

Wang, Tao

Working Paper

How do agents form macroeconomic expectations? Evidence from inflation uncertainty

Bank of Canada Staff Working Paper, No. 2024-5

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Wang, Tao (2024) : How do agents form macroeconomic expectations? Evidence from inflation uncertainty, Bank of Canada Staff Working Paper, No. 2024-5, Bank of Canada, Ottawa,
<https://doi.org/10.34989/swp-2024-5>

This Version is available at:

<https://hdl.handle.net/10419/297453>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

How Do Agents Form Macroeconomic Expectations? Evidence from Inflation Uncertainty

by Tao Wang

Financial Stability Department
Bank of Canada
TaoWang@bankofcanada.ca



Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

DOI: <https://doi.org/10.34989/swp-2024-5> | ISSN 1701-9397

©2024 Bank of Canada

Acknowledgements

This paper was originally circulated as "How Do Agents Form Inflation Expectations? Evidence from the Forecast Uncertainty." I thank Jonathan Wright, Chris Carroll, Greg Duffee, Corola Binder, Yuriy Gorodnichenko, Ricardo Reis, Shihan Xie, Rosen Valchev, Chenyu (Sev) Hou, Claudia Godbout, Oleksiy Kryvtsov, Michael Boutros, Olena Kostyshyna, Chris Sutherland, Thomas Pugh, Cars Hommes, Michael Boutros, Cosmin Ilut, the anonymous referee, participants of JHU Macro Brownbag Seminar, SNDE symposium, the Brownbag seminar at Bank of Canada, and IAAE annual meeting for invaluable comments. Declarations of interest: none.

Abstract

This paper studies the behaviors of *uncertainty* through the lens of several popular models of expectation formation. The full-information rational expectations model (FIRE) predicts that both the ex ante uncertainty and the variance of ex post forecast errors are equal to the conditional volatility of shocks to the fundamentals. Incomplete-information models such as Sticky Expectation (SE) and Noisy Information (NI) and non-rational models such as Diagnostic Expectations (DE) predict distinctive rankings of these moments. The paper also shows that uncertainty provides additional parametric restrictions to favor SE over NI as a model of information rigidity, although both predict similar aggregate patterns of forecast errors and disagreement.

Topics: Business fluctuations and cycles; Inflation and prices; Monetary policy and uncertainty
JEL codes: D84, E31, E71

Résumé

Cette étude analyse les comportements de *l'incertitude* à travers le prisme de différents modèles répandus de formation des anticipations. Le modèle des anticipations rationnelles fondées sur une information complète prédit que l'incertitude *ex ante* et la variance des erreurs de prévision *ex post* sont égales à la volatilité conditionnelle des chocs sur les déterminants des facteurs fondamentaux. Les modèles avec information incomplète (p. ex., les modèles avec information rigide ou avec information bruitée) et les modèles non rationnels (p. ex., le modèle des anticipations diagnostiques) prédisent des classements distinctifs de ces moments statistiques. L'étude montre également que l'incertitude conduit à des contraintes paramétriques supplémentaires qui rendent le modèle avec information rigide préférable au modèle avec information bruitée comme modèle de rigidité de l'information, même si les deux prédisent des profils globaux similaires d'erreurs de prévision et de divergence.

Sujets : Cycles et fluctuations économiques; Incertitude et politique monétaire; Inflation et prix
Codes JEL : D84, E31, E71

1 Introduction

Macroeconomists have developed different models of expectation formation to account for observed patterns of survey expectations that are inconsistent with the benchmark of full-information rational expectations (FIRE). Some of the models, such as Sticky Expectations (SE)¹ and Noisy Information (NI),² feature incomplete information and information rigidity in the form of sluggish responsiveness to aggregate shocks, as documented by [Coibion and Gorodnichenko \(2012, 2015\)](#). Other models, such as Diagnostic Expectation (DE) ([Bordalo et al., 2018](#)), implies overreaction to the news at the individual level.³ A hybrid model of DE and NI ([Bordalo et al., 2020](#)) has been developed to match the disagreement in addition to individual overreactions.

The estimation of these models⁴ has been primarily focused on the patterns of forecast errors, revision, and disagreement in the survey data. What remains underexplored in this literature are the behaviors of *uncertainty*, which, here, strictly refers to the dispersion of density forecasts. This paper first shows that, according to these non-FIRE models, forecasters' uncertainty differs from the fundamental volatility of the shocks for several reasons, such as incomplete information and imperfect responses to new information. Then it uses distinctive predictions regarding uncertainty by different models to identify the exact model of expectation formation.

Consider FIRE, first, for some intuition on why various models of expectation formation have distinctive predictions about the rankings and parameter restrictions regarding various forecasting moments. With complete information and an identical model as assumed, forecast uncertainty only reflects the conditional variance of the unforecastable component of the inflation. Therefore, the ex-ante uncertainty would be exactly equal to the variance of ex-post forecast errors and to the size of the conditional volatility of inflation.

In contrast, there is additional uncertainty of the variable to the conditional volatility given a perfectly updated information set in the two models featuring incomplete information or information rigidity, Sticky Expectations (SE) and Noisy Information (NI), but for different reasons.

In SE, the extra uncertainty arises because of lagged updating of information. It results in a higher uncertainty than the volatility of the unrealized shock. Meanwhile, the variance of ex-post forecast errors is reduced compared to FIRE as information rigidity attenuates the average forecast responses to shocks to inflation. These patterns are observed in both the New York

¹[Mankiw and Reis \(2002\)](#); [Carroll \(2003\)](#); [Reis \(2006\)](#).

²[Lucas \(1972\)](#); [Woodford \(2001\)](#); [Maćkowiak and Wiederholt \(2009\)](#).

³[Kohlhas and Walther \(2021\)](#) proposes an extended model of NI allowing for multiple unobserved components to reconcile the coexistence of under- and overreaction.

⁴For instance, [Coibion and Gorodnichenko \(2012, 2015\)](#); [Bordalo et al. \(2020\)](#).

Fed's *Survey of Consumer Expectations* (SCE) and the Federal Reserve Bank of Philadelphia's *Survey of Professional Forecasters* (SPF). In contrast, the additional uncertainty in NI comes from the noisiness of the information and endogenously determines agents' degree of reaction to the news in the Kalman filtering problem. The model predicts both uncertainty and variance of forecast error to be greater than the conditional volatility of the inflation due to the presence of noises. However, it accommodates a flexible relative size between the two depending on the noisiness of signals. The model-consistent noisiness of signals hence becomes a quantitative question.

Different from models of information rigidity, the canonical model of Diagnostic Expectation (DE), although assuming that agents extrapolate in the mean forecast of the variable, keeps the uncertainty equal to that in FIRE. At the same time, DE predicts an attenuated variance of forecast errors due to the mean-reverting nature of the overreaction to persistent shocks. A hybrid of Diagnostic Expectations and Noisy Information (DENI), proposed by [Bordalo et al. \(2020\)](#), entails both overreaction mechanisms and dispersed noisy information. Uncertainty is unambiguously larger than the conditional volatility of the variable due to the presence of noisy information. Meanwhile, the variation of forecast error could be either attenuated or amplified, depending on the parameter values of the model.

In addition to the model differentiation, uncertainty also provides extra moment restrictions to estimate model-specific parameters. The second part of this paper structurally estimates all models based on cross-moment predictions that jointly target different moments, forecast errors (FE), cross-sectional disagreements (Disg), and uncertainty (Var).⁵ The estimations further favor SE over NI mechanisms as a source of information rigidity in inflation expectations.

In particular, my estimates of the SE model report a sensible updating rate of around one-third per period for both types of agents, which is aligned with many estimates in the literature. In contrast, the estimated noisiness of public and private signals in NI are unrealistically high and unstable, i.e., at a ballpark value of 3 percentage points or higher relative to an unconditional standard deviation of inflation of 0.8 in headline CPI or 0.4 in core PCE in the sample period. The intuition behind the poor fit of canonical NI is that Kalman filtering requires agents to efficiently decide their responsiveness to new information based on the prior uncertainty and noisiness of information. Therefore, effectively, only extremely imprecise signals could result in persistent information rigidity as we see in the data.

In addition to the evidence for rigidity, my estimates of DE and DENI do suggest a coexistent overreacting mechanism at the individual level, i.e., a non-zero fraction of agents in the economy

⁵Some other contemporaneous papers also structurally estimate theories on expectation formation based on single or multiple moments of surveyed expectations, such as [Giacomini et al. \(2020\)](#); [Xie \(2023\)](#); [Bordalo et al. \(2020\)](#); [Farmer et al. \(2021\)](#); [Ryngaert \(2017\)](#). However, none of these studies explore the behaviors of uncertainty.

have a positive degree of overreaction $\hat{\theta}$. This is consistent with the findings in the literature showing the coexistence of rigidity and overreaction (Angeletos et al. (2021); Kohlhas and Walther (2021)).

In addition to the benchmark estimation based on the assumption that inflation follows a stationary AR(1) process with constant volatility, I also extend the analysis to consider an alternative process of stochastic volatility (SV) as formulated in Stock and Watson (2007). The alternative assumption flexibly accommodates components of different persistence and time-varying volatility. I also alter the estimation across low and high inflation periods (before and after 2020) to examine the possible state-dependence of expectation formation. This relates to several studies showing that the information rigidity and degree of underreaction tend to be lower in periods of high volatility.⁶

Through the lens of SE, both households and professionals have increased the updating rate of new information in the recent period. The estimates of DENI, however, seem to suggest divergent patterns of the two types of agents. In particular, professionals have shifted from overreaction on average and relatively precise individual information⁷ to underreaction and more dispersed information in the high inflation episode. Households, in contrast, have become overreactive as inflation has elevated in recent years, consistent with the intuitive pattern that news coverage and social interactions have significantly resulted in household attentiveness to inflation news.

Related Literature

This paper is related to four strands of literature. First, it is related to a series of empirical studies directly testing and estimating various theories on expectation formation using survey data. Early examples of such work include Mankiw and Reis (2002), Mankiw et al. (2003), Carroll (2003), and Branch (2004). More recent examples include Coibion and Gorodnichenko (2012, 2015); Coibion et al. (2018), which test common implications of various theories with different micro-foundations. In addition to testing particular sets of theories, there are also a number of papers showing that people’s expectations are driven by individual heterogeneity, such as socioeconomic characteristics, cognitive abilities, and experiences of macroeconomic histories (Malmendier and Nagel (2015), Das et al. (2017), and D’Acunto et al. (2019)⁸). Compared to reduced-form estimation based on a particular moment, this paper shares the spirit of Giacomini et al. (2020); Xie (2023); Valchev and Gemmi (2023) in carrying out a structural

⁶Coibion and Gorodnichenko (2015); Xie (2023); Weber et al. (2023).

⁷Consistent with the findings of Bordalo et al. (2018).

⁸See D’Acunto et al. (2023) for a thorough survey of the empirical evidence of heterogeneous inflation expectations and their drivers.

estimation of models of expectation formation using multiple moments in the survey, taking into account factors such as measurement error and the strategic incentives of forecasters. However, none of these studies directly use density forecasts or surveyed uncertainty. This is one theme on which this paper differs from the existing literature.

Second, this paper is related to the macroeconomic literature on measuring uncertainty, especially those using survey data.⁹ Various proxies of uncertainty that have often been used include ex-ante cross-sectional disagreement (Bachmann et al., 2013), *approximated* conditional volatility based on time-series forecasting (e.g., Jurado et al. (2015)), and ex-post forecast errors (Bachmann et al., 2013; Rossi and Sekhposyan, 2015). Some studies empirically evaluated the correlation between the aforementioned proxies and the uncertainty measured by the dispersion of density forecasts. Zarnowitz and Lambros (1987) made a clear conceptual distinction between disagreement and uncertainty and found a very low correlation between the two in an early sample of SPF. Follow-up studies (Rich and Tracy, 2010; D’Amico and Orphanides, 2008; Abel et al., 2016; Glas, 2020; Rich and Tracy, 2021) echoed such a finding, mostly based on SPF data, although Bomberger (1996); Giordani and Söderlind (2003); Lahiri and Sheng (2010) arrive at different conclusions. One point that was often not explicit in this literature is that the relationship between various ex-ante uncertainty, ex-post forecast errors, and disagreement depends on the mechanisms of expectation formation.¹⁰ My paper explicitly compares various models of expectation formation, which predicts distinctive relationships across these moments.

Third, Manski (2004), Delavande et al. (2011), Manski (2018), and many other papers have long advocated for eliciting probabilistic questions measuring subjective uncertainty in economic surveys. Although the initial suspicion concerning people’s ability to understand, use, and answer probabilistic questions is understandable, Bertrand and Mullainathan (2001) and other work have shown that respondents have the consistent ability and willingness to assign a probability (or “percent chance”) to future events. Armantier et al. (2017) provide a thorough discussion on designing, experimenting, and implementing consumer expectation surveys to ensure the quality of the responses.¹¹ Broadly speaking, the literature has argued that going beyond the revealed preference approach, the availability of survey data provides economists with direct information on agents’ expectations and helps avoid imposing arbitrary assump-

⁹Survey-based uncertainty measures are among various methods seen in the literature, such as using news texts (Bloom, 2009), econometric methods (Jurado et al. (2015)), and market derivatives (e.g., VIX index), as summarized in Cascaldi-Garcia et al. (2023). Besides, Binder (2017) creates a novel measure using *household* survey data based on the insight from cognitive science that people tend to round numbers when facing higher uncertainty.

¹⁰Gambetti et al. (2023) shows that under dispersed information, uncertainty measured as forecast error variance entails both fundamental volatility and dispersion of the information.

¹¹Others include Van der Klaauw et al. (2008) and Delavande (2014), etc. See Bassetti et al. (2023) for a complete survey on methods of extracting information from density forecasts and their macroeconomic applications.

tions. This insight holds for not only point forecast but also, and even more importantly, for uncertainty, because for any economic decision made by a risk-averse agent, not only the expectation but also the perceived risks matter a great deal.

Finally, the literature that has been originally developed under the theme of forecast efficiency (Nordhaus, 1987; Davies and Lahiri, 1995; Clements, 1997; Faust and Wright, 2008; Patton and Timmermann, 2012) provides a framework for analyzing the dynamics of uncertainty that is useful for the purpose of this paper. The focus of the forecasting efficiency literature is evaluating forecasters' performance and improving forecasting methodology, but it can be adapted to test the theories of expectation formation of different types of agents. This is especially relevant to this paper, where I focus on uncertainty.

2 Theoretical Benchmark and Basic Facts

2.1 Full-information rational expectation (FIRE)

Assume the underlying data generating process of y_t is AR(1) with a persistence parameter $0 < \rho < 1$ and i.i.d. shock ω_t whose time-invariant volatility is σ_ω .

$$y_t = \rho y_{t-1} + \omega_t, \quad \omega_t \sim N(0, \sigma_\omega^2) \quad (1)$$

FIRE benchmark assumes that all agents perfectly observe y_t at time t and understand the true process of y . Therefore, the individual forecast is $\rho^h y_t$, which is shared by all agents. Therefore, it is also equal to the average forecast.

Both individual and population forecast errors are simply the realized shocks between $t + 1$ and $t + h$.

$$\overline{FE}_{t+h|t}^* = - \sum_{s=1}^h \rho^{s-1} \omega_{t+h+1-s} \quad (2)$$

I use the superscript of $*$ to denote all the moments according to FIRE. It is easy to see that the forecast error is orthogonal to information available till time t . This provides a well-known null hypothesis of FIRE.¹²

¹²Another well-known prediction of FIRE is that forecast errors of non-overlapping horizon are not correlated, namely, $Cov(\overline{FE}_{t+h|t}^*, \overline{FE}_{t+s+h|t+s}^*) = 0 \quad \forall s \geq h$. This is not the case within h periods, as the realized shocks

The unconditional variance of h -period-ahead FE, or equivalently, the expected value of its square (due to zero unconditional mean), is equal to the following. (\bullet indicates that it is unconditional on t .)

$$\overline{FE}_{\bullet+h|\bullet}^{*2} = \sum_{s=1}^h \rho^{2(s-1)} \sigma_{\omega}^2 \quad (3)$$

The uncertainty about future y simply comes from uncertainty about unrealized shocks between t and $t + h$. With the same model in mind (Equation 1) and the same information y_t , everyone's uncertainty is equal to the weighted sum of the future volatility before its realization (Equation 4), which is exactly equal to the variance of forecast errors, $\overline{FE}_{\bullet+h|\bullet}^{*2}$.

$$\overline{\text{Var}}_{\bullet+h|\bullet}^* = \sum_{s=1}^h \rho^{2(s-1)} \sigma_{\omega}^2 \quad (4)$$

Another testable implication of rationality lies in the revision of uncertainty. Moving from t to $t + 1$, for instance, the revision in uncertainty is simply a negative constant independent of time (Equation 5). There is an unambiguous reduction in uncertainty (or information gain in the forecasting literature) as more and more shocks have been realized. In the most intuitive case, from the one-step-ahead forecast at $t - 1$ to that of t , i.e., $h = 1$, the variance drops by exactly the resolution of the uncertainty of ω_t , which is σ_{ω}^2 , to zero uncertainty.

$$\overline{\text{Var}}_{t+h|t+1}^* - \overline{\text{Var}}_{t+h|t}^* = -\rho^{2(h-1)} \sigma_{\omega}^2 \quad (5)$$

Lastly, FIRE predicts a zero disagreement, and it is so regardless of the behaviors of forecast errors and uncertainty: $\overline{Disg}_{\bullet+h|\bullet}^* = 0$.

2.2 Density surveys of inflation

Both forecast errors and disagreement are easily computed with cross-sectional surveyed expectations. The uncertainty is only available in density forecasts. This paper uses two datasets of density forecasts of inflation by professionals and households, where respondents are asked to assign probabilities to various ranges of values of future inflation.

in overlapping periods enter both forecast errors.

The *Survey of Professional Forecasters* (SPF) collects professionals’ individual density forecasts of core CPI and core PCE inflation since 2007.¹³ In each quarter, density forecasts of fourth-quarter-to-fourth-quarter inflation in the current year and next year are elicited. As a result, both quarterly revisions and annual revisions can be calculated. This makes it possible to directly test the implications of the revisions in uncertainty. In addition, because the forecasts are reported in all quarters of a year regarding Q4/Q4 inflation, the forecast horizons change within a year.

The New York Fed’s *Survey of Consumer Expectations* (SCE), which started in 2013, also asks households to report their distribution forecasts of 1-year- and 3-year-ahead inflation for various ranges of values each month.¹⁴ This allows for comparing the 3-year-ahead forecast at time $t - 3$ with the 1-year-ahead forecast at $t - 1$. Since the maximum duration for households to stay in the panel is 12 months (for about one-third of the households), forecast revisions can only be examined at the population level. One major difference between SCE and SPF is that the former elicits fixed-horizon expectations, while the latter elicits fixed-event ones.¹⁵

Converting expressed density forecasts based on externally divided histograms into an underlying subjective distribution requires a density estimation. I closely follow [Engelberg et al. \(2009\)](#)’s method to estimate the density distribution of each respondent in SPF.¹⁶ For SCE, I directly use the estimates by the New York Fed ([Armantier et al., 2017](#)), following the same method.

To address potential biases from extreme answers or data errors, I remove outliers from both the top and bottom 1 percent of the SPF sample’s mean forecast and uncertainty estimates and 3 percent of the SCE’s sample, respectively.¹⁷

2.3 Higher average uncertainty than variance of forecast errors

Table 1 reports the empirically computed moments of professionals and households. For the benchmark assumption of AR(1) process of inflation, I restrict the sample to be before March 2020, in which inflation remained stably low. This was also the period for which most of the empirical results showing the deviations from FIRE in inflation expectations were established.

Regardless of the realized inflation volatility, the survey moments alone exhibit patterns

¹³Previous studies used an extended sample of the density forecast of the GDP deflator starting from 1968 in the predecessor of SPF, or the NBER-ASA Economic Outlook Survey.

¹⁴The survey respondents are guaranteed to assign probabilities to all bins that sum to one as a feature of the online survey design.

¹⁵See [Clements et al. \(2023\)](#) for a detailed discussion of these differences in survey structure.

¹⁶See the [Python program](#) with detailed steps of estimation.

¹⁷This entails dropping 6,528 observations for mean forecasts and 5,096 observations for uncertainty, out of a total of 68,887 observations.

Table 1: Moments of Inflation and Expectations

	SPF	SCE
InfAV	0	0
InfVar	0.219	1.282
InfATV	0.194	1.206
FE	0.125	1.812
FEVar	0.136	0.935
Disg	0.161	2.805
Var	0.213	1.749

This table reports the moments of demeaned realized inflation and inflation expectations, respectively, for the SPF’s core CPI forecast and the SCE’s headline CPI forecast. They include the mean ($InfAV$), variance ($InfVar$), and auto-covariance ($InfATV$) of the inflation; the mean (FE), variance ($FEVar$), and auto-covariance ($FEATV$) of population forecast error; the mean disagreement ($Disg$); and the mean uncertainty (Var). The sample period is 2007M1–2020M3 for SPF and 2013M1–2020M3 for SCE. For SCE moments, both disagreement and uncertainty are computed using the regression residuals of individual mean forecasts and uncertainty after controlling for individual fixed effects.

inconsistent with the predictions of FIRE: ex-ante uncertainty is sizably larger than the average squared ex-post forecast error in both SPF and SCE, in addition to a positive disagreement.

Additional patterns emerge when an estimated inflation process is considered. The quarterly core CPI inflation during the sample period had persistence of $\hat{\rho} = 0.99$ and volatility $\hat{\sigma}_\omega = 0.23$. The conditional volatility of four-quarter-ahead annualized inflation, $\sum_{s=1}^4 \rho^{2(s-1)} \sigma_\omega^2$, is therefore approximately equal to 0.162. This should be equal to the size of uncertainty and variance of forecast error under FIRE. In the data, however, quarterly professional forecasts of the core CPI inflation have an uncertainty of 0.213, which is larger than conditional volatility. Meanwhile, the variance of forecast error, $FEVar = 0.136$, is smaller than the unconditional volatility. In addition, $Disg$ is 0.161, significantly positive instead of zero like in FIRE.

Household expectations exhibit different rankings of these moments. During the sample period of 2013–2020, the monthly headline CPI inflation is estimated to have a persistence parameter of 0.98 and volatility of 0.45.¹⁸ This, when translated into a 12-month-ahead annualized inflation forecast, implies a conditional volatility of 1.972. Meanwhile, household forecasts had a variance of forecast error of 0.935, a disagreement of 2.805, and uncertainty of 1.75. Both the variance of forecast error and uncertainty are smaller than the conditional volatility. Note that these values are computed based on residuals by first controlling for the individual fixed effects. This follows from the rich empirical evidence¹⁹ that individual-specific effects such as

¹⁸The higher volatility compared to that of core CPI is due to a higher frequency and inclusion of more volatile items in the CPI basket.

¹⁹Malmendier and Nagel (2015), Das et al. (2017), D’Acunto et al. (2019).

demographics and experience result in persistent differences in expectations.

We will revisit these stylized patterns through the lens of alternative theories of expectation formation discussed in Section 3.

2.4 Correlations of forecast errors, disagreement, and uncertainty

Despite the stark differences in magnitudes between professionals' and households' forecasting moments, both types of agents share common patterns in terms of the relationship across various moments. Figure 1a, 1b, and 1c plot the population uncertainty against expected inflation, squared forecast errors, and disagreements, respectively.

Both academic research (Friedman, 1977; Ball et al., 1990; Ball, 1992) and anecdotal evidence indicate that higher inflation or expected inflation is generally associated with greater inflation uncertainty. This relationship is partially reflected in the positive correlation observed between expected inflation and directly measured forecasting uncertainty among both households and professional forecasters, with correlation coefficients of 0.54 for SPF and 0.58 for SCE between 2007–2021. However, it is important to note that this correlation is primarily driven by recent dynamics since 2020, with the onset of higher inflation. Prior to 2020, the correlation coefficients were notably smaller (0.15 for SPF and -0.24 for SCE).

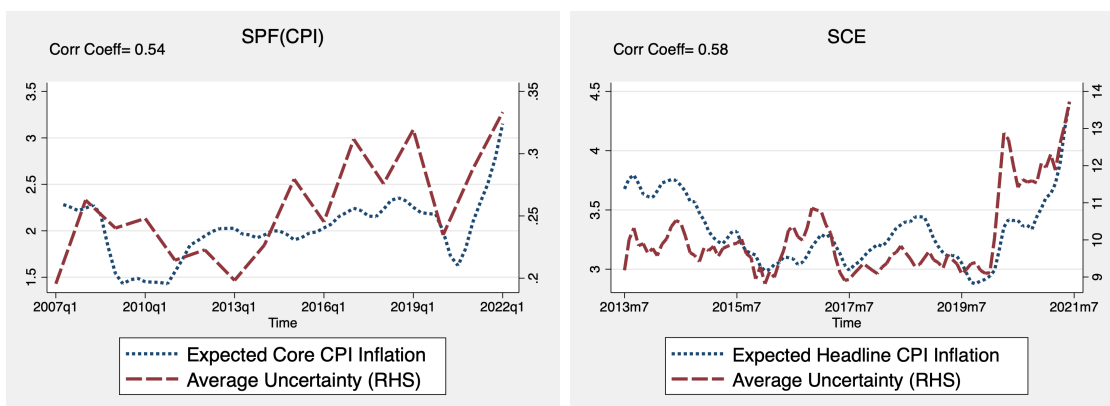
Figure 1b inspects the relationship between the size of the forecast error and uncertainty. According to the benchmark prediction under FIRE, the ex-ante forecast uncertainty is equal to the variance of ex-post forecast errors on average, as shown in Table 1. In the data, the correlation coefficients of the two are 0.26 and 0.42 for SPF core CPI forecasts and SCE's forecasts, respectively. Excluding the post-2020 sample yields even smaller correlation coefficients. The fact that the two are at most weakly correlated is inconsistent with the FIRE benchmark prediction.

Figure 1c plots the relationship between disagreement and uncertainty. A large body of empirical literature in macroeconomics uses disagreement, which is often more available than uncertainty, as a proxy of the latter. This practice implicitly assumes some form of deviation from the benchmark FIRE, because in FIRE, regardless of the inflation process, disagreement should always be zero, and it is therefore not correlated with the average uncertainty.²⁰ The empirical correlation between disagreement and uncertainty is indeed weakly positive, which is

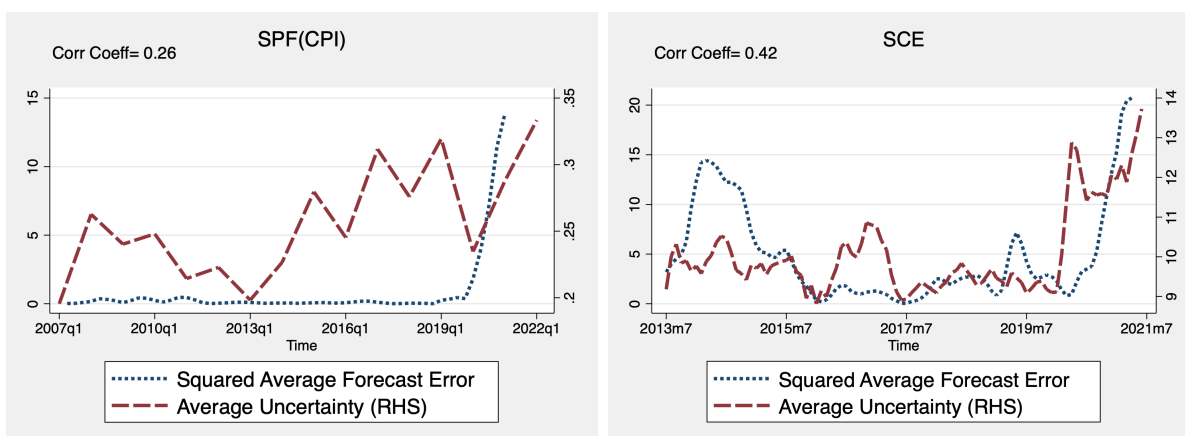
²⁰Several studies since Zarnowitz and Lambros (1987) found conflicting evidence about the correlation between disagreement and uncertainty. (Bomberger, 1996; Giordani and Söderlind, 2003; Rich and Tracy, 2010; Lahiri and Sheng, 2010; D'Amico and Orphanides, 2008; Binder, 2017; Glas, 2020; Rich and Tracy, 2021). Most of the comparisons are based on professional forecasters, with the exception of Binder (2017), which, by measuring the uncertainty based on rounding, found a high correlation between the two measures. More recently, Manski (2018) points out that much empirical work has confused dispersion with uncertainty.

Figure 1: Uncertainty and Other Moments

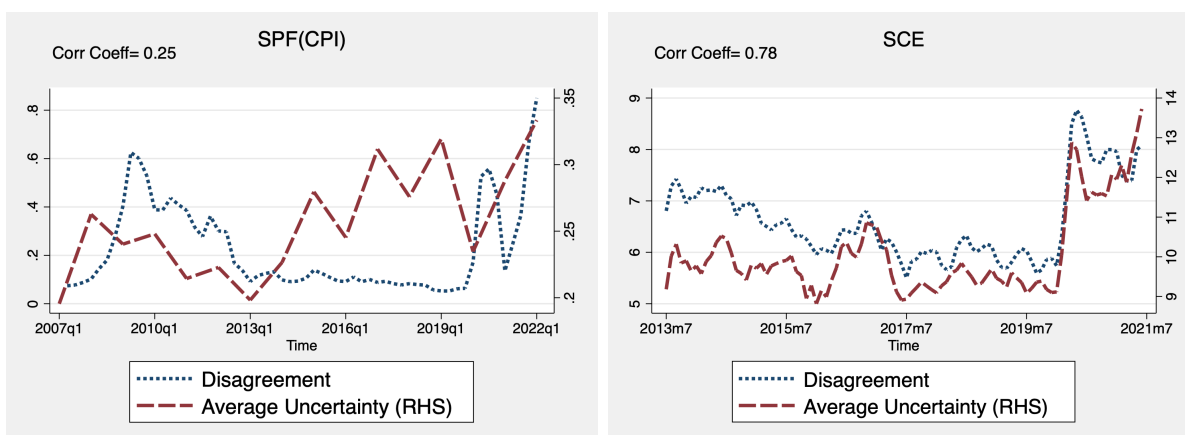
(a) Expected Inflation and Uncertainty



(b) Squared Forecast Errors and Uncertainty



(c) Disagreement and Uncertainty



Note: From left to right: SPF's forecasts of core CPI and SCE's household forecast of headline CPI. From top to bottom: uncertainty (dashed line) versus expected inflation (dotted line) with a correlation coefficient of 0.54 and 0.58, respectively; uncertainty (dashed line) versus square of the realized forecast errors (dotted line) with a correlation coefficient of 0.26 and 0.42, respectively; uncertainty (dashed line) versus disagreements (dotted line) with a correlation coefficient of 0.25 and 0.78, respectively.

0.25 for professionals and 0.78 for households. The positive correlation between the two was primarily driven by post-2020 dynamics, which saw rapidly rising inflation.

To summarize, compared to the FIRE prediction of a perfect correlation between ex-ante uncertainty and the size of ex-post forecast error in the benchmark FIRE, the empirical patterns of professional and household expectations exhibit divergent and time-varying behaviors of the two. In particular, the uncertainty is more correlated with expected inflation, disagreement, and size of forecast errors when the level of inflation is already high.

2.5 Inefficient revisions in uncertainty

Under the AR(1) process of inflation, FIRE predicts an unambiguous reduction in uncertainty as one approaches the date of realization, where the drop is exactly equal to the conditional volatility of the realized shocks σ_ω^2 .

Figure 2 plots the average revision in uncertainty in SPF and SCE. Since the individual-specific revisions are not available in SCE, I instead calculate the revisions in the average uncertainty across all respondents or, more specifically, the difference between 1-year-ahead uncertainty and 3-year-ahead forecasts made 2 years ago. In most of the years, both histograms suggest that revisions are left-skewed relative to zero. This implies that, on average, forecasters feel more certain about their nowcasts relative to their forecasts made before. However, over the entire sample, there is always a positive fraction of forecasters who revise uncertainty upward, which is inconsistent with the benchmark prediction with the AR(1) process of inflation. Positive uncertainty revisions were particularly common in the sample after 2020, a period with rapidly rising inflation caused by a combination of various demand and supply shocks.

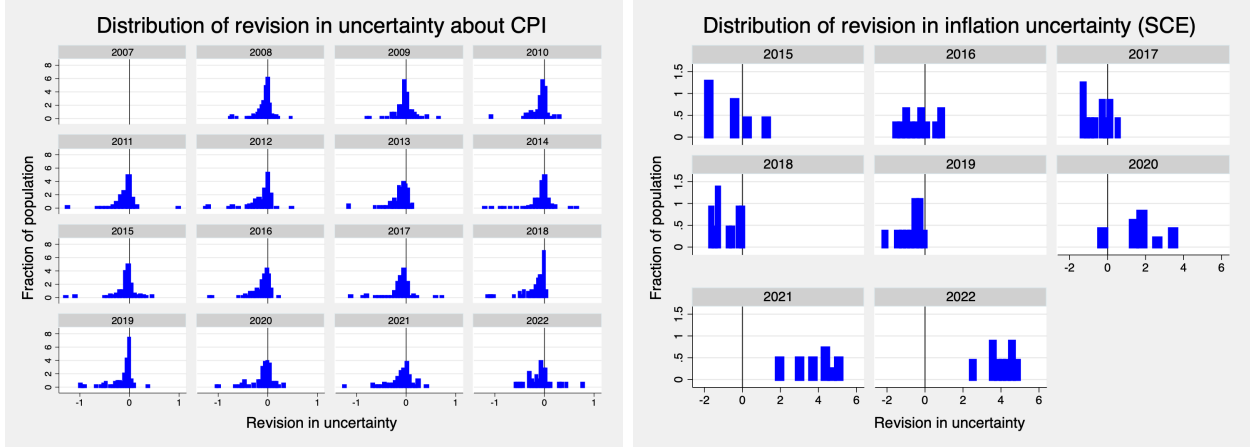
It is worth noting that the existence of positive uncertainty revisions does not necessarily reject the FIRE assumption in more general conditions. The pattern could also be reconciled by alternative models where the inflation volatility is stochastic instead of deterministic. With the former scenario, newly arrived information may cause an upward revision in the conditional perceived uncertainty of inflation, even though uncertainty in the period elapsed has resolved. Therefore, in the latter part of this paper, I explore the patterns of uncertainty in conjunction with the alternative assumption of stochastic volatility.

A formal test based on uncertainty takes the following form.²¹

$$\text{Var}_{i,t|t} - \text{Var}_{i,t|t-1} = \alpha^{\text{var}} + \beta^{\text{var}}(\text{Var}_{i,t-1|t-1} - \text{Var}_{i,t-1|t-2}) + \psi_t^{\text{var}} + \zeta_{i,t}^{\text{var}} \quad (6)$$

²¹This in the spirit of forecasting efficiency by Nordhaus (1987). It is an extension of revision tests on mean forecasts by Fuhrer (2018) to uncertainty.

Figure 2: Distribution of Uncertainty Revision



Note: The revisions in SPF (left) are calculated at the individual level and is computed as the difference between forecast uncertainty in quarter q about current-year Q4/Q4 inflation and that of the next-year Q4/Q4 inflation in quarter $Q - 4$. The revisions in SCE (right) are calculated at the population level, and it is computed as the difference between the average 1-year-ahead uncertainty at month m and the 3-year-ahead uncertainty at $m - 24$.

Equation 5 predicts that under FIRE, individual uncertainty revisions are all identical and equal to the innovation of the conditional volatility of inflation. This means that under FIRE the size of revisions in uncertainty is the same by all forecasters and hence should be fully absorbed by either the time-invariant constant α^{var} or the time-varying fixed effect ψ_t^{var} . Meanwhile, the auto-correlation coefficient β^{var} takes the value of zero under FIRE. A higher value of β^{var} indicates a slower speed of the drop in uncertainty, or forecast inefficiency, possibly due to information rigidity.²²

The two aforementioned regressions need to be adapted to be strictly consistent with the specific data structure in SPF and SCE. In particular, the revision in SPF is computed between the forecasts of the current-year Q4/Q4 inflation and the forecasts made four quarters before regarding the next-year Q4/Q4 inflation. The lagged revision, a measure of past information, was made four quarters before. For SCE, revisions and lagged revisions are relative to forecasts made 24 months before. This is critical, as revisions are expected to be correlated within the forecast horizon even under the assumption of FIRE. Furthermore, since individual revisions are not observed in SCE, I can only run regressions using average expectations and uncertainty. Hence, I cannot control for time fixed effects.

Table 2 reports regression results. The first column tests the mean revision against the null hypothesis of zero. For professional forecasters, the mean revisions in uncertainty are negative

²²This interpretation corresponds to Equation 10 for SE and Equation 16 for NI.

and statistically significant, confirming our observation from Figure 2 that forecasters are more certain about current inflation compared to their previous-year forecast.

The second to fourth column in the bottom panel of Table 2 show that revision in uncertainty of non-overlapping forecasts are serially correlated in both SPF and SCE. SPF forecasters' uncertainty revision from 1 year ago positively predicts current revisions in uncertainty. With the aggregate information innovation to be absorbed by time fixed effects and the constant, the coefficients of individual past revisions remain significant.

Households in SCE exhibit similar patterns of rigidity despite some differences with SPF professionals. In particular, the average revision in uncertainty between 3-year-ahead to 1-year-ahead forecasts is significantly negative when past revisions and squared realized forecast errors are controlled (InfExpFE_2). This implies that, on average, households are also more certain in their 1-year-ahead forecast than in their 3-year-ahead forecast. In addition, past uncertainty revisions are significantly correlated with current revisions, although negatively. Upward revisions 1 year ago are usually followed by downward revisions later. The coefficients remain significant even though I control for the size of realized average forecast errors over the past year. Higher realized squared forecast errors predict a larger uncertainty revision. This is different from the prediction of FIRE, according to which the two, on average, should be negatively correlated one by one, i.e., the resolution of forecast errors is equal to the reduction in uncertainty over the same period.

In summary, the empirical tests in this section use the uncertainty revision to show additional evidence for the deviation from FIRE in expectation formation. In particular, information rigidity of incorporating new information implies inefficiency of revisions in forecasts and a drop in uncertainty. The reduced-form results confirm this pattern. The next section compares a variety of candidate theories of expectation formation that may lead to such patterns.

3 Alternative Theories of Expectation Formation

3.1 Sticky expectations (SE)

The model of sticky expectations (Mankiw and Reis (2002), Carroll (2003)), regardless of its micro-foundation,²³ builds upon the assumption that agents do not update information instantaneously as they do in FIRE. One tractable assumption is that agents update their information with a homogeneous and time-independent probability, denoted by λ . Specifically, at any point of time t , each agent learns about the up-to-date realization of y_t with the probability of λ ;

²³For instance, Mankiw and Reis (2002) models SE as a result of individual attention choice subject to information cost, while Carroll (2003) models SE as a gradual diffusion of information among the population.

Table 2: Tests of Revision Efficiency Using Mean Revision and Uncertainty

	SPF CPI			SPF PCE			SCE				
	Mean revision	4q before	5q before	Mean revision	4q before	4-5 q before	Mean revision	24m before	25m before	26m before	
L4.InfExp_Var_rv	0.448*** (0.056)	0.456*** (0.058)	0.384*** (0.044)	0.395*** (0.044)	L24.InfExp_Var_rv	-0.679** (0.202)					
L5.InfExp_Var_rv	0.440*** (0.053)	0.406*** (0.042)	0.406*** (0.042)	0.406*** (0.042)	L25.InfExp_Var_rv	-0.740** (0.216)					
					L26.InfExp_Var_rv	-0.874*** (0.201)					
Constant	-0.091*** (0.000)	-0.049*** (0.008)	-0.048*** (0.005)	-0.079*** (0.000)	InfExp_FE2	0.510*** (0.101)	0.510*** (0.101)	0.568*** (0.092)	0.600*** (0.091)	0.600*** (0.091)	
					Constant	0.114 (0.180)	-0.051*** (0.007)	-0.050*** (0.003)	-1.121*** (0.247)	-1.237*** (0.204)	-1.366*** (0.237)
R2	0.047	0.196	0.248	0.054	R2	0.000	0.204	0.215	0.372	0.402	0.448
N	1529	1157	1021	1439	N	74	960	1091	50	49	48
Time FE	Yes	No	Yes	Yes	Time FE	No	Yes	Yes	No	No	No

Standard errors are clustered by date. *** p<0.001, ** p<0.01, and * p<0.05.

otherwise, they form the expectation based on the most recent up-to-date realization of $y_{t-\tau}$, where τ is the elapsed time since the last update.

The average forecast under SE is a weighted average of up-to-date rational expectations and lagged average expectations.²⁴ It follows that the average forecast errors are serially correlated, as described in Equation 7.

$$\overline{FE}_{t+h|t}^{se} = (1 - \lambda)\rho\overline{FE}_{t+h-1|t-1}^{se} + \lambda FE_{t+h|t}^* \quad (7)$$

The unconditional variance of the h -period-ahead forecast error (or its square) is proportional to that of the FIRE. It is also easy to confirm the former is always smaller than the latter as long as there is stickiness ($\lambda < 1$). Intuitively speaking, stickiness in expectations implies attenuated responses to inflation shocks compared to FIRE, hence a lower variation in forecast errors across time.

$$\overline{FE}_{\bullet+h|\bullet}^{se2} = \frac{\lambda^2}{1 - (1 - \lambda)^2\rho^2} FE_{\bullet+h|\bullet}^{*2} \leq FE_{\bullet+h|\bullet}^{*2} \quad (8)$$

Like average forecasts, average uncertainty at time t is also a weighted average of uncertainty to agents whose last updates took place in different periods of the past: $t - \tau \quad \forall \quad \tau = 0, 1, \dots, \infty$ (Equation 9). Since at any point in time there are agents who have not updated the recent realization of the shocks, and thus perceive higher uncertainty, the population uncertainty is unambiguously higher than in the case of FIRE. (See the Appendix for detailed derivations.)

$$\begin{aligned} \overline{\text{Var}}_{t+h|t}^{se} &= \sum_{\tau=0}^{+\infty} \underbrace{\lambda(1 - \lambda)^\tau}_{\text{fraction of non-updater until } t-\tau} \underbrace{\text{Var}_{t+h|t-\tau}^*}_{\text{uncertainty based on updating by } t-\tau} \\ &= \left(1 - \frac{\lambda\rho^{2h}}{1 - \rho^2 + \lambda\rho^2}\right) \frac{1}{1 - \rho^2} \sigma_\omega^2 \\ &\geq \overline{\text{Var}}_{t+h|t}^* \end{aligned} \quad (9)$$

For example, the average uncertainty regarding 1-period-ahead inflation ($h = 1$) is equal to $\frac{\sigma_\omega^2}{1 - (1 - \lambda)\rho^2}$, which collapses to σ_ω^2 when $\lambda = 1$ under FIRE and takes a larger value for any $0 < \lambda < 1$.

With respect to revision, the inefficiency of reducing uncertainty in SE takes the following

²⁴See Coibion and Gorodnichenko (2012) for detailed steps.

form at the aggregate level. Since not all agents incorporate the recently realized shocks, the revision in average uncertainty exhibits a serial correlation, as described in Equation 10. It is a weighted average of the resolution of uncertainty from the most recent shocks and its lagged counterpart.

$$\overline{\text{Var}}_{t+h|t+1}^{se} - \overline{\text{Var}}_{t+h|t}^{se} = (1 - \lambda)(\overline{\text{Var}}_{t+h|t}^{se} - \overline{\text{Var}}_{t+h|t-1}^{se}) - \lambda\rho^{2(h-1)}\sigma_\omega^2 \quad (10)$$

In particular, the second component is the information gain from the most recent realization of the shock, underweighted by $\lambda < 1$. The first component is the inefficiency sourced from the stickiness of updating. The higher the rigidity (lower λ), the smaller is the efficiency gain or uncertainty reduction compared to FIRE.

Lastly, SE also predicts non-zero disagreements and sluggish adjustment compared to FIRE. This is because of different lags in updating across populations.

In summary, SE predicts a higher average uncertainty and a lower expected forecast error square than their counterparts in FIRE, both of which should be identical to the conditional volatility of inflation under FIRE. In addition, SE predicts positive disagreements. These patterns are indeed observed in survey data, as reported in Table 1. Next, we move to other theories to see if such patterns are distinctive predictions by SE.

3.2 Noisy information (NI)

A class of models (Lucas (1972), Sims (2003), Woodford (2001), Maćkowiak and Wiederholt (2009), etc.), noisy information (NI hereafter) describes the expectation formation as a process of extracting the underlying variable y_t from a sequence of realized signals.

I reproduce the same signal structure as in Coibion and Gorodnichenko (2015) by assuming that an agent i observes two signals: $s_t^{pb} = y_t + \epsilon_t$, being a public signal common to all agents, and $s_{i,t}^{pr} = y_t + \xi_{i,t}$ being a private signal specific to the agent i . Both noises follow a normal distribution with zero mean and variances equal to σ_ϵ and σ_ξ , respectively.

Skipping the details of derivation, the average forecast error under NI is the following (Equation 11). Since private signals cancel out, on average, across agents, only a public signal, ϵ_t , directly enters the average forecast errors. Average population forecast error is a combination of lagged errors in nowcasting, partial responses to new public information, and forecast errors under FIRE.

$$\overline{FE}_{t+h|t}^{ni} = (1 - PH)\rho\overline{FE}_{t+h-1|t-1}^{ni} + \rho^h P_\epsilon \epsilon_t + FE_{t+h|t}^{*2} \quad (11)$$

P is the Kalman gain, a vector of size two that determines the degrees of reaction to signals. H is a vector of two ones, $[1, 1]'$. Notice, further, that the degree of reaction to new information P endogenously depends on, first, the prior nowcasting uncertainty, $\text{Var}_{i,t|t-1}^{ni}$, as of $t - 1$, and second, the noisiness of signals summarized by Σ^v .

$$P = [P_\epsilon, P_\xi] = \text{Var}_{i,t|t-1}^{ni} H' (H' \text{Var}_{i,t|t-1}^{ni} H + \Sigma^v)^{-1}$$

$$\Sigma^v = \begin{bmatrix} \sigma_\epsilon^2 & 0 \\ 0 & \sigma_\xi^2 \end{bmatrix} \quad (12)$$

Kalman filtering also updates the uncertainty recursively according to the following rule. The posterior uncertainty at time t is a linear function of prior uncertainty and noisiness of signals.

$$\begin{aligned} \text{Var}_{i,t|t}^{ni} &= (1 - PH)\text{Var}_{i,t|t-1}^{ni} \\ &= (1 - PH)(\rho^2 \text{Var}_{i,t-1|t-1}^{ni} + \sigma_\omega^2) \end{aligned} \quad (13)$$

The average square (or the unconditional variance) of the forecast errors is unambiguously greater than $FE_{t+h|t}^{*2}$ in FIRE, as shown in Equation 14. The simple reason for this is that the forecast errors under NI always come from not only the realized shocks to inflation but also nowcasting noises.

$$\overline{FE}_{\bullet+h|\bullet}^{ni2} = \frac{\rho^{2h} P_\epsilon^2 \sigma_\epsilon^2 + FE_{\bullet+h|\bullet}^{*2}}{(PH)^2} \geq FE_{\bullet+h|\bullet}^{*2} \quad (14)$$

The unconditional nowcasting variance can be solved as the steady-state value in Equation 13. Note that the average uncertainty non-linearly depends on the noisiness of the two signals σ_ϵ^2 and σ_ξ^2 , as well as the volatility of inflation, σ_ω .

The h -period-ahead forecasting uncertainty comes from both nowcasting uncertainty and volatility of unrealized shocks in the future (Equation 15).

$$\text{Var}_{t+h|t}^{ni} = \rho^{2h} \text{Var}_{t|t}^{ni} + \sum_{s=1}^h \rho^{2(s-1)} \sigma_{\omega}^2 \geq \text{Var}_{t+h|t}^* \quad (15)$$

Equation 13 also directly gives the revision in uncertainty from time $t - 1$ to t . The newly arrived information, albeit noisy, still brings about information gains, thus leading to an unambiguous drop in uncertainty. But because the signal is not perfect, i.e., Σ^v 's diagonal is non-zero, there is inefficiency in reducing uncertainty compared to in FIRE.

$$\text{Var}_{t|t}^{ni} - \text{Var}_{t|t-1}^{ni} = -PH\text{Var}_{i,t|t-1} < 0 \quad (16)$$

As a result of the NI mechanism, the revision in h -period-ahead uncertainty from $t - 1$ to t only partially reacts to the resolution of uncertainty from the newly realized shock ω_t in the past period.

NI also predicts non-zero disagreement in the presence of private signals. The size of the disagreement depends on, but is not a strictly increasing function of, the noisiness of the private signals. This is so because if the noisiness of private signals σ_{ξ} is much larger, say close to infinity, than that of the public signal σ_{ϵ} , agents will optimally not at all react to private signals. In this scenario, the disagreement will no longer increase with σ_{ξ} .

In summary, both the variance of forecast error and uncertainty are greater than conditional volatility of the inflation because of the presence of noisy signals. Meanwhile, their relative sizes are ambiguous. Disagreement is positive as long as there is dispersed information in the form of private signals and they are not too noisy. In addition, all three moments contain parametric restrictions about the noisiness of public and private signals. It is possible that, under a range of parameter values, NI generates moments that are consistent with the observed data from the survey. We leave this task for the structural estimation in Section 4.

3.3 Diagnostic expectations (DE)

Different from the previous two theories featuring informational rigidity, diagnostic expectation (Bordalo et al. (2018)) introduces an extrapolation mechanism in expectation formation that results in overreactions to the news (Bordalo et al. (2020)). Both SE and NI deviate from FIRE in terms of the information set available to the agents (the ‘‘FI’’ assumption), while DE deviates from FIRE in terms of the processing of an otherwise fully updated information set (the ‘‘RE’’ assumption).

Skipping over its micro-foundation, Equation 17 captures the essence of the DE mechanism. Each individual i 's h -period-ahead forecast consists of two components. The first component can be considered as a rational forecast based on the fully updated y_t . The second component corresponds to the potential overreaction to the unexpected surprises from $t - 1$ to t . The degree of overreaction is governed by the parameter θ . The premise of the DE model is that $\theta > 0$, which captures the fact that the agent overly responds to the realized forecast errors. The model collapses to the FIRE when $\theta = 0$. Meanwhile, as argued in [Bordalo et al. \(2020\)](#), a negative θ is not incompatible with an interpretation of underreaction if we treat DE as a more generalized model of expectation formation.

$$\bar{y}_{i,t+h|t}^{de} = \rho^h y_t + \theta_i (\rho^h y_t - \bar{y}_{i,t+h|t-1}^{de}) \quad (17)$$

There is no room for disagreement with a homogeneous degree of overreaction. To account for the possibility of a positive disagreement, I assume θ to be different across different agents. Therefore, I add the subscript i to the parameter. Since agents are equally informed about the realizations of the variable, the only room for disagreement to be positive is heterogeneous degrees of overreaction. To capture this, I assume θ_i to follow a normal distribution across the population, $N(\hat{\theta}, \sigma_\theta^2)$. So the DE model has two parameters, and disagreement increases with the dispersion of overreaction, σ_θ .

The average forecast takes exactly the same form, with the individual-specific θ_i replaced by the population average $\hat{\theta}$. Therefore, we focus on average forecast errors directly. The average forecast error under DE evolves as the following (see the Appendix for derivations):

$$\overline{FE}_{t+h|t}^{de} - FE_{t+h|t}^* = -\hat{\theta}\rho(\overline{FE}_{t+h-1|t-1}^{de} - FE_{t+h-1|t-1}^*) + \rho^h \hat{\theta} \omega_t \quad (18)$$

Intuitively, moving from $t - 1$ to t , the h -period-ahead FE exceeds that of FIRE by two components, one is the mean-reverting change from the overreacted forecast errors from $t - 1$; the other is the overreaction to the newly realized shock. Combined, they can be interpreted as the overreaction to the surprise compared to the expectation formed at $t - 1$.

The unconditional variance of h -period-ahead forecast errors is the most straightforward in the special case $h = 1$, which is equal to the following. It is smaller than the variance of forecast error in the FIRE benchmark and the conditional volatility of the inflation, σ_ω^2 .

$$\overline{FE}_{\bullet+1|\bullet}^{de2} = \frac{\sigma_\omega^2}{1 + \hat{\theta}^2 \rho^2} \quad (19)$$

Finally, as to the uncertainty, since the mechanism of extrapolation in DE does not change the agent’s perceived distribution of future shocks, the benchmark DE model forecasts uncertainty to remain the same as in FIRE.

$$\overline{Var}_{t+h|t}^{de} = \overline{Var}_{t+h|t}^* \quad (20)$$

In summary, under DE, the ex-ante uncertainty, which is identical to the conditional volatility of inflation, is greater than the square of ex-post forecast error. The variability of average forecast errors are smaller than that in FIRE because of its mean-reversion.

3.4 Diagnostic expectations (DE) augmented with noisy information (DENI)

[Bordalo et al. \(2020\)](#) embeds heterogeneous information in a standard DE model. Their motivations are to generate cross-sectional disagreement in forecasts and jointly produce the underreaction in consensus forecasts and overreaction at individual levels. The framework is essentially a hybrid of the NI and DE.²⁵ It maintains the assumption that agents overreact to new information at individual levels, but the information is no longer the real-time realization of the variable y_t itself but noisy signals of it, which we denote as $s_{i,t}$. I assume a more general signal structure than in [Bordalo et al. \(2020\)](#) to include both public and private signals, as assumed in Section 3.2.

Then the h -period-ahead forecast takes a recursive form as follows.

$$y_{i,t+h|t}^{deni} = y_{i,t+h|t-1}^{deni} + (1 + \theta)P^{deni}H(\rho^h s_{i,t} - y_{i,t+h|t-1}^{deni}) \quad (21)$$

$P^{deni} = [P_\epsilon, P_\xi]$ is the vector of Kalman gain as a function of nowcasting uncertainty $\text{Var}_{t|t}^{deni}$ and noisiness of signals $\sigma_\epsilon, \sigma_\xi$. With $\theta = 0$, Equation 21 becomes identical to that in NI with one private signal, where the forecast is a Kalman-gain-weighted average of new information and the prior information. Any $\theta > 0$ implies overreaction beyond the optimal Kalman gain.

²⁵It is called “the Diagnostic Kalman Filter” in [Bordalo et al. \(2020\)](#).

With such a mechanism, the average FE evolves as follows (see the Appendix for derivations). At any t , the deviation of FE from FIRE consists of its mean-reversion from $t - 1$, the overreaction to the Kalman-gain-weighted inflation shock, and the public noise.

$$\begin{aligned} \overline{FE}_{t+h|t}^{deni} - \overline{FE}_{t+h|t}^* &= -\theta\rho(\overline{FE}_{t+h-1|t-1}^{deni} - FE_{t+h-1|t-1}^*) \\ &\quad + \rho^h((1+\theta)P_\epsilon - 1)\omega_t + \rho^h(1+\theta)P_\epsilon\epsilon_t \end{aligned} \quad (22)$$

Taking $h = 1$ as an example, the unconditional variance of FE is equal to the following:

$$\overline{FE}_{\bullet+1|\bullet}^{deni2} = \frac{\sigma_\omega^2 + \rho^2(1+\theta)^2(1-P_\epsilon)^2\sigma_\omega^2 + \rho^2(1+\theta)^2P_\epsilon^2\sigma_\epsilon^2}{1 + \theta^2\rho^2} \quad (23)$$

The above equation collapses to that in FIRE when the public information is perfectly precise ($P_\epsilon = 1, \sigma_\epsilon = 0$) and there is no overreaction ($\theta = 0$). It collapses to DE when θ remains positive and $P_\epsilon = 1, \sigma_\epsilon = 0$. Compared to the DE model, in which the variation of average forecast error is attenuated, the NI model introduces a counteracting force that makes variation of FE possibly bigger than FIRE due to the existence of the noisy signals. Therefore, the relative size between the variance of FE in DE and FIRE is ambiguous.

Forecast uncertainty under DENI is identical to that in NI, because only the NI mechanism affects the behaviors of uncertainty.

$$\overline{Var}_{t+h|t}^{deni} = \overline{Var}_{t+h|t}^{ni} \quad (24)$$

Finally, the DENI model also predicts positive Disg for any noisy private information: $\sigma_\xi > 0$.

3.5 Comparing theories

Table 3 summarizes the predictions by different theories.

In addition, for each theory not only forecast error but also higher moments, disagreement, and uncertainty contain restrictions to identify the model parameters within each theory. I will use these moment conditions to estimate each theory in Section 4.

Table 3: Model Predictions

	FIRE	SE	NI	DE	DENI
Fact 1: Higher uncertainty than forecast error variance	No	Yes	?	Yes	?
Fact 2: Uncertainty greater than conditional volatility	No	Yes	Yes	No	Yes
Fact 3: Inefficient uncertainty revisions	No	Yes	Yes	No	Yes
Fact 4: Positive disagreement	No	Yes	Yes	Yes	Yes

4 Model Estimation and Sensitivity Analysis

4.1 SMM Estimation

Although reduced-form tests based on survey data often provide additional evidence for rejecting the null hypothesis of FIRE, there are two limitations with such an approach in terms of identifying differences among non-FIRE theories. First, the coefficient estimates from the reduced-form regression cannot always be mapped into a structural parameter of the particular model, especially when reported expectations and forecast horizons are at different time frequencies. Second, even if it does so, the tests fall short of simultaneously using all the restrictions across moments implied by a particular non-FIRE theory, as discussed in great detail in Section 3. In this section, I undertake a structural estimation that jointly accounts for cross-moment restrictions.

Since many of the moment conditions cannot be easily derived as a closed-form function of parameters, I adopt the simulated method of moment (SMM). In a nutshell, the estimation chooses the best set of model parameters by minimizing the weighted distances between the data moments and the model-simulated moments. For a given process of inflation, and a particular theory of expectation formation, the vector of the parameters estimates is defined as the minimizer of the following objective function:

$$\hat{\Omega}^o = \underset{\{\Omega^o \in \Gamma^o\}}{\operatorname{argmin}} (M_{\text{data}} - F^o(\Omega^o, H))' W (M_{\text{data}} - F^o(\Omega^o, H))$$

where Ω^o stands for the parameters of the particular pair of theories of expectation and inflation process, i.e., $o \in \{se, ni, de, deni\} \times \{ar, sv\}$. Γ^o represents the corresponding parameter space respecting the model-specific restrictions. M_{data} is a vector of the unconditional moments that is computed from data on expectations and inflation. F^o is the simulated model moments under the theory pair o . W is the weighting matrix used for the SMM estimation.

I report estimation results using the two-step feasible SMM approach, in which the inverse of the variance-covariance matrix from the first-step estimation using identity matrix is used as the W in the second step, which has been shown to give asymptotically efficient estimates of the model parameters.

Crucially, notice that the model-implied moments F^o are not just a function of model parameters Ω^o , but also a function of the corresponding information set available to the forecasters. I use H to represent the historical realizations of the variables in the agents' information set that are used as the inputs for forecasts.

It is also important to mimic the information set that was truly available to the agents at each point in time in history.²⁶ Therefore, I use the real-time data on historical inflation that was publicly available at each point in time instead of the historical data series released later, since the latter usually experiences many rounds of revisions over time. I obtained the data from the Real-Time Data Research Center hosted by the Philadelphia Fed.²⁷

The estimation is also specific to the choice of moments to be matched. I focus on the unconditional population moments across time in the sample as listed in Table 1. One of the rationales for focusing on population moments is to avoid measurement errors at the individual level, as shown in Juodis and Kučinskas (2023).

The model-implied moments also implicitly depend on the parameters of the inflation process for a given model. This point is illustrated well in Bordalo et al. (2020). For instance, the observed overreaction in DE is lower for an AR(1) process with higher persistence. In recovering the model parameters associated with expectation formation, it is important to take into account the information contained in expectation data regarding the process of inflation itself. To account for this, I undertake both two-step and joint estimation. The former is to first externally estimate the inflation process and then estimate expectation formation, separately, treating the inflation parameter as the *true* data-generating process of inflation. The latter refers to jointly estimating parameters of inflation and expectation.

These alternative specifications of the estimation also serve as a model sensitivity analysis with respect to the following criteria: (1) different choices of moments; (2) AR(1) and SV for the process of inflation; (3) two-step and joint estimation; (4) for both professionals and households. A reasonable theory of expectation formation ought to be relatively robust to these criteria. I discuss the findings in greater detail along these four dimensions next.

²⁶For the importance of using real-time data to study survey forecasts, see Faust and Wright (2008), Faust and Wright (2009), and so on.

²⁷<https://www.philadelphiafed.org/research-and-data/real-time-center/real-time-data/>.

4.2 Moments matching and parameter estimates

Table 4 presents the SMM estimates for professionals, as a benchmark. For each theory, I estimate the theory in two steps and jointly using expectations and inflation moments. Different rows within each panel report the estimates depending on various choices of moments used for estimation: forecast errors only (FE), forecast error and disagreement (FE+Disg), and the two plus uncertainty (FE+Disg+Var).

Table 4: SMM Estimates of Different Models: Professionals

SE								
Moments Used	Two-Step Estimate			Joint Estimate				
	$\hat{\lambda}$	ρ	σ_ω	$\hat{\lambda}$	ρ	σ_ω		
FE	0.36	0.99	0.23	0.18	0.97	0.11		
FE+Disg	0.28	0.99	0.23	0.22	0.95	0.14		
FE+Disg+Var	0.26	0.99	0.23	0.32	0.9	0.22		
NI								
Moments Used	Two-Step Estimate				Joint Estimate			
	$\hat{\sigma}_\epsilon$	$\hat{\sigma}_\xi$	ρ	σ_ω	$\hat{\sigma}_\epsilon$	$\hat{\sigma}_\xi$	ρ	σ_ω
FE	0	0.87	0.99	0.23	0	0.15	0.97	0.11
FE+Disg	1.5	2.26	0.99	0.23	1.48	2.33	0.97	0.11
FE+Disg+Var	2.64	3	0.99	0.23	3	3	0.94	0.16
DE								
Moments Used	Two-Step Estimate				Joint Estimate			
	$\hat{\theta}$	σ_θ	ρ	σ_ω	$\hat{\theta}$	σ_θ	ρ	σ_ω
FE	0.64	0.58	0.99	0.23	0.81	1.68	0.97	0.11
FE+Disg	0.27	2.2	0.99	0.23	0.38	2.1	0.9	0.2
FE+Disg+Var	0.42	2.1	0.99	0.23	0.33	2.1	0.9	0.23
DENI								
Moments Used	Two-Step Estimate				Joint Estimate			
	$\hat{\theta}$	$\hat{\sigma}_\xi$	ρ	σ_ω	$\hat{\theta}$	$\hat{\sigma}_\xi$	ρ	σ_ω
FE	0.76	0	0.99	0.23	0.82	0	0.97	0.11
FE+Disg	0.85	0.14	0.99	0.23	N/A	N/A	N/A	N/A
FE+Disg+Var	0.85	0.16	0.99	0.23	N/A	N/A	N/A	N/A

4.2.1 Cross-moment consistency of each theory

Among the four models under consideration, SE and DENI outperform the others in terms of the within-model robustness against targeted moments, as shown in the estimation of professional forecasts in Table 4.

For SE, the estimated quarterly updating rate λ is between 0.22 and 0.36 across different combinations of moments. The estimate is 0.36 when only FE moments are targeted. It

is smaller, 0.28 and 0.26, respectively, when disagreement and uncertainty are sequentially included. These imply that the information rigidity is revealed through both low- and high-order moments.

For DENI, the implied overreaction parameter θ is in the range of 0.76–0.85, suggesting the existence of overreaction mechanisms in the population. Meanwhile, the noisiness of the private signals σ_ξ is around 0.14–0.16 in percentage points.

In contrast, the NI estimates of σ_ϵ and σ_ξ are rather volatile across targeted moments. When only FE is targeted, the estimation points to a highly precise public signal and mildly noisy private signals. But when disagreement and uncertainty are included, the estimated noisiness of both signals significantly increases. They are often so large in magnitude that they hit the externally set upper bound of 3. These are highly noisy signals compared to the conditional standard deviation of inflation shocks $\sigma_\omega = 0.22$. Although qualitatively NI mechanisms accommodate patterns of information rigidity similar to SE, quantitatively the required noisiness of the signals is less sensible to interpret.

DE estimates are also sensitive to moment restrictions, although all the estimates confirm the existence of a positive mass of overreacting agents. $\hat{\theta}$ is estimated to range from 0.27 to 0.81 depending on the estimation specification. Using disagreement helps identify the population dispersion in the degree of overreaction σ_θ , which is estimated to be 2.2. This suggests a significant amount of heterogeneity in the degree of reaction to the news through the lens of DE.

4.2.2 Interactions between expectation formation and inflation process

In all four models, the estimated parameters vary when one jointly estimates expectation and inflation process parameters. With the benchmark AR(1) process, both the persistence of the shock to inflation ρ and the overall volatility of the inflation shock σ_ω determine the value of the corresponding moments under a particular model of expectation formation.²⁸ The differences between two-step estimation and joint estimation reveal such interdependence.

An alternative interpretation of the joint estimates is that they reveal possibly subjective models of inflation as perceived by the forecasters, which may be different from the one estimated from historical data retrospectively.²⁹ The joint estimation results seem to support such an interpretation. The estimates of SE, NI, and DE all produce very similar parameters of ex-

²⁸Afrouzi et al. (2023) emphasizes this point in the context of DE. They further address such non-identification challenges in lab experiments by exogenously altering the persistence of the data-generating process.

²⁹For instance, Jain (2019) uses professional forecasts of inflation to infer the perceived persistence of inflation. Macaulay and Moberly (2022) directly shows survey evidence for the heterogeneity of households in their perceived persistence of inflation.

pectation formation, yet rather different inflation persistence and volatility. The survey-implied persistence of inflation and conditional volatility are both smaller than those estimated solely based on inflation data. This implies that in addition to information rigidity and overreaction mechanisms in the canonical versions of these models, allowing for the possibility of a subjective model is necessary to fit the joint dynamics of inflation and expectations better.

4.2.3 Professionals versus households

Table 5: SMM Estimates of Different Models: Households

SE				
Moments Used	Two-Step Estimate			
	$\hat{\lambda}$	ρ	σ_ω	
FE	0.36	0.98	0.45	
FE+Disg	0.36	0.98	0.45	
FE+Disg+Var	0.36	0.98	0.45	
NI				
Moments Used	Two-Step Estimate			
	$\hat{\sigma}_\epsilon$	$\hat{\sigma}_\xi$	ρ	σ_ω
FE	0	1	0.98	0.45
FE+Disg	3	1.18	0.98	0.45
FE+Disg+Var	2.06	3	0.98	0.45
DE				
Moments Used	Two-Step Estimate			
	$\hat{\theta}$	σ_θ	ρ	σ_ω
FE	0.49	0.5	0.98	0.45
FE+Disg	1.91	5	0.98	0.45
FE+Disg+Var	1.03	5	0.98	0.45
DENI				
Moments Used	Two-Step Estimate			
	$\hat{\theta}$	$\hat{\sigma}_\xi$	ρ	σ_ω
FE	N/A	N/A	0.98	0.45
FE+Disg	-0.54	3	0.98	0.45
FE+Disg+Var	-0.35	2.43	0.98	0.45

Table 5 reports the estimates for households. The updating rate in SE is 0.36, implying around a one-third chance of updating per month. This is a slightly lower degree of stickiness than professionals. It is well documented in the literature that household expectations tend to be more inattentive to economic news than professionals.³⁰ But the SE results of our estimates show that the major differences are not simply due to the differences in updating rates of information.

³⁰See [Cornand and Hubert \(2022\)](#) for a recent discussion on this point.

NI estimates of households, when all moments are targeted, reveal extremely noisy signals. The public signals need to have a noisiness of 2 to 3 percentage points while private signals need to be between 1 to 3 percentage points.

DE estimates of households suggest a consistently positive and greater degree of overreaction, together with an extremely large dispersion of 5.0, which gives also a positive mass of underreactive agents.

Compared to the previous three models, DENI estimates of the household are significantly different from that of professionals. The average degree of overreaction $\hat{\theta}$ becomes negative, taking the value of -0.35 to -0.54, implying underreaction on average. In addition, the noisiness of private signals is estimated to still be extremely large, taking values of 2.43 to 3.

5 Inflation with Stochastic Volatility (SV)

This section considers an alternative data-generating process of inflation using the Unobservable Component/Stochastic Volatility (SV) model proposed by [Stock and Watson \(2007\)](#), which is arguably a more realistic stochastic process of inflation given the presence of macroeconomic shocks of varying persistence.

The extension to stochastic volatility achieves two objectives. First, a basic inflation process with constant volatility does not account for the observed time-varying pattern of forecast uncertainty nor its correlation with other moments such as disagreement and forecast error size, as shown in [Section 2.4](#). This is particularly relevant to account for the rapid rise of disagreement, forecast error, and uncertainty over the inflationary period since 2020.

Second, allowing for SV in the inflation process serves as a robustness test of various theories of expectation formation because it captures the sensitivity of these theories to the assumed underlying generating process of inflation. This extension provides a more comprehensive analysis of the relationship between inflation dynamics and expectation formation.

In particular, SV assumes that inflation consists of a permanent ζ and transitory component η .³¹ Time variations in the relative size of the volatility of two components σ_ζ^2 and σ_η^2 drive time variations of the persistence of inflation shocks. The logged volatility of the two components themselves follows a random walk subject to shocks μ_ζ and μ_η .

³¹For such a multi-component formulation of the inflation process, see [Kohlhas and Walther \(2021\)](#); [Farmer et al. \(2021\)](#).

$$\begin{aligned}
y_t &= \zeta_t + \eta_t, \quad \text{where } \eta_t = \sigma_{\eta,t} \nu_{\eta,t} \\
\zeta_t &= \zeta_{t-1} + z_t, \quad \text{where } z_t = \sigma_{z,t} \nu_{z,t} \\
\log \sigma_{\eta,t}^2 &= \log \sigma_{\eta,t-1}^2 + \mu_{\eta,t} \\
\log \sigma_{z,t}^2 &= \log \sigma_{z,t-1}^2 + \mu_{z,t}
\end{aligned} \tag{25}$$

The shocks to the level of the two components η_t and z_t , and those to their volatility, $\mu_{\eta,t}$ and $\mu_{z,t}$, are drawn from the following normal distributions, respectively. The only parameter of the model is γ , which determines the smoothness of the time-varying volatility.

$$\begin{aligned}
\nu_t &= [\nu_{\eta,t}, \nu_{z,t}] \sim N(0, I) \\
\mu_t &= [\mu_{\eta,t}, \mu_{z,t}]' \sim N(0, \gamma I)
\end{aligned} \tag{26}$$

I reproduce the estimates of [Stock and Watson \(2007\)](#) using the Markov Chain Monte Carlo algorithm covering the extended period till March 2023. The estimated time-varying permanent and transitory volatility of both core CPI and headline CPI is shown in [Figure 3](#). The figure depicts intuitive patterns of higher permanent volatility of inflation around the 2008 Great Recession and the COVID era since 2020.

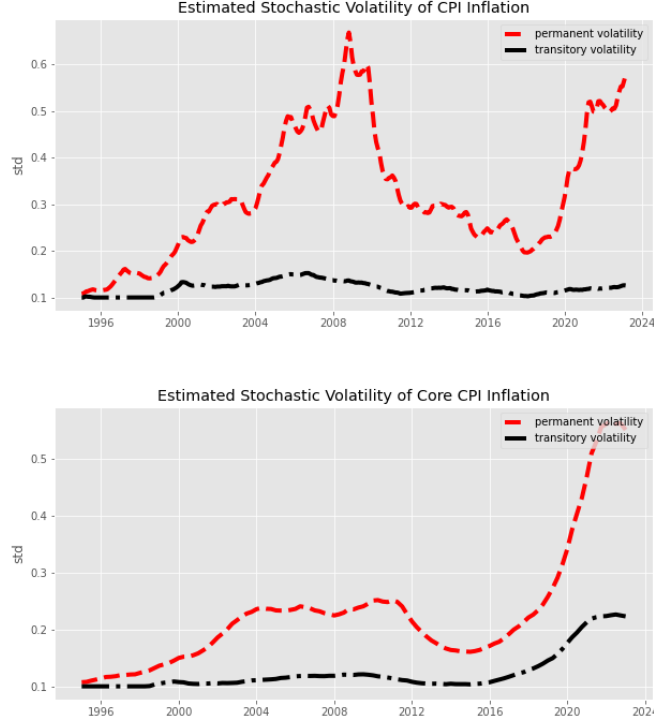
5.1 Model predictions under SV

The information set necessary for forecasting is different in SV from that in an AR(1) process. Consider first the benchmark case of FIRE. At the time t , the FIRE agent sees the most recent and past realization of all stochastic variables as of t , including y_t , ζ_t , η_t , $\sigma_{\eta,t}$, and $\sigma_{z,t}$. Using the superscript **sv* to denote the FIRE benchmark prediction under stochastic volatility, and suppressing the individual subscript i (because there is no disagreement in FIRE), the h -period-ahead forecast of inflation is equal to the contemporaneous realization of the permanent component, $\epsilon_t \equiv \zeta_t$.

$$\bar{y}_{t+h|t}^{*sv} = \zeta_t \tag{27}$$

Under FIRE, forecast error is simply the cumulative sum of unrealized permanent and transitory shocks from t to $t+h$, which is equal to the following, and disagreement is zero across agents in FIRE.

Figure 3: Stochastic Volatility of Inflation



Note: This figure plots the estimated stochastic volatility of permanent and transitory components of monthly headline CPI inflation (top) and quarterly core CPI inflation (bottom) using the same approach as in [Stock and Watson \(2007\)](#).

$$\overline{FE}_{t+h|t}^{*sv} = - \sum_{s=1}^h (\eta_{t+s} + z_{t+s}) \quad (28)$$

The h -step-ahead conditional variance, or the forecast uncertainty, is time-varying since the volatility is stochastic now. It is essentially the conditional expectation of the cumulative sum of future volatility given the current realizations of the component-specific volatility at t .

$$\begin{aligned} \overline{Var}_{t+h|t}^{*sv} &= \sum_{s=1}^h E_t(\sigma_{\eta,t+s}^2) + E_t(\sigma_{z,t+s}^2) \\ &= \sigma_{\eta,t}^2 \sum_{s=1}^h \exp^{-0.5s\gamma} + \sigma_{z,t}^2 \exp^{-0.5h\gamma} \end{aligned} \quad (29)$$

SESV Under the sticky expectation (SE), an agent whose most recent up-to-date update happened in $t - \tau$ has only seen the realizations of y , ζ , η , σ_η , and σ_z till $t - \tau$. The average forecast is hence the weighted average of all past realizations of the permanent component up to t .

$$y_{t+h|t-\tau}^{sesv} = \sum_{\tau=0}^{\infty} \lambda(1-\lambda)^\tau \zeta_{t-\tau} \quad (30)$$

The distribution of lagged updating is also reflected in the average forecast uncertainty. The population average uncertainty is a weighted average of FIRE uncertainty at $t, t-1 \dots t-\tau \dots t-\infty$ (Equation 31). The key difference in SV from AR(1) is that the average uncertainty exhibits a positive serial correlation under SV. Expectations being sticky further increase the positive serial correlation compared to that in FIRE due to the lag in updating the shocks to the volatility. The predictions regarding both forecast errors and disagreements under SV are the same as under the AR(1) model.

$$Var_{t+h|t}^{sesv} = \sum_{\tau=0}^{\infty} \lambda(1-\lambda)^\tau Var_{t+h|t-\tau}^{*sv} \quad (31)$$

NISV Under noisy information (NI), in order to forecast future y the agent at time t needs to form her best nowcast of the permanent component ζ_t , denoted as $\bar{\zeta}_{t|t}$, using noisy signals and Kalman filtering. We assume again that the noisy signals of ζ_t consist of a public signal s_t^{pb} and a private signal $s_{i,t}^{pr}$ containing noises around the true realization of ζ_t . Following a long tradition of modeling the signaling-extraction problem in this two-component context, we further assume the public signal $s_t^{pb} = y_t$, meaning the inflation realization itself is the public signal of the permanent component. Accordingly, the transitory shock η_t is equivalent to the realized noise of the public signal ϵ_t in the benchmark NI model with AR(1) process.

Then the average forecast is a Kalman-gain-weighted average of prior belief and new information:

$$y_{t+h|t}^{nisv} = \bar{\zeta}_{t|t} = (1 - P_t^{sv} H) y_{t+h-1|t-1} + P_t^{sv} H \zeta_t + P_{\eta,t}^{sv} \eta_t \quad (32)$$

In the above equation, Kalman gain $P_t^{sv} = [P_{\eta,t}^{sv}, P_{\xi,t}^{sv}]$ is a function of forecasting uncertainty $Var_{t|t-1}^{svni}$, the constant noisiness of private signal σ_ξ and that of public signal $\sigma_{\eta,t}$, which is also the time-varying volatility of the transitory component of the inflation.

What is different under time-varying volatility is that there is no steady-state Kalman gain and uncertainty that are independent of time because the underlying volatility of the variable is time-varying. This also implies that the rigidity induced by the noisiness of information is state-dependent. In each period, agents in the economy will update their forecasts based on the realized volatility. In periods with high (low) fundamental volatility, the Kalman gain from noisy signals is larger (smaller), thus the agents will be more (less) responsive to the new information. There is no such state-dependence of rigidity in the canonical SE.

The mechanisms of DE and DENI exactly mimic that under AR(1) except that the average volatility is time-varying now.

5.2 The role of stochastic volatility

Tables 6 and 7 report the estimates under SV, respectively, for the low-inflation pre-pandemic period and the extended sample covering the high-inflation era between 2020–2023. I juxtapose the two episodes to explore possible state-dependence of expectation formation, especially given the rapid rise in inflation in the post-2020 era.

The major finding from the estimates is that SV process significantly improves the within-model consistency across targeted moments for both types of agents and both sample periods. This is probably not surprising given that the two-component formulation proves to be a more realistic foundation of the inflation dynamics, as the previous literature established.

The improvement in model consistency is the most obvious in NI estimates of professionals in which the benchmark estimates under the AR(1) process produce unrealistically imprecise signals. With SV, the estimated noisiness of private signals falls into a more reasonable range of values, i.e., 0.21–0.68. It is more reasonable to assume that forecasters imperfectly observe the permanent component instead of inflation itself.

Despite this improvement in the cross-moment consistency for the low-inflation sample, NI's estimates remain extremely large once the sample includes the recent inflation episodes. In addition, the estimation of NI for households fails to converge in all specifications. This implies that the model has a rather poor fit to household expectations even if a more realistic inflation process is assumed.

Among all the theories, SE gives the closest parameter estimates to that of the benchmark AR(1). The updating rate is estimated to remain in the range of 0.2–0.36 for both households and professionals. This suggests that SE has very good consistency against the assumed inflation process in capturing the overall patterns of the survey expectations. This is consistent with the preliminary diagnosis simply based on the empirical rankings of moments that are the most consistent with SE predictions.

Table 6: SMM Estimates of Different Models under Stochastic Volatility: Professionals

Before March 2020		Till March 2023		
SE				
Moments Used	Two-Step Estimate	Two-Step Estimate		
	$\hat{\lambda}$	$\hat{\lambda}$		
FE	0.2	0.3		
FE+Disg	0.25	0.36		
FE+Disg+Var	0.36	0.36		
NI				
Moments Used	Two-Step Estimate	Two-Step Estimate		
	$\hat{\sigma}_{pb}$	$\hat{\sigma}_{pr}$	$\hat{\sigma}_{pb}$	$\hat{\sigma}_{pr}$
FE	0.68	0.24	2.3	3
FE+Disg	0.67	0.24	2.3	3
FE+Disg+Var	0.64	0.21	2.3	3
DE				
Moments Used	Two-Step Estimate	Two-Step Estimate		
	$\hat{\theta}$	σ_{θ}	$\hat{\theta}$	σ_{θ}
FE	-0.03	0.54	0.31	0.41
FE+Disg	-0.03	0.16	0.28	0.19
FE+Disg+Var	-0.04	0.16	0.31	0.19
DENI				
Moments Used	Two-Step Estimate	Two-Step Estimate		
	$\hat{\theta}$	$\hat{\sigma}_{pr}$	$\hat{\theta}$	$\hat{\sigma}_{pr}$
FE	0.64	0.47	-0.25	0.93
FE+Disg	0.82	0.26	-0.26	0.93
FE+Disg+Var	0.82	0.24	-0.26	0.93

Table 7: SMM Estimates of Different Models under Stochastic Volatility: Households

Before March 2020		Till March 2023		
SE				
Moments Used	Two-Step Estimate	Two-Step Estimate		
	$\hat{\lambda}$	$\hat{\lambda}$		
FE	0.27	0.36		
FE+Disg	0.2	0.27		
FE+Disg+Var	0.26	0.26		
NI				
Moments Used	Two-Step Estimate	Two-Step Estimate		
	$\hat{\sigma}_\epsilon$	$\hat{\sigma}_\xi$	$\hat{\sigma}_\epsilon$	$\hat{\sigma}_\xi$
FE	N/A	N/A	N/A	N/A
FE+Disg	N/A	N/A	N/A	N/A
FE+Disg+Var	N/A	N/A	N/A	N/A
DE				
Moments Used	Two-Step Estimate	Two-Step Estimate		
	$\hat{\theta}$	σ_θ	$\hat{\theta}$	σ_θ
FE	-0.09	0.58	-0.07	0.57
FE+Disg	0.29	0.57	0.47	1.07
FE+Disg+Var	0.29	0.57	0.28	1.07
DENI				
Moments Used	Two-Step Estimate	Two-Step Estimate		
	$\hat{\theta}$	$\hat{\sigma}_\xi$	$\hat{\theta}$	$\hat{\sigma}_\xi$
FE	-0.48	0.64	0.43	0.26
FE+Disg	-0.48	0.64	0.43	0.26
FE+Disg+Var	-0.48	0.64	0.43	0.26

Estimates of DENI under SV also point to a similar degree of overreaction of professionals ($\hat{\theta}$ around 0.6–0.82) and underreaction of households ($\hat{\theta}$ around -0.48) as in the benchmark AR(1) estimates. The estimated dispersion of private information also remains similar to SV assumptions.

5.3 Expectation formation when inflation is high

The benchmark estimation is based on the low-inflation sample period before 2020 where our assumption of the stationary AR(1) process of inflation is a reasonable one. But SV formulation naturally fits better the dynamics of inflation once we want to examine if the estimates change when the sample covers an episode of high inflation and volatility, as shown in Figure 3.

For most models, the parameter estimates change significantly between two sample periods. Such changes do not necessarily invalidate the model mechanisms but instead reflect possible state-dependence of expectation formation. The difference in estimates between the two sample periods does suggest that both professionals and households have altered their responsiveness to the inflation news.

In particular, for professionals, the SE estimates imply an on-average higher updating rate $\lambda = 0.36$. Households' updating rate is estimated to be higher than in the pre-2020 sample. Both exhibit less information rigidity in the form of SE. The estimates of the DE model corroborate this pattern. Both professionals and households exhibit a higher degree of overreaction. $\hat{\theta}$ changes from -0.03 to 0.3 for professionals and from 0.29 to 0.47 for households. This echoes the findings of [Coibion and Gorodnichenko \(2015\)](#) that the information rigidity is state-dependent. [Goldstein \(2023\)](#); [Pfäuti \(2023\)](#) found that inflation expectations exhibit less rigidity when inflation is elevated.

Different from the finding that information rigidity is lessened in the high inflation episode for both types of agents, DENI estimates depict a more divergent pattern between the two types. Professionals have on average turned to underreaction with more dispersed information: $\hat{\theta}$ changes from 0.82 to -0.26 , and σ_{ξ} changes from 0.24 to 0.93 . In contrast, households have turned to overreaction with less dispersed information: $\hat{\theta}$ changes from -0.48 to 0.43 , and σ_{ξ} changes from 0.64 to 0.26 .

5.4 The final assessment of different models

To summarize, Table 8 reports my evaluation of the four theories under consideration based on four sensitivity criteria laid out in the previous section. According to this evaluation, SE seems to capture the average behavior of expectations better than the other three theories, as

it constantly produces a stable range of estimates of the updating rate around 0.2 to 0.3 across all specifications.

NI, another theory that also features information rigidity and captures similar qualitative patterns as SE, exhibits less cross-moment consistency. The major weakness of the model is that it produces unrealistically large sizes of the parameters to match the rigidity of the data. This is per se not a rejection of the theory. It is indeed found that once the more realistic inflation process of SV is used, NI estimation produces much more consistent and sensible values of parameters for professionals. Despite this improvement for professionals, however, NI proves to be a poor model to fit the patterns of household expectations. Although the previous literature (Coibion and Gorodnichenko (2012, 2015)) treat SE and NI as two indistinguishable theories that both produce information rigidity, this paper shows that using information from uncertainty significantly disciplines the parameter choices and allows me to distinguish the two theories by their model sensitivity.

Compared to the two rigidity models, a modified canonical DE that allows for heterogeneous degrees of over/underreaction are estimated to reveal a large degree of heterogeneity ranging from overreaction to underreaction across individuals. In addition, the model estimates are rather sensitive along many dimensions. The estimates with the high inflation episode do suggest a shift toward an average degree of overreaction of both professionals and households.

A hybrid of DE and NI, as in Bordalo et al. (2020), which accommodates the coexistent overreacting mechanisms and dispersed noisy information, does improve the fit of the model and its robustness compared to DE. The estimates feature a reasonable degree of underreaction (overreaction) of households with dispersed information in the low-inflation environment but an opposite pattern once the high-inflation episode is included in the estimation.

Table 8: Evaluation of Different Models

Criteria	SE	NI	DE	DENI
Sensitive to moments used for estimation?	No	Yes	Yes	No
Sensitive to the assumed inflation process?	No	Yes	Yes	No
Sensitive to two-step or joint estimate?	No	No	No	Yes
Sensitive to the type of agents?	No	Yes	Yes	Yes

6 Conclusion

Most studies on expectation formation that document how it deviates from the FIRE benchmark have focused on the first moment, namely the mean forecasts and the cross-sectional dispersion of the forecasts. However, this paper has shown that the surveyed forecasting uncer-

tainty by professionals and households provides useful information for understanding the exact mechanisms of expectation formation. It not only provides additional reduced-form testing results of rejecting FIRE, such as persistent disagreements in forecasting uncertainty and its inefficient revisions, but also provides additional moment restrictions to any particular model of expectation formation, which helps identify differences across theories.

At least three lines of questions remain unresolved in this paper and require future research. First, this paper focuses on a selective list of models of expectation formation and inevitably omits others that are likewise proven to match certain aspects of surveyed inflation expectations, such as adaptive learning (Marcet and Sargent, 1989; Evans and Honkapohja, 2012), experience-based learning (Malmendier and Nagel, 2015), heterogeneous models (Patton and Timmermann, 2010; Farmer et al., 2021), and asymmetric attention (Kohlhas and Walther, 2021). It would be fruitful to explore the corresponding predictions of these models about uncertainty. Second, throughout the analysis we maintained the normality/symmetric assumptions of the shocks and ignored beliefs in tail events or even higher moments. It would be natural to explore how different theories of expectation formation may contain different predictions on tail beliefs. Finally, although this paper focuses only on macroeconomic expectations regarding inflation, it is worth asking if the belief formation regarding individual variables such as income bears similar mechanisms and matches the observed empirical patterns of surveyed expectations and risks.³²

³²Here are a few recent studies on income/wage/unemployment/job search expectations: Mueller et al. (2021); Wang (2022); Koşar and Van der Klaauw (2022); Jäger et al. (2022); Caplin et al. (2023).

References

- Abel, Joshua, Robert Rich, Joseph Song, and Joseph Tracy**, “The measurement and behavior of uncertainty: Evidence from the ECB survey of professional forecasters,” *Journal of Applied Economics*, April 2016, *31* (3), 533–550.
- Afrouzi, Hassan, Spencer Y Kwon, Augustin Landier, Yueran Ma, and David Thesmar**, “Overreaction in expectations: Evidence and theory,” *The Quarterly Journal of Economics*, 2023.
- Angeletos, George-Marios, Zhen Huo, and Karthik A Sastry**, “Imperfect macroeconomic expectations: Evidence and theory,” *NBER Macroeconomics Annual*, 2021, *35* (1), 1–86.
- Armantier, Olivier, Giorgio Topa, Wilbert Van der Klaauw, and Basit Zafar**, “An overview of the Survey of Consumer Expectations,” *Economic Policy Review*, 2017, (23-2), 51–72.
- Bachmann, Rüdiger, Steffen Elstner, and Eric R Sims**, “Uncertainty and economic activity: Evidence from business survey data,” *American Economic Journal: Macroeconomics*, 2013, *5* (2), 217–49.
- Ball, Laurence**, “Why does high inflation raise inflation uncertainty?,” *Journal of Monetary Economics*, 1992, *29* (3), 371–388.
- , **Stephen G Cecchetti, and Robert J Gordon**, “Inflation and uncertainty at short and long horizons,” *Brookings Papers on Economic Activity*, 1990, *1* (1990), 215–254.
- Bassetti, Federico, Roberto Casarin, and Marco Del Negro**, “Chapter 15 - Inference on probabilistic surveys in macroeconomics with an application to the evolution of uncertainty in the survey of professional forecasters during the COVID pandemic,” in Rüdiger Bachmann, Giorgio Topa, and Wilbert van der Klaauw, eds., *Handbook of Economic Expectations*, Academic Press, January 2023, pp. 443–476.
- Bertrand, Marianne and Sendhil Mullainathan**, “Do people mean what they say? Implications for subjective survey data,” *American Economic Review*, 2001, *91* (2), 67–72.
- Binder, Carola C**, “Measuring uncertainty based on rounding: New method and application to inflation expectations,” *Journal of Monetary Economics*, 2017, *90*, 1–12.
- Bloom, Nicholas**, “The impact of uncertainty shocks,” *Econometrica*, 2009, *77* (3), 623–685.

- Bomberger, William A**, “Disagreement as a measure of uncertainty,” *Journal of Money, Credit, and Banking*, 1996, 28 (3), 381–392.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer**, “Diagnostic expectations and credit cycles,” *The Journal of Finance*, 2018, 73 (1), 199–227.
- , – , **Yueran Ma, and Andrei Shleifer**, “Overreaction in macroeconomic expectations,” *American Economic Review*, 2020.
- Branch, William A**, “The theory of rationally heterogeneous expectations: Evidence from survey data on inflation expectations,” *The Economic Journal*, 2004, 114 (497), 592–621.
- Caplin, Andrew, Victoria Gregory, Eungik Lee, Søren Leth-Petersen, and Johan Sæverud**, “Subjective Earnings Risk,” Technical Report, National Bureau of Economic Research 2023.
- Carroll, Christopher D**, “Macroeconomic expectations of households and professional forecasters,” *The Quarterly Journal of Economics*, 2003, 118 (1), 269–298.
- Cascaldi-Garcia, Danilo, Cisil Sarisoy, Juan M Londono, Bo Sun, Deepa D Datta, Thiago Ferreira, Olesya Grishchenko, Mohammad R Jahan-Parvar, Francesca Loria, Sai Ma, Marius Rodriguez, Ilknur Zer, and John Rogers**, “What is certain about uncertainty?,” *Journal of Economic Literature*, June 2023, 61 (2), 624–654.
- Clements, Michael P**, “Evaluating the rationality of fixed-event forecasts,” *Journal of Forecasting*, 1997, 16 (4), 225–239.
- , **Robert W Rich, and Joseph S Tracy**, “Surveys of professionals,” in “Handbook of Economic Expectations,” Elsevier, 2023, pp. 71–106.
- Coibion, Olivier and Yuriy Gorodnichenko**, “What can survey forecasts tell us about information rigidities?,” *Journal of Political Economy*, 2012, 120 (1), 116–159.
- **and** – , “Information rigidity and the expectations formation process: A simple framework and new facts,” *American Economic Review*, 2015, 105 (8), 2644–2678.
- , – , **and Saten Kumar**, “How do firms form their expectations? New survey evidence,” *American Economic Review*, 2018, 108 (9), 2671–2713.
- Cornand, Camille and Paul Hubert**, “Information frictions across various types of inflation expectations,” *European Economic Review*, 2022, 146, 104175.

- D’Acunto, Francesco, Ulrike Malmendier, and Michael Weber**, “What do the data tell us about inflation expectations?,” in “Handbook of economic expectations,” Elsevier, 2023, pp. 133–161.
- D’Amico, Stefania and Athanasios Orphanides**, “Uncertainty and disagreement in economic forecasting,” *Finance and Economics Discussion Series*, July 2008, 2008 (56), 1–38.
- Das, Sreyoshi, Camelia M Kuhnen, and Stefan Nagel**, “Socioeconomic status and macroeconomic expectations,” Technical Report, National Bureau of Economic Research 2017.
- Davies, Anthony and Kajal Lahiri**, “A new framework for analyzing survey forecasts using three-dimensional panel data,” *Journal of Econometrics*, 1995, 1 (68), 205–227.
- Delavande, Adeline**, “Probabilistic expectations in developing countries,” *Annual Review of Economics*, 2014, 6 (1), 1–20.
- , **Xavier Giné, and David McKenzie**, “Measuring subjective expectations in developing countries: A critical review and new evidence,” *Journal of Development Economics*, 2011, 94 (2), 151–163.
- der Klaauw, Wilbert Van, Wändi Bruine de Bruin, Giorgio Topa, Simon Potter, and Michael F Bryan**, “Rethinking the measurement of household inflation expectations: preliminary findings,” *FRB of New York Staff Report*, 2008, (359).
- D’Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber**, “IQ, expectations, and choice,” Technical Report, National Bureau of Economic Research 2019.
- Engelberg, Joseph, Charles F Manski, and Jared Williams**, “Comparing the point predictions and subjective probability distributions of professional forecasters,” *Journal of Business & Economic Statistics*, 2009, 27 (1), 30–41.
- Evans, George W and Seppo Honkapohja**, *Learning and expectations in macroeconomics*, Princeton University Press, 2012.
- Farmer, Leland, Emi Nakamura, and Jón Steinsson**, “Learning about the long run,” Technical Report, National Bureau of Economic Research 2021.
- Faust, Jon and Jonathan H Wright**, “Efficient forecast tests for conditional policy forecasts,” *Journal of Econometrics*, 2008, 146 (2), 293–303.
- and – , “Comparing Greenbook and reduced form forecasts using a large realtime dataset,” *Journal of Business & Economic Statistics*, 2009, 27 (4), 468–479.

- Friedman, Milton**, “Nobel lecture: Inflation and unemployment,” *Journal of Political Economy*, 1977, 85 (3), 451–472.
- Fuhrer, Jeffrey C**, “Intrinsic expectations persistence: Evidence from professional and household survey expectations,” 2018.
- Gambetti, Luca, Dimitris Korobilis, John Tsoukalas, and Francesco Zanetti**, “Agreed and Disagreed Uncertainty,” *arXiv preprint arXiv:2302.01621*, 2023.
- Giacomini, Raffaella, Vasiliki Skreta, and Javier Turen**, “Heterogeneity, inattention, and Bayesian updates,” *American Economic Journal: Macroeconomics*, 2020, 12 (1), 282–309.
- Giordani, Paolo and Paul Söderlind**, “Inflation forecast uncertainty,” *European Economic Review*, December 2003, 47 (6), 1037–1059.
- Glas, Alexander**, “Five dimensions of the uncertainty–disagreement linkage,” *International Journal of Forecasting*, April 2020, 36 (2), 607–627.
- Goldstein, Nathan**, “Tracking inattention,” *Journal of the European Economic Association*, 2023, p. jvad022.
- Jäger, Simon, Christopher Roth, Nina Roussille, and Benjamin Schoefer**, “Worker beliefs about outside options,” Technical Report, National Bureau of Economic Research 2022.
- Jain, Monica**, “Perceived inflation persistence,” *Journal of Business & Economic Statistics*, 2019, 37 (1), 110–120.
- Juodis, Artūras and Simas Kučinskas**, “Quantifying noise in survey expectations,” *Quantitative Economics*, 2023, 14 (2), 609–650.
- Jurado, Kyle, Sydney C Ludvigson, and Serena Ng**, “Measuring uncertainty,” *American Economic Review*, 2015, 105 (3), 1177–1216.
- Kohlhas, Alexandre N and Ansgar Walther**, “Asymmetric attention,” *American Economic Review*, 2021, 111 (9), 2879–2925.
- Koşar, Gizem and Wilbert Van der Klaauw**, “Workers’ perceptions of earnings growth and employment risk,” *Unpublished paper*, 2022.
- Lahiri, Kajal and Xuguang Sheng**, “Measuring forecast uncertainty by disagreement: The missing link,” *Journal of Applied Econometrics*, June 2010, 25 (4), 514–538.

- Lucas, Robert E**, “Expectations and the neutrality of money,” *Journal of economic theory*, 1972, *4* (2), 103–124.
- Macaulay, Alistair and James Moberly**, “Heterogeneity in imperfect inflation expectations: Theory and evidence from a novel survey,” 2022.
- Maćkowiak, Bartosz and Mirko Wiederholt**, “Optimal sticky prices under rational inattention,” *American Economic Review*, 2009, *99* (3), 769–803.
- Malmendier, Ulrike and Stefan Nagel**, “Learning from inflation experiences,” *The Quarterly Journal of Economics*, 2015, *131* (1), 53–87.
- Mankiw, N Gregory and Ricardo Reis**, “Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve,” *The Quarterly Journal of Economics*, 2002, *117* (4), 1295–1328.
- , – , and **Justin Wolfers**, “Disagreement about inflation expectations,” *NBER Macroeconomics Annual*, 2003, *18*, 209–248.
- Manski, Charles F**, “Measuring expectations,” *Econometrica*, 2004, *72* (5), 1329–1376.
- , “Survey measurement of probabilistic macroeconomic expectations: Progress and promise,” *NBER Macroeconomics Annual*, 2018, *32* (1), 411–471.
- Marcet, Albert and Thomas J Sargent**, “Convergence of least squares learning mechanisms in self-referential linear stochastic models,” *Journal of Economic Theory*, 1989, *48* (2), 337–368.
- Mueller, Andreas I, Johannes Spinnewijn, and Giorgio Topa**, “Job seekers’ perceptions and employment prospects: Heterogeneity, duration dependence, and bias,” *American Economic Review*, 2021, *111* (1), 324–63.
- Nordhaus, William D**, “Forecasting efficiency: Concepts and applications,” *The Review of Economics and Statistics*, 1987, pp. 667–674.
- Patton, Andrew J and Allan Timmermann**, “Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion,” *Journal of Monetary Economics*, 2010, *57* (7), 803–820.
- and – , “Forecast rationality tests based on multi-horizon bounds,” *Journal of Business & Economic Statistics*, 2012, *30* (1), 1–17.
- Pfäuti, Oliver**, “The inflation attention threshold and inflation surges,” *arXiv preprint arXiv:2308.09480*, 2023.

- Reis, Ricardo**, “Inattentive consumers,” *Journal of Monetary Economics*, 2006, *53* (8), 1761–1800.
- Rich, Robert and Joseph Tracy**, “The relationships among expected inflation, disagreement, and uncertainty: Evidence from matched point and density forecasts,” *Review of Economics and Statistics*, February 2010, *92* (1), 200–207.
- **and —**, “A closer look at the behavior of uncertainty and disagreement: Micro evidence from the euro area,” *Journal of Money, Credit, and Banking*, February 2021, *53* (1), 233–253.
- Rossi, Barbara and Tatevik Sekhposyan**, “Macroeconomic uncertainty indices based on nowcast and forecast error distributions,” *American Economic Review*, 2015, *105* (5), 650–655.
- Ryngaert, Jane**, “What do (and don’t) forecasters know about US inflation?,” 2017.
- Sims, Christopher A.**, “Implications of rational inattention,” *Journal of monetary Economics*, 2003, *50* (3), 665–690.
- Stock, James H and Mark W Watson**, “Why has US inflation become harder to forecast?,” *Journal of Money, Credit, and Banking*, 2007, *39*, 3–33.
- Valchev, Rosen and Luca Gemmi**, “Biased surveys,” Technical Report, National Bureau of Economic Research 2023.
- Wang, Tao**, “Perceived versus calibrated income risks in heterogeneous-agent consumption models,” *Unpublished paper, Johns Hopkins University*, 2022.
- Weber, Michael, Bernardo Candia, Tiziano Ropele, Rodrigo Lluberas, Serafin Frache, Brent H Meyer, Saten Kumar, Yuriy Gorodnichenko, Dimitris Georgarakos, Olivier Coibion et al.**, “Tell me something I don’t already know: Learning in low and high-inflation settings,” Technical Report, National Bureau of Economic Research 2023.
- Woodford, Michael**, “Imperfect common knowledge and the effects of monetary policy,” Technical Report, National Bureau of Economic Research 2001.
- Xie, Shihan**, “An estimated model of household inflation expectations: Information frictions and implications,” *Review of Economics and Statistics*, 2023, pp. 1–45.
- Zarnowitz, Victor and Louis A Lambros**, “Consensus and uncertainty in economic prediction,” *Journal of Political Economy*, 1987, *95* (3), 591–621.

Appendix

Detailed derivation

SE

$$\begin{aligned}
\overline{Var}_t^{se}(y_{t+h}) &= \sum_{\tau=0}^{+\infty} \underbrace{\lambda(1-\lambda)^\tau}_{\text{fraction who do not update until } t-\tau} \underbrace{Var_{t|t-\tau}^*(y_{t+h})}_{\text{uncertainty of most recent update at } t-\tau} \\
&= \sum_{\tau=0}^{+\infty} \lambda(1-\lambda)^\tau \sum_{s=1}^{h+\tau} \rho^{2(s-1)} \sigma_\omega^2 \\
&= \sum_{\tau=0}^{+\infty} \lambda(1-\lambda)^\tau (1 + \rho^2 + \dots + \rho^{2(h+\tau-1)}) \sigma_\omega^2 \\
&= \sum_{\tau=0}^{+\infty} \lambda(1-\lambda)^\tau \frac{\rho^{2(h+\tau)} - 1}{\rho^2 - 1} \sigma_\omega^2 \\
&= \sum_{\tau=0}^{+\infty} \lambda(1-\lambda)^\tau \frac{\rho^{2(h+\tau)}}{\rho^2 - 1} \sigma_\omega^2 - \sum_{\tau=0}^{+\infty} \lambda(1-\lambda)^\tau \frac{1}{\rho^2 - 1} \sigma_\omega^2 \\
&= \sum_{\tau=0}^{+\infty} \lambda(1-\lambda)^\tau \rho^{2\tau} \frac{\rho^{2h}}{\rho^2 - 1} \sigma_\omega^2 - \sum_{\tau=0}^{+\infty} \lambda(1-\lambda)^\tau \frac{1}{\rho^2 - 1} \sigma_\omega^2 \\
&= \sum_{\tau=0}^{+\infty} \lambda((1-\lambda)\rho^2)^\tau \frac{\rho^{2h}}{\rho^2 - 1} \sigma_\omega^2 - \sum_{\tau=0}^{+\infty} \lambda(1-\lambda)^\tau \frac{1}{\rho^2 - 1} \sigma_\omega^2 \\
&= \sum_{\tau=0}^{+\infty} \lambda(1-\lambda)^\tau \frac{1}{\rho^2 - 1} \sigma_\omega^2 \\
&= \left(\frac{\lambda \rho^{2h}}{(1-\rho^2 + \lambda \rho^2)(\rho^2 - 1)} - \frac{1}{\rho^2 - 1} \right) \sigma_\omega^2 \\
&= \left(\frac{\lambda \rho^{2h}}{1 - \rho^2 + \lambda \rho^2} - 1 \right) \frac{\sigma_\omega^2}{\rho^2 - 1} \\
&= \left(\frac{\lambda \rho^{2h} - 1 + \rho^2 - \lambda \rho^2}{1 - \rho^2 + \lambda \rho^2} \right) \frac{\sigma_\omega^2}{\rho^2 - 1}
\end{aligned} \tag{33}$$

NI

The steady-state nowcasting uncertainty Var_{ss}^{ni} is solved using the updating equation (Equation 13).

$$\begin{aligned}
\overline{\text{Var}}_{t|t}^{ni} &= \overline{\text{Var}}_{t|t-1}^{ni} - \text{Var}_{t|t-1}^{ni} H' (H \text{Var}_{t|t-1}^{ni} H' + \Sigma^v)^{-1} H \overline{\text{Var}}_{t|t-1}^{ni} \\
&\rightarrow \overline{\text{Var}}_{t|t}^{ni} = \rho^2 (\overline{\text{Var}}_{t-1|t-1}^{ni} + \sigma^2) \\
&\quad - \rho^2 (\overline{\text{Var}}_{ss}^{ni} + \sigma^2) H' (H \rho^2 (\overline{\text{Var}}_{ss}^{ni} + \sigma^2) H' + \Sigma^v)^{-1} H \overline{\text{Var}}_{ss}^{ni} \\
&\rightarrow \overline{\text{Var}}_{ss}^{ni} = \rho^2 (\overline{\text{Var}}_{ss}^{ni} + \sigma_\omega^2) \\
&\quad - \rho^2 (\overline{\text{Var}}_{ss}^{ni} + \sigma_\omega^2) H' (H \rho^2 (\overline{\text{Var}}_{ss}^{ni} + \sigma_\omega^2) H' + \Sigma^v)^{-1} H \overline{\text{Var}}_{ss}^{ni}
\end{aligned} \tag{34}$$

DE

$$\begin{aligned}
FE_{i,t+h|t}^{de} &= y_{i,t+h|t}^{de} - y_{t+h} \\
&= \rho^h y_t - y_{t+h} + \theta_i (\rho^h y_t - y_{i,t+h|t-1}^{de}) \\
&= \rho^h y_t - y_{t+h} + \theta_i (\rho^h y_t - y_{t+h} - FE_{i,t+h|t-1}^{de}) \\
&= FE_{t+h|t}^* + \theta_i (\rho^h y_t - y_{t+h} - FE_{i,t+h|t-1}^{de}) \\
&= (1 + \theta_i) FE_{t+h|t}^* - \theta_i FE_{i,t+h|t-1}^{de} \\
&= (1 + \theta_i) FE_{t+h|t}^* - \theta_i (\rho FE_{i,t+h-1|t-1}^{de} - \omega_{t+h}) \\
&= (1 + \theta_i) FE_{t+h|t}^* - \theta_i \rho FE_{i,t+h-1|t-1}^{de} + \theta_i \omega_{t+h} \\
&= (1 + \theta_i) FE_{t+h|t}^* + \theta_i (\omega_{t+h} - \rho FE_{i,t+h-1|t-1}^{de}) \\
&= (1 + \theta_i) FE_{t+h-1|t}^* + (1 + \theta_i) (-\omega_{t+h}) + \theta_i (\omega_{t+h} - \rho FE_{i,t+h-1|t-1}^{de}) \\
&= (1 + \theta_i) FE_{t+h-1|t}^* - \omega_{t+h} - \theta_i \rho FE_{i,t+h-1|t-1}^{de} \\
&= FE_{t+h|t}^* + \theta_i FE_{t+h-1|t}^* - \theta_i \rho FE_{i,t+h-1|t-1}^{de} \\
&= FE_{t+h|t}^* + \theta_i (FE_{t+h-1|t}^* - \rho FE_{i,t+h-1|t-1}^{de}) \\
&= FE_{t+h|t}^* + \theta_i (\rho FE_{t+h-1|t-1}^* + \rho^h \omega_t - \rho FE_{i,t+h-1|t-1}^{de}) \\
&= FE_{t+h|t}^* - \theta_i \rho (FE_{i,t+h-1|t-1}^{de} - FE_{t+h-1|t-1}^*) + \theta_i \rho^h \omega_t
\end{aligned} \tag{35}$$

DENI

Current forecast error is

$$\begin{aligned}
\overline{FE}_{t|t}^{deni} &= \rho y_{t-1|t-1}^{deni} + (1 + \theta)P_\epsilon(s_t^{pb} - \rho y_{t-1|t-1}^{deni}) - y_t \\
&= \rho(\overline{FE}_{t-1|t-1}^{deni} + y_{t-1}) + (1 + \theta)P_\epsilon(s_t^{pb} - \rho y_{t-1|t-1}^{deni}) - y_t \\
&= \rho(\overline{FE}_{t-1|t-1}^{deni} + y_{t-1}) + (1 + \theta)P_\epsilon(y_t + \epsilon_t - \rho(\overline{FE}_{t-1|t-1}^{deni} + y_{t-1})) - \rho y_{t-1} - \omega_t \\
&= \rho\overline{FE}_{t-1|t-1}^{deni} + (1 + \theta)P_\epsilon(\rho y_{t-1} + \omega_t + \epsilon_t - \rho(\overline{FE}_{t-1|t-1}^{deni} + y_{t-1})) - \omega_t \\
&= \rho\overline{FE}_{t-1|t-1}^{deni} + (1 + \theta)P_\epsilon(\omega_t + \epsilon_t - \rho\overline{FE}_{t-1|t-1}^{deni}) - \omega_t \\
&= \rho\overline{FE}_{t-1|t-1}^{deni} - (1 + \theta)\rho\overline{FE}_{t-1|t-1}^{deni} + (1 + \theta)P_\epsilon(\omega_t + \epsilon_t) \\
&= -\theta\rho\overline{FE}_{t-1|t-1}^{deni} + ((1 + \theta)P_\epsilon - 1)\omega_t + (1 + \theta)P_\epsilon\epsilon_t
\end{aligned} \tag{36}$$

Furthermore, we know

$$\begin{aligned}
\overline{FE}_{t+h|t}^{deni} &= \rho^h \overline{FE}_{t|t}^{deni} + FE_{t+h|t}^* \\
\overline{FE}_{t+h-1|t-1}^{deni} &= \rho^h \overline{FE}_{t-1|t-1}^{deni} + FE_{t+h-1|t-1}^*
\end{aligned} \tag{37}$$

So,

$$\begin{aligned}
\overline{FE}_{t+h|t}^{deni} &= \rho^h \overline{FE}_{t|t}^{deni} + FE_{t+h|t}^* \\
&= \rho^h(-\theta\rho\overline{FE}_{t-1|t-1}^{deni} + ((1 + \theta)P_\epsilon - 1)\omega_t + (1 + \theta)P_\epsilon\epsilon_t) + FE_{t+h|t}^* \\
&= -\theta\rho(\overline{FE}_{t+h-1|t-1}^{deni} - FE_{t+h-1|t-1}^*) + \rho^h(((1 + \theta)P_\epsilon - 1)\omega_t + (1 + \theta)P_\epsilon\epsilon_t) + FE_{t+h|t}^* \\
&= -\theta\rho\overline{FE}_{t+h-1|t-1}^{deni} + \theta\rho FE_{t+h-1|t-1}^* + \rho^h(((1 + \theta)P_\epsilon - 1)\omega_t + (1 + \theta)P_\epsilon\epsilon_t) + FE_{t+h|t}^* \\
&= \theta\rho(FE_{t+h-1|t-1}^* - \overline{FE}_{t+h-1|t-1}^{deni}) + \rho^h(((1 + \theta)P_\epsilon - 1)\omega_t + (1 + \theta)P_\epsilon\epsilon_t) + FE_{t+h|t}^*
\end{aligned} \tag{38}$$

Rearranging it, we get

$$\overline{FE}_{t+h|t}^{deni} - FE_{t+h|t}^* = -\theta\rho(\overline{FE}_{t+h-1|t-1}^{deni} - FE_{t+h-1|t-1}^*) + \rho^h(((1 + \theta)P_\epsilon - 1)\omega_t + \rho^h(1 + \theta)P_\epsilon\epsilon_t) \tag{39}$$

Set $h = 1$, we get

$$\overline{FE}_{t+1|t}^{deni} - FE_{t+1|t}^* = -\theta\rho(\overline{FE}_{t|t-1}^{deni} - FE_{t|t-1}^*) + \rho((1 + \theta)P_\epsilon - 1)\omega_t + \rho(1 + \theta)P_\epsilon\epsilon_t \tag{40}$$

When $\theta = 0$, $P_\epsilon = 1$, and $\epsilon_t = 0$, the equation collapses to FIRE.

Which is equivalent to the following:

$$\begin{aligned}
\overline{FE}_{t+1|t}^{deni} + \omega_{t+1} &= -\theta\rho(\overline{FE}_{t|t-1}^{deni} + \omega_t) + \rho((1+\theta)P_\epsilon - 1)\omega_t + \rho(1+\theta)P_\epsilon\epsilon_t \\
\rightarrow \overline{FE}_{t+1|t}^{deni} &= -\theta\rho(\overline{FE}_{t|t-1}^{deni} + \omega_t) + \rho((1+\theta)P_\epsilon - 1)\omega_t + \rho(1+\theta)P_\epsilon\epsilon_t \\
\rightarrow \overline{FE}_{t+1|t}^{deni} &= -\theta\rho\overline{FE}_{t|t-1}^{deni} - \theta\rho\omega_t + \rho((1+\theta)P_\epsilon - 1)\omega_t + \rho(1+\theta)P_\epsilon\epsilon_t - \omega_{t+1} \\
&= -\theta\rho\overline{FE}_{t|t-1}^{deni} - \theta\rho\omega_t + \rho((1+\theta)P_\epsilon - 1)\omega_t + \rho(1+\theta)P_\epsilon\epsilon_t - \omega_{t+1} \\
&= -\theta\rho\overline{FE}_{t|t-1}^{deni} - (\rho(1+\theta)P_\epsilon - \rho - \theta\rho)\omega_t + \rho(1+\theta)P_\epsilon\epsilon_t - \omega_{t+1} \\
&= -\theta\rho\overline{FE}_{t|t-1}^{deni} - (\rho P_\epsilon + \rho\theta P_\epsilon - \rho - \theta\rho)\omega_t + \rho(1+\theta)P_\epsilon\epsilon_t - \omega_{t+1} \\
&= -\theta\rho\overline{FE}_{t|t-1}^{deni} - \rho(P_\epsilon + \theta P_\epsilon - 1 - \theta)\omega_t + \rho(1+\theta)P_\epsilon\epsilon_t - \omega_{t+1} \\
&= -\theta\rho\overline{FE}_{t|t-1}^{deni} + \rho((1+\theta)(1-P_\epsilon))\omega_t + \rho(1+\theta)P_\epsilon\epsilon_t - \omega_{t+1}
\end{aligned} \tag{41}$$

This means

$$\overline{FE}_{\bullet+1|\bullet}^{deni2} = \frac{\sigma_\omega^2 + \rho^2(1+\theta)^2(1-P_\epsilon)^2\sigma_\omega^2 + \rho^2(1+\theta)^2P_\epsilon^2\sigma_\epsilon^2}{1 + \theta^2\rho^2} \tag{42}$$

Cross-sectional dispersion of forecast uncertainty

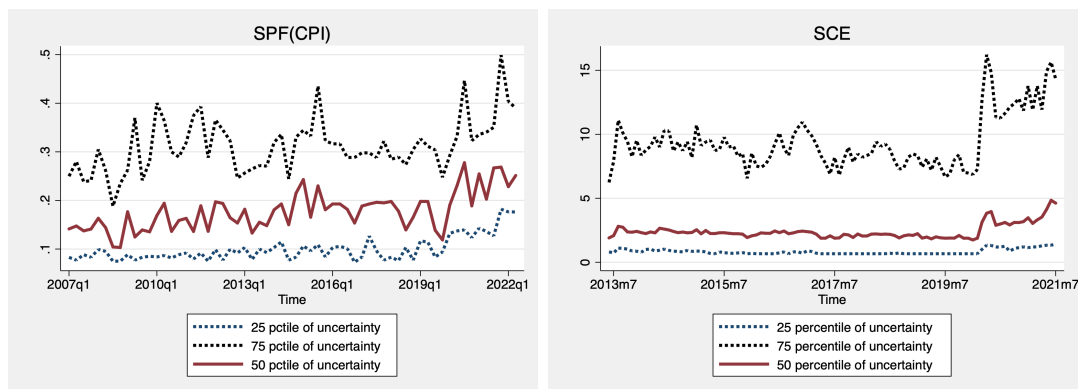
Persistent dispersion in expectations has been among the most commonly cited evidence that is inconsistent with the assumption of identical expectations predicted by FIRE (Mankiw et al. (2003)). A similar argument can be made with the dispersion in forecasting uncertainty, as FIRE predicts individuals share an equal degree of uncertainty.³³

Figure 4 plots the median uncertainty along with its 25/75 percentiles in both SCE and SPF. There is persistent dispersion in uncertainty across agents. The dispersion in uncertainty of households is much greater than that of the professionals. The IQR of the uncertainty of households is around 150-200 times(12–14 times in standard deviation terms) of that of professional forecasters.

One difference in the distribution of uncertainty between households and professionals is that the distribution of the former is more skewed toward the right (higher uncertainty), while professional forecasters disagree in uncertainty more symmetric around its cross-sectional mean.

³³In contrast, SE predicts that the uncertainty of individuals differs in that agents are not equally updated at a point in time. NI generates a homogeneous degree of uncertainty only under the stringent conditions of equal precision of signals and the same prior for uncertainty (Equation 15). DE predicts an equal degree of uncertainty across agents (Equation 20). Therefore, taken by the face value, the presence of dispersion of uncertainty across agents is not consistent with predictions of FIRE, or the canonical version of NI and DE.

Figure 4: Dispersion of Uncertainty



Another pattern worth discussing in Figure 4 is that there is a notable rise in the dispersion of uncertainty along the rapid rise in inflation in 2020, which was primarily driven by an increase in uncertainty reported in the upper end of the forecasts.

Reduced-form tests with forecast errors

The FE-based null-hypothesis of FIRE uses the moment restrictions on forecast errors. In plain words, the null hypotheses of the three tests are the following: First, since the forecasts are on average unbiased according to FIRE, forecast errors across agents should converge to zero in a large sample. Second, forecast errors of non-overlapping forecasting horizon are not serially correlated. Third, forecast errors cannot be predicted by any information available at the time of the forecast, including the mean forecast itself and other variables that are in the agent’s information set. This follows from Equation 2. In addition, I include what is called a weak version of the FE-based test, which explores the serial correlation of forecast errors in overlapping periods, i.e., 1-year-ahead forecasts within 1 year. The forecast errors are correlated to the extent of the realized shocks in the overlapping periods. So the positive serial correlation does not directly violate FIRE. But the correlation of overlapping forecast errors still contains useful information about the size of the realized shocks.

Individual-level data are used whenever possible, using the panel structure of both surveys. Since test 2 and 3 require individual forecasts in vintages that are more than 1 year apart while SCE only surveys each household for 12 months, the two tests are done with the population average expectations for SCE. Also, the regressions are adjusted accordingly depending on the quarterly and monthly frequency of SPF and SCE. Since these regressions are based on 1-year inflation in overlapping periods, Newy-West standard error is computed for hypothesis testing.

First, all three forecast series easily reject the null hypothesis of unbiasedness at the significance level of 0.1%. There are upward biases in both professional forecasts of core PCE

inflation and households' forecast of headline inflation,³⁴ while at the same time professionals underpredicted core CPI inflation over the entire sample period. This was primarily driven by the underprediction of inflation over the recent 2 years since the pandemic.

Second, the average point forecast 1 year ago predicts the forecast errors of both groups at the significance level of 0.1%. For headline CPI inflation, for instance, a one percentage point inflation forecast corresponds to 0.35 percentage points of the forecast errors 1 year later. Thus, test 2 in Table 9 easily rejects the second hypothesis test of FIRE that past information does not predict future forecast errors. This suggests that both types of agents inefficiently use all information when making the forecasts.

Third, forecast errors are positively correlated with the forecast errors 1 year ago, with a significant coefficient ranging from 0.35 to 0.572. A higher positive auto-correlation coefficient of forecast errors by households is consistent with the common finding that households are subject to more information rigidity than attentive professionals.

Lastly, test 4 in Table 9 presents a higher serial correlation of forecast errors produced within a year. For SPF forecasts, the serial correlation does not exist beyond two quarters, implying the relative efficiency of professional forecasts. For households, the forecast errors are more persistent over the entire year, in that current forecast errors are correlated with all past forecast errors over the past three quarters. Although the persistence of 1-year forecast errors within 1 year does not directly violate FIRE, the fact that households' forecast errors are more persistent than professionals' indicates that the former group is subject to a higher degree of rigidity than the latter one.

Reduced-form tests with mean revision

Table 10 reports the results for mean revisions (InfExp_Mean_rv). The first column of each panel regards the regression on a constant. Average revisions of forecasts of CPI and PCE by SPF inflation are both negative and significant, indicating an average downward revision over the sample period. The second to fourth columns in the upper panel examine the dependence of revisions on past information beyond forecast horizons. In particular, revisions are negatively correlated with the median SPF forecasts made four quarters before (InfExp_Mean_ct50) and are also serially correlated four or five quarters apart, and the coefficients are positive and significant. This implies that individual revisions in forecasts react to lagged information, some evidence against the null hypothesis of FIRE.

Similar to professionals, average household revisions in SCE also positively depend on past revisions made 2 years before. Since individual revisions are not available, time-fixed effect

³⁴Coibion et al. (2018) finds the same upward bias for firms' managers.

Table 9: Tests of Rationality and Efficiency Using Forecast Errors

	SPF CPI	SPF PCE	SCE
Test 1: Bias			
Constant	-3.021*** (0.242)	0.460*** (0.047)	1.673*** (0.008)
N	5510	1610	112668
Test 2: FE depends on past information			
Forecast 1-yr before	0.350*** (0.035)	0.460*** (0.047)	4.190*** (0.659)
Constant	-3.452*** (0.386)	-2.333*** (0.192)	-12.92*** (2.213)
N	3945	1610	84
R^2	0.828	0.826	0.311
Test 3: FE of non-overlapping forecast horizons are serially correlated			
Forecast Error 1-year before	0.350*** (0.035)	0.460*** (0.047)	0.572** (0.195)
Constant	0.314 (0.231)	-1.351*** (0.156)	-0.149 (0.445)
N	3945	1610	84
R^2	0.828	0.826	0.0957
Time FE	Yes	Yes	No
Test 4: Overlapping FE are serially correlated			
Forecast Error 1-q before	0.502*** (0.060)	0.551*** (0.075)	0.327*** (0.010)
Forecast Error 2-q before	0.0901 (0.064)	0.231*** (0.060)	0.341*** (0.024)
Forecast Error 3-q before	0.146* (0.065)	0.0693 (0.052)	0.333*** (0.023)
Constant	1.147*** (0.224)	-0.356*** (0.058)	0.509*** (0.035)
N	2971	1338	4432
R^2	0.890	0.903	0.243

Note: White standard errors reported in the parentheses of estimations. *** $p < 0.001$, ** $p < 0.01$, and * $p < 0.05$.

cannot be controlled to absorb all common contemporary innovations. Instead, I control for the average forecast error (InfExp_FE) that has just been realized in the same period, meant to capture innovations in the information set common to all forecasters. It is found to be negatively correlated with average revisions, implying information rigidity.³⁵

³⁵Coibion and Gorodnichenko (2012) show that information rigidity implies that past forecast errors and current revisions are positively correlated, where forecast errors are defined as the opposite sign as in this paper.

Table 10: Tests of Revision Efficiency Using Mean Revision

	SPF CPI				SPF PCE				SCE			
	Mean revision	Past info	4q before	5q before	Mean revision	Past info	4q before	4-5 q before	Mean revision	24m before	25m before	26m before
L4.InfExp_Mean_ct50		-0.512* (0.227)				-0.503 (0.280)		L24.InfExp_Mean_rv	0.233* (0.100)			
L4.InfExp_Mean_rv			0.211*** (0.047)				0.265*** (0.052)	L25.InfExp_Mean_rv		0.225* (0.103)		
L5.InfExp_Mean_rv				0.218*** (0.046)			0.265*** (0.054)	L26.InfExp_Mean_rv				0.254* (0.098)
								InfExp_FE	-0.330*** (0.046)	-0.327*** (0.047)	-0.331*** (0.047)	
Constant	-0.075*** (0.000)	0.994* (0.487)	-0.028*** (0.007)	-0.016 (0.007)	-0.036*** (0.000)	0.888 (0.528)	0.001 (0.007)	0.026*** (0.007)	-0.106 (0.054)	0.504*** (0.074)	0.510*** (0.080)	0.530*** (0.080)
R2	0.376	0.030	0.439	0.441	0.393	0.027	0.461	R2	0.000	0.595	0.600	0.616
N	1697	1697	1290	1149.000	1362	1362	1033	924.000	74	50	49	48
Time FE	Yes	No	Yes	Yes	Yes	No	Yes	Yes	No	No	No	No

Standard errors are clustered by date. *** p<0.001, ** p<0.01, and * p<0.05.