Intrinsic Expectations Persistence
Evidence from professional and household survey expectations

Jeff Fuhrer

Abstract

This paper examines the expectations behavior of individual responses in the Survey of Professional Forecasters, the University of Michigan's Survey Research Center survey of consumers, and the ECB Survey of Professional Forecasters. The paper finds that the most robust feature of all of these expectations measures is that respondents inefficiently revise their forecasts so as to very gradually incorporate new information. A key variable that is used to update forecasts is the lagged central tendency of expectations in the survey. This result holds for all of the surveys at all forecast horizons for inflation, unemployment, short-term interest rates, and real growth, and is quantitatively and statistically significant. It is robust to the inclusion of all of the real-time information available in these surveys. The paper examines the relationship between these results and those of Coibion and Gorodnichenko (2015), which suggest that aggregate surveys conform with a key prediction of the sticky information model of Mankiw and Reis (2002) and the noisy information model of Maćkowiak and Wiederholt (2009). This paper finds considerably less coherence with these models in the micro data. The paper also provides evidence that distinguishes this behavior from learning, suggesting that inefficient use of previous forecasts and gradual revision toward the central tendency are much more important quantitatively than least-squares learning in these expectations measures. Finally, this empirical regularity bears important implications for macroeconomic dynamics, as illustrated in the last section of the paper, as it provides a micro-based foundation for an earlier paper's finding that intrinsic persistence in expectations may be a key source of macroeconomic persistence (Fuhrer 2017).

JEL Classifications: E32, E52

Keywords: Intrinsic persistence, rational expectations, survey expectations

Jeff Fuhrer is an executive vice president and senior policy advisor at the Federal Reserve Bank of Boston. His e-mail address is jeff.fuhrer@bos.frb.org. Thanks to Susanto Basu, Olivier Coibion, Chris Foote, Yuriy Gorodnichenko, Ricardo Nunes, and Giovanni Olivei for helpful comments.

The views expressed herein are those of the author and do not indicate concurrence by the Federal Reserve Bank of Boston, or by the principals of the Board of Governors, or the Federal Reserve System.

This version: June 30, 2017

Expectations lie at the heart of all current macroeconomic models. Decisions about prices, capital goods, consumer durable goods, housing, life-cycle savings choices and monetary policy all inherently depend on expectations about future economic conditions. The idea that economic actors "look forward" or think about the future in making some economic decisions seems relatively uncontroversial. Exactly how they peer into the future is much less clear.

The rational expectations paradigm has been used widely in macroeconomic models for decades, and has served the discipline well due to its elegance and computational simplicity. However, few believe that the theory of rational expectations is to be taken literally. Whether it serves as a reasonable approximation to the expectations-formation behavior of firms and households is an empirical matter, and likely depends on the economic question at hand, on the agents in question, and on the economic circumstances. In tranquil times, many financial market participants likely use information quite efficiently. In their own domains, successful firms likely know enough about their environment to make near-rational decisions about inputs, pricing, and market strategy. It may be the case that in these instances, rational expectations works fairly well as a description of forward-looking behavior (although this too remains an empirical question).

But evidence is mounting that suggests that rational expectations may not be the best assumption to embed in macroeconomic models (see, for example, Fuhrer (2017), Trehan (2015), Fuster, Hebert and Laibson (2012), Adam and Padula (2011), and Roberts (1997)). The addition of many "bells and whistles" to DSGE models (habits, price indexation, complicated adjustment costs) as well as the ubiquitous presence of highly autocorrelated structural shocks, may be construed as evidence that these models are misspecified, perhaps due to the restrictions imposed by the rational expectations assumption. In addition, a number of papers have shown that the rational expectations implied by such models deviate significantly from measured expectations (Del Negro and Eusepi (2010) is one notable example). This finding could mean that the models are misspecified, even though rational expectations remains the valid assumption. Or it could be that the basic model structures are reasonable, but the expectations assumption causes the models to make strongly counterfactual predictions.

A number of papers have explored alternative expectations assumptions and their implications for economic outcomes, in both theoretical and empirical settings. A leading example is learning: see Adam (2005), the many papers of Evans and Honkapohja and their 2001 book, Milani (2007), Orphanides and Williams (2005), and Slobodyan and Wouters (2012). Milani (2007) shows that the introduction of adaptive learning significantly reduces the dependence of a particular DSGE

model on habit formation and price indexation to explain the persistence of macroeconomic time series. Slobodyan and Wouters (2012) find a notable reduction in the persistence of the estimated shocks that drive wages and prices; they also note that the expectations based on the "small forecasting models" in their paper bear a close resemblance to survey expectations. Others have posited models of information frictions to better explain macroeconomic dynamics, including the "sticky information" model of Mankiw and Reis (2002), and the "noisy information" models motivated by Sims' (2003, 2006) work on rational inattention, and implemented in Maćkowiak and Wiederholt (2009), for example.

It is striking that relatively few authors have examined in detail the expectations behavior of individual economic agents. Most of the empirical papers cited above use aggregated measures of expectations from available surveys and (in fewer cases) from financial asset prices. Exceptions include empirical work by Crowe (2010), Andrade and Le Bihan (2013), Paloviita and Viren (2013) and a vast theoretical literature that emphasizes the role of higher-order expectations (see especially Frydman and Phelps (2013) and the papers contained and cited therein). Gennaioli, Ma and Shleifer (2016) document the characteristics of surveys of CFO's expectations of earnings growth. They find that they are not well proxied by Tobin's Q or discount rates, that they are not rational (in the sense that they make errors that are predictable using information available to the CFOs at the time of prediciton), and that they do well in explaining both investment plans and realized investment. But few have attempted to characterize the underlying behaviors in the micro-data from the oft-cited aggregate surveys from the Survey of Professional Forecasters (SPF) and the University of Michigan's Survey Research Center survey of consumers.

This paper examines a rich set of micro-data evidence on the expectations behavior of firms and households, both in the U.S. and in the Euro Area. The paper is motivated by the observation that aggregated expectations from the SPF appear to improve significantly the performance of standard dynamic macroeconomic models (Fuhrer 2017). While that paper provides an internally consistent way of describing expectations behavior, it does not answer the fundamental question of why survey expectations appear to account for a significant portion of the persistence found in macroeconomic data. That is, apart from the theoretical mechanisms that commonly generate persistence in macroeconomic models (for example, persistence in marginal costs, habit formation, price indexation, costs of adjustment), expectations appear to add intrinsic persistence above and beyond these mechanisms, and in so doing, account for a large fraction of the persistence observed in macroeconomic time series.

To be a bit more precise about the macroeconomic observation, consider an inflation Euler equation that is widely used in many DSGE models:

$$\pi_t = (\beta - \omega)E_t\pi_{t+1} + \omega\pi_{t-1} + \gamma s_t + \varepsilon_t; \varepsilon_t = \frac{\eta_t}{1-\rho L},$$

where π is inflation, *s* is marginal cost, β is the discount rate, ε_t is the serially correlated shock to the equation with autocorrelation parameter ρ and *iid* innovation η_t, and E is understood to be the rational or model-consistent expectation of the next period's inflation rate. This Euler equation may be derived from a Calvo pricing model in which a fraction ω of price-setters who do not get the Calvo draw in period *t* choose to index their current prices to last period's inflation rate. A number of authors have found fairly sizable and significant estimates of ω in estimated versions of this equation (Christiano, Eichenbaum, and Evans (2005), Smets and Wouters (2007)). In addition, it is quite common to estimate sizable values for ρ, the parameter indexing the degree of autocorrelation in the structural shock ε_t.

However, if one instead uses survey measures of expectations in this equation—for example, the median forecast of inflation for period *t+1* from the Survey of Professional Forecasters—one finds that the data prefer an estimated value for ω that is much smaller and typically not statistically significantly different from zero. In addition, the estimated autocorrelations of the error term ε_t, while sizable in rational expectations implementations of the equation, are much smaller and not significantly different from zero. The same is true for other key equations in standard DSGE models: Structural add-ons that induce lagged dependent variables (habits in consumption, for example) diminish greatly in importance, and autocorrelated structural shocks become much less, if at all, autocorrelated.

What is happening in the estimates of these models with survey expectations? The expectations themselves have incorporated some inertia that was previously proxied by indexation, habits, and/or autocorrelated shock processes. For inflation, the expectations add persistence above and beyond the persistence that inflation inherits from the marginal cost process. For habits, the expectations capture much of the sluggish adjustment of consumption growth to shocks that were previously proxied by lagged consumption.[1] While Fuhrer (2017) documents this finding with

[1] Fuhrer (2000) is one of the earliest papers to document the strong empirical significance of habit formation in monetary policy models.

aggregate data, this paper aims to understand the underlying expectation behaviors that give rise to this kind of persistence in measures of expectations.

The paper uses the individual responses in the SPF, the ESPF and the Michigan Survey of Consumers to better understand the sources of inertia in expectations data. The SPF comprises a few thousand observations on a few hundred firms over the past 30 to 45 years (depending on the variable studied), while the Michigan survey contains over 500,000 observations on tens of thousands of households since 1978. The ESPF begins in 1999, surveys about 100 firms and like the SPF contains several thousand observations per expectations variable. The structures of the datasets differ: Whereas many firms in the SPF and ESPF participate in the survey for many years, if not decades, the Michigan survey samples a household once and then, for a subset of respondents, once again, six months later. The ability to observe individual respondents' forecasts over time is an advantage for the questions this paper aims to investigate. While both surveys afford such across-time comparisons to a certain extent, the SPF and the ESPF are much richer in this dimension.

Although firms' and households' expectations differ in some respects, they share one key feature. The forecast revisions exhibit what appears to be a significant inefficiency that bears important implications for macroeconomic dynamics: while forecasters revise forecasts in response to new information, such as that revealed in the lagged central tendency of forecasts (and other variables), they appear to inefficiently slow revisions in the face of new information by linking forecast revisions to their own forecasts for the same variable made in the previous period.[2]

Two possible rationales for this observation derive from the models of sticky or noisy information mentioned above. In these frameworks, forecast revisions could be linked to past forecasts, either because forecasters have not yet updated their information sets, or because they reduce the weight on news received, because it is not clear how much signal is reflected in the news. Coibion and Gorodnichenko (2015) provide tests of aggregate expectations that appear to generally conform with these models. We will examine implications of these models below, and conclude that the aggregate results in Coibion and Gorodnichenko are strongly contradicted in the micro data.[3]

One variable that all forecasters appear to incorporate in their revisions is the lagged median of individual forecasts. This information is not available to forecasters at time *t-1*, so using it to update time *t* forecasts is entirely reasonable, as it serves as a handy aggregator of diverse views on

[2] Earlier papers that examined the properties of forecast revisions for limited sets of forecasters include Berger and Krane (1985) and Nordhaus (1987).

[3] Coibion and Gorodnichenko (2015) are careful to point out that their key test—that forecast errors should be related only to forecast revisions—holds only on average across forecasters.

the variables of interest. This result is related to but quite distinct from the "epidemiological" phenomenon found in Carroll (2003), whereby in the aggregate, household forecasts are found to converge over time to the forecasts of professionals. Here, the individual forecasters within the cross-section of household or professional forecasts link their forecasts to previously observed aggregate forecasts from the same sector.

However, updating the forecasts to reflect new information, even if the information was generated in period *t-1,* does not explain the tendency to revise forecasts gradually in the presence of this new information, inefficiently linking the revision in one's *t*-period forecast for a given horizon to the *t-1*-period forecast for the same horizon. As suggested above, one model that might imply such sluggish updating is Mankiw and Reis's (2002) sticky information framework. Agents in that framework either (a) update their information and form a rational forecast, or (b) do not update their forecasts at all. We will show that it is uncommon for professional forecasters not to update their information sets from quarter to quarter. Rather, in the presence of updated information, they update inefficiently, slowing the incorporation of new information into forecasts by anchoring the revision to previous forecasts. Households may well update infrequently, but they are similarly shown to update quite inefficiently. The noisy information model bears similar implications, and is similarly rejected in the micro data.

Another obvious input to individual forecasts is the lagged realization of the variable of interest. It will be shown that the micro data exhibit a much stronger response to the lagged viewpoint forecast than to any of the lagged (real-time) actual data. In fact, inefficient adjustment to new information will be shown to be a much stronger feature of the data than classic adaptive least-squares learning, which generally takes the form of updated OLS projections of expectations on lagged observable data. To this point, the paper provides more formal evidence comparing least-squares learning and intrinsic expectations persistence, and finds the latter to be both quantitatively and statistically much more important in determining expectations behavior.

Of course, such sluggish anchoring of individual forecasts to lagged aggregate information imparts additional persistence to the expectations, beyond the persistence that would otherwise be a component of the variables they wish to forecast. Thus, the pervasiveness of this kind of expectations behavior may bear important implications for explaining the persistence of aggregate macro time series. The rational expectations assumption can build into expectations only those characteristics that the model implies for all variables. The empirical results in this paper suggest that actual expectations add significant persistence of their own to the system. The final section of the

paper explores the extent to which such an expectations mechanism affects the dynamics of key macroeconomic variables in a simple DSGE model.

While much work remains to be done in characterizing such expectations behavior from a theoretical perspective, the implications of these findings for macroeconomic modeling are significant. If expectations at the micro level are indeed persistent in the way described above—above and beyond the persistence of the variables they use to forecast inflation—then expectations will add their own "intrinsic persistence," in the sense articulated in the context of standard inflation models in Fuhrer (2006, 2011). It will therefore be reasonable to assume that some portion of the persistence observed in key macroeconomic time series arises from this "intrinsic expectations persistence," a finding that is consistent with the macro-survey findings referenced above. This suggests that other sources of persistence that are common in DSGE models and the like may be (at least in part) an artifact of the misspecification of expectations in those models. This assumption is tested in the empirical work in Fuhrer (2017), and illustrated in the context of stylized models below.

The paper concludes by providing some suggestive macro-modeling exercises that highlight the role that persistent expectations can play in the macroeconomy.

1. Evidence from professional forecasters

We begin by examining the expectations formed by the (presumably) more-sophisticated actors in the economy, namely those who make their living forecasting macroeconomic aggregates such as unemployment, inflation, interest rates and growth. To be sure, not all of the firms surveyed in the SPF or the ESPF are large firms with extensive staff and a long track record of forecasting and forecast model-building. However, as compared to the expertise that is likely embodied in the average household, it seems reasonable to assume that this group of forecasters is relatively sophisticated.

Tables 1a and 1b provide some summary statistics describing key features of the SPF and ESPF samples. Figure 1shows the duration and timing of each forecaster's participation in the SPF survey from 1981:Q3 to the most recent survey in the sample.[4] A few forecasters are in the survey for two decades or more; quite a few participate for only a few years. The mean and median forecasts for selected years suggest that the distribution of forecasts is not strongly skewed in one direction or the other. The sample is roughly evenly split between financial and nonfinancial firms.

[4] We focus on this sample as it represents the period over which the consumer price index (CPI) is collected for the survey. This variable has the advantage that the survey collects both its lagged values and long-term forecasts of it.

Others have written about the forecasting accuracy of the SPF and other forecasts, although that is not the focus of this paper (see, for example, Batchelor (1986), Bryan and Gavin (1986), Mehra (2002), and Thomas (1999)). For more details on the SPF, Michigan and ESPF data, see the links to the sources in Appendix A.[5] For reference, Table 16 provides the results of efficiency tests for the individual forecasts, using real-time actual data to compute forecast errors, and testing the efficiency of these errors against real-time data available to the forecasters, as reported in the SPF forecast data set. It is not difficult to reject the null of efficiency, but we will examine in more detail a particularly striking form of inefficiency in what follows.

Properties of individual SPF forecasts

The first set of results examines the correlations among individual inflation forecasts made in period *t*, the forecasters' idiosyncratic (real-time) estimates of lagged inflation, measures of the previous period's central tendency of the SPF forecast for the same variable, and lagged individual forecasts (both lagged viewpoint date for the *t+1* forecast and the lagged one-period-ahead forecast).[6] Table 2 presents results from the first set of test regressions, which take the general form

$$\pi^i_{t+1,t} = a\pi^i_{t-1} + b\pi^i_{t+1,t-1} + cC(\pi^i_{t+1,t-1}) + d\pi^i_{t,t-1} + eZ^i_t + \delta_i + \varepsilon^i_t, \qquad (1.1)$$

where $\pi^i_{t+1,t}$ is the i^{th} forecaster's forecast of consumer price index (CPI) inflation for period *t+1* made in period *t*; π^i_{t-1} is the i^{th} forecaster's estimate of lagged inflation as of period *t*, $\pi^i_{t+1,t-1}$ is the i^{th} forecaster's forecast for the same horizon *t+1* made last period (*t-1*), $\pi^i_{t,t-1}$ is the i^{th} forecaster's forecast for period *t* made in period *t-1*, $C(\pi^{SPF}_{t+k,t-1})$ is a measure of the lagged central tendency of forecasts for the same variable for period *t+1* using the previous period's information set, here taken to be the median of the forecasts, Z^i_t is a vector of other forecaster-specific variables, which includes real-time individual estimates of lagged unemployment, output growth, and the Treasury bill rate, and δ_i denotes forecaster-specific fixed effects.[7] Standard errors are corrected for

[5] For many applications, including price-setting and investment behavior, it would be more appropriate to investigate the properties of firms' expectations. However, a consistent dataset that includes firms' numerical expectations of key macroeconomic variables does not exist for the United States. See Coibion, Gorodnichenko, and Kumar (2015) for an analysis of a set of New Zealand firms' expectations.

[6] Observations later in the sample show a considerably smaller dispersion of estimates of lagged inflation.

[7] We consider other proxies for the lagged central tendency of forecasts when we estimate revision regressions below.

heteroskedasticity, autocorrelation, and correlation among panels using the method developed in Driscoll and Kraay (1998).[8]

Note that one can think of regression (1.1) as embedding two types of change regressions. First, one can subtract the *t-1* forecast for period *t+1* from the both sides of the equation to obtain $(\pi^{i}_{t+1,t} - \pi^{i}_{t+1,t-1})$ on the left-hand side, the revision to the *t+1* forecast from one viewpoint date to the next. Second, one can subtract the *t-1* forecast for period *t* from the left-hand side to obtain $(\pi^{i}_{t+1,t} - \pi^{i}_{t,t-1})$, the difference in one-period forecasts, made from successive viewpoint dates. We will examine evidence below for both types of regressions, focusing primarily on the revisions.

The regression is estimated as a panel for the sample from 1981:Q4 to 2016:Q3. As indicated in Table 2, in these regressions, the strongest explanatory variables are the lagged central tendency of the distribution of forecasts and the individual forecasters' own lagged forecasts. Forecasters' estimates of lagged inflation often enter significantly, but with relatively small coefficients. Other lagged variables that might reasonably be reduced-form inputs to the forecast similarly enter with small and often insignificant coefficients. The estimated coefficient on the median forecast for *t+1* made in period *t-1* ranges from 0.28 to 0.73 across the specifications in the table. The coefficients on the lagged-viewpoint date forecasts range from 0.3 to 0.5. Other results with additional controls, not shown in this table, verify that this strong dependence on the lagged viewpoint date forecasts and the lagged central tendency of the previous period's forecast for the same period is robust to the inclusion of essentially any other variable in the forecast dataset.[9]

The right-hand columns of Table 2 show the same regressions for forecasts at horizons *t+2, t+3, t+4*. The results are the same. The bottom panel of the table replicates these same regressions for the unemployment forecasts from the SPF. Again, the lagged central tendencies and lagged viewpoint date forecasts are consistently correlated with the individual forecasts for all horizons. Here, the coefficients on the lagged central tendency range from 0.44 to 0.86, and the lagged viewpoint date forecast develop coefficients that range from 0.3 to 0.5.

[8] The data for the GDP deflator begin earlier, in 1968:Q4, but we focus on the CPI because (a) the SPF does not collect sufficient lags of the GDP deflator to form a lagged inflation measure, and (b) long-run inflation expectations are not collected for the GDP deflator. Despite these limitations, similar test regressions using the GDP inflation measure develop very similar results.

[9] For example, including current, *t+1* and *t+2* forecasts for unemployment, the Treasury bill, and output growth yields a coefficient on the lagged median forecast of 0.41 with a *p*-value of 0.000.

To help with interpretation of these results, it is useful to consider a simple framework for forecast revisions.[10] One can decompose the forecast of a variable *x* made at time *t* for forecast period *t+1* as equal to the forecast for the same variable and period made at period *t-1*, plus news about the variable that is received in period *t:*

$$x_{t+1,t} = x_{t+t,t-1} + News_t \tag{1.2}$$

Many of the regressors in equation (1.1) may be interpreted as news that becomes available in period *t* and is relevant to the forecast for *x* in period *t+1*—the estimates of lagged actual inflation, the lagged median of forecasts made in *t-1*, and other variables contained in *Z* and observed in *t*.[11] Efficient revisions imply that the coefficient in a regression of $x_{t+1,t}$ on $x_{t+1,t-1}$ should be one, and other variables' coefficients will reflect the way that forecasters use new information to revise their forecasts for period *t+1*. Equivalently, the forecast revision from period *t-1* to period *t* will reflect only news.

If the coefficient on $x_{t+1,t-1}$ differs significantly from one, so that equation (1.2) becomes

$$x_{t+1,t} = ax_{t+t,t-1} + News_t \; ; a < 1 \;, \tag{1.3}$$

then the revision from period *t-1* to period *t* responds inefficiently to the news received in period t:

$$x_{t+1,t} - x_{t+t,t-1} = (a-1)x_{t+t,t-1} + News_t$$

Put differently, the forecast from viewpoint date *t* now inefficiently over-weights information from period *t-1*. This particular inefficiency implies adding inertia to the expectations—that is an augmented dependence on lagged information—in a way that is independent of the inertia in the underlying process being forecasted. Section 7 will derive the additional persistence in the context of a multi-equation dynamic model.

In table 2, all of the coefficients on the lagged viewpoint-date forecasts develop coefficients that are quantitatively far from one. Table 2a presents results that more simply and directly test the efficiency of forecast revisions in this respect, using an augmented version of equation (1.3), as indicated at the top of the table. For all variables and all horizons, the hypothesis *a=1* is rejected overwhelmingly.[12]

---

[10] See Nordhaus (1987) for an exposition of the relationship between forecast revisions and efficiency.

[11] Here the "*News*" term subsumes the coefficient on the variables that constitute information, which would reflect the information content of those variables for forecasting *x*, although we do not assume that all of the information is incorporated efficiently, given the other results in the paper.

[12] The *p*-values are 0 to greater than ten decimal places. In fact, all of the variables and horizons reject the hypothesis $a = 0.65$ at least at the 0.01% level; most develop *p*-values of 0 to at least five decimal places.

While these simple regressions provide an interesting first look at the data, they suffer from the difficulty that it is not possible to control for all the possible inputs to the individual forecast. The lagged median forecast may enter simply because it proxies for a host of other—presumably common—information that becomes available in period *t,* and thus influences individual forecasts made in that period. The influence of common information in the individual forecasts will be explored in greater depth below.

An easier-to-interpret version of the regression casts it in terms of revisions, as suggested above. The revision explicitly differences out whatever information, idiosyncratic or common, was incorporated in the previous period's forecast of the variable. This does not assume efficiency, but it allows us to focus on the information that is used to update the forecasts from one viewpoint to the next.[13] Working with revisions may also be preferable to working with forecast errors, as it avoids having to make arbitrary decisions that are required to define "real-time" actual data. Subtracting the lagged-viewpoint forecast from both sides, one can write a restricted version of equation (1.1) in revisions form

$$\pi_{t+1,t}^{i,SPF} - \pi_{t+1,t-1}^{i,SPF} = \gamma[\pi_{t+1,t-1}^{i,SPF} - C(\pi_{t+1,t-1})] + a\pi_{t-1}^{i} + cZ_t^i + \delta_i + \varepsilon_t^i \ , \qquad (1.4)$$

which simply restricts equation (1.1) so that $b+c=1$. The forecast revision is thus a function of the discrepancy between the t-1 viewpoint forecast and the *t-1* central tendency, along with other variables. Table 2 shows the *p*-value for a test of the restriction $b+c=1$. For the inflation forecast, the test of this restriction rejects overwhelmingly, as indicated in the top panel. For the unemployment forecasts, the restriction fails to reject in all but one case. Thus for unemployment forecasts, the data cannot reject the hypothesis that the revision relationship is an appropriate representation of the forecast data.

For the remainder of the paper, we will estimate the regressions using the revision in the forecast as the dependent variable. However, because the restriction for equation (1.4) is rejected for the inflation data in the SPF, we will usually include the lagged central tendency of forecasts in the regressions, along with the discrepancy between the individual and the central tendency forecasts:

$$\pi_{t+1,t}^{i,SPF} - \pi_{t+1,t-1}^{i,SPF} = \gamma[\pi_{t+1,t-1}^{i,SPF} - C(\pi_{t+1,t-1})] + a\pi_{t-1}^{i} + bC(\pi_{t+1,t-1}) + cZ_t^i + \delta_i + \varepsilon_t^i \qquad (1.5)$$

Throughout, one can interpret the coefficient γ as the difference between a in equation (1.3) and 1. The total effect of the lagged central tendency on the revision is the sum of $-\gamma$ and b. When $b=0$, this means the sole influence of the central tendency is via an "error-correction" of the current

[13] Focusing on revisions also avoids the many difficulties that arise in working with forecast errors, as the appropriate definition of the "actual" data to use in computing the forecast error is fraught with difficulty.

forecast to the discrepancy between the previous forecast and the central tendency. When $b \neq 0$, the central tendency has an influence beyond that of a simple error-correction. That relationship may reflect common information among the individual forecasts that is revealed, with a lag, in the previous period's median forecast, and is incorporated in *t*-period forecasts.

Thus the coefficients on the lagged discrepancy in the revision regressions reveal the inefficiency with which expectations incorporate the new information contained in the lagged central tendency and other variables, or equivalently the degree of intrinsic persistence present in expectations due to this inefficiency.[14] The larger is the coefficient on the lagged discrepancy, the more weight the revision places on the lagged forecast, and the slower is the adjustment to new information. Table 3a reports the results from revision regressions from equation (1.5), where the variables are as defined for table 2. The table examines several candidates for the central tendency reference: (1) the median of all forecasts for period *t+1* made in period *t-1* (this is the measure used in table 2); (2) the average of forecasts for period *t+1* made in period *t-1* by the three forecasters with the lowest RMSE, computed real-time for the preceding 8 quarters; and (3) the forecasts for the same origin and horizon made by the forecasters who have been in the dataset longest, as a proxy for the largest and (perhaps) most-respected forecasters in the sample.

Regression (1.5) is estimated as a panel regression for the sample 1981:Q3 to 2016:Q3, with standard errors corrected as noted above.[15] The results show clearly that the lagged median of previous forecasts enters most reliably in inflation forecast revision regressions (the fourth column includes all three measures; it shows that the median dominates the other concepts). For the balance of the paper, we will use the median as the measure of the central tendency.[16] All regressions develop negative and precise estimates of γ: when forecaster *i*'s *t-1* period forecast of inflation in period *t+k* is above the central tendency of all *t-1* vintage forecasts, the i^{th} forecaster tends to quite gradually revise his next forecast for the same period toward the central tendency. Even more so than is the case for the regressions of forecast levels in Tables 2, this result appears quite robust across control variable sets and time periods. The right-hand columns of Table 3a, like their

[14] This relationship is obviously akin to the error-correction relationship between nonstationary variables. Note that in this case, the error-correction cannot really go both ways: It's not possible for the median forecast to error-correct toward all of the individual forecasts, but the converse can be true.

[15] The use of the longest-participating forecast members involves taking into account information that could not be known in the current quarter. However, it is meant to capture the idea that a few of the forecasters in the sample are large, nationally recognized forecasting firms, and thus tend to participate regularly and over a long period. The RMS forecast error measure is truly real time, with the smallest RMS error up to the regression date determining which forecasters are in this group.

[16] There is some evidence in favor of including the RMSE measure for unemployment forecasts, although it does not dominate the median.

counterparts in Table 2, show the results of forecast revision regressions for additional forecast horizons. For all forecast horizons, with all sets of controls, the coefficient on the lagged discrepancy varies between -0.53 and -0.63. The results are uniformly strong, suggesting that individual forecasters revise *all* of their forecasts in response to deviations from previous median forecasts, but do so inefficiently and only gradually.[17] The inclusion of the lagged median forecasts ($\pi^{Median}_{t+1,t-1}, \pi^{Median}_{t+2,t-1}, \pi^{Median}_{t+3,t-1}$) does not alter the size or significance of the response of revisions to the lagged discrepancy—the estimate of inefficiency in the forecast revision, $a-1$.[18]

Figure 2 displays a scatter plot of the left-hand-side variable (the forecast revision) against the lagged discrepancy (the first term on the right-hand side in (1.5)), and the negative correlation is clear. Figure 3 displays a histogram of the coefficients for equation (1.5) estimated for each forecaster in the sample. While there is clearly some heterogeneity in the "speed of adjustment" to new information, it is also clear that the mass of estimates is solidly centered between zero and minus one, with a modest standard error. The aggregate regression is not the artifact of a few outliers.

Table 3b provides parallel results for the unemployment forecasts from the SPF, using the revisions to the one- to three-quarter-ahead forecasts for the unemployment rate. Once again, the evidence of inefficient revisions that respond slowly to new information in the median forecast is strong, and changes little with the addition of other forecaster-specific controls. The right-hand columns display results for the longer forecast horizons, and the results are similarly strong. Regardless of the set of control variables, the revision in the forecast for period *t+k* between periods *t-1* and *t* always responds significantly and sizably to the lagged-viewpoint forecast and to the median of all forecasts last period. Regressions using the SPF's forecasts of the 3-month Treasury bill and real GDP growth, not shown, produce very similar results.

Figure 4 displays evidence on the time-variation in the key regression coefficient in Figure 2. The top panel shows the histogram of the estimates of $(a-1)$, using twenty-quarter rolling samples from 1981-2015. The coefficients all fall between -0.4 and -0.8, most commonly from -0.55 to -0.75. The second panel of the figure shows the estimated coefficients over time. The values are quite

[17] Because the quarterly forecasts extend out only four quarters, we are only able to compute lagged forecast revisions out to quarter t+3.

[18] Note that the discrepancies for horizons *t+2* and *t+3* are adjusted accordingly ($\pi^{i}_{t+2,t-1} - \pi^{Median}_{t+2,t-1}, \pi^{i}_{t+3,t-1} - \pi^{Median}_{t+3,t-1}$, respectively). Results for the four-quarter average forecast from *t* to *t+3* produce similar results—for example, the coefficient on the discrepancy is -0.46 for inflation, with *p*-value of 0.000.

stable from 1981 through 2000. There is a modest decline in the magnitude in the mid-2000s to about -0.4, but in more recent samples, the estimate has reverted to about -0.7. The standard errors on these coefficients, not shown, are about 0.01, so these fluctuations are statistically significant.

Table 4 provides results for the forecast difference variant of the regressions reported in Table 3. That is, the dependent variable is the change in the *k*-period-ahead forecast from quarter to quarter; for example, the *t+1* forecast made in period *t* minus the *t*-period forecast made in period *t-1*. Correspondingly, the discrepancy variable is defined as the difference between the *k*-period-ahead forecast made in period *t-1* and the central tendency of such forecasts made in period *t-1*. Additional columns in the table allow for other controls in the forecast dataset, and for increasing forecast horizons from *t*+1 through *t+3*, so that the dependent variables are the *t+k* forecast made in period *t* minus the *t+k-1* forecast made in period *t-1;* the discrepancies are adjusted accordingly.

Interpreting this regression requires some care: *k*-period-ahead forecasts for persistent variables should be correlated over time. The question is whether the changes in *k*-period-ahead forecasts reflect the same kind of gradual adjustment to new information that the revisions in table 3 exhibit. To clarify interpretation, consider a simple process for *x*[19]

$$x_t = \rho x_{t-1} + e_t \ .$$

The forecasts for *t* and *t+1* made in periods *t-1* and *t* respectively should be related by

$$x_{t+1,t} = x_{t,t-1} + \rho(x_{t,t} - x_{t-1}). \tag{1.6}$$

Like the forecast revision regression, the coefficient on the lagged forecast should be one, and the change in the one-period forecast should reflect news bearing on the change in *x* from period *t-1* to *t*. Many of the regression variables in Tables 2 and 3 may represent information that is relevant for assessing the change in *x* from *t-1* to *t*, but any relationship between the change in the forecast and $x_{t,t-1}$ represents an inefficiency, as there should be no new information about the change in *x* that is present in $x_{t,t-1}$. If the coefficient on $x_{t,t-1}$ in equation (1.6) is not one, then the regression of the change in *k*-period forecasts on $x_{t,t-1}$ will indicate an inefficient linking of the current *k*-period forecast to last period's, analogous to the issue with the revision forecast. In addition to the controls in the revision regressions, we include the forecast for the current period made in the current period, $x^i_{t,t}$, which will contain whatever information forecasters believe to be relevant in updating x_t.

The forecast change regressions displayed in table 4 show even stronger results than those in table 3: no matter the set of controls, the change in the *k*-period-ahead forecast responds strongly

[19] This derivation can be generalized by allowing *x* to depend on a vector of factors X; the logic remains the same.

and significantly to the discrepancy between the *k*-period-ahead forecast last period and the median of all *k*-period-ahead forecasts last period. This of course indicates that, in addition to updating forecasts with information about the change in inflation from period *t* to period *t-1* (some of which may be contained in the lagged median forecast), the forecasts are inefficiently tied to the previous *k*-period-ahead forecast. The effect is fairly consistent in magnitude and significance as the forecast horizon increases and as additional controls are added. This similarly implies that forecasts are adjusted gradually over time to incorporate the information in the lagged central tendency, but here it is not the revision in the forecast for a fixed forecast period that error-corrects, but the change in the *k*-quarter-ahead forecast for successive forecast periods. Such behavior also constitutes a source of intrinsic persistence in expectations. The macroeconomic implications of these results, and those for forecast revisions, are discussed in section 7.

The role of common information

As discussed above, it is possible that the forecasts are correlated with the lagged median forecast simply because the median forecast, not observed when forecasters submit their *t-1* forecasts, contains information that forecasters should use to update their forecasts. Of course, revisions to individual forecasts should not reflect the common information known to forecasters at the time of forecast. However, revisions to individual forecasts might reflect revisions to the common information known at the time of the forecast. To control for this possibility, Table 5 presents regressions of the individual forecast revisions on the lagged discrepancies from Table 3, adding the revision in the median forecast, which could reflect revisions due to changes in commonly held information. The last aggregate forecast revision that we know can be observed by individual forecasters is the change in the median forecast from viewpoint *t-2* to viewpoint *t-1*; this is the first added regressor in the table. As the results in the table indicate, while the lagged aggregate revision is sometimes significant, this addition has no impact on the key result from above: Individual forecasters continue to revise their forecasts gradually and inefficiently in response to the lagged discrepancy between their forecast and the median forecast.

But forecasters may also revise the current forecast based on revisions in common information for period *t* that is not observable to the econometrician. While the contemporaneous revision to the aggregate forecast cannot be observed by individual forecasters in real time, some of the information that it contains may be observed by forecasters at time *t*. Thus contemporaneous aggregate forecast revisions are included in the right-hand columns of Table 5 as a generous proxy

for contemporaneous revisions in unobserved common information. While the coefficients on this variable are larger and quite significant—estimated magnitudes fall between 0.8 and 0.9, with near-zero *p*-values—the coefficients on the individual forecast discrepancies are essentially the same as those using the lagged aggregate revision, and are qualitatively unchanged from the regressions that omit the aggregate revision. As a way of controlling for the fact that the contemporaneous revision is not observable to individual forecasters at the time it is collected, the final column of the upper panel of the table provides estimates in which the current aggregate revision is instrumented by lags of aggregate revisions for periods *t* and *t+1*. The results are virtually identical to the others.

The bottom panel of Table 5 replicates these results for the unemployment forecasts in the SPF. As with the inflation forecasts, the inclusion of lagged, contemporaneous or instrumented contemporaneous revisions has no effect on the correlation between the individual forecast revisions and the lagged discrepancy from the median forecast. If anything, the inclusion of controls for revisions in common information strengthens the key results from Table 3.

Note that the bottom panel of Table 4 presents the effect of common information revisions on the regressions in that table (for changes in forecasts, rather than revisions). Once again, the inclusion of either lagged or contemporaneous revisions in aggregate forecasts has no effect on the tendency for forecast changes to respond quite gradually to the lagged discrepancy between the same-horizon forecast and the central tendency of such forecasts.[20]

Table 6 presents regressions that add a host of additional revision variables. The revisions include revisions to the aggregate forecasts, both lagged and contemporaneous; revisions to individual lagged inflation, unemployment, Treasury bill and output growth estimates; revisions to current-period forecasts of the same four variables; and revisions to other forecast variables for other forecast horizons. The table essentially provides a way of decomposing the source of revisions for a given forecast from period *t-1* to period *t,* using all of the information in the forecast dataset. As the table indicates, none of these variables alter the conclusion that revisions respond inefficiently to new information, including any information newly revealed in the lagged central tendencies. The coefficients for the inflation variable are a bit smaller than in the baseline; the coefficients for the unemployment variable are the same size. The significance is not at all affected.

[20] As in the top panel of table 4, these regressions include the current real-time individual estimates of x_t, that is the forecasts $x^i_{t,t}$.

Given the "kitchen sink" nature of this regression, and the difficulty in interpreting it once one includes revisions of other forecast variables and horizons, this is a strong result.

Learning versus lagged central tendencies

A vast literature has examined the properties of models in which agents must learn about their economic environments, possibly converging to rational expectations equilibria over time (see the citations above). Can the results in this paper distinguish between anchoring to a lagged central tendency and learning behavior?

The answer appears to be "yes," although this is a tentative conclusion. Learning models typically posit least-squares or recursive least-squares learning, in which expectations are formed by time-varying projections of observables on lagged data. Such projections may be viewed as the reduced form for an expectations process that could converge, with sufficient observations and stability of the economic environment, to the restricted reduced form consistent with the rational expectations solution for the model economy (see the work pioneered by Evans and Honkapohja, as summarized in their landmark 2001 book).

Table 7 examines regressions that include the lagged discrepancy variables discussed above, along with individual real-time estimates of lagged macro variables, as a way of determining whether the results presented above are in some way a proxy for learning about the reduced-form projection of the variables of interest on lagged observables. The left-hand columns focus on inflation forecasts, and the right-hand columns focus on unemployment forecasts. The leading columns in these blocks simply reprise the results from above, which show that for the full sample, the inclusion of lagged actual variables does not change the dependence on the lagged discrepancy. The next sets of columns estimate these regressions over shrinking samples going forward in five-year blocks. These columns show that this feature of the forecasts is extremely stable over time. The results in Table 7 suggest strongly that the tendency to revise forecasts inefficiently, leading to intrinsic persistence in expectations, is quite distinct from the formation of expectations from lagged real-time realizations of inflation, unemployment, output or interest rates. The coefficient on the discrepancy variables remains uniformly negative and overwhelmingly significant. There is some evidence of a linkage from expectations to lagged and current real-time actuals, but these coefficients are generally smaller and less significant. The presence of these variables does not reduce the size of the response to the discrepancy, suggesting that learning and inefficiently gradual responses to new information remain distinct in these regressions.

Figure 5 presents results that allow period-by-period time-variation in the projections, which conforms more to the spirit of the learning literature. The figure shows estimated coefficients for rolling estimates of the equation from Table 7 for the revision to the one-quarter inflation forecast. The top panel shows the coefficient on the lagged discrepancy, and the bottom panel shows the coefficients on lagged real-time inflation. The coefficients are estimated precisely throughout. There is a modest amount of time-variation, but there is no evidence in these estimates that the tendency for forecasters to move their forecast toward the lagged central tendency is a proxy for least-squares learning projections on lagged observables.

Altogether, the results summarized in Tables 2–7 suggest that forecasters revise their current-period forecasts inefficiently, incorporating news (including the lagged central tendency of all forecasts) slowly. In so doing, they introduce intrinsic persistence to their forecasts, dramatically slowing their adjustment to new information. This finding holds for all forecast horizons for inflation, unemployment, and other forecasted variables in the SPF dataset. The result holds when including controls for lagged information, revisions to aggregate forecasts that might reflect revisions to unobserved common information, and revisions to estimates of lagged and current variables that might be used as inputs to individual forecasts.

The dependence of forecast revisions on lagged forecasts suggests dynamics in expectations that cannot be captured by full-information rational expectations models. The results presented in table 7 and in figure 5 suggest that this behavior is not a stand-in for least-squares learning. A richer information structure combined with sluggish incorporation of new information is required to motivate these findings; a simple example of such a structure is discussed in Section 6 below.

2. Evidence from the European SPF

The ESPF surveys are organized somewhat differently from the Philadelphia Fed's SPF. The available forecast horizons change during the history of the survey, which began in 1999. The forecasts employed in this paper include the current year and the one- and two-year ahead forecasts for inflation, unemployment, and output growth. The relationship between forecasts from quarter to quarter is not the same as in the SPF; the current forecast year remains the same for all four quarters of a calendar year, whereas the quarterly-focused SPF's current quarter changes with the survey quarter. As a consequence, some care must be taken in defining forecast revisions in the ESPF. More details on the ESPF may be found on the ECB website, referenced in the appendix.

Tables 8-10 provide estimation results for forecast revisions that parallel those for the SPF dataset. For each forecast variable (inflation, unemployment and output growth), we examine the predictability of the revision in the current-year and one-year-ahead forecast. As with the SPF forecasts, we are particularly interested in whether the revisions efficiently incorporate new information. To do so, we run regressions like those in tables 3, focusing on the correlation between the revisions and the discrepancy between the previous quarter's individual forecast and the median of all previous quarter's forecasts. As above, these regressions can provide evidence of inefficient revisions that imply sluggish adjustment to new information. Recognizing the difference in the timing convention between the SPF and the ESPF, we estimate regressions of the form

$$\pi_{y1,t}^{i,ESPF} - \pi_{yk,t-1}^{i,ESPF} = \gamma[\pi_{yk,t-1}^{i,ESPF} - C(\pi_{yk,t-1})] + b\pi_{t-1} + cZ_t^i + \delta_i + \varepsilon_t^i; k = 0,1 \qquad (2.1)$$

where now the revision denoted by $\pi_{yk,t}^{i,ESPF} - \pi_{yk,t-1}^{i,ESPF}$ refers to the change from last quarter to this quarter in the forecast for year *k* made by forecaster *i*. The discrepancy from last period denoted by $\pi_{yk,t-1}^{i,ESPF} - C(\pi_{yk,t-1})$ is the difference between the forecast for year *k* made last quarter by forecaster *i* and the central tendency of all forecasts for year *k* made last quarter. In this section, we consider only the median as the measure of central tendency. The ESPF does not collect individual forecasters' assessments of last quarter's/year's observations, so we use the real-time estimates of lagged inflation (and unemployment and real growth) in the regressions that follow. Of course, the observations for these real-time estimates do not vary across forecasters.

The control variables in Z_t^i differ from those in the US SPF, as the ECB survey collects what they call "assumption" variables for the price of oil, the exchange value of the euro relative to the dollar, the ECB policy rate assumption, and (for some observations) a labor cost measure. These "assumption" variables are collected for the same forecast horizons as the three main variables of interest. Tables 8-10 display simple versions of the test regression (2.1) which omit Z_t^i, as well as versions that include assumption variables, lagged revisions, lagged discrepancies, and current values of the forecasts for the other variables in the survey.[21] The regressions all span the available data for the Euro SPF from 1999:Q1 to 2016:Q1.

The robust conclusion from these results is the same as that for the US's SPF: Individual forecasters adjust their forecasts in this period to the information revealed in the median of all forecasts last period, but they do so gradually and inefficiently, tying current forecasts to previous

[21] An important difference between the ECB dataset and the Philadelphia Fed's SPF is that the former does not capture the real-time estimate of lagged inflation.

forecasts. The results are as strong as the U.S. results for inflation, with somewhat smaller coefficients for the unemployment rate. Table 11 includes the revisions to the aggregate (median) forecasts, in an attempt to control for the influence of common information on individual forecasts as with the SPF data. Again, the response to the lagged forecast discrepancy is unaffected by the inclusion of these strong proxies for revisions to common information.

3. Evidence from households

Table 12 provides evidence on the revisions of forecasts from the University of Michigan's Survey Research Center Survey of Consumers. This monthly survey is largely a cross-sectional survey of about 500 randomly selected households per month. However, a subsample (about one-fifth) of respondents is interviewed again six months later, and the unique identifiers assigned to each respondent allow us to track this subset of households from the first to the second interview. This limited panel feature of the data allows us to examine the revisions in inflation expectations.

Table 12 displays the results from the test regressions

$$\pi^{i,Mich}_{t+1y,t} - \pi^{i,Mich}_{t+1y,t-1} = a\pi_{t-1,t} + b[\pi^{Mich}_{t+1y,t-1} - C(\pi^{Mich}_{t+1y,t-1})] + cC(\pi^{Mich}_{t+1y,t-1}) + dZ^{i}_{t} + \delta_{i} + \varepsilon^{i}_{t}, \qquad (3.1)$$

where $\pi^{i,Mich}_{t+1y,t}$ is the i^{th} forecaster's one-year-ahead inflation expectation made in period *t* and $\pi^{i,Mich}_{t+1y,t-1}$ the corresponding expectation made in the previous period *t-1,* $\pi^{i}_{t-1,t}$ is the real-time estimate for lagged actual inflation for the vintage of data collected for period *t,* $C(\pi^{Mich}_{t+1y,t-1})$ is the median of all forecasters' one-year-ahead inflation forecasts made in period *t-1,* and *Z* represents a vector of other controls that include survey respondents' continuous and qualitative assessments of unemployment, family income, current and expected financial prospects, and general business conditions.[22]

The bottom panel of Table 12 provides the results of the simple test for forecast revision efficiency, as discussed above for the SPF forecasts. The sample spans 1978:Jan through 2016:Apr. The results for the test regression, for both the one-year and the five-year inflation forecasts, are unequivocal: The sub-sample of Michigan respondents does not use the information in their previous forecasts efficiently (the test $a = 1$ in the test regression $\pi^{i,Mich}_{t+1y,t} = a\pi^{i,Mich}_{t+1y,t-1} + bC(\pi^{Mich}_{t+1y,t-1}) + \varepsilon^{i}_{t}$ rejects with overwhelming significance).

Table 12 provides the results from regressions of the regression in equation (3.1), as in equation (1.4) above for the SPF data. Because the time dimension of individual survey participants'

[22] The assessments of one-year and five-year inflation and family income expectations are numeric; other variables are encoded according to better/worse/same or similar qualitative categories.

responses is limited, we examine in this table the extent to which the pooled-cross section results vary over time. With a sizable number of observations for each cross-section, we are also able to examine whether these revision regressions correspond only to times of economic tumult (recessions), or times of relative calm, or both.

Here again, the results are strong and consistent across controls and time periods. The respondents inefficiently use the information in their previous forecasts of inflation. The coefficient on the lagged discrepancy between individual forecasts and the median forecast varies narrowly between -0.68 and -0.72 for all of the specifications presented in the table, indicating a small coefficient on the lagged viewpoint date forecast and a sizable coefficient on the lagged median forecast. While it certainly seems plausible that Michigan responds do not produce efficient forecast revisions, it seems somewhat less plausible that households exhibit the kind of consistency that the SPF participants show in responding to previous periods' central tendencies. On the other hand, the number of observations is almost two orders of magnitude larger, so our confidence in the statistical significance of the results is high, even if the individual behaviors of household respondents may vary significantly around the estimated results.

Some may question the likelihood that the household respondents in the Michigan survey anchor their expectations to the previous central tendency. However, the revision results in Table 12 are based on the subset of survey participants who are re-sampled six months later. This subgroup may make some effort at that point to check the newspaper, the news, or the Internet to discover what people are saying about inflation, and they may revise their expectations toward that observation, as suggested by the regression results. This kind of "paying attention when it counts"—a variant of rational inattention models (see, for example, Sims 2006)—might suggest that consumers considering an important decision may also pay attention to prevailing forecasts/economic opinions/commentary at these key decision points.

4. **"Anchoring" inflation expectations**

Many economists embrace the notion that inflation expectations may be "well-anchored" to the central bank's inflation goal, especially in the context of a credible inflation-targeting monetary regime. By this, economists often mean that long-run inflation expectations do not deviate far from the central bank's announced inflation goal. In addition, they often assert that such anchored expectations provide a firm anchor for realized inflation, perhaps explaining why the variation of inflation in the wake of the Great Recession has been relatively small.

Note that in rational expectations models, if the price-setting agents know the central bank's target, their expectations will be perfectly anchored, in the sense that all well-behaved models that embed such a price-setting mechanism will converge to the central bank's goal. Of course, the rate of convergence will depend upon key parameters governing other aspects of the model, including the monetary authority, the consumption Euler equation, and so on. But one can envision an environment in which price-setters are uncertain about the central bank's goal, or about the central bank's commitment to a known goal. In this case, it is possible for long-run expectations to become un-anchored from the central bank's target. While most speak of "anchored expectations" with somewhat less specificity than this, it has nonetheless become a mantra of central bankers to speak about the importance of anchored expectations that assure an ultimate return of inflation to the central bank's inflation target.

If anchoring to long-run expectations is an important feature of inflation and inflation expectations, then the omission of this variable from the regressions above could bias the estimates presented in Tables 2–12. However, the SPF and Michigan datasets allow us to examine the extent to which short-run inflation expectations are anchored to long-run expectations. Figure 6 displays the median 10-year CPI inflation forecast from the SPF from the date it was first collected (1991:Q4) through mid-2016.

Table 13 presents results from regressions that augment those in Section 2 with the revision to the median 10-year CPI inflation forecast, which enters with a lag, as it would not be observable to all forecasters contemporaneously. The top panel of the table presents results from these regressions for the full sample. The long-run forecast revision typically does not enter significantly, but regardless, it does not alter the strong but sluggish reversion to the lagged discrepancies reported throughout. The bottom panel displays the same regressions for the period from 2000 to the present. While a few of the coefficients on the lagged 10-year forecast revision change in magnitude, none are significant, and the effects on the response to the lagged discrepancy are trivial.

The household data afford some opportunity to examine the question of anchoring as well. For most of the sample, a 5-year inflation forecast is collected by the SRC, so we use this as a proxy for the long-run forecast around which short-run expectations might be anchored. For expositional clarity, and because the 1- and 5-year expectations have a 20 percent overlap, we construct the

implied expectation for years 2–5 and use it as the long-run anchoring proxy.[23] As Table 14 shows, short-run expectations remain tied to the lagged central tendency regardless of which other regressors are included. There appears to be some linkage to the lagged median 2–5-year expectation, but the magnitude is modest. Whether this constitutes anchoring to the central bank's inflation goal or part of the solution to a filtering problem, in much the same way as the link to the 1-year expectation, is difficult to tell. Overall, then, while the evidence for sluggishly incorporating the information in lagged aggregate expectations remains strong, the evidence for anchoring to the long-run expectation is modest, at best.

5. **Sticky information?**

The important work of Coibion and Gorodnichenko (2015) finds high-level support in aggregate surveys of expectations for the sticky information model of Mankiw and Reis (2002), and for the noisy information model of Maćkowiak and Wiederholt (2009) and others. While the paper provides a host of useful empirical results, the key insight is that both models imply that forecast revisions are sufficient to explain forecast errors (in the sense that all other variables lose their significance in aggregate forecast error regressions). The logic follows directly from the definition of the sticky information setup (the noisy information case is discussed in the next section). The average expectation for variable x at date t will be a geometrically weighted average of the rational expectations formed at the current and all lagged viewpoint dates:

$$x_{t+1,t} = (1-\lambda)\sum_{k=0}^{\infty} \lambda^k E_{t-k} x_{t+1} \tag{5.1}$$

The average expectation as of date $t-1$ is given by a parallel equation

$$x_{t+1,t-1} = (1-\lambda)\sum_{k=0}^{\infty} \lambda^k E_{t-k-1} x_{t+1} \,, \tag{5.2}$$

which implies that the revision from the $t-1$ to the t period forecast is given by

$$R_{t+1} \equiv x_{t+1,t} - x_{t+1,t-1} = (1-\lambda)(E_t x_{t+1} - x_{t+1,t-1}) \,, \tag{5.3}$$

and the forecast errors are related only to the revision, as indicated in equation (5) of their paper

$$x_{t+1} - x_{t+1,t} = \nu_{t+1,t} + \frac{\lambda}{1-\lambda}[x_{t+1,t} - x_{t+1,t-1}] \,, \tag{5.4}$$

[23] The two- to five-year expectation is computed as one fourth the difference between five times the five-year expectation and the one-year expectation, i.e., $X^e_{t+2...5} = 0.25[5(X^e_{t+1,...,5}) - X^e_{t+1}]; X^e_{t+1,...,5} = 0.2[X^e_{t+1} + ... + X^e_{t+5}]$.

where $\nu_{t+1,t}$ is the rational expectations error defined as the difference between realized x_{t+1} and the rational expectation. These relations hold only over the mean forecast. As Coibion and Gorodnichenko emphasize, under the assumptions of the sticky information model, agents either do not revise at all, or they revise to the rational expectation, so it is only on average that equations (5.1) -(5.4) are expected to hold.

The evidence above, augmented by evidence in this section, suggests that the sticky information model is not a good approximation to expectations behavior in these surveys. First, the model suggests that in any given quarter, a significant number of agents do not update their information sets, so that their forecasts in period *t* equal those in period *t-1*. It is not credible that professional forecasters do not update their information sets for six months at a time. For households, this might well be a good approximation to their updating frequency, but then the premise that households that do update information sets make rational forecasts is suspect. Likely or not, we will test these propositions below.

To begin with, we can provide a crude measure of the fraction of professional forecasters and households who do not update their information set, using the fraction whose forecast revision is precisely zero (see Andrade and Le Bihan (2013) who examine the same issue for the European SPF dataset). Of course, at the quarterly frequency, some forecasters may well have fully updated their information set but, from time to time, they may judge that the information received is not sufficient for them to alter their forecasts.[24] So for the professionals, this fraction is likely biased upward from the true share who does not update their information set. Table 15 provides these shares. For one-quarter-ahead inflation forecasts from the SPF, about 18 percent of forecasters' revisions are zero. The number is about the same for unemployment rate forecasts. For the four-quarter average forecast, the primary horizon studied in Coibion and Gorodnichenko (2015), the fraction of unrevised forecasts drops quite a bit to about six or seven percent; equivalently, 93-94 percent of forecasters have revised their four-quarter forecasts from one quarter to the next, and it is likely that at least that many have updated their information sets. The difference between the fractions for the one-quarter and four-quarter average forecasts likely reflects the fact that while any one quarter's forecast might not be revised from one quarter to the next, the likelihood is small that none of the four quarterly forecasts is changed. Thus this number probably provides a better indication of whether forecasters update information from one quarter to the next. The numbers are

[24] This possibility is increased slightly by the fact that some of the forecasters in the survey always report forecasts to the nearest one-tenth of a percentage point.

similar but still noticeably higher for the Euro SPF forecasters, in the three right-hand panels. The Michigan survey participants, not surprisingly, have a higher incidence of zero revisions, at about one-third. Infrequent updating of information may indeed make more sense for households. Figure 7 displays the histogram of revisions to the 4-quarter inflation forecasts from the SPF.

Because the Coibion-Gorodnichenko test regression applies only to the average of forecasts, it is not replicated here on individual forecasts. However, the crux of the sticky information model is that agents who update their information sets should at that point form rational expectations with all the information available at that time. Thus, another simple test of the sticky information model is a regression of (real-time) forecast errors on information available at the time of the forecast to forecasters who update. Using the imperfect proxy of nonzero forecast revisions to identify information updaters, we regress forecast errors on *t*-period information, notably the forecast revisions and the lagged median forecast that has been used throughout. Forecast errors are defined relative to real-time actual data, using the convention that the "actual" is the real-time estimate of the variable at the appropriate forecast horizon, as of the data vintage eight quarters after the period the forecast was made. Table 16 provides the results of these regressions for both the SPF and the Michigan surveys. In both cases, lagged median forecasts, revisions, and other variables enter significantly, and the R-squareds for the SPF forecasts are sizable. The column that includes "additional *t*-period information" adds other individual forecast variables and lagged median forecasts, all of which are available to the forecasters. For these columns, the R^2 get fairly large, ranging from 0.14 to 0.29, which is a lot of individual forecast error variation explained. The Michigan forecasts similarly evince very significant coefficients on lagged median forecasts (and lagged individual forecasts, not shown); the R-squareds are even higher than those for the SPF inflation forecasts, which is striking given the noise in these household responses.[25]

Of course, because most all SPF forecasters update information frequently, the results presented in the previous sections also constitute a wealth of evidence rejecting the sticky information model, as all of these results also reflect grossly inefficient forecasts. Thus the results in the paper suggest an inefficient use of information by all forecasters, but that appears not to be well-represented as the outcome of agents who infrequently update their information sets, but form rational forecasts when they do. Evidence on the frequency of updating suggests the professionals

[25] The SPF forecast errors are defined relative to real-time data for the vintage of data eight quarters after the realization date, using the real-time data provided on the Survey of Professional Forecasters site. For the Michigan survey, we employ the same timing convention, using the Philadelphia Fed's eight-quarter forward real-time vintages for the monthly 12-month percentage change in the CPI.

are not surprisingly quite up-to-date on their macro information. Nonetheless, they use it inefficiently. About two-thirds of household revisions are non-zero after six months, suggesting the possibility of infrequent updating on their part. But even those who do revise their forecast show significant signs of inefficiency. For both these reasons, then, the sticky information model receives little support from the micro data.

6. Noisy information?

The results presented so far may map more neatly into a noisy information framework, in which agents receive noisy idiosyncratic signals about the variables they wish to forecast. In this case, they will not adjust completely to the news in current information, but will instead revise their forecast with some weight on the new information and some on their previous forecast, with the weights depending on their perceptions of the relative signal-to-noise ratios in the two inputs.

Following the simple framework in Coibion and Gorodnichenko (2015) but adapting for our notation and for one-period-ahead forecasts, we can derive some implications for the results in the paper. First, posit an autoregressive process for a variable

$$x_t = \rho x_{t-1} + \varepsilon_t; -1 \le \rho \le 1 \ . \tag{6.1}$$

This process may be readily generalized by allowing x to be a vector of variables, including lags of the vector x, and ρ a conformable matrix. Agents in the economy cannot (ever) observe x_t without noise, but instead receive a noisy signal y_t^i

$$y_t^i = x_t + \omega_t^i \ , \tag{6.2}$$

where ω_t^i is assumed *iid* across time and individuals. Under these circumstances, agents will compute forecasts for periods *t* and *t+h* as

$$\begin{aligned} x_{t,t}^i &= G y_t^i + (1-G) x_{t,t-1}^i \\ x_{t+h,t}^i &= \rho^h x_{t,t}^i \end{aligned} \ , \tag{6.3}$$

where G is the Kalman gain, based on the relative signal-to-noise ratios in y_t^i and $x_{t+1,t-1}^i$. These equations imply that the forecasts for period *t+1* made in periods *t-1* and *t* are

$$\begin{aligned} x_{t+1,t}^i &= \rho x_{t,t}^i = \rho[G y_t^i + (1-G)\rho x_{t-1,t-1}^i] \\ x_{t+1,t-1}^i &= \rho^2 x_{t-1,t-1}^i \end{aligned} \ , \tag{6.4}$$

which in turn implies, after some simplification, that the revision in the *t+1* forecast between viewpoint dates *t-1* and *t* is

$$x^i_{t+1,t} - x^i_{t+1,t-1} = \rho G(y^i_t - \rho x^i_{t-1,t-1}) \quad .^{26} \tag{6.5}$$

This forecast update equation depends on the Kalman gain and the difference between the newly-received signal for x_t and last period's forecast. When $G = 1$, the difference between these estimates of x_t is just the true news about x_t, which is ε_t, so the revision reduces to $\rho\varepsilon_t$, as in the simple examples in section 2. In the regressions in Tables 3-12 above, the weight on the lagged forecast is estimated to be negative, sizable, and remarkably significant, consistent with equation (6.5). Coibion and Gorodnichenko show that one can also use these definitions to derive a forecast error regression like equation (5.4) above, such that the average forecast errors are related only to the average forecast revisions. In this case, the coefficient on the forecast revisions may be interpreted as a simple function of the gain parameter. As they point out, the coefficient on different forecast variables will vary with the Kalman gain, which depends in turn on the signal-to-noise ratio of the variable and its persistence. But one can also show that the individual forecast errors in this noisy information setup are proportional to individual forecast revisions, plus the *iid* error term ω^i_t:

$$x_t - x^i_{t,t-1} = \frac{1}{G}[x^i_{t,t} - x^i_{t,t-1}] - \omega^i_t \quad . \tag{6.6}$$

In the results presented in tables 3-12 above, we can infer values of *G* in equation (6.5) from the estimated coefficients on the lagged discrepancies and the estimated persistence of the series being forecast. The following table shows the simple mapping from the estimated coefficient to the estimated gain, for different values of the coefficient and ρ:

Implied value of gain coefficient *G* for values of persistence and revision regression coefficient						
Coefficient on discrepancy	**Persistence (ρ)**					
	0.5	**0.6**	**0.7**	**0.8**	**0.9**	**0.95**
-0.3	1.2	0.83	0.61	0.47	0.37	0.33
-0.4	1.6	1.1	0.82	0.62	0.49	0.44
-0.5	2	1.4	1	0.78	0.62	0.55
-0.6	2.4	1.7	1.2	0.94	0.74	0.66
-0.7	2.8	1.9	1.4	1.1	0.86	0.78

[26] When G=1, $y^i_t = x_t = \rho x_{t-1} + \varepsilon_t$, and $x^i_{t-1,t-1} = x_{t-1}$, and in this case of course the forecast revision reduces to $\rho\varepsilon_t$, the news about x_t that is revealed in period *t*. This in turn is consistent with the definition of an efficient full-information revision in equation (1.2) above.

Some of these estimates are in the range of those in Coibion and Gorodnichenko, although their baseline estimate of G=0.46 implies either very high persistence (appropriate for the unemployment rate), or a smaller coefficient on the lagged forecast than we typically obtain.

While these comparisons may be of interest, it is still difficult to reconcile the noisy information story with the findings presented in Table 16, which encompass the test regression for this model in equation (6.6). Forecast errors should only be predictable on average across forecasters; individual forecasters should be making rational forecasts, conditional on their information sets. If it can be shown that individual forecast errors are inefficient, given information known to the individual forecasters, the model is violated. As can be seen in table 16, forecast errors are still quite predictable by a number of variables, in addition to forecast revisions.[27] Note that this table does not include an exhaustive list of all variables that are clearly in the forecaster's information sets. The test that all variables other than the forecast revision are insignificantly different from zero in these forecast error regressions rejects incredibly strongly, all with *p*-values of less than 0.000. Note too that the increment to the R^2 from including the additional *t*-period information is sizable, suggesting that most of the forecast error is not explained by the revisions, but by other information clearly available to the forecasters at the time the forecasts are made.

Overall, it seems fair to conclude that forecasters, both household and professional, do not make rational forecasts, even accounting for possible information frictions. They simply use information inefficiently. This is not an artifact of a staggered information or noisy information environment, as these models' predictions appear to be strongly violated at the micro level.

7. Implications for macroeconomic modeling

Here, we examine the macroeconomic implications of expectations that incorporate inefficient revisions in a model that conforms reasonably well to the results from the micro survey data. specifically, we construct a model in which the *t+1*-quarter expectation made in period *t* inefficiently uses the information in the expectation for quarter *t+1* made from expectation viewpoint *t-1,* and/or the lagged aggregate one-quarter-ahead expectation. The empirical results in Tables 2–12 provide evidence of both types of anchoring, and, as suggested above, there is a conceptual difference between the two inefficiencies.

[27] In this case, one would not restrict the sample to those forecasts that are revised from the previous viewpoint date. Replicating Table 16 for the full sample does not change the results.

We begin by showing algebraically how this kind of expectational inefficiency introduces inertia into a simple forward-looking model. Consider a model that comprises a New-Keynesian Phillips curve, augmented with an AR(1) process for the output gap:

$$\begin{aligned} \pi_t &= \beta E_t \pi_{t+1} + \gamma y_t + \varepsilon_t \\ y_t &= \rho y_{t-1} + u_t \end{aligned} \quad . \tag{7.1}$$

One can show that, employing the rational expectations solution for the model, the *t*-period expectation of inflation in period *t+1* is[28]

$$E_t \pi_{t+1} = \frac{\rho\gamma}{1-\rho\beta} y_t \ , \tag{7.2}$$

and the *t-1* period expectation for the same quantity is

$$E_{t-1} \pi_{t+1} = \frac{\rho^2\gamma}{1-\rho\beta} y_{t-1} \ , \tag{7.3}$$

which implies that the efficient revision to the rational forecast is

$$E_t \pi_{t+1} - E_{t-1} \pi_{t+1} = \frac{\rho\gamma}{1-\rho\beta} u_t \, , \tag{7.4}$$

which is proportional to the news about the output gap that is revealed in period *t*. If instead, agents in the model revise their expectations in conformity with the empirical results above, i.e.

$$E_t \pi_{t+1} = a E_{t-1} \pi_{t+1} + \frac{\rho\gamma}{1-\rho\beta} u_t ; a < 1 \ , \tag{7.5}$$

then the revision introduces additional lag-dependence to the *t*-period forecast;

$$E_t \pi_{t+1} = b^* y_{t-1} + \frac{\rho\gamma}{1-\rho\beta} u_t ; b^* = \frac{a\rho^2\gamma}{1-\rho\beta} \ , \tag{7.6}$$

(compare equation (7.6) to equation (7.4)). This illustrates the way in which inefficient revisions add persistence to expectations in a very stylized model with a single lag. In general, any lagged information that is used to form the (here rational) *t-1* period forecast will now enter the *t*-period forecast due to inefficient revision, adding persistence in this case to inflation that arises entirely through the expectations process.

We now examine a more fully-articulated DSGE model that embeds such expectations behavior throughout. The model includes a Phillips curve that mixes rational and inertial expectations

$$\pi_t = b \pi_{t+1,t}^{Agg} + (1-b) E \pi_{t+1} - \gamma \tilde{U}_t \, , \tag{7.7}$$

[28] See the derivation in Fuhrer (2006).

where $\pi^{Agg}_{t+1,t}$ is the aggregate inertial expectation for inflation in period *t+1* using information up to period *t,* and $\tilde{U}_t$ is the unemployment gap (or the output gap, or real marginal cost; for these purposes all of these driving variables are equivalent).[29] We add an "IS" curve of similar form

$$\tilde{U}_t = (1-b)U^{Agg}_{t+1,t} + bEU_{t+1} - \sigma(f_t - \pi^{Agg}_{t+1,t} - \overline{\rho}), \tag{7.8}$$

where the aggregate inertial expectation for the driving variable appears in parallel fashion to (7.7), f_t is the short-term nominal policy rate, and $\overline{\rho}$ is the short-term equilibrium real interest rate. The policy rate is determined by a conventional (non- inertial) policy rule

$$f_t = \overline{\rho} + \overline{\pi} + a_\pi(\pi_t - \overline{\pi}) - a_u\tilde{U}_t \ . \tag{7.9}$$

We can envision a set of *N* economic agents who form expectations as suggested by the empirical results in the paper,

$$\pi^i_{t+1,t} = \omega\pi^i_{t,t-1} + (1-\omega)\pi^i_{t+1,t-1} - c\tilde{U}_{t-1} + \varepsilon_{it}, i = 1,\ldots,N \tag{7.10}$$

and similarly for individual expectations of the unemployment/output gap

$$U^i_{t+1,t} = \omega U^i_{t,t-1} + (1-\omega)U^i_{t+1,t-1} + d(f_t - \pi^i_{t+1,t} - \overline{\rho}) + \eta_{it}, i = 1,\ldots,N \ . \tag{7.11}$$

The shocks ε_{it} and η_{it} reflect the idiosyncratic component of the i^{th} forecaster's forecasts, although in principle that component could also be modeled as idiosyncratic variations in the coefficients ω, *c,* and *d*. Equations (7.10) and (7.11) are very close analogues of the expectations regressions in Sections 2–4 above, in which individual expectations for period *t+1* depend on lagged central tendencies of period *t* and period *t+1* forecasts made in period *t-1*. For simplicity, we assume that the coefficients ω, *c* and *d* are the same across the *N* forecasters.[30] In this case, aggregation is trivial, and the aggregate version of equation (7.10) is[31]

$$\pi^{Agg}_{t+1,t} = \omega\pi^{Agg}_{t,t-1} + (1-\omega)\pi^{Agg}_{t+1,t-1} - c\tilde{U}_{t-1} \ . \tag{7.12}$$

The aggregate expectations $[\pi^{Agg}_{t,t-1}, \pi^{Agg}_{t+1,t-1}]$ are the simple averages of the individual expectations in equation (7.10) for the current and next period's inflation and unemployment gap, respectively, as of viewpoint date *t-1*. An analogous expression holds for the aggregate version of the unemployment/output gap expectations $[\tilde{U}^{Agg}_{t,t-1}, \tilde{U}^{Agg}_{t+1,t-1}]$ expressed for individual agents in (7.11).

---

[29] Of course the rational expectations are computed consistent with some fraction of expectations formation defined by $\pi^{Agg}_{t+1,t}$ in equation (7.12), as long as $b \neq 1$. When $b = 1$ as in Figure 7 below, the model depends completely on the rational expectation.

[30] Allowing for greater and perhaps systematic heterogeneity in expectations, as might be suggested by Figure 2, could impart additional dynamics to the system, but those enhancements lie beyond the scope of this paper.

[31] The use of multiple forecasters comport well with the empirical work in the preceding sections. However, for these purposes, we could just as well use a representative agent.

Importantly, none of the individual agents who form inertial expectations in the model know the true model, and none know the current value of the aggregate expectation. In addition, they do not attempt to form higher-order expectations (expectations of other agents' expectations). Such augmentations, while perhaps reasonable, would extend this simple example well beyond the scope of this paper. Equation (7.12) allows expectations to be formed inertially, with more weight on the lagged one-period-ahead expectation or the lagged two-period-ahead expectation, as the weight ω varies between zero and one. Equation (7.7) allows inflation to depend more or less on inertial versus rational expectations, as *b* increases and decreases in size respectively, and the same is true for the unemployment gap in equation (7.8).

Figure 8 examines the properties of this simple model (equations (7.7), (7.8), (7.9), and (7.12) by simulating a disinflation shock. That is, the model variables begin at a steady state with the equilibrium real rate and inflation at two percent, while the inflation target is set to 0 at the beginning of the simulation. The simulation traces the paths of the key model variables in response to this unexpected downshift in the inflation goal, for various values of the parameters ω and *b*. Inspection of equation (7.12) (and its unemployment gap cousin) suggests that, for values of ω like those estimated in the empirical section, this backward-referential expectations behavior can impart considerable persistence to output, inflation, and the policy rate. Figure 8 displays the quantitative implications of this intuition. The green line, which assumes rational expectations exhibits no persistence. The black, red and blue lines, which employ different weights on lagged *t* and *t+1* aggregate expectations (ω and (1-ω), respectively), exhibit considerable persistence in response to a disinflation shock. Thus all of the persistence in this model may be attributed to the contribution from inertial expectations of the type uncovered in the micro survey data.

The conclusion from this simple exercise is that if expectations are formed in a manner consistent with the micro evidence, such intrinsic expectations inertia can account for a sizable fraction of the persistence exhibited by the macroeconomic data. Whether the data suggest that this or other forms of persistence best account for the inertial responses that are present in aggregate data is a topic for additional research.

8. Conclusion

There is little question that expectations lie at the heart of much economic decision-making, and thus at the heart of models of the macroeconomy that hope to reflect such decision-making. How expectations are formed is an open research question. In earlier work, Fuhrer (2017) showed that

empirical estimates of a standard DSGE model preferred inertia in expectations over price indexation or habit formation as a mechanism to explain the persistence of aggregate time series for output, inflation, and interest rates. A question left open in that paper was why and how expectations might exhibit such inertia.

Through examination of data on individuals' and forecasting firms' forecasts, this paper suggests one possible reason for expectational inertia: Individual expectations exhibit significant inefficiency, particularly in the way in which they update information over time. In this paper, we document the inefficient updating to current information, especially the information revealed in previous aggregate expectations, across three well-known surveys of expectations. In doing so, forecasters and households build intrinsic inertia into the expectations process. Sections 5 and 6 examine the possibility that this apparent inefficiency is instead a manifestation of sticky or noisy information. The results in Tables 15-16 suggest that this is not the case. The reason is straightforward: Those models imply that those who update still do so rationally, given their information constraints. The regression results suggest that (a) most professional forecasters update quite frequently, which is not a surprise; (b) some households may not be updating their information sets frequently, also not a surprise; (c) those professional and household forecasters who appear to have updated still do not do so efficiently; and (d) forecast errors appear not to be consistent with a noisy information model, as a number of variables apart from the forecast revision hold significant explanatory power for the errors. Thus revisions are inefficient, but not because of sticky or noisy information.

The last section of this paper shows that building expectations with intrinsic persistence into a relatively standard (but admittedly simple) macroeconomic model can generate the kinds of impulse responses that are commonly found in macroeconomic VARs, without resorting to the bells and whistles that have been added to DSGE models in recent years—price indexation, habit formation, and autocorrelated structural shocks.

While the micro-data results appear quite robust, their implications for macroeconomic dynamics no doubt merit further investigation; this paper provides only simple examples of the possible implications of such expectations behavior in macro models. However, coupled with earlier work), this paper suggests that micro data-based expectations that exhibit intrinsic persistence due to significant inefficiencies might go far in explaining the persistence observed in macro data.

References

Adam, Klaus. 2005. "Learning to Forecast and Cyclical Behavior of Output and Inflation." *Macroeconomic Dynamics* 9(1): 1–27.

Adam, Klaus, and Mario Padula. 2011**. "**Inflation Dynamics and Subjective Expectations in the United States." *Economic Inquiry* 49(1): 13–25.

Andrade, Philippe and Hervé Le Bihan, 2013. "Inattentive professional forecasters," *Journal of Monetary Economics,* 60(8), pp. 967-982.

Banerjee, Abhijit 1992. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics* 107 (3): 797–817.

Batchelor, Roy A. 1986. "Quantitative v. Qualitative Measures of Inflation Expectations." *Oxford Bulletin of Economics and Statistics* 48(2): 99–120.

Batchelor, Roy A., and Pami Dua 1990. "Product Differentiation in the Economic Forecasting Industry," *International Journal of Forecasting* 6(3): 311—316.

Berger, Allen and Spencer Krane 1985. "The Informational Efficiency of Econometric Model Forecasts," *The Review of Economics and Statistics*, Vol. 67, No.1, pp. 128-134.

Bryan, Michael F., and William T. Gavin. 1986. "Models of Inflation Expectations Formation: A Comparison of Household and Economist Forecasts: Comment." *Journal of Money, Credit and Banking* 18(4): 539–544.

Carroll, Christopher D. 2003. "Macroeconomic Expectations of Households and Professional Forecasters." *The Quarterly Journal of Economics*118(1): 269–298

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113(1): 1–45.

Cipriani, Marco, and Antonio Guarino 2014. "Estimating a Structural Model of Herd Behavior in Financial Markets." *American Economic Review* 104(1): 224–251.

Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar, "How Do Firms Form Their Expectations? New Survey Evidence," NBER Working Paper No. 21092, April 2015.

Coibion, Olivier, Yuriy Gorodnichenko, 2015. " Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review* 105, pp. 2644–2678.

Crowe, Christopher, 2010. "Consensus Forecasts and Inefficient Information Aggregation," IMF Working Paper 10/178, July.

Del Negro, Marco, and Stefano Eusepi. 2010. **"**Fitting Observed Inflation Expectations." Federal Reserve Bank of New York Staff Report No. 476. New York: Federal Reserve Bank of New York.

Driscoll**,** John, and Aart Kraay, 1998**. "**Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data," *The Review of Economics and Statistics,* 80(4): 549–560.

Evans, George W., and Seppo Honkapohja. 2001. *Learning and Expectations in Macroeconomics* Princeton NJ: Princeton University Press.

Frydman, Roman, and Edmund S. Phelps, eds., 2013. *Rethinking Expectations: The Way Forward for Macroeconomics* Princeton NJ: Princeton University Press.

Fuhrer, Jeffrey. 2000. "Habit Formation in Consumption and Its Implications for Monetary Policy Models." *American Economic Review* 90 (3): 367–390.

Fuhrer, Jeffrey. 2006. "Intrinsic and Inherited Persistence." *International Journal of Central Banking* 2(3): 49–86.

Fuhrer, Jeffrey. 2011. "Inflation Persistence." Chapter 9 in *Handbook of Monetary Economics*, B. Friedman and M. Woodford, editors, North-Holland, San Diego, CA, pp. 423–483.

Fuhrer, Jeff. 2017. "Expectations as a Source of Macroeconomic Persistence: Evidence from Survey Expectations in Dynamic Macro Models." *Journal of Monetary Economics* **86,** April, pp. 22-35.

Fuster, Andreas, Benjamin Hebert, and David Laibson. 2012. "Natural Expectations, Macroeconomic Dynamics, and Asset Pricing." In *NBER Macroeconomics Annual* 26, ed. Daron Acemoglu and Michael Woodford, 1–48. Chicago: University of Chicago Press.

Galí, Jordi, and Mark Gertler 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44: 195–222.

Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer. 2016. "Expectations and Investment." NBER Macroeconomics Annual, Vol. 30 (2015): 379-442.

Maćkowiak, Bartosz and Mirko Wiederholt, 2009. "Optimal Sticky Prices under Rational Inattention," *American Economic Review* 99(3), 769-803.

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

Mehra, Yash P. 2002. "Survey Measures of Expected Inflation: Revisiting the Issues of Predictive Content and Rationality." Federal Reserve Bank of Richmond *Economic Quarterly* 88(3):17–36.

Milani, Fabio. 2007. "Expectations, Learning and Macroeconomic Persistence." *Journal of Monetary Economics* 54: 2065–2082.

Nordhaus, William. 1987. "Forecasting Efficiency: Concepts and Applications." *The Review of Economics and Statistics*, Vol. 69, No. 4, pp. 667-674.

Orphanides, Athanios, and John C. Williams. 2005. "Imperfect Knowledge, Inflation Expectations, and Monetary Policy." in *The Inflation-Targeting Debate*, ed. Ben S. Bernanke and Michael Woodford, 201–245. Chicago: University of Chicago Press.

Paloviita, Maritta, and Matti Viren, 2013. "Are Individual Survey Expectations Internally Consistent?" National Bank of Poland Working Paper No. 140.

Pesaran, M. Hashem and Martin R. Weale. 2006. "Survey Expectations," in *Handbook of Economic Forecasting*, Gaham Elliott, Clive W.J. Granger, and Allan Timmermann (eds.), North-Holland Press.

Roberts, John M. 1997. "Is Inflation Sticky?" *Journal of Monetary Economics* 39(2): 173–196.

Scharfstein, David, and Jeremy Stein. 1990. "Herd Behavior and Investment." *American Economic Review* 80(3): 465–479.

Sims, Christopher A. 2003. "Implications of Rational Inattention," *Journal of Monetary Economics* 50(3), 665-690.

Sims, Christopher A. 2006. "Rational Inattention: Beyond the Linear-Quadratic Case." *American Economic Review* 96(2): 158–163.

Slobodyan, Sergey, and Raf Wouters. 2012. "Learning in a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models." *American Economic Journal: Macroeconomics* 4(2): 65–101.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97(3): 586–606.

Thomas, Lloyd B., Jr. 1999. "Survey Measures of Expected U.S. Inflation." *Journal of Economic Perspectives* 13(4): 125–144.

Trehan, Bharat. 2015. "Survey Measures of Expected Inflation and the Inflation Process." *Journal of Money, Credit and Banking*, 47(1): 207–222.

Welch, Ivo**.** 1992. "Sequential Sales, Learning, and Cascades." *Journal of Finance* 47 (2): 695–732

Figure 1

Frequency of forecaster participation, CPI inflation 1-qtr. ahead

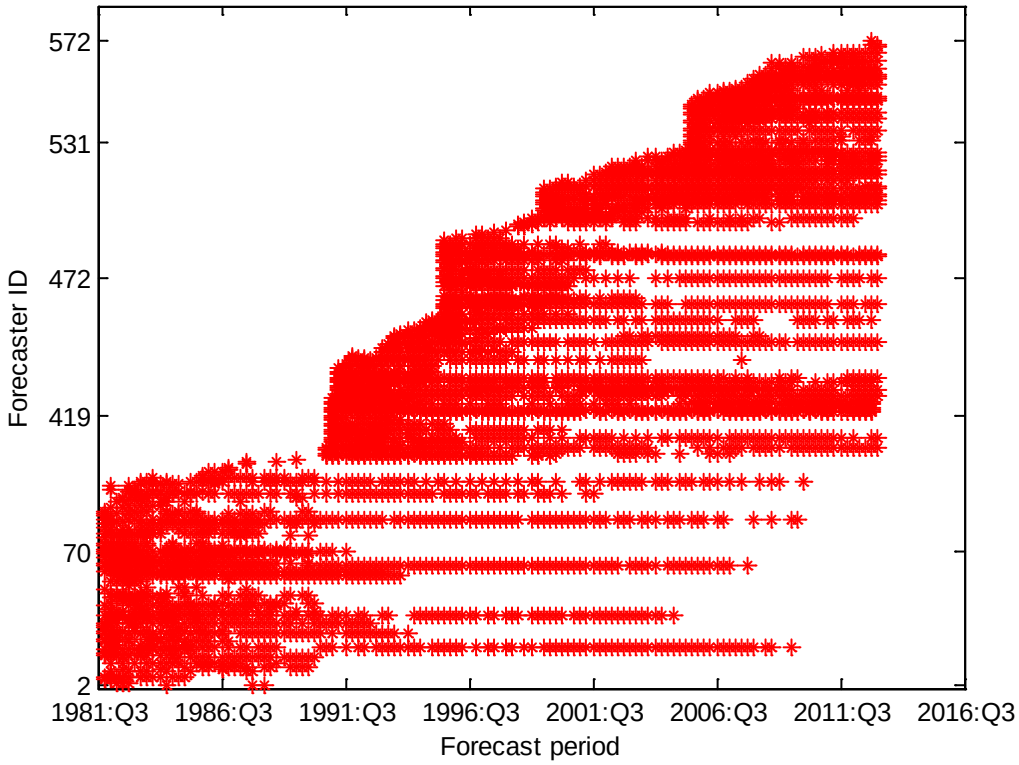


Figure 2

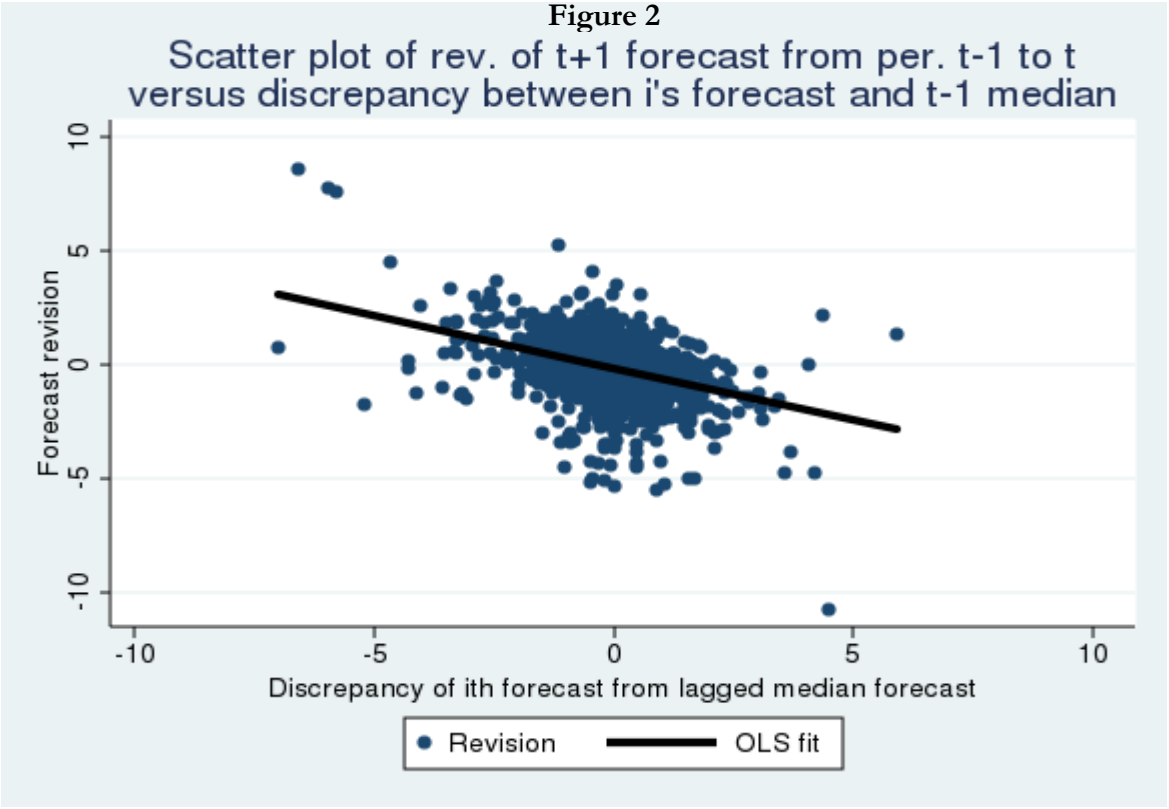


Figure 3
Distribution of individual
forecaster revision coefficients

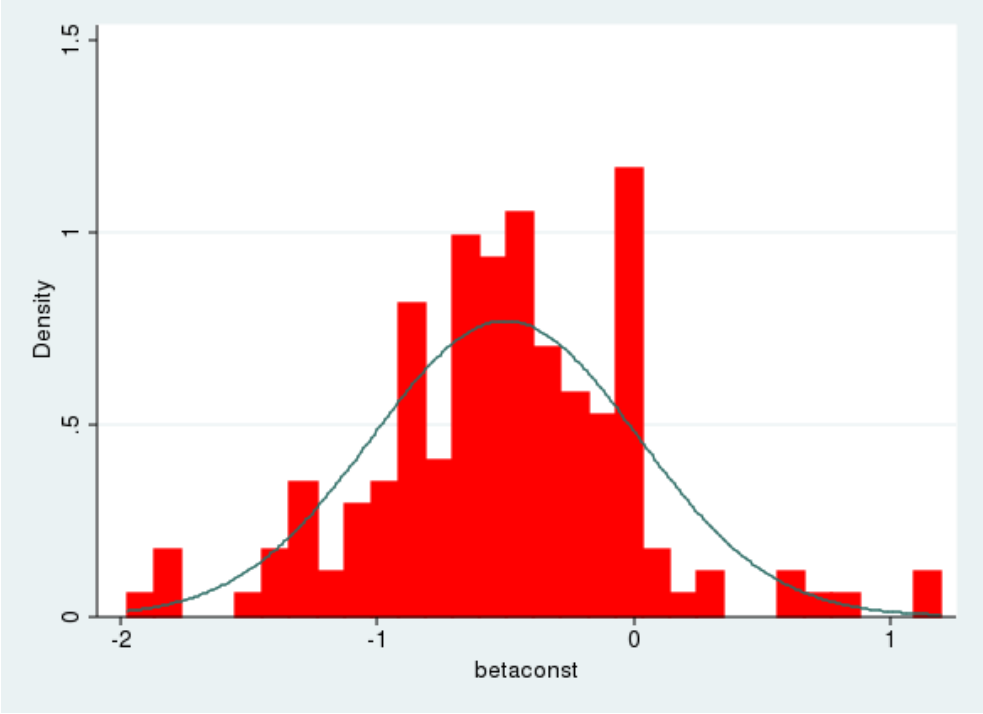


Figure 4

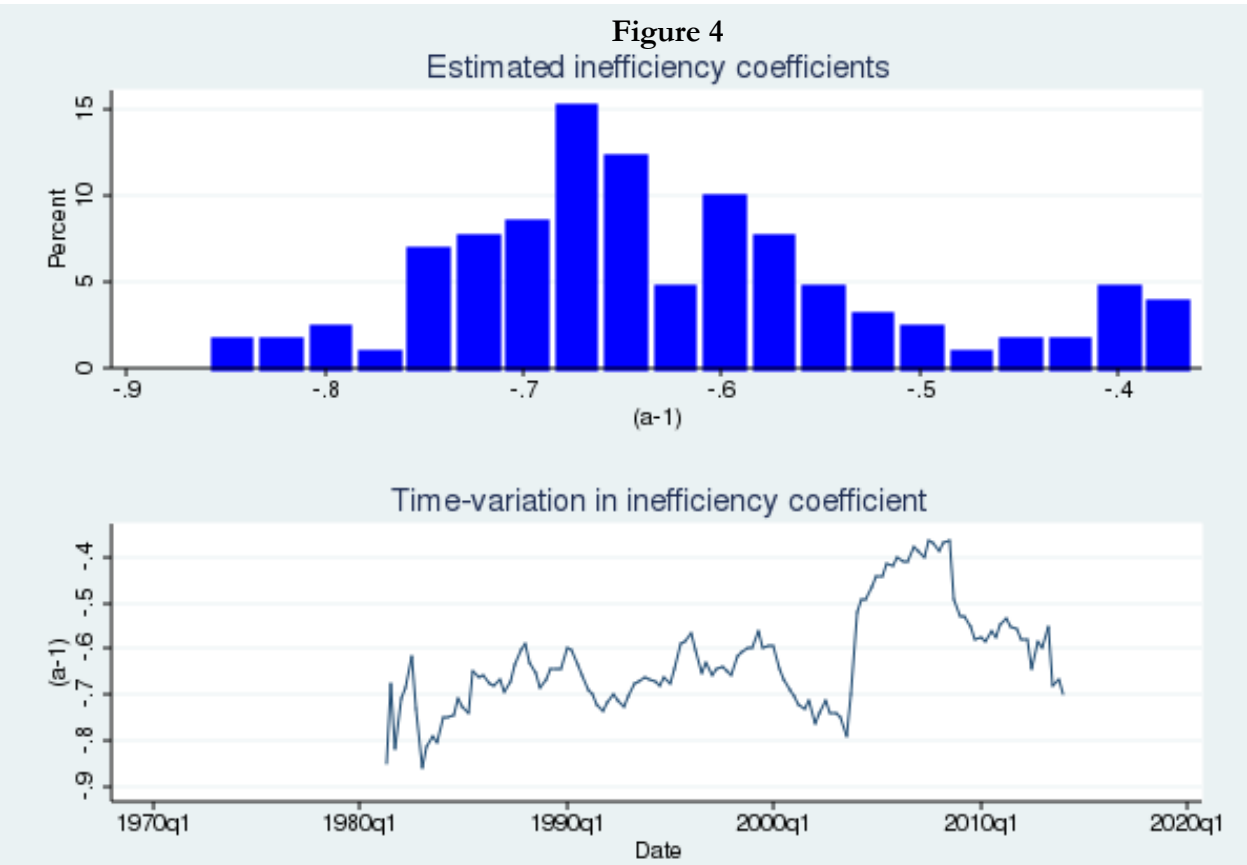
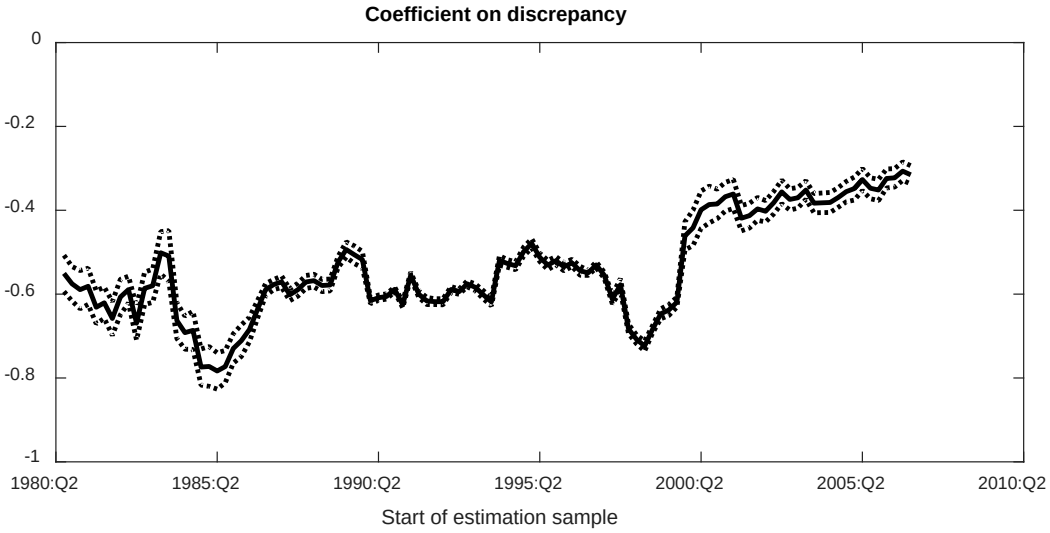

Estimated inefficiency coefficients
Percent
0
5
10
15
-.9
-.8
-.7
-.6
-.5
-.4
(a-1)
Time-variation in inefficiency coefficient
(a-1)
-.9
-.8
-.7
-.6
-.5
-.4
1970q1
1980q1
1990q1
2000q1
2010q1
2020q1
Date

Figure 5
Adding least-squares learning variables to the regressions
(40-quarter rolling regressions)



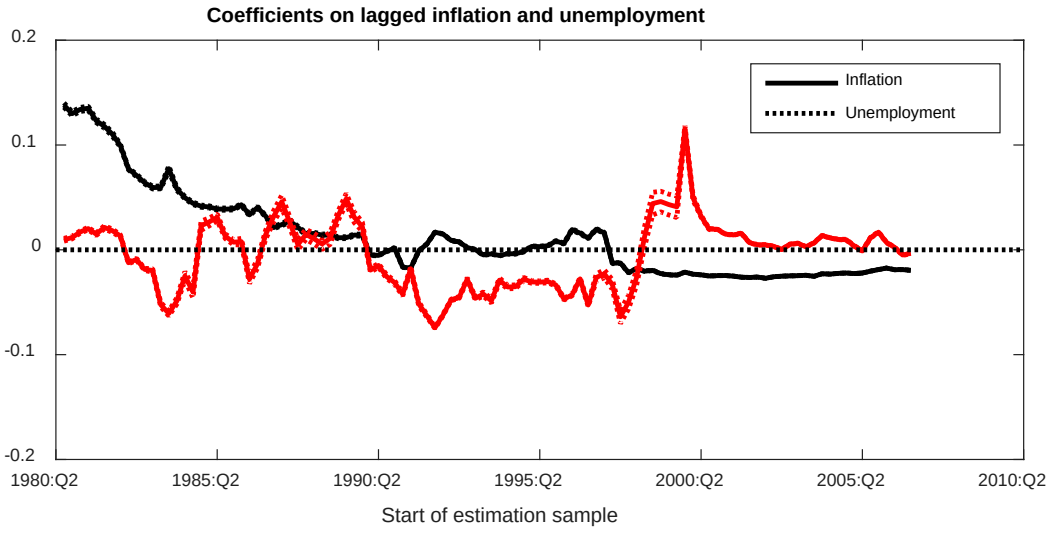


Figure 6
Median 10-year CPI inflation expectation, SPF

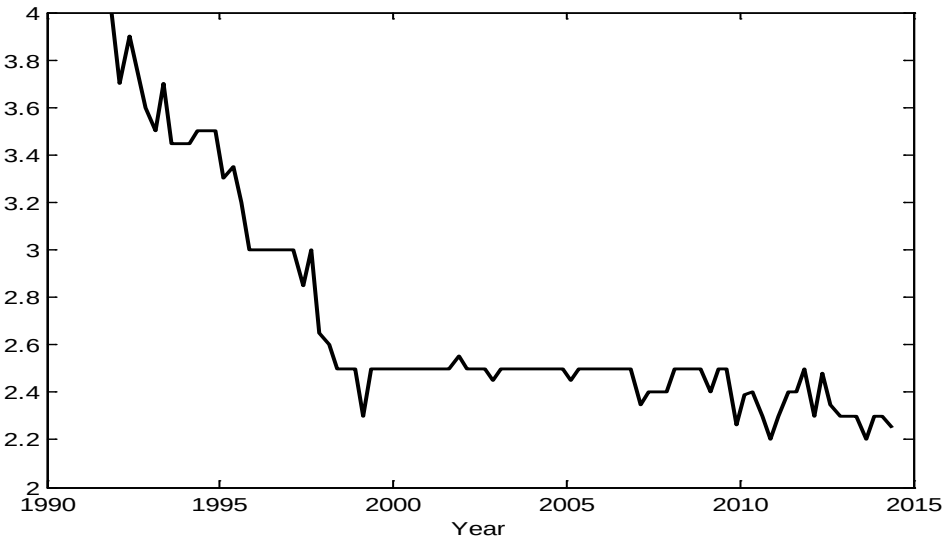


Figure 7

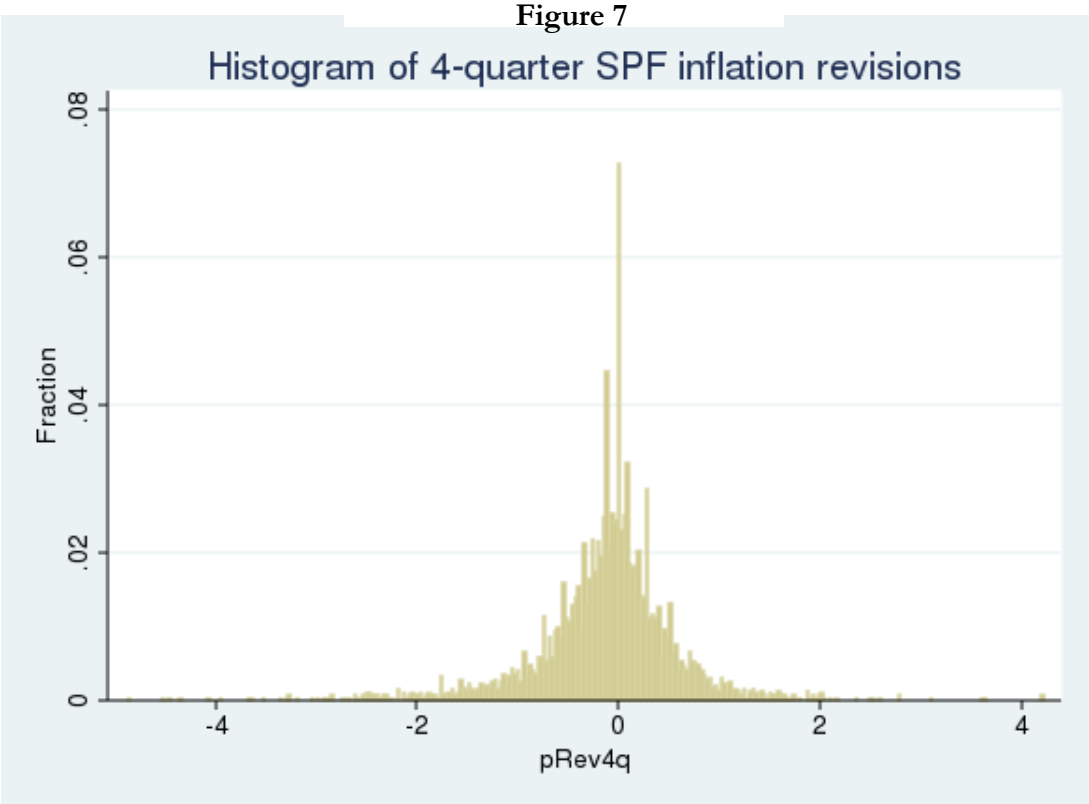


Figure 8
Simulation of simple model with and without inertial expectations effects

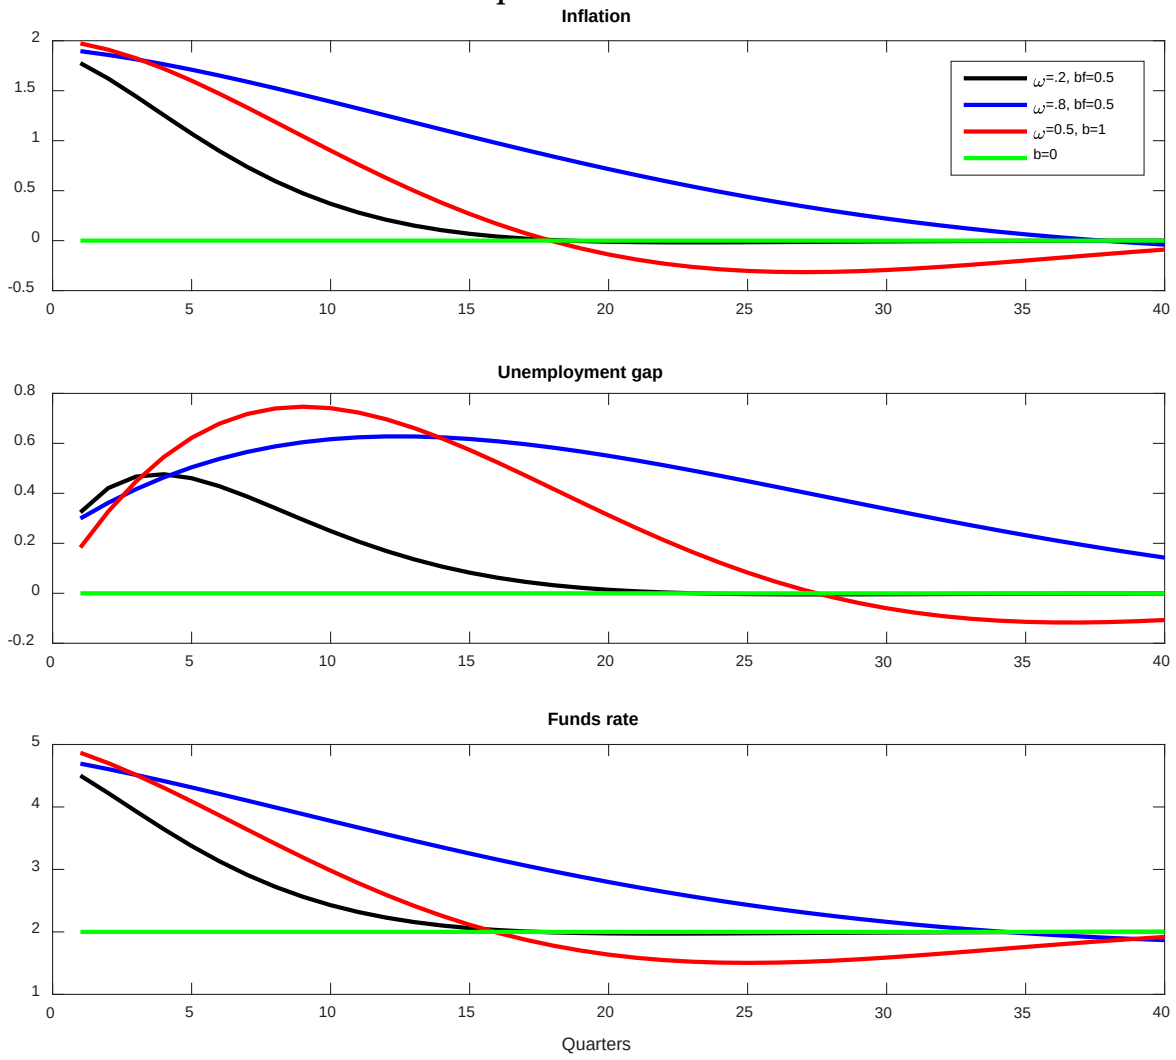


Table 1a Characteristics of SPF sample							
Forecaster participation (number of forecasts submitted, 1968-2013)		Central tendency of forecast (1-qtr. Ahead)					
Inflation, CPI							
N_t=126		1968:Q4		1981:Q3		2012:Q3	
Mean	15.0	Mean	Med.	Mean	Med.	Mean	Med.
Median	8.7	-		7.9	8.0	2.0	2.1
Min, max	1, 70						
Inflation, GDP deflator							
N = 177		1968:Q4		1981:Q3		2012:Q3	
Mean	9.5	Mean	Med.	Mean	Med.	Mean	Med.
Median	5.1	3.0	3.3	7.4	8.5	1.7	1.8
Min, max	1, 71						
Unemployment							
N = 177		1968:Q4		1981:Q3		2012:Q3	
Mean	9.4	Mean	Med.	Mean	Med.	Mean	Med.
Median	4.5	3.8	3.8	7.5	7.5	7.9	7.9
Min, max	1, 71						
Firm type (percentage, N_f=154)[1]							
Financial	45.8						
Nonfinancial	46.4						
Unknown	7.7						
[1] Firm type available only beginning in 1990:Q2 survey							

Table 1b Characteristics of ESPF sample							
Forecaster participation (number of forecasts submitted, 1968-2013)		Central tendency of forecast (1-year ahead)					
Inflation, CPI							
N_t=70		1999:Q1		2007:Q3		2015:Q4	
Mean	39	Mean	Med.	Mean	Med.	Mean	Med.
Median	43	1.3	1.4	2.0	2.0	1.05	1.1
Min, max	1, 69						
Output growth							
N = 70		1999:Q1		2007:Q3		2015:Q4	
Mean	39	Mean	Med.	Mean	Med.	Mean	Med.
Median	43	2.4	2.5	2.3	2.3	1.7	1.7
Min, max	1, 69						
Unemployment							
N = 70		1999:Q1		2007:Q3		2015:Q4	
Mean	39	Mean	Med.	Mean	Med.	Mean	Med.
Median	43	10.5	10.3	6.7	6.7	10.5	10.5
Min, max	1, 69						

Table 2
Inflation forecast dependence on lagged central tendency, with various controls and horizons

$$\pi^{i}_{t+k,t} = a\pi^{i}_{t-1} + b\pi^{i}_{t+k,t-1} + cC(\pi^{i}_{t+k,t-1}) + d\pi^{i}_{t,t-1} + eZ^{i}_{t} + \delta_{i} + \varepsilon^{i}_{t}$$

Variable	$\pi^{i}_{t+1,t}$ (*k=1*)					$\pi^{i}_{t+2,t}$ *(k=2)*	$\pi^{i}_{t+3,t}$ *(k=3)*	$\pi^{i}_{t+4,t}$ *(k=4)*
π^{i}_{t-1}	0.15 (0.007)	0.06 (0.050)	0.06 (0.048)	0.04 (0.029)	0.03 (0.095)	0.05 (0.022)	0.08 (0.000)	0.07 (0.009)
$\pi^{Median}_{t+k,t-1}$		0.73 (0.000)		0.38 (0.000)	0.28 (0.023)	0.37 (0.000)	0.40 (0.000)	0.45 (0.000)
$\pi^{i}_{t+k,t-1}$				0.41 (0.000)	0.43 (0.000)	0.47 (0.000)	0.39 (0.000)	0.32 (0.017)
$\pi^{i}_{t,t-1}$			0.36 (0.000)					
U^{i}_{t-1}					-0.01 (0.653)			
R^{i}_{t-1}					0.03 (0.305)			
ΔY^{i}_{t-1}					0.03 (0.026)			
Test: $\pi^{i}_{t+1,t-1} + \pi^{Median}_{t+1,t-1} = 1$ (b+c=1, *p*-value)				0.000	0.003	0.000	0.000	0.000
Adjusted R-squared	0.048	0.239	0.149	0.320	0.321	0.433	0.421	0.310
Observations	4800	4737	3750	3751	3520	4737	4716	4601
Unemployment forecast dependence on lagged central tendency								
Variable	$U^{i,SPF}_{t+1,t}$					$U^{i,SPF}_{t+2,t}$	$U^{i,SPF}_{t+3,t}$	$U^{i,SPF}_{t+4,t}$
U^{i}_{t-1}	0.93 (0.000)	0.14 (0.142)	0.35 (0.000)	0.08 (0.420)	0.33 (0.008)	-0.03 (0.771)	-0.05 (0.575)	-0.20 (0.352)
$U^{Median}_{t+k,t-1}$		0.86 (0.000)		0.61 (0.000)	0.44 (0.002)	0.59 (0.000)	0.56 (0.000)	0.76 (0.002)
$U^{i}_{t+k,t-1}$				0.31 (0.000)	0.21 (0.000)	0.44 (0.000	0.51 (0.000	0.44 (0.000)
$U^{i}_{t,t-1}$			0.60 (0.000)					
$\pi^{i,SPF}_{t-1}$					-0.00 (0.995)			
$R^{i,SPF}_{t-1}$					0.01 (0.345)			
$\Delta Y^{i,SPF}_{t-1}$					-0.08 (0.000)			
Test: $U^{i,SPF}_{t+k,t-1} + U^{Median}_{t+k,t-1} = 1$ (b+c=1, *p*-value)				0.491	0.009	0.740	0.541	0.420
Adjusted R-squared	0.903	0.934	0.916	0.940	0.956	0.920	0.910	0.900
Observations	7396	7283	5575	5573	3519	5550	5269	3712

Table 2a
Test of revision efficiency, all variables, all horizons, 1981-2016

$$x^i_{t+k,t} = ax^i_{t+k,t-1} + bMedian(x_{t+k,t-1}) + cx^i_{t-1}$$

	Inflation				**Unemployment**			
	k=1	k=2	k=3	k=1	k=1	k=2	k=3	k=1
$x^i_{t+k,t-1}$	0.41 (0.000)	0.47 (0.000)	0.39 (0.000)	0.43 (0.000)	0.31 (0.000)	0.44 (0.000)	0.51 (0.000)	0.21 (0.000)
$Median(x^i_{t+k,t-1})$	0.38 (0.000)	0.37 (0.000)	0.40 (0.000)	0.28 (0.023)	0.61 (0.000)	0.59 (0.000)	0.56 (0.000)	0.44 (0.002)
x^i_{t-1}	0.04 (0.029)	0.05 (0.000)	0.06 (0.000)	0.03 (0.095)	0.08 (0.420)	-0.03 (0.771)	-0.05 (0.575)	0.38 (0.001)
Other variables				Y				Y
Test: a=1 (p-value)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Observations	3751	3734	3649	3452	5573	5550	5269	3519
	Treasury bill rate				**Output growth**			
	t+1	t+2	t+3	t+1	t+1	t+2	t+3	t+1
$x^i_{t+k,t-1}$	0.26 (0.000)	0.33 (0.000)	0.47 (0.000)	0.26 (0.000)	0.26 (0.000)	0.27 (0.000)	0.27 (0.000)	0.24 (0.005)
$Median(x^i_{t+k,t-1})$	0.33 (0.142)	0.26 (0.214)	0.16 (0.218)	0.16 (0.481)	0.86 (0.000)	0.92 (0.000)	0.73 (0.000)	0.78 (0.000)
x^i_{t-1}	0.36 (0.116)	0.37 (0.061)	0.32 (0.010)	0.52 (0.022)	0.08 (0.003)	0.04 (0.194)	0.01 (0.610)	0.12 (0.000)
Other variables				Y				Y
Test: a=1 (p-value)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Observations	3702	3597	3593	3452	5500	5477	5172	3509

Table 3a

Response of forecast revisions to lagged discrepancies between individual forecasts and central tendency measures

$$\pi_{t+1,t}^{i,SPF} - \pi_{t+1,t-1}^{i,SPF} = \delta[\pi_{t+1,t-1}^{i,SPF} - \pi_{t+1,t-1}^{Median}] + a\pi_{t-1}^{i} + cZ_t^i + \delta_i + \varepsilon_t^i$$

Inflation results, 1981-2016

Variable	$\pi_{t+1,t}^{i} - \pi_{t+1,t-1}^{i}$								$\pi_{t+2,t}^{i} - \pi_{t+2,t}^{i}$	$\pi_{t+3,t}^{i} - \pi_{t+3,t-1}^{i}$
$\pi_{t+1,t-1}^{i} - \pi_{t+1\|t-1}^{Median}$	-0.58 (0.000)			-0.63 (0.000)	-0.58 (0.000)	-0.59 (0.000)	-0.56 (0.000)	-0.58 (0.000)		
$\pi_{t+1,t-1}^{i} - \pi_{t+1\|t-1}^{RMSE}$		-0.11 (0.000)		-0.06 (0.054)						
$\pi_{t+1,t-1}^{i} - \pi_{t+1\|t-1}^{Big}$			-0.33 (0.000)	0.02 (0.830)						
π_{t-1}^{i}					0.02 (0.128)	0.04 (0.029)	0.04 (0.035)	-0.04 (0.002)	0.05 (0.000)	0.06 (0.000)
$\pi_{t+1,t-1}^{Median}$						-0.21 (0.000)	-0.31 (0.001)	-0.20 (0.001)		
$\pi_{t+2,t-1}^{i} - \pi_{t+2\|t-1}^{Median}$									-0.53 (0.000)	
$\pi_{t+3,t-1}^{i} - \pi_{t+3\|t-1}^{Median}$										-0.61 (0.000)
$\pi_{t+2,t-1}^{Median}$									-0.16 (0.000)	
$\pi_{t+3,t-1}^{Median}$										-0.20 (0.000)
U_{t-1}^{i}							-0.01 (0.593)	-0.10 (0.263)		
R_{t-1}^{i}							0.04 (0.259)	0.01 (0.921)		
Additional controls								Y		
Adjusted R-squared	0.16	0.01	0.08	0.16	0.16	0.18	0.17	0.23	0.28	0.35
Observations	3762	1926	2729	1591	3751	3751	3508	3331	3734	3649

Estimation sample: 1981:Q3-2016:Q2

"Additional regressor" includes all lagged real-time variables, current and t+1-period forecasts of all variables, all revisions for other variables, all discrepancies for other variables, current and lagged revisions to aggregate forecasts.

Table 3b

Response of forecast revisions to lagged discrepancies between individual forecasts and central tendency measures, UNEMPLOYMENT Results, SPF, 1981-2016

$$U_{t+1,t}^{i} - U_{t+1,t-1}^{i} = \delta[U_{t+1,t-1}^{i} - U_{t+1,t-1}^{Median}] + aU_{t-1}^{i} + cZ_t^i + \delta_i + \mu_t + \varepsilon_t^i$$

Variable	$U_{t+1,t}^{i} - U_{t+1,t-1}^{i}$					$U_{t+2,t}^{i} - U_{t+2,t-1}^{i}$	$U_{t+3,t}^{i} - U_{t+3,t-1}^{i}$
$U_{t+1,t-1}^{i} - U_{t+1\|t-1}^{Median}$	-0.68 (0.000)	-0.33 (0.000)	-0.69 (0.000)	-0.75 (0.000)	-0.83 (0.000)		
U_{t-1}^{i}	0.01 (0.531)		0.08 (0.420)	0.08 (0.710)	-0.27 (0.000)	-0.03 (0.771)	-0.05 (0.575)

$U^{Median}_{t+1,t-1}$			-0.08 (0.491)	-0.07 (0.752)	-0.71 (0.000)		
$U^{i}_{t+2,t-1} - U^{Median}_{t+2\|t-1}$						-0.56 (0.000)	
$U^{i}_{t+3,t-1} - U^{Median}_{t+3,t-1}$							-0.49 (0.000)
$U^{Median}_{t+2,t-1}$						0.04 (0.740)	
$U^{Median}_{t+3,t-1}$							0.07 (0.541)
$\pi^{i}_{t-1,t}$				-0.01 (0.673)	-0.00 (0.880)		
$R^{i}_{t-1,t}$				0.01 (0.340)	0.01 (0.475)		
Additional controls					Y		
Adjusted R-squared	0.21	0.38	0.21	0.22	0.15	0.15	0.90
Observations	5573	2942	5573	3587	5269	5550	5269
Estimation sample: 1981:Q3-2016:Q2 "Additional regressor" includes all lagged real-time variables, current and t+1-period forecasts of all variables, all revisions for other variables, all discrepancies for other variables, current and lagged revisions to aggregate forecasts.							

Table 4

Regression of change in k-period-ahead forecast ($\pi^{i}_{t+k,t} - \pi^{i}_{t+k-1|t-1}$) on lagged discrepancy and other controls, 1981-2016

	k=1				*k=2*	*k=3*
$\pi^{i}_{t,t-1} - \pi^{Median}_{t|t-1}$	-0.80 (0.000)	-0.80 (0.000)	-0.82 (0.000)	-0.87 (0.000)		
$\pi^{i}_{t+1,t-1} - \pi^{Median}_{t+1|t-1}$					-0.82 (0.000)	
$\pi^{i}_{t+2,t-1} - \pi^{Median}_{t+2|t-1}$						-0.69 (0.000)
π^{i}_{t-1}		-0.05 (0.053)	0.01 (0.749)	-0.05 (0.002)	-0.00 (0.826)	0.05 (0.000)
$\pi^{Median}_{t,t-1}$			-0.33 (0.000)	-0.59 (0.000)		
$\pi^{Median}_{t+1,t-1}$					-0.26 (0.000)	
$\pi^{Median}_{t+2,t-1}$						-0.24 (0.000)
$\pi^{i}_{t,t}$				0.25 (0.000)	0.12 (0.000)	0.05 (0.141)
$U^{i}_{t,t}$				0.01 (0.665)		
$R^{i}_{t,t}$				0.06 (0.020)		
Adjusted R-sq.	0.28	0.29	0.30	0.42	0.33	0.31
Observations	3761	3750	3750	3504	3749	3720

Controlling for common information						
	k=1 change		*k=2* change		*k=3* change	
$\pi^{i}_{t,t-1} - \pi^{Median}_{t|t-1}$	-0.85 (0.000)	-0.85 (0.000)				
$\pi^{i}_{t+1,t-1} - \pi^{Median}_{t+1|t-1}$			-0.79 (0.000)	-0.79 (0.000)		
$\pi^{i}_{t+2,t-1} - \pi^{Median}_{t+2|t-1}$					-0.66 (0.000)	-0.67 (0.000)
$\pi^{Median}_{t+1,t-1} - \pi^{Median}_{t+1|t-2}$	-0.20 (0.026)					
$\pi^{Median}_{t+1,t} - \pi^{Median}_{t+1|t-1}$		0.42 (0.004)				
$\pi^{Median}_{t+2,t-1} - \pi^{Median}_{t+2|t-2}$			0.04 (0.689)			
$\pi^{Median}_{t+2,t} - \pi^{Median}_{t+2|t-1}$				0.64 (0.000)		
$\pi^{Median}_{t+3,t-1} - \pi^{Median}_{t+3|t-2}$					-0.22 (0.003)	
$\pi^{Median}_{t+3,t} - \pi^{Median}_{t+3|t-1}$						0.59 (0.000)
$\pi^{i}_{t,t}$	0.25 (0.000)	0.20 (0.000)	0.11 (0.000)	0.08 (0.005)	0.05 (0.102)	0.04 (0.215)
Adjusted R-squared	0.42	0.43	0.29	0.32	0.30	0.32
Observations	3469	3504	3471	3506	3448	3483

All regressions include lagged real-time unemployment, inflation, T-bill rate, as well as the lagged viewpoint date median forecast for the horizon indicated.

Table 5
The effect of common information
Response of forecast revisions to lagged discrepancies between individual forecasts and central tendency measures, controlling for revision in aggregate forecast, 1981-2016

$$\pi_{t+1,t}^{i,SPF} - \pi_{t+1,t-1}^{i,SPF} = \gamma[\pi_{t+1,t-2}^{Median} - \pi_{t+1|t-1}^{Median}] + \delta[\pi_{t+1,t-1}^{i,SPF} - C(\pi_{t+1,t-1})] + a\pi_{t-1}^{i} + cZ_t^{i} + \delta_i + \mu_t + \varepsilon_t^{i}$$

Inflation results

Variable	Lagged revision				Contemporaneous revision			
$\pi_{t+1,t-1}^{i} - \pi_{t+1|t-1}^{Median}$	-0.58 (0.000)	-0.57 (0.000)	-0.58 (0.000)	-0.55 (0.000)	-0.59 (0.000)	-0.60 (0.000)	-0.57 (0.000)	-0.57 (0.000)
$\pi_{t+1,t-1}^{Median} - \pi_{t+1|t-2}^{Median}$		0.11 (0.403)	0.17 (0.210)	0.21 (0.165)				
$\pi_{t+1,t}^{Median} - \pi_{t+1|t-1}^{Median}$					0.91 (0.000)	0.88 (0.000)	0.87 (0.000)	0.59 (0.011)
π_{t-1}^{i}	0.02 (0.128)	0.02 (0.359)	0.03 (0.097)	0.03 (0.168)	-0.01 (0.502)	0.00 (0.934)	0.00 (0.896)	
$\pi_{t+1,t-1}^{Median}$			-0.25 (0.000)	-0.38 (0.001)		-0.07 (0.004)	-0.08 (0.072)	
Additional forecast variables	N	N	N	Y	N	N	Y	Instrumented
Adjusted R-squared	0.16	0.16	0.18	0.17	0.29	0.29	0.28	-
Observations	3751	3710	3710	3473	3751	3751	3508	3720
* "Additional forecast variables" includes real-time estimates of lagged unemployment, Treasury bill rate.								

Unemployment results

Variable	Lagged revision				Contemporaneous revision			
$U_{t+1,t-1}^{i} - U_{t+1|t-1}^{Median}$	-0.68 (0.000)	-0.64 (0.000)	-0.67 (0.000)	-0.72 (0.000)	-0.66 (0.000)	-0.66 (0.000)	-0.70 (0.000)	-0.66 (0.000)
$U_{t+1,t-1}^{Median} - U_{t+1|t-2}^{Median}$		0.45 (0.000)	0.52 (0.000)	0.62 (0.000)				
$U_{t+1,t}^{Median} - U_{t+1|t-1}^{Median}$					0.96 (0.000)	0.96 (0.000)	0.99 (0.000)	1.01 (0.000)
U_{t-1}^{i}	0.01 (0.531)	-0.01 (0.276)	0.23 (0.000)	0.41 (0.000)	0.00 (0.572)	-0.01 (0.102)	-0.00 (0.846)	
$U_{t+1,t-1}^{i}$			-0.27 (0.000)	-0.45 (0.000)		0.02 (0.061)	0.00 (0.903)	
Additional forecast variables	N	N	N	Y	N	N	Y	Instrumented
Adjusted R-squared	0.21	0.36	0.39	0.46	0.77	0.77	0.80	-
Observations	5573	5504	5504	3571	5573	5573	3587	5514
* "Additional forecast variables" includes real-time estimates of lagged inflation, Treasury bill rate.								

Table 6

Effect of common information and all other revisions

Response of forecast revisions to lagged discrepancies between individual forecasts and central tendency measures, controlling for revision in aggregate forecast and in lagged and period-t estimates

$$\pi_{t+1,t}^{i,SPF} - \pi_{t+1,t-1}^{i,SPF} = \gamma[\pi_{t+1,t-2}^{Median} - \pi_{t+1|t-1}^{Median}] + \delta[\pi_{t+1,t-1}^{i,SPF} - C(\pi_{t+1,t-1})] + a\pi_{t-1}^{i} + cZ_t^i + \delta_i + \mu_t + \varepsilon_t^i$$

	$\pi_{t+1,t}^{i} - \pi_{t+1,t-1}^{i}$	$\pi_{t+2,t}^{i} - \pi_{t+2,t}^{i}$	$\pi_{t+3,t}^{i} - \pi_{t+3,t-}^{i}$	$U_{t+1,t}^{i} - U_{t+1,t-}^{i}$	$U_{t+2,t}^{i} - U_{t+2,t}^{i}$	$U_{t+3,t}^{i} - U_{t+3,t}^{i}$
$\pi_{t+k,t-1}^{i} - \pi_{t+k\|t-1}^{Median}$	-0.35 (0.000)	-0.36 (0.000)	-0.43 (0.000)	-0.40 (0.000)	-0.35 (0.000)	-0.37 (0.000)
$\pi_{t+k,t-1}^{Median} - \pi_{t+k,t-2}^{Median}$	-0.07 (0.440)	0.02 (0.867)	-0.15 (0.078)	0.18 (0.001)	0.30 (0.000)	0.27 (0.000)
Adjusted R-squared	0.197	0.233	0.265	0.631	0.580	0.550
Observations	2779	2761	2678	2813	2791	2699
Contemporaneous revisions to aggregate forecasts						
$\pi_{t+k,t-1}^{i} - \pi_{t+k\|t-1}^{Median}$	-0.58 (0.000)	-0.54 (0.000)	-0.55 (0.000)	-0.64 (0.000)	-0.57 (0.000)	-0.51 (0.000)
$\pi_{t+1,t}^{Median} - \pi_{t+1,t-1}^{Median}$	0.84 (0.000)	0.79 (0.000)	0.73 (0.000)	0.90 (0.000)	0.85 (0.000)	0.86 (0.000)
Adjusted R-squared	0.297	0.282	0.296	0.790	0.731	0.709
Observations	2779	2761	2678	2813	2791	2699

Additional variables include revisions of lagged inflation, unemployment, Treasury bill, output growth; Revisions to current period forecasts for the same; *t-1* viewpoint date forecast of inflation or output for period $t+k$; and *t*-period individual estimates of lagged inflation, unemployment, Treasury bill, and output growth.

Table 7
Learning versus lagged central tendencies
Response of forecast revisions to lagged discrepancies between individual forecasts and central tendency measure, with lagged real-time actual data

Sub-sample estimates

	$\pi^i_{t+1,t} - \pi^i_{t+1,t-1}$					$U^i_{t+1,t} - U^i_{t+1,t-1}$				
Sample	Full sample	1990-	1995-	2000-	2005-	Full sample	1990-	1995-	2000-	2005-
$\pi^i_{t+1,t-1} - \pi^{Median}_{t+1\|t-1}$	-0.57 (0.000)	-0.51 (0.000)	-0.50 (0.000)	-0.50 (0.000)	-0.50 (0.000)					
$U^i_{t+1,t-1} - U^{Median}_{t+1\|t-1}$						-0.70 (0.000)	-0.70 (0.000)	-0.71 (0.000)	-0.72 (0.000)	-0.75 (0.000)
$\pi^{Median}_{t+1,t} - \pi^{Median}_{t+1,t-1}$	0.88 (0.000)	0.86 (0.000)	0.86 (0.000)	0.85 (0.000)	0.87 (0.000)					
$U^{Median}_{t+1,t} - U^{Median}_{t+1,t-1}$						0.96 (0.000)	0.93 (0.000)	0.95 (0.000)	0.95 (0.000)	0.94 (0.000)
Observations	3430	2964	2512	1976	1499	3497	3080	2610	2056	1550
Additional controls include π^i_{t-1}, $\pi^i_{t+1,t-1}$ for inflation, $U^i_{t-1}, U^i_{t+1,t-1}$ for unemployment										

Table 8
Response of forecast revisions to lagged discrepancies between individual forecasts and central tendency measures, INFLATION Results, Euro SPF, 1999-2016

	Dependent variable (forecast revisions)					
Regressor	$\pi^{i}_{y1,t}-\pi^{i}_{y1,t-1}$	$\pi^{i}_{y2,t}-\pi^{i}_{y2,t-1}$	$\pi^{i}_{y1,t}-\pi^{i}_{y1,t-1}$	$\pi^{i}_{y2,t}-\pi^{i}_{y2,t-1}$	$\pi^{i}_{y1,t}-\pi^{i}_{y1,t-1}$	$\pi^{i}_{y2,t}-\pi^{i}_{y2,t-1}$
$\pi^{i}_{y1,t-1}-\pi^{Median}_{y1,t-1}$	-0.56 (0.000)		-0.59 (0.000)		-0.54 (0.000)	
$\pi^{i}_{y2,t-1}-\pi^{Median}_{y2,t-1}$		-0.47 (0.000)		-0.48 (0.000)		-0.50 (0.000)
π_{t-1}	0.18 (0.001)	0.05 (0.000)	0.19 (0.000)	0.05 (0.013)	0.23 (0.000)	0.08 (0.000)
Additional controls						
$\pi^{Median}_{yk,t-1}$			Y	Y	Y	Y
Unemployment discrepancy					Y	Y
Exogenous assumptions					Y	Y
Output and unemployment forecasts					Y	Y
Adjusted R-squared	0.19	0.23	0.29	0.25	0.45	0.32
Observations	3075	850	2891	828	1921	587

Table 9
Response of forecast revisions to lagged discrepancies between individual forecasts and central tendency measures, UNEMPLOYMENT Results, Euro SPF, 1999-2016

	Dependent variable (forecast revisions)					
Regressor	$U^{i}_{y1,t}-U^{i}_{y1,t-1}$	$U^{i}_{y2,t}-U^{i}_{y2,t-1}$	$U^{i}_{y1,t}-U^{i}_{y1,t-1}$	$U^{i}_{y2,t}-U^{i}_{y2,t-1}$	$U^{i}_{y1,t}-U^{i}_{y1,t-1}$	$U^{i}_{y2,t}-U^{i}_{y2,t-1}$
$U^{i}_{y1,t-1}-U^{Median}_{y1,t-1}$	-0.38 (0.000)		-0.24 (0.000)		-0.39 (0.000)	
$U^{i}_{y2,t-1}-U^{Median}_{y2,t-1}$		-0.37 (0.000)		-0.30 (0.000)		-0.41 (0.000)
U_{t-1}	-0.24 (0.118)	-0.04 (0.555)	-1.09 (0.000)	-0.96 (0.000)	-0.25 (0.002)	-0.07 (0.397)
Additional controls						
$U^{Median}_{yk,t-1}$			Y	Y	Y	Y
Inflation discrepancy					Y	Y
Exogenous assumptions					Y	Y
Output and unemployment forecasts					Y	Y
Adjusted R-squared	0.16	0.15	0.67	0.47	0.36	0.37
Observations	2929	779	2916	777	1921	578

Table 10
Response of forecast revisions to lagged discrepancies between individual forecasts and central tendency measures, OUTPUT GROWTH Results, Euro SPF, 1999-2016

	Dependent variable (forecast revisions)					
Regressor	$\Delta y^{i}_{y1,t} - \Delta y^{i}_{y1,t-1}$	$\Delta y^{i}_{y2,t} - \Delta y^{i}_{y2,t-1}$	$\Delta y^{i}_{y1,t} - \Delta y^{i}_{y1,t-1}$	$\Delta y^{i}_{y2,t} - \Delta y^{i}_{y2,t-1}$	$\Delta y^{i}_{y1,t} - \Delta y^{i}_{y1,t-1}$	$\Delta y^{i}_{y2,t} - \Delta y^{i}_{y2,t-1}$
$\Delta y^{i}_{y1,t-1} - \Delta y^{Median}_{y1,t-1}$	-0.69 (0.000)		-0.75 (0.000)		-0.78 (0.000)	
$\Delta y^{i}_{y2,t-1} - \Delta y^{Median}_{y2,t-1}$		-0.53 (0.000)		-0.55 (0.000)		-0.58 (0.000)
Δy_{t-1}	0.16 (0.164)	-0.02 (0.015)	0.04 (0.625)	-0.07 (0.000)	0.14 (0.068)	0.00 (0.671)
Additional controls						
$\Delta y^{Median}_{yk,t-1}$			Y	Y	Y	Y
Inflation discrepancy					Y	Y
Exogenous assumptions					Y	Y
Output and unemployment forecasts					Y	Y
Adjusted R-squared	0.15	0.13	0.30	0.24	0.42	0.34
Observations	2917	819	2813	802	1921	589

Table 11
Effect of common information: Response of forecast revisions to lagged discrepancies between individual forecasts and central tendency measures, Euro SPF, with revisions to aggregate forecast, 1999-2016

	Dependent variable (forecast revisions)					
Regressor	$\pi^{i}_{y1,t} - \pi^{i}_{y1,t-1}$	$\pi^{i}_{y2,t} - \pi^{i}_{y2,t-1}$	$U^{i}_{y1,t} - U^{i}_{y1,t-1}$	$U^{i}_{y2,t} - U^{i}_{y2,t-1}$	$\Delta y^{i}_{y1,t} - \Delta y^{i}_{y1,t-1}$	$\Delta y^{i}_{y2,t} - \Delta y^{i}_{y2,t-1}$
$X^{i}_{y1,t-1} - X^{Median}_{y1,t-1}$	-0.53 (0.000)		-0.52 (0.000)		-0.67 (0.000)	
$X^{i}_{y2,t-1} - X^{Median}_{y2,t-1}$		-0.47 (0.000)		-0.40 (0.000)		-0.56 (0.000)
$X^{Median}_{y1,t} - X^{Median}_{y1,t-1}$	0.94 (0.000)		0.95 (0.000)		0.98 (0.000)	
$X^{Median}_{y2,t} - X^{Median}_{y2,t-1}$		0.65 (0.000)		0.97 (0.000)		0.96 (0.000)
Additional controls						
$X^{Median}_{y1,t-1}$	-0.02 (0.448)	-0.12 (0.024)	0.03 (0.103)	0.07 (0.066)	-0.02 (0.179)	0.02 (0.402)
Adjusted R-squared	0.48	0.30	0.65	0.46	0.78	0.31
Observations	3075	850	2929	779	2917	819

Table 12
Regression of revision in 12-month inflation forecast (from current interview to 6-months previous) on discrepancy between last inflation forecast and lagged mean/median, as well as other controls
1978:Jan-2017:Apr

	Full sample									
	With lagged discrep.	With lagged median forecast	All indiv. con-trols	Add aggre-gate revs.	1985-forward	1995-forward	2000-forward	2005-forward	Recess. only	Non-recess.
$\pi^{Mich}_{1Y,t-1}$ - $Median(\pi^{Mich}_{1Y,t-1})$	-0.72 (0.000)	-0.72 (0.000)	-0.69 (0.000)	-0.69 (0.000)	-0.69 (0.000)	-0.70 (0.000)	-0.69 (0.000)	-0.69 (0.000)	-0.67 (0.000)	-0.69 (0.000)
$Median(\pi^{Mich}_{1Y,t-1})$		-0.42 (0.000)	-0.48 (0.000)	0.07 (0.055)	-0.84 (0.000)	-0.86 (0.000)	-0.93 (0.000)	-1.07 (0.000)	-0.60 (0.000)	-0.43 (0.000)
Revision to family income, 1-yr. expec.			0.00 (0.783)	0.00 (0.842)	0.00 (0.860)	0.00 (0.019)	0.00 (0.016)	0.00 (0.143)	0.00 (0.823)	0.00 (0.843)
Revision to 5-year inflation expec.			0.20 (0.000)	0.19 (0.000)	0.21 (0.000)	0.27 (0.000)	0.28 (0.000)	0.28 (0.000)	0.21 (0.000)	0.19 (0.000)
Aggregate revision				0.81 (0.000)						
Test of EC restriction	0.000	0.000	0.000	-	0.000	0.000	0.000	0.000	0.000	0.000
Adjusted R-squared	0.427	0.432	0.470	0.480	0.471	0.468	0.443	0.452	0.420	0.482
Observations	84217	84217	56912	56912	51564	40278	30834	22198	7117	49795

Simple test of revision efficiency

$$\pi^{Mich}_{kY,t} = a\pi^{Mich}_{kY,t-1} + bMedian(\pi^{Mich}_{kY,t-1}); k = 1,5$$

Test: $a = 1$

	One-year forecast		
	Coefficient		p-value of test $a = 1$
$\pi^{Mich}_{1Y,t-1}$ (a)	0.29 (0.000)	0.28 (0.000)	0.000
$Median(\pi^{Mich}_{1Y,t-1})$ (b)		0.60 (0.000)	0.000
	Five-year forecast		
	Coefficient		p-value of test $a = 1$
$\pi^{Mich}_{5Y,t-1}$ (a)	0.33 (0.000)	0.30 (0.000)	0.000
$Median(\pi^{Mich}_{5Y,t-1})$ (b)		0.77 (0.000)	0.000

Table 13
"Anchoring" regressions
SPF inflation forecast revisions, varying horizons
Revision regressions with the revision in the long-term (10-year) forecast, full sample

	Revision				Revision			
	t	*t+1*	*t+2*	*t+3*	*t*	*t+1*	*t+2*	*t+3*
$\pi^{i}_{t,t-1} - \pi^{Median}_{t\|t-1}$	-0.60 (0.000)				-0.66 (0.000)			
$\pi^{i}_{t+1,t-1} - \pi^{Median}_{t+1\|t-1}$		-0.48 (0.000)				-0.48 (0.000)		
$\pi^{i}_{t+2,t-1} - \pi^{Median}_{t+2\|t-1}$			-0.45 (0.000)				-0.45 (0.000)	
$\pi^{i}_{t+3,t-1} - \pi^{Median}_{t+3\|t-1}$				-0.54 (0.000)				-0.55 (0.000)
Lagged revision in 10-year aggregate forecast	-0.33 (0.563)	0.33 (0.096)	0.15 (0.395)	0.03 (0.906)	-0.53 (0.334)	0.30 (0.174)	0.10 (0.612)	-0.03 (0.884)
Other controls	N	N	N	N	Y	Y	Y	Y
Adjusted R-squared	0.11	0.10	0.15	0.18	0.20	0.12	0.17	0.23
Observations	2785	2784	2772	2704	2591	2590	2582	2539
Post-1999 sample								
	t	*t+1*	*t+2*	*t+3*	*t*	*t+1*	*t+2*	*t+3*
$\pi^{i}_{t,t-1} - \pi^{Median}_{t\|t-1}$	-0.61 (0.000)				-0.66 (0.000)			
$\pi^{i}_{t+1,t-1} - \pi^{Median}_{t+1\|t-1}$		-0.48 (0.000)				-0.47 (0.000)		
$\pi^{i}_{t+2,t-1} - \pi^{Median}_{t+2\|t-1}$			-0.44 (0.000)				-0.43 (0.000)	
$\pi^{i}_{t+3,t-1} - \pi^{Median}_{t+3\|t-1}$				-0.55 (0.000)				-0.56 (0.000)
Lagged revision in 10-year aggregate forecast	-1.09 (0.319)	0.23 (0.468)	0.00 (0.993)	-0.15 (0.573)	-0.81 (0.438)	0.37 (0.382)	0.13 (0.464)	-0.03 (0.890)
Other controls	N	N	N	N	Y	Y	Y	Y
Adjusted R-squared	0.10	0.09	0.13	0.17	0.22	0.12	0.18	0.22
Observations	1919	1919	1913	1872	1768	1768	1766	1748

Table 14

Michigan survey, one-year ahead inflation expectations
Test for "anchoring" to long-run (2- to 5-year) median expectations

	(1)	(2)	(3)	(4)	(6)
Lagged median 1-yr. expec.	0.84 (0.000)	0.79 (0.000)	0.59 (0.000)	0.59 (0.000)	0.53 (0.000)
Lagged median 2-5-yr. expec.	0.30 (0.000)	0.31 (0.000)	0.35 (0.000)	0.36 (0.000)	0.42 (0.000)
Unemp. controls		Y	Y	Y	Y
Income, financial controls			Y	Y	Y
In previous survey?				Y	Y
Interaction terms					Y
Adjusted R-squared	0.041	0.054	0.091	0.092	0.107
Observations	174601	174601	44305	44305	42811

Table 15

Percentage of forecasters whose revision equals zero

SPF (1981-2014)				Michigan (1978-2016)	Euro SPF (1999-2016)		
One-quarter		Four-quarter		One-year	0, 1, 2, 5-year (joint)		
Inflation	Unemp.	Inflation	Unemp.	Inflation	Inflation	Unemp.	Output growth
18.5	20.1	5.9	6.8	32.7	33.6	29.2	9.2

Table 16
Test regressions for sticky and noisy information models

$$Error^{i}_{t+k} \equiv x_{t+k} - x^{i}_{t+k,t} = +\alpha x^{Median}_{t-1} + \beta R^{i}_{t+k,t} + \gamma x^{i}_{t+k,t-1} \mid R^{i}_{t+k,t} \neq 0$$

$$R^{i}_{t+k,t} \equiv x^{i}_{t+k,t} - x^{i}_{t+k,t-1}$$

SPF forecasts

	Inflation errors					**Unemployment errors**				
	t	t+1	t+1	t+2	t+3	t	t+1	t+1	t+2	t+3
$Median(x^{i}_{t+k,t-1})$ $[\alpha]$	0.18 (0.093)	0.46 (0.168)	1.62 (0.295)	0.21 (0.308)	0.20 (0.218)	0.01 (0.013)	0.04 (0.018)	0.04 (0.018)	0.08 (0.013)	0.13 (0.011)
$R^{i}_{t+k,t}[\beta]$	0.04 (0.725)	0.58 (0.000)	0.90 (0.000)	0.67 (0.000)	0.55 (0.000)	-0.07 (0.307)	-0.20 (0.108)	-0.20 (0.108)	-0.34 (0.038)	-0.46 (0.008)
$x^{i}_{t+k,t-1}[\gamma]$	0.07 (0.120)	0.27 (0.000)	-0.07 (0.168)	0.35 (0.000)	0.39 (0.000)					
Additional *t*-period information			Y					Y		
Test, non-revision variables = 0			0.000					0.000		
R-squared	0.06	0.12	0.29	0.11	0.11	0.04	0.08	0.22	0.11	0.12
Observations	3325	3123	2859	2983	2867	3269	3267	2889	3226	3084
	Output growth errors					**Treasury bill errors**				
	t	t+1	t+1	t+2	t+3	t	t+1	t+1	t+2	t+3
$Median(x^{i}_{t+k,t-1})$ $[\alpha]$	-0.24 (0.015)	-0.16 (0.390)	-0.57 (0.036)	0.13 (0.709)	0.12 (0.807)	0.00 (1.000)	-0.30 (0.030)	-1.05 (0.199)	-0.42 (0.001)	-0.31 (0.009)
$R^{i}_{t+k,t}[\beta]$	0.25 (0.000)	0.14 (0.102)	-0.07 (0.075)	0.34 (0.000)	0.63 (0.000)	-0.02 (0.251)	0.13 (0.281)	0.59 (0.000)	0.09 (0.432)	0.03 (0.866)
$x^{i}_{t+k,t-1}[\gamma]$	0.15 (0.081)	0.07 (0.386)	0.13 (0.001)	0.09 (0.254)	0.33 (0.000)	0.02 (0.763)	0.37 (0.004)	0.73 (0.000)	0.58 (0.000)	0.54 (0.000)
Additional controls			Y					Y		
Test, non-revision variables = 0			0.000					0.000		
R-squared	0.05	0.01	0.25	0.03	0.13	0.01	0.05	0.14	0.11	0.12
Observations	3846	3847	3364	3826	3693	3251	3236	3011	3201	3065

Michigan Forecasts

	One-year inflation forecast errors (monthly, 12-month change)	
$Median(x^{i}_{t+k,t-1})[\beta]$	-0.20 (0.000)	0.09 (0.062)
$R^{i}_{t+k,t}[\gamma]$	-0.41 (0.000)	-0.39 (0.000)
Additional *t*-period information		Y
R-squared	0.293	0.347
Observations	60324	11388

Appendix A

Data sources
SPF, ESPF and Michigan Survey Data

All of the SPF survey data used in this study come from the Philadelphia Fed's website (http://www.phil.frb.org/research-and-data/real-time-Center/survey-of-professional-forecasters). The documentation for all of the series employed in this paper may be found here: (http://www.phil.frb.org/research-and-data/real-time-center/survey-of-professional-forecasters/spf-documentation.pdf) .

The ESPF data come from the European Central Bank's website http://www.ecb.europa.eu/stats/prices/indic/forecast/html/index.en.html . The documentation for all of the series in the paper may be found here: http://www.ecb.europa.eu/stats/prices/indic/forecast/shared/files/SPF_dataset_description.pdf

The individual responses for the Michigan survey are available upon request from the University of Michigan's Survey Research Center data archive, and may be found here: http://data.sca.isr.umich.edu/sda-public/cgi-bin/hsda?harcsda+sca