Predicting Forecaster Inattention

By William R. Bennett*

May 20, 2025

I provide new estimates of forecaster inattention and information rigidity related to the sticky information model of expectation formation. While most papers use aggregate-level regressions or model calibrations to estimate the amount of information rigidity, I employ a more granular estimation strategy using micro-level data from the US Survey of Professional Forecasters (SPF). I provide evidence that the true amount of forecaster inattention is much smaller than previously thought. I also document novel state-dependence and time series facts. My results imply that some other, stronger source of information rigidity must exist to account for the discrepancy between the aggregate- and micro-level results. A model accounting for information rigidity should not use sticky information as neither its only nor main mechanism, as doing so could result in incorrect model predictions.

I. Introduction

Expectations have always had an important role in economics. Uncertainty about conditions of the future affects many economic decisions. Individuals use the information at their disposal to make predictions about those conditions, and choose their actions accordingly. To make better predictions, individuals will often seek to obtain more information. However, information, like everything else, is never truly free. It comes at a cost, perhaps literally, metaphorically, or both. In the literal sense, information can come in a form that can be bought with money; people can buy newspapers, watch a news channel by paying for a TV or phone, pay tuition and other expenses for education, or purchase data-sets from companies and government entities. In the metaphorical sense, information must be acquired and understood by exerting some kind of effort. It requires using our limited time, physical energy, and mental faculties.

However, many macroeconomic models do not factor in the costs of information and assume that learning is frictionless. The standard way to model the expectation formation process in economics for the past half-century has been the full-information rational expectations (FIRE) paradigm.

*Department of Economics, University of Houston (email: wrbennett@uh.edu)

This means two distinct things. First, full-information requires agents to know all present and past conditions of their environment, and have uncertainty only about future conditions. Second, rational expectations requires agents to use all available information and not make systematically biased forecasts. Together, these assumptions produce several hypotheses, with the main one being that a FIRE agent's forecast error should not be consistently predictable: no variable should be useful in predicting an individual forecast error because the agent should have already accounted for that variable when making their forecast.

Since its inception, researchers have tested FIRE's hypotheses in several ways and many have rejected it, with several attempting to explain and model deviations from FIRE. Two importantand influential contributions to this area of research were Mankiw and Reis (2002) and Sims (2003), which popularized forecasting models with sticky information and noisy information, respectively. Both of these models relax the full-information part of FIRE by introducing some kind of information rigidity: either by having some agents sometimes not receive any information sometimes (sticky) or by having all agents always receive information but the information itself has some uncertainty (noisy), which hearkens back to Lucas (1972)'s model with signal extraction.[1]

From there, several papers have applied these and similar methods to more complicated macroeconomic models with a focus on the actual macroeconomy, such as Mankiw, Reis, and Wolfers (2007), Gorodnichenko (2008), Reis (2009), Carrillo (2012), Paciello and Wiederholt (2014), Tortorice (2018), Gelain et al. (2019), Carroll et al. (2020), Angeletos and Huo (2021), and Morales-Jiménez (2022). Many papers specifically focus on how information rigidities could replace or augment sticky price New Keynesian models, such as Khan and Zhu (2006), Kiley (2007), Klenow and Willis (2007), Korenok (2008), Maćkowiak and Wiederholt (2009), Coibion (2010), Dupor, Kitamura, and Tsuruga (2010), Knotek (2010), and L'Huillier (2020).

For example, several of these papers show how incorporating sticky information into a model can explain the sluggish updating of prices and wages, the delayed response of inflation to monetary policy, the flattening of the Phillips curve, inflation volatility fluctuations, and the strong effects of forward guidance. However, the degree of stickiness required to explain these puzzles is skeptically high: most of the above papers require a quarterly rigidity parameter above 0.5, implying at minimum 50% of agents making forecasts learn nothing about the economy every quarter. I find that the parameter is 5 times smaller, which casts doubt on sticky information being a reasonable way to explain the above puzzles.

[1]Throughout this paper I use the terms *sticky information*, *information rigidity*, and *forecaster inattention* interchangeably for verbal variety. They are not technically the same thing, but for the purposes of this paper and its results they can be treated as such.

The other category of papers (to which this paper belongs) stays focused on the forecasting aspects of the model and places less emphasis on how to embed it into a macroeconomic model, such as Morris and Shin (2006), Branch (2007), Coibion and Gorodnichenko (2012), Sarte (2014), Cavallo, Cruces, and Perez-Truglia (2017), Bergemann, Bonatti, and Smolin (2018), Coibion, Gorodnichenko, and Kumar (2018), Gaus and Sinha (2018), Kim and Kim (2019), Morikawa (2019), Bordalo et al. (2020), and Kohlhas and Walther (2021). These papers have shown how different information rigidities and forecaster heterogeneities can explain certain data but not others, thus providing conflicting evidence for and against the sticky information model, noisy information model, their generalizations, and other models of expectation formation.

This paper's contribution to this literature is threefold: first, I provide additional evidence against the sticky information model being a reasonable way of explaining forecasters' departures from FIRE; second, conditional on sticky information still being a (now weaker) explanation of non-FIRE behavior, I show how the parameter determining the stickiness of information is heterogeneous across variables and forecast horizons with greater precision than in prior research; three, sticky information is state-dependent and autoregressive, and therefore not constant like the standard model assumes.

My first contribution involves empirically showing that the literature's estimates of the stickiness parameter are inflated upward. A large majority of previous estimation strategies can broadly be separated into two groups. First, before the influential work of Coibion and Gorodnichenko (2012) and Coibion and Gorodnichenko (2015), the parameter was usually calibrated to match some moment of the paper's macroeconomic model, such as inflation persistence. Then those papers showed how to estimate various information rigidities, including information stickiness, using regressions with mean forecasts. Since then papers less focused on the macroeconomic model implications have used this approach.

However, using an approach similar to Andrade and Bihan (2013) and Giacomini, Skreta, and Turen (2020), I employ a micro-estimation strategy that does not require regressing aggregate variables. This approach involves counting the number of times forecasts were unchanged, with the fraction of times an individual forecaster's predictions were unchanged being a measure of that forecaster's inattention. Andrade and Bihan (2013) and to a lesser extent Giacomini, Skreta, and Turen (2020) use this micro-estimate approach to analyze two surveys of professional forecasters (from the European Central Bank and Bloomberg, respectively). I follow their approach by analyzing the US SPF, which has several more decades of forecasts and variables. I argue this is a better approach than Coibion and Gorodnichenko (2015)'s aggregate regression, which, due to a loss of observations from aggregation, produces larger standard errors of coefficients and allows for other conflicting interpretations of the estimated parameter.

My second contribution involves determining what predicts if a forecaster will be inattentive. The sticky information model assumes a constant stickiness parameter such that the probability of not having an information set update is always the same in each time period for each forecaster for each variable being forecast for each horizon forecast. I find evidence that each of these is untrue. While multiple papers mentioned above have also found evidence of this, I do so in a different way that comes with a much larger sample size. Previous research involved sub-setting their data-sets into different groups of time periods—for example, before, during, and after the Great Moderation and finding that information rigidity was lower before and after the Great Moderation and higher during. The interpretation is that in times of unusual economic activity, such as inflation being relatively high/low or volatile, people pay more attention to and care more about what is happening in the economy, leading them to find ways to update their information sets more and improve their predictions.

I instead estimate a distributed lag model and regress inattention on variables directly related to the forecaster. Specifically, I analyze how previous individual forecast errors, previous average forecast errors, and previous levels of the forecast variable can change the likelihood of an individual updating their forecast(s). In contrast to some of the literature, I find that variables being very different from their average does not impact the likelihood of forecasts updating. Previous individual and mean forecast errors do have an effect, which is a new discovery enabled by my micro-estimation strategy and relatively large dataset.

My last contribution is to construct time series for the average degree of information stickiness for multiple economic variables. These series, like my findings above, are only possible (have suitable power) because of my micro-estimation strategy and large dataset. I find that none of the series have a unit-root but do have positive autocorrelation, have their mean and variance declining over time because of structural breaks that occur around the start of end of multiple recessions, and have a significant impact on the forecasted variables in the short run.[2]

The rest of this paper is organized as follows. Section II describes the sticky information model, derives the model's predictions and implications, and discusses estimates of the stickiness parameter in the literature. Section III describes the data used and explains the construction of the stickiness variables. Section IV contains my empirical results, including my various estimates for constant yet heterogeneous information stickiness, my estimates of state-dependence for nonconstant information

[2]While not causal, the impulse response functions show increased forecaster inattention is associated with changes to most of the forecasted variables.

stickiness, and new stylized facts about the time series properties of information stickiness. Section V concludes. Section VI contains appendices.

II. The Sticky Information Model

The standard sticky information model assumes that in each time period, each individual forecaster updates their information set (relative to the previous period) with constant probability $(1-\lambda)$, where λ is the degree of information rigidity. A higher value of λ implies a lower probability of receiving new information. By updating, the forecaster acts as a FIRE agent, while forecasters who do not update receive no information (that is useful for forecasting) yet still make a rational forecast based off their limited information set.

The combination of forecasters who updated this period and those whose last update was at an earlier time creates disagreement in the mean forecast. The mean forecast across individuals at time t of the forecasted variable x at time $t+h$ is given by

$$f_t x_{t+h} = (1-\lambda)E_t x_{t+h} + \lambda f_{t-1} x_{t+h}. \tag{1}$$

The mean forecast is a weighted average of forecasts from individuals who updated at time t, $E_t x_{t+h}$, and forecasts from individuals who did not update and leave their forecasts unchanged from $t-1$, $f_{t-1}x_{t+h}$, with weights $(1-\lambda)$ and λ respectively.[3] The different time scripts on f and x are for the time period that the forecast is made and the time period of the variable forecasted, respectively. Full-information rational expectations are such that, for those who received information,

$$E_t x_{t+h} = x_{t+h} - \nu_{t+h,t} \tag{2}$$

where $\nu_{t+h,t}$ is the unpredictable FIRE error term and is therefore uncorrelated with any information or variables from period t or earlier. Combining equations 1 and 2 yields the predicted relationship between the ex post mean forecast error and the ex ante mean forecast revision

$$\overbrace{x_{t+h} - f_t x_{t+h}}^{\text{ex post mean forecast error}} = \frac{\lambda}{1-\lambda} \overbrace{(f_t x_{t+h} - f_{t-1} x_{t+h})}^{\text{ex ante mean forecast revision}} + \nu_{t+h,t}. \tag{3}$$

for horizon h.

Importantly, the coefficient on the forecast revision depends only on the degree of information

[3]Following the notation and timing of the data sources discussed later, forecasters make forecasts *of* time t (i.e. $h=0$) *at* time t but *before* x_t is revealed.

rigidity λ. In the special case of no information frictions, $\lambda = 0$, and the specification becomes FIRE and reduces to equation 2, i.e., the average forecast error is unpredictable. When $0 < \lambda < 1$, there is predictability in mean forecast errors that reflects the slow updating of information by some agents.[4] This inertia anchors the mean forecast to the previous period's, leading to a gradual adjustment of mean forecasts.

Using data from the US SPF, Coibion and Gorodnichenko (2015) estimate equation 3 with inflation as the main variable of interest and horizon $h = 3$ (at the quarterly frequency) and get a value for λ of 0.544. This result implies that about 50% of professional forecasters learn no useful information for predicting inflation 3 quarters ahead. They find that their estimate barely changes by adding additional control variables. They also estimate the equation 3 using the Livingston Survey with inflation as the variable of interest and horizon $h = 1$. Depending on the pooling of forecasters, they get a range for λ of 0.322 to 0.515. They also look at the Michigan Survey of Consumers, getting an estimate of 0.413, as well as financial markets inflation expectations constructed using a method developed in Haubrich, Pennacchi, and Ritchken (2008) and data from the Federal Reserve Bank of Cleveland, getting 0.599. All results listed above are statistically significant at the ten percent level, with all but two being significant at the five percent level.

From there, Coibion and Gorodnichenko (2015) continue to estimate more regressions (using the US SPF) similar to equation 3 but extended to include multiple variables and horizons. These results are comparable to the estimates previously stated. Interestingly, they find that they can reject the null hypothesis that the estimated λ's for the different variables are equal, but cannot reject the null hypothesis that the estimated λ's for the different horizons are equal (and find the horizon of 3 quarters insignificant and smaller than the shorter horizons). Later I will show that most horizons are indeed significantly different (if only slightly).

I will briefly state some additional estimates of information rigidity from relevant papers. The seminal Mankiw and Reis (2002) calibrate λ to match their macroeconomic model and calculate a value of 0.73 when their monthly data is adjusted to a quarterly frequency. Andrade and Bihan (2013) use the European Central Bank's Survey of Professional Forecasters to estimate inattention for inflation, unemployment, and GDP growth and get values of 0.28, 0.25, and 0.2 respectively. Finally, Giacomini, Skreta, and Turen (2020) construct a forecasting model with elements of both sticky and noisy information. Using the "Economic Forecasts ECFC" Survey of Professional Forecasters conducted by Bloomberg, they analyze monthly updates of US annual year-on-year inflation forecasts. For the purposes of their model, they calculate a stickiness parameter for every month of

[4] The case of $\lambda = 1$ is nonsensical, as it implies information sets are never updated and forecasts never change.

their data, with most of their values being between 0.157 and 0.415 (after conversion to a quarterly frequency).

Contrary to the calibration methods as in Mankiw and Reis (2002) and the aggregated regressions as in Coibion and Gorodnichenko (2015), Andrade and Bihan (2013) and Giacomini, Skreta, and Turen (2020)'s calculations of λ involve constructing an indicator variable capturing whether an agent updates their previous forecast, i.e. 1 if yes and 0 if no. In the empirical work that follows, I adapt the indicator variable approach to provide new micro-estimates for the amount of information rigidity in the US SPF. While the samples of Andrade and Bihan (2013) and Giacomini, Skreta, and Turen (2020) have short time dimensions, the US SPF goes back to the late 1960s. Combining the indicator variable approach with the large sample of the US SPF allows greater precision of estimates than other methods and datasets. It also makes it possible to investigate if sticky information has state-dependence and/or autoregressive properties.

III. Data Sources and Variables

The Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters (SPF) is a rolling, unbalanced panel survey. The Federal Reserve Bank has administered the survey every quarter since 1990q2. Before 1990q2, the American Statistical Association (ASA) and National Bureau of Economic Research (NBER) conducted it starting in 1968q4. The sample used in this paper ends with the survey round of 2023q3, so the sample size is 220 time periods. Each quarter, between 20 and 100 professional forecasters are asked for forecasts of several economic and financial variables.[5]

Following Coibion and Gorodnichenko (2015), focus is given to inflation, real GDP growth, industrial production growth, the unemployment rate, and the number of housing starts. The first three variables are quarter-over-quarter annual growth rates, unemployment is a percentage, and housing is the raw amount of starts (new constructions).[6] The variables and their forecasts are compared using real-time data to prevent reclassifications and redefinitions from affecting estimation since final values are not directly observable to the agents at the time they make their prediction.[7] [8]

In each time period t, forecaster i makes five forecasts of different horizons h for each variable x: $\{f_{it}x_{t+h}\}_{h=0}^{4}$. I denote a forecast for a certain horizon as *sticky* if it didn't change from the

[5] When discussing research using surveys of professional forecasters, it is often asked, "Who are these forecasters?", and "What incentives do they have to make forecasts?". The identities of the forecasters are not public knowledge (though their general profession is); the forecasters are volunteers and are not compensated; the accuracy of their forecasts does not impact the Federal Reserve Bank's decision to have them continue to take part in the survey. More details can be found in the survey documentation.

[6] See Appendix VI.A for information on the growth rate calculations.

[7] The real-time data is also published by the Federal Reserve Bank of Philadelphia.

[8] While the survey only asks for quarterly forecasts of the variables, unemployment and industrial production are monthly data. When calculating forecast errors for these variables, the quarterly forecast is compared to the average of the 3 months in that quarter.

previous time period, i.e. if $f_{it}x_{t+h} = f_{it-1}x_{t+h}$. I construct an indicator variable that captures this:

$$\lambda_{xith} = \begin{cases} 1 & \text{if } f_{it}x_{t+h} = f_{it-1}x_{t+h} \\ 0 & \text{if } f_{it}x_{t+h} \neq f_{it-1}x_{t+h} \end{cases}. \tag{4}$$

A value of 1 implies one of two things. First, the forecaster learned nothing relevant (for whatever reason) and rationally didn't change their prediction. Second, the forecaster actually did learn something relevant, but irrationally didn't change their prediction.[9] A generalization of the sticky information model would be needed to address potential irrationality. Following the literature, I assume the first case to be what is actually happening.

This stickiness indicator variable is the main object of interest in this paper. Regressing this on various other indicators (fixed effects) and controls provides new micro-estimates of the degree of forecaster inattention. Averaging these indicator lambda's across variables and horizons shows that inattention is heterogeneous and the coefficients on the control variables will demonstrate significant state-dependence.

A caveat to λ_{xith}'s construction is if $f_{it}x_{t+h}$ is missing from the dataset but $f_{it-1}x_{t+h}$ is not, than I set $\lambda_{xith} = 1$. This is based on the assumption that if a forecaster made a forecast last period but didn't make it this period it is because they received no new information warranting them to update. The results are robust to this assumption (they get insignificantly smaller), but the assumption is retained so that the variables using the horizon $h = 3$ have a comparable number of observations to the other horizons.[10]

IV. Estimation

My empirical results are divided into three categories. Section IV.A contains my new estimates for the degree of sticky information for multiple variables and horizons. Section IV.B uses control variables to test if a forecaster updating their forecast or not can be predicted. Section IV.C discusses numerous new properties of the autoregressive nature of sticky information.

[9]There is actually a third case; the forecaster learns something and rationally changes their prediction slightly, but rounding causes the reported number to be the same. The literature has found this to have a minuscule effect.

[10]A possible extension of this work would be to jointly model sticky and missing forecasts, including the case where a forecaster gives no forecasts for any horizons (which I currently treat as unusable observations for estimation) as both represent (potentially different) forms of inattention, but this is left to future research.

A. Main Results

To estimate the degree of sticky information of each variable and horizon, I run the following regression,

$$\lambda_{xith} = \boldsymbol{\alpha_x} + \boldsymbol{\alpha_i} + \boldsymbol{\alpha_h} + \varepsilon_{xith} \tag{5}$$

where λ_{xith} is the stickiness indicator variable, $\boldsymbol{\alpha_x}$ is a vector of fixed effects for the five forecasted variables, $\boldsymbol{\alpha_i}$ is a vector of fixed effects for the individual forecasters, and $\boldsymbol{\alpha_h}$ is a vector of fixed effects for the four horizons.[11]

The main estimates from the regression are in Table 1, while explicit information about the regression is shown in Table 2. The estimate for the average λ without accounting for which variable or horizon is being forecast is 0.098. This value for λ can be considered the average amount of inattention without controlling for anything (except individual forecaster fixed effects), and thus the value that researchers may wish to use in their models that require constant information rigidity.

Table 1—Forecaster Inattention Values

	λ
Mean	**0.098**
Inflation	0.057
GDP	0.022
Unemployment	0.176
Industrial Production	0.073
Housing	0.165
$h = 0$	0.086
$h = 1$	0.098
$h = 2$	0.104
$h = 3$	0.105

Notes: This table shows the estimated overall mean information stickiness and the average for the five variables and four horizons. These numbers are calculated using Table 1.2: the mean comes from regression 1, while the variables and horizons come from regressions 2 and 3 by adding a variable's or horizon's coefficient to the constant term, where the constant term is the estimate for inflation and $h = 0$, respectively.

The estimates for the average λ for inflation, GDP, unemployment, industrial production, and housing are all significantly different from each other, with GDP having the least amount of information rigidity and unemployment having the most. However, while horizons 0 and 1 are significantly different from each other and from the other horizons, horizons 2 and 3 are not significantly different from each other.[12] A possible explanation for this could be that forecasters, in this context,

[11] Time fixed effects have negligible impacts on estimates and standard errors. They excluded to allow the inclusion of the controls in equation 6, as having both would cause multicollinearity.

[12] There is no λ_{xit4} because this period was not forecasted in $t-1$.

divide the future into three main parts: the immediate future ($h = 0$, 0-3 months), the near future ($h = 1$, 3-6 months) and the far future ($h \geq 2$, 6+ months).

Table 2—Forecaster Inattention for Multiple Variables and Horizons

	(1) FE	(2) FE	(3) FE	(4) FE	(5) FE Logit
GDP		-0.035***		-0.035***	-1.058***
		(0.004)		(0.004)	(0.089)
Unemployment		0.120***		0.120***	1.316***
		(0.008)		(0.008)	(0.094)
Industrial Production		0.016***		0.016***	0.271***
		(0.005)		(0.005)	(0.077)
Housing		0.108***		0.108***	1.230***
		(0.011)		(0.011)	(0.112)
$h = 1$			0.012***	0.012***	0.153***
			(0.002)	(0.002)	(0.023)
$h = 2$			0.018***	0.018***	0.224***
			(0.002)	(0.002)	(0.025)
$h = 3$			0.019***	0.019***	0.236***
			(0.003)	(0.003)	(0.029)
Constant	0.098***	0.057***	0.086***	0.045***	
	(.)	(0.004)	(0.002)	(0.004)	
Observations	125,219	125,219	125,219	125,219	123,275
R-squared		0.044	0.001	0.045	
Number of id	294	294	294	294	

Notes: *FE* refers to individual forecaster fixed effects and *number of id* is the total number of distinct individual forecasters. For regressions 2 through 5, inflation and $h = 0$ are the excluded indicator variables. Regression 5 uses a conditional logit to allow for fixed effects. Robust standard errors are in parentheses, and *** p<0.01, ** p<0.05, * p<0.1.

Compared to the previously discussed stickiness estimates in the literature, these new estimates are relatively small. Recall that the smallest estimate in the literature for any variable or horizon is 0.157 (which also comes from micro-level estimation). My estimated unconditional average is 38% smaller than this literature minimum, and only two variables from my regressions, GDP and housing starts, have slightly larger values.

B. *State-Dependence Results*

I now transition from comparing my main estimates of stickiness with the literature's to discussing estimates of state-dependence. The standard sticky information model assumes a constant λ, however some rational inattention studies discussed in Section I document that information rigidities are correlated with aggregate economic activity. Specifically, Coibion and Gorodnichenko (2015) applies models from McConnell and Perez-Quiros (2000) and Gorodnichenko (2008) and find that the information rigidities of some variables have an inverse relationship with the volatility of GDP growth over time, and find these information rigidities decrease after a recession starts.

I do a similar exercise in a more micro-fashion; I test for state dependence by expanding equation 5 to include three time-varying control variables I *a priori* believe may affect a forecaster's inattention,

$$\lambda_{xith} = \boldsymbol{\alpha_x} + \boldsymbol{\alpha_i} + \boldsymbol{\alpha_h} + \boldsymbol{\beta controls_{xith}} + \varepsilon_{xith}, \tag{6}$$

where $\boldsymbol{controls_{xith}}$ is a vector containing the three control variables.

The first of these control variables is the standardized absolute value of an individual forecaster's forecast error for a variable from the previous period, $\widehat{|FE_{xit-1}|}$, where

$$\widehat{|FE_{xit-1}|} = \frac{\mathbf{E}\left[|FE_{xit-1}|\right] - |FE_{xit-1}|}{\sqrt{\mathbf{V}\left[|FE_{xit-1}|\right]}}, \tag{7}$$ [13]

and $|FE_{xit-1}|$ is the lagged absolute forecast error,

$$|FE_{xit-1}| = |f_{it-2}x_{t-1} - x_{t-1}|. \tag{8}$$ [14]

Similarly, the second control variable is the standardized average absolute value of forecast errors for a variable from the previous period, $\widehat{|FE_{xt-1}|}$, where

$$\widehat{|FE_{xt-1}|} = \frac{\mathbf{E}\left[|FE_{xt-1}|\right] - |FE_{xt-1}|}{\sqrt{\mathbf{V}\left[|FE_{xt-1}|\right]}}, \tag{9}$$

and $|FE_{xt-1}|$ is the lagged average absolute forecast error,

$$|FE_{xt-1}| = \frac{1}{n_{t-1}} \sum_{i=1}^{n_{t-1}} |FE_{xit-1}|, \tag{10}$$

and n_{t-1} is the number of forecasters last period.

The third and last control variable is the absolute standardized value of a variable from the previous period, $\widehat{|x_{t-1}|}$, where

$$\widehat{|x_{t-1}|} = \left| \frac{\mathbf{E}[x_{t-1}] - x_{t-1}}{\sqrt{\mathbf{V}[x_{t-1}]}} \right|. \tag{11}$$

The reasoning for including these three variables is as follows: if a forecaster's prediction was very different from the actual value last period, then that forecaster may attempt to gain more information than they did before to prevent another large error, suggesting that the coefficient on $\widehat{|FE_{xit-1}|}$

[13] $\mathbf{E}$ and $\mathbf{V}$ are the unconditional mean and variance operators, respectively. Here, unconditional means calculating the mean and variance using all individual forecasters and time periods in the sample.

[14] Squaring instead of taking the absolute value, for this and the other control variables, does not impact results.

should be negative. Similarly, if the (cross-sectional) average of the forecasters' (absolute) forecast errors was large last period, then individual forecasters may attempt to gain more information for a similar reason as the previous variable. This again suggests that the coefficient on $|\widehat{FE_{xt-1}}|$ should be negative, though perhaps less negative than the coefficient on $|\widehat{FE_{xit-1}}|$, because an individual likely cares more about their own errors than the group average errors. Finally, if the variable being forecast is very different from it's average value, such as low GDP growth potentially causing a recession or supply disruptions causing periods of high inflation, then forecasters may be more likely to acquire more information, since governments, businesses, and consumers care more about the forecasts of these variables during these time periods. Therefore the coefficient for $\widehat{|x_{t-1}|}$ should also be negative.

Results for this regression are shown in Table 3. In all specifications, the coefficients for $|\widehat{FE_{xit-1}}|$ and $|\widehat{FE_{xt-1}}|$ are both negative, as hypothesized. Also as predicted, the coefficient for $|\widehat{FE_{xit-1}}|$ is more negative than the coefficient for $|\widehat{FE_{xt-1}}|$. However, the coefficient for $\widehat{|x_{t-1}|}$ is not statistically significant once variable and horizon fixed effects are included, implying the overall state of the economy may not have as substantial an impact on forecaster inattention as previously thought.

My findings about the heterogeneity of stickiness parameters and their state-dependent nature provide new evidence against using the sticky information model as an explanation of non-FIRE behavior in forecasters. Because the estimates from the aggregate-level and micro-level regressions are so different and the micro-approach is directly addressing inattention while the aggregate approach is not, they imply that the standard aggregate regression to test for sticky information must be misspecified or have quantitatively significant omitted variable bias. Thus, while sticky information accounts for some of the predictability of forecast errors, some other factor must be at work to account for the large discrepancy in results.

C. Time Series Results

Having variables such as previous forecast errors predict sticky information is one form of state-dependence. Another form is if sticky information predicts itself. Analyzing forecaster inattention as a time series variable reveals interesting details about how inattention changes over time. For this section, I look at average inattention in a time period averaged over individuals and horizons: for each variable, this average stickiness is calculated as

$$\lambda_{xt} = \frac{1}{4n_t} \sum_{i=1}^{n_t} \sum_{h=1}^{4} \lambda_{xith}, \tag{12}$$

Table 3—Forecaster Inattention State-Dependence

	(1) FE	(2) FE	(3) FE	(4) FE Logit
GDP			-0.035***	-1.062***
			(0.004)	(0.091)
Unemployment			0.120***	1.310***
			(0.008)	(0.098)
Industrial Production			0.016***	0.265***
			(0.005)	(0.077)
Housing			0.108***	1.238***
			(0.011)	(0.113)
$h = 1$			0.012***	0.157***
			(0.002)	(0.023)
$h = 2$			0.018***	0.229***
			(0.002)	(0.025)
$h = 3$			0.019***	0.241***
			(0.003)	(0.029)
$\|\widehat{FE_{xit-1}}\|$		-0.009***	-0.008***	-0.121***
		(0.003)	(0.002)	(0.038)
$\|\widehat{FE_{xt-1}}\|$		-0.007***	-0.005**	-0.095***
		(0.002)	(0.002)	(0.034)
$\|\widehat{x_{t-1}}\|$		0.010***	-0.001	0.001
		(0.003)	(0.003)	(0.037)
Constant	0.098	0.091***	0.045***	
	(.)	(0.002)	(0.005)	
Observations	125,219	125,059	125,059	123,115
R-squared		0.002	0.047	
Number of id	294	294	294	

Notes: *FE* refers to individual forecaster fixed effects and *number of id* is the total number of distinct individual forecasters. For regressions 3 and 4, inflation and $h = 0$ are the excluded indicator variables. Regression 4 uses a conditional logit to allow for fixed effects. Robust standard errors are in parentheses, and *** p<0.01, ** p<0.05, * p<0.1.

where n_t^i is the number of forecasters in that quarter. I also analyze the unconditional average of the variables,

$$\lambda_t = \frac{1}{5}\sum_{x=1}^{5} \lambda_{xt}, \tag{13}$$

Figure 1 shows graphs of these time series.

The reason to aggregate sticky information here is threefold. First, I want to focus more on the general temporal patterns of inattention. By shutting down the heterogeneity of individuals and horizons, I get a broad look at how information stickiness changes over time. Second, because the data comes from an unbalanced panel, there is a nontrivial quantity of forecasters who either partake in the survey for only a small number of periods or for non-consecutive periods. This makes interpreting the results of individual time series difficult. Third, aggregation turns the *binary* variable λ_{xith} into *continuous* variables, λ_{xt} and λ_t. While not strictly necessary for analysis, working with sticky information as a continuous variable eases interpretation, allows for closer comparisons to previous results in the literature, and enables standard vector autoregression (VAR) and local projections estimations.

Figure 1—Forecaster Inattention Over Time

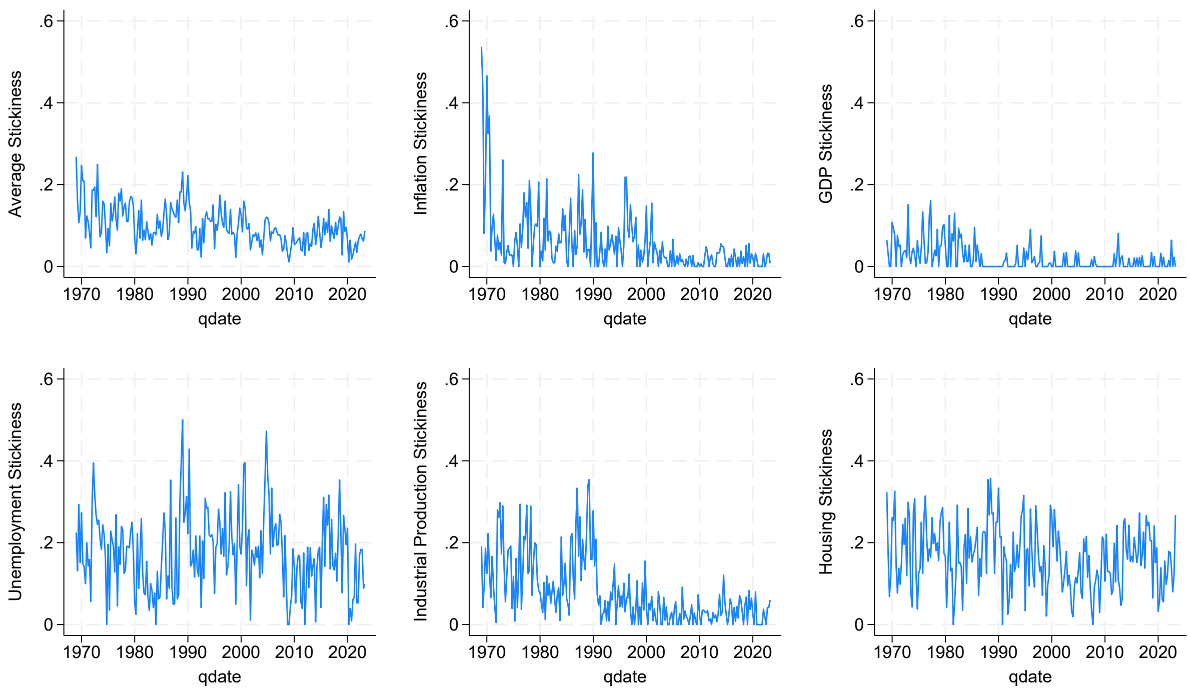

Notes: This figure plots, for each variable and their average, the average stickiness at each available point in time. *qdate* stands for *quarter-date*, because the data is at the quarterly frequency.

The first test I conduct on the time series of each variable is a unit-root test. Specifically, I conduct an Augmented Dickey-Fuller test with a constant (drift term) with 12 lags to determine if any of the variables contain a unit root. The specific regression for variable x is

$$\Delta\lambda_{xt} = \alpha + \beta\lambda_{xt-1} + \sum_{k=0}^{12} \zeta_k \Delta\lambda_{xt-k} + \varepsilon_{xt}. \tag{14}$$

The null hypothesis of the test is $\beta = 0$, i.e. the time series of the variable does contain a unit root, and the alternative hypothesis is $\beta < 0$, i.e. the time series of the variable does not contain a unit root. Table 4 shows the results of each test. All five variables reject the null of a unit root.[15]

Because I concluded that there are no unit-roots, I estimate multiple autoregressive (AR) pro-

[15]As discussed below, issues can arise because λ_{xt} has a lower and upper bound of zero and one, respectively. These bounds mean that the standard ADF regressions (and, later, the structural break tests) are misspecified, and the results may be incorrect. To my knowledge, there is not yet a standard way to test for unit roots in bounded time series data.

Table 4—Unit-Root Tests

	Test Statistics
Average	-2.360***
Inflation	-2.611***
GDP	-1.712**
Unemployment	-3.109***
Industrial Production	-1.946**
Housing	-3.517***

Notes: This table reports, for each variable and their average, the test statistic of an Augmented Dickey-Fuller (ADF) test. Each test has a drift term and 12 lags. *** p<0.01, ** p<0.05, * p<0.1.

cesses. I show lag selection results for autoregressive moving-average (ARMA) models and vector autoregression (VAR) models in Table 7 of Appendix VI.B, but only discuss AR(1) results in the main text for parsimony. Table 5 shows the results of the AR(1) estimations. Because there is the problem of the data having a lower-bound of zero and thus being zero-inflated, I estimate a Tobit (censored) AR(1) in addition to a regular AR(1).

Specifically, the AR(1) regression for variable x is

$$\lambda_{xt} = \alpha + \rho\lambda_{xt-1} + \varepsilon_{xt}, \tag{15}$$

while the Tobit AR(1) regression takes the AR(1) regression equation, interprets it as a latent regression model and λ_{xt} as a latent variable, and incorporates it into a lower- and upper-censored Tobit model,

$$\lambda^*_{xt} = \begin{cases} \lambda_{xt} & \text{if } 0 < \lambda_{xt} < 1 \\ 0 & \text{if } \lambda_{xt} \leq 0 \\ 1 & \text{if } \lambda_{xt} \geq 1 \end{cases}, \tag{16}$$

where λ^*_{xt} is the observed outcome.[16]

Regardless of the specification, the information rigidity of each variable displays significant persistence. However, by ignoring the zero lower bound, AR(1) estimates lower persistence values for all the variables except housing. The bound-incorporating Tobit AR(1) shows that persistence is actually higher. Furthermore, a joint test with the null hypothesis of all 5 of the persistence terms of the variables are equal is rejected at the 5% level for the AR(1) and at the 1% level for the Tobit AR(1). These results imply that information rigidity in the previous period may have an effect on

[16] $\lambda_{xt} = 1$ never occurs in this sample. The upper bound is included for completeness.

information rigidity today, and the strength of this effect varies depending on the economic variable in question.[17]

It is perhaps surprising that the tests indicate positive persistence. One could imagine that, *ceteris paribus*, if a forecaster has to sometimes be inattentive for some reason (e.g. there exist costs of acquiring information), they would want to alternate relatively quickly between being attentive and inattentive (paying and not paying the costs) so that they don't have several consecutive periods where they make poor forecasts.[18] Instead, positive persistence implies that, on average, forecasters "prefer" to have relatively long strings of good and bad forecasts. However, it is important to note that the estimated persistence terms, while positive, are fairly low, and average inattention displays significant fluctuations and jumps.

Table 5—AR(1) Estimations

	AR(1)		Tobit AR(1)	
	ρ	α	ρ	α
Average	0.510***	0.048***	0.510***	0.048***
	(0.055)	(0.006)	(0.055)	(0.006)
Inflation	0.453***	0.029***	0.495***	0.016**
	(0.107)	(0.005)	(0.111)	(0.007)
GDP	0.284***	0.014***	0.497***	-0.016***
	(0.093)	(0.002)	(0.144)	(0.006)
Unemployment	0.360***	0.112***	0.367***	0.110***
	(0.069)	(0.013)	(0.071)	(0.013)
Industrial Production	0.554***	0.034***	0.604***	0.024***
	(0.071)	(0.006)	(0.076)	(0.007)
Housing	0.305***	0.114***	0.308***	0.114***
	(0.063)	(0.012)	(0.064)	(0.012)

Notes: This table reports, for each variable and their average, the estimated persistence and constant terms of the AR(1) and Tobit AR(1) regressions. The Tobit AR(1) regressions are Tobit regressions where λ_{xt} is the dependent variable and λ_{xt-1} is the dependent variable. The lower- and upper-censoring limits are 0 and 1, respectively. Robust standard errors are in parentheses, and *** p<0.01, ** p<0.05, * p<0.1.

Now considering the long-run trend of inattention, Figure 1 visually shows that stickiness has decreased over time for inflation, GDP growth, industrial production, and the average. There are several periods where all forecasters update their forecasts (λ_{xt} hits its lower bound of zero), and the frequency of these periods has increased over time. The unemployment rate and housing starts, however, don't show a clear decline in stickiness. Furthermore, while not tested, the variances of sticky information for most variables appear to have declined over time as shown in Figure 2.

I formally test for declines with structural break tests. For each variable, I do a Bai and Perron (2003) sequential F-test for potentially multiple structural breaks at unknown dates with a sym-

[17] Or that the structural/latent cause of information rigidity in the previous period may have an effect on the structural/latent cause of information rigidity today.

[18] Assuming that having more relevant information makes forecasts better on average.

Figure 2—Forecaster Inattention Volatility Over Time

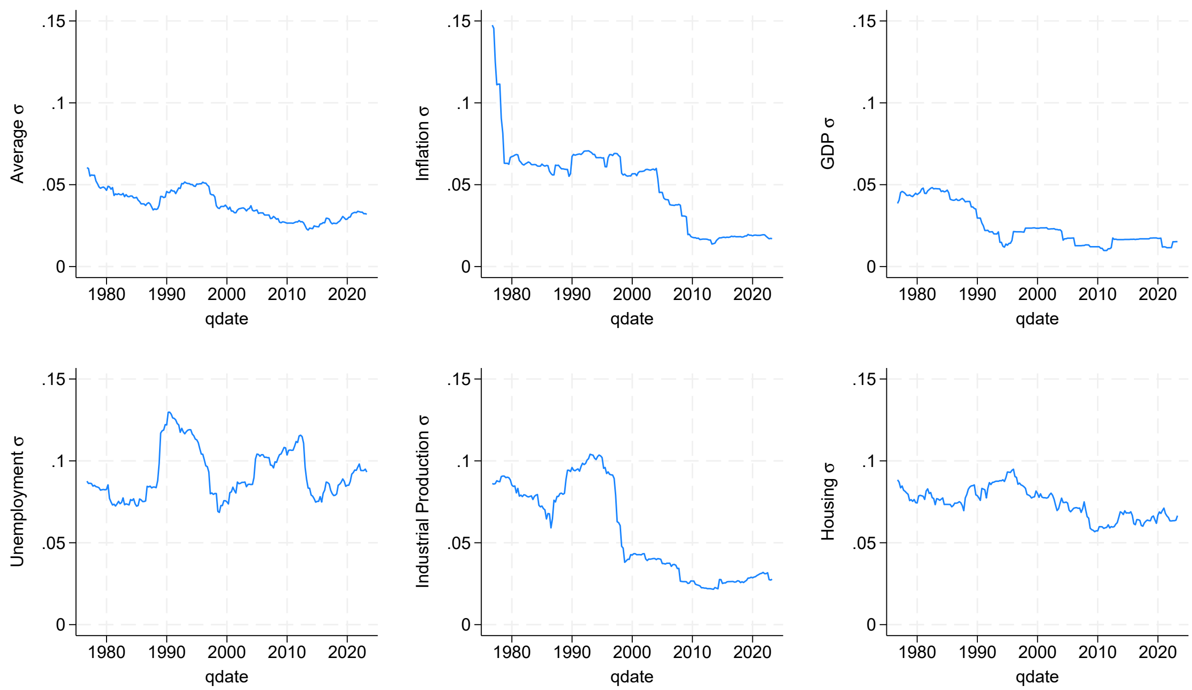

Notes: This figure plots, for each variable and their average, the rolling standard deviation of average stickiness at each available point in time; at time t, the standard deviation of average stickiness is calculated using only the values from $\{t-32, ..., t\}$. The first 32 time periods are not shown since their value would use fewer observations in their calculation than the other periods. *qdate* stands for *quarter-date*, because the data is at the quarterly frequency.

metric trimming of 15%. This will determine if any of the coefficients of the AR(1) processes, α and ρ, change, and if so at what dates these changes occurred. Table 6 shows the results of the tests for the variables with a break: the average, inflation, GDP, and industrial production. The first three are estimated to have only 1 break at 2000q4, 2001q1, and 1983q2, respectively. Industrial production has 2 breaks at 1990q3 and 2000q4. All the tests finding a break conclude that the only parameter that significantly changed is the constant term α. While the estimated value of the persistence term ρ is different before and after the breaks, the difference is not statistically significant. However, comparing the values of ρ in Table 5 and Table 6, I can see that ignoring the structural break caused the AR(1) to inflate its estimate of ρ.

These break dates approximately correspond to different recessions in US history. 1983q2 follows the two recessions of the early 1980s related to the oil crisis and Federal Reserve Chair Paul Volcker's monetary policy. Interestingly, this break is for GDP, not inflation. 2000q4 and 2001q1 are likely both related to the collapse of the dot-com bubble. Finally, 1990q3 is close to the 1990 oil price

Table 6—Structural Break Tests

	Break Dates	Cause	ρ	α
Average	2000q4	α	0.383, 0.356	0.072, 0.047
Inflation	2001q1	α	0.371, 0.075	0.050, 0.018
GDP	1983q2	α	-0.005, 0.063	0.049, 0.010
Industrial Production	1990q3, 2000q4	α	0.367, 0.242, 0.075	.056, 0.046, 0.018

Notes: This table reports, for each time series with a break, the estimated break dates from sequential F-tests for structural breaks at unknown dates, the parameter(s) that changed to cause the break (i.e. are statistically different before and after), and the values of the parameters before and after the breaks. The tests use a symmetric trimming of 15% of the observations from the start and end of each variable's time series.

shock.

Another potential reason for 1990q3 being a break date is the Federal Reserve Bank of Philadelphia took over administration of the survey in 1990q2. However, there is not a clear reason for this to be the cause of the break. Examples of the survey forms from 1981q2 and 1990q1, when the National Bureau of Economic Research administered the survey, are publicly available on the Federal Reserve Bank's website. Examples of the Federal Reserve Bank forms are available too; they have changed their survey form 8 times since they took over. The main difference between all the forms is mostly visual appearance, and not altered questions. The only reasonable thing that could explain the structural change is that the Federal Reserve Bank started reporting the values of the variables in the previous quarter in the survey. While perhaps a minor difference, this could constitute a non-negligible change in the structure of available information for the forecasters.

Finally, I test if a shock to the inattention of a variable has an effect on that variable over time. To do this I estimate the impulse response function of each variable due to a shock to that variable's information rigidity using the method of local projections from Jordà (2005). The Jordà method requires estimation of a series of regressions for each variable and horizon $h = 1, 2, ...H$,

$$x_{t+h-1} = \theta_h \lambda_{xt-1} + \sum_{k=2}^{L+1} x_{t+h-k} + \sum_{k=2}^{L+1} \lambda_{xt-k} + \varepsilon_{t+h-1}, \tag{17}$$

where θ_h is the parameter of interest and L is the number of past lags of the regression variables to act as controls.[19] There is a slight complication with the Jordà method, as it introduces serial correlation in the error terms caused by the successive leading of the dependent variable. I use the correction from Newey and West (1987) for my standard errors to account for this serial correlation.

The graphs for these impulse responses are shown in Figure 3.[20] The number of horizons used

[19] The use of horizons here is different from their use in the panel estimation.

[20] There is no statement of causality here, as there is likely endogeneity and omitted variable bias for which I can't control. Controlling for lags potentially helps but does not guarantee the removal of bias. This analysis only describes conditional

in the estimation is $H = 12$ and the lag length is $L = 2$. The inattention shocks have significant contemporaneous effects for all the variables except GDP. The effects persist for one additional quarter for the four affected variables before becoming insignificant. The only variable that is impacted for more than two quarters is unemployment, where the shock is felt for over a year.

Figure 3—Impulse Response Functions

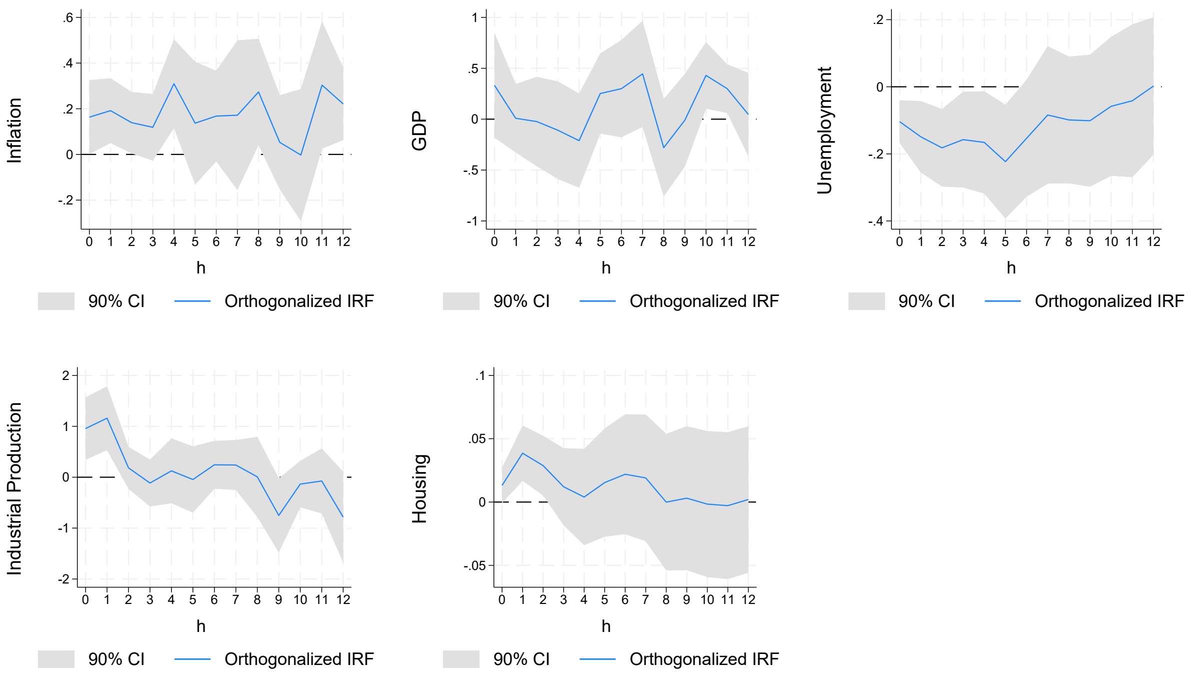

Notes: This figure plots, for each variable, the effect of a shock to a variable's information rigidity today on the variable today and on the future path (the next twelve quarters) of the variable assuming no other shocks. This effect is not causal and only describes correlations. The 90% confidence intervals are constructed from robust standard errors with the Newey-West correction.

V. Conclusion

This paper investigated the sticky information forecasting model and analyzed two aspects of it. The first aspect involved my new micro-estimates of the stickiness parameter being smaller than is needed to 1) explain the amount of information rigidity observed in the data, and 2) explain various monetary policy puzzles. Some other mechanism must be present and primary to explain these discrepancies. However, while the estimates I get are smaller than much of the literature's, they are still nonzero. This implies the sticky information model, while still potentially useful in

correlations over time.

certain macroeconomic models, may be less useful than currently believed.

This leads to the second aspect analyzed: when using a model with sticky information, a researcher should exercise caution when assuming the stickiness parameter is constant (state-independent). I find new evidence that periods of lower forecast errors, and thus potentially more stable economic environments, are accompanied with a higher degrees of information rigidity. As much of the literature has emphasized, this can have important ramifications for the macroeconomy. Potentially, this may be a cyclical dynamic. For example, low inflation and small forecast errors make agents pay less attention. This decrease in attention could cause agents to make decisions that lead to higher inflation and errors, which causes attention to increase, and ultimately causing inflation and errors to become low again.

VI. Appendix

A. *Growth Rate Transformations*

Following the SPF documentation, for inflation, real GDP growth, and industrial production growth, I focus on forecasts of the quarter-over-quarter annual growth rate calculated by the formula

$$g_{t+h|t-1} = 100\left(\left(\frac{X_{t+h|t-1}}{X_{t+h-1|t-1}}\right)^4 - 1\right), \quad h = 0, 1, ..., 4, \tag{18}$$

where $g_{t+h|t-1}$ is the forecast for quarter-over-quarter growth in period $t+h$ made on the basis of observations known through period $t-1$, and $X_{t+h|t-1}$ represents the forecast for the level. For inflation, X is the seasonally adjusted chain-weighted GDP price index; for real GDP growth, X is the seasonally adjusted chain-weighted real GDP; for industrial production growth, X is the seasonally adjusted index of industrial production.

There are two main reasons to focus on growth rates for these three variables. First, growth rates make the time series of the variables non-explosive.[21] Second, most of the literature uses growth rates, so I also use growth rates to make my results more comparable to theirs.

[21]Neither the unemployment rate nor housing starts are explosive; however there is still ongoing debate on whether they do or do not have a unit root. Since most of the literature focuses on the levels of these variables, I do also.

B. Lag Selection

An autoreggressive moving-average model with p autoregressive terms and q moving average terms, denoted ARMA(p,q), models the average information rigidity as

$$\lambda_{xt} = \alpha + \sum_{k=1}^{p} \rho_k \lambda_{xt-k} + \sum_{k=1}^{q} \theta_k \varepsilon_{t-k} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma^2). \tag{19}$$

An AR(p) model is an ARMA model with zero lagged error terms, i.e. an ARMA(p,0).

Table 7—ARMA and VAR Lag Selection

	AR(p)		ARMA(p,q)		VAR(p)	
	BIC	AIC	BIC	AIC	BIC	AIC
Average	4	4	(1,1)	(7,5)	NA	NA
Inflation	4	6	(1,1)	(8,5)	2	3
GDP	5	11	(1,3)	(7,7)	0	1
Unemployment	2	2	(1,1)	(1,1)	2	4
Industrial Production	3	15	(1,1)	(3,3)	3	4
Housing	1	1	(0,0)	(6,2)	2	5

Notes: This table shows, for each variable and their average, the number of lags recommended to use in a AR, ARMA, and VAR model based off of the Akaike information criterion (AIC) and the Bayesian information criterion (BIC). For ARMA estimation, the maximum AR and MA lag orders for which the information criteria are to be calculated is 8. For AR and VAR estimation, the maximum lag order is 8.

A vector autoregression model with p autoregressive terms, denoted VAR(p), models the average information rigidity and the value of the variable as

$$y_t = \alpha + \sum_{k=1}^{p} A_k y_{t-k} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \Omega), \tag{20}$$

where y_t is a 2×1 vector containing λ_{xt} and x_t, α is a 2×1 vector of constants, A_k is a 2×2 matrix of coefficients (for every $k = 0, ..., p$), ε_t is a 2×1 vector of error terms, and Ω is a 2×2 covariance matrix.

REFERENCES

Andrade, Philippe, and Hervé Le Bihan. 2013. "Inattentive Professional Forecasters." *Journal of Monetary Economics*, 60(8): 967–982.

Angeletos, George-Marios, and Zhen Huo. 2021. "Myopia and Anchoring." *American Economic Review*, 111(4): 1166–1200.

Bai, Jushan, and Pierre Perron. 2003. "Computation and Analysis of Multiple Structural Change Models." *Journal of Applied Econometrics*, 18(1): 1–22.

Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin. 2018. "The Design and Price of Information." *American Economic Review*, 108(1): 1–48.

Bordalo, Pedro, Nicola Gennalioli, Yueran Ma, and Andrei Shleifer. 2020. "Overreaction in Macroeconomic Expectations." *American Economic Review*, 110(9): 2748–2782.

Branch, William A. 2007. "Sticky Information and Model Uncertainty in Survey Data on Inflation Expectations." *Journal of Economic Dynamics and Control*, 31(1): 245–276.

Carrillo, Julio A. 2012. "How Well Does Sticky Information Explain the Dynamics of Inflation, Ouput, and Real Wages?" *Journal of Economic Dynamics and Control*, 36(6): 830–850.

Carroll, Christopher D., Edmund Crawley, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White. 2020. "Sticky Expectations and Consumption Dynamics." *American Economic Journal: Macroeconomics*, 12(3): 40–76.

Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia. 2017. "Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments." *American Economic Journal: Macroeconomics*, 9(3): 1–35.

Coibion, Olivier. 2010. "Testing the Sticky Information Phillips Curve." *Review of Economics and Statistics*, 92(1): 87–101.

Coibion, Olivier, and Yuriy Gorodnichenko. 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy*, 120(1): 116–159.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review*, 105(8): 2644–2678.

Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar. 2018. "How Do Firms Form Their Expectations? New Survey Evidence." *American Economic Review*, 108(9): 2671–2713.

Dupor, Bill, Tomiyuki Kitamura, and Takayuki Tsuruga. 2010. "Integrating Sticky Information and Sticky Prices." *The Review of Economics and Statistics*, 92(3): 657–669.

Gaus, Eric, and Arunima Sinha. 2018. "What does the yield curve imply about investor expectations?" *Journal of Macroeconomics*, 57: 248–265.

Gelain, Paolo, Nikolay Iskrev, Kevin J. Lansing, and Caterina Mendicino. 2019. "Inflation dynamics and adaptive expectations in an estimated DSGE model." *Journal of Macroeconomics*, 59: 258–277.

Giacomini, Raffaella, Vasiliki Skreta, and Javier Turen. 2020. "Heterogeneity, Inattention, and Bayesian Updates." *American Economic Journal: Macroeconomics*, 12(1): 282–309.

Gorodnichenko, Yuriy. 2008. "Endogenous Information, Menu Costs and Inflation Persistence." *National Bureau of Economic Research Working Paper 14184*.

Haubrich, Joseph G., George Pennacchi, and Peter H. Ritchken. 2008. "Estimating Real and Nominal Term Structures Using Treasury Yields, Inflation, Inflation Forecasts, and Inflation Swap Rates." *FRB of Cleveland Working Paper No. 08-10*.

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review*, 95(1): 161–182.

Khan, Hashmat, and Zhenhua Zhu. 2006. "Estimates of the Sticky-Information Phillips Curve for the United States." *Journal of Money, Credit, and Banking*, 38(1): 195–207.

Kiley, Michael T. 2007. "A Quantitative Comparison of Sticky Price and Sticky-Information Models of Price Setting." *Journal of Money, Credit, and Banking*, 39(1): 101–125.

Kim, Insu, and Young Se Kim. 2019. "Inattentive agents and inflation forecast error dynamics: A Bayesian DSGE approach." *Journal of Macroeconomics*, 62: 103–139.

Klenow, Peter J., and Jonathan L. Willis. 2007. "Sticky Information and Sticky Prices." *Journal of Monetary Economics*, 54: 79–99.

Knotek, Edward S. 2010. "A Tale of Two Rigidities: Sticky Prices in a Sticky Information Environment." *Journal of Money, Credit, and Banking*, 42: 1543–1564.

Kohlhas, Alexandre N., and Ansgar Walther. 2021. "Asymmetric Attention." *American Economic Review*, 111(9): 2879–2925.

Korenok, Oleg. 2008. "Empirical Comparison of Sticky Price and Sticky Information Models." *Journal of Macroeconomics*, 30(3): 906–927.

L'Huillier, Jean-Paul. 2020. "Consumer Imperfect Information and Endogenous Price Rigidity." *American Economic Journal: Macroeconomics*, 12(2): 94–123.

Lucas, Robert E. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory*, 4(2): 103–124.

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

Mankiw, N. Gregory, and Ricardo Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *The Quarterly Journal of Economics*, 117(4): 1295–1328.

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers. 2007. "Sticky Information in General Equilibrium." *Journal of the European Economic Association*, 5(2-3): 603–613.

McConnell, Margaret M., and Gabriel Perez-Quiros. 2000. "Output Fluctuations in the United States: What Has Changed since the Early 1980's?" *American Economic Review*, 90(5): 1464–1476.

Morales-Jiménez, Camilo. 2022. "The Cyclical Behavior of Unemployment and Wages under Information Frictions." *American Economic Journal: Macroeconomics*, 14(1): 301–331.

Morikawa, Masayuki. 2019. "Uncertainty over production forecasts: An empirical analysis using monthly quantitative survey data." *Journal of Macroeconomics*, 60: 163–179.

Morris, Stephen, and Hyun Song Shin. 2006. "Inertia of Forward-Looking Expectations." *American Economic Review*, 96(2): 152–157.

Newey, Whitney, and Kenneth West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55(3): 703–708.

Paciello, Luigi, and Mirko Wiederholt. 2014. "Exogenous Information, Endogenous Information, and Optimal Monetary Policy." *Review of Economic Studies*, 81(1): 356–388.

Reis, Ricardo. 2009. "Optimal Monetary Policy Rules in an Estimated Sticky-Information Model." *American economics Journal: Macroeconomics*, 1(2): 1–28.

Sarte, Pierre-Daniel. 2014. "When is Sticky Information More Information?" *Journal of Money, Credit, and Banking*, 46(7): 1345–1379.

Sims, Christopher A. 2003. "Implications of rational inattention." *Journal of Monetary Economics*, 50: 665–690.

Tortorice, Daniel L. 2018. "The business cycle implications of fluctuating long run expectations." *Journal of Macroeconomics*, 58: 266–291.