	2.94	6.51	3.04	3.10	2.93	0.95
Michigan	4.00	10.00	6.16	4.46	3.00	0.90
Experiments	0.70	1.77	1.63	3.45	3.43	0.20

Note: These statistics are computed over the full sample for each of the five datasets. The IQR is the Inter Quantile Ratio. The IDR is the Inter Decile Ratio, the distance between the 90th and 10th percentiles. The SD is the standard deviation. The IQR_N is a normalized IQR by the level of inflation forecasts and corresponds to the ratio of the IQR divided by the mean.

The fact that firms and households disagree more is in line with both our predictions and the work of Coibion et al. (2018a). Nevertheless, by providing a wider picture (including policymakers and participants to experiments in our dataset), our results put into perspective the findings of Link et al. (2021), according to which firms' expectations are less dispersed than those of households and more closely aligned with expert forecasts: it seems that firms are further away from professional forecasters than professional forecasters are from policymakers. Moreover, we refine the result of Mankiw et al. (2004, p. 216), who show a "correspondence between disagreement among policymakers and disagreement among professional economists". While their study is based on the range of FOMC forecasts, our results based on individual FOMC data confirm their findings.

In terms of the external validity of experimental inflation forecasts, we find that experimental data do not well reproduce the forecasting behavior of firms that they are supposed to mimic in the considered experiments. This could be surprising since students usually recruited to perform experimental tasks and people working in the industry present relatively comparable level of education. Nevertheless, the fact that the expectations of participants to experiments rather mimic the expectations of professional forecasters or policymakers may be due to the high stakes in the incentivized forecasting task (more in line with survey of professional

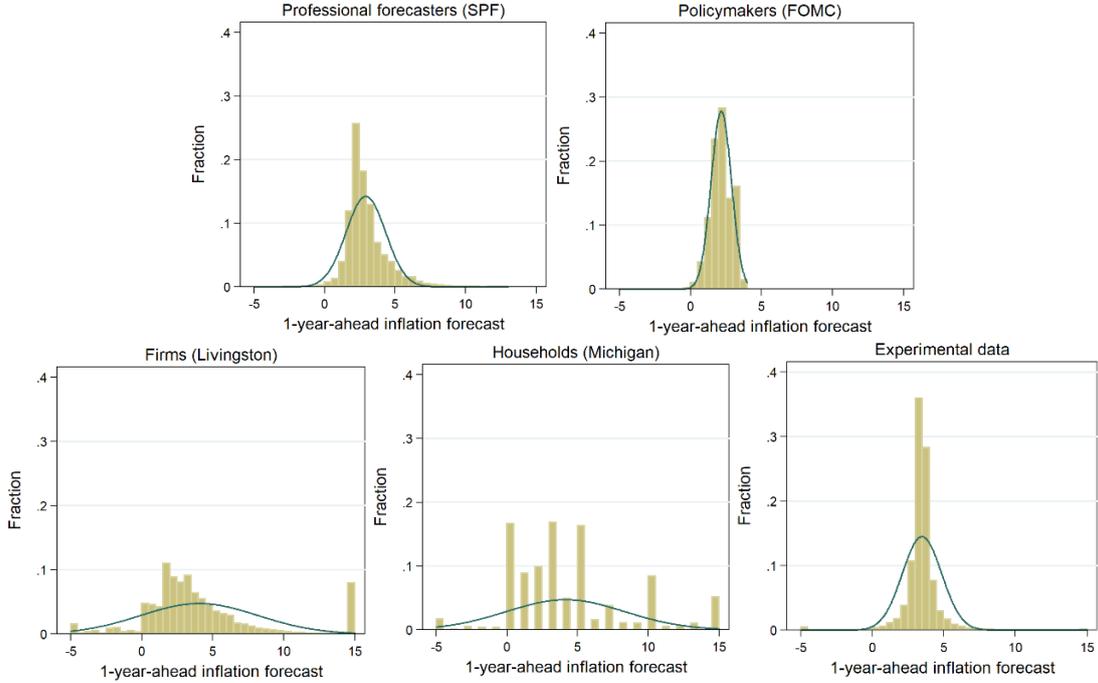
¹⁶ This normalization allows us to consider the potential issue that may arise due to the mechanical link between the dispersion of forecasts and the level of the underlying variable. Indeed, a high level of inflation may mechanically increase the dispersion of inflation forecasts due to the larger span of possible values. We also acknowledge the potential asymmetric distribution of forecasts in the vicinity of zero, that separates deflation and inflation, that can reduce forecast dispersion when inflation approaches zero.

¹⁷ Figure A1 in Appendix presents the cumulative density function of the five different agents' expectations.

¹⁸ Note that we observe some specificities for some categories on Figure 1: households focus on round numbers, firms are subject to threshold effects, and professional forecasters and policymakers' forecasts exhibit tails that are not fully symmetric.

forecasters or central bankers' forecasting tasks, whose incentives to form accurate forecasts are either financial or reputational) on the one hand and the nature of the data generating process (NK model) that central banks also use themselves to understand the working of the economy on the other hand. Indeed, in LtFEs, the NK data-generating process motivates the forecasts of participants and macroeconomic data information are made salient to participants. Salient macroeconomic data and information about the data-generating process may contribute to reduce disagreement.

Figure 1 - Distribution of inflation expectations



Note: These figures show the distribution of inflation forecasts for each dataset truncated at -5% and 15%, with the fraction that represents each bin on the y-axis. The blue line represents the normal density approximation.

4.2. Disagreement and inflation dynamics

We investigate whether disagreement is affected by inflation shocks. We regress our second measure of disagreement $IQR_t^{f(\pi)}$ on the conditional variance of inflation and we control for the level of inflation to capture the potential mechanical link between forecast dispersion and the level of the underlying variable:

$$IQR_t^{f(\pi)} = \alpha_1 + \beta_1 v_t^\pi + \beta_2 \pi_{t-1} + \eta_t \quad (6)$$

where v_t^π is the conditional variance of inflation, and β_1 (expected to be null under the assumption of no information friction) and β_2 are the estimated coefficients. Observe that in equation (6), contrary to equations (3), (4) and (5), we calculate a measure of dispersion for each date t , instead of over the whole sample of individual expectations. Equation (6) is estimated with OLS and heteroskedasticity-robust standard errors.

Table 3 presents the results from the estimation of equation (6) for all categories of agents.¹⁹ The conditional variance of inflation (CPI Cond. Var.) positively and significantly affects disagreement for the considered categories of agents, except for policymakers. This result is consistent with the prediction of the sticky information model. In terms of magnitude, disagreement reacts more to the conditional variance of inflation for firms and households, less so for participants to experiments and professional forecasters. Disagreement does not increase to a large extent for policymakers. However, the fact that the conditional variance of inflation has no effect on disagreement for policymakers has to be qualified by the small number of observations for this category of agents (three explanatory variables for 28 observations only). Figure A2 in Appendix illustrates these findings by exhibiting a clear, positive, relationship between dispersion and the conditional variance of inflation for *all* categories of agents. Tables A4, A5, A6 and A7 in Appendix show that the findings regarding the effect of the conditional variance on inflation on all categories of agents are robust to respectively taking IDR and SD as alternative measures of disagreement, considering a common time-sample for all field data, and controlling for the level of forecasts.

Table 3 – Drivers of disagreement – Eq. (6)

	(1)	(2)	(3)	(4)	(5)
	SPF	FOMC	Livingston	Michigan	Experiments
	IQR _t				
CPI CondVar	0.114*** [3.74]	0.031 [0.94]	0.495*** [3.16]	0.390*** [8.32]	0.299*** [6.12]
CPI Level	0.188*** [9.21]	0.002 [0.09]	0.233 [1.24]	1.461*** [19.22]	0.003 [0.08]
Constant	0.775*** [38.60]	0.243*** [13.72]	1.437*** [12.96]	4.675*** [87.92]	0.369*** [24.60]
N	156	28	145	509	5882
R2	0.40	0.10	0.20	0.60	0.32

Note: t-statistics in brackets. * p < 0.10, ** p < 0.05, *** p < 0.01. Parameters are estimated with Equation (6) using OLS. The conditional variance of inflation is estimated based on a GARCH(1,1) model and the inflation level is introduced with a lag. The dependent variable is the IQR computed per period for each dataset.

To summarize, we observe disagreement for all categories of agents, which suggests the presence of information frictions for all categories. Firms and households seem particularly affected by information frictions, while policymakers, professional forecasters and participants to experiments less so. Disagreement is linked to the frequency of forecast revision: if economic agents revise their forecasts infrequently, disagreement will necessarily increase. We now turn to this component of information frictions.

5. Frequency of inflation forecast revisions

We measure the frequency of forecast revision within our five categories of economic agents and proceed to comparisons across categories of agents.

5.1. Descriptive statistics

To measure the frequency of forecast revisions of our five categories of economic agents, we look at whether forecasts in t differ from forecasts in $t-1$ and compute how many subjects revise their forecast for each category of agents. We consider the frequency of forecast revisions with

¹⁹ Table A3 in the Appendix provides a robustness check (to collinearity issues) of these findings when regressions are conducted separately for the conditional variance of inflation and the level of inflation.

respect to the time unit that corresponds to each dataset and compare these frequencies of revisions across our different categories of agents directly. The frequency of forecast revisions provides an indication about the extent to which agents incorporate new information in their expectations. Following Andrade and Le Bihan (2013), the probability we estimate is:

$$P_i \left(f_{it,t+h}^\pi \neq f_{it-1,t+h}^\pi \right). \quad (7)$$

Table 4 presents the probability of forecast revision for the five considered categories of agents, its variance, the number N of observations, and for agents who revise their forecast, the average revision (Mean Rev.), the standard deviation of revisions (SD Rev.), and the number of agents who revised.²⁰

Table 4 - Descriptive statistics about frequency of forecast revisions - Eq. (7)

	(1)	(2)	(3)	(4)	(5)	(6)
	P_{Mean}	P_{SD}	N	Mean Rev.	SD Rev.	N_{Rev}
SPF	0.88	0.33	4,089	0.52	0.67	3,583
FOMC	0.96	0.19	241	0.40	0.41	232
Livingston	0.90	0.29	5,751	1.45	1.92	5,204
Michigan	0.75	0.43	91,390	4.58	5.44	68,500
Experiments	0.92	0.28	41,280	1.07	2.14	37,778

Note: These statistics are computed over the subsample N for which we observe 2 consecutive forecasts of a same individual for each of the five datasets. The P_{Mean} is the average frequency of revisions for each dataset. The P_{SD} is the standard deviation of the frequency of revisions. The Mean Rev. is the average magnitude of the revision for individuals who revised their forecasts (the subsample N_{Rev}) while the SD

Policy makers are those who show the highest frequency of revision, but they are also those who revise the least (i.e. they have the lowest mean revision). Firms as well as participants to laboratory experiments and professional forecasters are also revising frequently. The amount by which professional forecasters revise is lower than that of participants to experiments and firms. Households are those who revise the least frequently (though with a larger heterogeneity – larger variance – than other categories), but present the highest mean revision (again with much variance). These results suggest that households is the category that is the most subject to information frictions.²¹

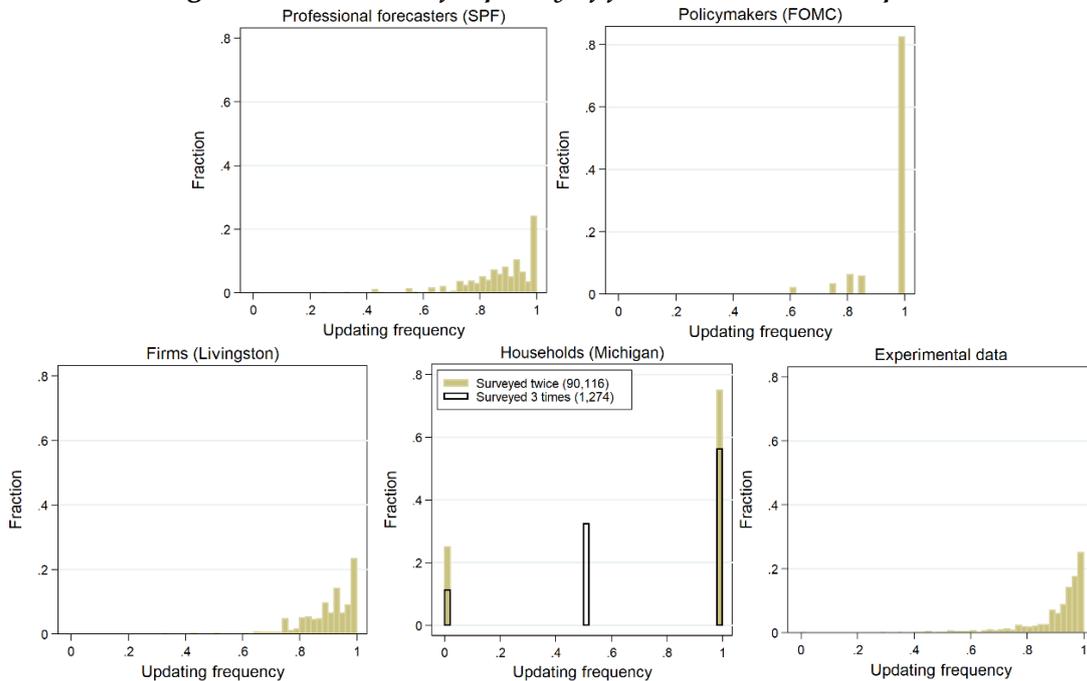
Figure 2 presents the distribution of individual frequency of forecast revision for each category of agents. It complements previous results in specifying whether the overall frequency of revision from Table 2 is reflecting individual behavior well (low heterogeneity within a category of agents) or not (large heterogeneity within a category of agents).

²⁰ Note that in the Livingston survey, participants form forecasts about the price level rather than its growth rate (i.e. the inflation rate). When it comes to exploring the frequency of *inflation* forecast revisions, the number of decimals in the computed growth rate matters. In order to exclude the potential upward bias in revisions induced by the decimals of a computed growth rate, we consider forecasts as unchanged if the change between two dates is less than 0.05 percentage point (see Clements, 2021, on forecasters' rounding behavior). We provide alternative measures of the frequency of forecast revisions using the raw data in the Appendix (Table A8).

²¹ Recall that we make comparisons across our categories of agents by using the frequency of forecast revisions with respect to the time unit that corresponds to each dataset. The benefit of such an approach is to consider whether agents revise their forecasts when they are asked for such forecasts. A drawback though is that agents may revise more frequently during this period, so our measure of the frequency of forecast revisions is an upper bound. We could instead divide the frequency of revisions for each category by the time unit to make thinner comparisons. For example, in the case of households, the frequency of 0.75 is computed on semi-annual data. On a quarterly basis, the frequency of revision would be 0.375.

In the case of central bankers, a vast majority (more than 80%) revise with probability 1, while some revise their forecast much less. Regarding professional forecasters, firms, and participants to experiments, about 30% revise with probability 1, while the large majority revises their forecast with a probability between 0.8 and 1.²² Note that the frequency with which professional forecasters revise their forecasts is more dispersed than that of other categories of agents. Finally, a disclaimer applies to the Michigan survey. In this survey, households are only observed twice or thrice and then are dropped out of the sample, so no comparison can be made with other categories of agents. However, it is possible to compare within this category what happens when households are surveyed twice versus thrice. More than 20% of households who are observed twice do not revise their forecast, while a bit less than 80% revise with probability 1. Both proportions are smaller when households are observed thrice, as about 30% of them revise with probability 0.5. Overall, we conclude that, except for the households (and for the reason above-mentioned), there is a relative homogeneity in the frequency of revision within each category of agents.

Figure 2 – Individual frequency of forecast revision – Eq. (7)



Note: These subfigures show the distribution of the average, by individual, of his/her frequency of forecast revisions. For the Michigan survey, because we observe individuals only 2 or 3 times, they can only not revise ($P=0$), revise 1 over 2 times ($P=0.5$) or revise all the time ($P=1$).

5.2. Frequency of forecast revisions and inflation dynamics

Andrade and Le Bihan (2013) also study whether the frequency of forecast revision depends on the state of the economy. Following their approach, we look at whether the frequency of forecast revision depends on inflation shocks for each of our categories of agents. We regress P_i over the conditional variance of inflation and we control for the level of inflation:

$$P_i \left(f_{it,t+h}^\pi \neq f_{it-1,t+h}^\pi \right) = \alpha_2 + \beta_3 v_t^\pi + \beta_4 \pi_{t-1} + \eta_t \quad (8)$$

²² Because we use a similar methodology to Andrade and Le Bihan (2013), our measures of the probability that professional forecasters revise their forecast are very comparable. According to Andrade and Le Bihan (2013), this probability stands between 0.6 and 0.9.

where β_3 (expected to be null under the assumption of no information friction) and β_4 are the estimated coefficients. Equation (8) is estimated with a probit model and heteroskedasticity-robust standard errors.

Table 5 presents the results from the estimation of equation (8) for all categories of agents.²³ A robust result across our five categories of agents is that inflation conditional volatility significantly and positively affects the frequency of forecast revision. This result is consistent with information frictions generated by noisy information models. The magnitude with which inflation conditional volatility affects this frequency varies across our different categories. In particular, it affects the frequency of forecast revisions to a larger extent for central bankers and then participants to experiments than for other categories of agents. Households are those who are the least affected. Figure A3 in Appendix illustrates these findings by exhibiting the link between the conditional variance of inflation and whether individuals revise or not their forecasts. It confirms that the conditional variance is higher when agents revise their forecasts.

Table 5 - Drivers of the frequency of forecast revision - Eq. (8)

	(1)	(2)	(3)	(4)	(5)
	SPF	FOMC	Livingston	Michigan	Experiments
	P_i	P_i	P_i	P_i	P_i
CPI CondVar	0.095***	0.997**	0.081**	0.033***	0.353***
	[3.11]	[2.17]	[2.41]	[7.46]	[3.16]
CPI Level	-0.020	0.332*	0.417***	0.074***	0.050**
	[-0.65]	[1.70]	[10.62]	[10.25]	[2.09]
Constant	1.160***	1.290***	1.426***	0.692***	1.575***
	[45.23]	[5.39]	[48.54]	[142.29]	[41.35]
N	4089	241	5751	91390	40558
Pseudo R2	0.00	0.10	0.05	0.00	0.09

Note: t-statistics in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Parameters are estimated with Equation (8) using a probit model. The conditional variance of inflation is estimated based on a GARCH(1,1) model and the inflation level is introduced with a lag. The dependent variable is a dummy variable that takes the value 1 when a given individual revise his/her forecasts.

To summarize, there is some heterogeneity in the frequency of forecast revision across the five categories of agents. Households and then firms appear to be the categories that are the most subject to information frictions.

6. Concluding remarks

Understanding the strength of information frictions in inflation expectations within and across different categories of economic agents is of particular importance for central banks whose aim is to manage inflation expectations. In this paper, we have compared disagreement in expectations and the frequency of forecast revisions among five categories of agents: households, firms, professional forecasters, policymakers and participants to laboratory experiments. We have provided evidence of disagreement among all categories of agents, although there is a strong heterogeneity across categories: while policymakers, professional forecasters and participants to experiments exhibit low disagreement, firms and households show strong disagreement. Inflation shocks play a key role on disagreement within all categories, except for policymakers. We have also documented a heterogeneous frequency of forecast revision across the five categories of agents, with policymakers revising more frequently than participants to experiments, firms and professional forecasters, who

²³ Table A9 in Appendix provides a robustness check (to collinearity issues) of these findings when regressions are conducted separately for the conditional variance of inflation and the level of inflation.

themselves revising much more frequently than households. Inflation shocks are found to be an important driver of forecast revision, suggesting that a shock on inflation plays the role of an attention shock.

By considering five categories of agents, our paper gives a broader view on information frictions than the sole comparison between firms and households on which the literature recently focused. Avoiding the magnifying glass effect associated with the opposition between firms and households, we observe that there is more difference between firms and households on the one hand and policymakers, professional forecasters, and participants to experiments on the other, than between firms and households themselves.

Our results provide useful insights in terms of macroeconomic theory, external validity of macroeconomic experiments, and central bank communication. First, consistently with macroeconomic theories, our results suggest that the assumption of very low information friction is reasonable for central banks, but much less so for other categories of agents. This heterogeneity across different groups has to be accounted for in macroeconomic models, as different categories of agents may respond to shocks, monetary policy or fiscal policy in a different manner. Moreover, it is particularly important to account for heterogeneity within categories of agents in macroeconomic models, especially for households, but also professional forecasters.

Second, while crucial for laboratory experiments to be useful for policymakers, the issue of the external validity of experimental inflation expectations has not been much studied. Our results question the external validity of experimental inflation forecasts. In terms of disagreement, the behavior of participants to experiments is closer to that of central bankers. In terms of frequency of forecast revision, the behavior of participants to experiments is relatively close to that of firms but the impact of shocks on this frequency yields a behavior that is between that of policymakers and professional forecasters. Overall, although participants to experiments are supposed to play the role of firms (since the NK model formalizing the economy is directly fed with participants' expectations in the place of firms' expectations), their behavior seems more aligned with those of central bankers and professional forecasters. The reason for such a discrepancy might find its roots in the incentives (closely aligned with those of professional forecasters) participants to experiments are faced with, the salience of the information they receive and the NK data generating process. This study could be used by experimenters to design experiments that mimic real-world features. More precisely, experimenters should insure that when stating their expectations, participants to experiment reach, on average, the disagreement and forecast revision properties found on firms' data. Similarly, if they intend to reproduce the expectations of households, professional forecasters or central bankers, the properties found for each category using field data should be mimicked. This would help reproduce stylized facts in the laboratory as a precondition for simulating the impact of alternative policy measures. While we do not expect usual undergraduate student participants in experiments to achieve the same forecasting performance as highly qualified professional central bankers, a way to mimic their performance in the laboratory could be to provide an appropriate training to participants, in addition to providing the correct incentives.

Third, we can derive some policy prescriptions from our results, in terms of central banks' communication. Indeed, which category of agents is more likely to be reached by central banks' communication depends on the frequency of forecast revision of each category. Those categories of agents whose frequency of forecast revision reacts most to inflation dynamics are likely to respond more to communication about monetary policy. However, we also find that there is strong disagreement among each category, meaning that the information released by

the central bank may not reach all agents in the same manner. Targeted communication towards groups of economic agents presenting the same characteristics (e.g. for firms, in terms of size or sector) might represent a useful tool.

References

- Adam K. (2007). "Experimental Evidence on the Persistence of Output and Inflation", *Economic Journal*, 117, 603-635.
- Andrade P., R. Crump, S. Eusepi and E. Moench (2016). "Fundamental Disagreement", *Journal of Monetary Economics*, 83, 106-128.
- Andrade P. and H. Le Bihan (2013). "Inattentive professional forecasters", *Journal of Monetary Economics*, 60(8), 967-982.
- Andre P., C. Pizzinelli, C. Roth and J. Wohlfart (2021). "Subjective Models of the Macroeconomy: Evidence from Experts and Representative Samples", *CESifo Working Paper*, No. 7850.
- Angeletos G.-M., Z. Huo and K. Sastry (2021). "Imperfect Macroeconomic Expectations: Evidence and Theory", *NBER Macroeconomics Annual 2020*, 35(1), 1-86.
- Angeletos G.-M. and C. Lian, (2016). "Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination", *Handbook of Macroeconomics 2*.
- Armantier O., W. Bruine de Bruin, G. Topa, W. van der Klaauw and B. Zafar (2015). "Inflation Expectations and Behavior: Do Survey Respondents Act on Their Beliefs?", *International Economic Review*, 56, 505-536.
- Armantier O., W. van der Klauw, G. Topa and B. Zafar (2016). "The Price is Right: Updating Inflation Expectations in a Randomized Price Information Experiment", *Review of Economics and Statistics*, 98(3), 503-523.
- Carlson J. and N. Valev (2003). "Sources of dispersion in consumer inflation forecasts", *Applied Economic Letters*, 10, 77-81.
- 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.
- Carroll C.D. (2003). "Macroeconomic Expectations of Households and Professional Forecasters", *Quarterly Journal of Economics*, 118(1), 269-298.
- Clements M. (2020). "Forecaster Efficiency, Accuracy and Disagreement: Evidence using Individual-Level Survey Data", *Journal of Money, Credit and Banking*, forthcoming.
- Clements M. (2021). "Rounding behaviour of professional macro-forecasters", *International Journal of Forecasting*, forthcoming.
- Coibion O. and Y. Gorodnichenko (2012). "What Can Survey Forecasts Tell Us About Informational Rigidities?", *Journal of Political Economy*, 120(1), 116-159.
- Coibion O. and Y. Gorodnichenko (2015). "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts", *American Economic Review*, 105(8), 2644-2678.
- Coibion O., Y. Gorodnichenko and S. Kumar (2018a). "How Do Firms Form Their Expectations? New Survey Evidence", *American Economic Review*, 108(9), 2671-2713.
- Coibion O., Y. Gorodnichenko, S. Kumar and J. Ryngaert (2018b). "Do You Know That I Know That You Know...? Higher-order Beliefs in Survey Data", *NBER Working Paper*, No. 24987.
- Cornand C. and F. Heinemann (2014). "Experiments on monetary policy and central banking", in "Experiments in Macroeconomics", J. Duffy (ed.), *Research in Experimental Economics* 17, Emerald Group Publishing.

- Cornand C. and F. Heinemann (2019). "Experiments on macroeconomics: methods and applications", in *Handbook of Research Methods and Applications in Experimental Economics*, A. Schram and A. Ule (eds.), Edward Elgar Publishing, 269-294.
- Cornand C. and P. Hubert (2020). "On the external validity of experimental inflation forecasts: A comparison with five categories of field expectations", *Journal of Economic Dynamics and Control*, 110, 103746.
- Cornand C. and C. K. M'baye (2018a). "Does inflation targeting matter? An experimental investigation", *Macroeconomic Dynamics*, 22(2), 362-401.
- Cornand C. and C. K. M'baye (2018b). "Band or Point Target? An Experimental Study", *Journal of Economic Interaction and Coordination*, 13(2), 283-309.
- Cukierman A. and P. Wachtel (1979). "Differential inflationary expectations and the variability of the rate of inflation: theory and evidence", *American Economic Review*, 69, 595-609.
- Dovern, J., U. Fritsche and J. Slacalek (2012). "Disagreement among Forecasters in G7 Countries", *Review of Economics and Statistics*, 94, 1081-96.
- Duffy J. (2008). "Experimental Macroeconomics", in S. Durlauf and L. Blume (Eds.), *New Palgrave Dictionary of Economics*, 2nd Ed., New York: Palgrave Macmillan.
- Duffy J. (2016). "Macroeconomics: A Survey of Laboratory Research", in J. Kagel and A. Roth (eds.), *Handbook of Experimental Economics 2*, Princeton University Press, 1-90.
- Hommes C. (2011). "The Heterogeneous Expectations Hypothesis: Some Evidence from the Lab", *Journal of Economic Dynamics and Control*, 35(1), 1-24.
- Hommes C. (2021). "Behavioral and Experimental Macroeconomics and Policy Analysis: A Complex Systems Approach", *Journal of Economic Literature*, 59(1), 149-219.
- Hommes C., D. Massaro and M. Weber (2019). "Monetary Policy under Behavioral Expectations: Theory and Experiment", *European Economic Review*, 118, 193-212.
- Kryvtsov O. and L. Petersen (2021). "Central Bank Communication That Works: Lessons from Lab Experiments", *Journal of Monetary Economics*, 117, 760-780.
- Link S., A. Peichl, C. Roth and J. Wohlfart (2021). "Information Frictions Among Firms and Households", *CESifo Working Paper*, No. 8969.
- Mankiw G. and R. Reis (2002). "Sticky information versus sticky prices: A proposal to replace the new Keynesian Phillips curve", *Quarterly Journal of Economics*, 117(4), 1295-328.
- Mackowiak B. and M. Wiederholt (2009). "Optimal Sticky Prices under Rational Inattention", *American Economic Review*, 99(3), 769-803.
- Mankiw G., R. Reis, and J. Wolfers (2004). "Disagreement about Inflation Expectations", in *NBER Macroeconomics Annual 2003*, M. Gertler and K. Rogoff (eds.), 209-248, Cambridge, MA: MIT Press.
- Petersen L. (2014). "Forecast Error Information and Heterogeneous Expectations in Learning to Forecast Macroeconomic Experiments", in "Experiments in Macroeconomics", J. Duffy (ed.), *Research in Experimental Economics*, 17, Emerald Group Publishing.
- Pfajfar D. and Žakelj B. (2018). "Inflation Expectations and Monetary Policy Design: Evidence from the Laboratory", *Macroeconomic Dynamics*, 22(4), 1035-1075.
- Sims C. (2003). "Implications of Rational Inattention", *Journal of Monetary Economics*, 50, 665-690.
- Souleles N. (2004). "Expectations, heterogeneous forecast errors, and consumption: micro evidence from the Michigan Consumer Sentiment Surveys", *Journal of Money, Credit and Banking*, 36, 39-72.

Appendix

A. Experimental designs

Table A1 is taken and adapted²⁴ from Cornand and Hubert (2020). It summarizes the experimental designs of the five considered experimental papers, in terms of data-generating processes, tasks of subjects, incentives, information sets and characteristics of subjects.

Table A1 - Characteristics of experimental designs

	Data generating process			Subjects' task	Incentives in ECU F stands for the absolute forecast error	Information set at date t : history of variables up to period $t-1$	Subjects: undergraduate students	
	Model	Output gap expectations	Shocks					
PZ	Qualitative description of the same NK model structure with relatively standard and comparable parameter values	Pre-programmed naïve	AR-1 process	Inflation forecasts in t for $t+1$	$\max \left\{ \frac{100}{1+F} - 20, 0 \right\}$	Inflation, Output gap, Interest rate, profits from forecasts	Group size: 9 Majors: economics, business	
CMa			iid distributions		$\max \left\{ \frac{160}{1+F} - 40, 0 \right\}$		In most treatments, information about the inflation target	Group size: 6 Majors: economics, business
CMb			variances may differ					Group size: 6 Majors: engineering, business
HMW		Part of subjects' task	Inflation and output gap forecasts in t for $t+1$	$\frac{100}{1+F}$			Group size: 6 Majors: economics, business	
Petersen							Information on forecast error, Period t 's interest rate and shock to the natural rate of interest, Expected shock size in $t+1$	Group size: 9 Students and non-students.

²⁴ We do not consider the aggregate data from Adam and we added half of the data from Petersen (2014).

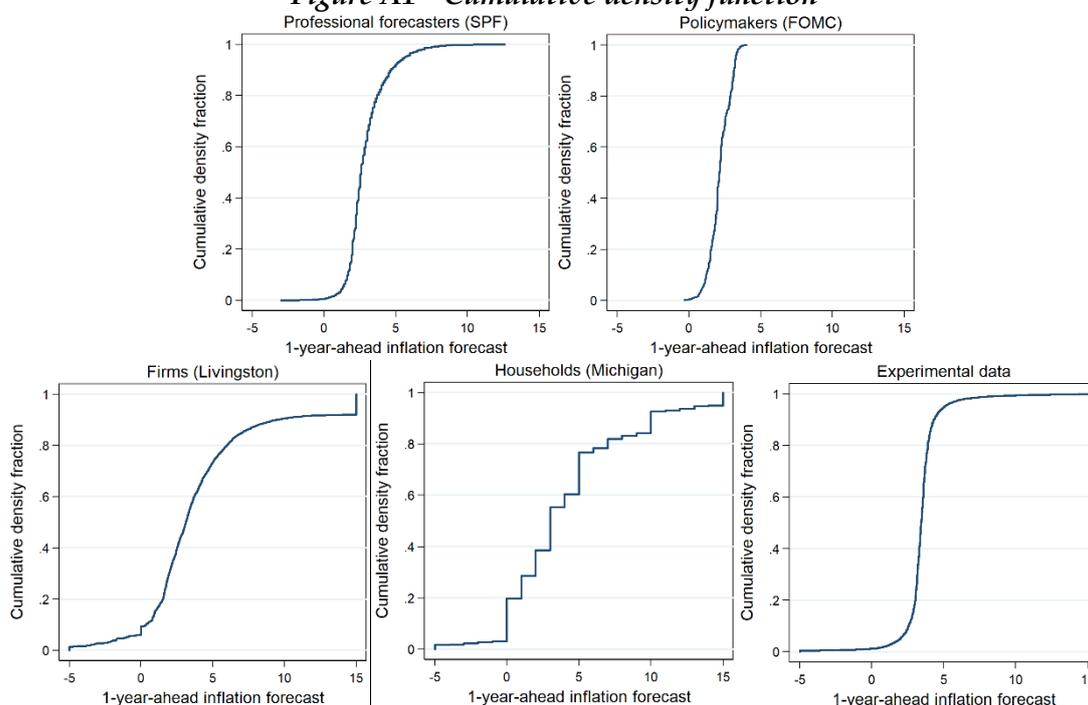
B. Further evidence

Table A2 - Alternative forecasts: Descriptive statistics about disagreement - Eq. (3), (4), (5)

		IQR	IDR	SD	Mean	Median	IQR_N
SPF	Fixed event	1.20	2.90	1.37	2.88	2.54	0.42
FOMC	Current year	1.25	2.00	0.81	2.27	2.25	0.55
FOMC	Next year	0.80	1.70	0.66	2.09	2.00	0.38
Livingston	Fixed event	2.09	4.65	2.05	3.54	2.97	0.59
Michigan	Lag forecast	5.00	18.00	26.25	10.70	3.00	0.47

Note: These statistics are computed over the full sample for each of the five datasets. The IQR is the Inter Quartile Ratio. The IDR is the Inter Decile Ratio, the distance between the 90th and 10th percentiles. The SD is the standard deviation. The IQR_N is a normalized IQR by the level of inflation forecasts and corresponds to the ratio of the IQR divided by the mean.

Figure A1 - Cumulative density function



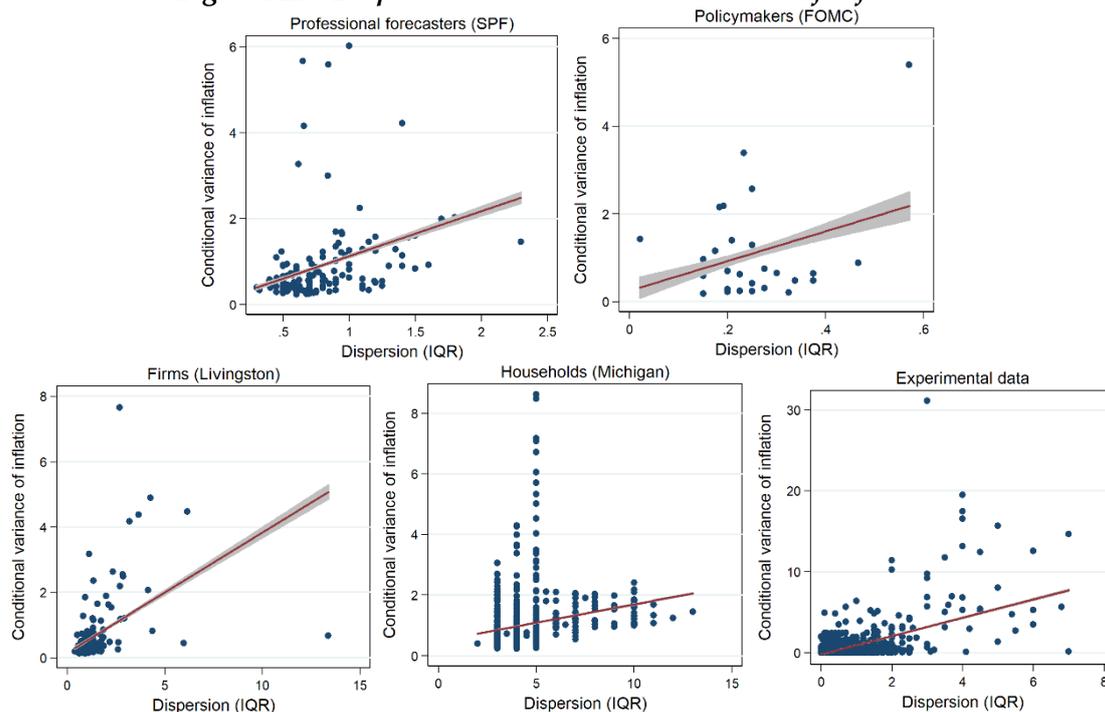
Note: These figures show the cumulative density function of all individual inflation forecasts for each dataset.

Table A3 - Estimation of Eq. (4) for each covariate separately

	SPF		FOMC		Livingston		Michigan		Experiments	
	IQR_t	IQR_t	IQR_t	IQR_t	IQR_t	IQR_t	IQR_t	IQR_t	IQR_t	IQR_t
CPI CondVar	0.105***		0.029		0.541***		0.375***		0.297***	
	[2.86]		[1.15]		[3.93]		[5.19]		[6.15]	
CPI Level		0.182***		-0.016		0.343*		1.455***		-0.015
		[6.86]		[-0.68]		[1.93]		[17.72]		[-0.33]
Constant	0.776***	0.770***	0.243***	0.250***	1.422***	1.450***	4.496***	4.673***	0.285***	0.400***
	[31.57]	[36.17]	[13.22]	[12.38]	[13.74]	[12.42]	[59.46]	[83.55]	[21.74]	[26.23]
N	156	156	28	28	145	145	509	509	5988	5882
R2	0.10	0.28	0.10	0.03	0.17	0.06	0.05	0.55	0.30	0.01

Note: t-statistics in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Parameters are estimated with a reduced version of Equation (6) using OLS. The conditional variance of inflation is estimated based on a GARCH(1,1) model and the inflation level is introduced with a lag. The dependent variable is the IQR computed per period for each dataset.

Figure A2 - Dispersion & Conditional variance of inflation



Note: These scatterplot figures show the link between the conditional variance of inflation at a given date and the dispersion (measured by the IQR) of inflation forecasts for the same date. The regression is plotted with a 95% confidence interval.

Table A4 – IDR measure as a dependent variable in Eq. (6)

	(1)	(2)	(3)	(4)	(5)
	SPF	FOMC	Livingston	Michigan	Experiments
	IDR _t				
CPI CondVar	0.392***	0.051	1.226***	0.950***	0.781***
	[4.54]	[1.05]	[6.40]	[9.09]	[7.91]
CPI Level	0.626***	-0.034	0.646***	2.784***	-0.025
	[3.52]	[-0.74]	[2.68]	[21.74]	[-0.30]
Constant	1.757***	0.582***	3.189***	9.947***	2.867***
	[28.75]	[16.45]	[21.38]	[118.27]	[8.07]
N	156	28	145	509	5882
R2	0.44	0.19	0.45	0.68	0.14

Note: t-statistics in brackets. * p < 0.10, ** p < 0.05, *** p < 0.01. Parameters are estimated with Equation (6) using OLS. The conditional variance of inflation is estimated based on a GARCH(1,1) model and the inflation level is introduced with a lag. The dependent variable is the IDR computed per period for each dataset.

Table A5 – SD measure as a dependent variable in Eq. (6)

	(1)	(2)	(3)	(4)	(5)
	SPF	FOMC	Livingston	Michigan	Experiments
	SD _t	SD _t	SD _t	SD _t	SD _t
CPI CondVar	0.172*** [5.03]	0.021 [0.91]	0.453*** [4.57]	0.248*** [4.09]	0.260*** [7.38]
CPI Level	0.232*** [6.20]	-0.007 [-0.43]	0.363*** [3.85]	1.636*** [32.66]	-0.011 [-0.40]
Constant	0.756*** [34.70]	0.219*** [15.59]	1.528*** [23.28]	5.375*** [99.08]	0.928*** [7.95]
N	156	28	145	509	5882
R2	0.49	0.14	0.42	0.59	0.13

Note: t-statistics in brackets. * p < 0.10, ** p < 0.05, *** p < 0.01. Parameters are estimated with Equation (6) using OLS. The conditional variance of inflation is estimated based on a GARCH(1,1) model and the inflation level is introduced with a lag. The dependent variable is the SD computed per period for each dataset.

Table A6 – Common time-sample for field data - Eq. (6)

	(1)	(2)	(3)	(4)	(5)
	SPF	FOMC	Livingston	Michigan	Experiments
	IQR _t	IQR _t	IQR _t	IQR _t	IQR _t
CPI CondVar	0.100*** [3.99]	0.031 [0.94]	0.145*** [4.15]	0.102*** [3.03]	.
CPI Level	0.076* [1.76]	0.002 [0.09]	0.118 [1.51]	-0.193 [-0.92]	.
Constant	0.687*** [40.20]	0.243*** [13.72]	0.938*** [19.10]	3.938*** [31.52]	.
N	72	28	36	216	.
R2	0.32	0.10	0.13	0.05	.

Note: t-statistics in brackets. * p < 0.10, ** p < 0.05, *** p < 0.01. Parameters are estimated with Equation (6) using OLS. The estimation is performed on common sample for all datasets, from 1992m1 to 2009m11. The conditional variance of inflation is estimated based on a GARCH(1,1) model and the inflation level is introduced with a lag. The dependent variable is the IQR computed per period for each dataset.

Table A7 – Controlling for the level of forecasts - Eq. (6)

	(1)	(2)	(3)	(4)	(5)
	SPF	FOMC	Livingston	Michigan	Experiments
	IQR _t	IQR _t	IQR _t	IQR _t	IQR _t
CPI CondVar	0.114*** [3.01]	-0.003 [-0.10]	0.482*** [3.38]	0.369*** [5.49]	0.297*** [6.16]
Forecast Level	0.103*** [6.68]	-0.057** [-2.54]	-0.061* [-1.80]	0.042*** [2.92]	0.001 [0.07]
Constant	0.417*** [9.30]	0.372*** [6.41]	1.609*** [9.47]	4.284*** [48.77]	0.283*** [6.66]
N	143	28	135	509	5988
R2	0.37	0.21	0.19	0.08	0.30

Note: t-statistics in brackets. * p < 0.10, ** p < 0.05, *** p < 0.01. Parameters are estimated with Equation (6) using OLS. The conditional variance of inflation is estimated based on a GARCH(1,1) model and the forecast level is introduced with a lag. The dependent variable is the IQR computed per period for each dataset.

Table A8 – Alternative forecasts: Descriptive statistics about frequency of forecast revisions

		P_{Mean}	P_{SD}	N	Mean Rev.	SD Rev.	N_{Rev}
SPF	Fixed event	0.85	0.36	3,265	0.43	0.62	3,307
FOMC	Current year	0.84	0.37	643	0.52	0.43	560
FOMC	Next year	0.75	0.44	158	0.30	0.26	173
Livingston	Fixed event	0.91	0.28	2,238	0.76	1.00	2,044
Livingston	No rounding	0.98	0.14	5,751	1.45	1.92	5,204
Michigan	Lag forecast	0.71	0.45	1,274	5.02	5.90	909

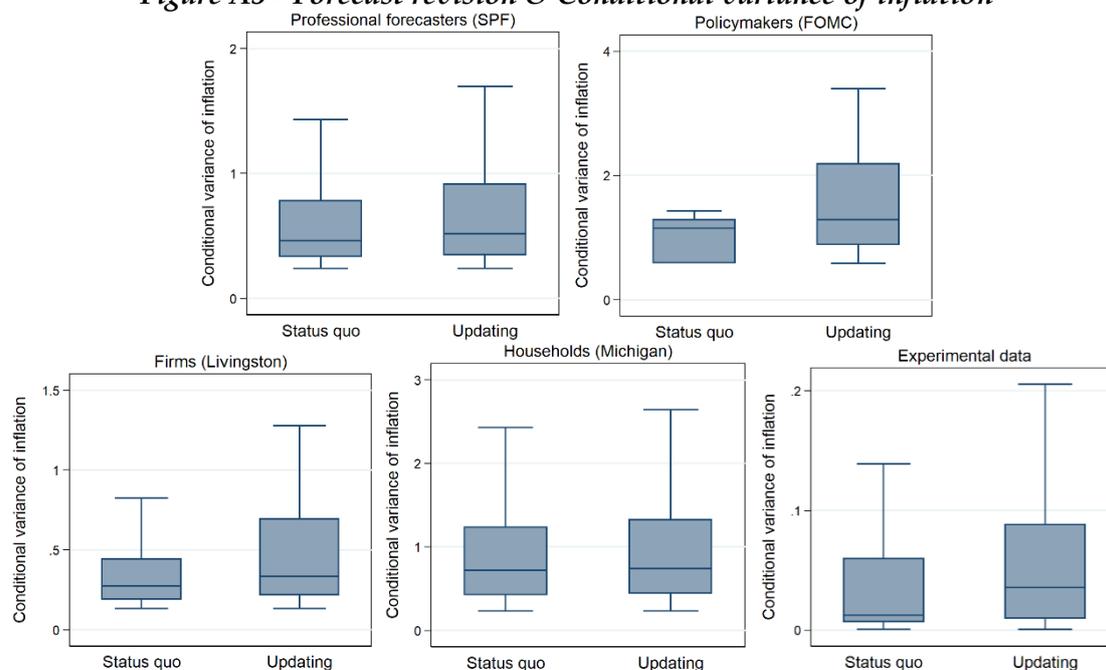
Note: These statistics are computed over the subsample N for which we observe 2 consecutive forecasts of a same individual for each of the five datasets. The P_{Mean} is the average frequency of revisions for each dataset. The P_{SD} is the standard deviation of the frequency of revisions. The Mean Rev. is the average magnitude of the revision for individuals who revised their forecasts (the subsample N_{Rev}) while the SD Rev. is the standard deviation of the magnitude of these revisions.

Table A9 – Estimation of Eq. (8) for each covariate separately

	SPF		FOMC		Livingston		Michigan		Experiments	
	P_i	P_i	P_i	P_i	P_i	P_i	P_i	P_i	P_i	P_i
CPI CondVar	0.096***		0.552**		0.126***		0.032***		0.336***	
	[3.17]		[2.34]		[3.67]		[7.11]		[3.36]	
CPI Level		-0.028		-0.107***		0.414***		0.072***		0.026**
		[-0.96]		[-3.48]		[11.45]		[10.12]		[2.55]
Constant	1.163***	1.153***	1.558***	1.801***	1.320***	1.415***	0.674***	0.691***	1.576***	1.612***
	[45.70]	[45.40]	[9.26]	[12.07]	[56.42]	[51.95]	[149.41]	[142.47]	[42.74]	[44.63]
N	4089	4089	241	241	5751	5751	91390	91390	41278	40559
Pseudo R2	0.00	0.00	0.07	0.01	0.01	0.05	0.00	0.00	0.09	0.09

Note: t-statistics in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Parameters are estimated with a reduced version of Equation (8) using a probit model. The conditional variance of inflation is estimated based on a GARCH(1,1) model and the inflation level is introduced with a lag. The dependent variable is a dummy variable that takes the value 1 when a given individual revise his/her forecasts.

Figure A3 – Forecast revision & Conditional variance of inflation



Note: These boxplot figures show the distribution of the conditional variance of inflation when individuals revise or not their inflation forecasts.