

What can survey forecasts tell us about information rigidities?

Olivier Coibion
College of William & Mary and NBER

Yuriy Gorodnichenko
UC Berkeley and NBER

This Draft: June 10th, 2011

Abstract: A lot. We derive a number of common and conflicting predictions from models in which agents face information constraints, then use surveys of forecasts from professional forecasters, consumers, firms and central bankers to assess their validity. Our key contribution is to document that in response to shocks, mean forecasts fail to completely adjust on impact, leading to statistically and economically significant deviations from the null of full information. The dynamic process followed by forecast errors after a shock is consistent with the predictions of models of information rigidities. In addition, we illustrate how the conditional responses of forecast errors and disagreement among agents can be used in conjunction with survey data of agents' forecasts to differentiate between some of the most prominent models of information rigidities.

Keywords: expectations, information rigidity, survey forecasts.

JEL: E3, E4, E5

We are grateful to anonymous referees, Monika Piazzesi, Alex Cukierman, Pierre-Olivier Gourinchas, Edward Knotek, Oleg Korenok, Maxym Kryshko, Bartosz Mackowiak, Peter Morrow, Ricardo Reis, David Romer, Mirko Wiederholt, Johannes Wieland and Justin Wolfers as well as seminar participants at the Kansas City Federal Reserve Bank, NBER, UC Berkeley, UC Davis, Michigan, ECB/CFS, VCU, and Yale for comments.

1 Introduction

How economic agents form their expectations has long been one of the most fundamental, and most debated, questions in macroeconomics. Indeed, the abandonment of adaptive expectations in favor of rational expectations was one of the defining features in the rebuilding of macroeconomics starting in the 1970s. Yet, even with the advent of rational expectations, research continued to emphasize the fact that, in forming their expectations, agents typically face constraints. For example, Lucas (1972) assumed agents could not observe all prices in the economy. Likewise, Kydland and Prescott (1982) assumed that agents could not differentiate in real time between transitory and permanent productivity shocks. Despite this early interest in the information problems faced by economic agents and their implications for aggregate dynamics, most modern macroeconomic models assume full-information rational expectations on the part of all agents. Yet recent work such as Mankiw and Reis (2002), Woodford (2001) and Sims (2003) has once more revived interest in better understanding the frictions and limitations faced by agents in the acquisition and processing of information.

This renewed interest in the expectations formation process has been spurred by several failures of full information models. For example, 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. Gorodnichenko (2006) shows a similar result in the context of economies with state-dependent pricing for which it is particularly hard to generate persistent and hump-shaped responses of inflation to nominal shocks. Similarly, Dupor, Han and Tsai (2009) show that the differential response of inflation to monetary policy and technology shocks is difficult to reconcile without information rigidities. In addition, departing from the assumption of full-information can account for some empirical puzzles. For example, Roberts (1997, 1998) and Adam and Padula (2003) 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. Similarly, Romer and Romer (2004) show that monetary policy shocks drawn from the Fed's Taylor rule conditional on its historical forecasts eliminate the price puzzle identified in previous work. Piazzesi and Schneider (2008), Gourinchas and Tornell (2004) and Bacchetta, Mertens, and van Wincoop (2009) all identify links between systematic forecast errors in survey forecasts and puzzles in various financial markets. Yet despite this resurgent focus on the nature of the expectations formation process, little empirical evidence exists on the size and nature of information rigidities.

This paper lays out a new set of stylized facts about the expectations formation process to address the two key issues: do agents have full information, and if not, how do we model their information

problem? We start by considering a large set of imperfect information models in which agents face different kinds of information rigidities and show that these models make clear predictions about the conditional response of agents' beliefs to economic shocks which can be used to characterize both the importance of and the nature of information rigidities faced by economic agents. In particular, we consider sticky information models a la Mankiw and Reis (2002) in which agents update their information sets infrequently as well as noisy information models as in Woodford (2001), Sims (2003) and Mackowiak and Wiederholt (2009a) in which agents are continuously updating their information but only observe noisy signals about the true state. In addition, we consider variants of the latter in which agents face strategic interactions in their forecasts, as in Morris and Shin (2002), have different priors about long-run means, as in Patton and Timmermann (2010), or face different signal-to-noise ratios and therefore process information at different rates.

Strikingly, all of these models make a common prediction: the average forecast across agents should respond more gradually to a shock than the variable being forecasted, i.e. the conditional response of the average forecast error across agents should be serially correlated and of the same sign as the forecasted variable. This is in direct contrast to the prediction under full-information rational expectations models in which agents would immediately process the new information such that the average forecast would respond by the same amount as the variable being forecasted. Using survey forecast data from U.S. professional forecasters, consumers, firms, and central bankers, we find robust evidence against the null of full-information consistent with the presence of information rigidities for each type of agent: forecast errors consistently move in the same direction as the variable being forecasted, be it inflation or unemployment, in response to a variety of macroeconomic shocks.

In addition to documenting pervasive and robust evidence consistent with information rigidities, we show that the underlying degree of information rigidity in each model can also be recovered from the conditional responses of forecast errors to shocks. The implied degrees of information rigidity are large and economically significant, particularly for inflation forecasts, and differ little across agent types or macroeconomic shocks conditioned over. In the context of sticky information models, for example, the estimated levels of information rigidity in inflation forecasts would imply that agents update their information sets less than once per year on average, while in the context of noisy information models the corresponding metric would be a weight consistently of or less than 0.2 on new information, i.e., it takes about three quarters to reduce the forecast error by a half. Furthermore, these estimates are unlikely to be driven by strategic considerations such as in Morris and Shin (2002). We replicate our analysis using inflation forecasts extracted from asset prices and document almost identical qualitative and quantitative results. Because financial market participants are unlikely to be trading on strategic interaction in

forecasts, this strongly suggests that this motive is not driving our empirical findings. As a result, the degree of information rigidity inherent in each type of agent's expectations formation process should have significant implications for macroeconomic dynamics and optimal policy.

We also derive a number of conflicting predictions from the different models of information rigidities to shed light on which forms of information constraints are most relevant to the expectations formation process for each type of agent. For example, as we demonstrate, the extensions of the baseline noisy information model to heterogeneity in priors about long-run means or signal strengths predict that forecast errors should be correlated with lags of the forecasted variable even after controlling for past forecast errors. This prediction receives no support in the data for any forecaster type and thus heterogeneity of agents in either their beliefs about long-run means or their signal-to-noise ratios are unlikely to play a prominent role in driving the expectations formation process for these agents. Second, we show that under sticky information, disagreement among agents should rise after any economic shock, whereas in noisy information models the amount of disagreement is independent of shocks (unless there is heterogeneity in signal-to-noise ratios). Consistent with noisy information models, we cannot generally reject the null of no response of disagreement to shocks but can reject the null that disagreement responds in the manner predicted under sticky information. Thus, these tests point to the basic noisy information model as the best characterization of the expectations formation process for professional forecasters, consumers, firms and central bankers alike.

We also show that the data are inconsistent with an alternative explanation for the gradual adjustment of forecasts and the presence of disagreement that does not rely on information rigidities. Specifically, Capistran and Timmermann (2009) argue that heterogeneity in loss-aversion across agents potentially explains these features of the data. However, we demonstrate that in their setting the conditional response of forecast errors should always be of the same sign, regardless of whether a shock raises or lowers a forecasted variable. We test this prediction by examining the conditional response of forecast errors to the absolute value of shocks rather than the levels. Whereas forecast errors respond to the levels of the shocks, we find little evidence of a consistent or significant response to the absolute values of the shocks thus indicating that information rigidities play a more important role in the expectations formation process than heterogeneous loss-aversion.

This paper is closely related to a number of recent papers on the expectations formation process and information rigidities. Korenok (2006), Kiley (2007), Klenow and Willis (2007), Coibion (2010), Dupor, Kitamura and Tsuruga (2010), and Knotek (2010) assess the potential empirical importance of sticky information for price-setting decisions whereas Mackowiak, Moench and Wiederholt (2009) compare the predictions of sticky information and noisy information models for the differential

persistence in sectoral price levels. Unlike these papers, our tests of the expectations formation process exploit the availability of survey data on agents' forecasts in a manner that *does not* hinge on auxiliary assumptions about the rest of the model, such as the nature of price-setting decisions. Mankiw, Reis and Wolfers (2004) also rely on survey data to assess how well the sticky information model can replicate some features of professional and consumer forecasts. Carroll (2003) proposes an epidemiological foundation for sticky information among consumers and tests it using survey data from professionals and consumers. Branch (2007) uses disagreement among consumers to distinguish between sticky information and other expectation models while Pesaran and Weale (2006) study more traditional tests of the rationality of survey data on expectations. Andrade and Le Bihan (2010) quantify the frequency at which professional forecasters change their forecasts while Coibion and Gorodnichenko (2010) apply a novel test of the expectations formation process to U.S. and international survey data from professional forecasters. We differ from this previous work in a number of ways. First, we consider a much larger set of theoretical models which deliver a number of new testable predictions to assess both the quantitative importance and the nature of information rigidities. Second, and in contrast to all previous work, we focus on the conditional response of forecasts to shocks. Third, we consider forecasts from a number of different kinds of economic agents, including professional forecasters, firms, consumers and central bankers.

The rest of the paper is structured as follows. In section 2, we present models of information rigidity and compare their predictions about conditional forecast errors and the response of the cross-sectional dispersion of beliefs after a shock. In section 3, we discuss our empirical methodology and data, and present benchmark results for forecasts of professional forecasters. Section 4 contains additional results for forecasts of consumers, firms and central bankers. Section 5 includes extensions and robustness checks of our results. Section 6 concludes.

2 Models of the Expectations Formation Process

In this section we lay out some key models of information rigidities and derive the implications of these models for the behavior of mean forecasts, forecast errors, and forecast disagreement. We first consider the sticky information model of Mankiw and Reis (2002) and the noisy information model as in Woodford (2001). We then present three extensions of the noisy information model: strategic interaction in forecasts, heterogeneous priors about long-run means, and heterogeneous signal-to-noise ratios. Finally, we consider a full-information model in which agents have heterogeneous loss-aversion.

2.1 Sticky Information

Reis (2006) considers the problem of a firm facing a fixed cost to acquiring and processing new information. In the presence of fixed costs, it becomes optimal for firms to update their information infrequently. Under certain conditions, Reis shows that the acquisition of information follows a Poisson process in which, each period, agents face a constant probability λ of not being able to update their information. We refer to λ as the degree of information rigidity for the sticky information model. Following Mankiw and Reis (2002), we assume that when agents update their information sets, they acquire complete information and form expectations rationally. In periods in which agents do not update their information sets, their expectations and actions continue to be based on their old information. Thus, agents who update their information sets in the same period have the same beliefs and forecasts about macroeconomic variables.

Suppose that inflation π_t is the variable of interest and it follows an AR(1) process:¹

$$\pi_t = \rho\pi_{t-1} + w_t \quad (1)$$

where $\{w_{t-s}\}_{s=0}^{\infty}$ is a sequence of shocks. The impulse response of inflation at time $t+k$ to a shock at time t is given by

$$\frac{d\pi_{t+k}}{dw_t} = \rho^k, \forall k \geq 0. \quad (2)$$

Denote the optimal h period ahead forecast for inflation at time t given agent i 's information with $\pi_{t+h|t}(i) \equiv E(\pi_{t+h}|I_{i,t})$. Since agents update their information at a Poisson rate λ , the mean forecast across agents at time t of inflation h periods ahead, which we denote with $\overline{\pi_{t+h|t}}$, is a weighted average of past (rational) expectations of the variable at time $t+h$:

$$\overline{\pi_{t+h|t}} = \bar{E}(\pi_{t+h|t}(i)) = (1-\lambda) \sum_{k=0}^{\infty} \lambda^k E_{t-k} \pi_{t+h}, \quad (3)$$

where $\bar{E}(\cdot)$ indicates that the expectation is taken over agents rather than time, $E_{t-k} \pi_{t+h} = \sum_{s=k}^{\infty} \rho^{s+k} w_{t-s}$ and thus

$$\overline{\pi_{t+h|t}} = \sum_{k=0}^{\infty} \rho^{k+h} w_{t-k} (1-\lambda^{k+1}). \quad (4)$$

The mean forecast depends on the average response of inflation, since when agents update their information sets, they acquire full information. Thus, after an inflationary shock, mean forecasts rise along with inflation. Because λ is less than one, the mean forecast under-reacts to a shock relative to the actual response of inflation and the coefficient $\rho^{k+h}(1-\lambda^{k+1})$ on shock w_{t-k} converges to ρ^{k+h} over time, so mean forecasts converge to the true value.

Given equations (4) and (1), the forecast error $FE_{t,t+h} \equiv \pi_{t+h} - \overline{\pi_{t+h|t}}$ obeys

¹ In Appendix A we show that results generalize for any MA(∞) process.

$$FE_{t,t+h} = \sum_{m=0}^{h-1} \rho^m w_{t+h-m} + \sum_{k=0}^{\infty} \rho^{k+h} w_{t-k} \lambda^{k+1}, \quad (5)$$

and consequently the impulse response of the forecast error to shocks is

$$\frac{dFE_{t+j,t+j+h}}{dw_t} = \rho^{j+h} \lambda^{j+1} = \left(\frac{d\pi_{t+j+h}}{dw_t} \right) \lambda^{j+1} \quad (6)$$

Forecast errors depend both on the inflation process after the shock, as well as the degree of information rigidity. When $\lambda = 0$, firms always update their information sets and the ex-post forecast error is zero on average. As the degree of information rigidity rises, conditional forecast errors will become increasingly persistent.

The impulse response for forecast errors above also makes clear that the convergence of the forecast error to the true value is independent of the volatility of the shock. Specifically, the response of the forecast error normalized by the response of inflation

$$\frac{dFE_{t+j,t+j+h}/dw_t}{d\pi_{t+j+h}/dw_t} = \lambda^{j+1} \quad (7)$$

is monotonically decreasing over time at a rate governed by the degree of information rigidity. Because agents must choose a certain average duration between information updates, this convergence rate is independent of the properties of the shock. In other words, two different kinds of shocks must yield the same convergence rate for mean forecast errors.

The sticky information model also makes predictions about the cross-sectional dispersion of beliefs across agents. Define $V_t \pi_{t+h|t}(i) = \overline{\text{var}}(\pi_{t+h|t}(i))$ to be the cross-sectional variance of h period ahead forecasts at time t for π where $\overline{\text{var}}(\cdot)$ denotes that variance is taken across agents. Then,

$$V_t \pi_{t+h|t}(i) = (1 - \lambda) \sum_{k=0}^{\infty} \lambda^k \left(E_{t-k} \pi_{t+h} - \overline{\pi_{t+h|h}} \right)^2. \quad (8)$$

The impulse response of the cross-sectional variance of h period ahead forecasts at time $t+j$ to a shock δ at time t is given by

$$\rho^{2(j+h)} \lambda^{j+1} (1 - \lambda^{j+1}) \delta^2 = \left(\frac{d\pi_{t+j+h}}{d\delta} \right)^2 \lambda^{j+1} (1 - \lambda^{j+1}) \delta^2. \quad (9)$$

As long as $\lambda > 0$, the dispersion, or degree of disagreement across agents, will rise in response to a shock, regardless of whether the shock is inflationary or disinflationary. Over time (assuming inflation does not explode), the dispersion will return to its steady-state level.

2.2 Noisy Information

Lucas (1972), Woodford (2001) and Sims (2003) develop models where economic agents filter the state of economic fundamentals from a series of signals contaminated with noise and hence we refer to this class of models as “noisy information”. In contrast to Mankiw and Reis (2002), agents *continuously* track variables and incorporate the most recent information into their decision making. The striking feature of

this class of models is that the dispersion of forecasts can be invariant to shocks in fundamentals. In this section, we present a simple model to illustrate the intuition behind this result.

Suppose that economic agents observe noisy signals about inflation $z_{it} = [y_{it} \ s_t]'$ with $y_{it} = \pi_t + v_{it}$ and $s_t = \pi_t + \eta_t$ where $v_{it} \sim iid N(0, \sigma_v^2)$ is an agent specific, private shock, and $\eta_t \sim iid N(0, \sigma_\eta^2)$ is a shock common for all agents. Also to simplify algebra, suppose that inflation follows an AR(1) process $\pi_t = \rho\pi_{t-1} + w_t$ where $w_t \sim N(0, \sigma_w^2)$.² Denote the optimal forecast for inflation at time t given agent i 's information at time k with $\pi_{t|k}(i) = E(\pi_t | z_{i,k}, z_{i,k-1}, \dots)$ and correspondingly $z_{t|k}(i) = E(z_{i,t} | z_{i,k}, z_{i,k-1}, \dots)$. Using properties of the Kalman filter, one can show that the forecast for the unobserved state π_t evolves as follows:

$$\begin{aligned} \pi_{t|t}(i) &= \pi_{t|t-1}(i) + P(z_{it} - z_{t|t-1}(i)) = (1 - PH)\pi_{t|t-1}(i) + PH\pi_t + P \begin{bmatrix} v_{it} \\ \eta_t \end{bmatrix} \\ &= (1 - PH)\rho\pi_{t-1|t-1}(i) + PH\pi_t + P \begin{bmatrix} v_{it} \\ \eta_t \end{bmatrix}, \end{aligned} \quad (10)$$

where $H = [1 \ 1]'$, $P = [P_\eta \ P_v] = \begin{bmatrix} \frac{\Psi\sigma_v^2}{\Psi(\sigma_v^2 + \sigma_\eta^2) + \sigma_v^2\sigma_\eta^2} & \frac{\Psi\sigma_\eta^2}{\Psi(\sigma_v^2 + \sigma_\eta^2) + \sigma_v^2\sigma_\eta^2} \end{bmatrix}$ is the gain of the Kalman filter (with $P_\eta, P_v \in (0,1)$) and the variance-covariance matrix for the one-step ahead forecast error for π_t is $\Psi = \rho \left\{ \Psi - \Psi H' (H \Psi H' + \text{diag}\{\sigma_v^2, \sigma_\eta^2\})^{-1} H \Psi \right\} \rho + \sigma_w^2$. Given multiple signals, we interpret $PH \in (0,1)$ (rather than P) as the degree of information rigidity. The gain of the filter does not vary across agents because all agents solve the same Ricatti equation and thus obtain the same gain P .

The average forecast for the current state of inflation given current information is then given by

$$\begin{aligned} \overline{\pi_{t|t}} &= \bar{E} \left(\pi_{t|t}(i) \right) = (1 - PH)\rho\bar{E} \left(\pi_{t-1|t-1}(i) \right) + PH\pi_t + P E_i \begin{bmatrix} v_{it} \\ \eta_t \end{bmatrix} \\ &= (1 - PH)\rho\overline{\pi_{t-1|t-1}} + PH\pi_t + P \begin{bmatrix} 0 \\ \eta_t \end{bmatrix} = \sum_{k=0}^{\infty} \{(1 - PH)\rho\}^{k-1} (PH\pi_t + P_\eta\eta_t) \\ &= \{1 - (1 - PH)\rho L\}^{-1} [\{1 - \rho L\}^{-1} PH w_t + P_\eta \eta_t] \end{aligned} \quad (11)$$

where L denotes the lag operator. Since $PH \in (0,1)$, the mean forecast moves in the same direction as actual inflation in response to a shock w_t but does so by a smaller amount than actual inflation. In other words, the mean forecast of inflation under-reacts to shocks relative to actual inflation.

Similar to the sticky information model, this model predicts that the average forecast error follows an AR(1) process:

$$\begin{aligned} FE_{t,t} &= \pi_t - \overline{\pi_{t|t}} = (1 - PH)(\pi_t - \rho\overline{\pi_{t-1|t-1}}) - P_\eta\eta_t \\ &= (1 - PH)\rho(\pi_{t-1} - \overline{\pi_{t-1|t-1}}) + (1 - PH)w_t - P_\eta\eta_t \\ &= (1 - PH)\rho FE_{t-1,t-1} + (1 - PH)w_t - P_\eta\eta_t. \end{aligned}$$

² Results for a general model AR(p) are available in Appendix A.

Because $PH \in (0,1)$, the forecast error moves in the same direction as the mean forecast does. Note that the forecast error in response to a shock w_t converges to zero with time as

$$\frac{dFE_{t+j,t+j}}{dw_t} = \{(1 - PH)\rho\}^j(1 - PH) \rightarrow 0$$

as $j \rightarrow \infty$. Thus, the impulse response of forecast errors under noisy information follows the same qualitative pattern as under sticky information. An additional similarity to the sticky information model is that the dynamics of the forecast error normalized by the inflation rate in the noisy information model are only a function of PH which is governed by the degree of information rigidity:

$$\frac{\partial(\pi_{t+j} - \overline{\pi_{t+j|t+j}})/\partial w_t}{\partial \pi_{t+j}/\partial w_t} = \frac{(1-PH)^{j+1}\rho^{j-1}}{\rho^{j-1}} = (1 - PH)^{j+1}, \quad j \geq 0. \quad (12)$$

These insights suggest that information rigidity measured by PH can be interpreted in two useful ways. First, $1 - PH$ captures the fraction of the signals incorporated contemporaneously into the revised estimate of the current unobserved fundamental π_t . Second, $1 - PH$ measures the persistence of beliefs in addition to the persistence determined by the fundamentals, which is equal to ρ in the present context. For instance, one may use $\ln(0.5)/\ln(1 - PH)$ to calculate the half-life for forecast errors after controlling for the dynamics of forecasted series which makes comparison of information rigidities across forecasted series feasible.

Using equation (10), we can derive the law of motion for the dispersion of forecasts across agents:

$$V_t \pi_{t|t}(i) = \overline{\text{var}} \left\{ (1 - PH)\rho \overline{\pi_{t-1|t-1}} + PH\pi_t + P \begin{bmatrix} v_{it} \\ \eta_t \end{bmatrix} \right\} = \{(1 - PH)\rho\}^2 V_{t-1} \pi_{t-1|t-1}(i) + P_v^2 \Sigma_v.$$

Note that $V_t \pi_{t|t}(i)$ does not depend on π_t and thus π_t does not affect the evolution of forecast dispersion. Intuitively, because agents continuously update their information sets, the disagreement in their forecasts arises only due to idiosyncratic differences in information sets induced by shocks v_{it} . Since the dispersion of v_{it} does not vary in response to shocks to fundamentals such as π_t , the forecast disagreement does not respond to π_t .³

In the context of the present model, one can show that the gain of the Kalman filter (information rigidity) is increasing (decreasing) in the signal-to-noise ratio and persistence of the fundamental process. More generally, one can show that if the agent's objective function (e.g., profit or utility) is more sensitive to certain types of fundamental shocks (e.g. technology) than to other types of fundamental shocks (e.g., monetary policy), then the reaction to sensitive shocks is stronger (see Mackowiak and

³ The dispersion of forecasts can respond to shocks in the noisy information model if shocks induce conditional heteroskedasticity $\Sigma_{v,t} = \Sigma_v(\pi_t)$, which is similar in spirit to heteroskedasticity analyzed in GARCH models. Cukierman and Wachtel (1979) present such a model. We focus on the model without heteroskedastic effects because it offers sharper predictions.

Wiederholt (2009b)). Thus, unlike the sticky information model, the noisy information model allows for a differential response of information acquisition to fundamental shocks. For example, agents may learn slowly the true state of monetary policy but may react quickly to shocks in technology.

2.3 Extensions

In this section, we examine several modifications of the noisy information model considered in the literature. We also explore alternative models that generate forecast disagreement—which is often interpreted as *prima facie* evidence for agents using differential information sets—because agents use different models to construct predictions.

2.3.1 Public signals and strategic interaction

In the baseline setup of the noisy information model we assume that there is no strategic interaction across agents. Morris and Shin (2002) show that introducing strategic interaction in noisy information models can change the qualitative behavior of these models. To explore the potential implications of strategic interaction, we follow Morris and Shin (2002) and suppose that there is an incentive to stay close to the average action (or average forecast) so that the objective function of agent i is to report forecast $\tilde{\pi}_{t|t}(i)$ such that it minimizes

$$E \left\{ \left(\pi_t - \tilde{\pi}_{t|t}(i) \right)^2 + R \left(\tilde{\pi}_{t|t}(i) - \overline{\tilde{\pi}_{t|t}} \right)^2 \middle| I_{i,t} \right\} \quad (13)$$

where the second term is the penalty for deviating from the average *reported* forecast $\overline{\tilde{\pi}_{t|t}}$, $I_{i,t}$ denotes the information set of agent i at time t . One can interpret $R > 0$ as capturing strategic complementarity. If $R = 0$, the objective function reduces to minimization of the mean squared error of forecasts. Note that, consistent with the practice of collecting survey measures of forecasts, $\overline{\tilde{\pi}_{t|t}}$ is not observed when an agent prepares his forecast and each agent guesses what other agents forecast.

The first order condition for the optimal reported forecast is

$$\tilde{\pi}_{t|t}(i) = \frac{1}{1+R} E \{ \pi_t | I_{i,t} \} + \frac{R}{1+R} E \{ \overline{\tilde{\pi}_{t|t}} | I_{i,t} \} = \frac{1}{1+R} \pi_{t|t}(i) + \frac{R}{1+R} E \{ \overline{\tilde{\pi}_{t|t}} | I_{i,t} \}. \quad (14)$$

Note that $\pi_{t|t}(i) \equiv E \{ \pi_t | I_{i,t} \}$ is the best forecast agent i can generate given his information set $I_{i,t}$. If we average $\tilde{\pi}_{t|t}(i)$ across agents and use repeated substitution in (14) as was done in Morris and Shin (2002) and Woodford (2001), we can express the average reported forecast as

$$F_t \equiv \overline{\tilde{\pi}_{t|t}} = \frac{1}{1+R} \sum_{k=0}^{\infty} \left(\frac{R}{1+R} \right)^k \bar{E}^{(k)} \{ \pi_t \} = \frac{1}{1+R} \overline{\pi_{t|t}} + \frac{R}{1+R} \overline{F_{t|t}} \quad (15)$$

where $\bar{E}^{(k)} \{ \pi_t \}$ is the k^{th} -order expectation of inflation, $\overline{F_{t|t}} \equiv \bar{E} (F_t | I_{i,t})$. Since F_t is not observed, agents infer F_t from a sequence of observed signals.

Following Woodford (2001), we guess and verify the law of motion for F_t and other relevant variables. Specifically, we conjecture that the state evolves according to

$$X_t \equiv \begin{bmatrix} \pi_t \\ F_t \\ u_t \end{bmatrix} = MX_{t-1} + m \begin{bmatrix} W_t \\ \eta_t \end{bmatrix} \quad (16)$$

and the measurement equations are given by⁴

$$z_{it} \equiv \begin{bmatrix} s_t \\ y_{it} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} X_t + \begin{bmatrix} 0 \\ v_{it} \end{bmatrix} = HX_t + \begin{bmatrix} 0 \\ v_{it} \end{bmatrix}. \quad (17)$$

Given the structure of the problem, we consider

$$M = \begin{bmatrix} \rho & 0 & 0 \\ G & Q & 0 \\ 0 & 0 & 0 \end{bmatrix} \text{ and } m = \begin{bmatrix} 1 & 0 \\ m_{21} & m_{22} \\ 0 & 1 \end{bmatrix}. \quad (18)$$

Given the conjectured structure of the system in equations (16) and (17), agent i uses the Kalman filter to infer the state. Specifically, the posterior estimate of the state by agent i is

$$\begin{aligned} X_{t|t}(i) &= X_{t|t-1}(i) + P \left(z_{it} - z_{i,t|t-1}(i) \right) \\ &= (I - PH)MX_{t-1|t-1}(i) + PHMX_{t-1} + PHm \begin{bmatrix} W_t \\ \eta_t \end{bmatrix} + P \begin{bmatrix} 0 \\ v_{it} \end{bmatrix} \end{aligned} \quad (19)$$

where P is the gain of the Kalman filter. We take the average of (19) across agents to find the law of motion for the average estimate of the current state:

$$\overline{X_{t|t}} = (I - PH)M\overline{X_{t-1|t-1}} + PHMX_{t-1} + PHm \begin{bmatrix} W_t \\ \eta_t \end{bmatrix} \quad (20)$$

where η_t does not wash out because it is a shock common across agents.

Define $\xi \equiv \begin{bmatrix} \frac{1}{1+R} & \frac{R}{1+R} & 0 \end{bmatrix}$ and note that one can write (15) as

$$\begin{aligned} F_t &= \xi \overline{X_{t|t}} = \xi(I - PH)M \begin{bmatrix} \overline{\pi_{t-1|t-1}} \\ \frac{1+R}{R} F_{t-1} - \frac{1}{R} \overline{\pi_{t-1|t-1}} \\ \overline{u_{t-1|t-1}} \end{bmatrix} + \xi PHMX_{t-1} + \xi PHm \begin{bmatrix} W_t \\ \eta_t \end{bmatrix} \\ &= \left\{ \rho \left(\frac{1}{1+R} - C_1 \right) + \frac{R}{1+R} G - \frac{1}{1+R} Q \right\} \overline{\pi_{t-1|t-1}} + QF_{t-1} + \rho C_1 \pi_{t-1} + C_1 W_t + C_2 \eta_t. \end{aligned} \quad (21)$$

where we used (15) to replace $\overline{F_{t-1|t-1}} = \frac{1+R}{R} F_{t-1} - \frac{1}{R} \overline{\pi_{t-1|t-1}}$ and we defined $C_1 \equiv \frac{P_{11} + P_{12} + R(P_{21} + P_{22})}{1+R}$

and $C_2 \equiv \frac{P_{11} + RP_{21}}{1+R}$. To be consistent with conjectures in (16) and (18), the coefficients in (21) must satisfy

$$G = \rho C_1, m_{21} = C_1, m_{22} = C_2, \left\{ \rho \left(\frac{1}{1+R} - C_1 \right) + \frac{R}{1+R} G - \frac{1}{1+R} Q \right\} = 0 \Rightarrow Q = \rho(1 - C_1).$$

Given the law of motion (16) and measurement equations (17), one can show that the covariance matrix for the one-step ahead forecast error Ψ solves the following Ricatti equation

⁴ Note that we introduce an additional state variable u_t to handle the correlation of common shocks η_t in the signal s_t and state F_t .

$$\begin{aligned}\Psi &\equiv E[(X_{t|t-1}(i) - X_t)(X_{t|t-1}(i) - X_t)'] = \\ &= M \left\{ \Psi - \Psi H' \left(H \Psi H' + \begin{bmatrix} 0 & 0 \\ 0 & \sigma_v^2 \end{bmatrix} \right)^{-1} H \Psi \right\} M' + m \begin{bmatrix} \sigma_w^2 & 0 \\ 0 & \sigma_\eta^2 \end{bmatrix} m\end{aligned}\quad (22)$$

and the gain of the Kalman filter is

$$P = \Psi H' \left(H \Psi H' + \begin{bmatrix} 0 & 0 \\ 0 & \sigma_v^2 \end{bmatrix} \right)^{-1} . \quad (23)$$

Note that M, m are functions of P and structural parameters $\{\rho, R, \sigma_w^2, \sigma_\eta^2, \sigma_v^2\}$. This is a nonlinear system of equations and thus it is difficult to derive general analytical results for how structural parameters affect properties of the system. However, similar to the baseline case without strategic interaction, one can demonstrate that $C_1, C_2 \in (0,1)$ so that the forecast error, mean forecast and actual inflation all move in the same direction in response to shock w_t . Because the only source of disagreement is the private signal y_{it} , disagreement across forecasters does not move in response to shocks, which is similar to the baseline model with noisy information. Also similar to the baseline case, forecast errors for *reported* forecasts follow an AR(1) process:

$$\begin{aligned}\widetilde{F}E_{t,t} &\equiv \pi_t - \overline{\pi_{t|t}} = \pi_t - F_t = [1 \quad -1 \quad 0]X_t = [1 \quad -1 \quad 0]MX_{t-1} + [1 \quad -1 \quad 0]m \begin{bmatrix} w_t \\ \eta_t \end{bmatrix} \\ &= \rho(1 - C_1)\widetilde{F}E_{t-1,t-1} + (1 - C_1)w_t - C_2\eta_t\end{aligned}\quad (24)$$

where we used $\frac{1+R}{R}(\overline{\pi_{t|t}} - F_t) = \overline{\pi_{t|t}} - \overline{F_{t|t}}$ which follows from (15). Thus, incorporating the possibility of strategic interaction in forecasts does not qualitatively alter the predictions of the noisy information model since it continues to predict that forecast errors will be serially correlated after shocks, of the same sign as the response of the variable being forecasted and asymptotically vanish. However, the presence of strategic interaction in forecasts implies that the persistence of the conditional response of forecast errors, when normalized by the response of the variable being forecasted, will no longer directly identify the underlying degree of information rigidity. Instead, these estimates will reflect a combination of strategic interaction and information rigidities. In particular, $R > 0$ can amplify the persistence of the serial correlation relative to what would arise solely from information rigidities. Note, however, that strategic interaction by itself is not enough to generate serial correlation in forecast errors. For example, if information is not noisy, every forecaster will report the rational expectations predictions which do not have serially correlated forecast errors.

2.3.2 Disagreement about means

Patton and Timmermann (2010) observe that one can also generate disagreement in forecasts if forecasters have different beliefs about the long-run behavior of the forecasted variables. Following Patton and Timmermann (2010), suppose that agents report

$$\check{\pi}_{t|t}(i) = \omega\mu_i + (1 - \omega)\pi_{t|t}(i) \quad (25)$$

where $\mu_i \sim (0, \sigma_\mu^2)$ is a prior of forecaster i , $\omega = \frac{\Psi}{\Psi + \kappa^2}$ is the shrinkage factor, and $\pi_{t|t}(i)$ is the forecast generated by the Kalman filter in the baseline noisy information model.⁵ Equation (25) suggests that even if agents observe the same signals, these agents will report different forecasts as a result of their heterogeneous priors. Given $\bar{E}\mu_i = 0$, $\bar{E}\check{\pi}_{t|t}(i) = \bar{\pi}_{t|t} = (1 - \omega)\bar{\pi}_{t|t}$, the mean forecast error follows

$$\begin{aligned} FE_{t,t} = \pi_t - \bar{\pi}_{t|t} &= (1 - PH)(1 - \omega)\rho FE_{t-1,t-1} + (1 - PH + PH\omega)w_t + \omega\rho PH\pi_{t-1} \\ &\quad - (1 - \omega)P_\eta\eta_t. \end{aligned} \quad (26)$$

As in the previous models, equation (26) implies that forecast errors should respond to a shock in the same direction as ex-post inflation. However, in contrast to the baseline model, the forecast error should also be correlated with the past level of inflation—a testable implication which we examine later in the paper.

2.3.3 Heterogeneous precision of signals

One can also generate disagreement in forecasts if agents receive signals of different precision so that the interpretation of the same signal will vary across agents. To simplify the argument, we take the baseline noisy information model and assume that i) $\rho = 1$; ii) there is no common signal s_t ; iii) the precision of signals v_{it} (i.e., σ_v^{-2}) varies across agents in such a way that the gain of the Kalman filter across agents is approximately distributed as $P_i \sim (P, \sigma_p^2)$ which is independent from v_{it} and w_t . In this setting, the inflation forecast for agent i is given by

$$\pi_{t|t}(i) = (1 - P_i)\pi_{t-1|t-1}(i) + P_i\pi_t + P_iv_{it} = P_i\sum_{k=0}^{\infty}(1 - P_i)^k(\pi_{t-k} + v_{i,t-k}). \quad (27)$$

Given this setup, we show in Appendix A that the mean forecast, mean forecast error and forecast disagreement should approximately follow:

$$\bar{\pi}_{t|t} = \{\sum_{k=0}^{\infty} A_k(P)\pi_{t-k-1}\} + (1 - P)\bar{\pi}_{t-1|t-1} + P\pi_t, \quad (28)$$

$$FE_t \equiv \pi_t - \bar{\pi}_{t|t} = (1 - P)FE_{t-1} + (1 - P)w_t + \{\sigma_p^2 \sum_{k=0}^{\infty} A_k(P)\pi_{t-k-1}\}, \quad (29)$$

$$\overline{\text{var}}(\pi_{t|t}(i)) \equiv \bar{E}(\pi_{t|t}(i) - \bar{\pi}_{t|t})^2 = Q_t^2\sigma_p^2 + \text{const}, \quad (30)$$

where

$$A_k(P) = (1 - P)^{k-1}(1 - P - kP),$$

⁵ Like Patton and Timmermann (2010), we also assume that $E(\mu_i v_{it}) = 0$, i.e., signals and priors are independent.

$$Q_t \equiv \{\sum_{k=0}^{\infty} A_k(P)\pi_{t-k}\}.$$

Although expressions in equations (28)-(30) are complex, several qualitative results stand out. First, this model yields the same qualitative prediction that forecasts will adjust more gradually to shocks than the variable being forecasted, leading to a sequence of serially correlated forecast errors after a shock. This follows from $P \in (0,1)$ and equation (29). Second, if one projects the mean forecast error on its lag, the error in this regression should be correlated with lags of inflation. This correlation arises from the sequence of non-zero $A_k(P)$ multiplying lags of inflation in equation (29). In contrast, our sticky information and noisy information models predict that there should be no such correlation. Third, equation (30) demonstrates that the cross-sectional dispersion of forecasts is time-varying and correlated with shocks to inflation, while the baseline model predicts no such variation or correlation. Furthermore, forecast dispersion should increase at the time of a shock to inflation since inflationary shocks enter (30) as squares and multiplied by non-zero constants.

2.3.4 *Asymmetric loss function*

Elliot, Komunjer, and Timmermann (2008) and Capistran and Timmermann (2009) show that even full-information agents can make biased forecasts if these agents have asymmetric loss functions over forecast errors. To the extent that agents have different asymmetries in the loss function, disagreement across agents can arise without resorting to information rigidities. Following Capistran and Timmermann (2009), suppose that agent i has a loss function over forecast errors $FE_{t,t-1} \equiv \pi_t - \pi_{t|t-1}$ which is of the LINEX form:

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

This loss function has the property that as ϕ goes to zero, the loss function converges to the standard mean squared error objective. When $\phi > 0$, agents dislike positive forecast errors more than negative forecast errors and vice-versa when $\phi < 0$.

Consider a general case where conditional on information at time t inflation is normally distributed $\pi_{t+1|t} \sim N(\mu_{t+1|t}, \sigma_{t+1|t}^2)$ with expected mean of $\mu_{t+1|t}$ and conditional variance $\sigma_{t+1|t}^2$. For example, in the context of our baseline noisy information model, $\mu_{t+1|t} = \rho\pi_t$ and $\sigma_{t+1|t}^2 = \Sigma_w$. Then the mean of the optimal forecast of inflation one period ahead conditional on time t information is given by $\overline{\pi_{t+1|t}} = \mu_{t+1|t} + \frac{1}{2}\sigma_{t+1|t}^2 E(\phi_i)$ where $E(\phi_i)$ is the average loss asymmetry across agents. Thus, aggregate forecasts can differ from the conditional mean if, on average, agents have asymmetric losses over forecast errors.

In addition, if the conditional mean is time-varying, forecast errors will have interesting dynamic properties. Again following Capistran and Timmermann (2009), suppose the conditional variance of inflation follows a standard GARCH(1,1) process such that $\sigma_{t+1|t}^2 = \alpha_0 + \alpha_1 w_t^2 + \beta_1 \sigma_{t|t-1}^2$ where w_t is the innovation to inflation (i.e., $\pi_{t+1} = \mu_{t+1|t} + w_{t+1}$), $\alpha_1 > 0$ and $\beta_1 \in (0,1)$. Then the average one-period ahead inflation forecast is given by

$$\overline{\pi_{t+1|t}} = \mu_{t+1|t} + \frac{1}{2}(1 - \beta_1)^{-1} \alpha_0 E(\phi_i) + \frac{1}{2} \alpha_1 E(\phi_i) \sum_{k=0}^{\infty} \beta_1^k w_{t-k}^2 \quad (31)$$

where σ_{π}^2 is the average variance of inflation. This implies that the impulse response of the forecast error $FE_{t,t+1} \equiv \pi_{t+1} - \overline{\pi_{t+1|t}}$ to a one-time inflation innovation δ at time zero is given by $-\frac{1}{2} \alpha_1 E(\phi_i) \beta_1^t \delta^2$. Note that the sign of the response of the forecast error is ambiguous, since $E(\phi_i)$ can be either positive or negative. Yet, the sign of the response of the forecast error is independent of whether the shock is inflationary or deflationary which contrasts with the predictions of the sticky and noisy information approaches.

Finally, the standard deviation of forecasts across agents is equal to $\frac{1}{4} \text{var}(\phi_i) \sigma_{t+1|t}^4$. To the extent that the conditional variance of inflation is time-varying, then disagreement across agents will also vary across time proportionally to the degree to inflation uncertainty. Assuming the same GARCH process for the inflation process as before, the impulse response of disagreement to a one-time innovation to inflation at time zero is $\frac{1}{4} \text{var}(\phi_i) \alpha_1^2 \beta_1^{2t} w_0^4$. As with sticky information, dispersion should rise after any innovation to inflation, be it positive or negative, since it is the squared innovation which affects dispersion. In addition, given the assumed GARCH process, the response of dispersion should be monotonically declining over time.

2.4 Taking stock

Table 1 presents a summary of predictions from various models we have considered above. All of the models with information rigidities share a common prediction: in response to an inflationary (disinflationary) shock, average forecast errors should be positive (negative) as agents fail to incorporate all of the relevant information into their forecasts. But asymptotically, forecast errors should go to zero as all of the information is acquired and processed. On the other hand, there are three key dimensions along which the models make differential predictions: i) whether mean forecast errors correlate with past inflation after conditioning on the mean forecast error in the previous period; ii) how quickly the response of mean forecast errors normalized by the response of the actual forecasted series to a given shock converges to zero and how this speed varies across shocks; iii) how disagreement of forecasts across agents responds to shocks. For example, while the baseline noisy information model predicts that the

speed of the response of forecast errors may differ across shocks and that disagreement across forecasters should not respond to shocks, the sticky information model predicts the same convergence rate of forecast errors and that disagreement responds to shocks. The key element of these tests that will allow us to differentiate between these models, all of which are consistent with the well-know presence of serially correlated forecast errors and unconditional disagreement across agents in the data, is the focus on conditional responses of forecast data to shocks.

3 Data, Methodology and Benchmark Results for Professional Forecasters

The theoretical predictions derived in section 2 are for the conditional responses of forecast errors and disagreement among agents to economic shocks. To assess the empirical validity of these models, we will consistently follow a two-step approach. In the first step, economic shocks are identified in a variety of ways suggested by the literature. In the second step, we generate responses of the relevant moments of agents' expectations to these shocks. In this section, we first discuss the shocks used in our analysis, then apply our approach to data for professional forecasters as a benchmark before turning to the expectations of other agents in subsequent sections.

3.1 Shocks

Because the predictions made by the models are all conditional on a macroeconomic shock, a key element of our analysis is the selection and identification of shocks. There is a long literature on identifying exogenous structural shocks to the economy, giving us a wide range of measures to consider. However, we document in Appendix B that to obtain informative estimates in small samples, the shocks must account for a sufficiently large fraction of the historical volatility of the variable being forecasted. As a result, we focus on the following three shocks from the previous literature because we have found these to be most important in accounting for the inflation volatility over our time sample:

- a) Technology shocks, identified using long-run restrictions as in Gali (1999).⁶
- b) Oil shocks, identified as in Hamilton (1996).⁷
- c) News shocks, identified as in Barsky and Sims (2011).⁸

⁶ To estimate technology shocks, we estimate a trivariate VAR(4) on quarterly data for the change in labor productivity, change in hours, and inflation rate of the GDP deflator. Labor productivity and hours are defined as in Gali (1999). The estimation sample covers 1952Q2 through 2007Q3. Technology shocks are identified from the restriction that only technology shocks have long run effect on productivity.

⁷ Hamilton (1996) identifies (WTI) oil price shocks as episodes when the oil price exceeds the maximum oil price over the last twelve months. When this is the case, the shock is the difference between the current price and the maximum over the last twelve months, and zero otherwise. We take logs of all prices. Data is quarterly from 1950 until the end of 2007.

Table 2 presents a variance decomposition of inflation volatility in terms of these three shocks from a VAR(4) over the time period of 1966-2007. The results, which are largely invariant to the ordering of the VAR, imply that technology shocks have played a key role in driving inflation volatility at all horizons, accounting for approximately 25% of the variance of inflation over this time period. Oil price shocks and news shocks each account for approximately 10% of inflation volatility at longer horizons.⁹

Also note that much of the inflation variance remains unexplained as a function of these shock measures. At the same time, the predictions derived in section 2 do not depend on specifics of a shock and thus any shock may be used to construct conditional responses. As a result, we will also consider “unidentified” innovations to inflation in our empirical tests of the models. We generate these innovations via the residuals v_t from the quarterly regression

$$\pi_t = c + \sum_{k=1}^4 \beta_k \pi_{t-k} + \sum_{s \in \{T, N, O\}} \sum_{j=0}^1 \mathcal{V}_j^s \hat{\varepsilon}_{t-j}^s + v_t \quad (32)$$

where $\hat{\varepsilon}^T, \hat{\varepsilon}^N, \hat{\varepsilon}^O$ are identified technology, news and oil price shocks respectively. While the unidentified inflation innovations represent the combined effects of different structural shocks, they account for a much larger component of the inflation volatility than other shocks and as such provide a useful complementary source of inflation variation for our analysis.

3.2 Inflation Forecasts from Professional Forecasters

As a benchmark for subsequent analysis, we first focus on inflation forecasts from the Survey of Professional Forecasters (SPF) for several reasons. Because professional forecasters are some of the most informed economic agents, one would expect any evidence of information rigidity on their part to be particularly notable. Second, predictions of professional forecasters are consistently available at a quarterly frequency. Third, professional forecasters make predictions of explicitly defined variables, such as the CPI or the GDP deflator (unlike consumer forecasts), so there is a well-defined relationship between their forecasts and ex-post values. Fourth, data on professional forecasters has been used in support of many of the theories presented in section 2. For example, Mankiw, Reis and Wolfers (2004) apply sticky information to professional forecaster data, Patton and Timmermann (2010) apply their model with heterogeneous priors about long-run means to professional forecaster data, and Capistran and Timmermann (2009) provide evidence of heterogeneous loss-aversion in professional forecasts. Finally,

⁸ The news shock is identified as the shock orthogonal to the innovation in current utilization-adjusted TFP that best explains variation in future TFP (the horizon is 10 years). Four variables are included in the benchmark system when the shock is identified: the corrected TFP series, non-durables and services consumption, real output, and hours worked per capita. We thank Eric Sims for sharing the shock series with us.

⁹ We also considered monetary policy shocks from a VAR as in Christiano et al. (1999), fiscal shocks from Romer and Romer (2011), uncertainty shocks from Bloom (2009) and confidence shocks from Barsky and Sims (2010). However, these shocks consistently accounted for less than 5% of inflation volatility making them unreliable shocks to use in the two-step procedure, as demonstrated in Appendix B.

we focus primarily on inflation forecasts because of the prominent role they play in macroeconomics, such as in price-setting decisions and the Phillips Curve, monetary policy decisions, ex-ante real interest rates used in consumption and investment decisions, etc.

The specific dataset that we use, the Survey of Professional Forecasters, is a quarterly survey of 9 to 40 professional forecasters. Forecasts are collected by the Philadelphia Federal Reserve in the middle of each quarter for a variety of macroeconomic variables at different forecasting horizons. We focus on forecasts of the GDP deflator over the next four quarters (GNP deflator prior to 1992). Panel A in Figure 1 presents time series of actual inflation, inflation forecasts, and the standard deviation for cross-section of forecasts in any given time period. Although forecasts track the actual inflation rate closely, the difference between actual and forecast inflation is fairly persistent. Forecast disagreement was high in the late 1970s and early 1980s but declined afterwards. There is also no obvious relationship between disagreement and the recessions dated by the National Bureau of Economic Research (NBER).

The persistence in forecast errors and, even more so, disagreement among professional forecasters has been emphasized by many as *prima facie* evidence against the null of full information rational expectations, with many of the models in section 2 considered as potential explanations. Mankiw, Reis and Wolfers (2004), for example, argue that a sticky information model can account for the level of disagreement among forecasters as well as the response of disagreement to the Volcker disinflation, while Andrade and Le Bihan (2010) show that European professional forecasters regularly do not change their forecasts. Disagreement among professional forecasters about the 10-year ahead inflation rate and the natural level of unemployment rate has been suggested as evidence for the model of Patton and Timmermann (2010) which emphasizes the potential importance of such heterogeneity about long-run means. Similarly, professional forecasters and policymakers frequently emphasize the importance of private information in formulating their forecasts. The Beige Book prepared before each FOMC meeting, for example, summarizes the anecdotal information collected by regional Federal Reserve Presidents from meeting with their business contacts in industry. Individual policymakers have also emphasized that such contacts affect their views about economic conditions.¹⁰ Berger, Ehrmann, and Fratzscher (2011) document that the geographical location of forecasters affects their ability to predict monetary policy decisions, consistent with the notion that forecasters place some weight on their contacts in industry, which are likely to come disproportionately from their geographic area. Thus, while there is

¹⁰ For example, Cleveland Fed President Pianalto stated in her October 1st, 2009 speech “For sure, this is a difficult time to be in the business of economic forecasting. To paraphrase one of my colleagues, we are looking at flawed data through the lens of imperfect models. To try to clarify my perspective on the economy, I also spend a lot of time talking with businesspeople—the heads of Fortune 500 companies, owners of small and medium-sized enterprises, and CEOs from large regional banks and small community banks.” See also Koenig’s November 2004 speech to the OECD.

likely some element of truth to each suggested theory of disagreement among forecasters, a key advantage of our approach is to determine if one source of heterogeneity is most important and can account for the conditional responses of both mean forecasts and disagreements to shocks.

3.3 Are Professional Forecasters Subject to Information Rigidities?

The first key prediction from models of information rigidities is that the conditional response of forecast errors to a shock should be of the same sign as the response of the variable being forecasted to the shock whereas the null of full information rational expectations is of an immediate and complete adjustment of forecasts to shocks and therefore zero forecast errors after any shock. We first present impulse responses of annual inflation $\pi_{t+4,t}$ from period t to period $t + 4$ to technology, news, oil price and unidentified shocks. To construct these impulse responses, we follow Romer and Romer (2004) and estimate

$$\pi_{t,t-4} = c + \sum_{k=1}^K \beta_k \pi_{t-k,t-4-k} + \sum_{j=0}^J \gamma_j \hat{\varepsilon}_{t-j}^s + \mu_t \quad (33)$$

where $s \in \{T, N, O, U\}$ denotes the type of shock, i.e. technology (T), news (N), oil prices (O), or unidentified (U). Inflation is measured using the GDP (GNP prior to 1992) deflator at the quarterly frequency. The lag lengths K and J , up to 2 years each, are selected via the Bayesian information criterion (BIC) but results are insensitive to using alternative information criteria for lag length selection. The time sample is 1976-2007, which reflects the starting date of SPF forecasts of year-ahead inflation after allowing for lags. We focus on the year-on-year inflation rate because this conforms to the forecast used for professional forecasters as well as other agents. Because professional forecasters at time t will be forecasting the inflation rate from time t to $t+4$, we drop the first four observations of the impulse response of annual inflation. As a result, the impulse responses correspond exactly to what forecasters are trying to predict.

The results for all four shocks are presented in the left column of Figure 2. The mean response of inflation to technology and news shocks is negative, whereas oil price and unidentified shocks are inflationary, by construction for the latter. The response of inflation to each shock converges gradually – and nearly monotonically – back to zero. Whereas much of the literature using impulse response analysis resorts to one-standard deviation confidence intervals, we present both one and two standard deviation confidence intervals.

Under the null of full-information rational expectations, forecasts should adjust to shocks by the same amount as future inflation, hence the response of forecast errors to these shocks should be zero. Models with information rigidities instead predict a non-zero response of mean forecast errors across agents to economic shocks. To assess these conflicting predictions, we estimate the following regression for each shock

$$\pi_{t,t-4} - \overline{\pi_{t,t-4|t-4}} = c + \sum_{k=1}^K \beta_k (\pi_{t-k,t-4-k} - \overline{\pi_{t-k,t-4-k|t-4-k}}) + \sum_{j=0}^J \gamma_j \hat{\varepsilon}_{t-j}^S + \mu_t \quad (34)$$

over the same sample as the inflation regression, again selecting K and J via the BIC, but now using the average forecast error across agents ($\pi_{t+4,t} - \overline{\pi_{t+4,t|t}}$) as a regressand. Figure 2 plots the implied impulse responses of mean forecast errors to the four shocks, again dropping the first four periods (since in the first four periods forecasters have not had the opportunity to observe the shock). For each shock, we can reject the null of no response of forecast errors to shocks at standard levels of statistical significance. Note that while we use generated regressors for shocks in our empirical specification, Pagan (1984) shows that under the null hypothesis that the coefficient on a generated regressor is zero, standard errors do not need to be adjusted for generated regressors. Since under the null of full-information rational expectations $\gamma_j = 0 \forall j$, our standard errors are valid. In addition, we show in section 5 that explicitly adjusting standard errors for the presence of generated regressors has negligible effects in this setting because the shocks are the residuals from the first stage rather than the fitted values. Thus, *one can strongly reject the null of full-information and this rejection goes exactly in the manner predicted by models of information rigidities: forecast errors are negative after disinflationary technology and news shocks but positive after inflationary oil price and unidentified shocks*. Also as predicted by models of information rigidities, forecast errors converge back to zero over time, as agents' information sets progressively incorporate and process the new information.

As discussed in section 2, serially correlated forecast errors could be observed even in the absence of information constraints. Capistran and Timmermann (2009) propose a model in which heterogeneous loss aversion across agents can lead to disagreement and serial correlation of forecast errors. However, we showed in section 2 that this model predicts that forecast errors should always be either positive or negative after any shock, regardless of whether the shock is inflationary or disinflationary. To assess this alternative potential explanation for our findings, we regress forecast errors on lags of themselves as well as contemporaneous and lagged absolute values of shocks:

$$\pi_{t,t-4} - \overline{\pi_{t,t-4|t-4}} = c + \sum_{k=1}^K \beta_k (\pi_{t-k,t-4-k} - \overline{\pi_{t-k,t-4-k|t-4-k}}) + \sum_{j=0}^J \gamma_j |\hat{\varepsilon}_{t-j}^S| + \mu_t. \quad (35)$$

If heterogeneous loss aversion, rather than information rigidity, is an important component of the forecasting decisions of professional forecasters, the conditional response of forecast errors to the absolute value of shocks should consistently be of the same sign across shocks. The third column of Figure 2 presents the implied impulse response of forecast errors from estimating equation (35): we find no evidence of a consistently positive or negative response to the absolute value of the shocks. The response is positive after technology shocks, but negative (and not statistically different from zero) for news and unidentified shocks. Since Hamilton's (1996) oil price shocks are all positive by construction, these shocks do not provide conflicting or contradictory evidence. Thus, while the response of forecast

errors to shocks all go precisely in the direction predicted by models of information rigidities, the response of forecast errors to the absolute value of shocks does not provide any evidence for the alternative hypothesis of heterogeneous loss-aversion being a primary determinant of the forecasting decisions of professional forecasters. This suggests that the clear pattern of conditionally correlated forecast errors in Figure 2 reflects information rigidities faced by professional forecasters.

3.4 Distinguishing Between Information Rigidities Faced by Professional Forecasters

To distinguish between models of information rigidities, we can first consider whether the response of inflation forecast errors to shocks is sensitive to past levels of inflation. Recall that both the sticky information and baseline noisy information models predict that the response of forecast errors to shocks is independent of past conditions, whereas the noisy information models with either heterogeneous priors about long-run means or heterogeneity in signal strength imply that the response of inflation forecast errors should be correlated with lagged levels of inflation. To assess these predictions, we consider the following regression:

$$\pi_{t+4,t} - \overline{\pi_{t+4,t}|t} = c + \beta(\pi_{t+3,t-1} - \overline{\pi_{t+3,t-1}|t-1}) + \gamma\pi_{t-1,t-5} + \mu_t. \quad (36)$$

In this specification, all of the structural shocks at time t are incorporated into the error term such that, because they are orthogonal with respect to information dated $t - 1$ and earlier, we can estimate this specification by OLS to assess whether $\gamma > 0$ as suggested by the noisy information models with heterogeneity in long-run means or signal strength. Using quarterly data from 1976 to 2007, we find

$$\begin{aligned} \pi_{t+4,t} - \overline{\pi_{t+4,t}|t} = & 0.05 + 0.88(\pi_{t+3,t-1} - \overline{\pi_{t+3,t-1}|t-1}) - 0.02\pi_{t-1,t-5} + \mu_t & (37) \\ & (0.09) \quad (0.04) & (0.03) \\ R^2 = & 0.77 \quad S.E. = 0.46 \quad p - val(\text{Box} - \text{Ljung} \text{ Q} - \text{stat}) = 0.93 \end{aligned}$$

where Newey-West HAC standard errors are in parentheses. The coefficient on lagged inflation is small and not statistically significantly different from zero. This is consistent with both sticky information and the baseline noisy information model but inconsistent with heterogeneous priors about long-run means or signal strength.

A second dimension along which we can examine the empirical evidence for different models of information rigidity is via the rate at which forecast errors converge to zero. In section 2, we showed that the convergence of forecast errors after a shock was a function of both the underlying degree of information rigidity and the persistence of the shock. However, because the convergence rate of the forecasted variable depends only on the persistence of the shock, one can recover an estimate of the degree of information rigidity by normalizing the impulse response of forecast errors by the impulse response of inflation. Thus, we construct what we call normalized impulse responses by taking the ratio

of the estimated impulse response of forecast errors from estimated equation (34) to the estimated impulse response of inflation from estimated equation (33). We then fit an AR(1) process to the normalized impulse response to assess its convergence rate, which corresponds to a direct estimate of the underlying degree of information rigidity from each model. This procedure yields estimates between 0.86 and 0.89 for technology, news and oil price shocks as well as for unidentified shocks. These estimates point to economically significant information rigidities. In the context of sticky information models, an estimate of 0.86 would correspond to forecasters updating their information sets every six to seven quarters. While high, this is in line with the estimates of the degree of information rigidity over the same time period using a sticky information Phillips Curve and data from professional forecasters in Coibion (2010). In the context of noisy information models, this implies a weight of 0.14 placed on new information relative to the previous forecast, which is close to the estimated value of 0.10 in Bordo et al. (2007) based on the behavior of professional forecasters during the Volcker disinflation. Furthermore, the p -value for the null hypothesis that the convergence rates are equal across shocks is 0.98. This is consistent with the prediction of sticky-information models. While noisy information models predict that convergence rates generally differ across shocks, our inability to reject the null of equality need not be interpreted as a rejection of these models but it does indicate that noisy information models which point to important differences in information acquisition rates across our shocks may be difficult to reconcile with the data.

A third dimension along which models of information rigidities make conflicting predictions is the predicted response of disagreement among forecasters to shocks. Under both sticky information and heterogeneous signal-noise ratios, disagreement should rise after any shock whereas the baseline noisy information model and the versions with strategic interaction or heterogeneous priors about long-run means predict instead that disagreement should be largely invariant to economic shocks. To assess whether disagreement responds to shocks, we estimate the following regression:

$$\sigma(\pi_{t+4,t|t}(i)) = c + \sum_{k=1}^K \beta_k \sigma(\pi_{t+4-k,t-k|t-k}(i)) + \sum_{j=0}^J \gamma_j |\hat{\varepsilon}_{t-j}^S| + \mu_t, \quad (38)$$

where $\sigma(\pi_{t+4,t|t}(i))$ is the cross-sectional standard deviation of time- t forecasts of year-ahead annual inflation from professional forecasters. We use the absolute value of shocks because both sticky information and the noisy information model with heterogeneous signal strength suggest that disagreement should be increasing after any shock, regardless of whether it is inflationary or disinflationary. We use the same time period of 1976-2007 and select K and J using the BIC. The results, presented in the fourth column of Figure 2, indicate no discernible evidence that disagreement responds in a statistically significant manner to these shocks. Thus, we cannot reject the null that disagreement is insensitive to shocks as suggested by the baseline noisy information model.

In addition, we include in these Figures the predicted response of disagreement to these shocks in the sticky information model. To do so, we note that under sticky-information, the average impulse response of the cross-sectional standard deviation of forecasts across agents should follow

$$\left(\frac{d\pi_t}{d\delta}\right)^2 \lambda^{j+1}(1 - \lambda^{j+1})\delta^2$$

where $\frac{d\pi_t}{d\delta}$ is the impulse response of inflation to a δ -sized innovation and λ is the degree of sticky information. Using the estimated impulse responses of inflation from equation (33) and the estimated degrees of information rigidity from the convergence rate of normalized forecast errors, we plot in Figure 2 the predicted response of disagreement to a one-unit shock under sticky-information. In each case, the predicted path of disagreement under sticky information is well above the confidence interval for the actual response of disagreement to these shocks. This indicates that not only do we fail to reject the null of no response of disagreement, but we also find that the predicted responses under sticky information consistently lie well outside the confidence intervals of the actual responses of disagreement to these shocks.

3.5 Professional Forecasters and Financial Market Forecasts

The results for professional forecasters strongly support the notion that they face information rigidities closely conforming to the predictions of models of noisy information: forecast errors respond in the same direction as the inflation rate after a shock, forecast errors are not predictable using lagged inflation conditional on lagged forecast errors, and disagreement among professional forecasters does not respond to the absolute value of shocks. In addition, the convergence rates of normalized forecast errors point to high levels of information rigidity. One potential caveat to the latter, however, comes from the possibility of strategic interaction as in Morris and Shin (2002). In their model, the effects of noisy information on individual forecasts can be compounded by the desire of individuals to either follow or deviate from the average forecast made by other agents. In the case of $R > 0$, i.e. there is strategic complementarity in forecasting, this can lead to forecasts adjusting more gradually than would be implied only by the underlying noise in public and private signals. In other words, the convergence rates of normalized forecast errors would combine the effects of information rigidity with the strategic interaction in forecasts, leading to an overestimate of the quantitative importance of noisy information.

To assess whether strategic interaction can account for the high estimated levels of information rigidity, we make use of two sets of inflation forecasts extracted from asset prices. The first set, which we take from the Cleveland Federal Reserve Bank, is constructed in Haubrich, Pennacchi, and Ritchken (2011) and is available starting in 1982 at the monthly frequency. These forecasts make use of the term structure of interest rates, inflation swaps and forecasts from the Blue Chip Economic Indicators to

control for time-varying risk premia in extracting the market's expectations of future inflation. The second set comprises of the one-year ahead inflation expectations constructed in Ang, Bekaert, and Wei (2008) who also exploit the term structure of interest rates to measure the market's expectation of inflation while allowing for regime switches and time-varying bond premia.¹¹ Unlike Haubrich, Pennacchi, and Ritchken (2008), their approach does not rely on survey forecasts at all. This measure is available over the same period as SPF forecasts.

Forecasts extracted from financial market forecasts provide a useful comparison because agents trading in these markets have "skin in the game" and as a result are unlikely to be subject to strategic interaction. In Appendix C, we replicate the impulse response analysis of forecast errors to shocks and the absolute values of shocks using both alternative sets of inflation forecasts and find results that conform closely to those obtained using professional forecasters' predictions, namely that forecast errors respond systematically to shocks in the same direction as inflation whereas the response of forecast errors to the absolute value of shocks does not reveal a consistent pattern of either positive or negative responses. This suggests that, like professional forecasters, the forecasts implicit in asset prices also embed information rigidities. In addition, we can estimate the persistence of normalized forecast errors to recover a measure of the underlying degree of information rigidity: these values (Table 4) range from 0.82 to 0.92, very similar to what obtained with professional forecasters. Because financial market participants are unlikely to be trading on expectations embodying much strategic interaction (if they were, there would be considerable arbitrage opportunities), this strongly suggests that the source of the high degree of information rigidity in professional forecasts is noisy information with strategic interaction contributing very little.

This finding is in line with the results of Ang, Bekaert, and Wei (2007) who document that forecasts of inflation from professional forecasters outperform financial market forecasts as well as those of variety of time series models. Coibion and Gorodnichenko (2010) also recover quantitatively indistinguishable estimates of information rigidity from professional forecasters and financial market participants using an alternative estimation strategy. Were strategic interaction a significant factor in the reporting of professional forecasts, one would expect their forecasts to be inferior to those extracted from asset prices. Instead, the data point to professional forecasts being at least as good as financial market forecasts with equivalent estimates of information rigidity, thereby indicating that strategic interaction in forecasts is unlikely to play a quantitatively important role in the formation of forecasts and expectations of these agents.

¹¹ We are grateful to Min Wei for sharing the inflation expectations from Ang, Bekaert, and Wei (2008) with us.

4 Consumers, Firms and Central Bankers

In this section, we apply our empirical tests of the expectations formation process to forecast data from consumers, firms, and central bankers. While professional forecasters provide a useful benchmark for assessing the potential importance of information rigidities, the economic significance of their forecasts for macroeconomic dynamics is ambiguous. Most macroeconomic models, for example, do not incorporate any specific role for these agents. Instead, it is the expectations of consumers, firms and central bankers which are at the center stage of most economic analyses, as illustrated for example by the New Keynesian models of Woodford (2003), in which these are the only agents in the analysis. As a result, this section considers the data used to analyze the nature of each agent type's expectations formation process then turns to applying the same empirical tests to these datasets as used for professional forecasters.

4.1 Data for Consumers, Firms and Central Bankers' Forecasts

For consumers, we rely on the Michigan Survey of Consumers (MSC). The MSC is a nationally representative survey of 500 to 1,300 consumers done quarterly since 1968 and monthly since 1978. Respondents are asked to report their expected inflation rate for the next twelve months. While most of the questions in the survey ask only for qualitative responses, the question about consumers' price expectations over the next twelve months asks for a numerical value. From consumers' answers, we can construct a measure of the average forecast of inflation over the next twelve months, analogous to the mean forecast in the SPF, as well as the cross-sectional standard deviation of forecasts to measure disagreement.

However, this data has several limitations relative to SPF forecasts. First, the question posed to consumers does not specify a price index, so it is unclear which price index is most appropriate to use to construct forecast errors. We use the annual change in the Consumer Price Index (CPI), but have verified that our results are robust to using the Personal Consumption Expenditure (PCE) index. More broadly, the absence of a specified price index in the survey question means that consumers could be making conceptually different forecasts. Some may indeed be forecasting an aggregate price level such as the CPI, while others may be forecasting price changes of their own consumption bundles. The latter could give rise to substantial heterogeneity in forecasts even in the absence of information rigidities since consumers with different ages, income levels, or preferences may consume very different bundles. A third limitation of this dataset is that consumers are answering a phone survey and therefore the survey responses could be somewhat contaminated with measurement error and the amount of heterogeneity in forecasts across consumers may be overstated.

For firms, we rely on the Livingston Survey, first established in 1946 by the columnist Joseph Livingston and managed by the Federal Reserve Bank of Philadelphia since 1990. This biannual survey collects forecasts from individuals in a variety of institutions such as academia, government organizations,

industry and banking. In December and June of each year, individuals provide their forecasts for a number of economic variables, including the CPI and the unemployment rate for future periods such as 6 or 12 months ahead. We use only the forecasts of individuals in commercial banking, consulting and business to represent the concept of firms' expectations, which yields an average of 27 forecasters per survey. From these individuals' forecasts, we can construct a measure of the mean forecast of inflation over the next 12 months (and the corresponding real-time forecast errors) analogous to that of the SPF, as well as the cross-sectional standard deviation of forecasts to measure disagreement. The only notable limitation of this dataset relative to the SPF is its more limited frequency: semi-annual rather than quarterly.

In addition, we consider forecasts from members of the Federal Open Market Committee (FOMC) of the Federal Reserve. These forecasts are formed as a component of the Federal Reserve's semiannual *Monetary Policy Reports* to Congress, submitted each February and July since 1979. As detailed in Romer and Romer (2008), forecasts are consistently available for nominal GNP/GDP growth, real GNP/GDP growth, a measure of inflation, and the unemployment rate, with additional variables released in recent years. July forecasts are for both the current and next year's values, whereas February forecasts prior to 2005 are for the current year only (starting in 2005, they report forecasts of current and next year values). While each FOMC member was required to submit a forecast, the Monetary Policy Reports only provide summary statistics for each variable. In particular, they report "central tendency" values, which shows the highest and lowest forecasts after dropping the extremes (commonly defined as the three highest and three lowest values, although this is not consistently made clear in the reports), and the "range" of forecasts listing the highest and lowest values. We construct a measure of mean forecast as the midpoint of the "central tendency" values. Because the underlying individual forecasts are not available for much of the sample, we approximate the cross-sectional standard deviation of forecasts by assuming that the underlying forecasts are normally distributed each period. Specifically, we treat the "range" as a measure of the 95 percent interval and the "central tendency" as a measure of the 68 percent interval of forecasts, and then use the average over the two implied standard deviations from each. The specific inflation measure being forecasted has changed over time: GNP price deflator until July 1988, CPI inflation from February 1989 until July 1999, the PCE index from February 2002 to February 2004, and the core PCE from July 2004 onwards. Because forecasts are made for calendar year values, forecasts from February and July meetings do not have identical time horizons. For February meetings, we use the forecast of the current year values, while for the July meeting, we use a weighted average of current and next year values (weights of 7/12 and 5/12 on current and subsequent year respectively). Ex-post values for the construction of forecast errors are defined in the same way.

Like the Livingston survey, one limitation of the FOMC forecasts is their semiannual frequency. In addition, the measure of disagreement available is not constructed directly from the underlying distribution of individual forecasts as was the case with professional forecasters, firms, and consumers. A third possible limitation is the extent to which FOMC members devote attention to making the forecasts themselves, since they are aware that only summary statistics will be released in the *Monetary Policy Report*. Romer and Romer (2008), for example, document that the central tendency of FOMC inflation forecasts has no additional predictive power over the Greenbook forecasts prepared by the staff of the Board of Governors of Federal Reserve, despite the fact that FOMC members are able to revise their forecasts after the Greenbooks are made available. This concern is likely to be mitigated by internal reputational considerations however, as well as the possibility of the individual forecasts being released to the public (as they now are with a 10-year lag, see Romer 2009). A fourth potential concern with this data is that FOMC members are asked to provide forecasts under what they view to be “appropriate” monetary policy, which may differ from what they perceive to be the most likely path of monetary policy. However, because monetary policy actions have only gradual effects on prices, differences in assumptions about the future path of monetary policy are unlikely to have significant effects on FOMC members’ forecasts of inflation over the course of the next few quarters.

Figure 1 presents time series for inflation, mean forecasts, and forecast disagreement for consumers, firms and FOMC members. Similar to the professional forecasters, mean forecasts track the actual inflation rate well but the differences between the two are persistent. In a similar vein, forecast disagreement for each type of agents exhibits a secular trend, with disagreement peaking in the early 1980s. The level of disagreement is highest for consumers while disagreement is similar for professional forecasters, firms, and FOMC members. Some of the cross-sectional heterogeneity in consumer forecasts is likely to reflect the different nature of this survey: namely that consumers have little time to think about their responses to a phone survey and therefore report values which may include significant noise as well as the absence of guidance in the question about what measure of inflation they are meant to forecast.

4.2 Empirical Results for Consumers, Firms and Central Bankers’ Forecasts

Using these forecasts from consumers, firms, and central bankers, we can apply the same empirical tests as used with the SPF to ascertain the nature of their expectations formation process. We focus on the period since 1976, to the extent that data is available, to ensure as common a sample across forecast types as possible while maintaining the late 1970s and early 1980s in the sample. We continue to allow for two years’ worth of lags in each empirical specification, as in section 3.3. As a first step, we estimate equation (34) for each type of forecast to study the response of forecast errors to technology, news, oil price and unidentified shocks. Figure 3 shows that in each case there is significant evidence of

information rigidities. After disinflationary technology and news shocks, the response of forecast errors are consistently negative and converge to zero over time as predicted by models of information rigidities. We can always reject the null of zero response of forecast errors using traditional one-standard deviation confidence intervals and can reject the null in five out of six cases using two-standard deviation confidence intervals. After inflationary oil price and unidentified inflation innovations, the responses of forecast errors are always positive and converging to zero as time passes, and we can reject the null of no response of forecast errors in every case using either one or two standard deviation confidence intervals. Hence, the results are again strongly supportive of the presence of information rigidities for these agents: *forecast errors consistently respond in the same direction as inflation to shocks.*

To assess whether these results are driven by heterogeneity in loss aversion rather than information rigidities, we estimate equation (35) using the absolute value of shocks rather than their levels. The results, presented in Figure 4, indicate that there is little evidence of a consistent response of forecast errors to the absolute values of the shocks. The null of no response can only be rejected in one case at the 5% level. This indicates that heterogeneity in loss aversion is again unlikely to be the primary source of serial correlation in forecast errors observed in the data, as was found with professional forecasters. Instead, the strong response of forecast errors to shocks but not to the absolute value of the shocks conforms closely to the prediction of models with information rigidities. Thus, the evidence from these empirical tests indicate that information rigidities are likely to be an important component of the expectations formation process for consumers, firms, and central bankers as well as professional forecasters.

In addition, we can assess which models of information rigidity might best characterize the nature of the expectations formation process for these different types of economic agents. First, we examine whether inflation forecast errors are systematically and positively correlated with lagged levels of inflation as suggested by the models of noisy information augmented with heterogeneity in single strength or priors about long-run means. As in section 3.4, we do so by regressing forecast errors on lags of themselves and lags of inflation using equation (36). The results, including those for professional forecasters and financial market forecasts, are displayed in Table 3. Using one lag of forecast errors and one lag of inflation, the results are inconsistent with the predictions of these models. The coefficients on lagged inflation are small and either not significantly different from zero or of the wrong sign. The latter, which occurs for FOMC members and to a lesser extent firms, is likely to be driven by time aggregation: when we replicate the estimate of equation (36) for professional forecasters at the semiannual frequency rather than the quarterly frequency, these also display a negative and statistically significant coefficient on lagged inflation, unlike the results at the quarterly frequency. Similar findings obtain with more general lag specifications of forecast errors and inflation. *Thus, these results suggest that heterogeneity about*

long-run means and heterogeneity in Kalman gains are unlikely to be an important component of the expectations formation process for consumers, firms, central bankers as well as financial market participants, consistent with that observed for professional forecasters.

As another way to distinguish between models of information rigidities, we also present impulse responses of disagreement among consumers, firms, and central bankers to the absolute value of shocks in Figure 5. For firms, there is little evidence that disagreement consistently rises after these shocks and the predicted responses are, in most cases, much lower than those predicted by the sticky information model. Very similar results are obtained for FOMC members, indicating that sticky information is unlikely to be the primary source of information rigidity for firms and central bankers. In the case of consumers, we uncover one case in which we can reject the null of no response of disagreement: after oil price shocks, disagreement among consumers' inflation forecasts rises and remains persistently positive even four or five years after an oil price shock. While the initial increase in disagreement is consistent with the sticky-information model, its persistence is much higher than predicted by the latter. The response of disagreement to unidentified shocks and technology shocks, on the other hand, are not statistically different from zero but well below the predicted response under sticky information. As a result, this evidence cannot readily be interpreted as supporting sticky information as the primary form of information rigidity underlying consumer forecasts.

The increase in disagreement after oil price shocks, nonetheless, is clearly at odds with the baseline prediction of noisy information models. One interpretation is that this result is purely statistical: given that we estimate 16 responses of disagreement to shocks, it should not be surprising to reject the null of no response for one of them even if the null hypothesis of no response is true. Alternatively, one could entertain “structural” interpretations of this result. For example, some consumers may forecast the price of their individual consumption bundles rather than an aggregate price index. To the extent that agents consume different quantities of gasoline and other energy products, oil price shocks could then lead to persistent disagreement among consumers about future inflation as a result of their different exposure to oil-related products. This could also reflect an element of rational inattention: different energy consumption across agents should lead to different incentives for agents to devote information processing abilities to tracking oil prices and therefore generate heterogeneity in the amount of noise in private signals. As demonstrated in section 2.3.3, this could reveal itself in disagreement responding to economic shocks. Another possibility is that the rise in disagreement could indeed reflect a sticky information component not present with respect to other shocks.

Finally, Table 4 presents the convergence rates of forecast errors (normalized by the responses of inflation to shocks) for each type of shock and agent. The first point to note is that all estimates are positive

and statistically significant at conventional levels. This reinforces the result that information rigidities are clearly present in the reported forecasts of professional forecasters, financial market participants, consumers, firms and central bankers. Second, the variation in point estimates of information rigidity across shocks and agents is relatively small: we cannot reject the null hypothesis that the estimates are identical across all agents and shocks.¹² Third, all of the point estimates imply economically significant degrees of information rigidity. For example, the average estimate of 0.82 across all specifications is very close to the 0.75 value assumed by Mankiw and Reis (2002) which delivered very persistent effects of monetary policy shocks stemming only from sticky-information in price-setting. In addition, one should note that for each type of agent, one cannot reject the null that the degree of information rigidity is identical across shocks. This result is consistent with the prediction of sticky information models in which the rate of information updating is common across shocks. While noisy information models predict that the rate of information processing can differ across shocks, the absence of sharp differences across shocks in Table 4 suggests that such heterogeneity in information processing rates across our shocks is unlikely to be economically significant.

The results in Table 4 also indicate that there do not appear to be significant differences in the rate of information acquisition and processing across agents. While similar degrees of information processing for firms, professional forecasters, and central bankers may not be particularly surprising, the fact that consumers appear to process information at a rate no lower than other agents is more at odds with common wisdom. Carroll (2003), for example, proposes an epidemiological model in which consumers gradually acquire information from professional forecasters by occasionally reading news reports. In such a model, the convergence rate of consumer forecasts to the full-information levels should be significantly slower than that of professional forecasters, contrary to what we obtain in Table 4.

To reconcile our results with Carroll (2003), we revisit the evidence that he provides using a longer time sample. First, he argues that the mean squared error (MSE) of SPF forecasts of future CPI inflation is substantially less than that of consumer forecasts. However, Carroll uses *core* CPI inflation to calculate forecast errors rather than the general CPI index. This is important since the SPF forecasts he uses are for the general CPI index, and consumers responding to the Michigan Survey are very unlikely to exclude food and energy prices when forecasting inflation. When we calculate MSE's of SPF and Michigan forecasts using the general CPI index, we find that consumer forecasts actually lead to lower MSE's than either the SPF forecasts of the CPI, both over the time period considered by Carroll (1981:3-2000:2) and the longer

¹² The p -values in Table 4 are constructed based on seemingly unrelated regressions. Note that while this procedure does not take into account generated regressors, these imply that we understate the uncertainty associated with estimates of information rigidity and therefore our p -values are, if anything, too low. Because we fail to reject the null of equality, the presence of generated regressors strengthens our point.

time sample now available (1981:3-2007:3).¹³ Second, Carroll uses Granger causality tests and finds that SPF forecasts Granger-cause consumer forecasts but that the reverse is not true. While we can reproduce his results over his time sample, over the longer sample the opposite is true: consumer forecasts Granger-cause SPF forecasts but not the reverse. Thus, Granger-causality tests yield little support for Carroll's model and appear to be exceedingly sensitive to time samples, lag lengths, etc.

The fact that the unconditional disagreement among consumers is much larger than for other agents would appear to be at odds with the fact that the estimated degrees of information rigidity are no larger for consumers. However, this fact can also suggest that some of cross-sectional heterogeneity in consumer forecasts is likely to reflect the different nature of this survey: namely that consumers have little time to think about their responses to a phone survey and therefore report values which can include significant noise as well as the absence of guidance in the question about what measure of inflation they meant to forecast. Because these types of errors are likely to average out across agents, the high unconditional level of disagreement among consumers need not be inconsistent with the relatively rapid response of mean forecasts to shocks.

5 Extensions and Robustness

In this section, we consider two factors related to our analysis in sections 3 and 4. First, we use forecasts of another macroeconomic variable, specifically the unemployment rate, rather than inflation forecasts. Second, we consider the use of alternative econometric specifications for the estimation of impulse responses of forecast errors to shocks.

5.1 Unemployment Forecasts

While all of our analysis has been focused on inflation forecasts, one can apply our empirical tests for information rigidities to other forecasts as well. In this section, we consider forecasts of the average unemployment rate over the next year. For professional forecasters in the SPF, we use the average and the standard deviation of the individual forecasts for unemployment over the next four quarters. For firms in the Livingston Survey, we use the average over the individual forecasts of the unemployment rate available two months prior to the forecast date, the forecast of the unemployment rate six months ahead, and the forecast of the unemployment rate twelve months ahead and construct the measure of disagreement in an analogous manner. For FOMC members, unemployment rate forecasts are available

¹³ Taylor (1999) and Mehra (2002) similarly conclude that Michigan consumer forecasts lead to similar, or even smaller, MSE's than the Survey of Professional Forecasters. We also found the same results using the Blue Chip Economic Indicators forecast of the CPI. Note that the start-date of 1981 reflects the availability of SPF forecasts of CPI inflation.

for the calendar year only, so quasi-year ahead mean forecasts and disagreement are constructed in the same way as for inflation. Ex-post unemployment rates are constructed to be consistent with the horizon of each set of forecasts. Note that we cannot apply our analysis to consumer forecasts of unemployment. While the Michigan survey includes a question about future unemployment, it only asks respondents whether they expect unemployment to rise, fall or stay the same which is insufficient to test the theoretical predictions of the models in section 2.

Figure 6 plots the mean forecasts of unemployment for firms, professional forecasters and FOMC members as well as the level of disagreement. As was the case with inflation forecasts, unemployment forecasts track the level of unemployment closely and differences between the two are persistent. Unlike the inflation forecasts, however, there has been little change in average levels of disagreement over time. There is also limited evidence that disagreement about unemployment is higher during recessions: while this may appear to be the case with the SPF, both firm and central bank forecasts fell during the 1991 recession and did not experience unusual increases during the 2001 recession.

To assess the conditional response of forecast errors and disagreement to shocks, we focus on unidentified shocks to the unemployment rate, constructed in an analogous manner to the inflation unidentified shocks. This choice reflects the fact that technology shocks, news shocks and oil price shocks account for little of the historical variation of unemployment in a VAR and, as demonstrated in Appendix B, our empirical tests require shocks to account for a non-trivial fraction of the variation in a forecasted time series in small samples. By construction, these unidentified shocks are associated with higher unemployment rates. Figure 7 plots the estimated impulse responses of forecast errors from professional forecasters, firms and FOMC members to these shocks. As was the case with inflation forecasts, we can strongly reject the null of a zero response of forecast errors. In addition, forecast errors are positive and converging to zero over time, thereby conforming closely to the prediction of models with information rigidities. Figure 7 also plots the response of forecast errors to the absolute value of the shocks, for which we cannot generally reject the null of no response. Hence, these results confirm the finding using inflation forecasts that heterogeneous loss-aversion is unlikely to be an important component of the expectations formation process for these agents.

Panel B in Table 3 reports estimates of equation (36) using unemployment forecasts. Similar to inflation forecasts, we find no evidence that forecast errors, conditional on past forecast errors, are correlated with the level of the unemployment rate: point estimates on past unemployment rates are small, often negative, and not statistically different from zero. This finding again suggests that heterogeneity about long-run means and heterogeneous learning rates are not quantitatively important components of the expectations formation process for these agents. Figure 7 plots the conditional response of

disagreement to shocks, which can also be used to distinguish among different models of information rigidities. For firms and central bankers, we find no evidence that disagreement is sensitive to the absolute value of shocks and the estimated responses lie well below the predicted levels from sticky information. For professional forecasters, there is a small positive response of disagreement to unidentified unemployment shocks, consistent with the fact that disagreement among professional forecasters rises during recessions. However, the estimated response is much smaller than what one would expect under sticky-information.

Finally, Panel B of Table 4 reports estimated degrees of information rigidity with respect to unemployment rates for each type of agent from estimating the persistence of the conditional response of forecast errors normalized by the conditional response of unemployment to unidentified shocks. These estimates range from a low of 0.46 to a high of 0.65 which is statistically significantly lower than magnitudes found for inflation forecasts although these values still imply economically significant degrees of information rigidity. Coibion and Gorodnichenko (2010) similarly document significant heterogeneity in the degree of information rigidity across forecasts of different macroeconomic variables from professional forecasters. Since the unemployment rate is much more persistent than inflation in our sample, the finding that informational rigidity is lower for unemployment than for inflation is consistent with the noisy information models. Indeed, these models generically predict that *ceteris paribus* agents should allocate more attention to more persistent variables as the cost of mistakes in terms of utility or mean squared errors for more persistent series is larger. On the other hand, such a finding is at odds with the prediction of the sticky-information model in section 2—in which the degree of information rigidity is common across variables—although Mankiw and Reis (2011) suggest that sticky information models could readily be extended to incorporate heterogeneity in attentiveness to different macroeconomic variables.

The evidence from forecasts of the unemployment rate is thus broadly in line with those found using inflation forecasts. First, forecast errors respond to shocks in the same direction as the variable being forecasted, which is predicted by all models of information rigidities considered in section 2, and we can strongly reject the null implied by full-information rational expectations. Second, we find little evidence that forecast errors respond systematically to the absolute value of shocks, as predicted by models with heterogeneity in loss-aversion. This strongly suggests that the gradual response of forecasts to shocks does indeed reflect information rigidities on the part of economic agents. Third, we find no evidence that forecast errors are correlated with past levels of the variable being forecasted, which indicates that heterogeneous beliefs about long-run means and heterogeneity in information acquisition rates across agents are unlikely to be economically significant components of the expectations formation process for these agents. Finally, the

conditional response of disagreement to shocks is rarely statistically different from zero and consistently is much lower than would be expected under sticky information, thereby pointing to the baseline noisy information model as the most adequate representation of the expectations formation process for professional forecasters, firms, and central bankers.

5.2 Robustness to Alternative Estimation Procedures

All of our results are based on a two-step procedure in which we first estimate shocks and then recover impulse responses of relevant variables to these shocks. This procedure has several advantages. First, it allows for flexibility in the first step by utilizing a number of different approaches to recover measures of structural shocks. Second, we are able to recover impulse responses of a number of different variables (e.g. forecast errors, disagreement) to both the levels of the shocks as well as the absolute value of the shocks. Nonetheless, there are a number of alternative procedures that could be used and in this section we investigate the robustness of our baseline findings to these alternative econometric specifications.

The most common procedure to recover impulse responses to structural shocks is to estimate a VAR. A VAR simultaneously estimates the shocks and impulse responses to the shocks, and therefore yields standard errors which incorporate the uncertainty associated with the estimation of the shocks. On the other hand, the VAR cannot readily recover responses to the absolute value of shocks, which we use in many empirical tests in the paper. Despite this, we can assess whether a VAR would yield qualitatively similar responses of forecast errors to the level of the shocks as our baseline specification. In the interest of space, we focus on the identification of technology shocks and the response of forecast errors from the SPF. We estimate a 4-variable VAR(4) including the change in labor productivity, the change in hours worked, the annual inflation rate, and the mean forecast error of annual inflation from the SPF using data from 1976 to 2007. Following Gali (1999), technology shocks are identified as those innovations which have permanent effects on labor productivity. Impulse responses of forecast errors to technology shocks from the VAR as well as our baseline specification are presented in Panel A of Figure 8. The point estimates for the impulse response are very similar to our baseline results, as are the standard errors. We can again strongly reject the null of zero response of forecast errors.

The fact that the standard errors from the VAR are not much larger than those from our two-step procedure may seem surprising given that the VAR takes into account the additional uncertainty surrounding the values of the shocks whereas our approach, which is valid under the null that the coefficients on the shocks are zero per Pagan (1984), does not. The main reason for this negligible difference is that the generated regressor in our empirical specification, i.e. the structural shock, is the residual from the first stage, not the fitted value. Recall that the adjustment is typically necessary because error terms in the first and second stage regressions could be correlated. However, in our case, by the

properties of least squares the generated regressor (which is an estimate of the error term in the first stage) is orthogonal to the error term in the second stage and, as a result, estimates in the first and the second stage regressions are approximately uncorrelated. To show how little generated regressors matter in this case, we apply the correction for generated regressors from Murphy and Topel (1985) to our two-step estimates of the effects of technology shocks on SPF inflation forecast errors. Panel B of Figure 8 presents both corrected and uncorrected confidence intervals: the effect of explicitly controlling for generated regressors is negligible.

We also consider alternatives to our procedure for estimating impulse responses in the second step. Our approach follows Romer and Romer (2004) who regress macroeconomic variables on lags of themselves and contemporaneous and lagged values of the shock. The basis for such a specification is that the reduced form of macroeconomic variables has a moving average representation which we approximate in the empirical specification. An alternative, as in Cochrane and Piazzesi (2002), is to estimate the moving average representation directly by regressing the variable of interest only on contemporaneous and lagged values of the shock, i.e.

$$x_t = c + \sum_{k=0}^K \varepsilon_{t-k} + v_t \quad (39)$$

where ε is the structural shock. Note that specification (39) a special case of our second step specification such as equation (32) in which coefficients on autoregressive terms are set equal to zero. We illustrate this alternative procedure in Panel C of Figure 8 for SPF inflation forecast errors after technology shocks. The point estimates are very close to our baseline ones but the standard errors are larger. This reflects the fact that the Romer and Romer (2004) approach allows for additional autoregressive terms which control for dynamics arising from other shocks through the AR terms. In short samples, this delivers more precise estimates.

Another approach to estimating impulse responses is the local projection method of Jorda (2005) which consists of running separate regressions for each horizon h of the impulse response of the form

$$x_{t+h} = c + \sum_{k=1}^K \beta_{h,k} Z_{t-k} + \gamma_h \varepsilon_t + v_{t+h} \quad (40)$$

where ε is the structural shock, Z is a vector of controls. The impulse response at the horizon h is the estimated coefficient γ_h . Applying this specification to the response of SPF inflation forecast errors to technology shocks yields almost identical results as the Cochrane and Piazzesi (2002) approach, see Panel D of Figure 8. The point estimates also similar to our baseline impulse response functions but the standard errors are again somewhat larger.

In short, these results indicate that our results are robust to using alternative empirical approaches to construct impulse responses to identified shocks. Using a VAR to estimate the shocks and responses in one step does not qualitatively alter our results largely because generated regressors do not qualitatively

affect our standard errors. Our two-step procedure has the additional advantage over VAR's of being able to deliver responses to both the levels and absolute values of shocks, which is necessary to test some of the theoretical predictions of the models. Similarly, alternative specifications for the second step lead to similar results: the response of SPF forecast errors to technology shocks is consistently negative, as predicted by models of information rigidities, and points to important deviations from the null of full-information rational expectations.

6 Conclusion

While there has been growing interest in integrating deviations from full-information in macroeconomic models, a key stumbling block has been the absence of robust evidence about the quantitative importance and nature of information rigidities faced by economic agents. Building on the predictions of a variety of models with information frictions, we document systematic evidence of a delayed response of mean forecasts to macroeconomic shocks for professional forecasters, consumers, firms, central bankers, and financial market participants consistent with the predictions of imperfect information models. Furthermore, the implied degrees of information rigidities are economically large and consistent with significant macroeconomic effects. This justifies the burgeoning interest in imperfect information models and provides a set of stylized facts that models should be consistent with.

In particular, the fact that information rigidities appear to be large for all agents suggests that future work should go beyond focusing on the effects of information rigidities on price setting decisions and work toward a systematic integration of these frictions into all components of macroeconomic models. Mankiw and Reis (2007) and Reis (2009), for example, take an important step in this direction by integrating information rigidities in consumption, wage-setting and price-setting decisions. The fact that FOMC members also appear to be subject to significant constraints on information processing indicates that incorporating rigidities on the part of the central bank is likely to be an important contribution as well. On the other hand, much of our empirical evidence suggests that noisy information models are likely to be the most appropriate characterization of the expectations formation process for professional forecasters, consumers, firms and central banks. There has as of yet been little work attempting to systematically incorporate this type of rigidity into all of the optimizing decisions in macroeconomic models. Mackowiak and Wiederholt (2009b), for example, is the first paper to integrate rational inattention on the part of both consumers and firms into a DSGE model, but like Mankiw and Reis (2007), they do not incorporate information rigidities on the part of the central bank into the model. This is likely to be a fruitful, if challenging, area for future work.

The results in the paper highlight the usefulness of survey data. The availability of direct measures of agents' forecasts allows us to assess the predictions of different models of the expectations formation process without having to take a stand on many auxiliary issues such as the nature of price-setting decisions. It is also clear that there are limitations to survey data. While forecasts from professional forecasters are available for long periods, different forecasting horizons, and a number of macroeconomic variables, surveys of forecasts for other agents are more limited. Consumer forecasts are particularly problematic: few questions require respondents to provide a quantitative answer but even those that do, such as the question about future prices, fail to specify a specific index to forecast. This leads to much more disagreement in the responses than is the case in other surveys. Forecasts from firms and central bankers are also limited: the Livingston survey is semi-annual and includes forecasts only from large institutions, whereas one would ideally like to have a representative survey of firms' expectations. FOMC forecasts are not readily available at the individual level, the variable being forecasted can change over time, and the frequency of the data is also limited. Nonetheless, the answer to the question of "what can survey forecasts tell us about information rigidities?" is "a lot." These surveys, combined with the models' theoretical predictions, yield robust evidence of information rigidities for all of the agents we consider, provide guidance as to how best to model their expectation formation process, and point to the importance of more work on integrating information rigidities into modern macroeconomic models to fully spell out their potential implications.

References

- Andrade, P., and H. Le Bihan, 2010. "Inattentive Professional Forecasters," mimeo.
- Adam, Klaus and Mario Padula, 2003. "Inflation Dynamics and Subjective Expectations in the United States," European Central Bank Working Paper 222.
- Ang, Andrew, Geert Bekaert, and Min Wei, 2007. "Do macro variables, asset markets, or surveys forecast inflation better?" *Journal of Monetary Economics* 54(4), 1163-1212.
- Ang, Andrew, Geert Bekaert, and Min Wei, 2008. "The Term Structure of Real Rates and Expected Inflation," *Journal of Finance* 63(2), 797-849.
- 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.
- Barsky, Robert, and Eric Sims, 2010. "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence," forthcoming *American Economic Review*.
- Barsky, Robert, and Eric Sims, 2011. "News Shocks and Business Cycles," forthcoming in *Journal of Monetary Economics*.
- Berger, Helge, Michael Ehrmann, and Marcel Fratzscher, 2011. "Geography or Skills: What Explains Fed Watchers' Forecast Accuracy of US Monetary Policy," *Journal of Macroeconomics* 33(3), 420-437.
- Bloom, Nicholas, 2009. "The Impact of Uncertainty Shocks," *Econometrica* 77(3), 623-685.
- Bordo, Michael, Christopher Erceg, Andrew Levin, and Ryan Michaels, 2007. "Three great American disinflations," *Proceedings*, Federal Reserve Bank of San Francisco.
- Branch, William A., 2007. "Sticky information and model uncertainty in survey data on inflation expectations," *Journal of Economic Dynamics and Control* 31, 245-276.

- Capistrán, Carlos, and Allan Timmermann, 2009. "Disagreement and Biases in Inflation Expectations," *Journal of Money, Credit and Banking* 41(2-3), 365-396.
- Carroll, Christopher D. 2003. "Macroeconomic Expectations of Households and Professional Forecasters," *Quarterly Journal of Economics* 118(1), 269-298.
- Christiano, Lawrence, Martin Eichenbaum, and Charles Evans. 1999. "Monetary Policy Shocks: What Have We Learned, and To What End?" In *Handbook of Monetary Economics* 1999, ed. John B. Taylor and Michael Woodford, 65-148. Amsterdam: Elsevier Science.
- Cochrane, John, and Monika Piazzesi, 2002. "The Fed and Interest Rates: A High-Frequency Identification," *American Economic Review P&P*, 92(2), 90-95.
- Coibion, Olivier, 2010. "Testing the Sticky Information Phillips Curve," *Review of Economics and Statistics* 92(1), 87-101.
- Coibion, Olivier, and Yuriy Gorodnichenko, 2010. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," NBER Working Paper No. 16537.
- Cukierman, Alex, and Paul Wachtel, 1979. "Differential Inflationary Expectations and the Variability of the Rate of Inflation: Theory and Evidence," *American Economic Review* 69(4), 595-609.
- Dopor, Bill, Tomiyuki Kitamura, and Takayuki Tsuruga, 2010. "Integrating Sticky Information and Sticky Prices," *Review of Economics and Statistics* 92(3), 657-669.
- Dopor, Bill, Jing Han, and Yi Chan Tsai, 2009. "What Do Technology Shocks Tell Us about the New Keynesian Paradigm?" *Journal of Monetary Economics* 56(4), 560-569. .
- Elliot, Elliott, Ivana Komunjer, and Allan Timmermann, 2008. "Biases in Macroeconomic Forecasts: Irrationality or Asymmetric Loss?" *Journal of the European Economic Association* 6(1), 122-157.
- Gali, Jordi, 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89(1), 249-271.
- Gorodnichenko, Yuriy, 2006. "Endogenous Information, Menu Costs, and Inflation Persistence," Manuscript.
- Gourinchas, Pierre-Olivier and Aaron Tornell, 2004. "Exchange-Rate Puzzles and Distorted Beliefs" *Journal of International Economics* 64, 303-333.
- Hamilton, James D., 1996. "This is what happened to the oil price-macroeconomy relationship," *Journal of Monetary Economics* 38, 215-220.
- Haubrich, Joseph G., George Pennacchi, and Peter Ritchken, 2011. "Inflation Expectations, Real Rates, and Risk Premia: Evidence from Inflation Swaps," FRB of Cleveland Working Paper 11-007.
- Jorda, Oscar, 2005. "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review* 95(1), 161-182.
- 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.
- Klenow, Peter J. and Jonathan L. Willis, 2007. "Sticky Information and Sticky Prices," *Journal of Monetary Economics* 54(S), 79-99.
- Knotek, Edward S., 2010. "A Tale of Two Rigidities: Sticky Prices in a Sticky Information Environment," *Journal of Money, Credit and Banking* 42(8), 1543-64.
- Korenok, Oleg, 2005. "Empirical Comparison of Sticky Price and Sticky Information Models." Working Paper 0501, Virginia Commonwealth University School of Business.
- Kydland, Finn E. and Edward C. Prescott, 1982. "Time to Build and Aggregate Fluctuations," *Econometrica* 50(6), 1345-1370.
- Lucas, Robert E., 1972. "Expectations and the Neutrality of Money," *Journal of Economic Theory* 4(2), 103-124.
- Mackowiak, Bartosz and Mirko Wiederholt, 2009a. "Optimal Sticky Prices under Rational Inattention," *American Economic Review* 99(3), 769-803.
- Mackowiak, Bartosz and Mirko Wiederholt, 2009b. "Business Cycle Dynamics under Rational Inattention," manuscript.

- Mackowiak, Bartosz, Emanuel Moench and Mirko Wiederholt, 2009. "Sectoral Price Data and Models of Price Setting" *Journal of Monetary Economics* 56(S), 78-99.
- 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 and Ricardo Reis, 2007. "Sticky Information in General Equilibrium," *Journal of the European Economic Association* 5(2-3) 603-613.
- Mankiw, N. Gregory, and Ricardo Reis, 2011. "Imperfect Information and Aggregate Supply," in *Handbook of Monetary Economics* edited by B. Friedman and M. Woodford, Elsevier-North Holland, vol. 3A, 182-230, 2011.
- Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers, 2003. "Disagreement about Inflation Expectations," *NBER Macroeconomics Annual 2003*, Cambridge, MIT Press, 18, 209-248, 2004.
- Mehra, Yash P., 2002. "Survey Measures of Expected Inflation: Revisiting the Issues of Predictive Content and Rationality," Federal Reserve Bank of Richmond *Economic Quarterly* 88(3), 17-36.
- Morris, Stephen and Hyun Song Shin, 2002. "Social Value of Public Information," *American Economic Review* 92(5), 1521-1534.
- Murphy, Kevin M., and Robert H. Topel, 1985. "Estimation and Inference in Two-Step Econometric Models," *Journal of Business and Economic Statistics* 3(4), 370-79.
- 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 Weale. 2006. "Survey Expectations," in *Handbook of Economic Forecasting*, G. Elliott, C.W.J. Granger, and A. 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), 173-96.
- 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.
- Romer, Christina and David H. Romer, 2004. "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review* 94(4), 1055-1084.
- Romer, Christina D., and David H. Romer, 2008. "The FOMC versus the Staff: Where Can Monetary Policymakers Add Value?" *American Economic Review* 98(2), 230-235.
- Romer, Christina and David H. Romer, 2011. "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks," *American Economic Review* 100(3), 763-801.
- Romer, David H., 2009. "A New Data Set on Monetary Policy: The Economic Forecasts of Individual Members of the FOMC," NBER Working Paper 15208.
- Sims, Christopher A., 2003. "Implications of Rational Inattention," *Journal of Monetary Economics*, 50(3), 665-690.
- Taylor, Lloyd B., Jr., 1999. "Survey Measures of Expected U.S. Inflation," *Journal of Economic Perspectives* 13(4), 125-144.
- Woodford, Michael, 2001. "Imperfect Common Knowledge and the Effects of Monetary Policy," published in P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, eds., *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund Phelps*, Princeton Univ. Press, 2002.
- Woodford, Michael, 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

Table 1: Summary of predictions of different models

Predictions:	Model						
	Full- Information Rational Expectations (FIRE)	Heterogeneous Loss-Aversion under FIRE	Sticky Info	Noisy-Information			
				Baseline	Strategic interaction	Model-Heterogeneity	
		Heterogeneity about Long-Run Means	Heterogeneity in Gains of the Kalman filter				
(1)	(2)	(3)	(4)	(5)	(6)	(7)	
Response of Forecast Errors to Shocks	No response	All positive or negative, asymptotically decline	Same direction as forecasted variable, asymptotically decline	Same direction as forecasted variable, asymptotically decline	Same direction as forecasted variable, asymptotically decline	Same direction as forecasted variable, correlated with past levels of forecasted variable.	Same direction as forecasted variable, correlated with past levels of forecasted variable.
Speed of Convergence of Normalized Forecast Errors to Shocks	Immediate Convergence	Same across shocks	Same across shocks	May differ across shocks	May differ across shocks	May differ across shocks	May differ across shocks
Response of Disagreement to Shocks	No response	Positive for any shock	Positive for any shock	No response	No response	No response	Positive for any shock

Notes: This table summarizes predictions of the models presented in Section 2.

Table 2: Decomposition of Inflation Volatility by Structural Shocks

Forecast Horizon	Share of Inflation Volatility Explained by Shocks:			
	Technology Shocks	News Shocks	Oil Price Shocks	Unexplained
1	28.5	1.4	0.3	69.8
4	25.9	5.8	9.6	58.7
8	23.7	7.2	11.4	57.7
12	22.8	7.8	11.4	58.0
20	22.0	8.2	11.7	58.1

Notes: Variance decomposition is based on a VAR(4) estimated on 1966Q1-2007Q3 sample.

Table 3: Sensitivity of Forecast Errors to Lagged Values.

Dependent variable FE_{t+h}	Survey measures								Financial markets			
	Professional Forecasters (SPF)		Consumers (MSC)		Firms (Livingston)		FOMC Members		Haubrich et al. (2011)		Ang et al. (2007)	
	Quarterly		Quarterly		Semi-Annual		Semi-Annual		Quarterly		Quarterly	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A: Inflation rate												
FE_{t+h-1}	0.89*** (0.05)	0.88*** (0.04)	0.76*** (0.08)	0.76*** (0.08)	0.83*** (0.15)	0.82*** (0.13)	0.40*** (0.11)	0.39*** (0.10)	0.76*** (0.08)	0.75*** (0.08)	0.85*** (0.07)	0.87*** (0.06)
π_{t-1}	-0.07 (0.12)	-0.02 (0.03)	0.01 (0.15)	0.01 (0.03)	-0.12 (0.13)	-0.09** (0.04)	-0.13** (0.06)	-0.13*** (0.03)	-0.20 (0.13)	-0.08 (0.08)	0.36 (0.24)	-0.03 (0.04)
π_{t-2}	0.05 (0.11)		-0.00 (0.15)		0.04 (0.13)		0.01 (0.07)		0.16 (0.14)		-0.40 (0.27)	
N	123	123	124	124	59	59	53	54	98	98	127	127
R ²	0.77	0.77	0.58	0.58	0.63	0.63	0.39	0.40	0.56	0.55	0.77	0.75
Panel B: Unemployment rate												
FE_{t+h-1}	0.84*** (0.07)	0.82*** (0.07)			0.54*** (0.11)	0.57*** (0.12)	0.57*** (0.12)	0.52*** (0.13)				
UR_{t-1}	-0.16 (0.23)	-0.02 (0.03)			0.06 (0.18)	-0.04 (0.04)	-0.15 (0.22)	-0.06 (0.07)				
UR_{t-2}	0.14 (0.22)				-0.11 (0.16)		0.09 (0.20)					
N	123	123			59	59	53	54				
R ²	0.68	0.68			0.38	0.37	0.35	0.32				


Notes: The table presents least-squares estimates for specification (36) and augmented versions of specification (36) for inflation rate π_t and unemployment rate UR_t . The dependent variable in Panel A is the forecast error for inflation $FE_{t+h} \equiv \pi_{t+h,t} - \overline{\pi_{t+h,t}|t}$. The dependent variable in Panel B is the forecast error for unemployment rate $FE_{t+h} \equiv UR_{t+h,t} - \overline{UR_{t+h,t}|t}$. Newey-West robust standard errors are in parentheses. ***, **, * indicates statistical significance at 1, 5 and 10 percent.

Table 4: Convergence Rates of Forecasts.

	Results using Survey Data					Results using Forecasts from Financial Markets	
	Professional Forecasters	Consumers	Firms	FOMC Members	<i>p</i> -value for equality across agents	Haubrich et al. (2011)	Ang et al. (2007)
Panel A: Inflation Rate							
Technology Shocks	0.86 (0.05)	0.80 (0.10)	0.89 (0.09)	0.86 (0.08)	0.91	0.87 (0.10)	0.90 (0.09)
News Shocks	0.89 (0.05)	0.81 (0.10)	0.89 (0.08)	0.88 (0.09)	0.89	0.87 (0.09)	0.90 (0.09)
Oil Price Shocks	0.88 (0.05)	0.74 (0.07)	0.86 (0.07)	0.59 (0.13)	0.04	0.82 (0.10)	0.87 (0.07)
Unidentified Shocks	0.88 (0.05)	0.74 (0.09)	0.89 (0.06)	0.87 (0.08)	0.44	0.90 (0.10)	0.92 (0.08)
<i>p</i> -value for equality across shocks	0.98	0.90	0.98	0.19	0.60		
Panel B: Unemployment Rate							
Unidentified Shocks	0.46 (0.07)	-	0.65 (0.04)	0.52 (0.07)	0.15		

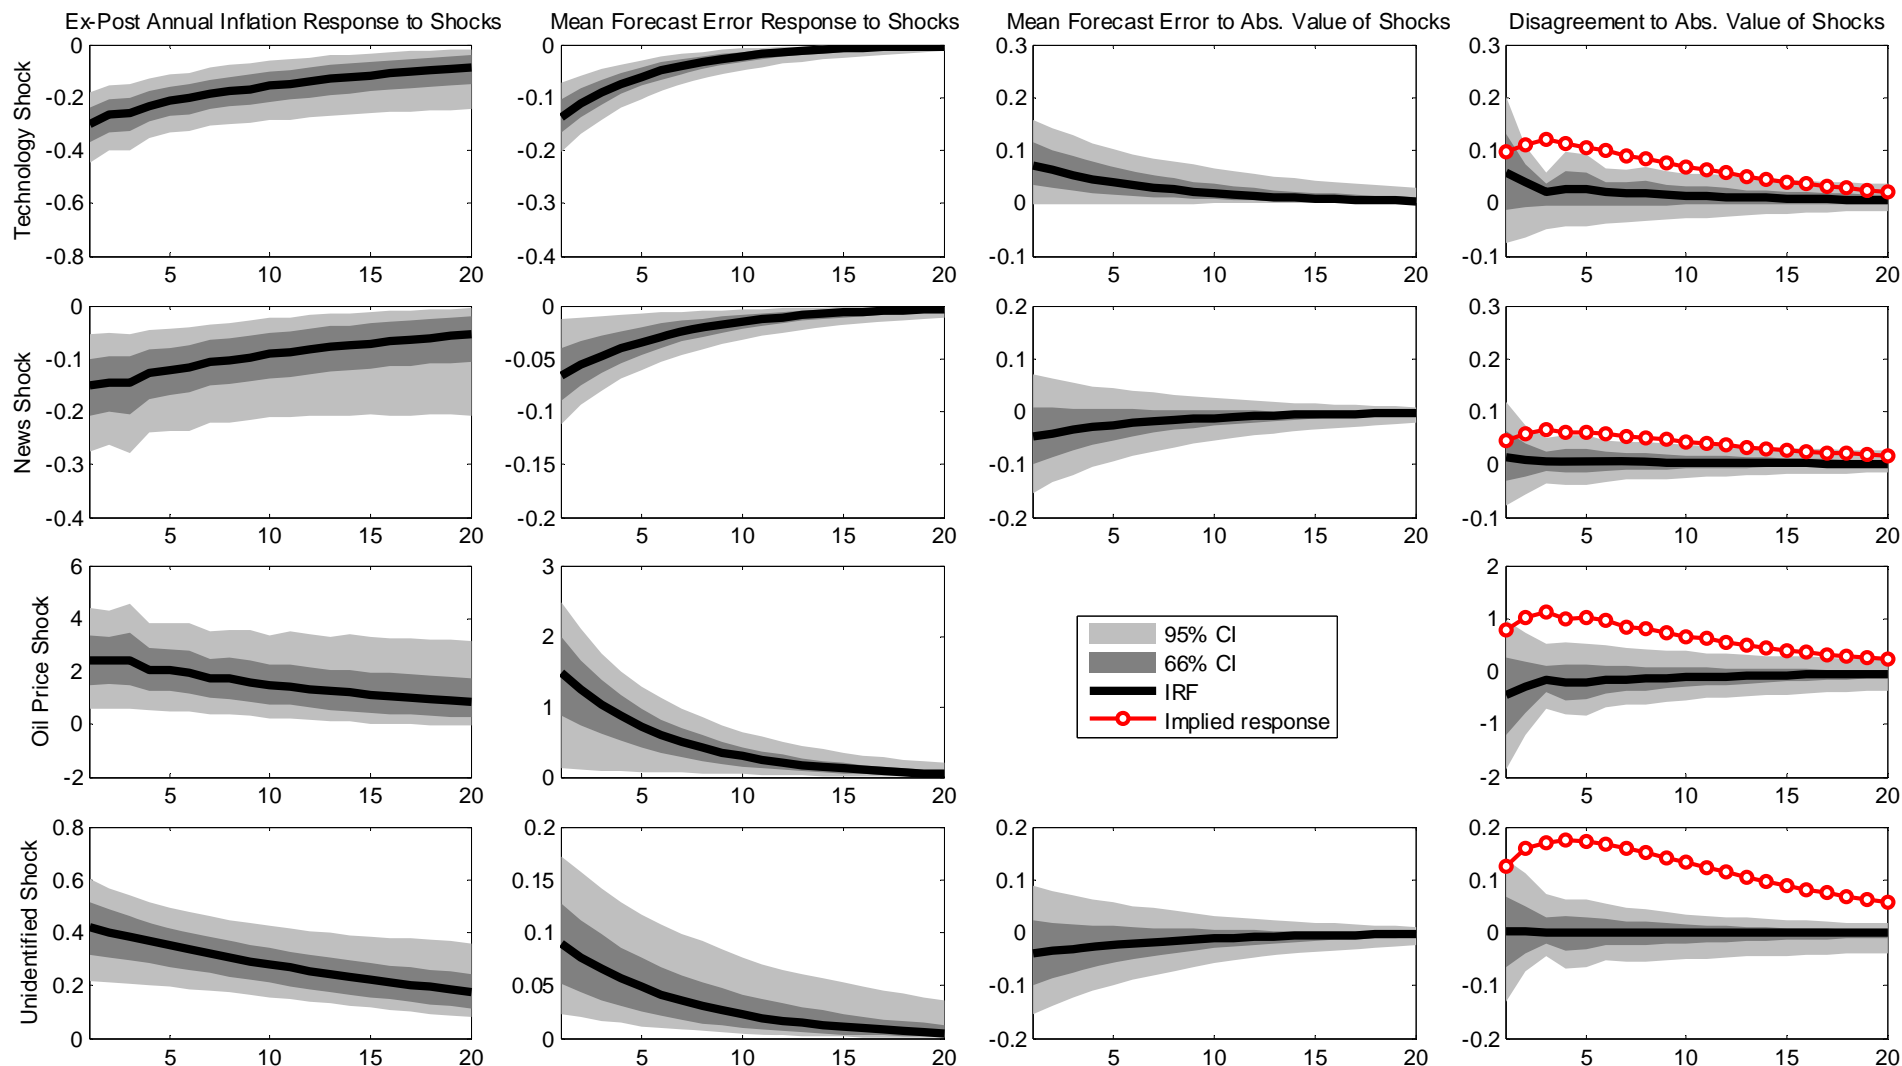
Notes: All estimates are of the persistence of response of forecast errors normalized by the response of the forecasted variable to shocks, 1976-2007 or as available. Newey-West robust standard errors are in parentheses. All estimates are converted to quarterly frequency for comparison.

Figure 1. Time series: Inflation rate.



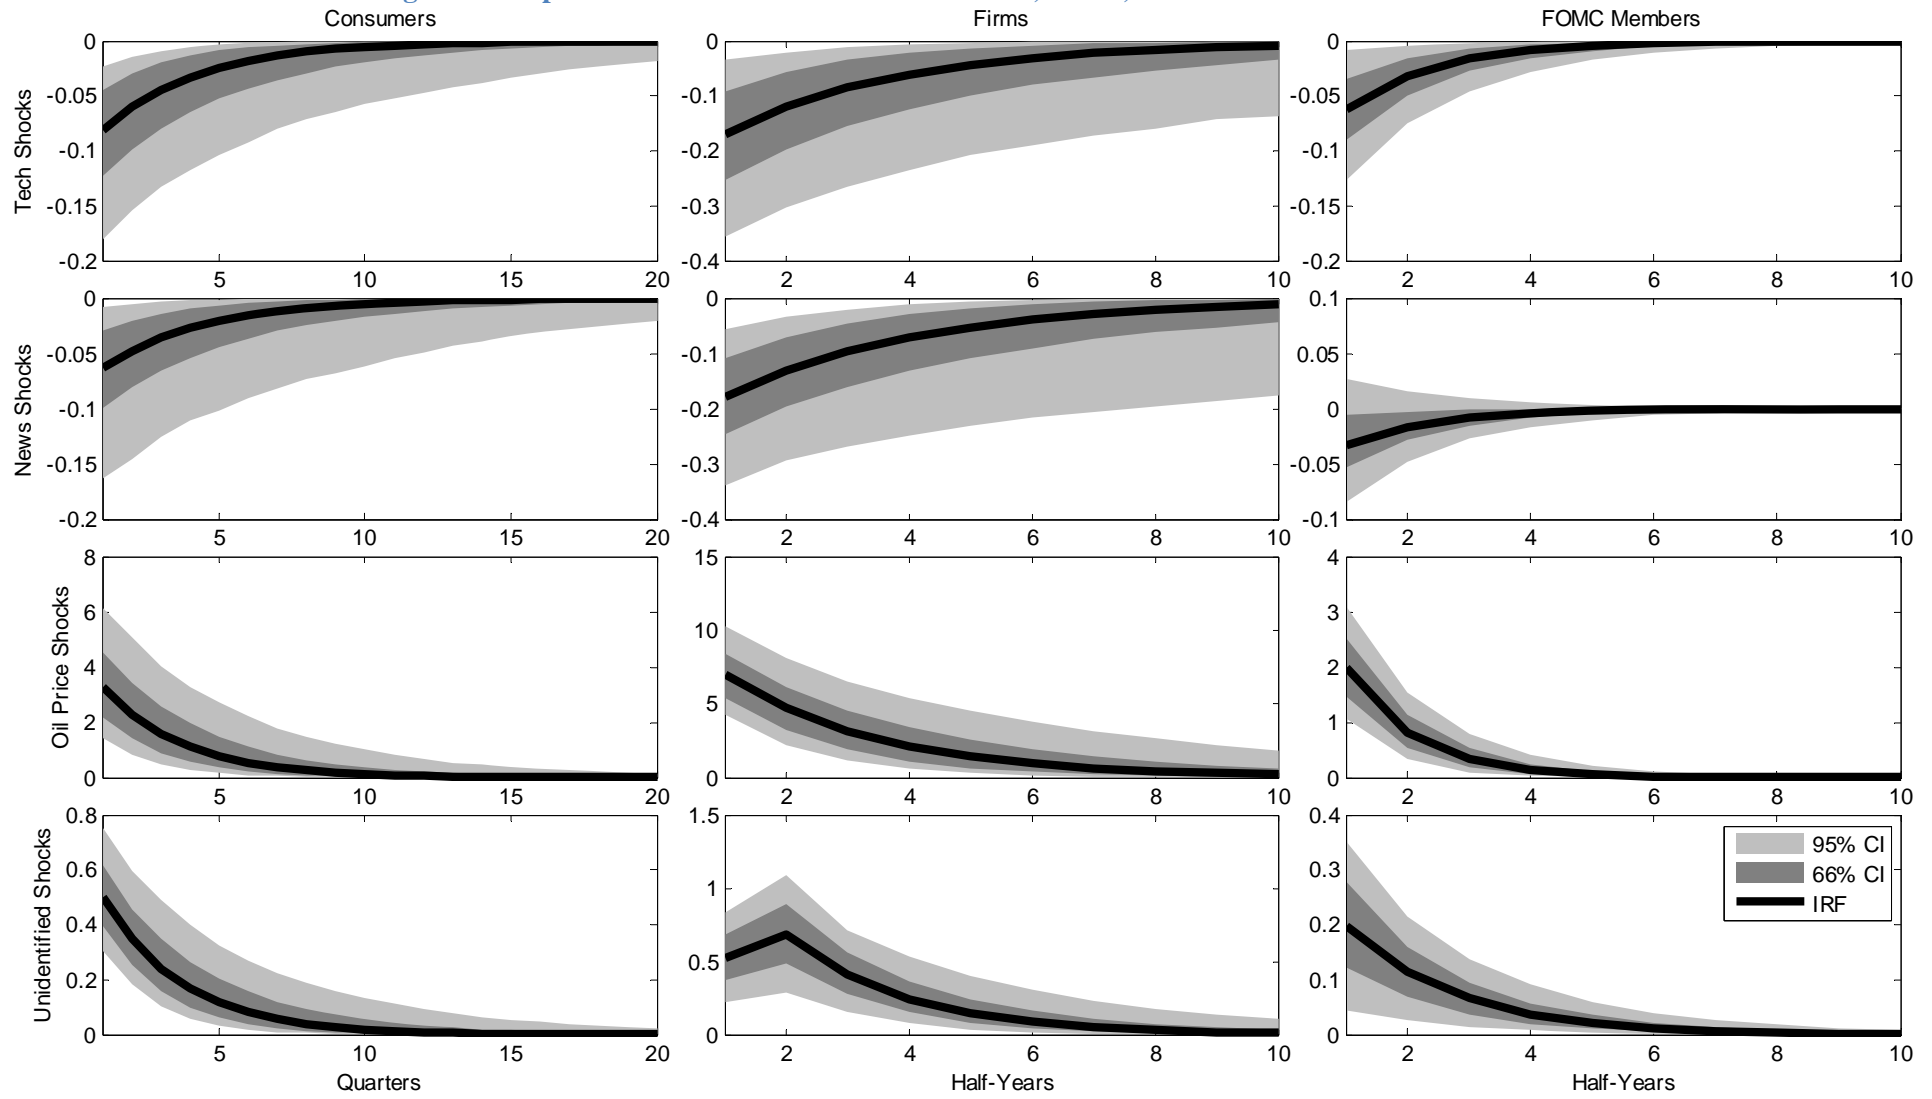
Notes: Each panel presents time series of actual inflation, mean one-year-ahead forecast, and forecast disagreement. The disagreement is measured as the standard deviation of the cross-section of reported forecasts. The actual inflation is based on GDP deflator for panel A and on the Consumer Price Index for panels B and C, and the different inflation rates forecasted by the FOMC in Panel D. Shaded regions indicate recessions dated by the NBER.

Figure 2. Baseline Results for Professional Forecasters.




Notes: The figure reports impulse responses to a unit shock computed from estimated specification (32) (column 1), (34) (column 2), (35) (column 3), and (38) (column 4). Each row shows responses to a given structural shock. The red line with circles in column 4 presents the implied response of forecast disagreement from the sticky information model given the information rigidity estimated for a given agent and shock. Standard errors for impulse responses are computed using parametric bootstrap.

Figure 3. Response of Forecast Errors of Consumers, Firms, and FOMC Members to Shocks.



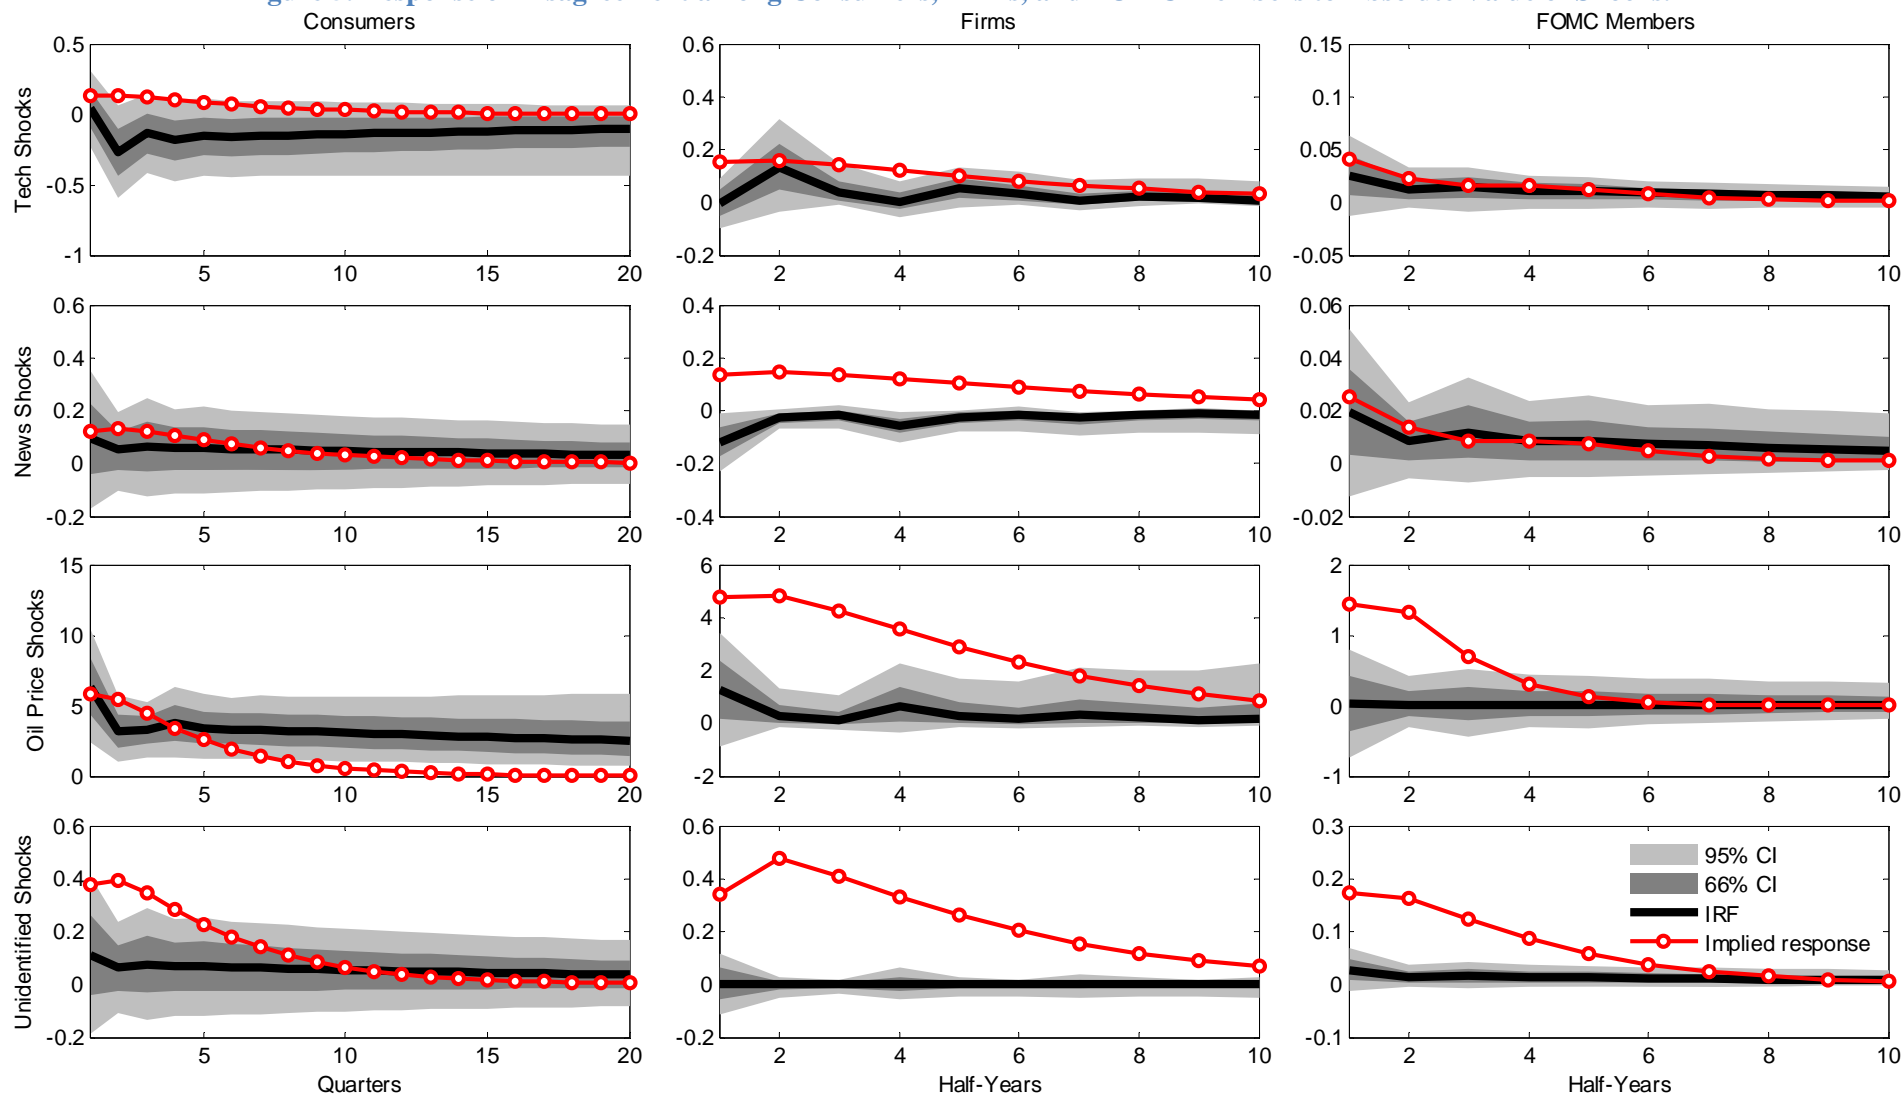
Notes: The figure reports impulse responses to a unit shock computed from estimated (34). Each row shows responses to a given structural shock. Standard errors for impulse responses are computed using parametric bootstrap. The left column is based on the forecasts reported in the Michigan Survey of Consumers. The middle column is based on firms' forecasts reported in the Livingston survey. The right column is based on the forecasts of FOMC members.

Figure 4. Response of Forecast Errors of Consumers, Firms, and FOMC Members to Absolute Value of Shocks.



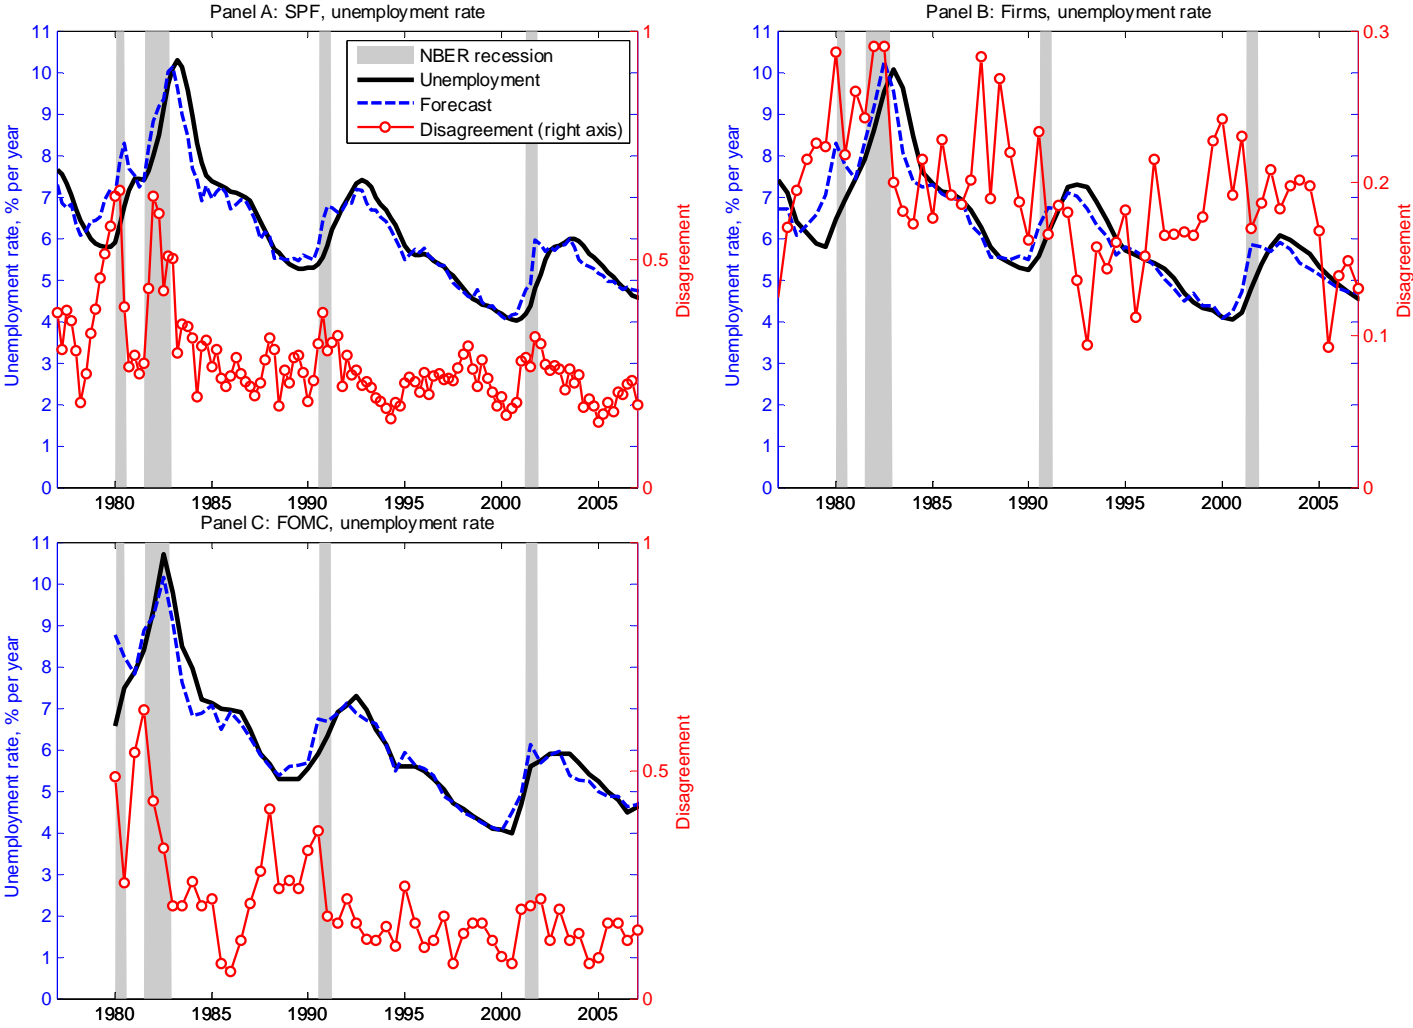
Notes: The figure reports impulse responses to a unit shock computed from estimated specification (35). Each row shows responses to a given structural shock. Standard errors for impulse responses are computed using parametric bootstrap. The left column is based on the forecasts reported in the Michigan Survey of Consumers. The middle column is based on firms' forecasts reported in the Livingston survey. The right column is based on the forecasts of FOMC members.

Figure 5. Response of Disagreement among Consumers, Firms, and FOMC Members to Absolute Value of Shocks.




Notes: The figure reports impulse responses to a unit shock computed from estimated specification (38). Each row shows responses to a given structural shock. The red line with circles presents the implied response of forecast disagreement from the sticky information model given the information rigidity estimated for a given agent and shock. Standard errors for impulse responses are computed using parametric bootstrap. The left column is based on the forecasts reported in the Michigan Survey of Consumers. The middle column is based on firms' forecasts reported in the Livingston survey. The right column is based on the forecasts of the FOMC members.

Figure 6. Time series: Unemployment rate.




Notes: Each panel presents time series of actual unemployment rate, mean one-year-ahead forecast, and forecast disagreement. The disagreement is measured as the standard deviation of the cross-section of reported forecasts. Shaded regions indicate recessions dated by the NBER.

Figure 7. Response of Forecast Errors and Disagreement of Professional Forecasters, and FOMC Members for the Unemployment Rate.



Notes: The figure reports impulse responses to a unit shock computed from estimated specification (34) (row 1), (35) (row 2), and (38) (row 3). Each column shows responses for a given type of agents. The left column is based on the forecasts reported in the Survey of Professional Forecasters. The middle column is based on firms' forecasts reported in the Livingston survey. The right column is based on the forecasts of FOMC members. The red line with circles in row 3 presents the implied response of forecast disagreement from the sticky information model given the information rigidity estimated for a given agent and shock. Standard errors for impulse responses are computed using parametric bootstrap.

Figure 8. Robustness check: alternative ways to compute IRFs and standard errors.



Notes: The figures show the sensitivity of estimated impulse response responses and associated standard errors to using alternative approaches. All results are for the mean forecast errors in the Survey of Professional Forecasters and the technology shock. Panel A compares impulse responses for specification (34) (two-step) and a one-step approach when a VAR used to identify technology shocks is augmented with the SPF mean forecast errors for inflation. Panel B shows how confidence bands are affected by the Murphy and Topel (1985) correction for generated regressors. Panel C shows the impulse response based on the Cochrane and Piazzesi (2002) approach. Panel D shows the impulse response based on the local projections approach developed in Jorda (2005). The vector of controls Z is the same as the set of variables and lags in the VAR.

Appendix A. Derivations and Extensions

1. Sticky information model.

In this section, we consider the properties of the sticky information model when the data generating process is a general MA(∞) rather than AR(1) as was presented in the text.

Suppose the reduced form process for inflation is MA(∞): $\pi_t = \sum_{k=0}^{\infty} \theta_k w_{t-k}$. Thus, impulse response of π to shock w_t is given by

$$\frac{d\pi_{t+h}}{dw_t} = \theta_h \quad \forall h \geq 0. \quad (1.1)$$

We are interested in two responses in particular: i) impulse response of forecast errors; ii) impulse response of dispersion.

The average forecast across agents at time t for inflation time $t+h$ is given by

$$\overline{\pi_{t+h|t}} = (1 - \lambda) \sum_{k=0}^{\infty} \lambda^k E_{t-k} \pi_{t+h} \quad (1.2)$$

where $E_{t-k} \pi_{t+h} = \sum_{s=k}^{\infty} \theta_{s+h} w_{t-s}$ so the average forecast follows

$$\begin{aligned} \overline{\pi_{t+h|t}} &= (1 - \lambda) \sum_{k=0}^{\infty} \lambda^k E_{t-k} \pi_{t+h} = (1 - \lambda) \sum_{k=0}^{\infty} \lambda^k [\theta_{k+h} w_{t-k} + \theta_{k+1+h} w_{t-k-1} + \dots] \\ &= \sum_{k=0}^{\infty} \theta_{k+h} w_{t-k} (1 - \lambda^{k+1}) \end{aligned}$$

The process for the forecast error is

$$\begin{aligned} FE_{t,t+h} &= \pi_{t+h} - \overline{\pi_{t+h|t}} = \sum_{k=0}^{\infty} \theta_k w_{t+h-k} - \sum_{k=0}^{\infty} \theta_{k+h} w_{t-k} (1 - \lambda^{k+1}) \\ &= \sum_{s=0}^{h-1} \theta_s w_{t+h-s} + \sum_{k=0}^{\infty} \theta_{k+h} w_{t-k} \lambda^{k+1}, \end{aligned} \quad (1.3)$$

where the first term corresponds to horizon effect and the second is the information rigidity delay. Then the impulse response of the forecast error to shocks is

$$\frac{dFE_{t+j,t+j+h}}{dw_t} = \theta_{j+h} \lambda^{j+1} = \lambda^{j+1} \left(\frac{d\pi_{t+j+h}}{dw_t} \right) \quad (1.4)$$

which corresponds to our result in the paper. Note also that we can then directly recover an estimate of λ by taking the ratio of the impulse response of the forecast error to that of the inflation response:

$$\frac{dFE_{t+j,t+j+h}}{dw_t} / \frac{d\pi_{t+j+h}}{dw_t} = \lambda^{j+1}. \quad (1.5)$$

To derive the impulse response of the forecast disagreement to a shock to inflation, it's useful to denote disagreement as a function of a given history of shocks $\mathbf{w}_t \equiv \{w_{t-s}\}_{s=0}^{\infty}$:

$$V_t \pi_t(\mathbf{w}_t) = (1 - \lambda) \sum_{k=0}^{\infty} \lambda^k (E_{t-k} \pi_t - \overline{\pi_{t|t}})^2$$

where

$$\begin{aligned} E_{t-k} \pi_t - \overline{\pi_{t|t}} &= \sum_{s=0}^{\infty} \theta_s w_{t-s} - \sum_{k=0}^{\infty} \theta_k w_{t-k} (1 - \lambda^{k+1}) \\ &= - \sum_{s=0}^{k-1} \theta_s w_{t-s} + \sum_{m=0}^{\infty} \theta_m w_{t-m} \lambda^{m+1} \end{aligned}$$

and therefore

$$\begin{aligned} (E_{t-k} \pi_t - \overline{\pi_{t|t}})^2 &= \sum_{s=0}^{k-1} \sum_{j=0}^{k-1} \theta_s \theta_j w_{t-s} w_{t-j} + \sum_{s=0}^{\infty} \sum_{j=0}^{\infty} \theta_s \theta_j w_{t-s} w_{t-j} \lambda^{s+1} \lambda^{j+1} \\ &\quad - 2 \sum_{s=0}^{k-1} \sum_{j=0}^{\infty} \theta_s \theta_j w_{t-s} w_{t-j} \lambda^{j+1} \end{aligned}$$

so that the variance of forecasts is $V_t \pi_t(\mathbf{w}_t) = V_t^A(\mathbf{w}_t) + V_t^B(\mathbf{w}_t) + V_t^C(\mathbf{w}_t)$ where

$$\begin{aligned} V_t^A(\mathbf{w}_t) &= \sum_{i=0}^{\infty} \sum_{j=0}^{\infty} \theta_i \theta_j w_{t-i} w_{t-j} \lambda^{i+1} \lambda^{j+1} \\ V_t^B(\mathbf{w}_t) &= (1 - \lambda) \sum_{k=0}^{\infty} \lambda^k \left[\sum_{s=0}^{k-1} \sum_{j=0}^{k-1} \theta_s \theta_j w_{t-s} w_{t-j} \right] \\ V_t^C(\mathbf{w}_t) &= -2(1 - \lambda) \sum_{k=0}^{\infty} \lambda^k \left[\sum_{s=0}^{k-1} \sum_{j=0}^{\infty} \theta_s \theta_j w_{t-s} w_{t-j} \lambda^{j+1} \right] \end{aligned}$$

To construct the average response of the variance of forecasts to a shock, we define a new history of shocks $\tilde{\mathbf{w}}_t$ which is identical to \mathbf{w}_t except for one observation $\tilde{w}_{t-m} = w_{t-m} + \delta$. It follows that

$$\begin{aligned}
V_t^A(\tilde{\mathbf{w}}_t) - V_t^A(\mathbf{w}_t) &= \theta_m \delta \lambda^{m+1} \sum_{s=0}^{\infty} \theta_s \tilde{w}_{t-s} \lambda^{s+1}, \\
V_t^B(\tilde{\mathbf{w}}_t) - V_t^B(\mathbf{w}_t) &= (1 - \lambda) \theta_m \delta \sum_{k=0}^{\infty} \lambda^k \left(\sum_{s=0}^{k-1} \theta_s \tilde{w}_{t-s} \right) = \theta_m \delta \sum_{s=0}^{\infty} \theta_s \tilde{w}_{t-s} \lambda^{s+1}, \\
V_t^C(\tilde{\mathbf{w}}_t) - V_t^C(\mathbf{w}_t) &= -2(1 - \lambda) \theta_m \delta \lambda^{m+1} \sum_{k=0}^{\infty} \lambda^k \left(\sum_{s=0}^{k-1} \theta_s \tilde{w}_{t-s} \right) = -2\theta_m \delta \lambda^{m+1} \sum_{s=0}^{\infty} \theta_s \tilde{w}_{t-s} \lambda^{s+1}.
\end{aligned}$$

Thus, the difference in the cross-sectional variance of forecasts from the two histories is

$$V_t \pi_t(\tilde{\mathbf{w}}_t) - V_t \pi_t(\mathbf{w}_t) = \sum_{S=\{A,B,C\}} V_t^S(\tilde{\mathbf{w}}_t) - V_t^S(\mathbf{w}_t) = \theta_m \delta (1 - \lambda^{m+1}) \sum_{s=0}^{\infty} \theta_s \tilde{w}_{t-s} \lambda^{s+1}.$$

The average response of the cross-sectional variance of forecasts to a shock δ at time $t - m$ is then

$$E[V_t \pi_t(\tilde{\mathbf{w}}_t) - V_t \pi_t(\mathbf{w}_t)] = \theta_m^2 \lambda^{m+1} (1 - \lambda^{m+1}) \delta^2 = \left(\frac{d\pi_{t+m}}{d\delta} \right)^2 \lambda^{m+1} (1 - \lambda^{m+1}) \delta^2.$$

2. Noisy information model.

This section extends the basic noisy information model from AR(1) to AR(p) case. In the context of AR(1), we have shown four key properties of the basic noisy information model. First, mean forecasts, forecast errors and forecasted series move in the same direction after a shock. Second, forecast errors asymptotically converge to zero, that is, agents eventually learn the true state of the fundamentals. Third, forecast disagreement is not a function of shocks to fundamentals. Fourth, one can recover the degree of information rigidity by dividing the impulse response of forecast errors by the impulse response of the forecasted series. In this appendix, we show that these properties generalize to AR(p) processes.

Consider the following state-space representation of the inflation process (i.e., the fundamental) and observed signal about inflation:

$$\text{State: } z_t \equiv \begin{bmatrix} \pi_t \\ \vdots \\ \pi_{t-p} \end{bmatrix} = \begin{bmatrix} b_1 & \dots & \dots & b_p \\ 1 & 0 & 0 & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \ddots & 1 & 0 \end{bmatrix} z_{t-1} + w_t = B z_{t-1} + w_t \quad (1.6)$$

where $w_t \sim iid N(0, \Sigma_w)$ with $\Sigma_w = \begin{bmatrix} \sigma_w^2 & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & 0 \end{bmatrix}$ is a shock to fundamental.

Measurement: $y_{it} = H z_t + v_{it}$ where $H = [1 \ 0 \ \dots \ 0]$ and $v_{it} \sim iid N(0, \Sigma_v)$ is the agent specific shock which is uncorrelated across agents. Without loss of generality we omit signals contaminated with noise shocks common across agents. We assume that $E(w_t v'_{it}) = 0$, that is shocks to fundamentals w_t and measurement error shocks v_{it} are independent.

Denote the one step-ahead forecast error for the forecast $z_{t|t-1}(i) \equiv E(z_t | y_{i,t-1}, y_{i,t-2}, \dots)$ in the Kalman filter with

$$\Psi \equiv \Sigma_z(t|t-1) \equiv E\left((z_{t|t-1}(i) - z_t)(z_{t|t-1}(i) - z_t)' \right).$$

We can find Ψ from the Ricatti equation:

$$\Psi = B\{\Psi - \Psi H'(H\Psi H' + \Sigma_v)^{-1} H\Psi\}B' + \Sigma_w.$$

Denote the gain of the Kalman filter with

$$P = \Psi H'(H\Psi H' + \Sigma_v)^{-1}.$$

The forecast for the unobserved state z_t evolves as follows:

$$\begin{aligned}
z_{t|t}(i) &= z_{t|t-1}(i) + P(y_t - y_{t|t-1}(i)) = z_{t|t-1}(i) + P(Hz_t - Hz_{t|t-1}(i)) \\
&= Bz_{t-1|t-1}(i) + P(HBz_{t-1} + Hw_t + v_{it} - HBz_{t-1|t-1}(i)) \\
&= (I - PH)Bz_{t-1|t-1}(i) + PHBz_{t-1} + PHw_t + Pv_{it}
\end{aligned}$$

where $y_{t|t-1}(i) = Hz_{t|t-1}(i)$. Consequently, the forecast for y_{it} evolves as

$$y_{t|t}(i) = Hz_{t|t}(i) = H(I - PH)Bz_{t-1|t-1}(i) + HPHBz_{t-1} + HPHw_t + HPv_{it}. \quad (1.7)$$

Note that $H\Psi H'$ is a scalar and $HP = H\Psi H'(H\Psi H' + \Sigma_v)^{-1} \in [0,1]$. More specifically, $HP = P_1$ which shows how strongly an innovation in w_t is translated into a revised estimate of the state and in this sense $1 - P_1$ measures the degree of informational rigidity. The mean forecast evolves as

$$\overline{y_{t|t}} = H(I - PH)B\overline{z_{t-1|t-1}} + HPHBz_{t-1} + HPHw_t. \quad (1.8)$$

The forecast error for the estimate of the inflation rate is

$$FE_{t,t} = y_t - \overline{y_{t|t}} = H(I - PH)B\{z_{t-1} - \overline{z_{t-1|t-1}}\} + (I - HP)Hw_t. \quad (1.9)$$

Property #1 follows from (1.6), (1.8), (1.9), and $HP \in [0,1]$, $(I - HP)H \in [0,1]$.

Note that we can compute the mean forecast error for z_t as follows:

$$z_t - \overline{z_{t|t}} = (I - PH)B(z_{t-1} - \overline{z_{t-1|t-1}}) + (I - PH)w_t = D(z_{t-1} - \overline{z_{t-1|t-1}}) + (I - PH)w_t \quad (1.10)$$

where $D \equiv (I - PH)B$. **Property #2** follows from Anderson and Moore (1979, Chapter 4). They show that if π_t is covariance stationary, eigenvalues of $(I - PH)B$ are less than one in absolute value and therefore $(z_t - \overline{z_{t|t}})$ shrinks to zero asymptotically.

Property #3 follows trivially from (1.7) since the only source of disagreement is v_{it} which does not depend on w_t .

Property #4 takes a more complex form. As we noted in the text, there are two useful ways to measure and interpret information rigidity in the context of noisy information models: i) fraction of the signal contemporaneously incorporated in the estimate of the current unobserved fundamental; ii) persistence of forecast errors after controlling for persistence of the fundamental. In the AR(1) model, these two interpretations coincide. In AR(p) models, these two metrics may differ mainly because the dynamics of the fundamental is more complex and there is no universal measure of persistence of the fundamental in this context. For example, one may use the half-life, the size of the largest eigenvalues of the companion matrix, the sum of AR(p) coefficients, etc. and the ranking of processes in terms of persistence can vary with the choice of the measure. Consequently, one needs to take a stand on what metric he or she uses.

For the first interpretation, note that (1.10) is a VAR(1) process. When B is estimated from (1.6), one can recover P_1 which captures how much an estimate of the unobserved fundamental is revised in light of observed signals by post-multiplying an estimate of $(I - PH)B$ with the inverse of an estimate of B .

For the second interpretation in the context of AR(1), we divided the impulse response of mean forecast errors by the impulse response of the forecasted series. For the AR(1) case, the result at horizon h was $(1 - P)^{h+1}$. Note that this result can be also obtained by raising the eigenvalue of the companion matrix for forecast errors (which is equal to $(1 - P)\rho$) to power h then divided by the eigenvalue of the companion matrix for the forecasted series (which is equal to ρ) raised to power h and then multiplied by the size of the contemporaneous forecast error (which is equal to $(1 - P)$ after a unit shock in the forecasted series). In other words, $\frac{\{(1-P)\rho\}^h}{\rho^h}(1 - P)$. Because the dynamics of forecast errors and actual series is described by one eigenvalue for each series, one has a simple metric of persistence of forecast errors after controlling for persistence of forecasted series.

As suggested above, in the case of AR(p) one does not have such simple metric because the dynamics of forecast errors as well as forecasted series is governed by multiple eigenvalues. However, for empirically plausible parameterizations of AR(p), one eigenvalue dominates other eigenvalues and thus impulse responses are likely to be governed by the dominant eigenvalue after a few periods after a shock. In this case, the second metric of information rigidity based on the division of one impulse response by another can provide a close analogue to the metric obtained in the AR(1) case.¹

As an illustration, we estimate an AR(4) process for inflation using the U.S. GDP deflator and then vary the noise/signal ratio to see how an estimate of the informational rigidity based on our approximation performs relative

¹ The quality of this approximation can be improved by dropping the first k periods in the impulse response.

to the ratio of the largest eigenvalue for forecast errors (from matrix D) divided by the largest eigenvalue for the forecasted series (from matrix B). Appendix Figure 1 shows that the approximation performs well: the dynamics of the ratio of the impulse responses is closely approximated by dynamics governed by the ratio of the largest eigenvalues for matrixes D and B .

3. Heterogeneous precision of signals.

In this section we derive the dynamics of forecast errors and disagreement in the model where agents have heterogeneous precision of signals. Since the dynamics is highly nonlinear, we will use approximations to characterize the properties of the dynamics and we will assume that shocks are small. To simplify the argument, we take the baseline noisy information model and assume that i) $\rho = 1$; ii) there is no common signal s_t ; iii) the precision of signals v_{it} (i.e., σ_v^{-2}) varies across agents in such a way that the gain of the Kalman filter across agents is approximately distributed as $P_i \sim (P, \sigma_p^2)$ which is independent from v_{it} and w_t . Hence, the inflation forecast for agent i is given by

$$\pi_{t|t}(i) = (1 - P_i)\pi_{t-1|t-1}(i) + P_i\pi_t + P_iv_{it} = P_i \sum_{k=0}^{\infty} (1 - P_i)^k (\pi_{t-k} + v_{i,t-k}).$$

The dynamics of the average forecast in the economy is:

$$\begin{aligned} \overline{\pi_{t|t}} &\equiv \bar{E} \left(\pi_{t|t}(i) \right) = \bar{E} \left\{ P_i \sum_{k=0}^{\infty} (1 - P_i)^k (\pi_{t-k} + v_{i,t-k}) \right\} \\ &= \bar{E} \left\{ P_i \sum_{k=0}^{\infty} (1 - P_i)^k \pi_{t-k} \right\} + \bar{E} \left\{ P_i \sum_{k=0}^{\infty} (1 - P_i)^k (v_{i,t-k}) \right\} = \bar{E} \left\{ \sum_{k=0}^{\infty} P_i (1 - P_i)^k \pi_{t-k} \right\} \\ &= \bar{E} \left\{ \sum_{k=0}^{\infty} \left[\begin{array}{l} P(1 - P)^k + ((1 - P)^k - kP(1 - P)^{k-1})(P_i - P) \\ + \frac{1}{2} (k(k - 1)P(1 - P)^{k-2} - 2k(1 - P)^{k-1})(P_i - P)^2 + h. o. t. \end{array} \right] \pi_{t-k} \right\} \\ &= \sum_{k=0}^{\infty} P(1 - P)^k \pi_{t-k} + \left\{ \sum_{k=0}^{\infty} ((1 - P)^k + kP(1 - P)^{k-1}) \pi_{t-k} \bar{E}[(P_i - P)] \right\} \\ &\quad + \frac{1}{2} \left\{ \sum_{k=0}^{\infty} [(k(k - 1)P(1 - P)^{k-2} - 2k(1 - P)^{k-1}) E_i(P_i - P)^2] \pi_{t-k} \right\} + h. o. t. \\ &= \sum_{k=0}^{\infty} P(1 - P)^k \pi_{t-k} + \frac{1}{2} \sigma_p^2 \left\{ \sum_{k=0}^{\infty} [(k(k - 1)P(1 - P)^{k-2} - 2k(1 - P)^{k-1})] \pi_{t-k} \right\} + h. o. t. \end{aligned}$$

We can also write the dynamics of the average forecast for inflation as follows:

$$\begin{aligned} \overline{\pi_{t|t}} &= \bar{E} \left(\pi_{t|t}(i) \right) = \bar{E} \left((1 - P_i)\pi_{t-1|t-1}(i) + P_i\pi_t + P_iv_{it} \right) = \bar{E} \left((1 - P_i)\pi_{t-1|t-1}(i) \right) + P\pi_t \\ &= \bar{E} \left((1 - P_i)(\pi_{t-1|t-1}(i) - \overline{\pi_{t-1|t-1}}) + \overline{\pi_{t-1|t-1}} \right) + P\pi_t \\ &= \bar{E} \left((1 - P_i)(\pi_{t-1|t-1}(i) - \overline{\pi_{t-1|t-1}}) \right) + (1 - P)\overline{\pi_{t-1|t-1}} + P\pi_t. \end{aligned}$$

The hardest term in this expression is $\bar{E} \left((1 - P_i)(\pi_{t-1|t-1}(i) - \overline{\pi_{t-1|t-1}}) \right)$ which has a complex dynamic structure.

Specifically, we can find that

$$\begin{aligned} &\bar{E} \left((1 - P_i)(\pi_{t-1|t-1}(i) - \overline{\pi_{t-1|t-1}}) \right) = \\ &= \bar{E} \left((1 - P_i)(P_i \sum_{k=0}^{\infty} (1 - P_i)^k (\pi_{t-k-1} + v_{i,t-k-1}) - \overline{\pi_{t-1|t-1}}) \right) \\ &= \bar{E} \left((1 - P_i)(\sum_{k=0}^{\infty} P_i (1 - P_i)^k \pi_{t-k-1} - \overline{\pi_{t-1|t-1}}) \right) \\ &= \bar{E} \left((1 - P_i) \left(\left\{ \sum_{k=0}^{\infty} \left[\begin{array}{l} P(1 - P)^k + ((1 - P)^k - kP(1 - P)^{k-1})(P_i - P) \\ + \frac{1}{2} (k(k - 1)P(1 - P)^{k-2} - 2k(1 - P)^{k-1})(P_i - P)^2 + h. o. t. \end{array} \right] \pi_{t-k-1} \right\} - \overline{\pi_{t-1|t-1}} \right) \right) \\ &= \bar{E} \left((1 - P_i) \left(\left\{ \sum_{k=0}^{\infty} \left[\begin{array}{l} P(1 - P)^k + ((1 - P)^k - kP(1 - P)^{k-1})(P_i - P) \\ + \frac{1}{2} (k(k - 1)P(1 - P)^{k-2} - 2k(1 - P)^{k-1})(P_i - P)^2 + h. o. t. \end{array} \right] \pi_{t-k-1} \right\} - \overline{\pi_{t-1|t-1}} \right) \right) = \end{aligned}$$

$$\begin{aligned}
&= \bar{E} \left((1 - P_i) \left(\left\{ \sum_{k=0}^{\infty} \left[+\frac{1}{2} (k(k-1)P(1-P)^{k-2} - 2k(1-P)^{k-1})((1-P_i) - (1-P))^2 + h.o.t. \right] \pi_{t-k-1} \right\} - \overline{\pi_{t-1|t-1}} \right) \right) \\
&= \bar{E} \left((1 - P_i) \left(\left\{ \sum_{k=0}^{\infty} \left[+\frac{1}{2} (k(k-1)P(1-P)^{k-2} - 2k(1-P)^{k-1})((1-P_i) - (1-P))^2 + h.o.t. \right] \pi_{t-k-1} \right\} \right. \right. \\
&\quad \left. \left. - \left\{ \sum_{k=0}^{\infty} P(1-P)^k \pi_{t-k-1} + \frac{1}{2} \sigma_p^2 \left\{ \sum_{k=0}^{\infty} [(k(k-1)P(1-P)^{k-2} - 2k(1-P)^{k-1})] \pi_{t-k-1} \right\} \right\} \right) \right) \\
&= \bar{E} \left((1 - P_i) \left(\left\{ \sum_{k=0}^{\infty} \left[+\frac{1}{2} (k(k-1)P(1-P)^{k-2} - 2k(1-P)^{k-1})((1-P_i) - (1-P))^2 \right] \pi_{t-k-1} \right\} \right. \right. \\
&\quad \left. \left. - \left\{ \frac{1}{2} \sigma_p^2 \left\{ \sum_{k=0}^{\infty} [(k(k-1)P(1-P)^{k-2} - 2k(1-P)^{k-1})] \pi_{t-k-1} \right\} \right\} \right) \right) + h.o.t. \\
&= \bar{E} \left((1 - P_i) \left(\left\{ \sum_{k=0}^{\infty} \left[+\frac{1}{2} (k(k-1)P(1-P)^{k-2} - 2k(1-P)^{k-1})((1-P_i) - (1-P))^2 \right] \pi_{t-k-1} \right\} \right) \right) \\
&\quad - \bar{E}(1 - P_i) \left\{ \frac{1}{2} \sigma_p^2 \left\{ \sum_{k=0}^{\infty} [(k(k-1)P(1-P)^{k-2} - 2k(1-P)^{k-1})] \pi_{t-k-1} \right\} \right\} + h.o.t. \\
&= -\bar{E}[(1 - P_i) - (1 - P)]^2 \left\{ \sum_{k=0}^{\infty} [(1 - P)^k - kP(1 - P)^{k-1}] \pi_{t-k-1} \right\} \\
&\quad + \bar{E}[(1 - P_i) - (1 - P)]^3 \left\{ \sum_{k=0}^{\infty} \left[\frac{1}{2} (k(k-1)P(1-P)^{k-2} - 2k(1-P)^{k-1}) \right] \pi_{t-k-1} \right\} \\
&\quad + (1 - P) \left\{ \frac{1}{2} \sigma_p^2 \left\{ \sum_{k=0}^{\infty} [(k(k-1)P(1-P)^{k-2} - 2k(1-P)^{k-1})] \pi_{t-k-1} \right\} \right\} \\
&\quad - (1 - P) \left\{ \frac{1}{2} \sigma_p^2 \left\{ \sum_{k=0}^{\infty} [(k(k-1)P(1-P)^{k-2} - 2k(1-P)^{k-1})] \pi_{t-k-1} \right\} \right\} + h.o.t. \\
&= -\bar{E}[(P_i - P)^2] \left\{ \sum_{k=0}^{\infty} [(1 - P)^k - kP(1 - P)^{k-1}] \pi_{t-k-1} \right\} + h.o.t. \\
&= -\sigma_p^2 \left\{ \sum_{k=0}^{\infty} [(1 - P)^k - kP(1 - P)^{k-1}] \pi_{t-k-1} \right\} + h.o.t. \\
&= -\sigma_p^2 \left\{ \sum_{k=0}^{\infty} [(1 - P)^k - kP(1 - P)^{k-1}] \pi_{t-k-1} \right\} + h.o.t. \\
&= -\sigma_p^2 \left\{ \sum_{k=0}^{\infty} [(1 - P)^{k-1} (1 - P - kP)] \pi_{t-k-1} \right\} + h.o.t. \\
&= -\sigma_p^2 \left\{ \sum_{k=0}^{\infty} A_k(P) \pi_{t-k-1} \right\} + h.o.t.
\end{aligned}$$

where $A_k(P) = (1 - P)^{k-1}(1 - P - kP)$. In short,

$$\overline{\pi_{t|t}} = -\sigma_p^2 \left\{ \sum_{k=0}^{\infty} A_k(P) \pi_{t-k-1} \right\} + (1 - P) \overline{\pi_{t-1|t-1}} + P \pi_t.$$

Now we are ready to find the dynamics of the mean forecast error:

$$\begin{aligned}
FE_t &\equiv \pi_t - \overline{\pi_{t|t}} = \pi_t - \left\{ \sum_{k=0}^{\infty} A_k(P) \pi_{t-k-1} \right\} + (1 - P) \overline{\pi_{t-1|t-1}} + P \pi_t \\
&= (1 - P) (\pi_{t-1} - \overline{\pi_{t-1|t-1}}) + (1 - P) w_t - \left\{ \sum_{k=0}^{\infty} A_k(P) \pi_{t-k-1} \right\} \\
&= (1 - P) FE_{t-1} + (1 - P) w_t + \sigma_p^2 \left\{ \sum_{k=0}^{\infty} A_k(P) \pi_{t-k-1} \right\}. \tag{1.11}
\end{aligned}$$

Therefore, if we run regression $FE_t = (1 - P)FE_{t-1} + error$ we should find that the error should be correlated with lags of inflation.

The cross-section dispersion of forecasts is then

$$\begin{aligned}
\bar{E}(\pi_{t|t}(i) - \overline{\pi_{t|t}})^2 &= \bar{E}(\left\{ \sum_{k=0}^{\infty} P_i(1 - P_i)^k (\pi_{t-k} + v_{i,t-k}) \right\} - \overline{\pi_{t|t}})^2 \\
&= \bar{E}(\left\{ \sum_{k=0}^{\infty} P_i(1 - P_i)^k \pi_{t-k} \right\} - \overline{\pi_{t|t}})^2 + \bar{E}(\left\{ \sum_{k=0}^{\infty} P_i(1 - P_i)^k v_{i,t-k} \right\})^2 \\
&= \bar{E}(\left\{ \sum_{k=0}^{\infty} P_i(1 - P_i)^k \pi_{t-k} \right\} - \overline{\pi_{t|t}})^2 + \sigma_v^2 \sum_{k=0}^{\infty} \bar{E}\{(P_i(1 - P_i)^k)^2\} \\
&= \bar{E} \left(\left(\left\{ \sum_{k=0}^{\infty} \left[+\frac{1}{2} (k(k-1)P(1-P)^{k-2} - 2k(1-P)^{k-1})(P_i - P)^2 + h.o.t. \right] \pi_{t-k} \right\} \right. \right. \\
&\quad \left. \left. - \left\{ \sum_{k=0}^{\infty} P(1 - P)^k \pi_{t-k} + \frac{1}{2} \sigma_p^2 \left\{ \sum_{k=0}^{\infty} [(k(k-1)P(1-P)^{k-2} - 2k(1-P)^{k-1})] \pi_{t-k} \right\} + h.o.t. \right\} \right)^2 \right) + const =
\end{aligned}$$

$$\begin{aligned}
&= \bar{E} \left(\left(\left\{ \sum_{k=0}^{\infty} \left[\begin{array}{c} ((1-P)^k - kP(1-P)^{k-1})(P_i - P) \\ + \frac{1}{2}(k(k-1)P(1-P)^{k-2} - 2k(1-P)^{k-1})(P_i - P)^2 \\ + h.o.t. \end{array} \right] \pi_{t-k} \right\} \right)^2 + const \right. \\
&\quad \left. - \left\{ \frac{1}{2} \sigma_p^2 \left[\sum_{k=0}^{\infty} [(k(k-1)P(1-P)^{k-2} - 2k(1-P)^{k-1})] \pi_{t-k} \right] + h.o.t. \right\} \right)^2 \\
&= \bar{E} \left(\left(\sum_{k=0}^{\infty} [(1-P)^k - kP(1-P)^{k-1}](P_i - P) \pi_{t-k} + h.o.t. \right)^2 + const \right) = \\
&= \bar{E} \left((P_i - P) \left\{ \sum_{k=0}^{\infty} [(1-P)^k - kP(1-P)^{k-1}] \pi_{t-k} \right\} + h.o.t. \right)^2 + const = \\
&= \bar{E} \left((P_i - P) \left\{ \sum_{k=0}^{\infty} [(1-P)^k - kP(1-P)^{k-1}] \pi_{t-k} \right\} + h.o.t. \right)^2 + const.
\end{aligned}$$

Let

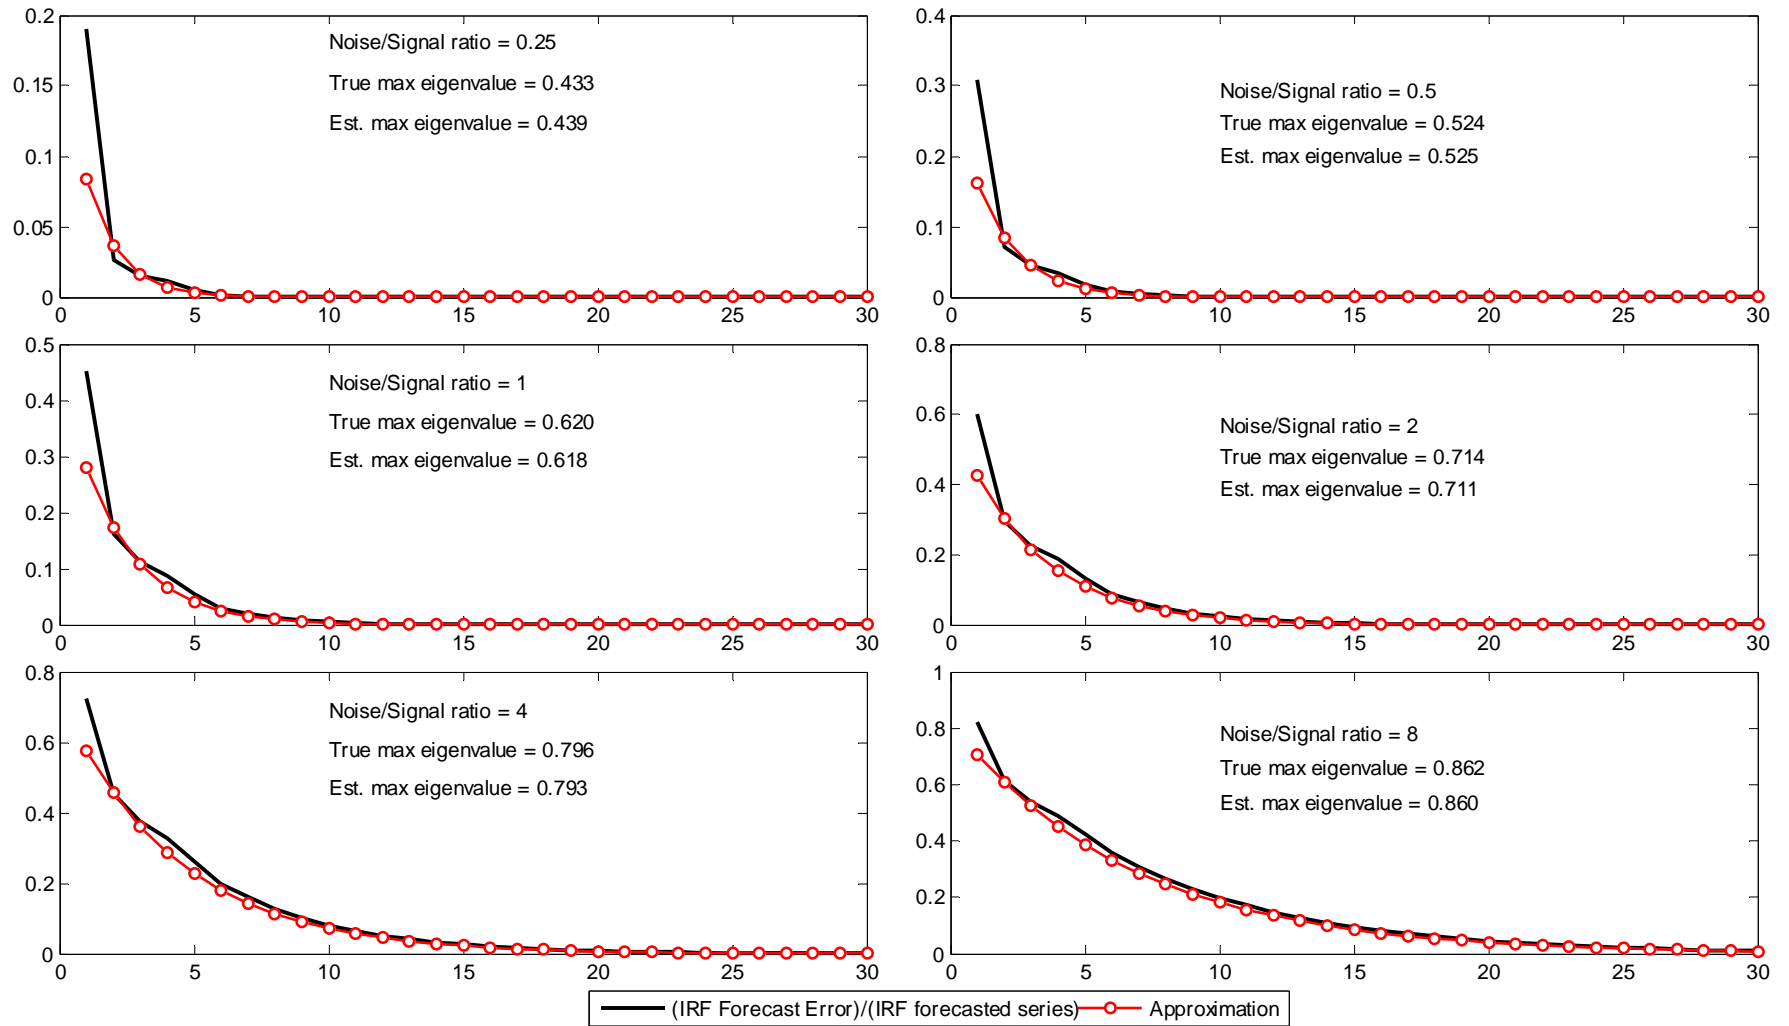
$$Q_t \equiv \left\{ \sum_{k=0}^{\infty} [(1-P)^k - kP(1-P)^{k-1}] \pi_{t-k} \right\} = \sum_{k=0}^{\infty} A_k(P) \pi_{t-k}$$

so that

$$\bar{E}(\pi_{t|t}(i) - \bar{\pi}_{i|t})^2 \approx \bar{E}((P_i - P)Q_t)^2 + const = Q_t^2 \sigma_p^2 + const. \quad (1.12)$$

Since Q_t varies over time, the disagreement is also time varying.

Appendix Figure 1. Dynamics of the ratio of impulse response of the mean forecast error to the impulse response of the forecasted variable.



Appendix B: Monte Carlo Experiments

In this appendix, we examine the ability of our empirical strategy to recover the sticky-information data generating process. Specifically, we develop a Monte Carlo simulation which replicates our two primary tests—the conditional responses of forecast errors and forecast dispersions—in response to shocks which differ in their quantitative magnitudes.

In each simulation, we let inflation follow an AR(1) process $\pi_t = \rho\pi_{t-1} + w_t$ where the shock ε_t is the sum of two independent innovations: $w_t = w_t^{(1)} + w_t^{(2)}$. Agents form expectations rationally but update them infrequently following a Poisson process in which $1 - \lambda$ is the probability of updating their information set each period, as in the sticky information model of Mankiw and Reis (2002). Thus, the mean forecast and the forecast dispersion of next-period ($h=1$) inflation at time t follow

$$\overline{\pi_{t+h|t}} = \sum_{k=0}^{\infty} \rho^{k+h} w_{t-k} (1 - \lambda^{k+1}),$$

$$V_t \pi_{t+h|t}(i) = (1 - \lambda) \sum_{k=0}^{\infty} \lambda^k (E_{t-k} \pi_{t+h} - \overline{\pi_{t,t+h}})^2.$$

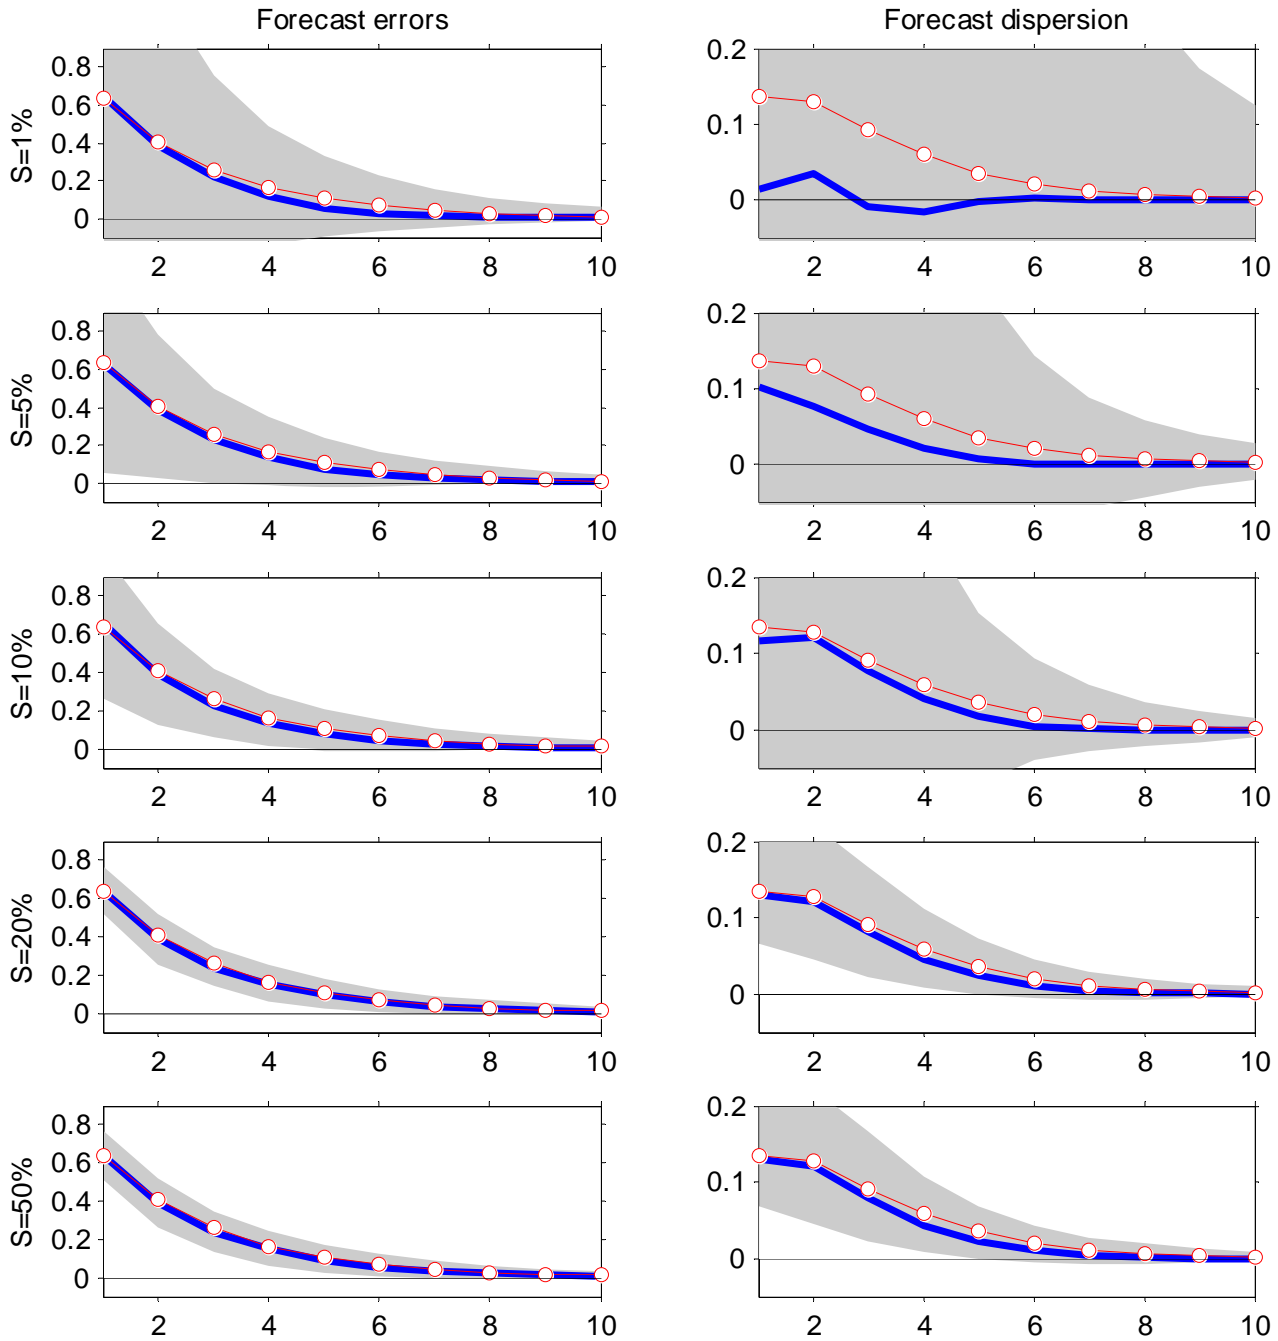
We simulate this model 10,000 times, with each simulation having 1,150 times periods. We set $\rho = 0.85$ and the variance of the total shock to inflation to be $\sigma_w^2 = 1.005$ to match estimates of an AR(1) process for GDP deflator inflation from 1979Q1 to 2007Q3.² Following Mankiw and Reis (2002), we set the degree of information rigidity λ to 0.75, implying that agents update their information once a year on average. In each simulation, we estimate the response of forecast errors to innovations $w_t^{(1)}$, as done in the paper, using the final $T=150$ periods of the simulation. The key parameter that varies across simulations is the fraction of the inflation variance accounted for by the innovation used to derive impulse responses ($w_t^{(1)}$). We consider five values: $S=1\%$, 5% , 10% , 20% and 50% .

Results of the Monte Carlo exercise are presented in Appendix Figure 2. We find that the mean estimated response for forecast errors is close to the response predicted by the model for all choices of S . For the response of forecast dispersion, small values of S ($S < 5\%$) tend to yield a fairly poor match of the theoretical response. When we increase the sample size T , the discrepancy between the mean estimated and model-predicted responses for forecast dispersion vanish. Hence, our approach can asymptotically recover the true responses of the data generating process. However, if structural innovations account for a small fraction of the inflation variance (below 10%), precisely estimating the conditional response of forecast dispersion could be too demanding on the data in short samples. Hence, we put more weight on oil, news and technology shocks and also investigate the behavior of expectations in response to unidentified shocks. We also find that the coverage rates for confidence intervals are close to nominal sizes when $S > 5\%$.

The key message of these Monte Carlo experiments is that in small samples our approach can consistently and precisely estimate the responses of forecast errors and forecast dispersion to quantitatively important shocks and it is less successful when shocks explain only a small fraction (below 5-10%) of variation.

² Results are similar when we use alternative ARMA models for the inflation.


Appendix Figure 2. Monte Carlo Simulations of Conditional Forecast Error and Conditional Forecast Dispersion Responses to Innovations.



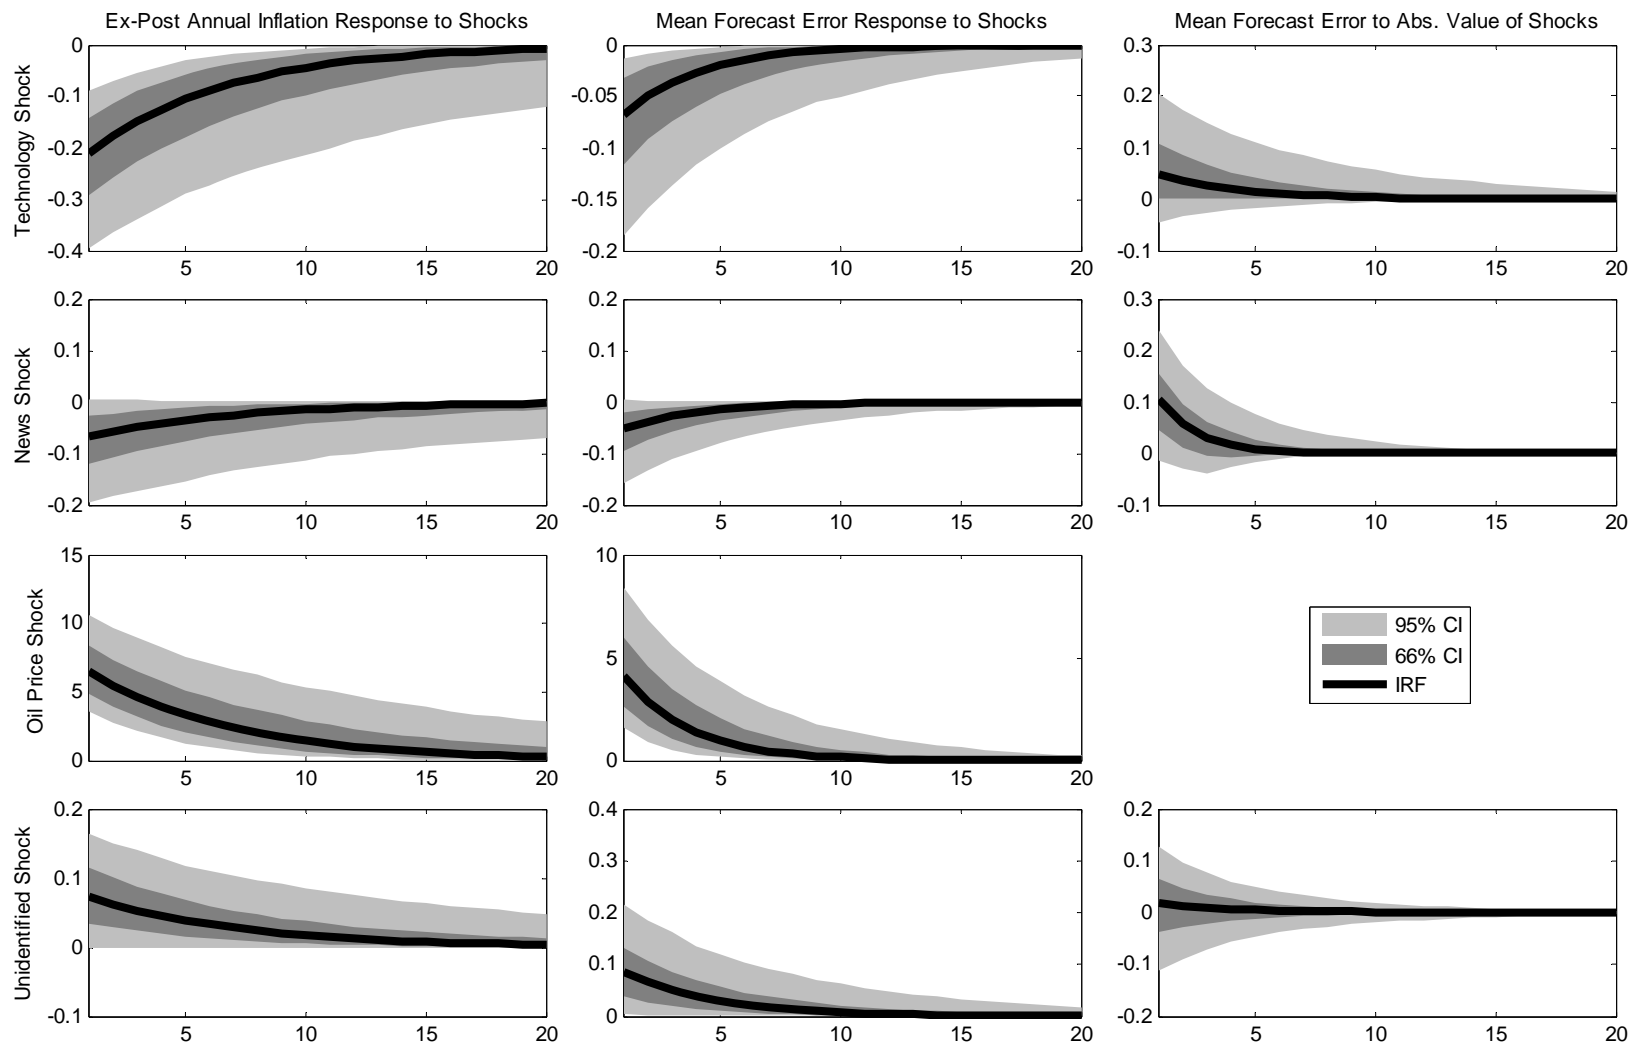
Notes: Thin solid red line with circles is the theoretically predicted response in the sticky information model. Thick solid blue line is the mean estimated response in the Monte Carlo simulations. Shaded region is the 90% distribution of the simulated responses. S is the fraction of the inflation variance accounted for by the innovation used in the estimation. Each experiment has 10,000 simulations.

Appendix C: Time series and impulse response for inflationary expectations extracted from asset prices

Appendix Figure 3. Time series of actual inflation and inflationary expectations extracted from asset prices.




Appendix Figure 4. Results for Financial Market Forecasts of Haubrich et al. (2011).



Note: The figure reports impulse responses to a unit shock computed from estimated specification (34) (column 1), (36) (column 2), and (37) (column 3). Each row shows responses to a given structural shock. Standard errors for impulse responses are computed using parametric bootstrap.

Appendix Figure 5. Results for Financial Market Forecasts of Ang et al. (2008).



Note: The figure reports impulse responses to a unit shock computed from estimated specification (34) (column 1), (36) (column 2), and (37) (column 3). Each row shows responses to a given structural shock. Standard errors for impulse responses are computed using parametric bootstrap.