Federal Reserve Bank of Cleveland Working Paper Series

The Distributional Predictive Content of Measures of Inflation Expectations

James Mitchell and Saeed Zaman

Working Paper No. 23-31

November 2023

Suggested citation: Mitchell, James and Saeed Zaman. 2023. "The Distributional Predictive Content of Measures of Inflation Expectations." Working Paper No. 23-31. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202331.

Federal Reserve Bank of Cleveland Working Paper Series
ISSN: 2573-7953

Working papers of the Federal Reserve Bank of Cleveland are preliminary materials circulated to stimulate discussion and critical comment on research in progress. They may not have been subject to the formal editorial review accorded official Federal Reserve Bank of Cleveland publications.

See more working papers at: www.clevelandfed.org/research. Subscribe to email alerts to be notified when a new working paper is posted at: https://www.clevelandfed.org/subscriptions.

The Distributional Predictive Content of Measures of Inflation Expectations*

James Mitchell† Saeed Zaman‡

November 28, 2023

Abstract

This paper examines the predictive relationship between the distribution of realized inflation in the US and measures of inflation expectations from households, firms, financial markets, and professional forecasters. To allow for nonlinearities in the predictive relationship we use quantile regression methods. We find that the ability of households to predict future inflation, relative to that of professionals, firms, and the market, increases with inflation. While professional forecasters are more accurate in the middle of the inflation density, households' expectations are more useful in the upper tail. The predictive ability of measures of inflation expectations is greatest when combined. We show that it is helpful to let the combination weights on different agents' expectations of inflation vary by quantile when assessing inflationary pressures probabilistically.

Keywords: inflation expectations measures, inflation, density forecasts, quantile predictive regressions, non-Gaussian models, nonlinearities

JEL Codes: C15, C53, E3, E37

*We thank Ricardo Reis and participants at ISF 2022, MEG 2023, and the Cleveland Fed for helpful comments. The views expressed herein are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Federal Reserve System.

†Research Department, Federal Reserve Bank of Cleveland, P.O. Box 6387, Cleveland, OH 44101-1387; Email address: james.mitchell@clev.frb.org

‡Research Department, Federal Reserve Bank of Cleveland; Email address: saeed.zaman@clev.frb.org

1 Introduction

Economic agents' expectations of future inflation play an important role both in monetary policy deliberations and in academic research studying the dynamics of inflation. This is because of the recognition, informed by a large body of research, that short-run expectations of future inflation influence wage- and price-setting behavior, which in turn affect future spending and realized inflation.[1] Policymakers closely monitor both agents' short-run and long-run expectations of future inflation.[2] Short-run expectations are an important input when assessing the monetary policy stance, as they feed directly into the calculation of short-run real interest rates; long-run inflation expectations are helpful in assessing the degree to which expectations are anchored around the central bank's inflation target.

Forward-looking measures of inflationary expectations are in turn a key component in new Keynesian dynamic stochastic general equilibrium (DSGE) models, the workhorse macroeconomic model. But there is growing appreciation that expectations are not rational and that agents devote more resources to forming their expectations at times of greater uncertainty; for example, see Coibion and Gorodnichenko (2015). This suggests that the utility of measures of expected inflation may be higher precisely at times when uncertainty about the future path of inflation is higher and, accordingly, policymakers are looking especially hard for a signal about inflationary prospects. Practically, for policymakers concerned with forecasting future inflation and understanding the drivers of inflationary expectations, often with an eye on anchoring longer-run expectations, direct estimates of agents' subjective inflationary expectations are therefore an important resource.[3] Increasingly, as

[1] Werning (2022) emphasizes that it is short-run inflation expectations (i.e., the horizon over which prices are sticky) that are most relevant for wage- and price-setting behavior.

[2] As one example of this practice (taken from https://www.clevelandfed.org/research/economists/mester-loretta-j/sp-20220926-inflation-inflation-expectations-and-monetary-policymaking-strategy), which also serves to motivate this paper, in a September 26, 2022 speech President Mester of the Cleveland Fed explains that "One difficulty in moving from theory to practice is that inflation expectations are not directly observable. So we look at a number of measures, which differ by type of agent and time horizon.... A clear signal is not always forthcoming because the inflation expectations of different groups of agents can behave differently from one another, even within groups there can be variation, and the literature has not firmly established whose expectations are most important for inflation dynamics."

[3] There is now a considerable empirical literature that uses direct measures of expectations to forecast inflation. Ang et al. (2007) find that measures of short-run expectations from both households and professionals, as elicited by surveys, are hard to beat using model-based forecasts, although Trehan (2015) finds that predictive performance has deteriorated more recently. Incorporating longer-run estimates of inflationary expectations within time-series models has also been found to improve their accuracy; e.g., see Faust and Wright (2013), Zaman (2013), Chan et al. (2018), and Tallman and Zaman (2020). Piger and Rasche (2008) find that expectations measures "trump" the gap, in terms of explaining actual inflation dynamics.

reviewed in Weber et al. (2022), these data are available at the micro-level and for different types of economic agents. At times, the inflation expectations of these different types of agents can show marked deviations. This raises questions for policymakers about whose inflationary expectations they should track, how much weight they should attach to them, and whether the predictive content of alternative measures of expectations varies with the prevailing level of inflation.

To illustrate this uncertainty, panel (a) of Figure 1 plots one-year-ahead (averaged) expectations of inflation for households, professionals (leading business economists), firms, and as extracted from financial market data.[4] The bottom panel of Figure 1 plots their corresponding long-run expectations of inflation.[5] Figure 1a shows that, since the early 2000s, households' expectations have been both higher and more volatile than the expectations from professional forecasters. Expectations from the market are also much lower, although quite volatile. A similar picture is seen when looking at agents' long-run expectations plotted in panel (b). However, there is less volatility, with the expectations of professionals anchored around 2 percent since 2000.

To further evidence the underlying heterogeneity across agents, as emphasized in work, for example, by Binder (2015), Figures 2a, 2b, and 2c break down households' expectations averaged by age, education, and income groups. These figures show that less educated and lower-income households tend to have higher and more volatile expectations. The plots also reveal that older households tend to have less volatile, but generally higher, expectations than younger households. The higher volatility in inflation expectations of the younger households is consistent with the findings in Malmendier and Nagel (2016), who document that younger people revise their expectations "more strongly" than older people in response to inflation surprises. In a similar vein, Pedemonte et al. (2023) find that since the inflation surge of 2021, older individuals have higher inflation expectations, consistent with their forming their expectations by drawing on their memory of the higher inflation they experienced in the 1970s.

It is these empirical features, seen in Figures 1 and 2, that motivate this paper. We show that a useful way of understanding these heterogeneities is to interpret them through the lens of a (reduced-form) model of the full conditional distribution of inflation. When using these expectations data to

[4]These data are introduced and defined in Section 2.

[5]The *average* across households is the median, in keeping with how these data are commonly reported and used. The *average* across the professionals, firms, and financial markets is the mean, again consistent with common practice.

Figure 1: Inflation Expectations Measures Across Different Agents

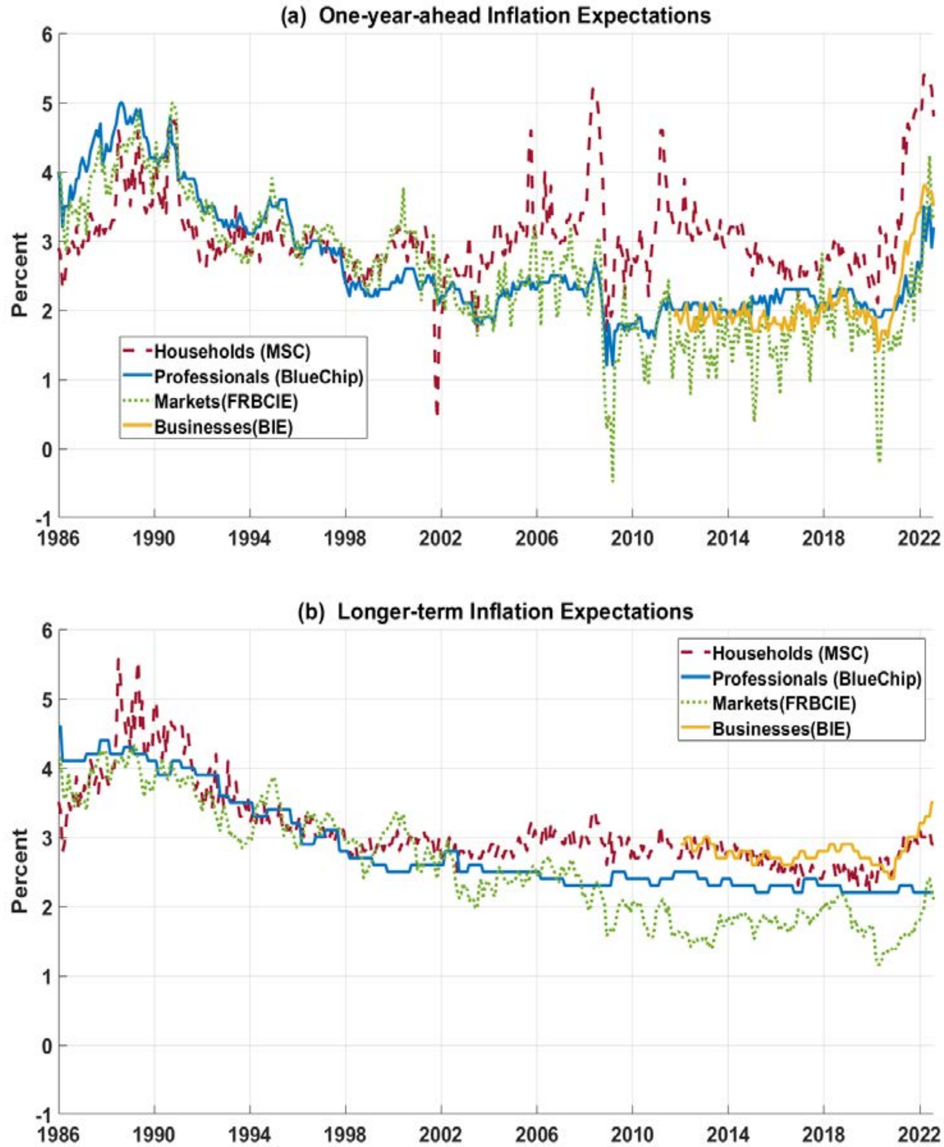

Notes: Panel (a) plots average one-year-ahead inflation expectations of households (from the Michigan Survey of Consumers), professionals (from Blue Chip), markets (from the Cleveland Fed model), and firms/businesses (from the Atlanta Fed business inflation expectations survey: inflation expectations are measured as the expected change to unit costs). Panel (b) plots the corresponding longer-term expectations; expectations of households are defined as the median expected inflation rate during the next five years; for professionals we use the 7- to 10-year-ahead CPI consensus forecast; for markets we use the expected inflation rate over the next 10 years; and for firms/businesses we use the mean expected change to unit costs per year over the next 5 to 10 years.

Figure 2: One-Year-Ahead Inflation Expectations Across Household Demographics

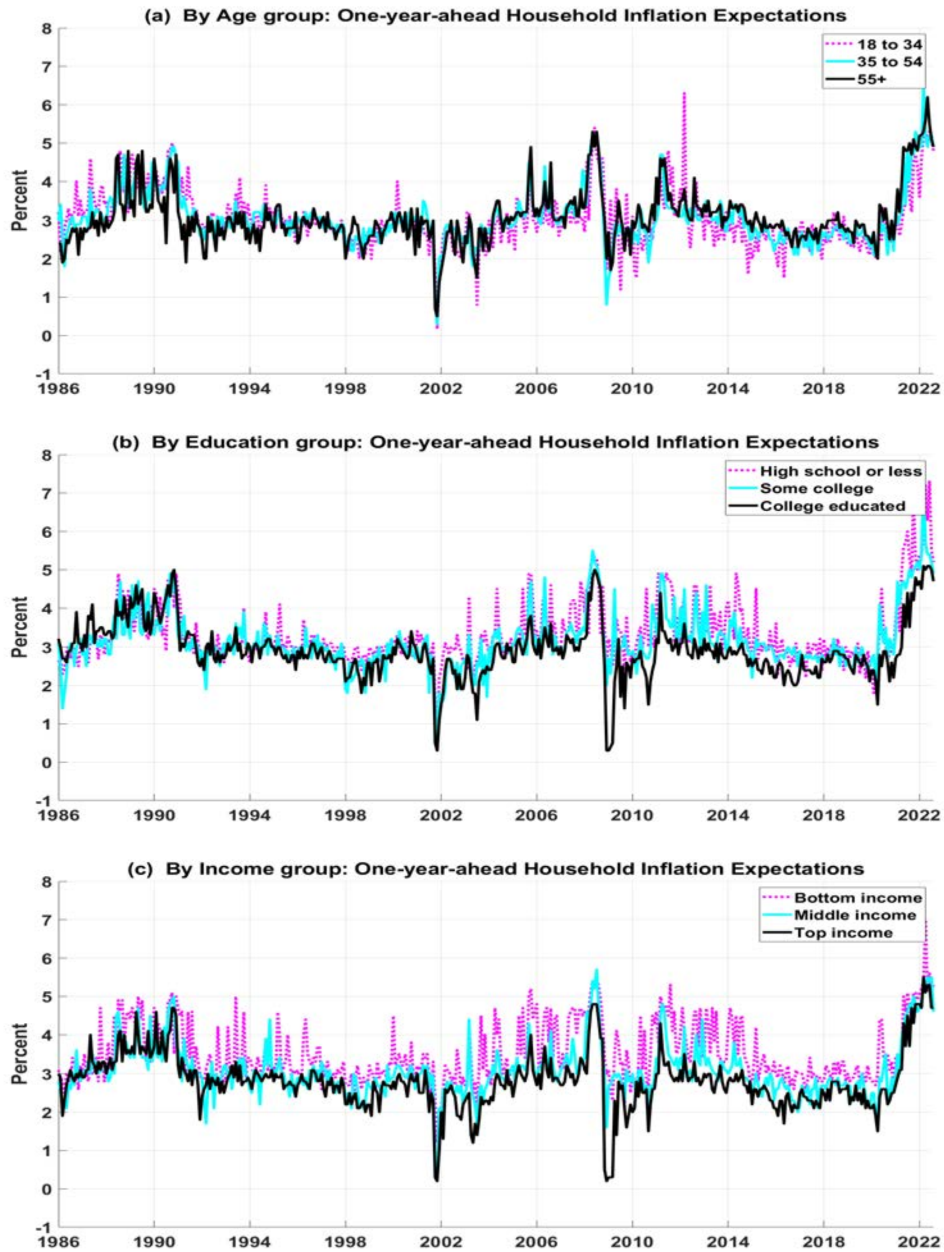

forecast inflation we find important nonlinearities. Specifically, we use predictive quantile regressions (QRs) to show that households' expectations, in particular, have more predictive content for future inflation when inflation is high (relative to its mean) rather than when it is low. This empirical finding is consistent with "rational (in)attention," namely, that households find it beneficial to invest more in the construction of their forecasts when inflation is high.[6] If this nonlinearity is ignored, as is common when forecasting the conditional mean alone, we show that one would mistakenly conclude that households' expectations are less informative than those of professionals. This explains why previous research that focuses on point forecasts (that is, conditional mean estimates), such as Carroll (2003), Ang et al. (2007), Trehan (2015), Meyer et al. (2021), and Verbrugge and Zaman (2021), has found that households forecast less well than professional forecasters. Our results show that while this is true when forecasting the center of the inflation density, it does not hold in the tails.

Although our focus is on short-run expectations, we also investigate the joint predictive content of both short- and long-run inflation expectations in forecasting inflation one year through three years ahead. We might expect long-run expectations to better capture agents' confidence in the central bank's commitment to price stability than provide accurate forecasts of short-run inflation. In contrast, given the putative long-and-variable lags of monetary policy, short-run expectations of inflation should be relatively unaffected by beliefs about the conduct of monetary policy and, thus, should better forecast near-term inflation. In any case, our empirical approach lets the data decide how much weight when forecasting inflation probabilistically to place both on alternative agents' expectations and on short- versus long-run expectations. Importantly, our proposed quantile-based combination approach lets these weights vary by quantile. It provides a way to reconcile the heterogeneous expectations from different agents to form "optimal" (specifically, continuous ranked probability score (CRPS) minimizing) combined density forecasts of inflation. Policymakers are known to track and scrutinize a variety of expectations measures, as evidenced by the recent development of the Fed's Index of Common Inflation Expectations (CIE); see Ahn and Fulton (2021). But rather than, like the CIE, extract from these alternative expectations series a composite measure that, in effect, attaches a high weight to those expectations measures that happen to correlate highly with others, our approach is to combine these different measures to produce CRPS minimizing density forecasts of future inflation itself.

The plan of the remainder of this paper is as follows. Section 2 explains the data on actual

[6] See Maćkowiak et al. (2023) for a review of rational inattention. Nonlinearities in the relationship between expectations and subsequent outcomes have also been emphasized in rational expectations models when there are regime shifts; see Hajdini and Kurmann (2023).

inflation and expected inflation that we use. Section 3 then describes the predictive QR approach used to model the relationship between the density for realized inflation and the various measures of expected inflation. Section 4 reports the in-sample empirical results evidencing, notably, the nonlinear relationship between inflation and households' inflation expectations. Section 5 reports supporting out-of-sample results. Section 6 concludes. Online appendices contain supplementary empirical results as referenced in the main paper.

2 Inflation Expectations and Realizations Data

We make use of data on monthly realized inflation and inflation expectations. The inflation expectations of households are measured by the University of Michigan Surveys of Consumers over the next one year and over the next five years.[7] Inflation expectations of professionals are measured by Blue Chip Economic Indicators. We consider Blue Chip's one-year-ahead and 7- to 10-year-ahead consensus forecasts of consumer price index (CPI) inflation. Inflation expectations of firms over the next year are measured by the Federal Reserve Bank of Atlanta's business inflation expectations survey. But as this survey begins much later (in 2011) than our other expectations measures, impeding time-series analysis, we do not consider it further in the main paper and instead analyze it separately in the online appendix (as referenced below when we discuss results for the other measures). A financial-market-based measure of inflation expectations is estimated using the model developed by the Federal Reserve Bank of Cleveland.[8] We use this model's reported estimates of expected inflation over the next one year, next five years, and next 10 years. Our primary focus is on use of the aggregated expectations data for each type of agent (households, professionals, firms, and markets). However, in the case of household expectations, as anticipated above, we also consider expectations data by age, income, and education level. These data again come from the University

[7]The University of Michigan (Surveys of Consumers) releases preliminary estimates of households' expectations toward the beginning of each month based on the responses of approximately 420 respondents. Toward the end of each month it then releases a "final" estimate based on approximately 600 household respondents. In this paper, we use the final estimate. For more details, see https://data.sca.isr.umich.edu/faq.php.

[8]The Cleveland Fed model (Haubrich et al. (2012)) estimates inflation expectations using data that include nominal yields from US Treasury securities, survey forecasts, and inflation swap rate data. Specifically, the model estimates and factors out an inflation risk premium from financial market data, partly on the basis of the inflation expectations of professional economists (Blue Chip Economic Indicators and the Survey of Professional Forecasters). The removal of this risk premium is very important. Previous work (Ang et al. (2007), Bauer and McCarthy (2015), and Mertens (2016)) finds that the forecast accuracy of "raw" financial-market-based inflation expectations is significantly inferior because of this risk premium.

of Michigan Surveys of Consumers.

We then examine the predictive relationship of each of these inflation expectations series with realized CPI inflation (one year, two years, and three years ahead). We repeat the examination replacing CPI inflation with its three main disaggregates: core CPI, food CPI, and energy CPI inflation. Accordingly, we collect and compute year-over-year inflation rates (that is, 12-month trailing rates) of CPI inflation and its three disaggregates using seasonally unadjusted data. In the main part of this paper we focus on results for CPI inflation, but when relevant, we summarize findings for the three disaggregates.[9]

Our main analysis, because of data limitations, is over the sample period January 1986 through August 2022. It therefore misses the so-called Great Inflation period. However, inflation expectations data for households and financial markets, albeit not for professionals, are available prior to 1986. So in a supplementary set of exercises just for households and financial markets, we consider earlier samples that include the Great Inflation as well as the aforementioned more recent sample that contains data for firms' expectations. These exercises serve as robustness checks to ascertain if our empirical results differ over both shorter and longer sample periods. In general, we find results to be robust.[10]

3 Empirical Approach: Predictive Quantile Regressions

We follow a growing literature in empirical macroeconomics in using quantile regressions (QRs) to model the relationship between the full distribution of the variable of interest (in our case, future inflation) and a set of driving variables (for example, inflation expectations measures) flexibly. Adrian et al. (2019) [henceforth ABG] found quantile regressions useful in revealing nonlinearities in how financial conditions affect GDP growth. López-Salido and Loria (2020) extend ABG's analysis of growth-at-risk to consider "inflation-at-risk." As emphasized by Manzan (2015), the attraction of QRs in our context is that the relative informativeness of expectations for future inflation may well

[9] All data, with the exception of household expectations by demographic group, are downloaded via Haver Analytics (Wolters Kluwer Legal and Regulatory Solutions U.S. (Blue Chip)).

[10] Recently, there has been increased attention paid to collecting households' (e.g., Cleveland Fed's consumers and COVID-19 survey: https://www.clevelandfed.org/indicators-and-data/consumers-and-covid-19) and firms' (e.g., Cleveland Fed's SoFIE: https://www.clevelandfed.org/indicators-and-data/survey-of-firms-inflation-expectations) expectations data at the micro-level. But many of these surveys do not extend far enough back in time to permit historical time-series analysis of the sort undertaken in this paper.

vary by quantile of the conditional density.

Specifically, we consider predictive QRs of the following form that relate the τ-th quantile of (subsequently observed at horizon, h) realized inflation, π_{t+h}, to π_t^e, a $d-$dimensional vector of conditioning variables including our measures of inflationary expectations (both individually and combined across agents), lagged actual inflation, and an intercept:

$$Q_\tau(\pi_{t+h}|\pi_t^e) = \pi_t^{e\prime}\beta_\tau, \tag{1}$$

where $\tau \in [0.05, 0.10, ..., 0.90, 0.95]$ and h is the forecast horizon. The importance of including lagged inflation in the set of conditioning variables is that if expectations are purely adaptive and determined by past inflation, then measures of inflationary expectations should not provide value-added in explaining future inflation. We refer to expectations determined by lagged inflation as "naïve" expectations.[11]

Following Gaglianone and Lima (2012) and Korobilis (2017), we do not model quantiles in the extreme tails, less than 0.05 and greater than 0.95, and instead rely on the fitted density (discussed below) to quantify the extreme tails of the density. Estimating extreme quantiles with small samples is well known to be challenging, with Chernozhukov (2005) suggesting the use of extremal methods. Note that, following ABG, we focus on QR models with time-invariant parameters.

The QR slope, β_τ, is chosen to minimize the weighted absolute sum of errors:

$$\hat{\beta}_\tau = \arg\min_{\beta_\tau} \sum_{t=1}^{T-h} (\tau.\mathbf{1}_{(\pi_{t+h} \geq \pi_t^{e\prime}\beta_\tau)} \left|\pi_{t+h} - \pi_t^{e\prime}\beta_\tau\right| + (1-\tau).\mathbf{1}_{(\pi_{t+h} \leq \pi_t^{e\prime}\beta_\tau)} \left|\pi_{t+h} - \pi_t^{e\prime}\beta_\tau\right|),\ \tau \in (0,1), \tag{2}$$

where $\mathbf{1}(.)$ denotes an indicator function.

We focus on the predictive relationship between measures of inflationary expectations and infla-

[11] We note that due to differences in the timing of the elicitation of the survey responses across the different survey measures and to account for the publication lag of the CPI inflation data, the lag of inflation that enters the quantile regressions varies across the survey measures. To be more specific, CPI data for the previous month are typically released in the second week of the current month. Both the household survey expectations (final estimate) and financial markets expectations derived from the Cleveland Fed model are elicited (or computed) after this CPI release, whereas the expectations of the professionals (from Blue Chip) and of the businesses (Atlanta Fed) are elicited earlier in the month, that is, prior to the CPI release. Therefore, for the latter two cases, the CPI data are lagged two months. For example, in the QRs relating future inflation to household and market expectations elicited in August 2021, the lagged CPI data relate to July 2021, and for professionals and businesses expectations, the lagged CPI data relate to June 2021.

tion.[12] Our modeling methodology lets us determine if and how different measures of inflationary expectations, both individually and in combination, vary by quantile in their informativeness for future inflation. When we include within π_t^e all of the different agents' expectations, estimation of (2) across τ amounts to a combination of the alternative measures of inflationary expectations in a manner that delivers "optimal" density forecasts for inflation. Specifically, with the density forecasts constructed from the set of quantile forecasts, as described below, and the quantile forecasts by construction minimizing the (in-sample) quantile score (the tick loss function seen in equation (2)), the constructed combined density forecast by minimizing the quantile score at each quantile is approximately (as the number of quantiles $\tau \in (0,1) \to \infty$) minimizing the (in-sample) continuous ranked probability score (CRPS). That is, the (combined) density forecast is constructed from a set of quantile regressions that each "optimally" combine the different agents' expectations at a given quantile; cf. Giacomini and Komunjer (2005).[13] The CRPS is a popular measure of density fit that is the integral of the quantile scores (see Gneiting and Raftery (2007) and Gneiting and Ranjan (2011)). The optimal combination weights on the competing measures of inflationary expectations in principle then vary by quantile. This is to reflect the possibility that while some measures of inflationary expectations may receive a high weight in one region of the inflation density, they may receive a lower weight in another region.

We consider two ways of constructing density forecasts from the quantile forecasts:

$$\hat{Q}_\tau(y_{t+h}|x_t) = x_t'\hat{\beta}_\tau. \tag{3}$$

First, we follow ABG and fit a skewed-t density to the quantile forecasts at $\tau = 0.05, 0.25, 0.75, 0.95$.[14] Second, to acknowledge that there is no reason to assume that the predictive density for inflation is skewed-t, we follow Mitchell et al. (2022) and construct the density forecast from the 19 quantile forecasts ($\tau = 0.05, 0.1, ..., 0.95$) nonparametrically. To contrast the ABG densities, we label these densities "NP" (nonparametric).

[12] A related literature considers whether additional variables, notably some measure of the output gap as motivated by the Phillips curve (for example, see Binder (2015) and López-Salido and Loria (2020)) also contribute to our understanding of inflation and inflation risks.

[13] Berrisch and Ziel (2021) formalize the conditions for CRPS minimization using QR. See also Aastveit et al. (2022) for a related approach that assigns density forecast combination weights based on quantile scores.

[14] We gratefully make use of ABG's replication code available at https://www.openicpsr.org/openicpsr/project/113169/version/V1/view.

We compare both of these potentially asymmetric density forecasts, constructed allowing for possible nonlinearities between inflation expectations and subsequent realizations of actual inflation, with a linear Gaussian benchmark model. This model (labeled "OLS/Normal") uses OLS to regress realized inflation on the measures of inflation expectations and centers the point forecast for inflation on the conditional mean forecast. The OLS/Normal density is then constructed around this conditional mean forecast by assuming Gaussianity with the standard deviation of the density set equal to the standard deviation of the residual from this linear predictive regression.[15]

We undertake both in-sample and out-of-sample analyses, aware that in-sample results need not hold out-of-sample. The in-sample analysis in Section 4 is primarily conducted over the maximum available sample period, 1986m1-2022m8. However, as explained in the data section above, we do experiment with different sample periods for robustness. The out-of-sample forecast analysis in Section 5 starts with data up to and including May 2000 being used to forecast inflation one year ahead; this means that the first forecast error is computed by comparing the forecast against the May 2001 inflation realization. The final forecast is made in August 2021 for 2022 (one year ahead), with the two- and three-year-ahead forecasts using data up to August 2020 and August 2019, respectively.

4 The Historical Predictive Relationship Between Inflation and Measures of Expected Inflation

4.1 Empirical Evidence for Nonlinearity

To characterize the conditional predictive distribution of inflation with respect to the various expectations measures and to test for nonlinearities, Figure 3 plots the (in-sample) estimated slopes of the QRs of (realized) inflation on each of the four expectations measures individually. In each panel of Figure 3, the OLS slope is indicated, along with 95 percent confidence bands generated under the null hypothesis that the true relationship between inflation and the expectations measure is linear. Figure 3 shows that linearity is rejected for household expectations alone. The regression slopes increase dramatically for households as the quantile increases, suggesting that households' expectations are more informative for tail inflation outcomes than they are when inflation falls toward the

[15] We abstract from parameter estimation error.

middle of its conditional density. Interestingly, as inflation drifts into the left tails of its distribution, the regression slopes decrease sharply as the quantile decreases. The uncovering of this nonlinear relationship of households' inflation expectations with realized inflation is consistent with evidence that when forming their inflation expectations, households put greater weight on price increases than price decreases (see Cavallo et al. (2017) and D'Acunto et al. (2021)). This suggests that households respond strongly to positive inflation surprises but more weakly to negative surprises.[16]

In contrast, both professional forecasters' expectations and those from financial markets have flatter but downward-sloping QR lines, suggesting a stronger relationship between their expectations and inflation when inflation is low. Naïve expectations, however, do relate more strongly to actual inflation when inflation is high. This is consistent with the findings in Wolters and Tillmann (2015) and Tsong and Lee (2011), who also find upward-sloping QR lines when they estimate quantile autoregressions in realized inflation.

Supplementary analysis reported in the online appendix repeats these QRs for each of the three main disaggregates of inflation (core CPI, energy CPI, and food CPI).[17] It finds that household expectations are especially informative about future energy prices and, to a lesser degree, movements in food prices. This suggests that it is the observed nonlinearity between household expectations and energy prices that drives the nonlinear relationship seen in Figure 3 for aggregate one-year-ahead CPI inflation. This is consistent with a growing body of research documenting the strong association between household inflation expectations and changes in gasoline and food prices (for example, see Campos et al. (2022), Weber et al. (2022), and Verbrugge and Zaman (2021)). Households appear to pay close attention to these highly visible prices when forming their expectations. Breaking households down into different demographic groups reveals similar evidence of nonlinearity, with the degree of nonlinearity stronger for older, less than college educated, and lower-income groups (see online appendix Figures A.8, A.9, and A.10). Turning to the professionals, the link between inflation expectations and one-year-ahead energy inflation is stronger in the left tail. This suggests

[16] As Figures A.1 and A.2 in the online appendix show, this nonlinearity is not simply a feature of the inflation data seen since the COVID pandemic. As the Michigan data do benefit from a longer sample period, we can re-estimate the QR in Figure 3 over sample periods from 1978 through both 2007 (so over the Great Inflation period, but pre-global financial crisis) and through the present day. In both cases, we again observe clear nonlinearities, with the QR coefficients displaying an upward slope similar to that in Figure 3. This robustness of Figure 3 to the sample period also extends to the professionals' and the financial-market-based expectations. Figure A.3 repeats Figure 3 but on a sample ending in 2007. Again we find that only household expectations have a nonlinear relationship with inflation.

[17] See Figures A.5, A.6, and A.7.

that professionals take a greater signal from realizations of low energy-price inflation when forming their expectations about future inflation, perhaps due to a belief that the shocks driving energy prices down will prove more persistent than positive shocks.

4.2 Density Forecast Accuracy

To further assess the marginal and combined predictive power of the different agents' inflationary expectations, Table 1 reports two popular measures of the in-sample density fit of each of the four expectations measures both individually and when combined. The benefits of combination are well-established, including across different measures of inflation expectations (for example, see Ahn and Fulton (2021) and Campbell et al. (2023)). Density fit is measured in two ways. First, we report the CRPS, the aforementioned measure of how good forecasts are in matching observed outcomes across the entire inflation distribution, averaged over the sample. Second, we report the tail-weighted CRPS of Gneiting and Ranjan (2011) to focus on accuracy in the tails of the realized inflation density. Computationally, we construct the densities from the quantile forecasts in the two ways discussed in Section 3 above.

Recalling that lower CRPS estimates indicate greater forecast accuracy, Table 1 shows that there are gains in forecast accuracy, both for CRPS as a whole and for the tail-weighted CRPS, when: (1) combining information across all the expectations measures; (2) constructing predictive densities allowing for nonlinear and non-Gaussian features, given OLS/Normal is consistently beaten; and (3) that the choice between ABG and NP is not so important, although NP always delivers slightly more accurate forecasts. We also observe that, of the different measures of expectations considered, naïve expectations are the least accurate. Conditioning a density forecast of inflation on any of the other expectations measures, or a combination of them, always improves accuracy, suggesting that agents' expectations are not simply backward-looking.[18] Given our finding that both ABG and NP produce similar densities, in the interests of brevity henceforth we confine discussion to the NP density forecasts.

The analysis in Table 1 summarizes forecast accuracy averaged across both quantiles and time.

[18]Results presented in the online appendix confirm that these three results also hold for core CPI inflation.

Figure 3: Estimated Quantile Regression Coefficients for the Different Agents

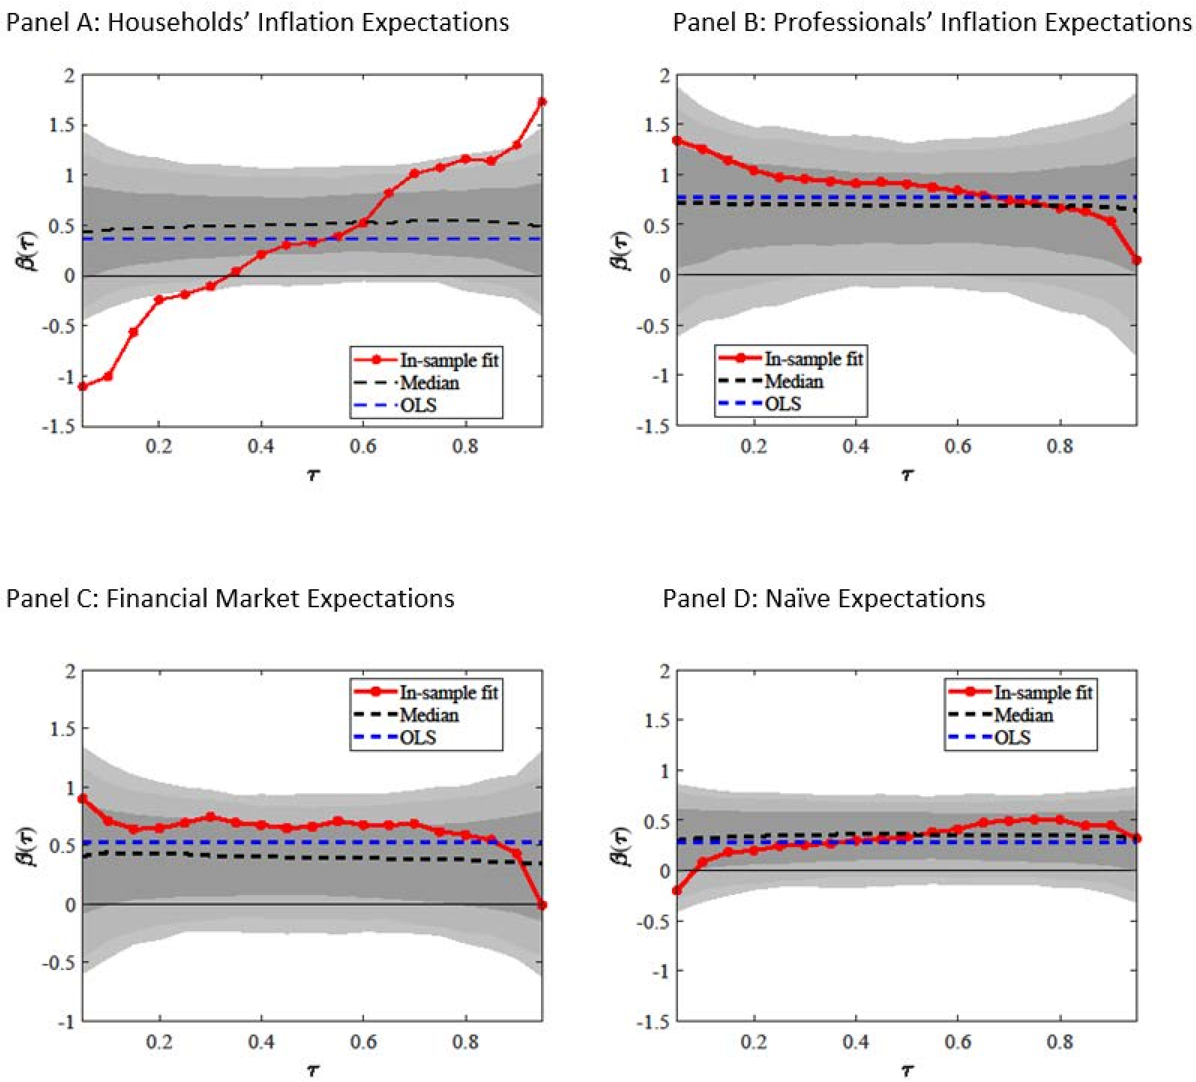

Note: Sample period: 1986m1-2022m8. The figure plots by quantile, τ, the in-sample estimated coefficients corresponding to predictive QRs of 12-month-ahead CPI inflation on current inflation expectations measures (of 12-month-ahead inflation) from each of the different agents. In each case, lagged inflation and an intercept are also included in the QR. Naïve expectations refers to use of the current realized value of inflation as the 12-month-ahead expectation. 95 percent confidence bands (based on 1,000 bootstrapped samples) are constructed following ABG and correspond to the null hypothesis that the true data-generating process is a general bivariate linear model of CPI inflation and inflation expectations (i.e., a VAR model with 12 lags).

To get a sense of how predictive performance may vary by quantile, Figure 4 plots the relative quantile scores for each of the four inflation expectations measures in panels (a) through (d). Each panel

Table 1: Density Forecast Accuracy Across Different Agents

(a) CRPS

Predictors(s): Constant + Lagged inflation + Expectations of:	ABG QR	NP QR	OLS/Normal
Households	0.782	0.779	0.820
Markets	0.762	0.756	0.785
Professionals	0.716	0.713	0.749
Naïve	0.803	0.801	0.821
Combined	0.694	0.688	0.747

(b) Tail-weighted CRPS

Predictors(s): Constant + Lagged inflation + Expectations of:	ABG QR	NP QR	OLS/Normal
Households	0.178	0.172	0.192
Markets	0.177	0.173	0.186
Professionals	0.165	0.164	0.179
Naïve	0.185	0.185	0.192
Combined	0.153	0.152	0.178

Note: Sample period 1986m1-2022m8. Panel (a) reports the CRPS averaged over the whole sample, panel (b) reports the average tail-weighted CRPS. Results given for the density forecasts constructed using the QR method of ABG, the nonparametric (NP) QR method of Mitchell et al. (2022), and assuming a linear Gaussian relationship (labeled "OLS/Normal"). "Combined" involves combining, by quantile, the different agents' expectations and then constructing the density forecast from the combined quantile forecasts.

reports the relative quantile score, that is, the average quantile score of each specific agent's expectations relative to the combined expectations (of households, professionals, markets, and naïve). So, values greater than one (on the y-axis) suggest that, for a given quantile (on the x-axis), combined expectations are more accurate than individual expectations on average over the sample period. The higher the ratio, the greater the gains of combined expectations. Regions where the line is colored green indicate that these gains, at the given quantile, are statistically significant at the 10 percent level, as judged by a Diebold and Mariano (1995) and West-type (1996) test of equality of the quantile scores. As is evident from the four panels, the predictive accuracy of the combination is superior to all four measures of expectations individually, and the gains are generally statistically significant across all quantiles. The predictive gains of the combination are smaller in the right tails compared to the left tails, with the exception of professionals, for whom the gains from the combination are comparable across the two tails.

Figure 4: Density Forecast Accuracy by Quantile: Quantile Scores for the Different Agents

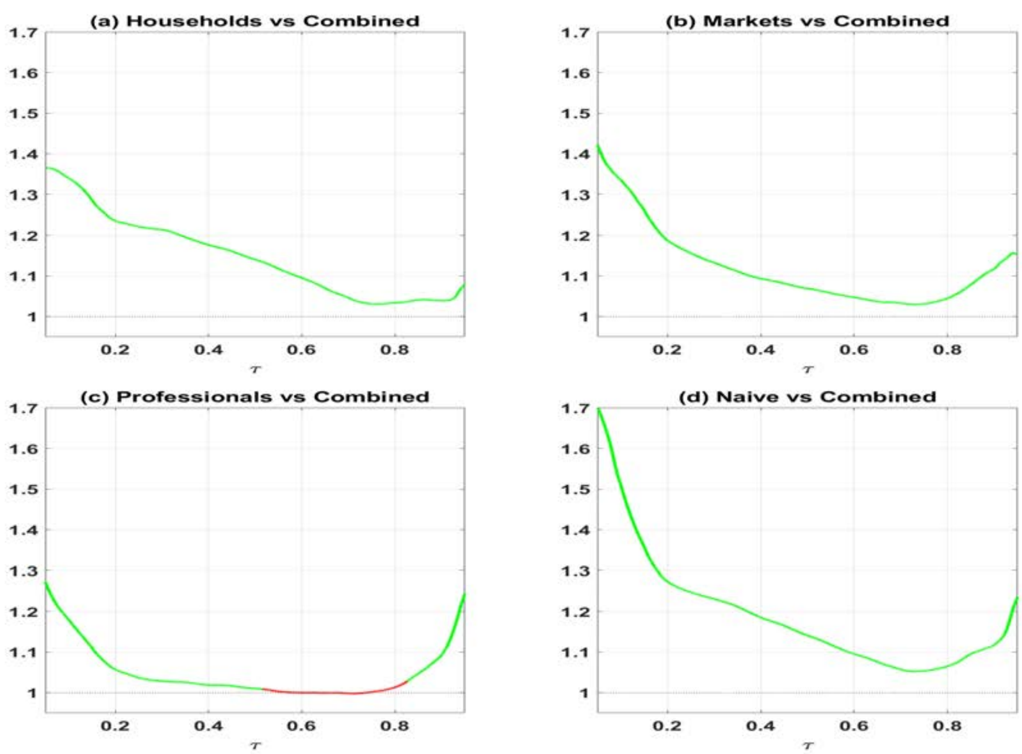

Note: Sample period 1986m1-2022m8. Each panel reports the relative quantile score, that is, the average quantile score of each specific agent's expectations relative to the optimally combined expectations (of households, professionals, markets, and naïve) denoted "Combined." Ratios greater than one (on the y-axis) indicate that, for a given quantile (on the x-axis), combined expectations are more accurate than individual expectations. Regions where the line is colored green indicate that gains (or losses), at the given quantile, are statistically significant at the 10 percent level, as judged by a Diebold and Mariano (1995) and West-type (1996) test of equality of the quantile scores.

Figure 5 then compares, for each of the expectations measures individually, the quantile scores of the QR density (constructed using NP) relative to those of the OLS/Normal density. This time a ratio greater than one indicates that the QR density is more accurate than OLS/Normal at a given quantile. For each of the four measures of expectations, we see gains to relaxing the Gaussianity

Figure 5: Density Forecast Accuracy by Quantile: Relative Quantile Scores for OLS/Normal vs. QR(NP)

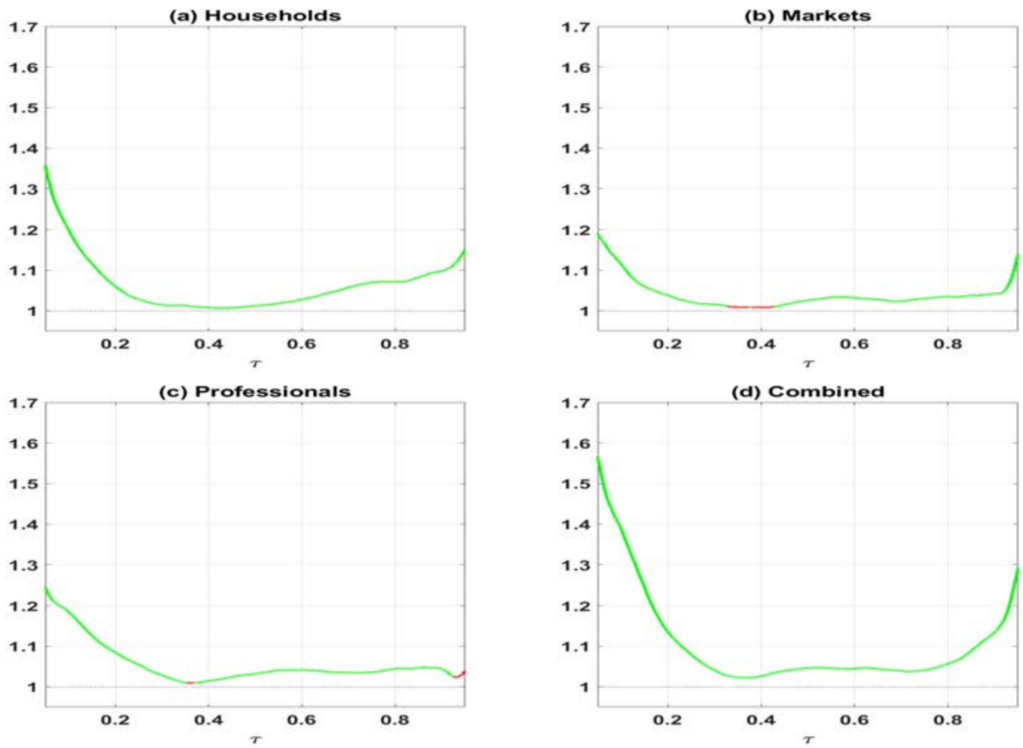

Note: Sample period 1986m1-2022m8. Each panel reports the relative quantile scores of the QR density (constructed using NP) relative to those of the OLS/Normal density. Values greater than one (on the y-axis) indicate that the QR density is more accurate than OLS/Normal at a given quantile. Regions where the line is colored green indicate that the gains (or losses), at the given quantile, are statistically significant at the 10 percent level, as judged by a Diebold and Mariano (1995) and West-type (1996) test of equality of the quantile scores.

assumption. These gains in forecast accuracy are statistically significant at most quantiles.

4.3 Features of the Density Forecasts for Inflation Conditional on Inflationary Expectations

To understand what empirical features these improved densities exhibit, in Figure 6 we plot the 5, 25, 50, 75, and 95 percent quantiles of the one-year-ahead CPI predictive densities conditional on household (panel (a)), market (panel (b)), and professional (panel (c)) measures of inflationary expectations. In each case, the density also conditions on lagged realizations of inflation. Comparison with the quantiles from the analogous density conditional only on lagged inflation (shown in panel (d): denoted "naïve") then reveals how much variation is introduced into the predictive densities by conditioning on agents' expectations.

Figure 6 reveals that each of the three direct expectations measures introduces extra temporal variation into the quantiles. For example, focusing on the 5 and 95 percent quantiles, the standard deviations (over the sample) of these two quantiles for household expectations are 0.7 and 1.0 compared to 0.3 and 0.4, respectively, for naïve. The corresponding standard deviations at the 5 and 95 quantiles for markets are 1.0 and 0.7 and for professionals 1.2 and 0.4, respectively. These estimates suggest that at the 95th quantile household expectations introduce the most variation, followed by market-based expectations. But at the 5th quantile, it is the expectations of professionals that introduce the most variation, followed by the market-based measure.

Figure 7 plots the combined predictive densities, that is, the combination across all the short-term expectations measures (households, markets, professionals, and naïve). It is striking how uncertainty, as evidenced by the width of the intervals, has increased particularly since the oil shocks of the mid-2000s and in the aftermath of the global financial crisis. But this trend increase has been associated with considerably more-temporary variation: the standard deviation (over time) of the 5 and 95 percent quantiles from the combined density increases to 1.2 and 1.0, respectively. As we show next, at higher quantiles the QR model that includes all of the expectations measures assigns a high weight to households' expectations. At lower quantiles, it assigns higher weight to the professionals. This explains the high standard deviation at both quantiles.

Figure 8 draws out additional properties of the predictive densities plotted in Figures 6 and 7. The top panel in Figure 8 plots estimates of "uncertainty," defined as the difference between the 5 percent and 95 percent quantiles. There is strong evidence of time-variation in these uncertainty estimates across all expectations measures. From 1986 through the late 1990s, households' expectations have the highest uncertainty, perhaps stemming from their exposure to the Great Inflation of the 1970s and early 1980s. But from the late 1990s, market-based expectations exhibit more uncertainty. Since COVID-19, however, households, markets, and professionals agree that there is more uncertainty associated with future inflation.

Panel (b) of Figure 8 plots skewness (asymmetry) estimates, extracted from the density forecasts, over time. There is strong evidence of time-varying and, in general, increasing positive skew for each of the agent-based density forecasts. This suggests that all agents believe that upside risks to inflation are not only greater than downside risks but also that they have been growing over time. When the different agents' density forecasts are combined, while we see much more variation over time in the degree of skewness of the (more accurate) combined density forecast, skewness is again generally positive and increasing, especially since the early 2000s. The large discrepancy relative to the skewness estimates from naïve expectations indicates how different backward- and forward-looking estimates of inflation expectations can be in terms of their implications for future inflation. Panel (c) of Figure 8 plots the kurtosis of the predictive densities over the evaluation sample. Again, despite some volatility, we see kurtosis generally increasing over time. This is consistent with the view that when forming their expectations, agents had a growing awareness that inflation could surprise in the tails of the distribution. The kurtosis estimates tend to be greater than 3, confirming the impression from looking at the skewness estimates that it is best not to view agents as forming Gaussian densities for future inflation.

Figure 6: One-Year-Ahead Predictive Densities for CPI Inflation Conditional on Each of the Expectations Measures

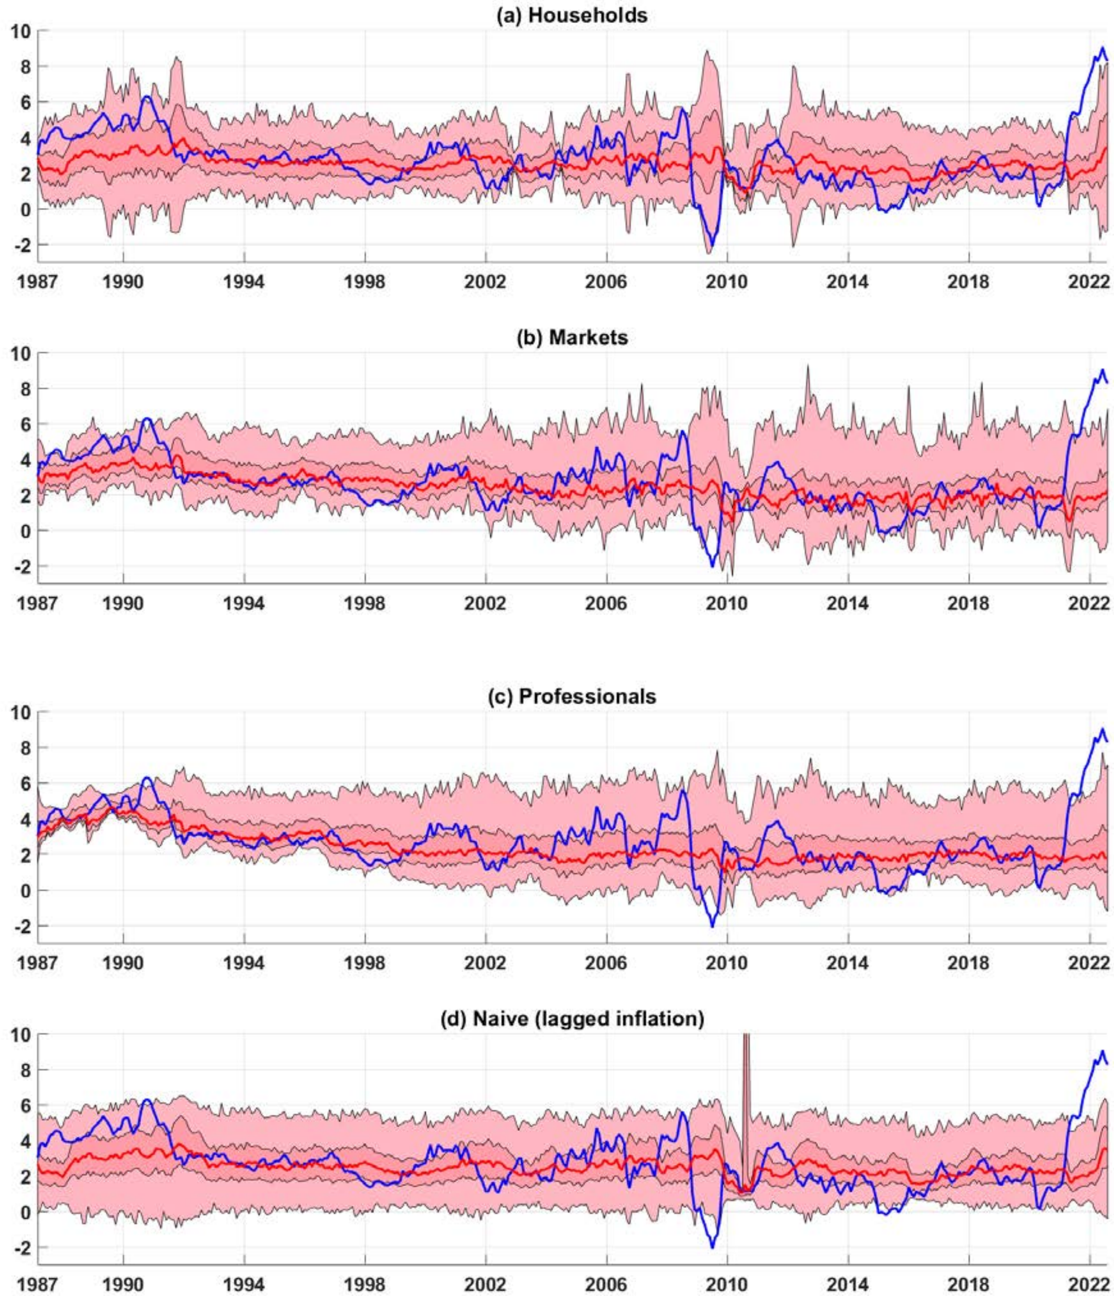

Note: Evaluation sample: 1987m3-2022m8. Dates shown on the x-axis correspond to the forecast evaluation dates, i.e., dates when the CPI realizations are observed. Blue line corresponds to the realization for CPI inflation. Plotted are the 5, 25, 50 (solid red), 75, and 95 percent quantiles of the density forecast for year-ahead inflation, conditional on lagged inflation and either inflation expectations of households, financial markets, or professionals. Naïve density is conditional on lagged inflation only. Density constructed using NP.

Figure 7: One-Year-Ahead Combined Predictive Densities for CPI Inflation

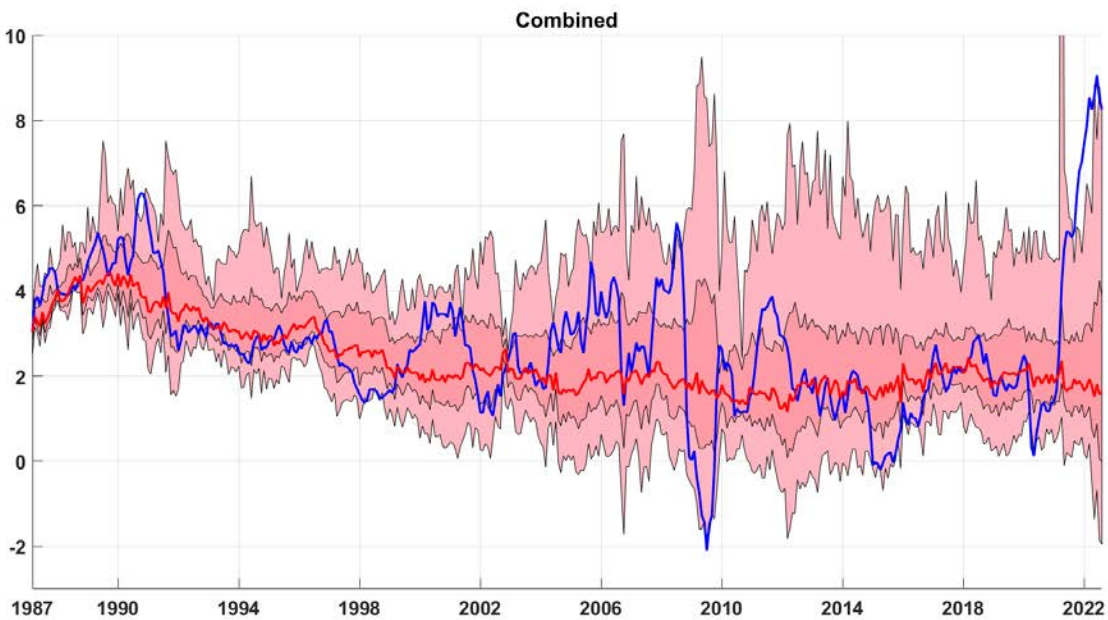

Note: Evaluation sample: 1987m3-2022m8. Dates shown on the x-axis correspond to the forecast evaluation dates, i.e., dates when the CPI realizations are observed. Blue line corresponds to the realization for CPI inflation. Plotted are the 5, 25, 50 (solid red), 75, and 95 percent quantiles of the density forecast for year-ahead inflation, conditional on lagged inflation, expectations of households, financial markets, and professionals. Density constructed using NP.

Figure 8: Properties of One-Year-Ahead Density Forecasts for CPI Inflation

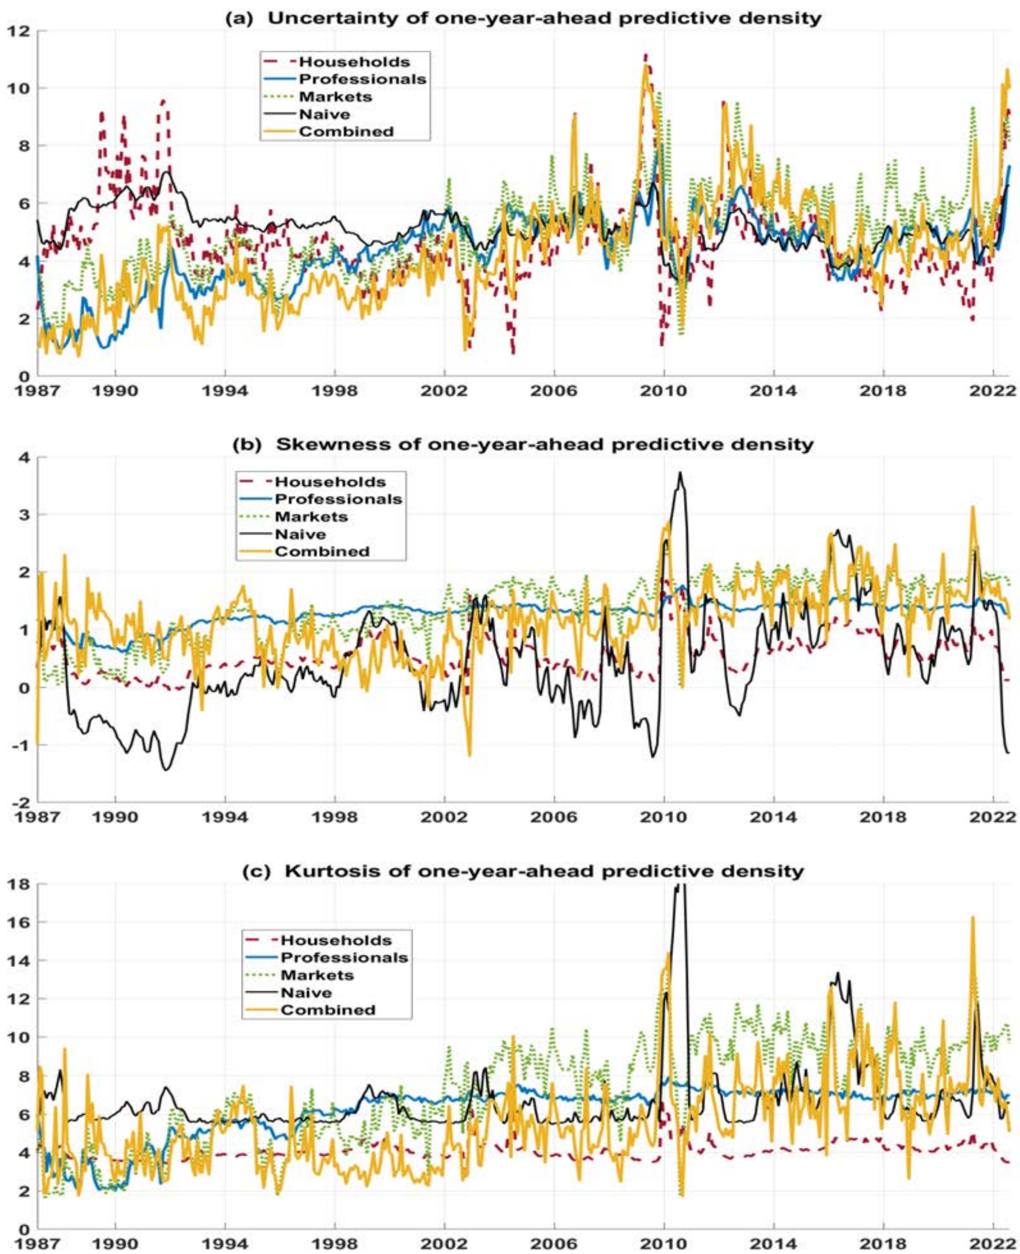

Note: Evaluation sample: 1987m3-2022m8. Dates shown on the x-axis correspond to the forecast evaluation dates, i.e., dates when the CPI realizations are observed. Uncertainty estimates plotted in panel (a) of the figure represent the width between the 95th and 5th quantiles of the one-year ahead density forecasts. Densities are constructed using NP. "Combined" denotes the optimal combination of households, professionals, markets, and naive expectations.

4.4 Combination Weights

To understand the relative informativeness of the different agents' expectations measures when forecasting future inflation probabilistically, panel (a) in Figure 9 plots the in-sample combination weights from both our preferred QR approach and from OLS/Normal. This panel reveals that the optimal weights on the different agents vary by quantile and strongly differ from what would be obtained if linearity was incorrectly assumed (cf. OLS). While the weight on the professional forecasters is highest at lower quantiles, in the upper right tail the weight on households rises sharply.

Anticipating our out-of-sample analysis in Section 5, panel (b) plots the recursively estimated combination weights at the 95 percent quantile. This involves estimating the underlying QRs on expanding windows of data, mimicking real-time use. At the 95 percent quantile, we see that households have received a higher weight than professionals since 2007. The weight on households increases sharply in the post-pandemic era, with the increase in realized and expected inflation.

Panel (c) of Figure 9 turns to the 5 percent quantile, again plotting the temporal evolution of the recursively computed optimal combination weights. Contrasting the right tail, in the left tail we see that professionals consistently receive the highest weight, implying that their expectations are more informative about low inflation outcomes.

These three "facts" suggest that metrics such as the aforementioned Federal Reserve Board's CEI measure (Ahn and Fulton (2021)), which combines measures of inflation expectations to construct a composite measure, are sensibly distilling information across alternative estimates of future inflation. But these facts indicate that weighting different measures of inflationary expectations equally across the density is not optimal empirically, especially when forecasting high inflation outcomes.

4.5 More on the Heterogeneity of Expectations

The following two sub-sections summarize (with the results presented in the online appendix) whether: (a) the predictive content of households' expectations varies across demographic groups; and (b) there is value-added in longer-run expectations of inflation.

Figure 9: Combination Weights

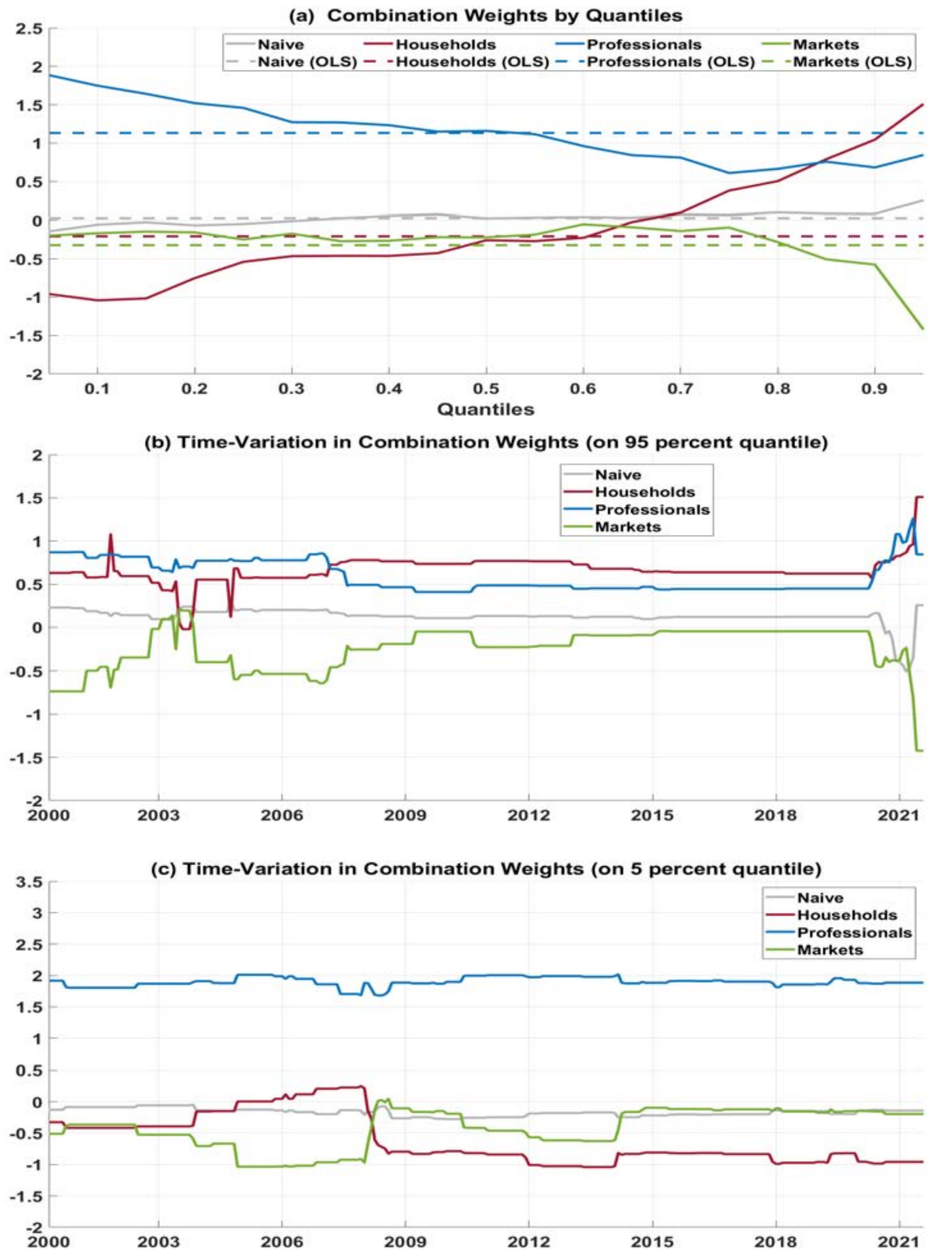

Notes: In panel (a), the combination weights are the quantile coefficients from the predictive QRs estimated using data from 1986 through August 2021. In panels (b) and (c), the time-varying weights are the coefficients from the predictive QRs recursively re-estimated with expanding windows of data from May 2000 through August 2021. This is the same sample period used in our out-of-sample forecasting exercise. The x-axis in panels (b) and (c) refers to the forecast origin dates (May 2000 through August 2021).

4.5.1 Household Expectations by Various Demographic Groups

Binder (2015) finds that expectations from high-income, college-educated, male, and working-age households have the greatest weight in linear Phillips curve models for inflation. Our empirical analysis (see Figure A.11 in the online appendix) confirms that this is the case for middling quantiles (the 50th through 60th quantiles). But at most other quantiles the expectations of specific sub-groups of households are less accurate than when household expectations are aggregated. In addition, we find that at higher quantiles (in the right tail), less educated households have a stronger predictive relationship with future inflation than aggregated household expectations or other sub-groups. As with our earlier findings for professional economists, this is consistent with the view that more educated households form their inflation expectations conditioning on the view that the central bank is committed to returning inflation to low values, and hence their expectations are less informative as inflation deviates from the central bank's (implicit or explicit) target.

4.5.2 Marginal Value of Long-Run Survey Expectations

A growing body of research has shown the utility of long-run survey expectations in improving model-based density forecasts of future inflation (for example, see Chan et al. (2018); Tallman and Zaman (2020); Bańbura et al. (2021)). We compare the predictive performance of QR-based inflation density forecasts with and without long-run survey expectations. We find (see Figure A.12 in the online appendix) that for each type of agent (households, markets, and professionals) long-run survey expectations improve predictive accuracy across the inflation distribution.[19]

5 Out-of-Sample Predictive Power

In this section, we report findings when we extend our analysis out-of-sample (OOS). Specifically, we repeat the exercises performed thus far using the full sample (that is, in-sample) on a recursively expanding window of data from May 2000 to August 2021 to mimic real-time forecasting. Importantly, this OOS exercise involves recursively generating one-year-ahead predictive inflation densities. For brevity, results of the OOS exercise are presented in the online appendix and we summarize the

[19]Similar gains are seen when forecasting inflation both 24 and 36 months ahead (see Figures A.16 and A.17 in the online appendix).

main takeaways here.

The density forecast accuracy comparison reported in Table A.1 echoes the main takeaways from our in-sample analysis. Forecast accuracy gains, both for the CRPS as a whole and for the tail-weighted CRPS, are seen when: (1) we optimally combine information across the different expectations measures; and (2) we construct predictive densities allowing for nonlinear and non-Gaussian features. We find that of the alternative expectations measures considered, naïve expectations are the least accurate. Conditioning a density forecast of inflation on any of the other expectations measures, or a combination of them, always improves accuracy. As is the case in-sample, incorporating information from the agents' long-run survey expectations further improves predictive accuracy (see Figure A.15). However, the magnitude of the gain is small compared to the in-sample results that benefit from a longer sample period.

Finally, we note that when we use a Rossi and Sekhposyan (2019) test to assess the absolute calibration of the OOS predictive densities using the probability integral transforms, we again find (see Figure A.14 in the online appendix) that households' expectations deliver the most accurate forecasts of future inflation in the upper tail of the inflation distribution. In contrast, the accuracy of professional forecasters deteriorates in the upper tail. This serves to reinforce the central result in this paper that when modeling and forecasting inflation, it pays to combine alternative expectations measures, but in a manner that lets the weight on households' expectations increase with inflation.

6 Conclusions

Extending Binder (2015) and Coibion and Gorodnichenko (2015), who implicitly focus on modeling the conditional mean of inflation, we show that the ability of household expectations of inflation to predict future inflation, relative to that of professional forecasters, firms, and market-based measures, increases with inflation. Some households are even better than others. Acknowledging this nonlinearity leads to more accurate density forecasts, especially in the tails. It also delivers conditional density forecasts for inflation that are highly non-Gaussian.

Our empirical results have implications for models of expectations formation in macroeconomics. They support the view that economic agents form expectations subject to information frictions, as in sticky-information or noisy-information models. Our results are consistent with the view that for professional forecasters these frictions are both smaller and constant across the inflation distribution. Instead, for households, these frictions decrease as inflation increases, such that households' expectations of (high) inflation become more informative as the costs of inattention rise.

Econometrically, building on the growing literature in empirical macroeconomics that finds quantile regression is a helpful way of modeling nonlinearities, our paper provides a simple data-based approach of combining different agents' expectations of inflation when assessing inflationary pressures in probabilistic terms. It shows that strategies, like that used in the Federal Reserve Board's Index of Common Inflation Expectations, that implicitly weight different measures of inflationary expectations equally across the density may not be optimal, especially when forecasting high-inflation outcomes. Our results imply that policymakers should be especially attuned to households' expectations of inflation when modeling and forecasting in high-inflation environments.

References

Aastveit, Knut Are, Saskia ter Ellen, and Giulia Mantoan (2022). "Quantile density combination: An application to US GDP forecasts." mimeo, Norges Bank. URL https://www.clevelandfed.org/-/media/project/clevelandfedtenant/clevelandfedsite/events/conference-on-real-time-data/knut-are-aastveit-paper.pdf.

Adrian, Tobias, Nina Boyarchenko, and Domenico Giannone (2019). "Vulnerable growth." *American Economic Review*, 109(4), pp. 1263–89. doi:10.1257/aer.20161923.

Ahn, Hie Joo and Chad Fulton (2021). "Research data series: Index of common inflation expectations." Technical report, FEDS Notes. Washington: Board of Governors of the Federal Reserve System. doi:10.17016/2380-7172.2873.

Ang, Andrew, Geert Bekaert, and Min Wei (2007). "Do macro variables, asset markets, or surveys forecast inflation better?" *Journal of Monetary Economics*, 54(4), pp. 1163–1212. doi:10.1016/j.jmoneco.2006.04.006.

Bauer, Michael D. and Erin McCarthy (2015). "Can we rely on market-based inflation forecasts?" *FRBSF Economic Letter*. URL https://ideas.repec.org/a/fip/fedfel/00070.html.

Bańbura, Marta, Danilo Leiva-León, and Jan-Oliver Menz (2021). "Do inflation expectations improve model-based inflation forecasts?" *European Central Bank Working Paper*. doi:10.2139/ssrn.3940806.

Berrisch, Jonathan and Florian Ziel (2021). "CRPS learning." *Journal of Econometrics*, p. 105221. doi:10.1016/j.jeconom.2021.11.008.

Binder, Carola Conces (2015). "Whose expectations augment the Phillips curve?" *Economics Letters*, 136, pp. 35–38. doi:10.1016/j.econlet.2015.08.013.

Campbell, Colin, Anthony M. Diercks, Steven Sharpe, and Daniel Soques (2023). "The swaps strike back: Evaluating expectations of one-year inflation." Finance and Economics Discussion Series 2023-061, Board of Governors of the Federal Reserve System. doi:10.17016/FEDS.2023.061.

Campos, Chris, Michael McMain, and Mathieu Pedemonte (2022). "Understanding which prices affect inflation expectations." *Federal Reserve Bank of Cleveland, Economic Commentary*, 2022(06). doi:10.26509/frbc-ec-202206.

Carroll, Christopher D. (2003). "Macroeconomic expectations of households and professional forecasters." *The Quarterly Journal of Economics*, 118(1), pp. 269–298. doi:10.1162/00335530360535207.

Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia (2017). "Inflation expectations, learning, and supermarket prices: Evidence from survey experiments." *American Economic Journal: Macroeconomics*, 9(3), pp. 1–35. doi:10.1257/mac.20150147.

Chan, Joshua C.C., Todd E. Clark, and Gary Koop (2018). "A new model of inflation, trend inflation, and long-run inflation expectations." *Journal of Money, Credit and Banking*, 50(1), pp. 5–53. doi:10.1111/jmcb.12452.

Chernozhukov, Victor (2005). "Extremal Quantile Regression." *The Annals of Statistics*, 33(2), pp. 806–839. doi:10.1214/009053604000001165.

Coibion, Olivier and Yuriy Gorodnichenko (2015). "Information rigidity and the expectations formation process: A simple framework and new facts." *American Economic Review*, 105(8), pp. 2644–78. doi:10.1257/aer.20110306.

D'Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber (2021). "Exposure to grocery prices and inflation expectations." *Journal of Political Economy*, 129(5), pp. 1615–1639. doi:10.1086/713192.

Diebold, F. X. and R. S. Mariano (1995). "Comparing predictive accuracy." *Journal of Business & Economic Statistics*, 13, pp. 253–263. doi:10.1080/07350015.1995.10524599.

Faust, Jon and Jonathan H. Wright (2013). "Forecasting Inflation." In Graham Elliott and Allan Timmermann, editors, *Handbook of Economic Forecasting*, volume 2, pp. 2–56. Elsevier. doi:10.1016/B978-0-444-53683-9.00001-3.

Gaglianone, Wagner Piazza and Luiz Renato Lima (2012). "Constructing Density Forecasts from Quantile Regressions." *Journal of Money, Credit and Banking*, 44(8), pp. 1589–1607. doi:10.1111/j.1538-4616.2012.00545.x.

Giacomini, Raffaella and Ivana Komunjer (2005). "Evaluation and combination of conditional quantile forecasts." *Journal of Business & Economic Statistics*, 23(4), pp. 416–431. doi:10.1198/073500105000000018.

Gneiting, Tilmann and Adrian E Raftery (2007). "Strictly proper scoring rules, prediction, and estimation." *Journal of the American Statistical Association*, 102(477), pp. 359–378. doi:10.1198/016214506000001437.

Gneiting, Tilmann and Roopesh Ranjan (2011). "Comparing density forecasts using threshold- and quantile-weighted scoring rules." *Journal of Business & Economic Statistics*, 29(3), pp. 411–422. doi:10.1198/jbes.2010.08110.

Hajdini, Ina and Andre Kurmann (2023). "Predictable forecast errors in full-information rational expectations models with regime shifts." Working papers, SSRN. doi:10.2139/ssrn.4124427.

Haubrich, Joseph, George Pennacchi, and Peter Ritchken (2012). "Inflation expectations, real rates, and risk premia: Evidence from inflation swaps." *Review of Financial Studies*, 25(5), pp. 1588–1629. doi:10.1093/rfs/hhs003.

Korobilis, Dimitris (2017). "Quantile Regression Forecasts of Inflation Under Model Uncertainty." *International Journal of Forecasting*, 33, pp. 11–20. doi:10.1016/j.ijforecast.2016.07.005.

López-Salido, J. David and Francesca Loria (2020). "Inflation at risk." Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System. doi:10.17016/FEDS.2020.013.

Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt (2023). "Rational inattention: A review." *Journal of Economic Literature*, 61(1), pp. 226–73. doi:10.1257/jel.20211524.

Malmendier, Ulrike and Stefan Nagel (2016). "Learning from inflation experiences." *The Quarterly Journal of Economics*, 131(1), pp. 53–87. doi:10.1093/qje/qjv037.

Manzan, Sebastiano (2015). “Forecasting the distribution of economic variables in a data-rich environment.” *Journal of Business & Economic Statistics*, 33(1), pp. 144–164. doi:10.1080/07350015.2014.937436.

Mertens, Elmar (2016). “Measuring the Level and Uncertainty of Trend Inflation.” *The Review of Economics and Statistics*, 98(5), pp. 950–967. doi:10.1162/REST_a_00549.

Meyer, Brent, Nicholas B. Parker, and Xuguang Sheng (2021). “Unit cost expectations and uncertainty: Firms’ perspectives on inflation.” Federal Reserve Bank of Atlanta Working Paper 2021-12a. doi:10.29338/wp2021-12a.

Mitchell, James, Aubrey Poon, and Dan Zhu (2022). “Constructing density forecasts from quantile regressions: Multimodality in macro-financial dynamics.” Working Paper 22-12R, Federal Reserve Bank of Cleveland. doi:10.26509/frbc-wp-202212r.

Pedemonte, Mathieu, Hiroshi Toma, and Estaban Verdugo (2023). “Aggregate Implications of Heterogeneous Inflation Expectations: The Role of Individual Experience.” Working Paper 23-04, Federal Reserve Bank of Cleveland. doi:10.26509/frbc-wp-202304.

Piger, Jeremy M. and Robert H. Rasche (2008). “Inflation: Do expectations trump the gap?” *International Journal of Central Banking*, 4(4), pp. 85–116. URL https://www.ijcb.org/journal/ijcb08q4a3.htm.

Rossi, Barbara and Tatevik Sekhposyan (2019). “Alternative tests for correct specification of conditional predictive densities.” *Journal of Econometrics*, 208(2), pp. 638–657. doi:10.1016/j.jeconom.2018.07.008.

Tallman, Ellis and Saeed Zaman (2020). “Combining survey long-run forecasts and nowcasts with BVAR forecasts using relative entropy.” *International Journal of Forecasting*, 36(2), pp. 373–398. doi:10.1016/j.ijforecast.2019.04.024.

Trehan, Bharat (2015). “Survey measures of expected inflation and the inflation process.” *Journal of Money, Credit and Banking*, 47(1), pp. 207–222. doi:10.1111/jmcb.12174.

Tsong, Ching-Chuan and Cheng-Feng Lee (2011). "Asymmetric inflation dynamics: Evidence from quantile regression analysis." *Journal of Macroeconomics*, 33(4), pp. 668–680. doi:10.1016/j.jmacro.2011.08.003.

Verbrugge, Randal and Saeed Zaman (2021). "Whose inflation expectations best predict inflation?" *Federal Reserve Bank of Cleveland, Economic Commentary*, 2021(19). doi:10.26509/frbc-ec-202119.

Weber, Michael, Francesco D'Acunto, Yuriy Gorodnichenko, and Olivier Coibion (2022). "The subjective inflation expectations of households and firms: Measurement, determinants, and implications." *Journal of Economic Perspectives*, 36(3), pp. 157–84. doi:10.1257/jep.36.3.157.

Werning, Iván (2022). "Expectations and the rate of inflation." Working Paper 30260, National Bureau of Economic Research. doi:10.3386/w30260.

West, Kenneth D. (1996). "Asymptotic inference about predictive ability." *Econometrica*, 64(5), pp. 1067–84. doi:10.2307/2171956.

Wolters, Maik H. and Peter Tillmann (2015). "The changing dynamics of US inflation persistence: a quantile regression approach." *Studies in Nonlinear Dynamics & Econometrics*, 19(2), pp. 161–182. doi:10.1515/snde-2013-0080.

Zaman, Saeed (2013). "Improving inflation forecasts in the medium to long term." *Federal Reserve Bank of Cleveland, Economic Commentary*, (2013-16). doi:10.26509/frbc-ec-201316.

A Online Appendix for "The Distributional Predictive Content of Measures of Inflation Expectations" by James Mitchell and Saeed Zaman

A.1 Sample Robustness

A.1.1 The Relationship Between Household Expectations and Realized Inflation: 1978-2022 and 1978-2007

Figure A.1: Estimated Quantile Regression Coefficients: 1978-2022

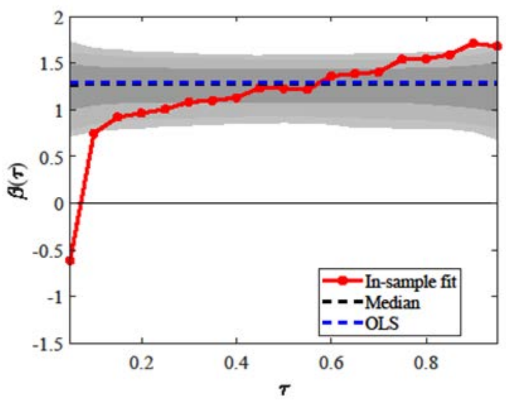

Note: Sample period 1978m1-2022m8. The figure plots the in-sample estimated coefficients corresponding to quantile regressions of 12-month-ahead CPI inflation on households' current inflation expectations measures (of 12-month-ahead inflation). 95 percent confidence bands for linearity are constructed following ABG.

Figure A.2: Estimated Quantile Regression Coefficients: 1978-2007

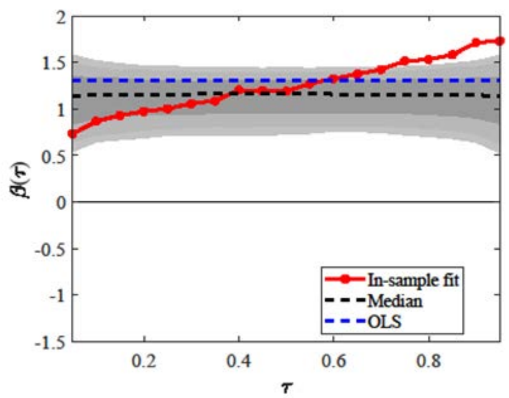

Note: Sample period 1978m1-2007m12. The figure plots the in-sample estimated coefficients corresponding to quantile regressions of 12-month-ahead CPI inflation on households' current inflation expectations measures (of 12-month-ahead inflation). 95 percent confidence bands for linearity are constructed following ABG.

A.1.2 Results over the Sample 1986-2007

Figure A.3: Estimated Quantile Regression Coefficients

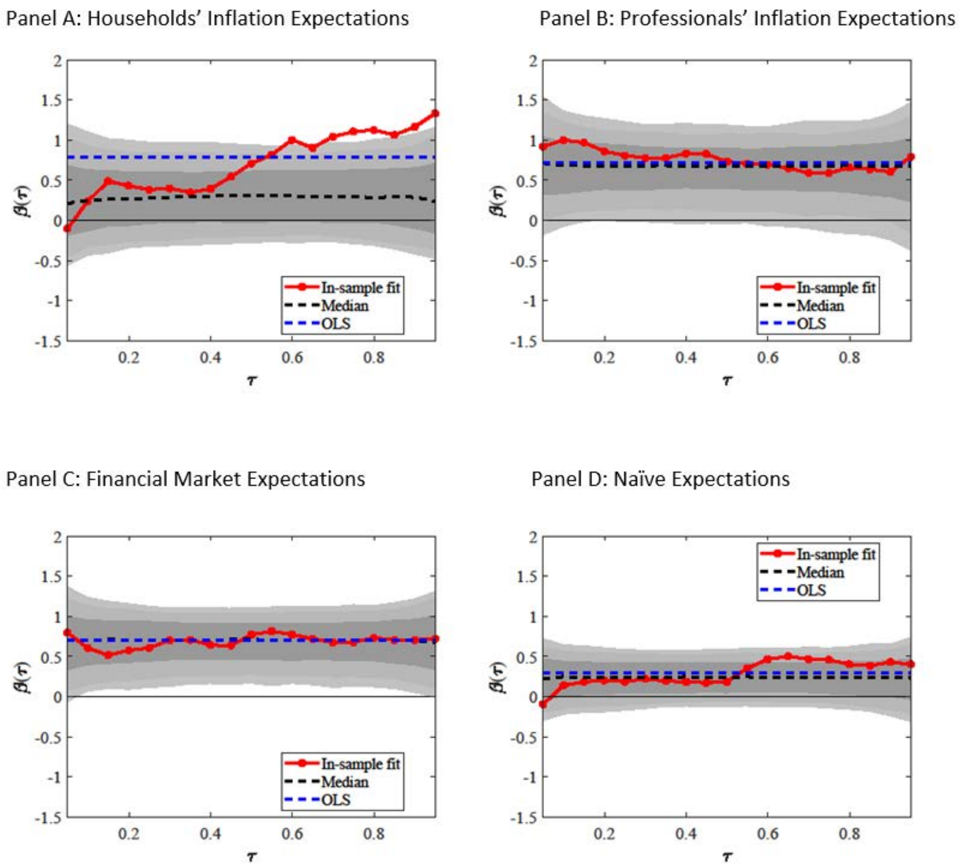

Note: Sample period 1986m1-2007m12. The figure plots the in-sample estimated coefficients corresponding to quantile regressions of 12-month-ahead CPI inflation on current inflation expectations measures (of 12-month-ahead inflation). Naïve expectations refers to using the current value of inflation as the 12-month-ahead expectation. 95 percent confidence bands (based on 1,000 bootstrapped samples) are constructed following ABG and correspond to the null hypothesis that the true data-generating process is a general bivariate linear model of CPI inflation and inflation expectations (i.e., a VAR model with 12 lags).

A.1.3 Results over the Sample 2011 - 2022 with Measures of Business Expectations Included)

Figure A.4: Estimated Quantile Regression Coefficients: Comparison including the Atlanta Fed's Business Expectations

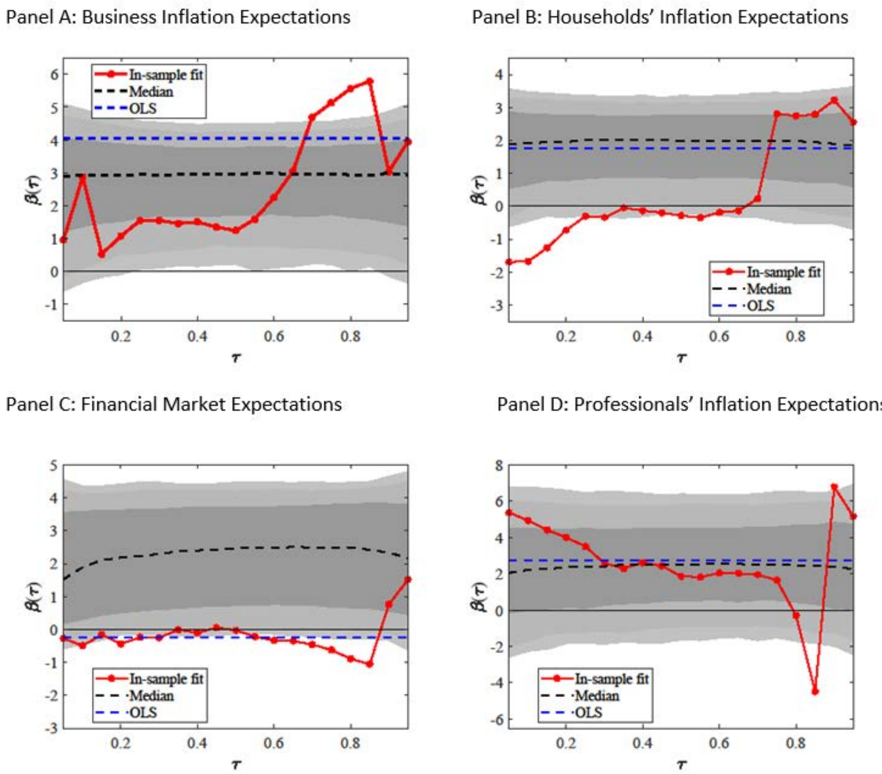

Note: Sample period: 2011m12-2022m8. The figure plots the in-sample estimated coefficients corresponding to quantile regressions of 12-month-ahead CPI inflation on current inflation expectations measures (of 12-month ahead inflation). Naïve expectations simply refers to using the current value of inflation as the 12-month-ahead expectation. 95 percent confidence bands (based on 1,000 bootstrapped samples) are constructed following ABG and correspond to the null hypothesis that the true data-generating process is a general bivariate linear model of CPI inflation and inflation expectations (i.e., a VAR model with 12 lags).

A.2 Results for the Disaggregates of CPI Inflation

A.2.1 Core CPI Inflation

Figure A.5: Estimated Quantile Regression Coefficients for Core Inflation

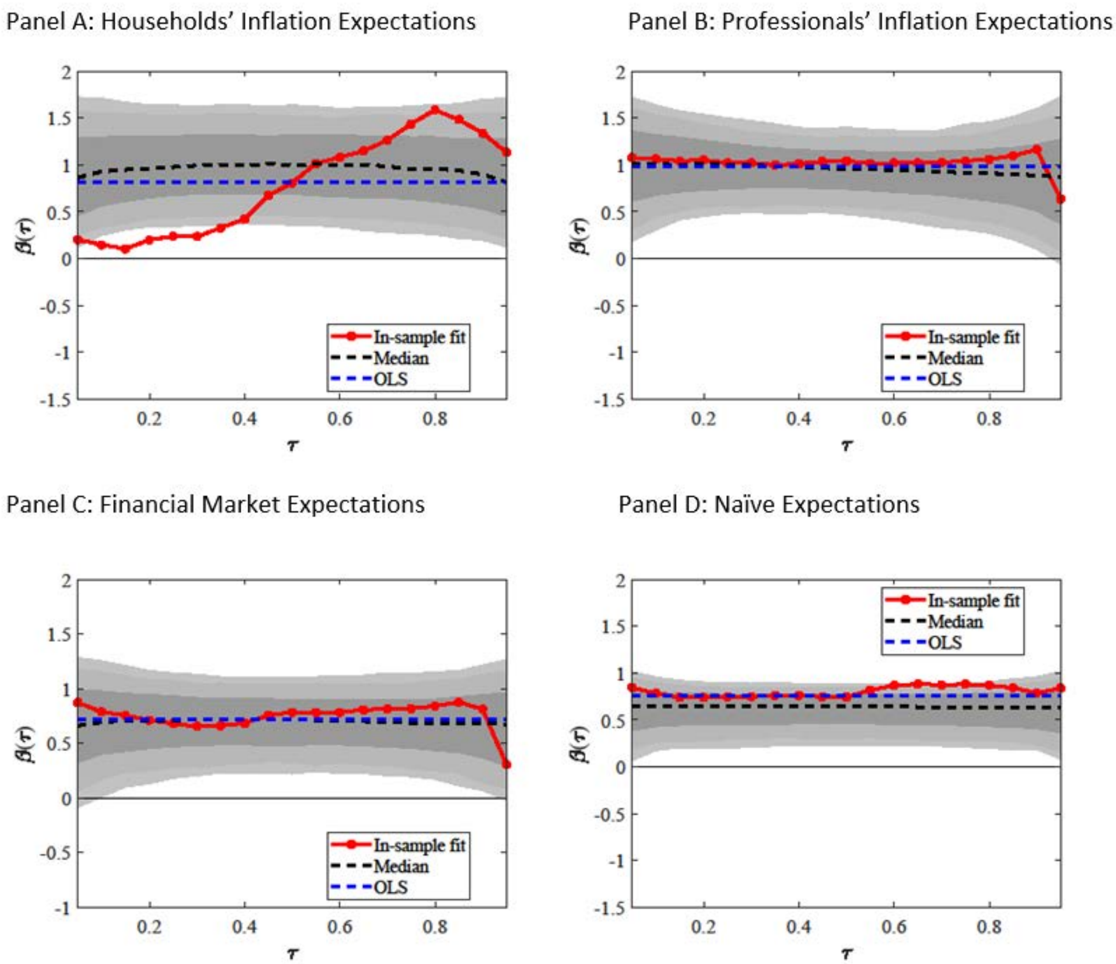

Note: Sample period 1986m1-2022m8. The figure plots the in-sample estimated coefficients corresponding to quantile regressions of 12-month-ahead core CPI inflation on current inflation expectations measures (of 12-month ahead inflation). Naïve expectations refers to use of the current value of inflation as the 12-month-ahead expectation. 95 percent confidence bands (based on 1,000 bootstrapped samples) are constructed following ABG and correspond to the null hypothesis that the true data-generating process is a general bivariate linear model of CPI core inflation and inflation expectations (i.e., a VAR model with 12 lags).

A.2.2 Food CPI Inflation

Figure A.6: Estimates of Quantile Regression Coefficients for Food Inflation

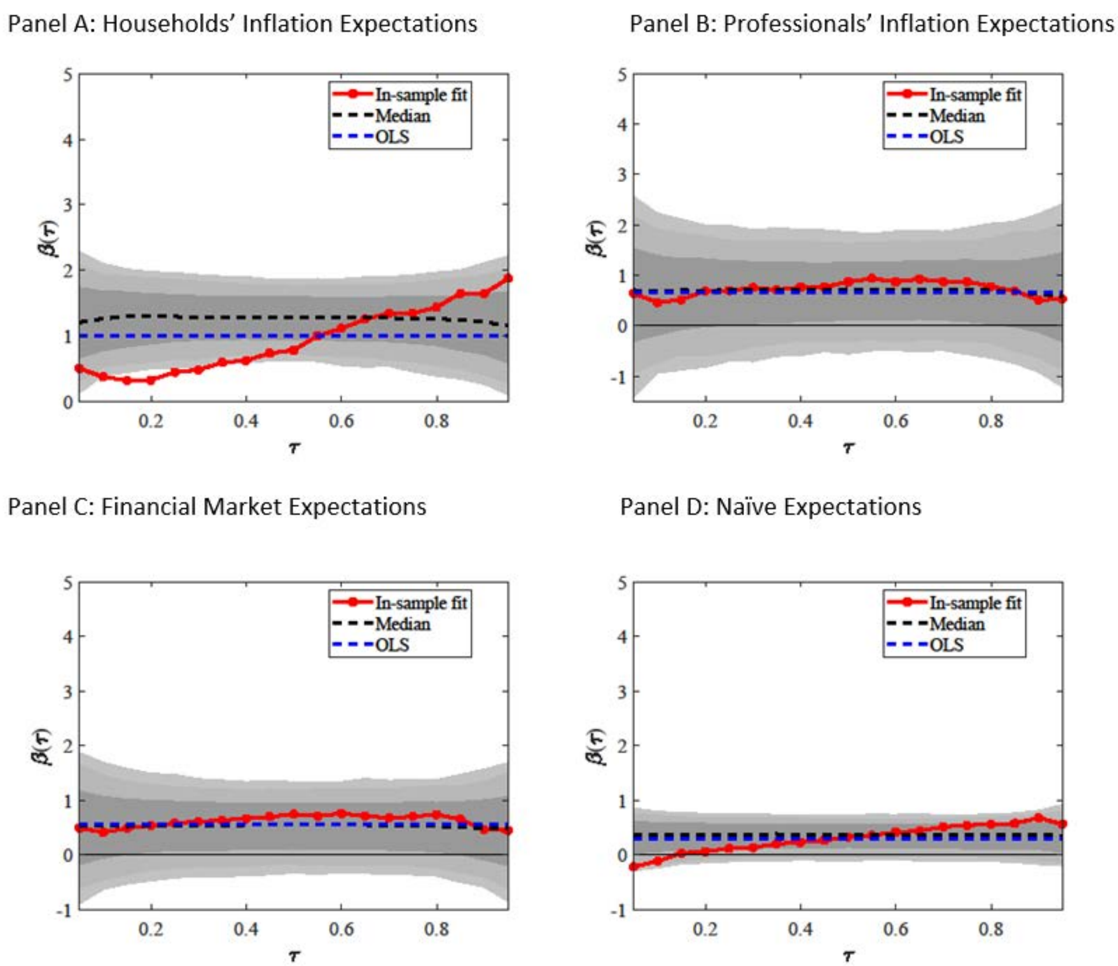

Note: Sample period: 1986m1-2022m8. The figure plots the in-sample estimated coefficients corresponding to quantile regressions of 12-month ahead food price inflation on current inflation expectations measures (of 12-month ahead inflation). Naïve expectations refers to use of the current value of inflation as the 12-month-ahead expectation. 95 percent confidence bands (based on 1,000 bootstrapped samples) are constructed following ABG and correspond to the null hypothesis that the true data-generating process is a general bivariate linear model of food price inflation and inflation expectations (i.e., a VAR model with 12 lags).

A.2.3 Energy CPI Inflation

Figure A.7: Estimated Quantile Regression Coefficients for Energy Inflation

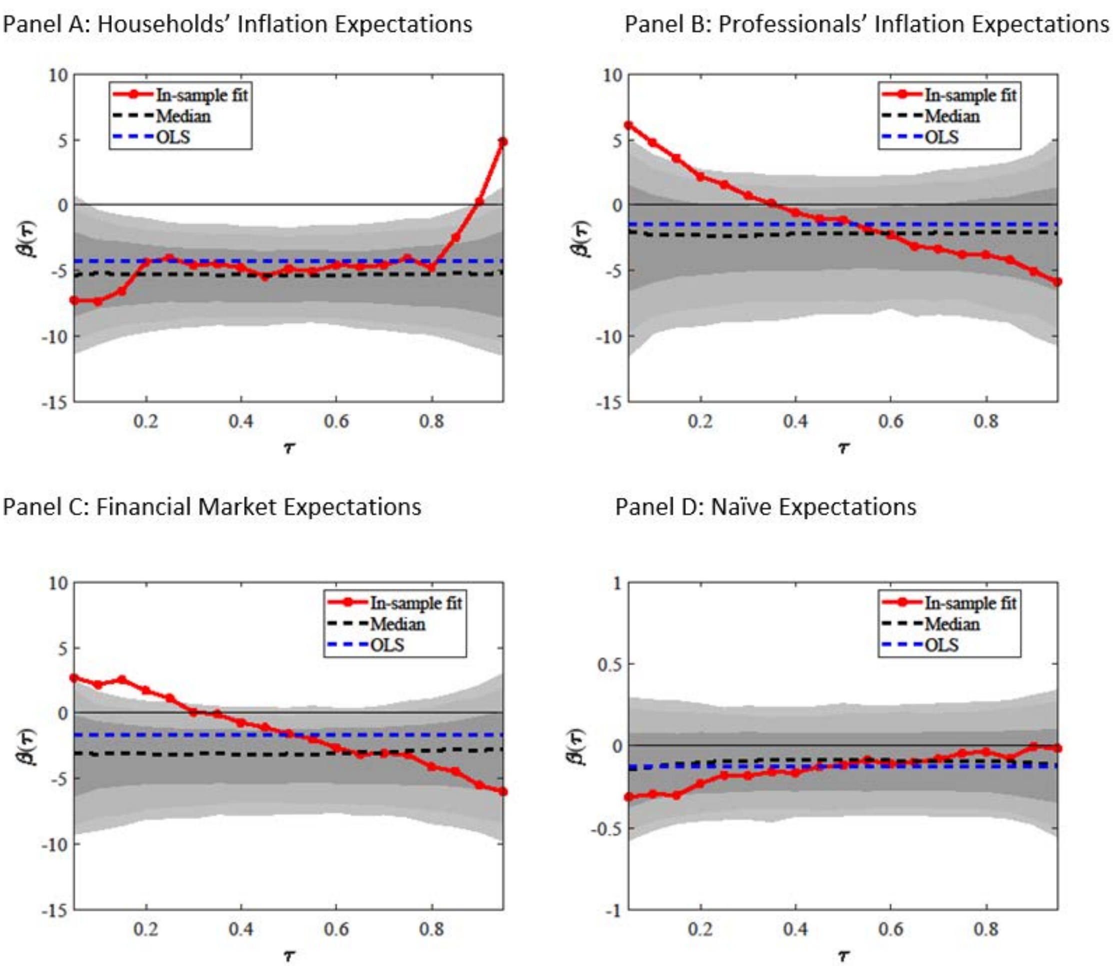

Note: Sample period: 1986m1-2022m8. The figure plots the in-sample estimated coefficients corresponding to quantile regressions of 12-month ahead energy CPI inflation on current inflation expectations measures (of 12-month ahead inflation). Naïve expectations refers to use of the current value of energy inflation as the 12-month-ahead expectation. 95 percent confidence bands (based on 1,000 bootstrapped samples) are constructed following ABG and correspond to the null hypothesis that the true data-generating process is a general bivariate linear model of energy inflation and inflation expectations (i.e., a VAR model with 12 lags).

A.3 Results for Household Demographic Groups

A.3.1 Breakdown by Age

Figure A.8: Estimated Quantile Regression Coefficients by Household Age

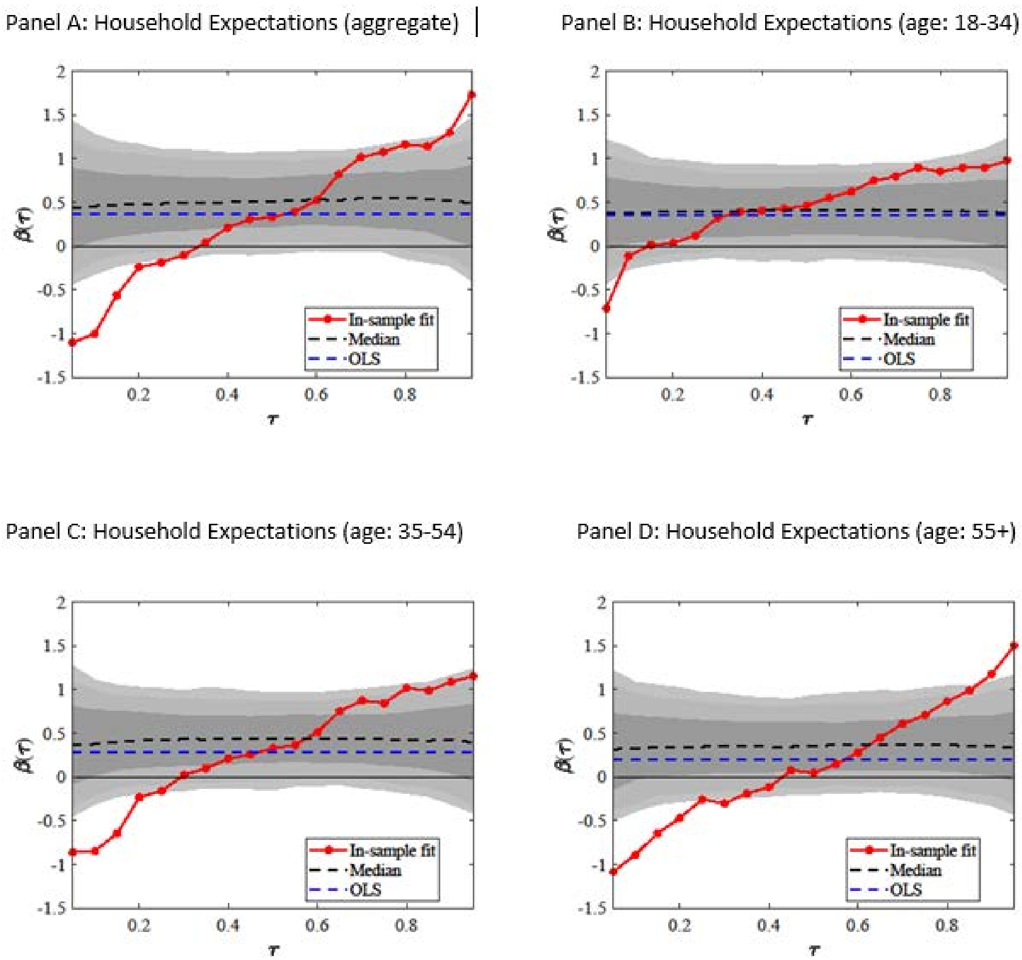

Note: Sample period: 1986m1-2022m8. The figure plots the in-sample estimated coefficients corresponding to quantile regressions of 12-month-ahead CPI inflation on current inflation expectations measures (of 12-month-ahead inflation). 95 percent confidence bands (based on 1,000 bootstrapped samples) are constructed following ABG and correspond to the null hypothesis that the true data-generating process is a general bivariate linear model of CPI inflation and inflation expectations (i.e., a VAR model with 12 lags).

A.3.2 Breakdown by Education

Figure A.9: Estimated Quantile Regression Coefficients by Household Education

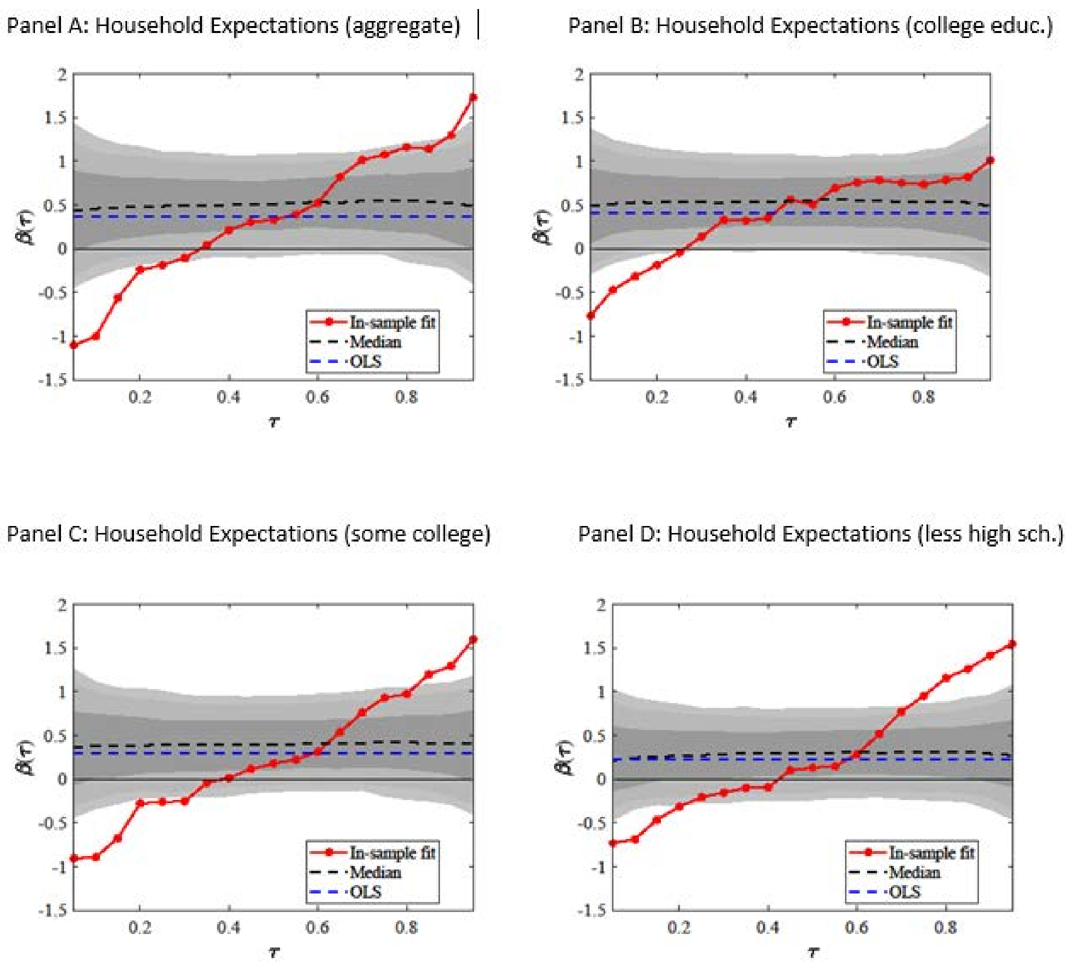

Note: Sample period 1986m1-2022m8. The figure plots the in-sample estimated coefficients corresponding to quantile regressions of 12-month-ahead CPI inflation on current inflation expectations measures (of 12-month-ahead inflation). 95 percent confidence bands (based on 1,000 bootstrapped samples) are constructed following ABG and correspond to the null hypothesis that the true data-generating process is a general bivariate linear model of CPI inflation and inflation expectations (i.e., a VAR model with 12 lags).

A.3.3 Breakdown by Income

Figure A.10: Estimated Quantile Regression Coefficients by Household Income

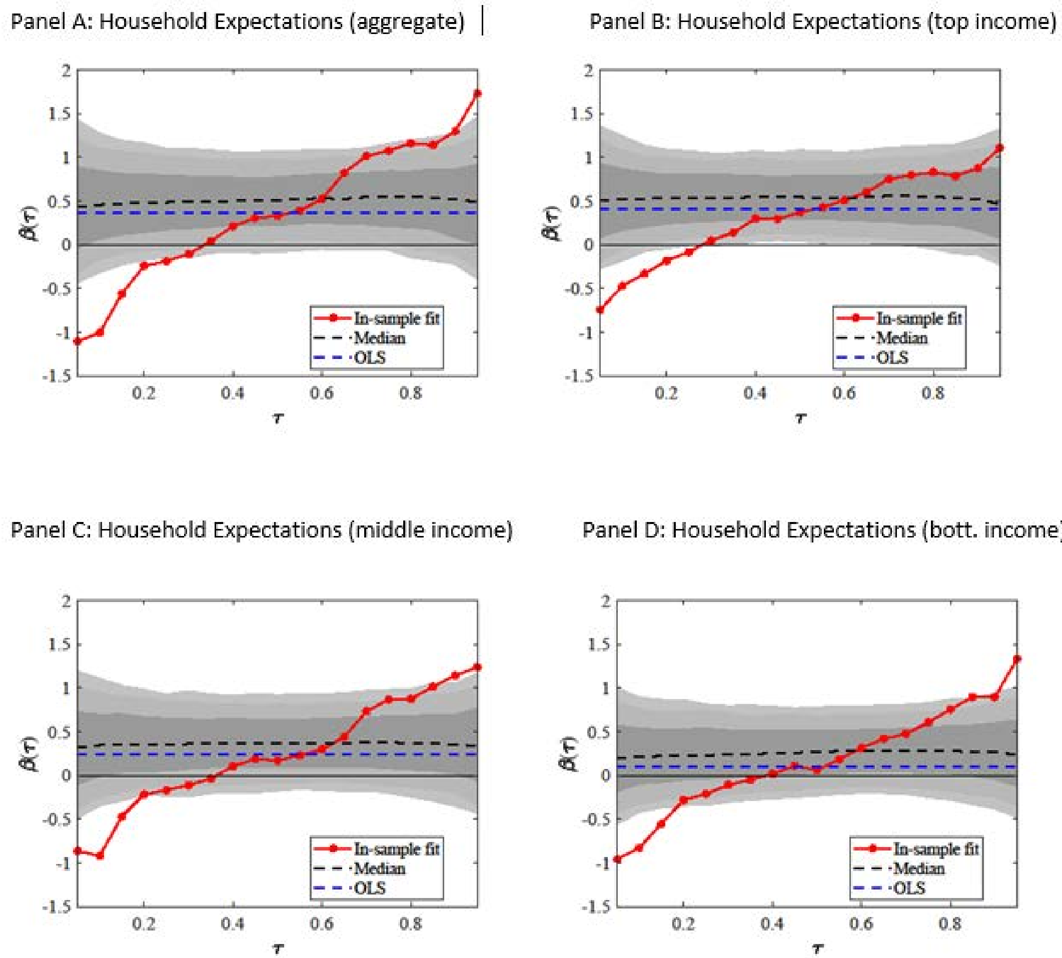

Note: Sample period 1986m1-2022m8. The figure plots the in-sample estimated coefficients corresponding to quantile regressions of 12-month-ahead CPI inflation on current inflation expectations measures (of 12-month-ahead inflation). 95 percent confidence bands (based on 1,000 bootstrapped samples) are constructed following ABG and correspond to the null hypothesis that the true data-generating process is a general bivariate linear model of CPI inflation and inflation expectations (i.e., a VAR model with 12 lags).

A.4 Household Inflation Expectations Across Various Demographic Groups

Figure A.11 compares the accuracy of the density forecasts from the QR model when household expectations are aggregated with those from specific demographic groups. A quantile score ratio greater than one indicates that the QR density constructed using aggregated expectations is more accurate than the density constructed from the expectations of a specific demographic group. Results show that the predictive density from the aggregate expectations measure is generally more accurate than the density from specific demographic groups. Exceptions are: (1) expectations of the less educated have a stronger predictive content for future inflation in the extreme right tails (see panels a and b); and (2) expectations of the college educated, top income, and younger working age have stronger predictive content for future inflation at the middle quantiles (see panels c, f, and g). This latter result is in line with the findings of Binder (2015), who only models the conditional mean of realized inflation.

Figure A.11: Density Accuracy by Demographic Groups: Relative Quantile Score

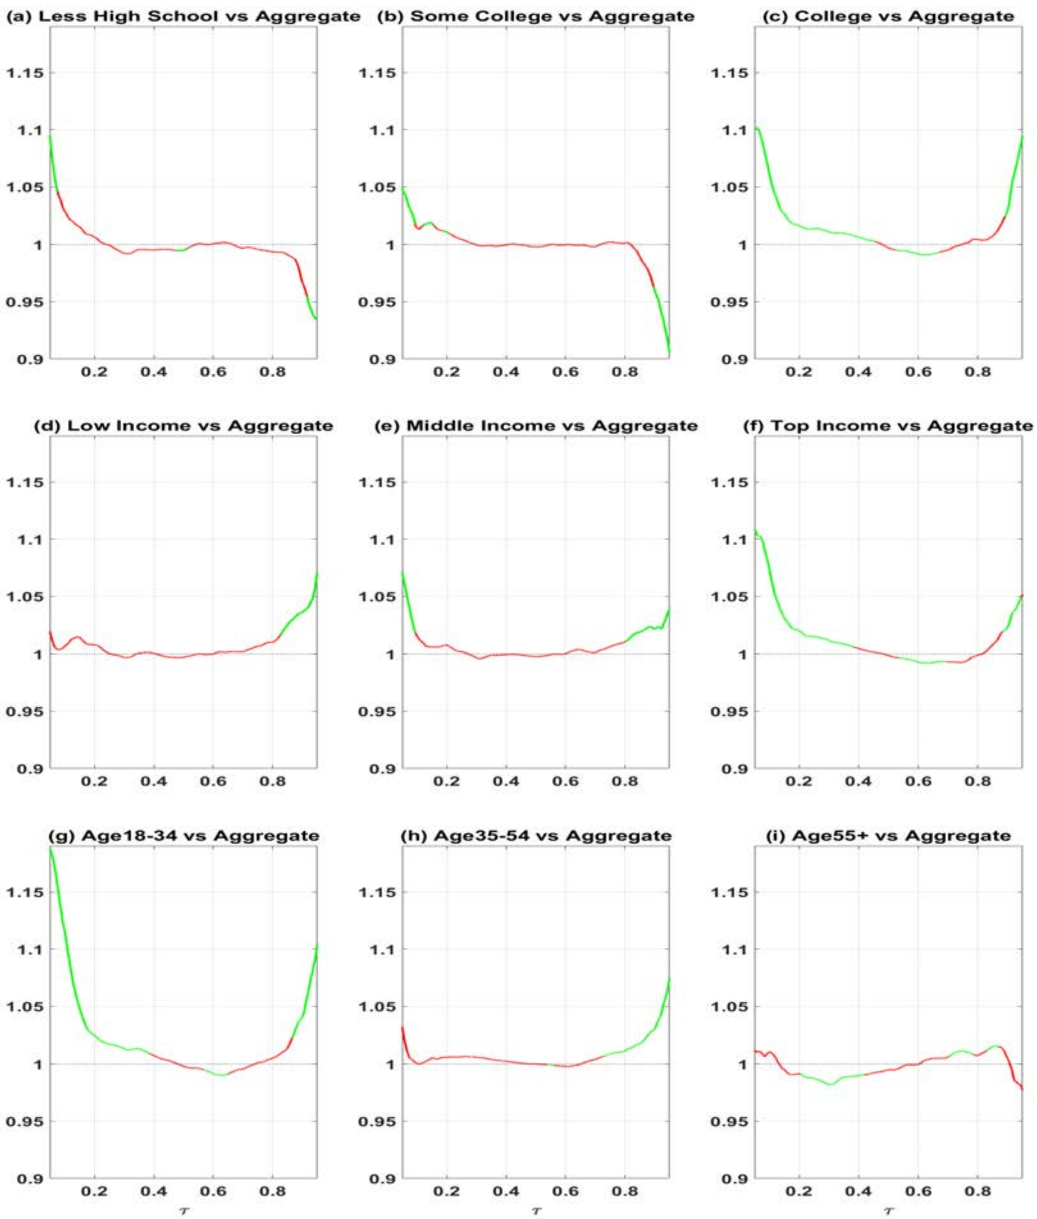

Note: Sample period spans 1986m1-2022m8. The plots are the relative quantile scores: quantile score from the QR model with household expectations from specific demographic groups relative to QR model with aggregate household expectations. A quantile score ratio greater than one (on the y-axis) indicates that, for a given quantile (on the x-axis), the QR density having aggregated expectations is more accurate than the density from the expectations of a specific demographic group. The higher the ratio, the greater the gains of aggregate household expectations. Regions where the line is colored green indicate that gains (or losses), at the given quantile, are statistically significant at the 10 percent level, as judged by a Diebold and Mariano (1995) and West-type (1996) test of equality of the quantile scores.

A.5 Marginal Value of Long-Run Survey Expectations

Figure A.12 compares, by quantile, the predictive accuracy of the QR-based inflation forecasts with and without long-run expectations for households (panel a), markets (panel b), professionals (panel c), and for the combined expectations measure (panel d). As in previous figures, ratios are computed as the quantile score with both short-run and long-run expectations relative to the quatile score from the specification with short-run expectations only. A ratio less than one thus indicates that the prediction from the specification with long-run expectations is more accurate.

As can be seen for all the expectations measures in Figure A.12, there are significant benefits to conditioning on long-run survey expectations, with the specification with household expectations gaining the most (especially at lower quantiles). For the most part, these gains are statistically significant. Similar accuracy gains are obtained when predicting future inflation both 24-months and 36-months out (see Figures A.16 and A.17).

Figure A.12: Value-Added of Long-Run Survey Expectations: Relative Quantile Score

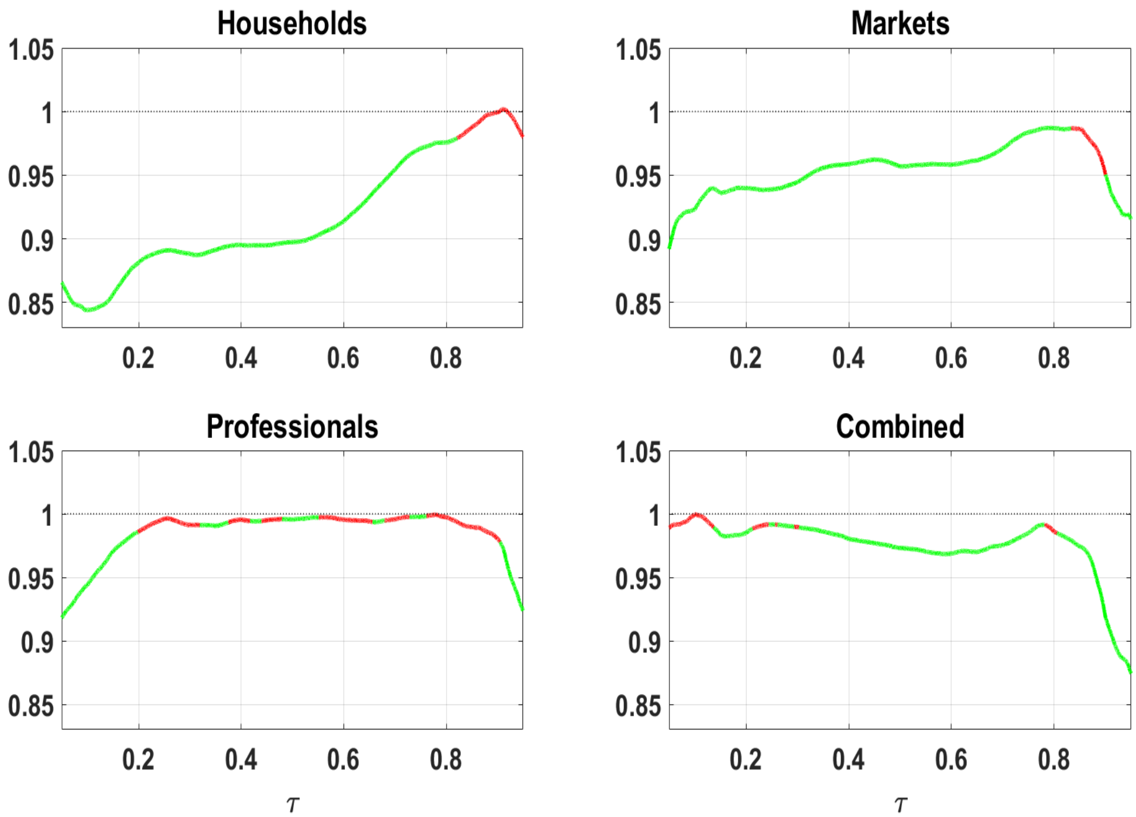

Note: Sample period 1986m1-2022m8. The plots are the relative quantile scores for each of the four inflation expectations measures. Each panel reports the relative quantile score, that is, the average quantile score of each specific agent's short-run and long-run expectations relative to the agent's short-run expectations only. Values less than one (on the y-axis) suggest that, for a given quantile (on the x-axis), incorporating information from an agent's short-run and long-run expectations is more accurate than information from an agent's short-run expectations on average over the sample period. The lower the ratio, the greater the gains from using information from agent's long-run expectations. Regions where the line is colored green indicate that gains (or losses), at the given quantile, are statistically significant at the 10 percent level, as judged by a Diebold and Mariano (1995) and West-type (1996) test of equality of the quantile scores.

A.6 Out-of-Sample (OOS) Forecasting Results

A.6.1 OOS: Density Forecast Accuracy Comparison

Table A.1: Density Forecast Accuracy Across Different Agents

(a) CRPS

Predictors(s): Constant + Lagged inflation + Expectations of:	ABG QR	NP QR	OLS/Normal
Households	1.039	0.965	0.980
Markets	1.048	0.955	0.968
Professionals	0.981	0.922	0.928
Naïve	0.964	0.974	0.983
Combined	1.020	0.894	0.914

(b) Tail-weighted CRPS

Predictors(s): Constant + Lagged inflation + Expectations of:	ABG QR	NP QR	OLS/Normal
Households	0.249	0.217	0.232
Markets	0.256	0.229	0.234
Professionals	0.245	0.223	0.225
Naïve	0.254	0.227	0.230
Combined	0.248	0.203	0.224

Note: Out-of-sample period 2000m5-2022m8. Panel (a) reports the CRPS averaged over the whole sample, panel (b) reports the average tail-weighted CRPS. Results given for the density forecasts constructed using the QR method of ABG, the nonparametric (NP) QR method of Mitchell et al. (2022), and assuming a linear Gaussian relationship (labeled "OLS/Normal"). "Combined" involves combining, by quantile, the different agents' expectations and then constructing the density forecast from the combined quantile forecasts.

A.6.2 OOS: One-Year Ahead Density Forecasts of CPI Inflation

Figure A.13: One-Year-Ahead Density forecast for CPI Inflation

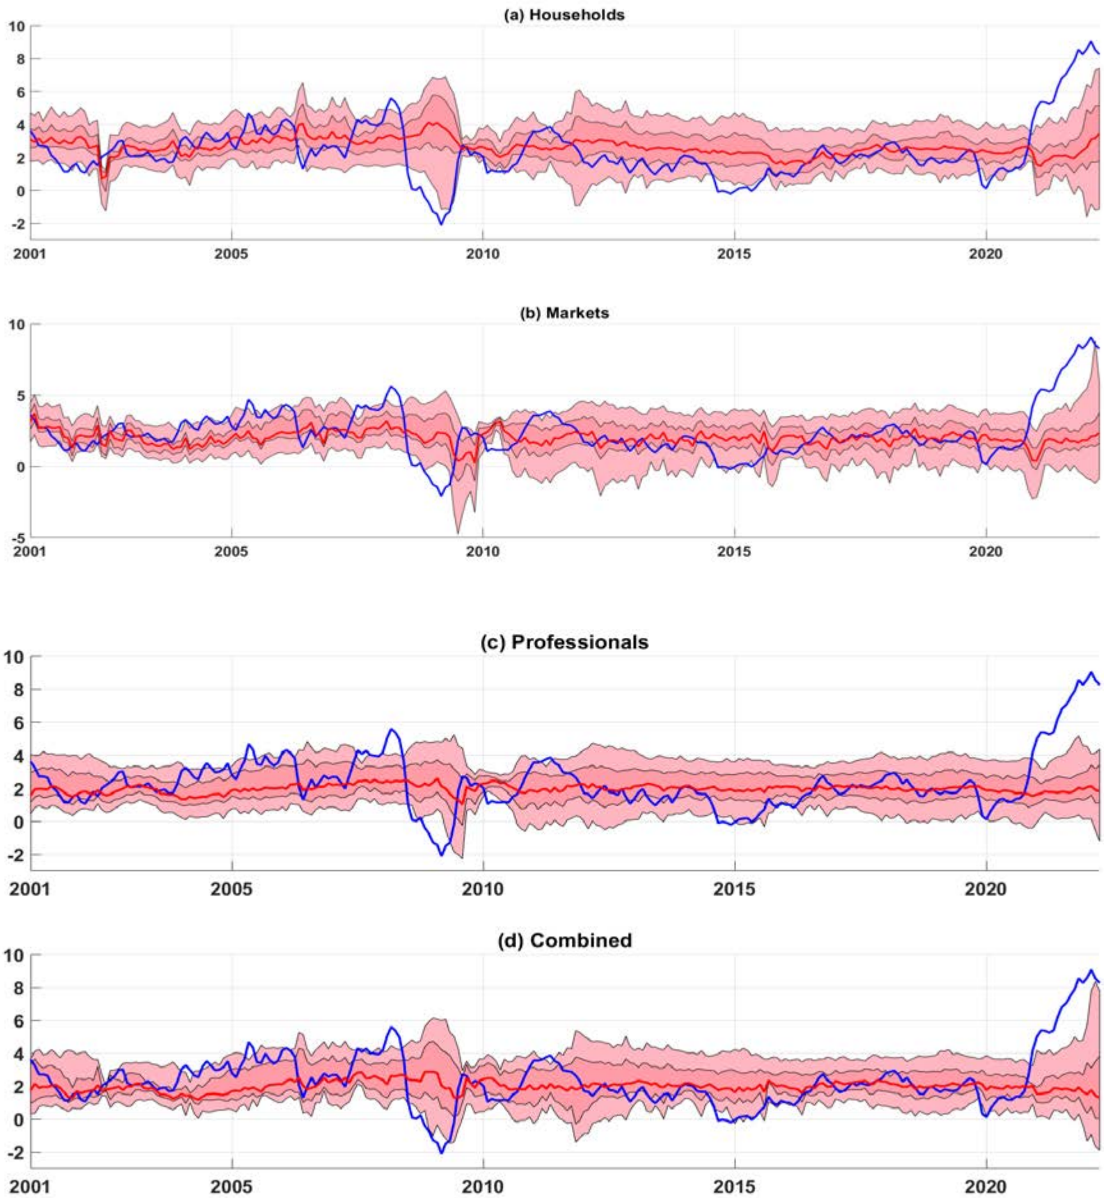

Note: Out-of-sample period: 2000m5-2022m8. Blue line corresponds to realized CPI inflation. Plotted are the 5, 25, 50 (solid red), 75, and 95 percent quantiles of the density forecast for year-ahead inflation, conditional on lagged inflation and either households, financial markets, professionals, or all three measures of inflationary expectations when combined. Density constructed using NP.

A.6.3 OOS: Calibration Diagnostics

Figure A.14: The Empirical Cumulative Distribution of the Probability Integral Transforms (PITs)

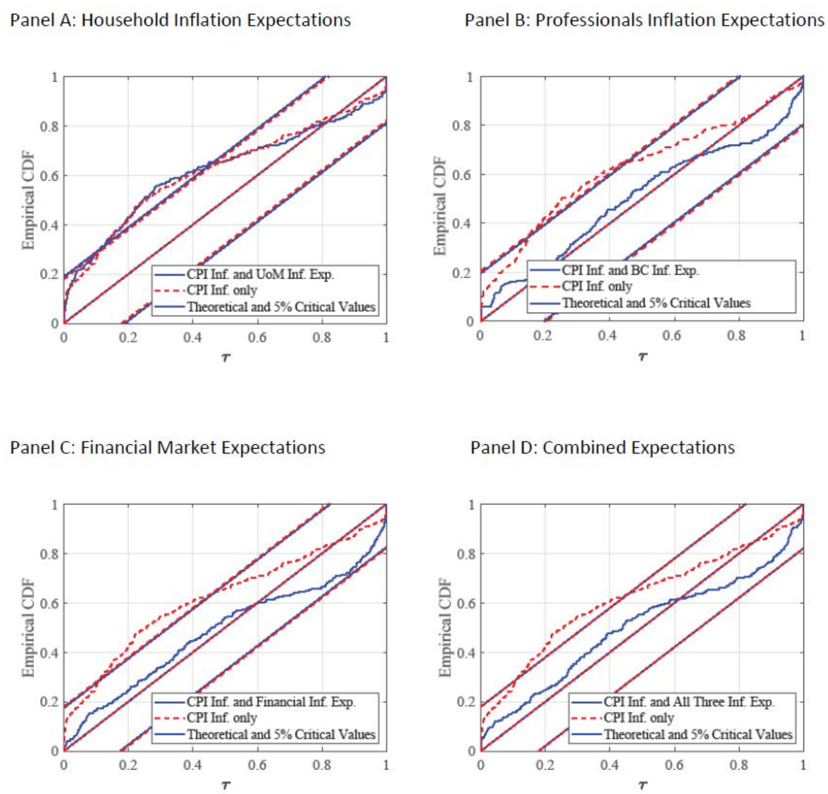

Note: Out-of-sample period: 2000m5-2022m8. The figures show the empirical CDF of the PITs (blue line) from the QR models with the indicated measure(s) of inflation expectations, the empirical CDF of the PITs (dashed red line) from the QR models with lagged inflation only (naïve expectations), and the CDF of the PITs under the null hypothesis of correct calibration (the 45-degree line), and the 5 percent critical value bands of the Rossi and Sekhposyan (2019) test.

A.6.4 OOS: Marginal Value of Long-Run Survey Expectations

Figure A.15: Marginal Value of Long-Run Survey Expectations

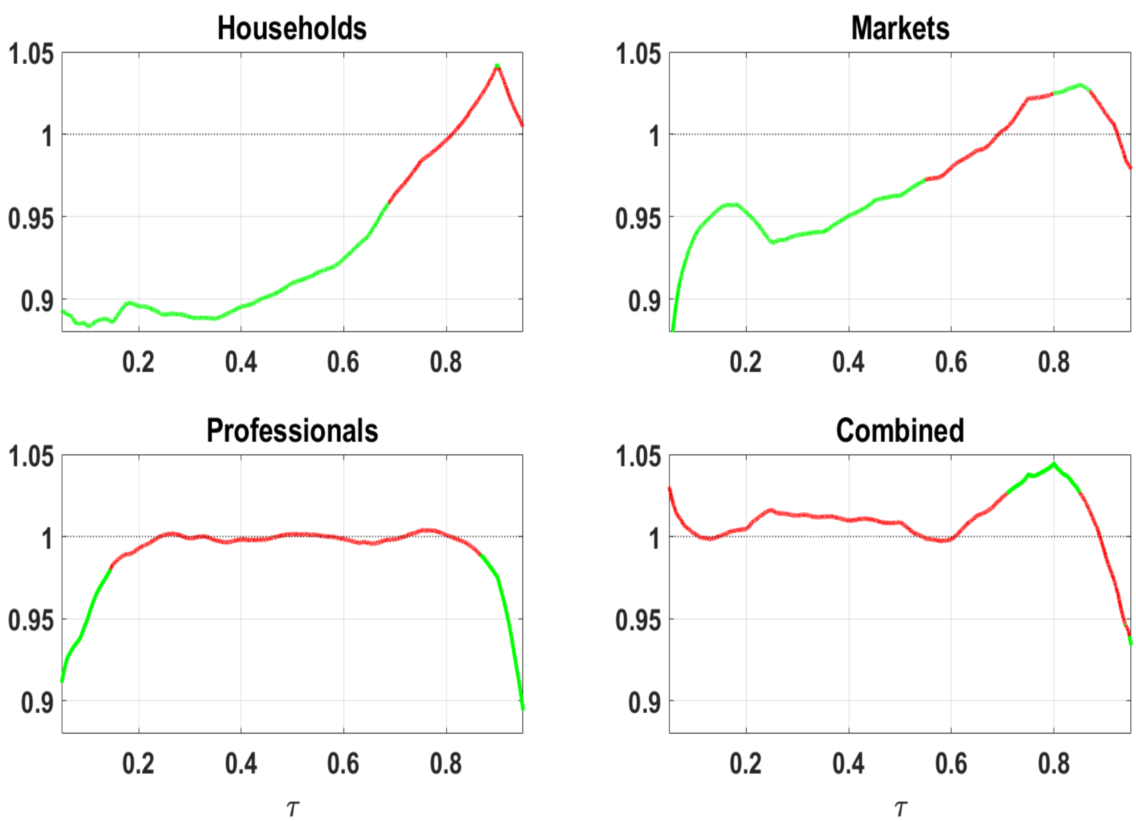

Note: Sample period 2000m5-2022m8. The plots are the relative quantile scores for each of the four inflation expectations measures. Each panel reports the relative quantile score, that is the average quantile score of each specific agent's short-run and long-run expectations relative to the agent's short-run expectations only. Values lower than one (on the y-axis) indicate that, for a given quantile (on the x-axis), incorporating information from an agent's short-run and long-run expectations is more accurate than information from an agent's short-run expectations on average over the sample period. The lower the ratio, the greater the gains from using information from an agent's long-run expectations. Regions where the line is colored green indicate that gains (or losses), at the given quantile, are statistically significant at the 10 percent level, as judged by a Diebold and Mariano (1995) and West-type (1996) test of equality of the quantile scores.

A.7 Robustness Across Forecast Horizons: 24 and 36 Months Ahead

Table A.2: Longer-Horizon Density Forecast Accuracy Metrics

(a) CRPS: 24 months

Predictors(s): Constant + Lagged inflation + Expectations of	NP QR	OLS/Normal
Households	0.826	0.848
Markets	0.780	0.810
Professionals	0.738	0.769
Naïve	0.835	0.850
Combined	0.727	0.769

(b) CRPS: 36 months

Predictors(s): Constant + Lagged inflation + Expectations of	NP QR	OLS/Normal
Households	0.816	0.839
Markets	0.787	0.815
Professionals	0.767	0.795
Naïve	0.823	0.844
Combined	0.750	0.786

Note: Sample period 1986m1-2022m8. Panels (a) and (b) report the CRPS averaged over the whole sample when forecasting 24 and 36 months ahead. Results provided for the density forecasts are constructed using the nonparametric (NP) QR method of Mitchell et al. (2022), and assuming a linear Gaussian relationship (labeled "OLS/Normal"). "Combined" involves combining, by quantile, the different agents' expectations and then constructing the density forecast from the combined quantile forecasts.

Figure A.16: Marginal Value of Long-run Survey Expectations: 24 Months Ahead

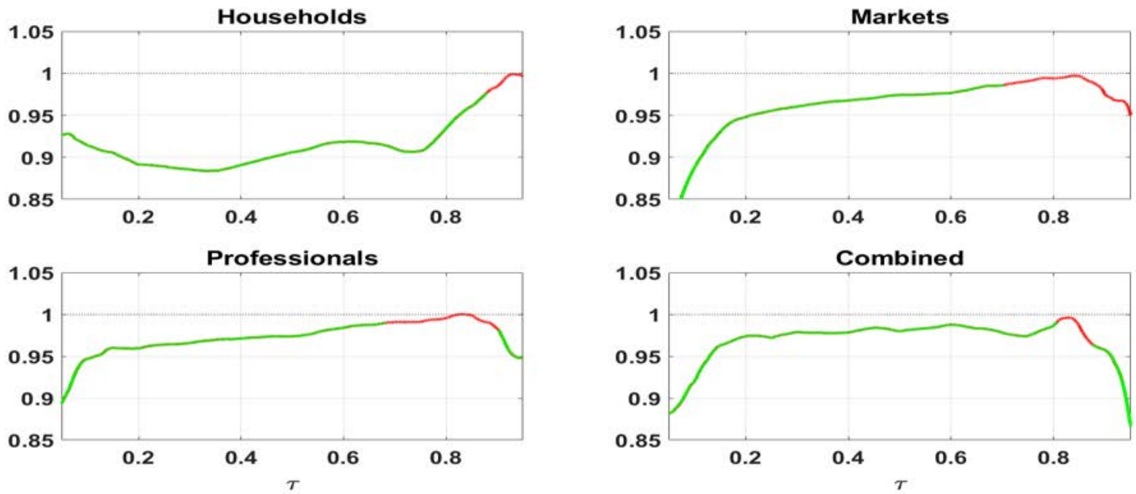

Figure A.17: Marginal Value of Long-Run Survey Expectations: 36 Months Ahead

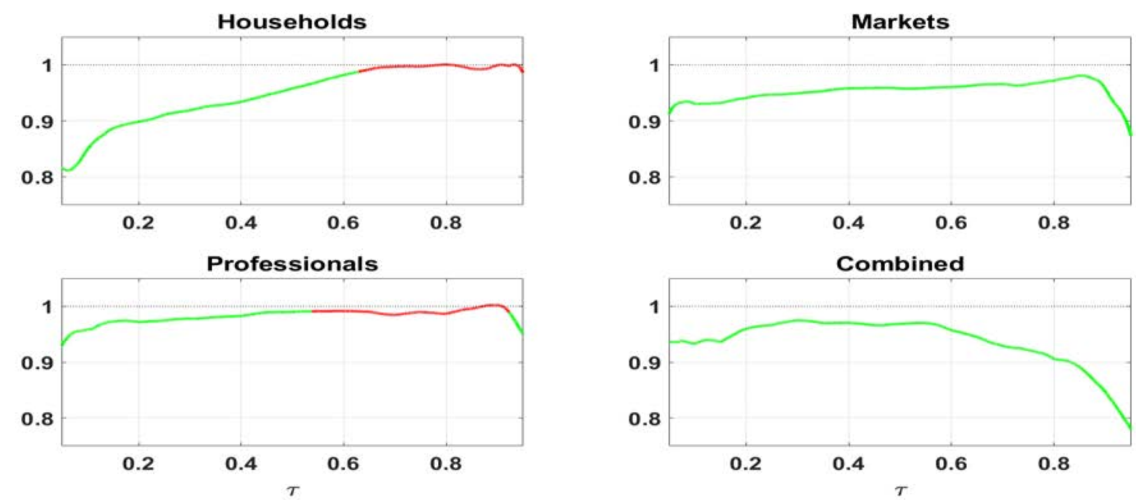

Note: Sample period 1986m1-2022m8. The plots are the relative quantile scores for each of the four inflation expectations measures. Each panel reports the relative quantile score, that is, the average quantile score of each specific agent's short-run and long-run expectations relative to the agent's short-run expectations only. Ratios less than one (on the y-axis) indicate that, for a given quantile (on the x-axis), incorporating information from an agent's short-run and long-run expectations is more accurate than information from an agent's short-run expectations on average over the sample period. Regions where the line is colored green indicate that gains (or losses), at the given quantile, are statistically significant at the 10 percent level, as judged by a Diebold and Mariano (1995) and West-type (1996) test of equality of the quantile scores.