Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts

By OLIVIER COIBION and YURIY GORODNICHENKO*

We propose a new approach to test the full-information rational expectations hypothesis which can identify whether rejections of the null arise from information rigidities. This approach quantifies the economic significance of departures from the null and the underlying degree of information rigidity. Applying this approach to U.S. and international data of professional forecasters and other agents yields pervasive evidence consistent with the presence of information rigidities. These results therefore provide a set of stylized facts which can be used to calibrate imperfect information models. Finally, we document evidence of state-dependence in the expectations formation process. (*JEL* codes: E3, E4, E5)

---

* Coibion: Department of Economics, University of Texas-Austin, 2225 Speedway, BRB 3.156, C3100, Austin, TX 78712 (email: ocoibion@austin.utexas.edu); Gorodnichenko: Department of Economics, University of California-Berkeley, 530 Evans Hall #3880, Berkeley, CA 94720 (email: ygorodni@econ.berkeley.edu). We are grateful to Bob Archibald, Christopher Crowe, Zeno Enders, Ulrich Fritsche, Pierre-Olivier Gourinchas, Ed Knotek, Ricardo Reis, Javier Reyes, David Romer and Chris Sims for helpful comments as well as seminar participants at the Bank of France, CESifo Conference on Macroeconomics and Survey Data, College of William and Mary, Duke University, George Washington U., Columbia, IMF, Minneapolis Fed, Michigan, Michigan State, NBER Summer Institute, NC State, Notre Dame, Richmond Fed, U. of Arkansas, UC Berkeley, UC Irvine, UNC, U. of Richmond, UT Austin, U. of Virginia, and Yale. We thank Yury Yatsynovich for excellent research assistance. Part of this work was done while Coibion was a visiting fellow at the IMF, whose support was greatly appreciated. Gorodnichenko thanks the NSF and the Sloan Foundation for financial support. The views expressed in this paper are those of the authors and do not necessarily represent those of the IMF or any other institution. We have no relevant or material financial interests that relate to the research.

Expectations matter. How much to consume or save, what price to set, and whether to hire or fire workers are just some of the fundamental decisions underlying macroeconomic dynamics that hinge upon agents' expectations of the future. Yet how those expectations are formed, and how best to model this process, remains an open question. From the simple automatons of adaptive expectations to the all-knowing agents of modern full-information rational expectations models, macroeconomists have considered a wide variety of frameworks to model the expectations formation process, yielding radically different results for macroeconomic dynamics and policy implications. Recent work on rational expectations models with information frictions such as Mankiw and Reis (2002), Woodford (2002), and Sims (2003) has emphasized how information rigidities can account for otherwise puzzling empirical findings but these same frictions can also lead to policy prescriptions that differ from those under models with full information.[1] Despite a growing body of work studying the implications of possible departures from full-information rational expectations, the empirical evidence against this assumption underlying most modern macroeconomic models has been limited. In particular, while statistical evidence against the null is commonly uncovered, the economic significance of these rejections remains unclear.

Building from the predictions of rational expectations models with information rigidities, we propose a novel approach to test the null of full-

[1] For example, Ball, Mankiw and Reis (2005) show that price-level targeting is optimal in sticky-information models whereas inflation targeting is optimal in a sticky-price model. Paciello and Wiederholt (2014) document how rational inattention as in Sims (2003) alters optimal monetary policy. Mankiw and Reis (2002) argue that the observed delayed response of inflation to monetary policy shocks is not readily matched by New Keynesian models without the addition of information rigidities or the counterfactual assumption of price indexation. Roberts (1997, 1998) and Adam and Padula (2011) demonstrate that empirical estimates of the slope of the New Keynesian Phillips Curve have the correct sign when conditioning on survey measures of inflation expectations while this is typically not the case under the assumption of full-information rational expectations. Gourinchas and Tornell (2004), Piazzesi and Schneider (2008), and Bachetta, Mertens and van Wincoop (2009) all identify links between systematic forecast errors in survey forecasts and puzzles in various financial markets.

information rational expectations in a way that sheds new light on possible departures from the null. Our baseline specification relates ex-post mean forecast errors to the ex-ante revisions in the average forecast across agents and possesses multiple advantages over traditional tests of full-information rational expectations (FIRE). First, we rely on the predictions of theoretical models of information rigidities to guide our choice of the relevant regressors. Second, models of information rigidities make specific predictions about the sign of the coefficient on forecast revisions, so that our specification provides guidance not only about the null of FIRE but also about alternative models. As a result, our framework can help determine whether rejections of the null should be interpreted as rejecting either the rationality of expectations or the full-information assumption. Third, we show that the coefficient on forecast revisions maps one-to-one into the underlying degree of information rigidity and therefore our approach can provide a metric by which to assess the economic significance of departures from the null of FIRE.

Two theoretical rational expectations models of information frictions motivate our empirical specification. In the sticky-information model of Mankiw and Reis (2002), agents update their information sets infrequently as a result of fixed costs to the acquisition of information. The degree of information rigidity in this model is then the probability of not acquiring new information each period. The second class of models we consider consists of noisy-information models such as Woodford (2002), Sims (2003), and Mackowiak and Wiederholt (2009). Here, agents continuously update their information sets but, because they can never fully observe the true state, they form and update beliefs about the underlying fundamentals via a signal extraction problem. Forecasts are a weighted average of agents' prior beliefs and the new information received, where the weight on prior beliefs can be interpreted as the degree of information rigidity. Strikingly, both models predict the same relationship between the average ex-post forecast errors across agents and the average ex-ante forecast revision such that the coefficient on forecast revisions depends only on the degree of information rigidity in each model. This predictability of the average forecast error across agents from forecast revisions is an emergent property in both models, i.e. a property which arises only from the aggregation process and not at the individual level.

The resulting empirical specification can be applied to study information rigidities for a variety of economic agents such as consumers, firms, and financial market participants for whom forecast data are available. As a first step, we focus on inflation forecasts from the U.S. Survey of Professional Forecasters (SPF) for two reasons. First, inflation forecasts have received the most attention in the literature so that these results are more readily comparable to previous work. Second, because professional forecasters are some of the most informed economic agents, they can provide a conservative benchmark for assessing potential deviations from full-information rational expectations. From 1969-2014, we can strongly reject the null of FIRE and find that the estimated coefficient on forecast revisions is positive, consistent with the prediction of rational expectations models incorporating information rigidities. Additional coefficient restrictions implied by these models cannot be rejected and past information incorporated in other economic variables loses much of its predictive power for ex-post mean forecast errors once we control for the forecast revision. This indicates that rejections of the null are unlikely to be driven by departures from rationality (such as adaptive expectations) and instead reflect deviations from the assumption of full-information. Furthermore, the implied degree of information rigidity is high: in the context of sticky-information models, it implies an average duration of six to seven months between information updates, while in noisy-information models it implies that new information receives less than half of the weight that it would under full-information relative to prior beliefs.

In addition, we document that qualitatively similar results obtain for different kinds of economic agents, such as academics, commercial banks, and non-financial businesses, as well as for consumers and financial market-based inflation expectations. This implies that information rigidities are present not just amongst professional forecasters but also for firms and consumers. Given that the estimated degree of information rigidity in inflation forecasts is relatively high across different types of economic agents, information rigidities are likely to play a pervasive role in macroeconomic dynamics. The prevalence of information rigidities across agents also suggests that the estimated levels of information rigidities are unlikely to be driven by either strategic behavior on the part of professional forecasters or reputational considerations. As a result, our empirical estimates provide a new set of stylized facts which can be used for the calibration of models with information rigidities.

To further verify that our results are indeed driven by information rigidities, we derive testable predictions from a number of competing hypotheses which could potentially account for the predictability of forecast errors. For example, if forecasters are heterogeneous in the degree of loss-aversion with respect to their forecast errors, then predictability in forecast errors can arise even in the absence of information rigidities as in Capistran and Timmermann (2009). However, we show that such a model would imply a negative correlation—rather than positive as observed in the data—between ex-post forecast errors and ex-ante forecast revisions. We similarly derive testable predictions from models in which agents place a different weight on new information, as they would if they held different beliefs about underlying parameter values, or hold heterogeneous views about long-run means of macroeconomic variables. In each case, we find that these models yield counterfactual predictions about the predictability of forecast errors.

Using professional forecasts for a number of additional macroeconomic variables, both in the U.S. and across eleven additional countries, we provide further evidence of pervasive information rigidity. First, pooled estimates across macroeconomic variables confirm the finding of predictability of forecast errors coming from ex-ante forecast revisions with the signs predicted by models of information rigidities. Across datasets, we also find robust evidence that the degree of information rigidity varies systematically across macroeconomic variables and that this cross-sectional variation is consistent with the theoretically predicted determinants of noisy-information models: the persistence of a variable and the signal-noise ratio can account for about 15-30 percent of the variation in the estimated degree of information rigidity across countries and macroeconomic variables in the Consensus Economics dataset. Since the canonical sticky-information model assumes a common rate of information updating across variables, these results suggest that subsequent work with the sticky-information model should explore how such heterogeneity can arise in the context of infrequent information updating.

Because our empirical specification allows us to recover estimates of the underlying degree of information rigidity, we can also characterize whether the degree of information rigidity varies in response to economic conditions, as the incentives for agents to collect and process additional information change. For example, macroeconomic volatility declined significantly after the early to mid-1980s during the Great Moderation. According to models with information rigidities, such a decline in volatility should result in a higher degree of inattention. We study the low-frequency time variation in the estimated degree of information rigidity among U.S. professional forecasters and find evidence that accords remarkably well with this intuition: the degree of information rigidity fell consistently throughout the 1970s and early 1980s when macroeconomic volatility was high, reaching a minimum in the early 1980s. The degree of information rigidity subsequently rose over the course of the Great Moderation, as macroeconomic volatility declined. We also document higher frequency endogenous variation in information rigidities. For example, the degree of information rigidity declines significantly during U.S. recessions, which points to state-dependence in the expectations formation process. Hence, agents appear to adjust the resources devoted to the collection and processing of information in response to economic conditions, consistent with state-dependence in the information updating process as in Gorodnichenko (2008) and Mackowiak and Wiederholt (2012).

This paper is closely related to recent empirical work trying to ascertain the nature of the expectations formation process. For example, Mankiw, Reis and Wolfers (2004) assess whether a sticky-information model can replicate some stylized facts about the predictability of forecast errors by professional forecasters while Andolfatto, Hendry and Moran (2007) consider whether noisy-information with respect to the inflation target of the central bank can account for observed deviations from FIRE. Khan and Zhu (2006), Kiley (2007), and Coibion (2010) assess the validity of sticky-information using estimates of its predicted Phillips curve. One advantage of our approach is that we can directly recover an estimate of the degree of information rigidity without having to make auxiliary assumptions about the model, such as the nature of price-setting decisions. Sarte (2013) also uses surveys to quantify sticky-information in firm forecasts, but our approach allows us to assess both sticky-information and noisy-information models. Coibion and Gorodnichenko (2012) study the evidence for sticky-information and noisy-information models but do so by estimating the response of forecast errors and disagreement to structural shocks whereas our approach does not require the identification of any shock. In the same spirit, Andrade and LeBihan (2013) provide evidence for both sticky and noisy-information in the European Survey of Professional forecasters, Branch (2007) compares the fit of sticky-information and model-switching characterizations of the expectations formation process while Carroll (2003) tests an epidemiological model of expectations in which information diffuses over time from professional forecasters to consumers. However, these papers focus almost exclusively on inflationary expectations whereas we utilize forecasts for a wide variety of macroeconomic variables as well as cross-country data, allowing us to more fully characterize the nature and quantitative importance of information rigidities faced by economic agents. Finally, our paper is closely related to the long literature on the rationality of both individual and consensus forecasts. In this context, our results using models of information rigidities provide a new rationalization for the otherwise puzzlingly weak evidence against rational expectations found at the individual level relative to that observed in average forecasts (Pesaran and Weale 2006). If agents form their expectations rationally subject to information frictions, predictability in forecast errors will follow from the aggregation of forecasts across agents, even if no such predictability exists at the individual level.

The paper is structured as follows. Section I presents the predicted relationship between ex-post mean forecast errors and ex-ante mean forecast revisions in sticky-information and noisy-information models, baseline results from professionals' forecasts of inflation, as well as tests of competing explanations. Section II expands the set of forecasts used to different kinds of agents, macroeconomic variables and countries. Section III presents evidence on the extent to which the degree of information rigidity varies in response to low-frequency and business-cycle-frequency changes in macroeconomic conditions. Section IV concludes.

I. Forecast Errors, Forecast Revisions and Information Rigidities

In this section, we present two models of information rigidities and derive their respective predictions for the relationship between ex-post mean forecast errors and ex-ante mean forecast revisions. We document evidence consistent with these predictions using U.S. inflation forecasts of professional forecasters and argue that alternative explanations are unlikely to account for these findings.

A. Sticky-Information Model

Mankiw and Reis (2002) propose a model of inattentive agents who update their information sets each period with probability $(1-\lambda)$ but acquire no new information with probability λ, so that λ can be interpreted as the degree of information rigidity and $1/(1-\lambda)$ is the average duration between information updates. When agents update their information sets, they acquire full-information and have rational expectations. Reis (2006) shows how this time-dependent updating of information sets can occur when firms face a fixed cost to updating their information. The average time *t* forecast across agents (F_t) of a variable x at time $t+h$ is a weighted average of current and past full-information rational expectations forecasts (E_{t-j}) of the variable being forecasted such that

(1) $$F_t x_{t+h} = (1-\lambda)\sum_{j=0}^{\infty} \lambda^j E_{t-j} x_{t+h}.$$

The average forecast at time $t-1$ can similarly be written as

(2) $$F_{t-1} x_{t+h} = (1-\lambda)\sum_{j=0}^{\infty} \lambda^j E_{t-1-j} x_{t+h}$$

which implies that the current average forecast is just a weighted average of the previous period's average forecast and the current rational expectation of variable x at time $t+h$

(3) $$F_t x_{t+h} = (1-\lambda)E_t x_{t+h} + \lambda F_{t-1} x_{t+h}.$$

Full-information rational expectations are such that

(4) $$E_t x_{t+h} = x_{t+h} - v_{t+h,t}$$

where $v_{t+h,t}$ is the full-information rational expectations error and is thus uncorrelated with information dated *t* or earlier.

Combining (3) and (4) yields the predicted relationship between the ex-post mean forecast error across agents and the ex-ante mean forecast revision

(5) $$x_{t+h} - F_t x_{t+h} = \frac{\lambda}{1-\lambda}(F_t x_{t+h} - F_{t-1} x_{t+h}) + v_{t+h,t}.$$

Importantly, the coefficient on the forecast revision depends only on the degree of information rigidity λ. In the special case of no information frictions, $\lambda = 0$ and the specification collapses to equation (4), i.e. the average forecast error is unpredictable using information dated t or earlier. This predictability in forecast errors reflects the slow updating of information by some agents ($\lambda > 0$). While those who update their information after a shock move immediately to the full-information rational expectations belief, other agents do not change their information at all. This anchors the mean forecast to the previous period's, leading to a gradual adjustment of mean forecasts and predictability of average forecast errors. Because this canonical sticky-information model implies a single rate of information acquisition, equation (5) holds for any macroeconomic variable and any forecasting horizon, including horizons of multiple periods. In addition, this specification holds regardless of the structure of the rest of the model.

B. Noisy-Information Model

We also consider models in which agents know the structure of the model and underlying parameter values, continuously update their information sets, but never fully observe the state. This class of models includes most famously the Lucas (1972) islands model but also a wide variety of limited information settings considered in the literature. For example, Kydland and Prescott (1982) assume that the level of technology reflects both permanent and transitory shocks but that agents cannot separately identify these two components. More recently, Woodford (2002) considers an environment in which firms observe aggregate demand subject to idiosyncratic errors, which combined with strategic complementarity in price-setting, can account for the persistent effect of monetary policy shocks. Suppose that a macroeconomic variable follows an AR(1) process:[2]

(6) $$x_t = \rho x_{t-1} + v_t, \qquad 0 \le \rho \le 1,$$

where v_t is an i.i.d. normally distributed innovation to x_t. Agents cannot directly observe x_t but instead receive a signal y_{it} such that

(7) $$y_{it} = x_t + \omega_{it}$$

[2] We consider more general data-generating processes in sections I.D.1 and II.B but focus here on the simpler case for analytical tractability.

where ω_{it} represents normally distributed mean-zero noise which is i.i.d. across time and across agents. Each agent i then generates forecasts (conditional expectations) $F_{it}x_{t+h}$ given their information sets via the Kalman filter

(8) $$F_{it}x_t = Gy_{it} + (1-G)F_{it-1}x_t,$$
$$F_{it}x_{t+h} = \rho^h F_{it}x_t,$$

where G is the Kalman gain which represents the relative weight placed on new information relative to previous forecasts. When the signal is perfectly revealing about the true state, $G = 1$, while the presence of noise induces $G < 1$. Thus, $(1-G)$ can be interpreted as the degree of information rigidity in this model.[3] The Kalman gain also corresponds to the average reduction in the variance of contemporaneous forecast errors relative to the variance of one-step ahead forecast errors.

After averaging across agents and rearranging, the following relationship between ex-post mean forecast errors and ex-ante mean forecast revisions holds:

(9) $$x_{t+h} - F_t x_{t+h} = \frac{1-G}{G}(F_t x_{t+h} - F_{t-1}x_{t+h}) + v_{t+h,t}$$

where $v_{t+h,t} = \sum_{j=1}^{h}\rho^{h-j}v_{t+j}$ is the rational expectations error and F_t denotes the average forecast across agents at time t.[4] Thus, while individuals form their forecasts rationally conditional on their information set, the ex-post mean forecast error across agents is systematically predictable using ex-ante mean forecast revisions. The predictability of average forecast errors in the noisy-information model reflects the gradual adjustment of beliefs by all agents to new information. Because agents do not know whether the new information reflects noise or

[3] Crowe (2010) explores alternative information structures (e.g., include public signals and allow forecasters to observe past average forecasts).

[4] The presence of common noise would introduce another component to the error term, dated time t and uncorrelated with information from $t-1$ and earlier. In this case, our baseline empirical specification cannot be estimated by OLS. However, one can show that the bias in OLS will be downward (see Appendix A), such that our estimates will present *lower* bounds on the degree of information rigidity.

innovations to the variable being forecasted, they adjust their beliefs only gradually in response to shocks to fundamentals. This makes average forecast errors predictable with respect to ex-ante average forecast revisions. This specification is identical to equation (5) from the sticky-information model, when $(1 - G)$ is interpreted as the degree of information rigidity, and applies for any forecast horizon h or forecasts over multiple horizons. In contrast to equation (5) under sticky-information, the coefficient on forecast revisions need not be the same for different macroeconomic variables. Instead, the coefficient will vary with the determinants of the Kalman gain, e.g. the persistence of the series and the signal-noise ratio.

C. A New Approach for Assessing the Nature of the Expectations Formation Process

The sticky-information and noisy-information models both point to the same relationship between ex-post mean forecast errors and ex-ante mean forecast revisions such that the coefficient on forecast revisions maps one to one into the underlying degree of information rigidities. This relationship can be readily estimated for a given macroeconomic variable x, mean forecasts across agents Fx and forecasting horizon h using the following empirical specification:

(10) $$x_{t+h} - F_t x_{t+h} = c + \beta(F_t x_{t+h} - F_{t-1} x_{t+h}) + error_t$$

While this is just a special case of the more general test of FIRE commonly employed in the literature in which the forecast error is regressed on a subset of the information available to agents at the time the forecast was made, it addresses several important shortcomings of traditional tests. First, the relevant regressor to use for testing the predictability of forecast errors is specified by the theory. Second, when traditional tests identify a rejection of the null hypothesis of FIRE, this rejection is not directly informative about other theories of the expectations formation process in the absence of a clear theoretical mapping from the theory to the empirical tests. In contrast, our specification is informative not just about the null hypothesis of FIRE but also about models with information rigidities. Third, statistical rejections of the null hypothesis of FIRE in the standard test do not directly address the economic significance of departures from FIRE. Specification (10), on the other hand, allows us to map estimates of β directly into the underlying degrees of information rigidity (λ under sticky-information and $1-G$ under noisy-information) and, hence, can help assess the economic significance of any rejections of the null hypothesis of FIRE.

As a first step to applying our approach, we follow much of the literature on survey measures of expectations and focus on historical forecasts of U.S. annual inflation from the Survey of Professional Forecasters (SPF). Inflation expectations have received disproportionate attention because of their importance in measuring ex-ante real interest rates, their role in expectational Phillips curves, and for monetary policy. Professional forecasts, while not typically included in macroeconomic models, are useful not only due to their historical availability but also because, as some of the most informed economic agents in the economy, they provide a conservative benchmark for assessing possible deviations from the null of FIRE. The SPF is a quarterly survey of approximately 30-40 professional forecasters currently run by the Philadelphia Fed. GDP/GNP deflator inflation forecasts are available starting in 1968Q4 at horizons ranging from the current quarter to four quarters ahead. We focus for now on forecasts of year-on-year annual inflation, where e.g. $\pi_{t+3,t}$ refers to the average inflation rate over the current (t) and next three quarters.[5] Forecast errors are constructed using forecasts made at time t and real-time data available one year after the period being forecasted over. We use real-time data to measure ex-post variables because final data may reflect reclassifications and redefinitions such that the final values are not directly comparable to the historical forecasts made by agents (Croushore 2010).

[5] There are missing values in four-quarter ahead forecasts in 1969Q1-Q3, 1970Q1, and 1974Q4. We treat these periods as having missing values for estimates using year-ahead inflation forecasts.

Because the predictability of ex-post forecast errors from ex-ante forecast revisions in the sticky-information and noisy-information models obtains only when averaging across agents, we focus on mean forecasts across professional forecasters.[6] It is important to emphasize that the predicted relationship between forecast errors and forecast revisions should not be expected at the individual level. Under sticky-information, agents either do not update their information, and hence do not revise their forecasts, or else update their forecasts to the FIRE, in which case their forecast errors are uncorrelated with their forecast revision. Since the decision to update information is time-dependent and therefore orthogonal to current economic conditions, there should be no predictability in an individual's forecast errors arising from their forecast revisions. Under noisy-information, the use of the Kalman gain by agents ensures that individual forecast errors should also be unpredictable on average given the agent's information set, since the latter includes their forecast revision. Hence, the predictability of the average ex-post forecast errors across agents using ex-ante forecast revisions is an emergent property of the aggregation across individuals, not a property of the individual forecasts.

The relationship between average year-ahead inflation forecast errors across agents and average forecast revisions in both sticky-information and noisy-information models can be expressed as

$$(11) \qquad \pi_{t+3,t} - F_t\pi_{t+3,t} = c + \beta\left(F_t\pi_{t+3,t} - F_{t-1}\pi_{t+3,t}\right) + error_t$$

where $\beta > 0$ if information rigidities are present and $error_t$ is the rational expectations error which is orthogonal to information dated t and earlier, so equation (11) can be estimated by OLS. From 1969-2014, we find $\hat{\beta} = 1.19\ (\text{s.e.} = 0.50)$ as shown in Panel B of Table 1. As a result, we reject the null

[6] The mean forecast in the SPF can also change because of variation in the composition of participants over time. We find almost identical results when we use mean forecasts and forecast revisions constructed only from those forecasters participating in the two adjacent surveys used to construct forecast revisions each quarter.

of FIRE at the 5 percent level of statistical significance in a manner that is directly informative about the expectations formation process. First, the rejection of the null goes exactly in the direction predicted by models of information rigidities, so that this finding presents direct evidence in favor of these models. Second, because β maps into the degree of information rigidity from each model, we can extract an estimate of information frictions. In the context of sticky-information models, $\hat{\lambda} = \hat{\beta}/(1+\hat{\beta}) \approx 0.54$ would imply that agents update their information sets every six to seven months on average. This magnitude of sticky-information should significantly affect macroeconomic dynamics and optimal policy decisions, as documented in Reis (2009). Alternatively, one can interpret this estimate of β under noisy-information models as implying that agents put a weight of less than one-half on new information and more than one-half on their previous forecasts (i.e., $\hat{G} = 1/(1+\hat{\beta}) \approx 0.46$). This is in line with the rational inattention model of Mackowiak and Wiederholt (2009) in which such magnitudes of inattention have profound macroeconomic effects.[7] Thus, our approach implies that information frictions are economically and statistically significant.

We can also test theoretical restrictions implied by these models. For example, both sticky-information and noisy-information models predict a constant of zero in equation (11), which we cannot reject in the data. If we estimate equation (11) omitting the constant, the estimated value of β and standard errors are essentially unchanged. Second, these models predict that the coefficients on the contemporaneous forecast and on the lagged forecast are equal in absolute

[7] Mackowiak and Wiederholt (2009) derive a relationship between the average variance of contemporaneous forecast errors for a variable, its persistence, volatility and the optimal amount of attention devoted to that variable in a rational inattention setting. Using the corresponding values for U.S. inflation and SPF inflation forecasts, their approach yields an implied Kalman gain of 0.42 and therefore a degree of information rigidity of 0.58. See Appendix H. Also note that the implied degrees of information rigidity are lower than those found in Coibion and Gorodnichenko (2012). This result could reflect the presence of common noise and therefore of a downward bias in estimates of information rigidity in our present approach.

value. To implement this additional test, we decompose the forecast revision into two terms as follows

(12) $\pi_{t+3,t} - F_t\pi_{t+3,t} = c + \beta_1 F_t\pi_{t+3,t} + \beta_2 F_{t-1}\pi_{t+3,t} + error_t$.

Under models of information rigidities, we expect $\beta_1 > 0$, $\beta_2 < 0$, and $\beta_1 + \beta_2 = 0$. Estimating equation (12) from 1969-2014, we find $\hat{\beta}_1 = 1.21\ (\mathrm{s.e.} = 0.50)$ and $\hat{\beta}_2 = -1.23\ (\mathrm{s.e.} = 0.50)$. The signs on both coefficients conform to the theoretical predictions of models of information rigidities, and we cannot reject the null that the sum of the two coefficients is equal to zero (p-value=0.62). The results thus provide additional evidence consistent with the notion that the expectations formation process of professional forecasters is subject to information constraints.

A third restriction from models of information rigidities is that, while the average ex-post forecast error should be predictable using ex-ante average forecast revisions across agents, no other variable should have any additional predictive power for forecast errors. This is in the same spirit as traditional tests of FIRE but now conditional on forecast revisions. To assess this prediction, we focus on four specific macroeconomic variables which previous work (e.g. Mankiw, Reis and Wolfers 2004, Pesaran and Weale 2006) has identified as having significant predictive power for ex-post inflation forecast errors: lagged annual inflation ($\pi_{t-1,t-4}$), lagged quarterly interest rates (3-month Tbills), lagged quarterly changes in oil prices (WTI spot price), and the lagged quarterly unemployment rate. As a first step, we present estimates of the traditional test, i.e.

(13) $\pi_{t+3,t} - F_t\pi_{t+3,t} = c + \gamma F_t\pi_{t+3,t} + \delta z_{t-1} + error_t$.

in which ex-post forecast errors are regressed on observable variables. We first include the contemporaneous forecast of inflation as a RHS variable in equation (13) then augment this with each of the additional variables discussed. The results, presented in Panel A of Table 1, confirm that these variables have predictive power for average ex-post inflation forecast errors over this time sample: lagged inflation, interest rates, changes in oil prices and unemployment rates all have statistically significant coefficients pointing to predictability of ex-post forecast errors.

We then assess whether these variables retain their predictive power for ex-post forecast errors after conditioning on forecast revisions, i.e. we estimate equation (11) augmented with each of these variables. When controlling for either inflation, interest rates, or changes in oil prices, the coefficient on forecast revisions is qualitatively unchanged, while the coefficients on these additional variables are no longer statistically different from zero. Hence, once one controls for forecast revisions, the predictive power of these three variables is eliminated, as predicted by models of information rigidities. In the case of unemployment, however, there is additional predictive power even after controlling for forecast revisions, although the coefficient on the unemployment rate is cut by approximately 40 percent. This finding suggests that deviations from FIRE may exist above and beyond those captured by simple models of information rigidities and further exploration of these deviations is a fruitful avenue for future research.

[INSERT TABLE 1 HERE]

It should be emphasized that the degree of information rigidity, like the degree of nominal rigidity in typical New Keynesian models, is not a structural parameter. Rather, it should depend on underlying economic conditions. Reis (2006), for example, shows that the rate of information updating in sticky-information models depends on the volatility of macroeconomic variables. The same result applies with respect to the Kalman gain in noisy-information models. As a result, we investigate in sections II and III the extent to which the degree of information rigidity varies cross-sectionally as well as over time. As shown in Appendix G, if the underlying degree of information rigidity varies over the sample, then our baseline procedure produces a weighted average of underlying degrees of information rigidity, where the weights assigned to different periods or variables reflect the relative variances in their forecast revisions. Our baseline estimates should therefore be interpreted as capturing an average degree of information rigidity over this time period, which provides a useful benchmark both for the calibration of economic models as well as for subsequently assessing the extent of cross-sectional and time-variation in information rigidity.

Another noteworthy feature of these results is that the predictability in forecast errors holds over a period of forty-five years. As discussed in Croushore (2010), previous work has documented that rejections of the null of full-information rational expectations are much more common over short samples in which specific episodes, such as the Volcker disinflation, can have a disproportionate influence on measuring the predictability of forecast errors. The predictability of ex-post forecast errors from ex-ante forecast revisions over the longer sample of 1969 to 2014 is therefore particularly striking.

Finally, the predictability of average ex-post forecast errors across agents from ex-ante forecast revisions should not be interpreted as a form of irrationality on the part of agents. Both sticky-information and noisy-information models have, as their foundation, agents forming rational expectations subject to information constraints (e.g. Reis 2006, Sims 2003, Mackowiak and Wiederholt 2009). The predictability of average ex-post forecast errors that results from the aggregation process across agents is an emergent property that does not obtain at the individual level. Indeed, previous tests of the rationality of forecasts have commonly found much weaker rejections, if any, of the null of full-information rational expectations at the individual level relative to the average forecast (Pesaran and Weale 2006). This is precisely what one would expect from models of information rigidities.

D. Extensions and Alternative Interpretations

We consider a number of extensions of models with information rigidities to assess whether these qualitatively and quantitatively affect our baseline predictions. Specifically, we first extend the noisy-information model along three dimensions: a more general process for the variable being forecasted, heterogeneity in priors about long-run means, and heterogeneity in signal strength. In addition, we consider alternative explanations proposed in the literature to account for the predictability of forecast errors that do not appeal to information rigidity: heterogeneity in loss aversion and forecast smoothing on the part of professional forecasters. None of these extensions and alternative explanations finds support in the data, as also found in the different framework of Coibion and Gorodnichenko (2012).

D.1 Generalized Noisy-Information Model

While the predictability of average ex-post forecast errors from forecast revisions in the sticky-information model does not depend on the specific data-generating process for the variable being forecasted, the equivalent prediction in the context of the noisy-information model requires the additional assumption of an AR(1) process. In this section, we consider the implications of a more general process. First, suppose that the variable being forecasted x follows an AR(p) such that $\boldsymbol{z}_t = [x_t \ldots x_{t-p+1}]'$ and $\boldsymbol{z}_t = \boldsymbol{B}\boldsymbol{z}_{t-1} + \boldsymbol{H}'v_t$ where $\boldsymbol{H} = [1\ 0 \ldots 0]$ and $v_t \sim iid\ N(0, \sigma_v^2)$. Each agent i observes signal $y_{it} = \boldsymbol{H}\boldsymbol{z}_t + \omega_{it}$ where $\omega_{it} \sim iid\ N(0, \Sigma_\omega)$ is the agent-specific shock which is i.i.d. across agents and time. We assume that $E(v_t\omega_{it}') = 0$, that is shocks to fundamentals v_t and measurement error shocks ω_{it} are independent. Agent i's forecast for the unobserved state is $F_{it}\boldsymbol{z}_t = F_{it-1}\boldsymbol{z}_t + \boldsymbol{G}(y_{it} - F_{it-1}y_t)$ where $\boldsymbol{G}$ is the $(p \times 1)$ gain of the Kalman filter. We show in Appendix B that the average forecast error for x at horizon h follows

$$(14) \quad x_{t+h} - F_t x_{t+h} = \beta_{11}(F_t x_{t+h} - F_{t-1}x_{t+h}) + \cdots + \beta_{1p}\big(F_t x_{t+h-(p-1)} - F_{t-1}x_{t+h-(p-1)}\big) + error_t$$

where β_{ij} is the (i^{th},j^{th}) element of $\boldsymbol{\beta} \equiv \boldsymbol{B}^h[\{\boldsymbol{GH}\}^+(\boldsymbol{I} - \boldsymbol{GH})](\boldsymbol{B}^h)^{-1}$, $\boldsymbol{A}^+$ denotes generalized inverse of matrix $\boldsymbol{A}$, and $error_t$ is the rational expectations error. Note that this expression is quite close to the baseline prediction, except for the fact that predictability in forecast errors now obtains not just from the revision in the forecasts at horizon h but also from the contemporaneous revisions to forecasts at shorter horizons due to the higher-order dynamic process. An additional difference is that the coefficients on forecast revisions may reflect not just the degree of information rigidity but also the specific AR(p) parameters in $\boldsymbol{B}$. While the effect of the latter can be quantitatively small since information rigidities are premultiplied by $\boldsymbol{B}^h$ then post-multiplied by its inverse, coefficient estimates on forecast revisions may not depend only on information rigidities as was the case with AR(1).

An appealing feature of the generalized noisy-information model is that we can test empirically whether higher-order dynamics are important in characterizing the nature of the expectations formation process. To this end, we estimate equation (14) with forecast revisions for different horizons. Because SPF forecasts are at quarterly horizons ranging from t to $t+4$, we focus on forecasts of quarterly inflation two quarters ahead, which allows us to consider up to AR(3) specifications.[8] We rely on Bayesian information criterion (BIC) to identify the best-fitting specification because, if e.g. an AR(1) approximates the true model, forecast revisions at other horizons will be highly correlated, and point estimates at different horizons will be imprecise. The results, shown in Table 2, point to an AR(1) representation of inflation as the preferred specification, consistent with the description of the model in section I.B.

One can also consider VAR(p) representations of the data-generating process, such that the dynamics of e.g. inflation also depend on the dynamics of other macroeconomic variables. The predictability of forecast errors for one variable will then depend not just on forecast revisions for that variable but also upon forecast revisions of the other variables in the dynamic system. For example, if we are interested in forecasts of variable x^a whose dynamics are

[8] Four-quarter ahead forecasts have missing values in the early part of the sample, thus making the samples for AR(4) specifications smaller and not comparable across specifications.

determined jointly with x^b in a two-variable VAR(1), Appendix B shows that the predictability of ex-post forecast errors for x^a would follow

(15) $$x_{t+h}^a - F_t x_{t+h}^a = \beta_{11}(F_t x_{t+h}^a - F_{t-1} x_{t+h}^a) + \beta_{12}\left(F_t x_{t+h}^b - F_{t-1} x_{t+h}^b\right) + error_t.$$

Because the SPF includes forecasts of additional variables since 1969 (output growth, unemployment, housing starts and industrial production), we can use the same procedure as with the AR(p) to assess whether the predictability of inflation forecast errors is better represented as a VAR process. The results in Table 2 illustrate that the simpler AR(1) specification is preferred to any VAR representation of the data-generating process, again consistent with the representation of the noisy-information model in section I.B. Of course, these results should not be interpreted as implying that an AR(1) is the best representation of inflation dynamics. Rather, the results suggest that this parsimonious description of the data is sufficient to adequately characterize the predictability of forecast errors in terms of forecast revisions.

[INSERT TABLE 2 HERE]

D.2 Heterogeneity in Signal-Noise Ratios

Another extension of the noisy-information model allows agents to have different signal-noise ratios and therefore place different weight on new information received. An appealing feature of this setup is that two agents can receive the exact same signal but, because of heterogeneity in their Kalman gains, will adjust their forecasts by different amounts. Because heterogeneity in the Kalman gains quickly reduces the tractability of the model, we impose the following restrictions to simplify exposition: 1) $\rho = 1$ such that the variable being forecasted is a random walk, 2) Kalman gains are distributed normally across agents with variance σ_G^2 and mean G. In such a setting, we show in Appendix C that the predictability of the forecast error is given by

$$x_{t+h} - F_t x_{t+h} = \frac{(1-G)}{G}(F_t x_{t+h} - F_{t-1} x_{t+h}) + \frac{\sigma_G^2}{G}\left\{\sum_{k=0}^{\infty} A_k(G) x_{t-k-1}\right\}$$
$$+ error_t$$

where G is the average gain across agents and $A_k(G)$ are coefficients defined in Appendix C. The key prediction that follows from heterogeneity in signal strength is that average forecast errors will be predictable using not only forecast revisions but also lagged values of the variable being forecasted. But as documented in Table 1, there is no evidence that average inflation forecast errors in the U.S. are predictable using lagged inflation once one conditions on forecast revisions. Hence, heterogeneity in signal strength does not appear to be a quantitatively significant source of information rigidity in our data.

D.3 Heterogeneity in Beliefs about Long-Run Means

We also consider a variation on the baseline noisy-information model in which agents hold different priors about long-run values for economic variables, as in Patton and Timmermann (2010). Specifically, given the same setup as in section I.B, we follow Patton and Timmermann (2010) and assume agents report forecasts $F_{it}x_{t+h} = \omega\mu_i + (1-\omega)E[x_{t+h}|\Omega_{it}]$ where conditional expectations $E[x_{t+h}|\Omega_{it}]$ are formed using the Kalman filter and the agent specific signals, as in section I.B, $\omega \in [0,1]$ is the shrinkage factor, and μ_i is agent *i*'s belief about the long-run value of *x* with μ_i being zero mean across agents. We show in Appendix D that in this case the predictability of average forecast errors is given by

(16) $$x_{t+h} - F_t x_{t+h} = \frac{1-G}{G}(F_t x_{t+h} - F_{t-1} x_{t+h}) + \omega\rho^{h+1} x_{t-1} + error_t.$$

There are two differences relative to the baseline prediction from the noisy-information model in section I.B. First, ex-post forecast errors should be predictable using not only the ex-ante forecast revisions but also past values of the variable being forecasted. Second, the error term now includes a time-*t* component, such that OLS is invalid. To assess the quantitative importance of heterogeneity about long-run means, we regress annual year-ahead inflation forecast errors on forecast revisions and lagged annual inflation using two quarters of lagged changes in log oil prices as instruments. As documented in Panel A of Table 3, these instruments are strong predictors of forecast revisions. The results from estimating equation (16) by IV yield no evidence of predictability in forecast errors coming from lagged inflation once one conditions on forecast revisions. This suggests that heterogeneity about long-run means is not playing a quantitatively significant role in accounting for the predictability of inflation forecast errors.

[INSERT TABLE 3 HERE]

D.4 Heterogeneity in Loss-Aversion

While each of the preceding interpretations of the predictability of ex-post forecast errors stemmed from the presence of information rigidities, previous research has suggested that such predictability could arise even under full-information rational expectations. Capistran and Timmermann (2009), for example, present a model in which forecasters have asymmetric loss functions with heterogeneity in the degree of loss-aversion. This heterogeneity combined with the presence of GARCH dynamics can account for predictable forecast errors despite forecasters having identical and complete information. To assess whether this theory can account for the predictability of forecast errors in terms of forecast revisions, we derive the covariance between the two predicted by this theory.

Following Capistran and Timmermann (2009), agents face a "Linex" loss function over their forecast errors $FE_{t,t-1} = \pi_t - F_{t-1}\pi_t$ given by

$$L\left(FE_{t,t-1}; \phi_i\right) = \left[\exp\left(\phi_i FE_{t,t-1}\right) - \phi_i FE_{t,t-1} - 1\right] / \phi_i^2.$$

Positive ϕ_i imply that agents dislike positive forecast errors more than negative ones, and vice-versa, while $\phi_i = 0$ yields the standard mean-squared-error (MSE) objective. Suppose that inflation follows: $\pi_t = \rho\pi_{t-1} + v_t$ where v_t is not serially correlated but potentially heteroskedastic. Specifically, $\sigma_t = \alpha_0 +$

$\alpha_1 v_{t-1}^2 + \beta\sigma_{t-1}$ so that $v_t \sim (0, \sigma_t)$ as in Capistran and Timmermann (2009). In this setting, we show in Appendix E that the sign of the covariance between the average forecast error and average forecast revision across agents is *negative*. This is because large innovations to inflation (in absolute value) will lead agents to significantly raise (lower) their inflation forecast when the average $\phi > 0$ ($\phi < 0$) relative to the conditional expectation because of the asymmetry in the loss function. As a result, an upward forecast revision will tend to be associated with a subsequent negative forecast error. Thus, the model implies a predictability of average forecast errors from ex-ante forecast revisions, but in the wrong direction relative to our empirical estimates.

D.5 Forecast Smoothing

An alternative explanation for predictable forecast errors by professional forecasters is that they engage in forecast smoothing for reputational considerations. For example, forecasters typically provide not just a numerical forecast to clients but also a qualitative interpretation of recent economic developments. To preserve their reputation, forecasters may try to avoid drastic short-run changes in their forecasts. As a result, their forecast errors could be predictable even in the absence of information frictions.

Suppose forecasters face the following problem: at time t, given a forecast from time $t-1$, the forecaster needs to choose a sequence of forecasts (in expectation) of a variable x at time $t+h$

$$min \sum_{j=0}^{h} \gamma^j E_t \left[\left(x_{t+h} - F_{t+j} x_{t+h} \right)^2 + \alpha \left(F_{t+j} x_{t+h} - F_{t+j-1} x_{t+h} \right)^2 \right],$$

where γ is the discount factor. As we show in Appendix F, the first-order condition, after imposing FIRE, can be written as

$$(17) \quad x_{t+h} - F_t x_{t+h} = -(1+\alpha\gamma)(F_t x_{t+h} - F_{t-1} x_{t+h}) + \alpha(F_{t-1} x_{t+h} - F_{t-2} x_{t+h}) + error_t$$

where the error term $error_t$ is correlated with information at time t but not earlier. Forecast smoothing also yields predictability of ex-post forecast errors, but the coefficient on the contemporaneous forecast revision is negative while that on the lagged forecast revision is positive. These coefficients reflect the fact that, while forecasters will smooth the change in their forecasts to minimize adjustment costs, they must also take into account how their choice of forecasts will affect the subsequent period's adjustment cost. As a result, forecast errors should be predictable using both contemporaneous and past forecast revisions for a given horizon. In this respect, forecast smoothing yields predictions which differ from those of information rigidities and which can be assessed empirically.

We regress ex-post inflation forecast errors on both current and lagged forecast revisions using two lags of log changes in oil prices as instruments to account for the correlation between the error term and time-t variables. However, because the horizons in the SPF data are not sufficient to construct lagged forecast revisions of year-on-year inflation at overlapping horizons, we use quarterly forecasting horizons ranging from the current quarter (h = 0) to two quarters ahead (h = 2) and present pooled estimates over these three horizons in Panel B of Table 3. We find no evidence of predictive power for the lagged forecast revision, and the coefficient on the contemporaneous forecast revision is positive, as predicted by models of information rigidities, rather than negative as predicted by forecast smoothing. Hence, the empirical evidence again supports the notion that the predictability of forecast errors across agents is driven by the presence of information rigidities such as sticky-information or noisy-information.[9]

[9] Two additional pieces of evidence support the notion that reputational considerations are not the source of forecast smoothing. First, SPF forecasts are no worse at forecasting inflation (in terms of mean squared error) than households or financial markets (Appendix Table 1). This is important because neither households nor financial markets should have reputational incentives as household surveys are anonymous while forecast smoothing would leave money on the table for financial market participants. Second, when we regress ex-post inflation on the ex-ante forecasts of SPF and either household or financial market forecasts (or both), we cannot reject the hypothesis that the

II. Information Rigidities across Agent Types, Macroeconomic Variables, and Countries

The results using U.S. inflation forecasts of professional forecasters are consistent with persistent deviations from full-information rational expectations and with the presence of economically significant information rigidities. In this section, we expand the scope of these results by examining forecasts of other economic agents and variables as well as cross-country evidence on information rigidities.

A. Information Rigidity across Agents

While professional forecasters provide a useful benchmark to assess the null of full-information rational expectations, the quantitative importance of their expectations is unclear. Thus, we turn to surveys of other agents to assess whether the information rigidities identified for professional forecasters appear to be a more general phenomenon. The Livingston Survey provides semiannual individual inflation forecasts from academic institutions, commercial banks, and non-financial firms, among others.[10] Thus, we can use this alternative source of data to assess information rigidities for academics and agents in the private sector, both financial and non-financial. The Livingston survey includes individual forecasts of the Consumer Price Index (CPI) in 6 months and in 12 months so we can apply our empirical specification at the 6-month forecasting horizon. Panel A of Table 4 presents the results from 1969 to 2014 for the mean forecasts across all forecasters as well as using the mean forecasts across subsets of professional forecasters. Academics have the smallest estimated coefficient on forecast revisions—different from zero only at the 10 percent significance level—while those of forecasters from

---

coefficient on SPF forecasts is one while that on other forecasts is zero (Appendix Table 1). Hence, professional forecasters do not seem to be underutilizing information in the interest of protecting their reputations relative to other agents.

[10] The categories of forecasters also include investment banks, government forecasters, the Federal Reserve, labor organizations, and "other." We do not look at these in detail because of how few forecasters there are in each of these groups over time while the theoretical predictions from sticky- and noisy-information models apply specifically to the mean forecasts across agents. The Livingston data is available on the website of the Federal Reserve Bank of Philadelphia.

commercial banks and non-financial institutions are slightly larger. For forecasts of firms, as well as when we average across all forecasters, the estimated level of information rigidity is positive and statistically significant, indicating that information rigidities are not limited to professional forecasters in the SPF. Furthermore, when converted to quarterly equivalents, the quantitative magnitudes are close to those reported for the SPF in the previous section.

We also consider two additional types of agents: consumers and financial market participants. For the former, we rely on the Michigan Survey of Consumers. Each month, the University of Michigan surveys 500-1,500 households and asks them about their expectation of price changes over the course of the next year. For the latter, we use inflation expectations from the Cleveland Fed based on the method developed in Haubrich, Pennacchini, and Ritchken (2008) which uses the term structure of interest rates and inflation swaps to extract measures of market expectations of CPI inflation at multiple yearly horizons starting in 1982. A drawback of both sources of expectations data is that they are only available at a forecasting horizon of one year and therefore revisions in forecasts over identical horizons are not available. Thus, we replace the forecast revision with the change in the year-ahead forecast, yielding the following quarterly specification

$$\pi_{t+4,t+1} - F_t\pi_{t+4,t+1} = c + \beta\left(F_t\pi_{t+4,t+1} - F_{t-1}\pi_{t+3,t}\right) + error_t \tag{18}$$

where $\pi_{t+4,t+1}$ denotes the inflation rate between $t+4$ and $t+1$. The error term now consists of the rational expectations forecast error, as in equation (11), and $\beta(F_{t-1}\pi_t - F_{t-1}\pi_{t+4})$ because the forecasts do not have perfectly overlapping time horizons across periods. As a result, this specification cannot be estimated by OLS. Instead, we estimate this specification by IV, using as an instrument the (log) change in the oil price. Because oil prices have significant effects on CPI inflation, they are statistically significant predictors of contemporaneous changes in inflation forecasts for all three measures of inflation expectations and can account for an important share of their volatility (Table 4).

[INSERT TABLE 4 HERE]

In addition to estimates using consumer and financial market forecasts, we also report estimates of equation (18) using SPF forecasts of CPI inflation at equivalent forecasting horizons. This allows us to assess the extent to which the additional error from non-overlapping horizons—and the need to use instruments—affects the empirical estimates. Using a common time sample of 1982 to 2014, we find a coefficient on forecast revisions of 1.1 for professional forecasters, a finding that closely mirrors our previous results. This indicates that our GMM procedure and instruments can adequately address the additional complications introduced by non-overlapping horizons and that our baseline results are insensitive to the measure of inflation. For consumers, the point estimate is smaller but also highly statistically significant which conforms to the findings of Coibion and Gorodnichenko (2012). For market expectations, the point estimate is higher than for professional forecasters but also less precisely estimated. This is likely a result of the reduced predictive power of oil price innovations on market-based forecast revisions, as can be seen in the first-stage fit. In short, estimates of the predictability of inflation errors from forecast revisions point to the presence of information rigidities across consumers, financial market participants, and firms as well as professional forecasters.

B. Information Rigidity across Horizons and Variables

While much of the empirical literature on the expectations formation process has focused on year-ahead inflation forecasts, our approach can be applied to other macroeconomic variables and horizons. In the canonical sticky-information model, for example, agents update their information sets infrequently, but when they do so, they acquire full-information rational expectations. As a result, there is a single parameter governing the frequency of updating information which is common across macroeconomic variables *and* forecasting horizons. Thus, a testable implication of the canonical sticky-information model is that the estimated degree of information rigidity is invariant to the forecasting horizon and the variable being forecasted. In noisy-information models of the type presented in section I, the degree of information rigidity will also be invariant to the horizon of the forecasts. However, the coefficient on forecast revisions for a given macroeconomic variable in noisy-information models will be governed by the Kalman gain associated with that variable, which will depend on factors such as the persistence of the series and the strength of the signal observed with respect to that macroeconomic variable. In this section, we first test the common prediction that the estimated degree of information rigidity is invariant to the forecasting horizon. Secondly, we test the null hypothesis of the sticky-information model that the degree of information rigidity is equal across macroeconomic variables.

To implement these tests, we exploit the fact that the Survey of Professional Forecasters (SPF) contains quarterly forecasts for four additional macroeconomic variables going back to 1968Q4. Besides the GDP price deflator, these include real GDP, industrial production, housing starts, and the unemployment rate.[11] Each of these variables is available at multiple forecasting horizons, ranging from forecasts of the current quarter to 4 quarters ahead. While the SPF includes forecasts up to 4 quarters ahead, the horizon is limited to 3 quarters in the empirical specification because forecast revisions call for an additional forecasting horizon, e.g. when h=3, the forecast revision is $F_t x_{t+3} - F_{t-1} x_{t+3}$. To construct forecast errors, we use real-time values available one year after the relevant time horizon. For the first three series, forecasts of annualized quarterly percent changes are constructed from the underlying mean forecasts of the levels. Starting in 1981, the SPF also includes forecasts of 8 additional macroeconomic variables: the 3-month Treasury bill (Tbill)

[11] Output is measured by GNP prior to 1992 and GDP thereafter.

rate, the AAA interest rate, real consumption expenditures, real residential investment, real non-residential investment, real federal government expenditures, real state/local government expenditures, and the overall CPI. For each NIPA series and CPI inflation, we construct forecasts of annualized quarterly percent changes and use real-time data to construct forecast errors, while the two interest rates are measured in levels. The forecast horizons usable in our approach again run from $h = 0$ to $h = 3$.

We estimate the degree of information rigidity pooled across variables, pooled across horizons, or pooled across both variables and horizon to assess the extent to which there is heterogeneity in estimated degrees of information rigidity across each dimension. Specifically, we estimate pooled regressions

(19) $$x_{k,t+h} - F_t x_{k,t+h} = c + \beta\left(F_t x_{k,t+h} - F_{t-1} x_{k,t+h}\right) + error_{k,h,t}$$

where x_k indicates which macroeconomic variable is included and h denotes the specific forecasting horizon ranging from 0 (forecasts of the current quarter) to 3 (forecasts for 3 quarters ahead). As illustrated in Appendix G, pooled estimates of (19) provide weighted averages of the underlying degrees of information rigidity where the weights reflect relative variances of forecast revisions. Under the null of the sticky-information model, the estimate of information rigidity should be the same across variables and horizons, so pooling serves to increase the precision of the estimates. Under noisy-information models, we can also pool across horizons to increase the precision of estimates for each variable. But because the degree of information rigidity may vary across variables, pooling across variables will deliver a weighted average of underlying degrees of information rigidity.

We first consider results when we pool across variables for each forecasting horizon. The left-hand figure in Panel A of Figure 1 presents pooled estimates of information rigidity at each forecasting horizon for the five variables available since 1968, while the right-hand figure in Panel A presents equivalent results using the thirteen variables consistently available since 1981. We also show in each panel the estimated level of information rigidity pooled across all variables and horizons within each time period to provide a benchmark for assessing the extent of cross-sectional heterogeneity. Standard errors are constructed as in Driscoll and Kraay (1998), which allows for both cross-sectional and serial correlation in the errors as well as heteroskedasticity in the errors. In each case, the estimate of information rigidity pooling across all variables and horizons is positive and statistically different from zero.[12] We find that one cannot reject the null hypothesis that the estimated degree of information rigidity is invariant across forecasting horizons (*p*-values of 0.14 and 0.12 for 1968 variables and 1981 variables respectively). Hence, the null hypothesis implied by both sticky-information and noisy-information models in terms of forecasting horizons cannot be rejected.

When we pool across horizons for each macroeconomic variable, however, we observe much more heterogeneity. The left-hand figure in Panel B of Figure 1 presents estimates of information rigidity for each variable available since 1968, while the right-hand figure in Panel B presents estimates for each variable available since 1981. We can reject the null hypothesis that the degree of information rigidity is the same for all variables: *p*-values for the null are 0.06 for the five variables since 1968 and less than 0.001 for the thirteen variables available since 1981. The presence of heterogeneity in estimated levels of information rigidity across variables is inconsistent with the baseline sticky-information model in which agents are assumed to update their expectations about all macroeconomic variables simultaneously. This suggests that future work employing the sticky-information model should allow for differential updating rates across macroeconomic variables,

[12] In Appendix Table 2, we present additional results for the specifications pooling across all variables and horizons. We show for example that the estimates are invariant to controlling for cross-sectional fixed effects, as well as controlling for both cross-sectional and time fixed-effects. We also present estimates of equation (19) in which we allow for different coefficients on contemporaneous and lagged forecasts. In each case, we cannot reject the null hypothesis implied by models of information rigidities that the sum of the two coefficients is equal to zero.

in the same spirit as Mackowiak and Wiederholt's (2009) model of rational inattention.

[INSERT Figure 1 HERE]

While these estimates have a clear interpretation in the sticky-information model since the degree of information rigidity is invariant across variables and horizons, estimating the degree of information rigidity for each variable separately could be more problematic in the context of noisy-information models if agents form forecasts of different macroeconomic variables jointly. To see this, suppose we have a multivariate process for the k-by-1 vector of variables $\boldsymbol{z}$ such that agents receive signals about each variable which are i.i.d. across time and agents. Then, as shown in Appendix B, average *contemporaneous* forecast errors across agents will be predictable from average forecast revisions of all variables according to

(20) $$\boldsymbol{z}_t - \boldsymbol{z}_{t|t} = \boldsymbol{M}\left(\boldsymbol{z}_{t|t} - \boldsymbol{z}_{t|t-1}\right) + \ error$$

where $\boldsymbol{M}$ is related to the Kalman gain used by agents to update their forecasts by $\boldsymbol{G} = (\boldsymbol{M} + \boldsymbol{I})^{-1}$ when $\boldsymbol{z}$ follows a VAR(1).[13]

In this setting, characterizing the degree of information rigidity is less direct since the Kalman gain is a matrix and predictability in forecast errors can come not only from forecast revisions of that variable but also from forecast revisions of other variables. Nonetheless, the diagonal elements of the Kalman gain still have a direct interpretation in line with those of section I: they represent the average decrease in the variance of forecast errors for each variable between time $t-1$ and t from all new information, since $\boldsymbol{\Omega} = (\boldsymbol{I} - \boldsymbol{G})\boldsymbol{\Psi}$ where $\boldsymbol{\Omega} \equiv \text{var}\left(\boldsymbol{z}_t - \boldsymbol{z}_{t|t}\right)$ and $\boldsymbol{\Psi} \equiv \text{var}\left(\boldsymbol{z}_t - \boldsymbol{z}_{t|t-1}\right)$, which has the same interpretation as $1 - \boldsymbol{G}$ in the univariate model of section I.

[13] If z follows a higher-order VAR, then equation (20) will contain additional terms representing average revisions to the backcasts of agents (e.g. $\boldsymbol{z}_{t-1|t} - \boldsymbol{z}_{t-1|t-1}$). Because these backcasts are not observable in survey data, we cannot test this dimension of the data directly.

To assess whether allowing for the joint forecasting of multiple variables affects our estimates of information rigidity, we first consider whether forecast revisions of other variables improve the predictability of forecast errors for each variable in the SPF available since 1968. Table 5 presents values of the Bayesian information criterion for AR and VAR estimates of equation (20). For each variable, the BIC selects AR(1) as the preferred specification, suggesting that forecast revisions of other variables have little predictive power for the forecast errors of each macroeconomic variable. Furthermore, if we consider estimates of the system (20), information criteria suggest that a diagonal M matrix is statistically preferred to one with non-zero off-diagonal entries. This again suggests that there is little additional predictive power coming from forecast revisions of other variables and therefore that one can estimate the degree of information rigidity using single equation methods for each variable.

[INSERT TABLE 5 HERE]

While there is therefore little statistical evidence to suggest a non-diagonal M matrix, one can alternatively estimate the unrestricted system of equations (20) and extract the implied estimates of the diagonal entries of the Kalman gain to compare them to the univariate estimate at the same forecasting horizon (h=0). We provide these results in Table 6. For all but one variable, the single equation and systems estimates are quite similar (we cannot reject the null of equality between the two) although the estimates from the systems approach are much less precise.[14] Only for industrial production is there a statistically significant

[14] The estimates for output growth, unemployment and industrial production are particularly less precise in the systems approach than with single equation estimates. This likely reflects collinearity issues: forecast revisions in output growth, industrial production and unemployment are highly correlated with one another.

difference. Again, we can reject the null of equal degrees of information rigidity across macroeconomic variables, as was the case with single equation estimates.[15]

Jointly, these two results—the absence of statistical evidence for important non-diagonal entries in $\boldsymbol{M}$ and the similarity of the systems and single equation estimates—suggest that our baseline single equation approach is empirically suitable to recover the degree of information rigidity in each macroeconomic variable even in more general noisy-information setups than those of section I.

[INSERT TABLE 6 HERE]

C. Cross-Sectional Variation and the Determinants of Information Rigidity

While the evidence strongly rejects the null of equal amounts of information rigidity across macroeconomic variables, a feature of the data which is at odds with canonical sticky-information models, the fact that heterogeneity in information rigidity exists across macroeconomic variables does not imply that noisy-information models can account for this cross-sectional variation. In the simple noisy-information model of section I.B, the degree of information rigidity depends on the Kalman gain, which is itself a function of the persistence of the underlying macroeconomic process as well as the precision of the signal received by economic agents. More persistent processes imply, holding all else constant, that agents should put a higher weight (G) to current signals. A more precise signal naturally implies that agents should place relatively more weight on the current signal than on past forecasts. Thus, noisy-information models imply that the degree of information rigidity should be decreasing in the persistence of the series being forecasted and increasing in the amount of noise in the signal.

To test whether the degree of information rigidity associated with macroeconomic variables follows these predictions requires forecasts for a wide

[15] This rejection of the null of equality across coefficients is not driven by industrial production. The result holds if we estimate the system with or without industrial production.

cross-section of macroeconomic variables. We have constructed a dataset of quarterly forecasts from the international survey of professional forecasters done by Consensus Economics. This dataset covers twelve countries: the G-7 countries of U.S., U.K., France, Germany, Italy, Japan and Canada as well as Spain, Norway, the Netherlands, Sweden and Switzerland. Data for the G-7 countries span 1989 to 2010 while data for other countries begin primarily in 1994. For each country, forecasts for five macroeconomic variables are available: consumer price inflation, real GDP growth, interest rates, industrial production growth and real consumption growth. Forecasts are available for the current quarter and for the subsequent 5-6 quarters.

This dataset therefore provides significant cross-sectional variation over countries and macroeconomic variables. Pooling across all variables and horizons within each country to provide a summary statistic of information rigidity by country, Panel A of Figure 2 illustrates that there is significant variation in the average degree of information rigidity across countries. The countries with the highest degrees of information rigidities are Spain and Japan, while the lowest are Canada and Norway. All of the estimates are statistically significantly positive so we can reject the null of FIRE for every country and this rejection of the null goes exactly in the direction predicted by models of information rigidities. Panel A of Figure 2 also plots the estimated degree of information rigidity pooled across all countries, variables and horizons to provide a benchmark for assessing the amount of cross-country heterogeneity. This cross-country pooled estimate is almost identical to that obtained for the U.S.[16] Furthermore, the country-specific estimate

[16] The estimate pooled across countries is almost identical to the average across country-specific estimates. In Appendix Table 2, we also show that the degree of information rigidity pooled across countries is largely invariant to controlling for cross-sectional and time fixed effects. Furthermore, we again cannot reject at the 5 percent level (and at the 10 percent when we control for fixed effects) the null that the coefficients on current and lagged forecasts sum to zero. Controlling for the introduction of the Euro also does not affect the estimates of information rigidity.

for the U.S. using Consensus Economic forecasts is remarkably close to those found using the SPF forecasts. Panel B presents pooled estimates of the coefficient on forecast revisions pooled across horizons and countries for each macroeconomic variable. As was the case just within the U.S., we find significant differences across macroeconomic variables (we can reject the null hypothesis of equality across variables with a *p*-value of 0.0003). Jointly, there is therefore a significant amount of variation in information rigidity across countries and macroeconomic variables.

[INSERT FIGURE 2 HERE]

We exploit this cross-sectional variation in the Consensus Economics cross-country data to assess whether the observed heterogeneity in information rigidity is consistent with the predictions of noisy-information models. First, for each country j and macroeconomic variable i in the Consensus Economics survey of professional forecasters, we fit an autoregressive process which yields an estimate of both the persistence of the variable ($\rho_{i,j}$) and the volatility of its innovations ($\sigma_{i,j}$). Second, we generate a measure of the noise associated with each series from the standard deviation of revisions to this series.[17] Third, we construct a measure of the noise-signal ratio ($\kappa_{i,j}$) by taking the ratio of the measure of the noise to the standard deviation of the innovations to the variable from the first step. Given these measures of the predictors of information rigidity, we assess their importance by regressing our estimates of the coefficients on forecast revisions for each country-macroeconomic variable pair, pooled across forecasting horizons, in the cross-country Consensus Economics dataset

(21) $$\beta_{i,j} = c + \gamma_1 \rho_{i,j} + \gamma_2 \kappa_{i,j} + error_{i,j}$$

[17] Specifically, for each time period, we take the difference between measures of the variable available two quarters and four quarters later, then compute the standard deviation of these revisions across the entire sample. Alternative time horizons for measuring revisions yield the same qualitative results. Real-time data, including revisions over the course of a year, are included in the Consensus Economics dataset. These measures reflect common noise, rather than idiosyncratic noise faced by forecasters, but unfortunately no measures of the latter are available. As a result, we take a measure of the degree of common noise as a proxy for the degree of the noise in private signals.

where i denotes a specific variable, j denotes the country, and $\beta_{i,j}$ is the estimated coefficient on forecast revisions for each country-variable pair in the cross-country dataset.

The results are presented in Table 7. The coefficients on the persistence are consistently negative across specifications.[18] The coefficient on the noise in the signal is positive, as expected, but not significantly different from zero. The latter finding is sensitive to outliers (Appendix Figure 2). As a result, we also consider estimates of (20) based on robust S-regressions, which automatically identify and account for outliers, and the results point to a positive and statistically significant effect of the noise-signal ratio, as predicted by the noisy-information model. Strikingly, this simple specification can account for about 15-30 percent of the heterogeneity in information rigidities. Thus, not only are the theoretical predictions of noisy-information models qualitatively consistent with the observed heterogeneity in information rigidities across countries and variables, but this model can also *quantitatively* account for a considerable share of the observed cross-sectional variation.

[INSERT TABLE 7 HERE]

III State-Dependence in Information Rigidities

The previous section presents evidence that the varying degrees of information rigidity associated with macroeconomic variables are well-explained by the persistence and noise-signal ratios of these variables. However, the precision of an agent's signal can be thought of as choice variable when agents have the ability to devote more resources to collecting and processing information as in the

[18] Under the null hypothesis that $\gamma_1$and γ_2 are zero, Pagan (1984) shows that standard errors on generated regressors are asymptotically valid. We found that our inference is not affected when we correct standard errors as in Murphy and Topel (1985). Furthermore, to the extent that the persistence of macroeconomic variables depends on the degree of information rigidity, this potential endogeneity will tend to bias our estimate of γ_1 upward. This reflects the fact that if the persistence of macroeconomic variables is increasing in the degree of information rigidity, then this would tend to push γ_1 up.

rational inattention models of Sims (2003) and Mackowiak and Wiederholt (2009). In this section, we investigate whether the degree of information rigidity responds to changing economic conditions. We focus on two particular dimensions. First, we assess whether the degree of information rigidity responds to low-frequency variation in economic volatility, as exemplified by the Great Moderation. Second, we characterize the extent to which information rigidity responds to business cycle fluctuations. For both, we document evidence of state-dependence in the degree of information rigidity, consistent with the reallocation of information collection and processing resources in light of economic conditions as suggested by rational inattention theories.

A. Information Rigidities and the Great Moderation

McConnell and Perez-Quirós (2000) and others have documented a substantial decrease in macroeconomic volatility both in the U.S. and other developed countries since the early to mid-1980s. Figure 3 plots the time-varying standard deviation of real GDP growth for the U.S., for example, which is rising throughout the 1970s, peaks in the very early 1980s, then exhibits a very sharp decline in the mid-1980s, declining by more than half relative to the average level during the 1970s. While the source of this phenomenon remains a point of contention, one explanation emphasizes the changes in monetary policy put in place under Volcker, either in terms of a stronger endogenous response to macroeconomic fluctuations as in Clarida, Gali, and Gertler (2000) or because of the Volcker disinflation as in Coibion and Gorodnichenko (2011). At the same time, there is only mixed evidence that microeconomic volatility declined over this time period. For example, Davis et al. (2006) report that the volatility of employment has fallen since the 1970s for non-publicly traded firms while Comin and Mulani (2004) and Comin and Philippon (2005) show that volatility increased for publicly traded firms over the same period. Furthermore, volatility at the household level appears to have been trending up over time (see Davis and Kahn (2008) for a review). As a result of the reduction in the volatility of macroeconomic variables relative to microeconomic variables, one might expect that economic agents would choose to allocate relatively more resources to tracking micro rather than macro-level shocks since these shocks became quantitatively more important for profits and utility as in Mackowiak and Wiederholt (2009, 2012). Thus information rigidity should have increased with the arrival of the Great Moderation.

To explore this hypothesis, we estimate equation (10) for each quarter separately using SPF data and then compute non-parametrically a local average of the estimated β's to provide a sense of the low frequency variation in the degree of information rigidities. Figure 3 plots the dynamics of the local averages of β as well as associated standard errors.[19] The figure shows that information rigidities were falling from the late 1960s to the early 1980s as the volatility of macroeconomic variables was rising. The minimum level of information rigidity is reached in the early 1980s, which closely matches the start of the Great Moderation identified in McConnell and Perez-Quirós (2000). The estimated degree of information rigidity then rose systematically over the course of the Great Moderation until the mid-2000s, before starting to decline after the onset of the Great Recession. The changes in the level of information rigidities over time are statistically and economically significant, especially when one compares the mid-1980s to the late 2000s, although estimates during the Great Recession period are noticeably less precise, which likely reflects the unusually high volatility of forecast errors during this time period.[20]

[INSERT FIGURE 3 HERE]

---

[19] We get nearly identical results if we estimate rolling regressions of the degree of information rigidity for each variable and then average across these estimates.

[20] Another potential source of rising information rigidity since the early 1980s is the decline in inflation persistence over this time period (Stock and Watson 2007). This is unlikely to be the sole factor, however, since the rise in estimated degrees of information rigidity since the 1980s holds even when inflation forecasts are excluded from the cross-section of variables.

This significant low-frequency variation in the estimated coefficients on forecast revisions suggests that one should be wary of treating information rigidities at the macroeconomic level as a structural parameter since these rigidities can vary over time in response to changes in macroeconomic conditions. Specifically, more tranquil times should be ceteris paribus associated with greater information rigidities, as suggested by the endogenous inattention model of Branch et al. (2009). The rising degree of inattention over the course of the 1990s and 2000s also implies that the same sized shock would have larger real effects towards the end of the sample because information rigidities, like nominal rigidities, amplify the response of the economy to a given set of shocks. Thus, this observation suggests an additional mechanism, along with increased risk-taking on the part of financial market participants, through which the Great Moderation may have contributed to the severity of the Great Recession. The figure also suggests that the high volatility of the 1970s was associated with a gradual increase in attention on the part of economic agents, making the economy become progressively less sensitive to any given shock. This decline in information rigidity may therefore also have contributed to onset of the Great Moderation. More generally, the figure points to the possibility of low-frequency cycles in volatility arising from the endogenous response of information rigidities to volatility and the feedback effect of changing information rigidity on volatility.

B. Information Rigidities over the Business Cycle

Our results in the previous section indicate that calm times are associated with stronger information rigidities. In light of this evidence, one may expect that recessions, as periods of increased volatility, should be times when economic agents update and process information faster than in expansions since the (relative) cost of ignoring macroeconomic shocks in recession rises. Gorodnichenko (2008), for example, shows in a theoretical model that the acquisition of information endogenously increases shortly after the occurrence of an aggregate shock as economic agents face increased uncertainty about the current state of the economy and consequently find it beneficial to devote more resources to learning about current macroeconomic conditions. Using the U.S. estimates of β computed for each quarter separately as in section III.A, we consider the following econometric specification

(22) $\beta_t = \alpha + \sum_{j=0}^{J} \phi_j I_{t-j}^{REC} + error_t$,

where I_t^{REC} is a dummy variable equal to one in the first quarter of each recession, as identified by the NBER, and zero otherwise. By varying index j, we construct a sequence of estimated ϕ_j which may be interpreted as an impulse response of information rigidities to a recession. To smooth the path of coefficients ϕ_j, we fit a polynomial distributed lag model with the polynomial order equal to 4 and $J = 20$. We consider estimates over the entire sample as well as over the subsample that excludes the Great Recession period to assess the sensitivity to this unique period. Figure 4 shows the path of the estimated ϕ_j over four years after the economy slides into a recession for each sample. In each case, we assume that the economy starts at an average level of information rigidity which is equal to $\hat{\alpha}$. When we exclude the Great Recession, information rigidities are initially high but, as time passes, information rigidities become less severe to the point where we cannot reject the null of FIRE one year after the start of the recession. The degree of information rigidity then starts to recover to the level observed before the start of a recession.[21] When the Great Recession period is added, the estimates display significantly more uncertainty, reflecting the high volatility of forecast errors during this period. The apparent sensitivity of information rigidity to business cycle conditions is consistent with a broad class of models of information rigidities in which the acquisition of information is state-dependent, such as Gorodnichenko

[21] Loungani, Stekler and Tamirisa (2013) apply this methodology to real GDP growth forecasts of professional forecasters in 46 countries and similarly find reduced rates of information rigidities after recessions and banking crises.

(2008), and suggests that modeling the evolution of inattention could play an important role in accounting for business cycle dynamics.[22]

[INSERT FIGURE 4 HERE]

IV. Concluding remarks

Building from the predictions of models of information rigidities, we provide a new test of the null of full-information rational expectations which is informative about the economic significance of departures from the null as well as the models that can account for these departures. The core of the proposed approach is a tight theoretical link between ex-post mean forecast errors and ex-ante mean forecast revisions. Applying this approach to professional forecasters in the U.S. and other industrialized countries, we document widespread rejections of full-information rational expectations in exactly the direction predicted by models of information rigidities. Consistent with these models, when one takes into account forecast revisions, other macroeconomic variables lose much of their ability to predict forecast errors. One interpretation of our results is that commonly observed rejections of the null of full-information rational expectations most likely reflect deviations from full-information rather than departures from rational expectations. Indeed, economic agents in the sticky-information or noisy-information models are fully rational but operate subject to information frictions. The estimates also point to economically significant estimates of information rigidities, thereby providing support for the recent body of work studying the integration of information frictions into macroeconomic models. The approach developed here has a number of advantages relative to previous work, such as not requiring the identification of shocks, greater tractability, and the ability to investigate variation in information

[22] The substantial decrease in information rigidity during recessions could also be consistent with models in which agents have ambiguity aversion, as in Epstein and Schneider (2008) and Ilut (2013). In these models, rational agents face uncertainty about the true data generating process and place more weight on the less favorable data generating processes. As a result, these models point to information rigidity being insensitive to good news and decreasing strongly in response to bad news, such as recessions.

rigidity across variables and time. In addition, our approach nests previous tests and offers new tests of models of information rigidities which find support in the data.

While we have focused primarily on professional forecasters, this approach can be applied to other economic agents. For example, we document qualitatively similar results using the inflation forecasts from the Michigan Survey of Consumers, businesses in the Livingston Survey, as well as forecasts extracted from financial market prices. The pervasiveness of information rigidity across agents suggests that future work on imperfect information should strive to incorporate information frictions on the part of different agents, as in Reis (2009). Our empirical results can provide a disciplining device for such models, in terms of stylized facts to be matched, and can be helpful in calibrating the degrees of information rigidity for different kinds of agents.

In addition, one can apply our approach to study the implications of different policies on the expectations formation process. For example, we document that the Great Moderation, frequently attributed to the monetary policy changes enacted by Fed Chairman Paul Volcker, was associated with a pronounced and persistent increase in the degree of information rigidity for professional forecasters. This finding suggests a new mechanism through which, along with increased risk-taking behavior on the part of financial market participants, the Great Moderation may have played a role in generating the Great Recession. Our empirical specification can also help quantify the effect of policy changes on the expectations formation process, thereby providing a more theoretically grounded notion of otherwise vague concepts such as “anchored” expectations. Importantly, this approach can be applied to study a wide variety of policies such as inflation targeting or exchange rate regimes to shed light on one of the key mechanisms via which these policies are supposed to affect dynamics: the expectations formation process.

References

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

Andolfatto, David, Scott Hendry, and Kevin Moran, 2007. “Are Inflation Expectations Rational?” *Journal of Monetary Economics* 55(2), 406-422.

Andrade, Philippe, and Herve Le Bihan, 2013. “Inattentive Professional Forecasters,” *Journal of Monetary Economics* 60(8), 967-982.

Bacchetta, Philippe, Elmar Mertens, and Eric van Wincoop. 2009. “Predictability in Financial Markets: What Do Survey Expectations Tell Us?” *Journal of International Money and Finance* 28(3), 406-426.

Ball, Lawrence, N. Gregory Mankiw, and Ricardo Reis, 2005. “Monetary Policy for Inattentive Economies,” *Journal of Monetary Economics* 52(4),703-25.

Branch, William A., 2007. “Sticky information and model uncertainty in survey data on inflation expectations,” *Journal of Economic Dynamics and Control* 31, 245-276.

Branch, William A., John Carlson, George W. Evans, and Bruce McGough, 2009. “Monetary Policy, Endogenous Inattention, and the Volatility Trade-off,” *Economic Journal* 119(534) 123-157.

Capistran, Carlos and Allan Timmermann, 2009. “Forecast Combination with Entry and Exit of Experts,” *Journal of Business and Economic Statistics* 27, 429-440.

Carroll, Christopher D. 2003. “Macroeconomic Expectations of Households and Professional Forecasters,” *Quarterly Journal of Economics* 118(1), 269-298.

Clarida, Richard, Jordi Galí, and Mark Gertler, 2000. “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory,” *Quarterly Journal of Economics* 115(1), 147-180.

Coibion, Olivier and Yuriy Gorodnichenko, 2011. “Monetary Policy, Trend Inflation, and the Great Moderation: An Alternative Interpretation,” *American Economic Review* 101(1), 341–370.

Coibion, Oliver and Yuriy Gorodnichenko, 2012. "What can survey forecasts tell us about informational rigidities?" *Journal of Political Economy* 120(1) 116-159.

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

Comin, Diego and Sunil Mulani, 2004. "Diverging Trends in Aggregate and Firm-level Volatility." *Review of Economics and Statistics* 88(2), 374-383.

Comin, Diego and Thomas Philippon, 2005. "The Rise in Firm-Level Volatility: Causes and Consequences," *NBER Macroeconomics Annual*, 167-228.

Croushore, Dean, 2010. "An Evaluation of Inflation Forecasts from Surveys Using Real-Time Data," *Berkeley Electronic Journal of Macroeconomics: Contributions* 10:10.

Crowe, Christopher W., 2010. "Consensus Forecasts and Inefficient Information Aggregation," IMF Working Paper No. 10/178.

Davis, Steven, John Haltiwanger, Ron Jarmin, and Javier Miranda, 2006. "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms," NBER WP #12354.

Davis, Steven J. and James A. Kahn, 2008. "Interpreting the Great Moderation: Changes in the Volatility of Economic Activity at the Macro and Micro Levels," *Journal of Economic Perspectives* 22(4), 155–180.

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

Epstein, Larry G. and Martin Schneider, 2008. "Ambiguity, Information Quality and Asset Pricing," *Journal of Finance* 63(1), 197-228.

Gorodnichenko, Yuriy, 2008. "Endogenous Information, Menu Costs, and Inflation Persistence," NBER Working Paper 14184.

Gourinchas, Pierre-Olivier and Aaron Tornell, 2004. "Exchange-Rate Puzzles and Distorted Beliefs" *Journal of International Economics* 64, 303-333.

Haubrich, Joseph, George Pennacchi, and Peter Ritchken, 2008. "Estimating Real and Nominal Term Structures using Treasury Yields, Inflation, Inflation Forecasts and Inflation Swap Rates," FRB Cleveland WP 08-10.

Ilut, Cosmin, 2013. "Ambiguity Aversion: Implications for the Uncovered Interest Rate Parity Puzzle," *American Economic Journal: Macroeconomics* 4(3), 33-65.

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.

Kydland, Finn and Edward Prescott, 1982. "Time to Build and Aggregate Fluctuations," *Econometrica* 50(6), 1345-70.

Loungani, Prakash, Herman Stekler, and Natalia T. Tamirisa, 2013. "Information Rigidity in Growth Forecasts: Some Cross-Country Evidence," *International Journal of Forecasting* 29(4), 605-621.

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

Mackowiak, Bartosz and Mirko Wiederholt, 2008. "Business Cycle Dynamics under Rational Inattention," *mimeo*.

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

Mackowiak, Bartosz and Mirko Wiederholt, 2012. "Information Processing and Limited Liability," *American Economic Review* 102(3), 30-34.

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

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers, 2004. "Disagreement about Inflation Expectations," *NBER Macroeconomics Annual 2003,* 209-248.

Mankiw, N. Gregory and Ricardo Reis, 2010. "Imperfect Information and Aggregate Supply," in: Benjamin M. Friedman & Michael Woodford (ed.), *Handbook of Monetary Economics*, volume 3, chapter 5, pages 183-229, Elsevier.

McConnell, Margaret M. and Gabriel Perez-Quirós. 2000. "Output Fluctuations in the United States: What Has Changed since the Early 1980s?" *American Economic Review* 90(5), 1464-76.

Murphy, Kevin M. and Robert H. Topel, 1985. "Estimation and Inference in Two-Step Econometric Models," *Journal of Business and Economic Statistics* 3(4), 370-379.

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

Pagan, Adrian, 1984. "Econometric Issues in the Analysis of Regressions with Generated Regressors," *International Economic Review* 25(1), 221-247.

Patton, Andrew J. and Allan Timmermann, 2010. "Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion," *Journal of Monetary Economics* 57(7), 803-820.

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

Piazzesi, Monika and Martin Schneider, 2008. "Bond Positions, Expectations, and the Yield Curve," Working Paper 2008-02, Federal Reserve Bank of Atlanta.

Reis, Ricardo, 2006. "Inattentive Producers," *Review of Economic Studies* 73(3), 793-821.

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

Roberts, John M., 1997. "Is Inflation Sticky?" *Journal of Monetary Economics* 39(1), 173-196.

Roberts, John M., 1998. "Inflation Expectations and the Transmission of Monetary Policy," Finance and Economics Discussion Series 1998-43, Board of Governors of the Federal Reserve System.

Sarte, Pierre-Daniel, 2013. "When Is Sticky Information More Information?" *Journal of Money, Credit and Banking*, forthcoming.

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

Stock, James H. and Mark W. Watson. 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking* 39, s1: 13 – 33.

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

Woodford, Michael, 2002. "Imperfect Common Knowledge and the Effects of Monetary Policy," published in P. Aghion et al., eds., *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps,* Princeton Univ. Press.

Woodford, Michael, 2009. "Information-Constrained State-Dependent Pricing," *Journal of Monetary Economics* 56(S), S100-S124.

Table 1. Tests of the Inflation Expectations Process

Dependent variable Forecast error $\pi_{t+3,t} - F_t\pi_{t+3,t}$	Additional Control: z_{t-1}				
	None	Inflation	Average quarterly 3-month Tbill rate	Quarterly change in the log of the oil price	Average unemployment rate
	(1)	(2)	(3)	(4)	(5)
Panel A: $\pi_{t+3,t} - F_t\pi_{t+3,t} = c + \gamma F_t\pi_{t+3,t} + \delta z_{t-1} + error_t$					
Constant	-0.181	-0.045	-0.091	-0.181	1.449**
	(0.248)	(0.223)	(0.236)	(0.221)	(0.676)
$F_t\pi_{t+3,t}$	0.059	-0.299**	0.210*	0.045	0.095
	(0.085)	(0.148)	(0.111)	(0.078)	(0.085)
Additional Control: z_{t-1}		0.318**	-0.125*	1.603**	-0.281**
		(0.147)	(0.066)	(0.763)	(0.117)
Observations	178	178	178	178	178
R^2	0.010	0.109	0.054	0.046	0.148
Panel B: $\pi_{t+3,t} - F_t\pi_{t+3,t} = c + \beta\left(F_t\pi_{t+3,t} - F_{t-1}\pi_{t+3,t}\right) + \delta z_{t-1} + error_t$					
Constant	0.002	-0.074	0.151	-0.021	1.134**
	(0.144)	(0.174)	(0.175)	(0.146)	(0.546)
$F_t\pi_{t+3,t} - F_{t-1}\pi_{t+3,t}$	1.193**	1.141**	1.196**	1.125**	1.062**
	(0.497)	(0.458)	(0.504)	(0.499)	(0.465)
Additional Control: z_{t-1}		0.021	-0.029	0.576	-0.178**
		(0.050)	(0.031)	(0.608)	(0.076)
Observations	173	173	173	173	173
R^2	0.195	0.197	0.201	0.200	0.249

Notes: The table reports coefficient estimates for the specified equations at the top of each panel. The additional controls (z) are lagged by one quarter. Newey-West standard errors are in parentheses.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

Table 2. Testing for Higher Order Dynamics

	Order, p		
	1	2	3
AR(p)	1.075	1.089	1.097
VAR(p)			
Output growth rates, GY	1.104	1.128	1.178
Housing starts, HS	1.109	1.145	1.188
Unemployment rate, UE	1.108	1.157	1.198
Industrial production, IP	1.110	1.148	1.191
GY, HS, UE, IP	1.199	1.325	1.455

Notes: The table reports BIC statistics associated with different specifications of equation (14) in the text for AR(p) dynamics or equation (15) for VAR(p) dynamics. Lower values of the BIC indicate preferred specifications. See section I.D.1 in the text for details.

Table 3. Tests of Alternative Interpretations of Forecast Error Predictability

Dependent variable	OLS	IV	
Forecast error	(1)	(2)	(3)
Panel A: Test Heterogeneity in Beliefs about Long-Run Means			
$F_t\pi_{t+3,t} - F_{t-1}\pi_{t+3,t}$	1.193**	1.907***	2.095**
	(0.497)	(0.605)	(0.878)
π_{t-1}			-0.050
			(0.077)
Observations	173	173	173
1^{st} stage F-stat		14.07	9.24
Panel B: Test Forecast Smoothing			
$F_t\pi_{t+h} - F_{t-1}\pi_{t+h}$	0.62*	2.20***	2.23***
	(0.33)	(0.57)	(0.75)
$F_{t-1}\pi_{t+h} - F_{t-2}\pi_{t+h}$			-0.05
			(0.31)
Observations	529	529	529
1^{st} stage F-stat		28.45	13.56

Notes: The table reports estimates of equations (16) and (17) in Panels A and B respectively by OLS and instrumental variables (IV). In Panel A, the dependent variable is year-ahead forecast error for inflation $\pi_{t+3,t} - F_t\pi_{t+3,t}$. In Panel B, the dependent variable is *h*-period ahead forecast error for inflation $\pi_{t+h} - F_t\pi_{t+h}$. Panel B pools across horizons *h*=0 through *h*=2. In both panels, the instrumental variables are the first two lags of log change in oil prices. *1^{st} stage F-stat* shows the first stage fit. Column (3) in Panel A reports estimates of specification (16). Column (3) in Panel B reports estimates of specification (17). Driscoll-Kraay (1998) standard errors are in parentheses.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

Table 4. Information Rigidity in Inflation Forecasts by Forecaster Type.

Dependent variable Forecast error $\pi_{t+1,t} - F_t\pi_{t+1,t}$	Panel A: Livingston survey			
	Academic Institutions	Commercial Banks	Non-Financial Businesses	All Forecasters
	(1)	(2)	(3)	(4)
Forecast revision	0.476*	0.935***	0.572**	1.063***
	(0.242)	(0.197)	(0.251)	(0.251)
Sample	1969-2014	1969-2014	1969-2014	1969-2014
Observations	91	91	91	91

Dependent variable Forecast error $\pi_{t+4,t+1} - F_t\pi_{t+4,t+1}$	Panel B: Instrumental variable regression		
	Survey of Professional Forecasters (SPF)	Michigan Survey of Consumers (MSC)	Financial markets (FIN)
	(1)	(2)	(3)
Forecast revision	1.110***	0.705***	1.495*
	(0.401)	(0.260)	(0.833)
s.e.e.	1.155	1.258	1.651
1^{st} stage F-stat	28.22	57.76	7.912
Observations	127	127	126

Notes: Panel A reports estimates of equation (10) using inflation forecasts from the Livingston Survey at the biannual frequency. Columns (1)-(3) report OLS estimates using subsets of the forecasters while column (4) reports estimates using all forecasters in the survey. Panel B reports quarterly estimates of specification (18) with inflation forecasts by instrumental variables (IV). All results are for the CPI inflation rate (forecast errors, forecast revisions). Estimation sample is 1982-2014. The instrumental variable is the time-*t* change in oil prices. *1^{st} stage F-stat* shows the first stage fit. Newey-West robust standard errors are in parentheses.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

Table 5. Testing for Higher Order Dynamics

		Order, *p*		
		1	2	3
Variable-by-variable approach				
Inflation	AR(p)	1.075	1.089	1.097
	VAR(p)	1.198	1.324	1.454
Output	AR(p)	2.267	2.302	2.331
	VAR(p)	2.339	2.501	2.618
Industrial production	AR(p)	3.570	3.600	3.630
	VAR(p)	3.619	3.760	3.872
Unemployment rate	AR(p)	-1.181	-1.152	-1.131
	VAR(p)	-1.179	-1.030	-0.883
Housing starts	AR(p)	-3.816	-3.807	-3.772
	VAR(p)	-3.761	-3.642	-3.492
System approach				
	AR(p)	-0.090	-0.031	0.042
	VAR(p)	0.056	0.598	1.128

Notes: The table reports BIC statistics associated with different specifications of equation (20) in the text for AR(p) dynamics or equation (20) for VAR(p) dynamics. Lower values of the BIC indicate preferred specifications. In the system approach the measure of the fit is given by $\log|\mathbf{\Omega}|$ where $|\mathbf{\Omega}|$ is the determinant of the covariance matrix of residuals $\mathbf{\Omega}$. The AR(p) specification in the system approach corresponds to the system where only own lags of a variable are included as regressors. See section II.B in the text for details.

Table 6. Information Rigidity in Univariate and Multivariate Approaches.

	Univariate approach				Multivariate approach		P-value: equality of IR based on multivariate and univariate approaches
	Slope from the regression		Implied degree of information rigidity $1-G=\frac{\beta}{\beta+1}$		Implied degree of information rigidity $1-G$		
variable	estimate	s.e.	estimate	s.e.	estimate	s.e.	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Inflation	0.51	0.24	0.34	0.11	0.50	0.15	0.11
Output	0.34	0.19	0.30	0.09	0.16	0.41	0.74
Industrial production	0.28	0.15	0.22	0.09	-0.83	0.52	0.04
Unemployment rate	0.20	0.05	0.17	0.03	0.13	0.08	0.60
Housing starts	0.35	0.06	0.26	0.03	0.25	0.07	0.96

Notes: Standard errors are robust to heteroskedasticity and serial correlation. P-value of equality of estimated information rigidity across variables estimated by the multivariate approach is 0.033.

Table 7. The Macroeconomic Determinants of Information Rigidities

Dependent variable: estimated coefficient on forecast revisions for country-variable pairs	Revisions in data releases as a measure of noise		
	OLS	OLS exclude outliers	Robust regression
	(1)	(2)	(3)
Persistence of Series, $\rho_{i,j}$	-0.802***	-0.833***	-0.628*
	(0.282)	(0.248)	(0.317)
Noise-Signal Ratio, $\kappa_{i,j}$	0.153	0.464**	0.506*
	(0.304)	(0.201)	(0.255)
Observations	60	56	60
R-squared	0.153	0.326	

Notes: The table reports estimated specification (21). The persistence of each series ($\rho_{i,j}$) is estimated as the sum of AR(4) coefficients. The standard deviation of the difference between first and final data releases is taken as a measure of noise in the series. In column (1), four observations are identified as outliers: consumption growth rates for Italy, France, Germany and Japan. These outliers are dropped in estimation in column (2). In column (3), robust S-regressions are run with no dummies for outliers and all available observations included. Robust standard errors are in parentheses.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

Figure 1: Estimates of Information Rigidity by Horizon and Macroeconomic Variable

Panel A: Testing Equality across Forecasting Horizons

U.S. SPF Variables: 1968-2014

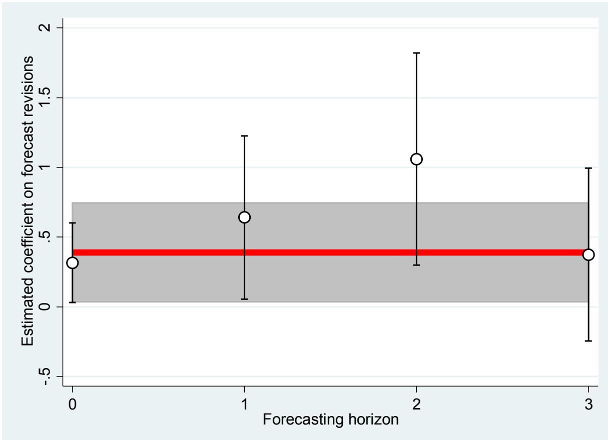

U.S. SPF Variables: 1981-2014

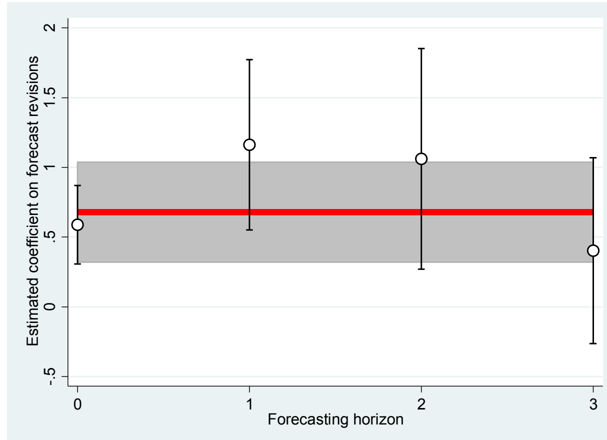

Panel B: Testing Equality across Macroeconomic Variables

U.S. SPF Variables: 1968-2014

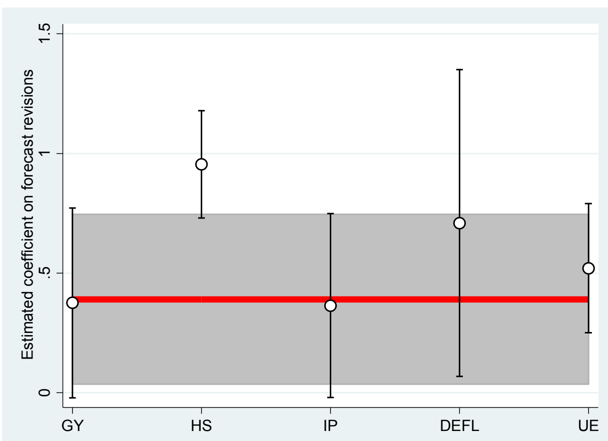

U.S. SPF Variables: 1981-2014

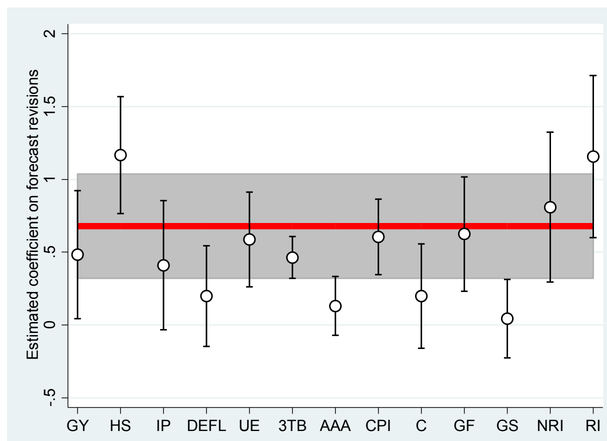

Notes: The figure plots estimated coefficient β on forecast revisions for different forecast horizons (Panel A) and macroeconomic variables (Panel B) in specification (10). Each circle presents a point estimate and whiskers show the 95 percent confidence interval. The solid red line is the point estimate of the coefficient on forecast revisions in specification (14) on pooled (across variables) sample with the shaded region showing the associated 95 percent confidence interval. All standard errors are Driscoll and Kraay (1998). GY = real GDP growth rate, HS = Housing starts, IP = Growth rate of industrial production index, DEFL = Inflation rate for GDP deflator, UE = Unemployment rate, 3TB = 3 month treasure bill interest rate, AAA = Interest rate on AAA debt, CPI = Inflation rate for the consumer price index, C = Consumption growth rate, GF = growth rate of federal government consumption expenditures, GS =

Growth rate of state government consumption expenditures, NRI = Growth rate of non-residential investment; RI = growth rate of residential investment.

Figure 2: International Evidence on Information Rigidities

Panel A: Cross-Country Estimates

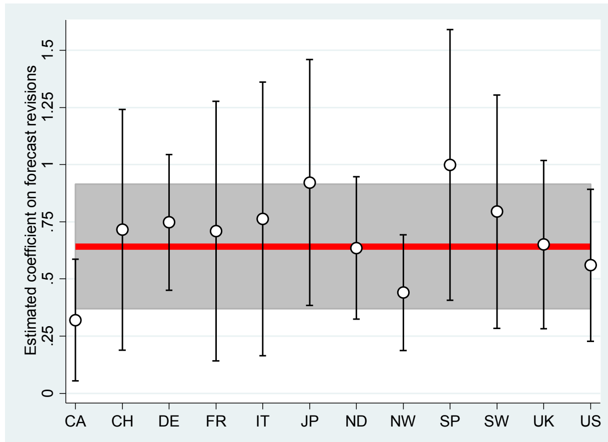

Panel B: Cross-Variable Estimates

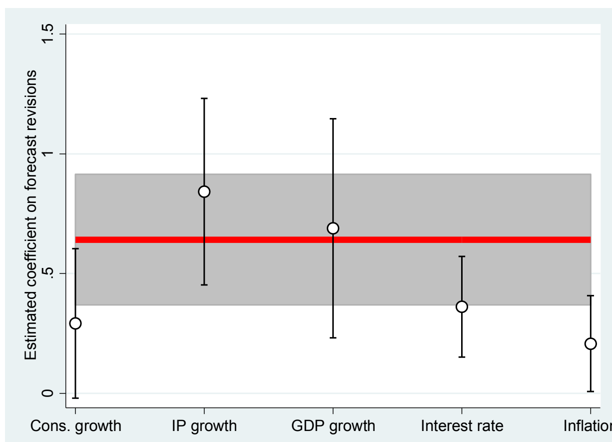

Notes: Each figure plots estimated coefficient β on forecast revisions in specification (10) from pooled specifications. In Panel A, estimates are for each country, pooled across variables and horizons. In Panel B, estimates are for each macroeconomic variable, pooled across countries and horizons. Each circle presents a point estimate for a given country and whiskers show the 95 percent confidence interval. The solid red line is the point estimate of the coefficient on forecast revisions pooled across all countries, variables and horizons with the shaded region showing the associated 95 percent confidence interval. All standard errors are Driscoll and Kraay (1998). In Panel A: CA = Canada, CH = Switzerland, DE = Germany, FR = France, IT = Italy, JP = Japan, ND = Netherlands, NW = Norway, SP = Spain, SW = Sweden, UK = United Kingdom, US = USA.

Figure 3: Information Rigidity and the Great Moderation

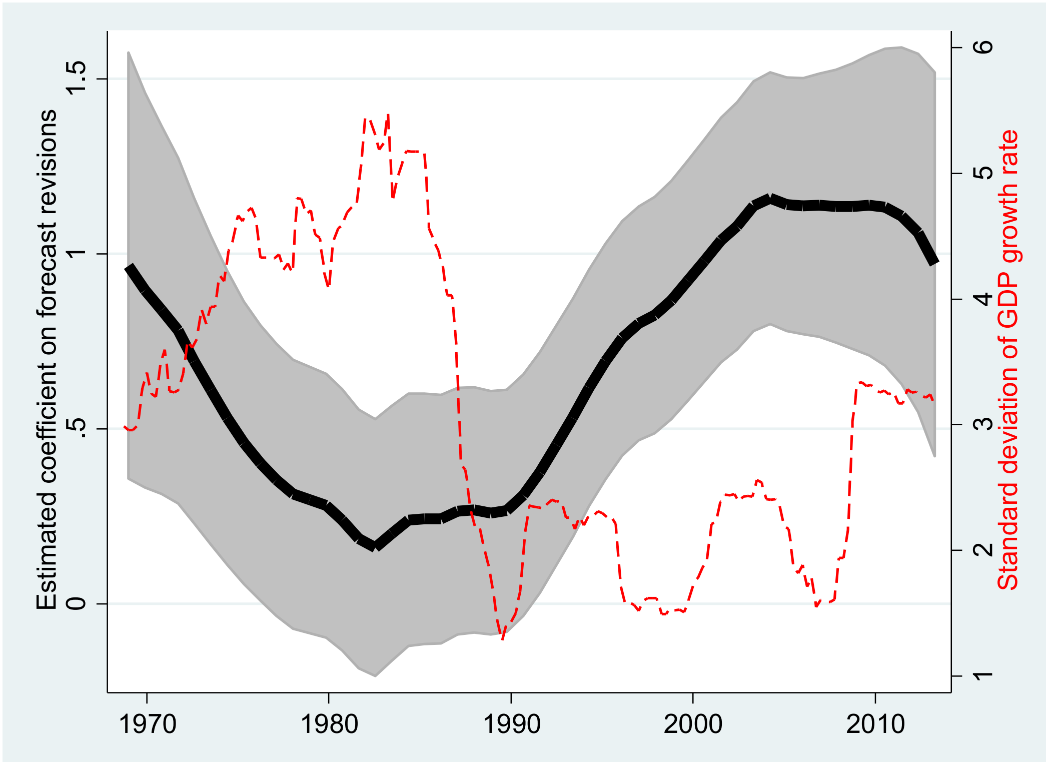

Notes: The figure plots the time series of two variables. The first is the standard deviation of the U.S. real GDP growth rate (annualized) over a five year moving window (red dash line; right axis). The second is the smoothed coefficient β_t on forecast revisions in specification (10) estimated for each quarter separately on the SPF data (black thick solid line; left axis). The shaded region is the 90 percent confidence interval. The smoother is a local average which uses Epanechnikov kernel with bandwidth equal to 3.5.

Figure 4: Information Rigidity during a Business Cycle

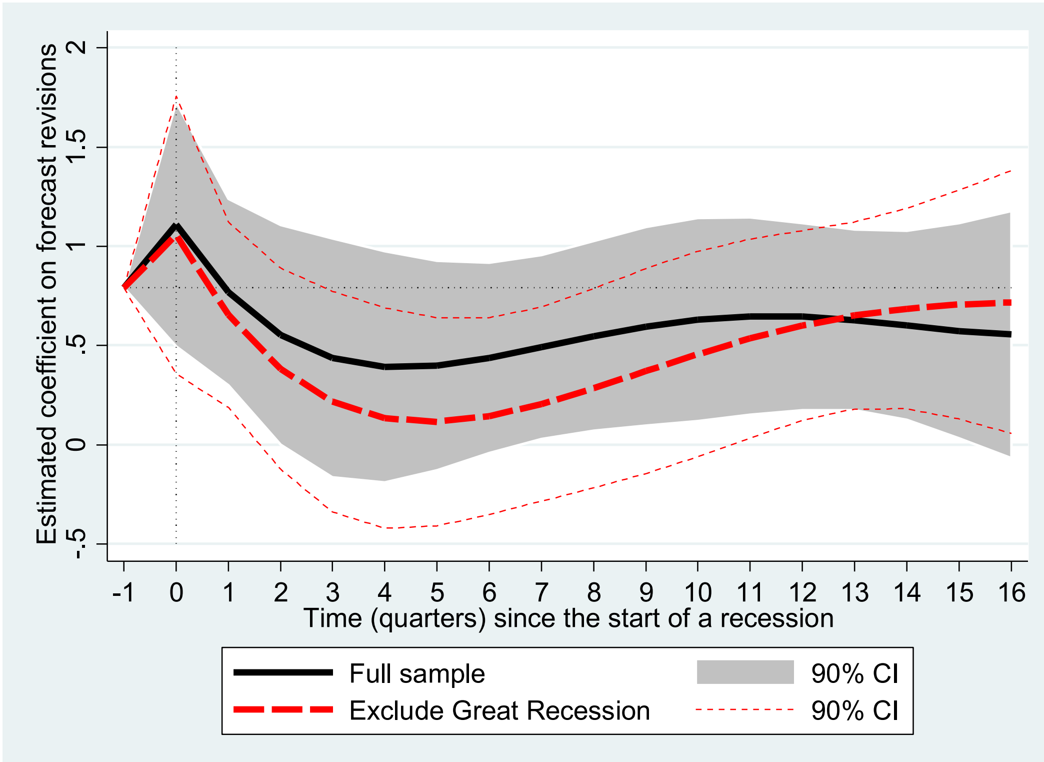

Notes: The figure plots the response of the coefficient β_t on forecast revisions in specification (10) estimated for each quarter separately on the SPF data. The response is estimated as in specification (22). The black line shows the response estimated on the full sample, while the red, thick, dashed line shows the response estimated on the sample that excludes the Great Recession. The response is normalized to be at the average value of the coefficient β_t one period before a recession starts. The shaded region is the 90 percent confidence interval for the full sample. The thin, red, dashed lines show the 90 percent confidence interval for the sample that excludes the Great Recession. The horizontal, thin, dashed line shows the average value of the coefficient β_t. The vertical, thin, dashed line shows the time when economy moves into a recession.