What can survey forecasts tell us about informational rigidities?

Olivier Coibion
College of William & Mary

Yuriy Gorodnichenko
UC Berkeley and NBER

This Draft: October 25th, 2008

Abstract: This paper uses three different surveys of economic forecasts to assess both the support for and the properties of informational rigidities faced by agents. Specifically, we track the impulse responses of mean forecast errors and disagreement among agents after exogenous structural shocks. Our key contribution is to document that in response to structural shocks, mean forecasts fail to completely adjust on impact, leading to statistically and economically significant deviations from the null of full information: the half life of forecast errors is roughly between 6 months and a year. Importantly, the dynamic process followed by forecast errors following structural shocks is consistent with the predictions of models of informational rigidities. We interpret this finding as providing support for the recent expansion of research into models of informational rigidities. In addition, we document several stylized facts about the conditional responses of forecast errors and disagreement among agents that can be used to differentiate between some of the models of informational rigidities recently proposed.

Keywords: expectations, information rigidity, survey forecasts.

JEL: E3, E4, E5

We are grateful to Alex Cukierman, Ed Knotek, Oleg Korenok, Maxym Kryshko, Peter Morrow, David Romer, and Mirko Wiederholt as well as seminar participants at the Kansas City Federal Reserve Bank 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 economics 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 (2003) and Sims (2002), 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, the observed delayed response of inflation to monetary policy shocks is not readily matched by New Keynesian models without the addition of informational rigidities or the counterfactual assumption of price indexation.[1] Similarly, the differential response of inflation to monetary policy and technology shocks is difficult to reconcile without informational rigidities.[2] In addition, departing from the assumption of full-information can also account for some empirical puzzles. For example, 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.[3] 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 Bachetta et al (2008) all identify links between systematic forecast errors in survey forecasts and puzzles in various financial markets.[4]

Despite this resurgent focus on the nature of the expectations formation process, little consensus exists on how best to model the acquisition and processing of information by agents. In large part, this reflects a lack of convincing empirical evidence on the matter. First, the evidence against the assumption of full-information rational expectations is sparse, often fragile, and most importantly of unclear

[1] See Mankiw and Reis (2002).
[2] See Dupor et al (2007b).
[3] See Roberts (1997, 1998) and Adam and Padula (2003).
[4] See also Kim and Orphanides (2005).

economic significance. Second, there is even less empirical evidence available to distinguish among competing models of informational rigidities, so macroeconomists seeking to include informational frictions in their models have a multitude of options to consider but little basis upon which to choose a particular specification.

The most traditional approach to testing expectations data relies on testing the null of full information using survey forecasts, typically of inflation, by regressing ex-post inflation on mean inflation forecasts. As documented in Pesaran and Weale (2006), this literature has yielded only mixed results, typically finding departures from full information over short time samples but not over longer periods. In addition, this approach has little insight to offer about the dynamic properties of the expectations formation process beyond a test of the null of full information.[5] A second strand of empirical evidence has focused on evaluating the *implications* of models with informational rigidities relative to those without.[6] These results have also been mixed and are in any case difficult to rely on for making general conclusions about the expectations formation process since this type of evidence is sensitive to the estimation approach as well as to auxiliary modeling assumptions. A third approach documents the costs of collecting and processing information (e.g. Zbaracki et al (2004)). This approach is instrumental in establishing the fact the economic agents face significant informational constraints but cannot address how informational constraints affect agents' choices and aggregate dynamics. Finally, the fourth approach studies the properties of disagreement among agents to make inferences about their underlying approaches to expectations formation. For example, Mankiw Reis and Wolfers (2004) note that disagreement among agents is inconsistent with full information but not with models in which agents face informational frictions.[7] They argue that a sticky-information model can reproduce many of the features of expectations data they consider. Branch (2007) similarly makes use of data on disagreement to compare the fit of the sticky information model and a model in which agents may endogenously use heterogeneous forecasting rules.

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 make use of three data sets containing survey measures of inflation forecasts: the Michigan

[5] See Pesaran and Weale (2006) and Mankiw, Reis, and Wolfers (2003) for surveys of this literature.
[6] See Korenok (2005), Andres et al (2005), Kiley (2007), Coibion (2007), and Dupor et al (2007a).
[7] It is possible to have forecast dispersion when fully rational agents share the *same* information set but use different forecasting models or have different objective functions. However, Jonung (1981) and others find that agents disagree about current and *past* inflation and perceptions of past inflation are a strong predictor of inflationary expectations so that heterogeneity in expectations is driven to a significant extent by perceptions about current and past conditions. Hence, a bulk of disagreement seems to arise from differential information sets. Probably more importantly, there is potentially no discipline on what rules agents could use and consequently models with heterogeneous forecasting rules can rationalize any outcome. Thus, we do not formally try to isolate the contribution of heterogeneous forecasting rules to disagreement because we could not formulate a set of falsifiable hypotheses to rule out this theory.

Survey of Consumers, the Survey of Professional Forecasters, and the Blue Chip Economic Indicators. Each survey provides mean forecasts of inflation over the next year, as well as measures of the cross-sectional dispersion of agents' forecasts of future inflation. Unlike the previous literature, we study the *conditional* responses of forecast errors and forecast dispersion to identified structural shocks. With this novel approach we can hope to have more power in distinguishing hypotheses of how expectations are formed. As we discuss later, different models of expectation formation deliver sign restrictions on first and second moments of impulse response functions that can be used to assess the validity of these models. In contrast, these models tend to agree on implications for unconditional moments of expectations (e.g., positive dispersion of forecasts). In this respect, our approach is similar to standard methods of applied macroeconomics that utilize conditional responses of variables to shocks (i.e., impulse response functions) to study and estimate the behavior of empirical models.

Our first stylized fact is that *forecasts fail to adjust one-for-one with the variable being forecasted after structural shocks.* Instead, we find systematic patterns of serially correlated conditional forecast errors, particularly after technology and oil price shocks.[8] Thus, after inflationary (deflationary) shocks, one observes a predictable sequence of serially correlated positive (negative) inflation forecast errors.[9] Over time, forecast errors monotonically converge to zero. This result not only contradicts the null of full-information but does so in exactly the manner predicted by standard models of informational rigidities. In addition, we find that these deviations from full-information are not only statistically significant but also of economically significant magnitudes. For example, the half-lives of conditional forecast errors are between six months and a year. Such a delayed adjustment of average expectations in response to structural shocks is large enough to have important repercussions for macroeconomic dynamics. We interpret these results as demonstrating that informational rigidities are an important component of the expectations formation process for both consumers and professional forecasters. Given the crucial role played by forward-looking behavior in macroeconomics, this has broad implications for understanding and modeling business cycles.

Having documented this novel evidence for informational rigidities in the data, we then turn to the question of how to differentiate between competing models of expectations formation. To do so, we first present contrasting implications of two general types of informational rigidities. The first is the

---

[8] Our baseline structural shocks are monetary policy, technology, and oil price shocks. These are the shocks that explain the largest component of the variance of inflation out of the ones we consider, between thirty and fifty percent jointly depending on the time horizon, which is important for our ability to distinguish models of expectation formation in the data. We also consider other exogenous shock measures, such as fiscal shocks, information shocks, and alternative measures of monetary policy shocks and find similar qualitative (but much less precise) results

[9] We show that similar results hold for forecasts of unemployment after the information shocks of Beaudry and Portier (2006).

sticky information model of Mankiw and Reis (2002). In this model, agents can update their information only infrequently, but when they do so they acquire complete information about the current and past states.[10] Thus, at any moment in time, this model implies that a fraction of agents will be relying on outdated information. This generates a distribution of information sets across agents based on the last date at which agents acquired new information. An appealing feature of this form of rigidity is that a single form of rigidity can help explain inertia in different macroeconomic variables. These sticky-information models are also able to deliver inertial inflation in response to monetary policy shocks and rapid adjustment of inflation to technology shocks (Mankiw and Reis (2007) and Reis (2008)).

In the second class of models of informational rigidities, which we group under the heading of imperfect information, agents cannot observe the current state perfectly and must thus form a belief about the current state based on the variables that they observe.[11] For example, Woodford (2001) considers a model in which agents observe a noisy signal about the current state. These agents use a Kalman filter to *continuously* update their beliefs about the current state in the face of new signals.[12] Sims (2003) argues that agents face limited information processing capacities, and must thus endogenously allocate their attention to different variables.[13] Applying Sims' approach to price-setting decisions, Paciello (2007) and Mackoviak and Wiederholt (2008a) show that under rational inattention, firms will pay more attention to technology shocks than monetary policy shocks because the former are more volatile and have greater effect on profits. As a result, firms change prices more rapidly after a technology shock than a monetary policy shock.

These two classes of models share certain important characteristics. First, both models consist of rational expectations agents subject to a specific source of informational rigidity, unlike earlier ad hoc models with adaptive expectations. Second, each model implies that after a shock the average forecast across agents will trail in terms of magnitudes and timing the variable being forecasted, but in each case forecast errors will converge to zero over time. Third, both approaches are consistent with disagreement among agents, i.e. agents have different beliefs about the current and future states. Finally, both are consistent with the differential response of inflation to technology and monetary policy shocks, making them particularly appealing approaches to consider.

[10] The microfoundations of sticky information are developed in Reis (2006a) for firms and Reis (2006b) for consumers. He shows that when agents face fixed costs to acquiring information, they will update their information sets infrequently and, under certain conditions, exactly in the way proposed in Mankiw and Reis (2002).
[11] Loosely speaking, if one thinks of sticky information as resulting from fixed costs to acquiring information, imperfect information models can be thought of as resulting from convex costs to acquiring information, leading to continuous but incremental information acquisition.
[12] This strand of literature goes back to Lucas (1972) which has an early version of a macroeconomic model with imperfect information.
[13] Sims (2003) illustrates how rationally inattentive agents behave much like imperfect agents but differ in that the signal to noise ratio is endogenous.

Despite these similarities, the two models also make conflicting predictions. For example, in sticky information models, agents update their information sets infrequently and independently of the type of shock hitting the economy. Thus, the convergence rate of forecast errors should be identical across shocks after controlling for the persistence of the inflation rate.[14] In imperfect information models, on the other hand, the convergence rate of forecast errors will depend on the quantitative importance of different shocks. For example, firms should devote more attention to gathering information about shocks that affect profits more. Thus, in general, shocks that play an important role in affecting business cycles should have a faster convergence rate of forecast errors than quantitatively unimportant shocks since agents should pay more attention to more important shocks. Secondly, in the sticky-information model, disagreement, as measured by the cross-sectional variance of forecasts, is predicted to increase after a shock. In the linear imperfect information model that we consider, there should be no increase in disagreement across agents after macroeconomic shocks. In summary, these models do offer crisp differential predictions that can be used to assess the empirical validity of the models.

Consequently, we study the differences in the conditional responses of percentage forecast errors across agents and shocks and uncover two new findings. Our second stylized fact is that *conditional forecast errors converge at similar rates across different shocks*. In other words, we find little evidence that forecasts adjust more or less rapidly to different shocks. In particular, mean inflation forecasts converge to actual inflation just as rapidly after monetary policy shocks as after technology shocks. While consistent with sticky information models, this result is harder to reconcile with imperfect information models, in which one would typically expect agents to pay more attention to quantitatively important shocks than others. Our third stylized fact is that *conditional forecast errors converge at similar rates across agents*. Thus, we find no evidence that professional forecasters' mean forecasts converge more rapidly to actual inflation than consumers' forecasts. This result is not necessarily inconsistent with either imperfect or sticky information models, but it is challenging for epidemiological models of information diffusion, as in Carroll (2003). This approach assumes that information spreads gradually from professional forecasters to consumers, a proposition that is inconsistent with our finding that mean professional forecasts converge no quicker to true values than consumer forecasts.

Finally, following Mankiw, Reis and Wolfers (2004) and Branch (2007), we study the cross-sectional dispersion of forecasts across agents. But unlike these authors, we focus on the conditional response of disagreement among agents to structural shocks. This delivers our fourth stylized fact: *structural shocks do not appear to lead to any discernible increase in disagreement*. This result is remarkably consistent with simple linear imperfect information models but not sticky information models.

[14] More precisely, the percentage forecast errors should converge at a common rate determined entirely by the rate of information updating.

The one exception that we identify is that oil price shocks tend to raise disagreement among agents, although this result is quite sensitive to the specification.

The structure of the paper is as follows. In section 2, we present two models of informational 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 present the survey measures that form the backbone of our empirical analysis as well as the shocks measures we utilize. Section 4 presents the baseline empirical results, along with a Monte Carlo exercise. Section 5 includes robustness checks of our results and an extension to unemployment forecasts. Section 6 concludes.

2 Two Models of Information Rigidity

2.1 Sticky Information

Reis (2006a) 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 informational 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.

Denoting the mean forecast across agents at time t of a variable π_{t+j} j periods ahead as $F_t\pi_{t+j}$, we have

$$F_t\pi_{t+j} = (1-\lambda)\sum_{k=0}^{\infty}\lambda^k E_{t-k}\pi_{t+j} \tag{1}$$

The mean forecast is a weighted average of past (rational) expectations of the variable at time $t+j$. Now suppose that the economy is initially in a steady-state, so that all expectations of π are equal to the steady-state value, normalized to zero for simplicity. At time $t=0$, a shock occurs and affects π in a deterministic way (i.e. yields a sequence of inflation rates $\{\pi_{t+j}\}$). There are no other shocks. The impulse response of the average forecast across agents follows

$$\frac{\partial F_t\pi_{t+j}}{\partial shock} = (1-\lambda^{t+1})\pi_{t+j} \quad \forall t \geq 0 \tag{2}$$

The mean forecast depends on the ex-post value of inflation, since when agents update their information sets, they acquire full information. Thus, after an inflationary shock, mean forecasts rise along with inflation. Note that the coefficient is converging to one over time, so mean forecasts converge to the true value. But because the coefficient is less than one, forecast errors will be non-zero and persistent. Defining the forecast error *j* periods ahead as $FE_{t,t+j} \equiv \pi_{t+j} - F_t\pi_{t+j}$, its impulse response is given by

$$\frac{\partial FE_{t,t+j}}{\partial shock} = \lambda^{t+1}\pi_{t+j} \quad \forall t \geq 0 \tag{3}$$

Forecast errors depend both on the inflation process after the shock, as well as the degree of informational rigidity. Note that when $\lambda = 0$, firms always update their information sets and the forecast error is always zero. As the degree of informational 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 percentage deviation of the mean forecast from the true value ($FE_{t,t+j} / \pi_{t+j}$) is monotonically decreasing over time at a rate governed the degree of informational 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 shocks must yield the same convergence rate for mean forecasts.

The sticky information model also makes predictions about the cross-sectional dispersion of beliefs across agents. Define $V_t\pi_{t+j}$ to be the cross-sectional variance of forecasts at time *t* of *π* in *j* periods. Then,

$$V_t\pi_{t+j} = (1-\lambda)\sum_{k=0}^{\infty}\lambda^k \left[E_{t-k}\pi_{t+j} - F_t\pi_{t+j} \right]^2 \tag{4}$$

In response to a shock at time $t = 0$ out of the steady-state, the impulse response of the cross-sectional variance of forecasts is given by

$$\frac{\partial V_t\pi_{t+j}}{\partial shock} = (1-\lambda^{t+1})\lambda^{t+1}\pi_{t+j}^2 \quad \forall t \geq 0 \tag{5}$$

As long as $\lambda > 0$ and $\pi_{t+j} > 0$, the dispersion, or degree of disagreement across agents, will rise in response to a shock. Over time (assuming inflation converges), the dispersion will return to its steady-state level.

We can summarize the predictions of the sticky-information model as follows:

1. Conditional forecast errors respond to shocks with the same sign as the predicted variable and converge to zero over time.
2. The convergence rate of percentage forecast errors is common across shocks.

3. The dispersion of forecasts should increase after a shock.

2.2 Imperfect Information

Lucas (1972), Cukierman and Wachtel (1979), Woodford (2001) and Sims (2003) develop models where economic agents filter the state of economic fundamentals from a series of signals contaminated with idiosyncratic, agent-specific noise. In contrast to Mankiw and Reis (2002), agents *continuously* track variables and incorporate most recent information into their decision making. The striking feature of this class of models is that the dispersion of forecasts does not vary in response to a shock in fundamentals. In this section, we present a simple model to illustrate the intuition behind this result.

Suppose that economic agents observe signals about inflation $\omega_{it} = \pi_t + v_{it}$ where π_t is the aggregate level of inflation and $v_{it} \sim iid\ N(0, \Sigma_v)$ is an agent specific shock. Also without loss of generality, suppose that inflation evolves as a random walk $\pi_t = \pi_{t-1} + w_t$ where $w_t \sim iid\ N(0, \Sigma_w)$.[15] Denote the optimal forecast for inflation at time *t* given agent *i*'s information at time *s* with $\pi_{i,t|s}$. Using properties of the Kalman filter, one can show that

$$\pi_{i,t|t} = \pi_{i,t|t-1} + G(\omega_{it} - \pi_{i,t|t-1}), \tag{6}$$

where $G \in (0,1)$ is the gain of the Kalman filter. Note that the gain of the filter does not vary across agents because all agents solve the same Ricatti equation $P = [P - P(P + \Sigma_w)^{-1} P] + \Sigma_v$, where *P* is the variance of the one-step ahead forecast $\pi_{i,t|t-1}$, and thus obtain the same gain $G = P(P + \Sigma_v)^{-1}$.

Note that the average forecast for the current state of inflation given current information is equal to

$$\begin{aligned} F_t\pi_{i,t|t} &\equiv E_i\pi_{i,t|t} = (1-G)E_i\pi_{i,t|t-1} + GE_i\omega_{it} = (1-G)E_i\pi_{i,t|t-1} + G(E_i\pi_t + E_i v_{it}) \\ &= (1-G)E_i\pi_{i,t|t-1} + G\pi_t = (1-G)E_i\pi_{i,t-1|t-1} + G\pi_t \end{aligned} \tag{7}$$

where $E_i(\cdot)$ indicates that expectation is taken over agents rather than time and the last equality follows from $\pi_{i,t+j|t} = \pi_{i,t|t}$. Then the impulse of the average forecast of inflation to a shock in inflation is $\pi_{t+j|t+j} = G\sum_{k=0}^{j}(1-G)^k \pi_t = [1-(1-G)^j]\pi_t \quad \forall j \geq 0$ where we omit $(1-G)^{j+1} E_i \pi_{i,t|t-1}$ by conditioning on the initial state of beliefs. Since for each agent the forecast for inflation is $\pi_{i,t+j|t} = \pi_{i,t|t}$, the impulse of the average *j*-step ahead forecast has the same properties as the average forecast for the current state.

[15] This can readily be extended to non-random walk behavior. Our conclusions also do not change if we allow shocks $v_{i,t}$ to be correlated across agents.

Similar to the sticky information model, this model predict that the average forecast will have serially correlated ex-post error equal to $(1-G)\pi_t$. Because $G<1$, the forecast error moves in the same direction as the mean forecast does. Note that the forecast error converges to zero with time as $\pi_{t+j|t+j}-\pi_t=(1-G)^j\pi_t \to 0$ with $j\to\infty$.

Using equation (6), we can derive the law of motion for the dispersion of forecasts across agents:

$$V_t\pi_{i,t|t}=(1-G)^2V\pi_{i,t|t-1}+G^2V(\pi_t+v_{it})=(1-G)^2V\pi_{i,t|t-1}+G^2\Sigma_v=(1-G)^2V\pi_{i,t-1|t-1}+G^2\Sigma_v\,.$$

Note that $V\pi_{i,t|t}$ does not depend on π_t. This means that the motion of forecast dispersion does not vary with π_t. Again using the fact that *j*-step ahead forecast is based only on the forecast for the current state (here, $\pi_{i,t+j|t}=\pi_{i,t|t}$), we conclude that the dispersion of the *j*-step ahead forecasts does not vary with π_t. 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 dispersion does not respond to π_t.[16]

Note that the speed of reacting to signals ω_{it} increases in the volatility of fundamentals and decreases in the volatility of the idiosyncratic shock v_{it}. Likewise, 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 e.g. Mackoviak and Wiederholt (2008a)). Thus, unlike the sticky information model, the imperfect 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.

We can summarize the predictions of the imperfect information model as follows:

1. Conditional forecast errors respond to shocks with the same sign as the predicted variable and converge to zero over time.
2. The convergence rate of percentage forecast errors can be different across shocks.
3. The dispersion of forecasts should not increase after a shock.

Observe that the imperfect information model makes the same qualitative predictions about forecast errors and mean forecasts as the sticky information model. In contrast to the sticky information model, the

[16] Dispersion of forecasts can respond to shocks in the imperfect 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.

imperfect information model implies that dispersion of forecasts does not respond to fundamental shocks and the speed of response to shocks can vary across shocks.

3 Data Description

We have described three predictions of each model of informational rigidities. We will test these predictions using three data sets of inflation surveys. The first is the Michigan Survey of Consumers (MSC).[17] The MSC is a nationally representative monthly survey of 500 to 1,300 consumers. Respondents are asked to report their expected inflation rate for the next twelve months. The MSC is collected over the corresponding month. The survey question on inflation forecast begins in January 1978. The second source is the Survey of Professional Forecasters (SPF).[18] This is a quarterly survey of 9 to 40 professional forecasters. The SPF collects quarterly forecasts (for the current quarter and the next four quarters) for a number of macroeconomic variables. We focus on forecasts of average GDP Deflator inflation over the next year for consistency with the MSC forecasts. SPF forecasts are collected in the middle of the quarter. The third dataset is the Blue Chip Economic Indicators (BCEI).[19] This is a collection of forecasts from professional forecasters for similar set of macroeconomic variables as the SPF. The consensus BCEI forecasts are available monthly starting in April 1980 but dispersion measures can only be constructed starting in July 1984. For the BCEI, we focus on forecasts of the CPI. For each dataset, we extract a mean forecast of average inflation over the next year and a measure of the cross-sectional dispersion of those forecasts. We discuss the construction of our data in more detail in Appendix A.

Figure 1 plots the mean forecasts and forecast dispersion series from each survey in Panels A and B respectively. The mean forecasts are highly correlated with each other. We measure dispersion as the log of the cross-sectional standard deviation of inflation forecasts. The standard deviation of MSC is much higher than for professional forecasters.[20] All three measures point to decreasing levels of disagreement since the early 1980s. Table 1 presents correlations of these measures with macroeconomic variables. All three mean inflation forecast measures are highly correlated with inflation and with each other, and, to a lesser extent, with the unemployment rate. They are also slightly negatively correlated with the growth rate of real GDP. The cross-correlation among dispersion measures is lower than for

[17] The data is available at http://www.sca.isr.umich.edu/. We drop all survey responses suggesting that inflation or deflation may exceed 49%.
[18] The data is available at http://www.philadelphiafed.org/econ/spf/index.cfm.
[19] This data is proprietary.
[20] Note that one cannot directly compare the level of dispersion of BCEI forecasters to other measures because of slight differences in the construction of this series. See Appendix A for details.

mean forecasts. Nonetheless, these unconditional correlations indicate an apparently strong link between the degree of disagreement among agents and macroeconomic conditions, a point emphasized by Mankiw, Reis and Wolfers (2003). Interestingly, since the mid-1980s, the correlation between disagreement about inflation and the unemployment rate exceeds the correlation of disagreement with the inflation rate for two of the three survey measures.

Because the predictions made by the two 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. We considered the following shocks from the literature:

a) Monetary policy shocks, identified from a VAR or a la Romer and Romer (2004)
b) Technology shocks, identified using long-run restrictions as in Gali (1999).
c) Oil shocks, identified as in Hamilton (1996).
d) Information shocks, identified as in Beaudry and Portier (2006).
e) Confidence shocks, identified as in Barsky and Sims (2008).
f) Fiscal shocks, from Romer and Romer (2007).

Appendix B discusses the details of each approach. Table 2 presents the cross-correlation matrix for these shocks. Most shock measures are largely uncorrelated with one another, consistent with their interpretations as exogenous structural shocks to the economy. The highest correlations are between the confidence shock of Barsky and Sims and the information shock of Beaudry and Portier, which indicates that there is some overlap between the two shock series, as well as between the confidence shock and oil price shock.

Our selection criterion for shocks is to focus primarily on the three shocks which account for the largest component of the variance of inflation. Table 3 presents a variance decomposition of inflation into the orthogonalized components of these shocks.[21] The most important shock in terms of explaining the volatility of inflation appears to be technology shocks, as identified using long-run restrictions by Gali (1999), accounting for approximately twenty-five to thirty-five percent of the variance of inflation. The next two most quantitatively important shocks are monetary policy and oil price shocks. Monetary policy shocks account for up to five percent of the variance of inflation, while oil price shocks account for up to twenty percent at long time horizons. Each of the remaining three shocks, information, confidence, and fiscal, accounts for less than five percent of the variance of inflation. It is worth noting that despite our including of six shock measures, these jointly account for only about half of the variance of inflation, leaving much of the volatility unexplained by these structural shocks.

[21] The ordering has no qualitative effect on the results since the shocks are already largely orthogonal to one another.

Given these shocks, we can consider the impulse response of inflation to each shock as well as the predicted response of forecast errors and dispersion measures from the models.[22] For our three baseline shocks, monetary policy, technology and oil shocks, we plot the impulse response of annual inflation based on the whole sample, as well as predicted responses from the sticky information model in Figure 2. In response to monetary policy shocks, inflation follows the well-known delayed response, reaching its minimum approximately two years after the interest rate increase. In response to such a shock, the sticky information model predicts a negative forecast error, as mean forecasts should decline more slowly than actual inflation. Forecast dispersion, on the other hand, should rise as disagreement increases between those who have observed the shock and those who have not. Disagreement disappears as more and more agents learn about the shock. In response to a positive technology shocks, the response of inflation is much more rapid and more precisely estimated. Forecast errors are predicted to decline by the sticky information model while dispersion should rise. Finally, oil price shocks lead to an increase in inflation. With informational rigidities, this yields a prediction of positive forecast errors after the shock. The prediction of imperfect information models would be a nearly identical path of forecast errors, but no response of dispersion after each shock.

4 Empirical Analysis

We begin the empirical analysis with a preliminary check on the response of mean forecasts to the exogenous shocks. The question here is whether forecasts move in a manner which is generally consistent with the ex-post path of inflation. We estimate the following equation for each type of shock k and for each measure of mean year-ahead (h = 4 or 12) inflation forecasts:

$$F_t\pi_{t+h} = c + \sum_{i=1}^{I}\beta_i F_{t-i}\pi_{t-i+h} + \sum_{j=0}^{J}\gamma_j \varepsilon_{t-j}^{k} + v_t \tag{8}$$

This equation is similar to the specification estimated in Romer and Romer (2004).[23] We then plot the impulse responses to a unit increase in the exogenous shock, along with 95% confidence intervals in

[22] Impulse responses of inflation to monetary policy and technology shocks come from including inflation in each VAR, while the response to oil shocks (following Kilian (2007)) comes from estimating an AR(2) process for inflation augmented with current oil shocks plus 6 lags of the oil shock. All standard errors are bootstrapped. Inflation measures used are the PCE price index for monetary shocks, GDP deflator for technology shocks, and CPI for oil shocks. The distribution of predicted forecast errors and dispersion responses are based on the bootstrapped distribution of inflation responses.

[23] Note that our procedure has two steps (estimate shocks and then regress the variable of interest on these estimated shocks) and one generally has to appropriately adjust standard errors in the second step, which corresponds to our equation (8) and similar specifications considered later in the text. This adjustment, however, is not necessary in our analysis. Our null hypothesis is that the shocks have no effect on the mean forecast, forecast error, and forecast dispersion, i.e. $\gamma_j = 0$ for all j. Pagan (1984) shows that under this null it is not necessary to adjust standard errors.

Figure 3.[24] Lag lengths (*I* and *J*) are selected based on the BIC allowing for a maximum lag length of a year and a half for each specification.

In response to monetary policy shocks, the MSC and SPF both have responses of mean forecasts which are very imprecisely estimated. The response of mean forecasts from the BCEI is negative and marginally insignificant at the 5% level.[25] In response to technology shocks, we see a much more pronounced response of inflation forecasts: for all three survey measures, forecasts of future inflation fall and the decline is statistically significant for two of the survey measures. In response to oil shocks, we also see increases in inflation forecasts, consistent with the inflation process after such shocks. Importantly, agents do adjust their forecasts in response to shocks and thus we may be confident that the shocks, particularly technology and oil price shocks, contain genuinely new information and that they create meaningful variation in information sets suitable for our analysis.

4.1 Response of Forecast Errors to Shocks

The responses of mean forecasts tell us little about informational rigidities but simply indicate that mean forecasts go in the same direction as actual inflation. This result would be expected to hold under full information as well as the models with informational rigidities considered here. A more relevant test for the importance of informational rigidities is looking at the response of forecast errors conditional on an exogenous shock. The presence of informational rigidities, be they sticky information or imperfect information, implies that we should observe persistent conditional forecast errors whereas full information implies that forecast errors will be zero after the time of the shock. Forecast errors are generated using real-time data of inflation measures.[26] To estimate the response of forecast errors, we run the following regression:

$$\pi_t - F_{t-h}\pi_t = c + \sum_{i=1}^{I} \beta_i (\pi_{t-i} - F_{t-i-h}\pi_{t-i}) + \sum_{j=0}^{J} \gamma_j \varepsilon_{t-j}^k + v_t \tag{9}$$

[24] Confidence intervals for impulse responses are generated based on 1,000 draws from the asymptotic distribution of parameter estimates of equation (8), from which we extract our distribution of impulse responses.

[25] Note that the fact that the mean forecasts do not significantly respond to monetary policy shocks does not mean that economic agents ignore these shocks. The response of the actual inflation to monetary policy shocks is fairly weak in the short run and thus one may expect that mean forecasts should have a weak response too. It is likely that with these weak responses we cannot clearly discern the effect of monetary policy shocks.

[26] For both the MSC and the BCEI, we use the CPI as the measure of inflation. For SPF, we use the GNP deflator prior to 1992, the implicit GDP price deflator from 1992 to 1995, and the chained GDP price index starting in 1996. This reflects the different measures forecasted by SPF over time. Each inflation series is real-time data. Specifically, we use the level of inflation that was available six months after the inflation date. This is to avoid identifying revisions in price indexes as forecast errors.

where h = 4 for quarterly data and h = 12 for monthly data. We plot the impulse responses and associated standard errors for forecast errors starting a year after the shock.[27] The results are presented in Figure 4.

For monetary policy shocks, there is only limited evidence of serially correlated forecast errors. MSC and BCEI forecast errors are negative, as would be expected under informational rigidities, but these are (marginally for MSC) insignificant at the 5% level. SPF forecast errors are positive, but these are statistically insignificant as well. For technology shocks, on the other hand, forecast errors are persistently negative and statistically significant for all 3 survey measures. This indicates significant departures from full information, since inflation forecasts are responding more slowly than actual inflation. With oil shocks, the evidence again points to informational rigidities: forecast errors are all positive and two out of three are statistically significant. Finally, all forecast errors converge to zero, so agents' forecasts correctly converge to true values over time.

These results point to the potential importance of informational rigidities. In response to technology and oil price shocks, we find clear evidence that forecasts fail to adjust one-for-one with ex-post inflation. This is particularly important since these two shocks appear to play such an important role in explaining inflation volatility. On the other hand, the results for monetary policy shocks are mixed. This could be interpreted in one of at least three ways. One may take this as evidence as suggesting that monetary policy shocks are observed (and fully processed) by all and are thus not subject to informational rigidities. A second interpretation would simply question whether monetary policy shocks are shocks at all, as done in Cochrane (1994). A third interpretation is that because monetary policy shocks play such a small role in explaining the volatility of inflation (and hence inflation expectations) and because the contemporaneous and short-run responses of actual inflation to monetary policy shocks are weak, there is little hope of being able to clearly discern the response of forecast errors to this shock. In section 4.4, we argue based on Monte Carlo simulations that this last interpretation is not unjustified.

4.2 Persistence of Conditional Percentage Forecast Errors

In this section, we consider another set of predictions of the two models. Sticky information implies that for a common set of agents, percentage forecast errors after a shock will decline monotonically at the same rate for all shocks. The persistence of conditional percentage forecast errors in this model hinges only on the degree of informational rigidity (λ). For imperfect information agents, on the other hand, the weight placed on a signal correlated with a shock depends positively on the variance of the shock (or more generally on the benefit of paying attention to the shock, as in rational inattention models). Thus,

[27] The response of forecast errors in the first year reflects only variation in inflation, since the date of the forecasts precedes that of the shock.

one would expect according to this model to find more persistence in the conditional forecast errors for shocks which have little effect on inflation than on technology shocks which account for the largest component of inflation volatility. In this section, we estimate the persistence of conditional percentage forecast errors in response to different shocks and ask whether this persistence is common across shocks and agents.

To do so, we first run the following regression:

$$\frac{(\pi_t - F_{t-h}\pi_t)}{\pi_t} = c + \sum_{i=1}^{I} \beta_i \frac{(\pi_{t-i} - F_{t-i-h}\pi_{t-i})}{\pi_{t-i}} + \sum_{j=0}^{J} \gamma_j \varepsilon_{t-j}^{k} + v_t \tag{10}$$

again selecting I and J using the BIC criterion, using one-year ahead forecasts, and individually for each shock and survey measure. Using the coefficients estimated from equation 10, we extract the impulse response of the percentage forecast error to the shock. The persistence of the conditional percentage forecast error is measured using the parameter θ which minimizes the squared distance between the estimated impulse response and the geometrically declining process governed by θ (i.e. θ^j for $j \geq 0$).[28] Results are presented in Table 4.[29]

First, *we find little evidence that the persistence of conditional forecast errors differs across shocks for each survey measure*. For consumers, the persistence of forecast errors after monetary policy shocks is somewhat lower than that after technology and oil price shocks. For SPF, the persistence is almost identical across shocks while for BCEI forecasts, the persistence of forecast errors is somewhat lower after oil price shocks than after monetary and technology shocks. Overall, the persistence of conditional forecast errors across shocks is quite similar for each survey measure. The half-life of the forecast error responses is about two-four quarters. Strikingly, forecast errors do not converge to zero any more gradually for monetary policy shocks than for other shocks. This result is surprising because monetary policy shocks account for a much smaller fraction of inflation volatility than technology and oil price shocks. One would expect imperfect information agents to pay more attention to technology and oil prices and therefore to update their forecasts more rapidly after these types of shocks. This is exactly the intuition underlying the rational inattention models of Paciello (2007) and Mackowiak and Wiederholt (2008) which seek to explain why inflation responds more rapidly to technology shocks than monetary policy shocks. Our results are inconsistent with a slower rate of information acquisition after monetary policy shocks than for other shocks. These results are in fact broadly consistent with the sticky

[28] We normalize the initial value of the impulse response of conditional forecast errors to be one. Standard errors are extracted by drawing from the empirical distribution of parameter estimates of (10), calculating impulse responses and solving for θ for each impulse response

[29] Note that although impulse responses may be imprecisely estimated, we estimate the parameter θ quite precisely because we utilize the information contained in the full path of the response.

information model of Mankiw and Reis which imposes a common rate of information updating which is independent of the type of shock.

Secondly, *we find no evidence that consumers acquire information in a less efficient manner than professional forecasters*. Our estimates indicate that forecast errors by consumers are actually less persistent than those of SPF forecasters (though not statistically significantly so) and approximately as persistent as those of BCEI forecasters. Carroll (2003) proposed and found empirical evidence in the U.S. for an epidemiological model of gradual information diffusion from professional forecasters to consumers which implied that mean forecasts of consumers would respond more slowly than that of professional forecasters. Dopke et al (2008) provide similar evidence for selected European countries. Our estimates indicate, on the other hand, that the persistence of consumers' conditional forecast errors is no higher than that of professional forecasters in response to the monetary, technology, and oil price shocks. Furthermore, we find the same pattern for unidentified shocks (see section 5.4). Hence, consumers respond to all shocks as rapidly as professional forecasters. This result contradicts the findings of Carroll (2003) in the context of the structural shocks which seem to account for approximately fifty percent of inflation volatility.

4.3 Response of Forecast Dispersion to Shocks

Now we turn to the response of the dispersion of inflation forecasts to the same shocks. The following estimation equation is used:

$$\ln\sigma(F_t\pi_{t+h}) = c + \sum_{i=1}^{I}\beta_i \ln\sigma(F_{t-i}\pi_{t-i+h}) + \sum_{j=0}^{J}\gamma_j\,|\,\varepsilon_{t-k}^{k}\,| + v_t \tag{11}$$

where $\ln\sigma(F_t\pi_{t+h})$ is the logarithm of the cross-sectional standard deviation of inflation forecasts over the next year.[30] We include the absolute value of the shock because the predicted response of forecast dispersion in the sticky information model is invariant to whether the shock is positive or negative.

Figure 5 plots the impulse responses of the forecast dispersion measures for each survey in response to the exogenous shocks, along with 95% confidence intervals. In response to monetary policy and technology shocks, we find no evidence that forecast dispersion rises after these shocks for any of the

[30] Taking logarithm serves two purposes. First, the distribution of forecast dispersion is highly skewed and rare observations in the right tail of the distribution can greatly distort the estimates. The logarithmic transformation makes the dispersion bell-shaped, attenuates the adverse effects of episodes with extreme variability in forecasts, and hence improves the finite sample properties of our estimates. Second, we can interpret impulse responses as percent deviations from steady state forecast dispersion.

survey measures.[31] In response to oil price shocks, the dispersion of forecasts of consumers displays a statistically significant and highly persistent increase after the shock. No such response is apparent for the Survey of Professional Forecasters or the BCEI. Thus, for professional forecasters, our results are consistent with imperfect information: professional forecasters can adequately be modeled as continuously tracking economic indicators without observing the underlying state perfectly. However, for consumers, the results are mixed. While monetary and technology shocks do not lead to significant changes in forecast dispersion, oil price shocks seem to lead to highly persistent increases in disagreement. Thus, although we cannot unanimously favor either representation of the informational rigidities facing consumers, the preponderance of evidence suggests that *disagreement in forecasts does not respond to structural shocks*.

4.4 Monte Carlo Simulations

While the conditional responses of forecast errors to technology and oil price shocks are remarkably consistent with the presence of informational rigidities, the evidence of serial correlation in forecast errors after monetary policy shocks is much weaker. One reason why this could occur is that monetary policy shocks account for a very small fraction of the inflation volatility, making it difficult to precisely estimate the response of forecast errors in short time sample. In this section, 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

$$\pi_t = \rho\pi_{t-1} + \varepsilon_t$$

where the shock ε_t is the sum of two independent innovations: $\varepsilon_t = \varepsilon_t^{(1)} + \varepsilon_t^{(2)}$. 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 of next-period inflation at time *t* follows

$$F_t\pi_{t+j} = (1-\lambda)\sum_{k=0}^{\infty}\lambda^k E_{t-k}\pi_{t+1}$$

and the cross-sectional distribution of forecasts of next-period inflation is

[31] It must be noted that this result is sensitive to the time period for monetary policy shocks. If one uses data starting in the late 1970s, monetary policy shocks (in absolute value) lead to a *permanent* increase in dispersion of SPF forecasts. We choose to focus on the post 1984 period for monetary policy shocks because FFR shocks are extremely noisy between 1979 and 1982, reflecting the fact that the Fed abandoned targeting the federal funds rate over this period. Thus, we are concerned that the identified FFR shocks prior to 1984 are not adequate measures of monetary policy innovations. None of the other results are sensitive to the time period.

$$V_t\pi_{t+j} = (1-\lambda)\sum_{k=0}^{\infty}(E_{t-k}\pi_{t+j} - F_t\pi_{t+j})^2 .$$

We simulate this model 1,000 times, with each simulation having 150 times periods. We set $\rho = 0.85$ and the variance of the total shock to inflation to be $\sigma_\varepsilon^2 = 1.005$ to match estimates of an AR(1) process for GDP deflator inflation from 1979Q1 to 2007Q3. Following Mankiw and Reis (2002), the degree of informational rigidity λ is set to 0.75, implying that agents update their information once a year on average. In each simulation, we estimate the response of forecast errors and forecast dispersion to innovations $\varepsilon_t^{(1)}$, as done in sections 4.1 and 4.3, using the final 100 periods of the simulation.[32] The key parameter that varies across simulations is the fraction of the inflation variance accounted for by the innovation used to derive impulse responses ($\varepsilon_t^{(1)}$). We consider four values: S =1%, 5%, 10% and 20%.

Results of the Monte Carlo exercise are presented in Figures 6 and 7. When the fraction of the inflation variance accounted for by the innovation is ten percent and greater, our estimation yields very precise estimates of the conditional response of forecast errors. At five percent, the distribution of responses is wider, but more than 95% of impulse responses have the correct sign. However, once the fraction of the inflation variance diminishes below five percent, it becomes much more difficult to precisely estimate the response of the forecast error and the 95% confidence interval now includes zero. We view this as a being a plausible interpretation for why our estimation fails to uncover a significant response of forecast errors to monetary policy shocks. Because these innovations account for such a small fraction of the inflation variance, precisely estimating the conditional response of forecast errors could be tenuous in such short time samples.

Turning to the conditional response of forecast dispersion to the innovations in Figure 7, the results are much less precise than with forecast errors. First, median estimates differ substantially from the theoretical response of dispersion for all fractions of inflation variance considered. This reflects the fact that the dispersion process is not as well approximated with an ARMA process as forecast errors. However, while the point estimates consistently differ from the true values, the estimation does consistently and correctly find a positive response of dispersion to the innovations, even when the fraction of the inflation variance accounted for by the innovation is quite small. Thus, to the extent that our analysis is restricted to looking for whether dispersion rises after innovations or not, the estimation approach used in section 4.3 should consistently yield the correct answer.

[32] In estimating the conditional response of forecast errors, we use an AR(1) process augmented with the shock series $\varepsilon_t^{(1)}$. One can readily verify that forecast errors follow this process given the AR(1) process of inflation used in the model. For forecast dispersion, we use an AR(1) process augmented with the absolute value of $\varepsilon_t^{(1)}$ and 12 lags of this variable.

5 Robustness Analysis

In this section, we investigate the robustness of our results to several issues. First, we reproduce our results using VARs for the estimation. Second, we allow for an alternative set of shocks. Third, we generate a composite shock that aggregates our three baseline structural shocks and study the response of forecast errors and forecast dispersion to this composite shock. Fourth, we present the response of forecast errors and dispersion to the unidentified component of inflation innovations. Finally, we reproduce our baseline estimates for forecasts of the unemployment rate rather than inflation.

5.1 Joint Estimation of Impulse Responses

Our baseline results rely on single equation estimates of the impulse responses of forecast errors and forecast dispersion to structural shocks. This approach is valid if our shock measures are truly orthogonal to one another. However, as indicated in Table 3, some of the shocks are correlated with one another. One way of controlling for this is to estimate the impulse responses using a VAR.[33] We consider a 4-variable VAR with our three shock measures ordered first (FFR shocks, then technology shocks, then oil price shocks) followed by either forecast errors or forecast dispersion. Using the BIC to select lag lengths, impulse responses of forecast errors and forecast dispersion from orthogonalized innovations to the shock measures are presented in Figures 8 and 9 respectively.[34]

The results are remarkably consistent with the single equation estimates. Point estimates of the conditional response of forecast errors are always in the direction predicted by models of informational rigidities. For monetary policy shocks, however, the responses of forecast errors are again insignificantly different from zero, as found with the single equation estimates. All other responses of forecast errors are statistically significantly different from zero, with the one exception being the response of consumer forecast to technology shocks. Thus, we conclude that our baseline results identifying strong evidence for informational rigidities in inflation forecasts are robust to joint estimation of the impulse responses.

Turning to the impulse responses of forecast dispersion to the orthogonalized innovations in the absolute value of the shock measures, the results point more strongly toward imperfect information than

[33] Another way would be to use single equation estimation for each survey measure but including all the shocks (and lags) on the RHS. The VAR approach has the additional advantage of purifying the structural shock of possible autocorrelation.

[34] For forecast errors, we again use $\pi_t - F_{t-h}\pi_t$ as the dependent variable, but present impulse responses starting a year after the shock. This is because we only care about the behavior of forecast errors for periods when the date of the forecasts is greater than or equal to the date of the innovation. All estimates are computed using quarterly data on the post-1984 period. When estimating the response of forecast dispersion, we again use the absolute value of shocks, as described in section 4.3. The 95% confidence interval is computed by drawing 1,000 times from the distribution of parameter estimates and from the variance covariance matrix of residuals and then computing new impulse responses for each draw.

the single equation estimates. In every case, we cannot reject the null that dispersion is invariant to the structural shocks. Nonetheless, it must be noted that the increase in dispersion of consumers after oil price shocks is only marginally insignificant. Thus, while imperfect information now appears to be a better representation for consumers than sticky information, the evidence is weaker than for professional forecasters.

5.2 Other Shock Measures

In this section, we reproduce the forecast error and forecast dispersion tests for three alternative shocks. First, we use the Romer and Romer (2004) measures of innovations to the FFR. These are available monthly from April 1970 until December 1996. Second, we use the information shocks of Beaudry and Portier (2006), which capture changes in stock prices that have no contemporaneous effect on TFP. These are designed to represent expected changes in TFP and cause temporary decreases in inflation. Third, we use the fiscal shocks of Romer and Romer (2007) which provide a sequence of exogenous tax changes and have temporary deflationary effects.[35] All three shocks are deflationary, so both models of informational rigidities would predict a sequence of negative forecast errors after the shock.

The results for the forecast errors are provided in Figure 10. While most of the forecast errors go in the expected direction, almost none of the responses are statistically significantly different from zero. Given the strong evidence for informational rigidities found for the more quantitatively important (for inflation) technology and oil price shocks, it is unclear whether one should interpret this lack of a response of forecast errors as being consistent with full information or as reflecting the fact that, because these shocks account for such a small component of the variance of inflation, it is unlikely that one will be able to precisely estimate the response of forecast errors in small samples (as argued in section 4.4). The results for the response of dispersion in forecasts to the absolute value of each exogenous shock are provided in Figure 11. Overall, there is little evidence of changes in disagreement after these structural shocks. Nonetheless, given the lack of a clear response of forecast errors to these alternative structural shocks, we are wary of placing too much weight on the impulse responses of dispersion to these same shocks.

[35] We do not present results for confidence shocks because these seem to have to no discernible effects on inflation. Thus, one could not tell whether informational rigidities would imply positive or negative conditional responses of forecast errors. We find these shocks lead to positive responses of forecast errors for MSC, (marginally) negative for BCEI, and insignificant for SPF. There is no response of dispersion for any of the surveys to the absolute value of the shocks.

5.3 Composite Shock

The Monte Carlo exercise of section 4.4 illustrates that when shocks explain only a small fraction of the variance of inflation, the conditional response of forecast errors and forecast dispersion will tend to be imprecisely estimated. Indeed, this seemed to be the case for forecast errors after monetary policy shocks in the baseline estimation. As an alternative approach designed to minimize this issue, we develop a composite shock which is an aggregate of the individual identified shocks and present responses of forecast errors and dispersion to this shock.

To do so, we first run an AR(4) on CPI inflation the residuals of which we interpret as inflation innovations. We then regress these innovations on the monetary policy shocks, the technology shocks and the oil price shocks. The predictable component of the inflation innovation we define as the composite shock. By construction, these innovations are inflationary. Figure 12 presents the conditional response of forecast errors and forecast dispersion to the composite shock, estimated as in sections 4.1 and 4.3, for all three survey measures. For every survey measure, the response of forecast errors is positive and highly significant. Each survey delivers a positive sequence of monotonically declining forecast errors, as predicted by models of informational rigidities. As found in the baseline results, the response of consumers to shocks is at least as quick as that of professional forecasters. The response of dispersion to the (absolute value of) composite shock is a little more puzzling. For consumers, dispersion is unresponsive to the shock, as was the case after monetary policy and technology shocks but not oil price shocks. For BCEI forecasters on the other hand, the response is marginally statistically significant and positive, whereas in response to individual shocks, there was no evidence of changes in dispersion after any of the shocks. For the SPF, we get the same result as before; dispersion appears to be invariant to the composite shock. Results of this check are in broad agreement with our baseline findings reported in section 4.

5.4 Unidentified Inflationary Innovations

While all of our analysis has focused on the conditional response of forecast measures to identified shocks, Table 3 indicates that only half of the inflation volatility is accounted for by these shocks. One could, of course, consider additional shock measures (markup shocks, preference shocks, etc) from the literature. Instead, we now consider the impulse response of forecast errors and forecast dispersion to unidentified inflation innovations. Specifically, we estimate an AR(4) process for inflation, then regress the residuals on the monetary policy shocks, technology shocks, and oil price shocks. The residuals from this second regression we call the unidentified innovations to inflation.[36] By construction, these innovations are inflationary, and should thus lead to positive forecast errors if informational rigidities are present.

Figure 13 presents the conditional response of forecast errors and forecast dispersion to these unidentified inflation innovations, using the same approach as in sections 4.1 and 4.3. For all three survey measures, the response of forecast errors is positive and statistically significant, although less precisely estimated than in the case of the composite shock. All of the impulse responses also monotonically decline to zero. We interpret this as indicating that the evidence for informational rigidities is not limited to the structural innovations considered before but extends more generally to all innovations to inflation, as long as these account for a large enough fraction of the inflation variance. The response of dispersion to these unidentified innovations is consistent across survey measures: we find no evidence that dispersion changes after these shocks.

5.5 Unemployment Forecasts

While almost all of the literature of survey expectations has focused on forecasts of inflation, one can also test for informational rigidities using real variables. In this section, we focus on year-ahead forecasts of the unemployment rate which are available from the Survey of Professional Forecasters (starting in 1974Q4) and the Blue Chip Economic Indicators forecasts (starting in June 1981). As with inflation, we choose to focus on the shocks which are most important in explaining fluctuations in the unemployment rate. Table 5 presents a variance decomposition of the unemployment rate based on running a VAR(4) of all the shocks followed by the unemployment rate over the period of 1980Q1 until 2006Q4. At horizons of two years and up, the most important shocks appear to be the information shocks of Beaudry and Portier (40-50%), followed by the confidence shocks of Barsky and Sims (3-6%) and the oil price shocks of Hamilton (3-6%).

The impulse responses of UE forecast errors and forecast dispersion from the SPF and BCEI surveys to these three shocks are presented in Figure 14.[37] It is important to note that oil price shocks lead to temporary increases in the unemployment rate while information and confidence shocks lead to decreases in the unemployment rate.[38] Thus, models of informational rigidities predict that conditional forecast errors should be positive after oil price shocks and negative after information and confidence

[36] Note that these innovations include the alternative shocks from section 5.2. Controlling for these shocks as well has no effect on the results.

[37] Forecast errors are computed using real-time data for unemployment rates. Specifically, we use the unemployment rate that was available six months following the date of the variable. This real-time data comes from the Philadelphia Federal Reserve Bank.

[38] This is based on impulse responses of the unemployment rate to each of the shocks. We do not present these responses in the interest of space, but all are statistically significant.

shocks. As can be seen in Figure 14, the evidence for informational rigidities based on forecasts of the UE rate is more mixed than for inflation forecasts. After oil price shocks, forecast errors are not statistically different from zero. After information shocks, we find much clearer evidence for informational rigidities for both survey measures, as the conditional response of the forecast errors is negative and statistically significant. With confidence shocks, there is no discernible response of forecast errors for SPF forecasters, and the response of the forecast error for BCEI forecasters is actually of the wrong sign. This would imply that BCEI forecasters overreact in lowering their forecasts of the unemployment rate after a confidence shock.

Following the results of the Monte Carlo exercise of section 4.4, we interpret these results as saying that for the one shock that explains a significant fraction of the volatility in the unemployment rate (information shocks), there is clear evidence of informational rigidities since forecast errors decrease on impact before converging back to zero. Because the other shocks explain such a small component of the unemployment volatility, we are hesitant to place much weight on the absence of a response of the forecast errors to these shocks.

Finally, turning to the response of forecast dispersion to the absolute value of the structural shocks, we find no evidence that these shocks lead to increases in disagreement among forecasters. This is consistent with imperfect information models but inconsistent with the sticky information model. One should be wary of putting too much weight on the results for oil price shocks and confidence shocks since these shocks fail to deliver strong evidence for informational rigidities based on the response of forecast errors. However, in the case of information shocks, which yield responses of forecast errors consistent with the presence of informational rigidities, the fact that dispersion is not increasing after a shock is quite consistent with imperfect information models but cannot be readily reconciled with delayed information updating.

6 Conclusion

Recent work utilizing informational rigidities offers potential explanations for some important macroeconomic puzzles such as the differential response of inflation to technology and monetary policy shocks and persistent hump-shaped response of inflation to nominal shocks and provides a strong amplification and propagation mechanism for business cycle models. Yet direct empirical evidence for informational rigidities has been limited, thereby raising fundamental questions about the validity of these models. Previous research on survey measures has focused on testing the null of full information by regressing forecast errors on lagged variables. While these tests are informative, they have yielded mixed results and do not address the issue of the dynamic behavior of forecast errors after shocks. Our results instead paint a clear picture: after structural shocks, agents fail to adjust their forecasts by a sufficient amount, inducing a non-zero response of forecast errors. As time goes by, forecast errors converge monotonically to the full information outcome. This evidence is particularly strong for structural shocks which have a significant quantitative effect on the variable being forecasted, whether it be inflation or unemployment. We interpret these results as providing a robust empirical basis for models of informational rigidities that has previously been sorely lacking.

The conditional response of forecast errors to structural shocks also uncovers some stylized facts about informational rigidities across shocks and agent types. First, we find no evidence that consumers' forecasts adjust more gradually than professional forecasts after a shock. Second, the convergence rate of forecast errors is just as rapid for monetary policy shocks as it is for technology shocks. The first fact is inconsistent with epidemiological models of information diffusion from professional forecasters to consumers (as in Carroll (2004)). The second fact is consistent with sticky information models in which agents acquire full information about the state on an infrequent basis. Third, our results on the conditional response of the dispersion of forecasts across agents to shocks on the other hand point primarily to imperfect information as the most adequate representation of informational rigidities. For professional forecasters, with only one exception, there is no evidence that disagreement among agents increases after structural shocks, as predicted by the sticky information model. On the other hand, consumer disagreement appears sensitive to oil prices, which is inconsistent with simple linear imperfect information models.

Of course, there are other types of informational rigidities that have been considered in the literature, such as models of learning (see Milani (2007) for a recent application) or model-switching behaviors (as in Branch (2007)) which one could also analyze. Integrating conditional heteroskedasticity into these models may also help clarify the source of falling forecast dispersion between the volatile 1970s and more recent time periods. Further study of the behavior of disagreement among agents is likely to be fruitful for better understanding the nature of informational rigidities affecting agents.

References

Adam, Klaus and Mario Padula, 2003. "Inflation Dynamics and Subjective Expectations in the United States," European Central Bank Working Paper 222.

Andres, Javier, David Lopez-Salido, and Edward Nelson, 2005. "Sticky-Price Models and the Natural Rate Hypothesis," *Journal of Monetary Economics* 52(5), 1025-1053.

Barsky, Robert B. and Eric R. Sims, 2008. "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence." Manucript.

Bacchetta, Philippe, Elmar Mertens, and Eric van Wincoop. 2008. "Predictability in Financial Markets: What Do Survey Expectations Tell Us?" Forthcoming in *Journal of International Money and Finance.*

Bernanke, Ben S. and Alan S. Blinder, 1992. "The Federal Funds Rate and the Channels of Monetary Transmission," *American Economic Review* 82(4), 901-921.

Beaudry, Paul and Franck Portier, 2006. "News, Stock Prices and Economic Fluctuations," *American Economic Review* 96(4), 1293-1307.

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

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

Cochrane, John H., 1994. "Shocks," *Carnegie-Rochester Conference Series on Public Policy* 41, 295-364.

Coibion, Olivier and Yuriy Gorodnichenko, 2007. "Strategic Interaction among Heterogeneous Price-Setters in an Estimated DSGE Model," Manuscript.

Coibion, Olivier, 2007. "Testing the Sticky Information Phillips Curve," forthcoming in *Review of Economics and Statistics*.

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.

Döpke, Jörg, Jonas Dovern, Ulrich Fritsche, and Jirka Slacalek, 2008. "Sticky Information Phillips Curves: European Evidence," *Journal of Money, Credit and Banking*, forthcoming.

Dupor, Bill, Tomiyuki Kitamura, and Takayuki Tsuruga, 2007a. "Integrating Sticky Information and Sticky Prices," forthcoming in *Review of Economics and Statistics*.

Dupor, Bill, Jing Han, and Yi Chan Tsai, 2007b. "What Do Technology Shocks Tell Us about the New Keynesian Paradigm?" Manuscript.

Erceg Christopher J. and Andrew T. Levin, 2003. "Imperfect credibility and inflation persistence," *Journal of Monetary Economics* 50(4), 915-944.

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.

Jonung, Lars, 1981. "Are Perceptions of Inflation Rational? Some Evidence for Sweden," *American Economic Review* 78(5), 1080-1087.

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

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

Kilian, Lutz., 2007. "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?" Forthcoming: *Review of Economics and Statistics.*

Kim, Don H. and Athanasios Orphanides, 2005. "Term Structure Estimation with Survey Data on Interest Rate Forecasts," Finance and Economics Discussion Series Working Paper 2005-48, Board of Governors of the Federal Reserve System.

Knotek, Edward S., 2007. "A Tale of Two Rigidities: Sticky Prices in a Sticky Information Environment," Federal Reserve Bank of Kansas City Working Paper No. 06-015.

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, 2008a. "Business Cycle Dynamics under Rational Inattention," Manuscript.

Mackowiak, Bartosz and Mirko Wiederholt, 2008b. "Optimal Sticky Prices under Rational Inattention," forthcoming in *American Economic Review*.

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).

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers, 2003. "Disagreement about Inflation Expectations," *NBER Macroeconomics Annual 2003,* Cambridge, MIT Press, 18, 209-248, 2004.

Milani, Fabio. 2007. "Expectations, Learning, and Macroeconomic Persistence," *Journal of Monetary Economics*, vol. 54(7), pages 2065-2082.

Orphanides, Athanasios and John C. Williams, 2003. "Imperfect Knowledge, Inflation Expectations, and Monetary Policy," NBER Working Paper 9884.

Paciello, Luigi, 2007. "The Response of Prices to Technology and Monetary Policy Shocks under Rational Inattention," Manuscript.

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

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

Reis, Ricardo, 2006b. "Inattentive Consumers," *Journal of Monetary Economics, 53 (8), 1761-1800*

Reis, Ricardo, 2008. "A Few Model-Based Answers to Monetary Policy Questions in the United States and the Euro-Area," Manuscript.

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 and David H. Romer, 2007. “The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks,” NBER Working Paper No. 13264.
Sims, Christopher A., 2003. “Implications of Rational Inattention,” *Journal of Monetary Economics*, 50(3), 665-690.
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 S. Phelps,* Princeton Univ. Press, 2002.
Woodford, Michael, 2008. “Information-Constrained State-Dependent Pricing,” Manuscript.
Zbaracki, Mark J., M. Bergen, S. Dutta, M.J., D. Levy, and M. Ritson, 2004. "Managerial and Customer Dimensions of the Costs of Price Adjustment: Direct Evidence from Industrial Markets," *Review of Economic and Statistics* 86(2), 514-533 .

Appendix A: Description of survey measures and construction of data series.
Michigan Survey of Consumers

The samples for the Surveys of Consumers are statistically designed to be representative of all American households, excluding those in Alaska and Hawaii. Each month, a minimum of 500 interviews are conducted by telephone. After a series of questions, consumers are asked, "During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?" If consumers answer that prices go up or down, the surveyor asks the follow-up question "By about what percent do you expect prices to go (up/down) on the average, during the next 12 months?" The follow-up question first appeared in the questionnaire in January 1978. The data were taken from http://www.sca.isr.umich.edu/main.php. In our calculations of the mean forecast and the standard deviation of forecasts, we excluded responses suggesting that inflation or deflation can be above 49%. This truncation eliminates extreme implausible responses. See "Procedure to Estimate Price Expectations" (http://www.sca.isr.umich.edu/documents.php?c=i) for more details on the sensitivity of mean and variance to different truncation thresholds. No answers are imputed to respondents who had difficulties predicting inflation.

Survey of Professional Forecasters

The forecasts for the Survey of Professional Forecasters are provided by the Federal Reserve Bank of Philadelphia. The survey began in 1968:Q4. The data set contains the 31 economic variables currently included in the survey. Some variables were added in 1981:Q3 and later in 2003:Q4 and in 2007:Q1. Quarterly forecasts of CPI are annualized quarter-over-quarter percent changes. Forecasts for the quarterly and annual level of the GDP price index are seasonally adjusted. The base year varies. Prior to 1996, GDP implicit deflator. Prior to 1992, GNP deflator. Annual forecasts are for the annual average. For every survey period, we compute forecast dispersion as the log of the standard deviation of the individual forecasts. See http://www.philadelphiafed.org/econ/spf/ for mode details.

Blue Chip Economic Indicators

Blue Chip Economic Indicators surveys economic forecasters at approximately 50 banks, corporations, and consulting firms. Blue Chip Economic Indicators started to collect inflation forecasts in the second quarter of 1980 and reported consensus (median) forecast for next four to seven quarters. We use the consensus forecast in our estimations. Blue Chip Economic Indicators also reported individual forecasts for the current and the next year. We constructed the standard deviation of one-year-ahead forecasts by taking a weighted average of the cross-sectional standard deviation of current-year forecast and the standard deviation of next-year forecasts. The weights correspond to the next 12 months (i.e. for January survey, we use weight of 11/12 on this year and 1/12 on next year, for February survey, we use weight of 10/12 on this year and 2/12 on next year). Because the number of forecasters included in monthly surveys has a strong seasonal pattern, we seasonally adjust the standard deviation measure using the Census X12 multiplicative methodology before taking logs. The dispersion measure runs from July 1984 to November 2007.

Appendix B: Description of Identified Shocks

R&R FFR shocks are shocks to monetary policy identified in Romer and Romer (2004). These shocks are identified as deviations of fed funds rate from its target rate, conditional on Greenbook forecasts. Data is available at http://elsa.berkeley.edu/~cromer/index.shtmlfrom Jan. 1966 to Dec. 1996.

R&R fiscal shocks are exogenous tax changes as a share of GDP, from Romer and Romer (2007) who use narrative approach to identify exogenous changes in taxes. Data is quarterly from 1947Q1 to 2006Q2.

Oil price shocks are taken from Hamilton (1996) who identifies (WTI) oil price shocks as episodes when 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.

Monetary policy shocks identified from a VAR are based on a five variable vector autoregression with twelve lags estimated on the data running from January 1957 to November 2007. Variables included in the VAR are personal consumption expenditures deflator, index of industrial production, change in the index of commodity prices, unemployment rate and Fed Funds rate. Following Bernanke and Blinder (1992), monetary policy shock is identified recursively using Cholesky decomposition and ordering the fed funds rate last. This restriction amounts to assuming that monetary policy shocks do not affect contemporaneously the variables ordered before the Fed Funds rate.

Technology shocks are identified as in Gali (1999) who uses long-run restrictions. The estimation sample covers 1952Q2 through 2007Q3. Labor productivity and hours are defined as in Gali (1999). Technology shocks are identified from the restriction that only technology shocks have long run effect on productivity. To estimate the impulse response of inflation to a technology shock, 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.

Information shocks are identified as in Beaudry and Portier (2006). We use the short-run restrictions imposed on the residuals of a bivariate VAR(4) which includes total factor productivity (TFP) and the S&P500 price index. As discussed in Beaudry and Portier (2006), the short run restriction imposes that TFP does not respond to an information shock contemporaneously while the stock market does. The estimation sample covers 1951Q1 through 2006Q4.

Confidence shocks are identified as in Barsky and Sims (2008). We use a trivariate VAR(4) with the relative forecast for economic conditions for next five years reported in the Michigan Survey of Consumers (Table 16), log of real consumption of nondurables and services, and the log of real GDP. The confidence shock is identified using recursive ordering of the variables where expectations about economic conditions are ordered first. The survey measure begins in January 1978 and ends in August 2007. We use the quarterly average in the VAR.

Table 1: Correlations of Survey Measures with Macroeconomic Variables

Panel A: Mean Forecasts of Inflation

	MSC	SPF	BCEI
MSC	1.00		
SPF	0.89	1.00	
BCEI	0.89	0.98	1.00
Unemployment	0.37	0.63	0.60
RGDP Growth	-0.24	-0.10	-0.18
Inflation	0.87	0.86	0.85

Panel B: Dispersion of Inflation Forecasts

	MSC	SPF	BCEI
MSC	1.00		
SPF	0.45	1.00	
BCEI	0.47	0.58	1.00
Unemployment	0.68	0.51	0.42
RGDP Growth	-0.21	0.04	-0.04
Inflation	0.31	0.24	0.26

Note: Correlations are over common samples. For Panel A, data is from 1980Q2-2007Q1. For Panel B, data is from 1984Q3-2007Q1. Inflation is measured using log-difference of GDP Deflator.

Table 2: Correlation of exogenous macroeconomic shocks

	VAR FFR	Gali Technology	Beaudry Information	Barsky Confidence	Romers Fiscal	Hamilton Oil
FFR	1	0.01	-0.07	0.01	-0.11	-0.03
Tech	0.01	1	-0.10	0.17	-0.05	-0.19
Information	-0.07	-0.10	1	0.30	0.09	-0.16
Confidence	0.01	0.17	0.30	1	-0.02	-0.29
Fiscal	-0.11	-0.05	0.09	-0.02	1	0.03
Oil	-0.03	-0.19	-0.16	-0.29	0.03	1

Note: The sample runs from 1979Q1 until 2004Q3.

Table 3: Variance Decomposition of Inflation

Quarters	FFR	Tech	Info	Confidence	Fiscal	Oil	Inflation
1	0.3	27.6	0.3	0.8	2.4	1.6	67.0
2	4.4	34.6	0.6	1.9	1.8	1.4	55.2
3	5.0	34.9	2.4	1.7	1.8	2.1	52.1
4	4.6	33.0	2.3	2.9	1.7	2.3	53.3
8	3.3	27.3	2.1	2.5	1.2	15.3	48.3
12	3.0	25.1	2.2	2.4	1.3	18.7	47.2
16	2.9	24.3	2.3	2.3	1.3	19.7	47.3
20	2.8	24.0	2.2	2.3	1.3	20.0	47.3

Note: Data is from 1980Q1 until 2006Q2. 4 lags in VAR, Cholesky decomposition as ordered in table. Variables are: 1) FFR: monetary policy shocks, 2) Tech: technology shocks identified a la Gali (1999), 3) Info: information shocks a la Beaudry and Portier (2006), 4) Confidence: Confidence shocks a la Barsky and Sims (2008), 5) Fiscal: fiscal shocks from Romer and Romer (2007), 6) Oil: oil price shocks from Hamilton (1996), 7) Inflation: annualized quarterly log change in the implicit GDP deflator.

Table 4: Persistence of Conditional Percentage Forecast Errors

	MSC	SPF	BCEI
FFR Shock	0.48	0.78	0.59
	(0.10)	(0.11)	(0.21)
Technology Shock	0.65	0.75	0.62
	(0.07)	(0.09)	(0.21)
Oil Price Shock	0.65	0.79	0.49
	(0.09)	(0.09)	(0.27)
Unidentified Shock	0.76	0.79	0.82
	(0.09)	(0.08)	(0.14)

Note: These are estimates of θ, the persistence of conditional forecast errors, as described in section 4.2. All estimates correspond to quarterly frequency. Standard errors are in parentheses. Standard errors are computed using 1,000 Monte Carlo simulations of the impulse responses based on draws from the estimated asymptotic distribution of parameters in equation (10). The unidentified shock is described in section 5.4.

Table 5: Variance Decomposition of the Unemployment Rate

Quarters	FFR	Tech	Info	Confidence	Fiscal	Oil	UE
1	0.1	2.5	0.0	3.2	3.8	1.5	88.9
2	0.0	4.2	2.7	6.5	2.4	1.7	82.5
3	0.6	3.9	9.3	7.2	1.8	1.9	75.4
4	0.7	3.7	17.9	5.7	1.2	2.5	68.2
8	1.2	1.4	38.9	3.2	1.4	2.6	51.3
12	0.8	1.2	45.2	4.7	1.5	3.3	43.3
16	0.7	1.3	47.1	5.7	1.3	4.1	39.8
20	0.6	1.3	47.4	6.4	1.2	4.5	38.6

Note: Data is from 1980Q1 until 2006Q2. 4 lags in VAR, Cholesky decomposition as ordered in table. Variables are: 1) FFR: monetary policy shocks, 2) Tech: technology shocks identified a la Gali (1999), 3) Info: information shocks a la Beaudry and Portier (2006), 4) Confidence: Confidence shocks a la Barsky and Sims (2008), 5) Fiscal: fiscal shocks from Romer and Romer (2007), 6) Oil: oil price shocks from Hamilton (1996), 7) UE is the unemployment rate. See Appendix B for details on construction of shock measures.

Figure 1: Plots of Survey Measures of Inflation Forecasts

Panel A: Mean Forecasts of Inflation over Next Year

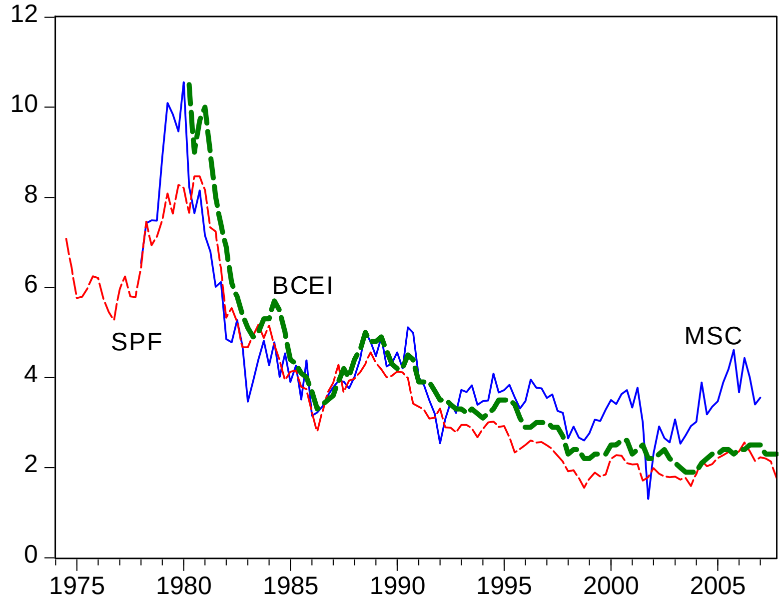

Panel B: (Log) Cross-sectional Standard Deviation of Inflation Forecasts over Next Year

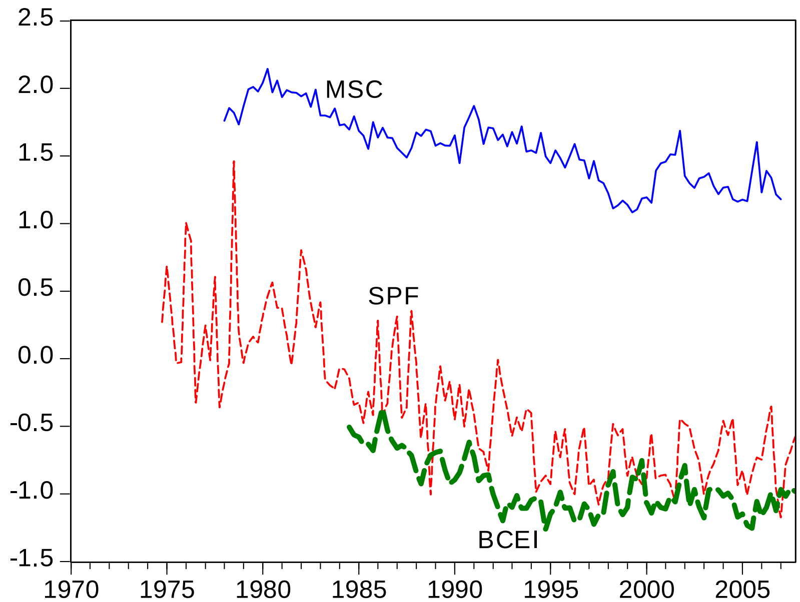

Figure 2: Impulse Response of Inflation to Shocks and Predicted Responses of Conditional Forecast Errors and Forecast Dispersion to Shocks

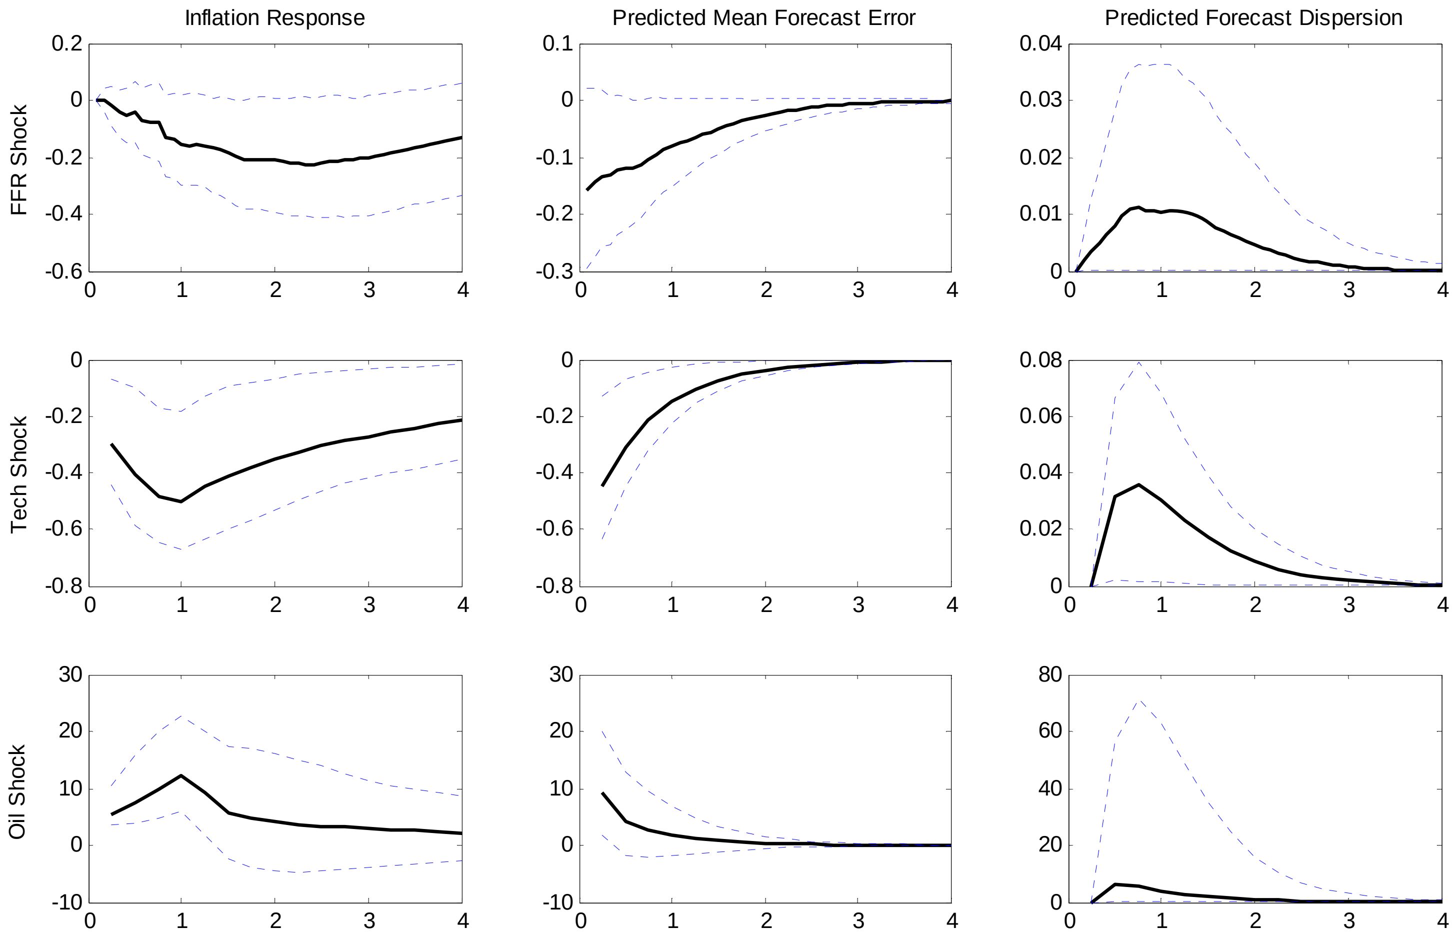

Note: Horizontal axis shows time in years. Inflation responses to oil and technology shocks are based on quarterly data. Inflation response to the FFR shock is based on monthly data. Technology shocks are identified using Gali (1999) long run restrictions. FFR shocks are identified recursively from a five variable VAR with inflation, unemployment rate, inflation in the commodity price index, index of industrial production and fed funds rate. Oils shocks are taken from Hamilton (1996). Predicted responses are from a sticky information model. We assume that agents update their information once a year on average, following Mankiw and Reis (2002), which implies λ=0.75. See section 3 for details.

Figure 3: Response of Mean Forecasts to Baseline Shocks

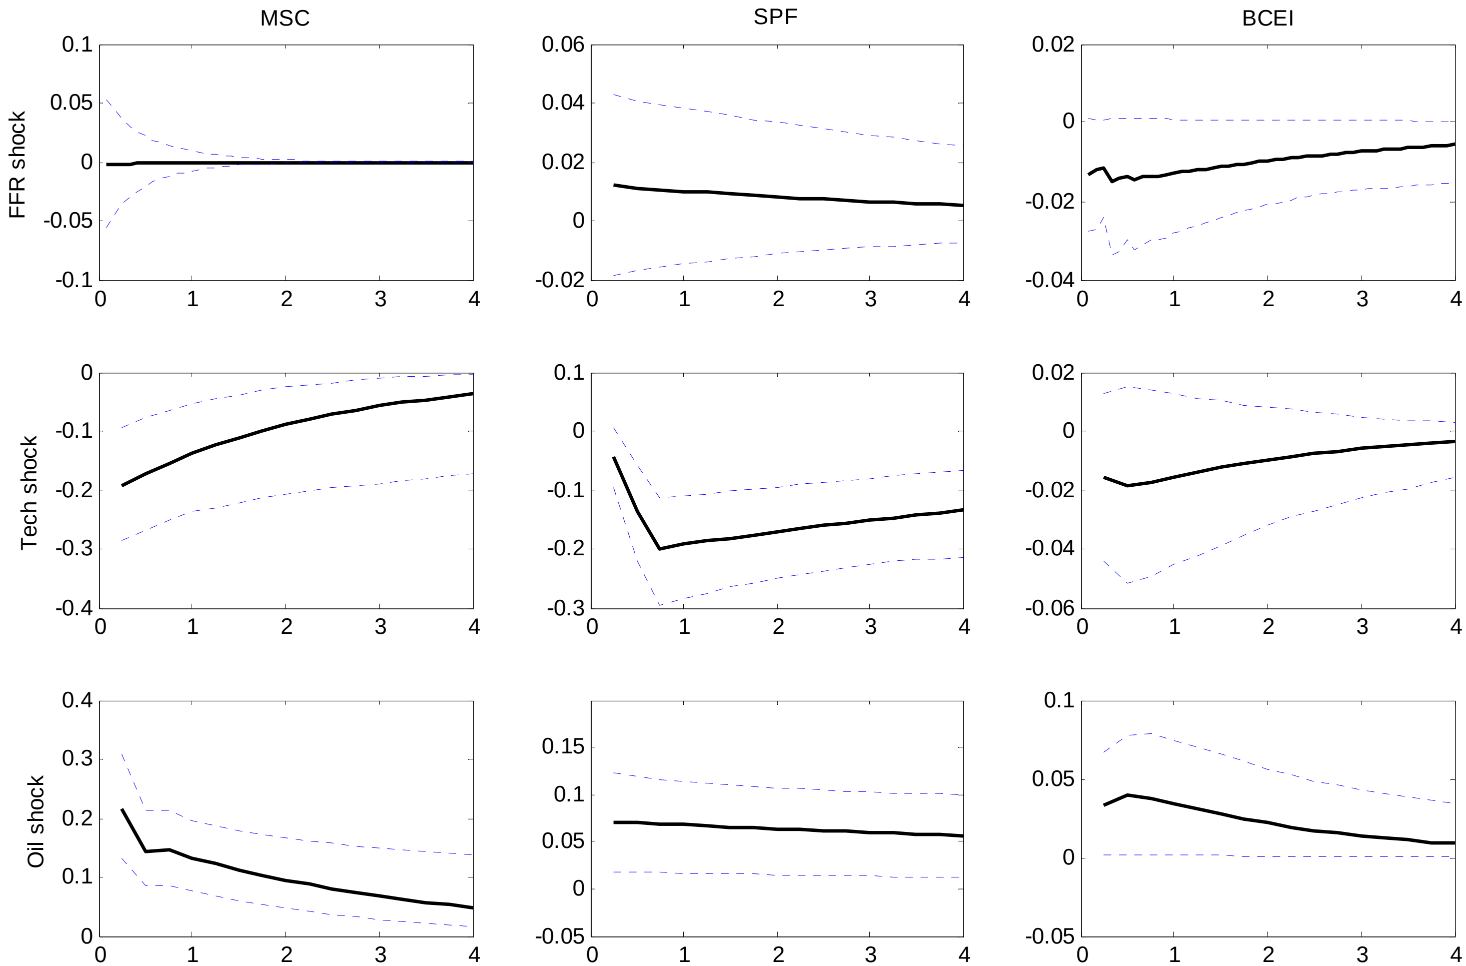

Note: Horizontal axis shows time in years. Technology shocks are identified using Gali (1999) long run restrictions. FFR shocks are identified recursively from a five variable VAR with inflation, unemployment rate, inflation in the commodity price index, index of industrial production and fed funds rate (monthly data). Oils shocks are taken from Hamilton (1996). Responses of MSC and BCEI forecasts to FFR shocks are based on monthly data. Responses in all other panels are based on quarterly data. Dashed lines are 95% confidence intervals. See section 4 for details.

Figure 4: Response of Forecast Errors to Baseline Shocks

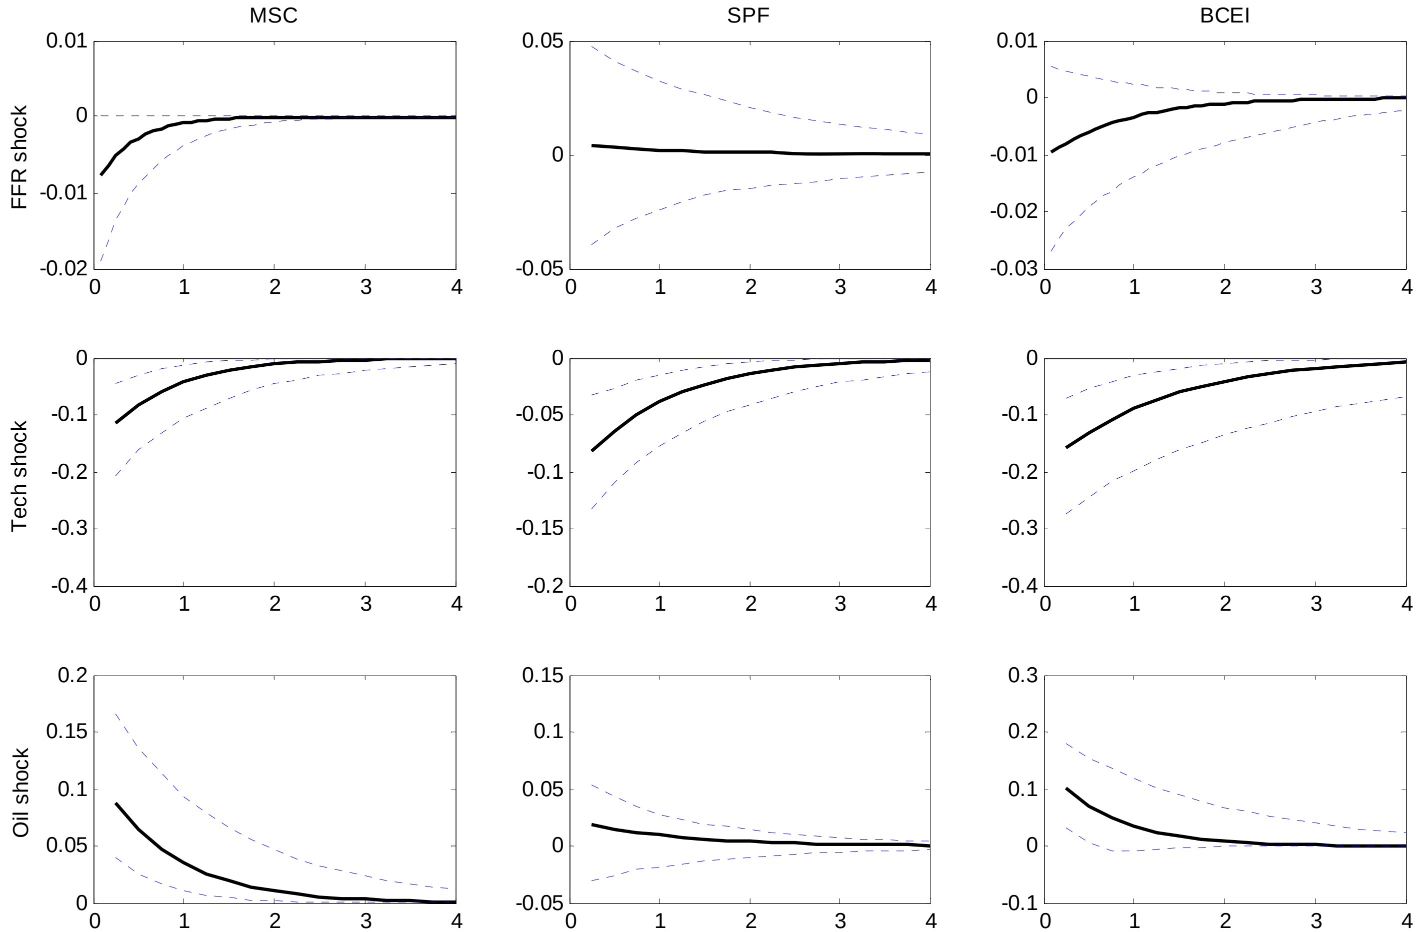

Note: Horizontal axis shows time in years. Technology shocks are identified using Gali (1999) long run restrictions. FFR shocks are identified recursively from a five variable VAR with inflation, unemployment rate, inflation in the commodity price index, index of industrial production and fed funds rate (monthly data). Oils shocks are taken from Hamilton (1996). Responses of MSC and BCEI forecasts to FFR shocks are based on monthly data. Responses in all other panels are based on quarterly data. Dashed lines are 95% confidence intervals. Responses of forecast errors begin one year after innovation. See section 4.1 for details.

Figure 5: Response of Forecast Dispersion to Baseline Shocks

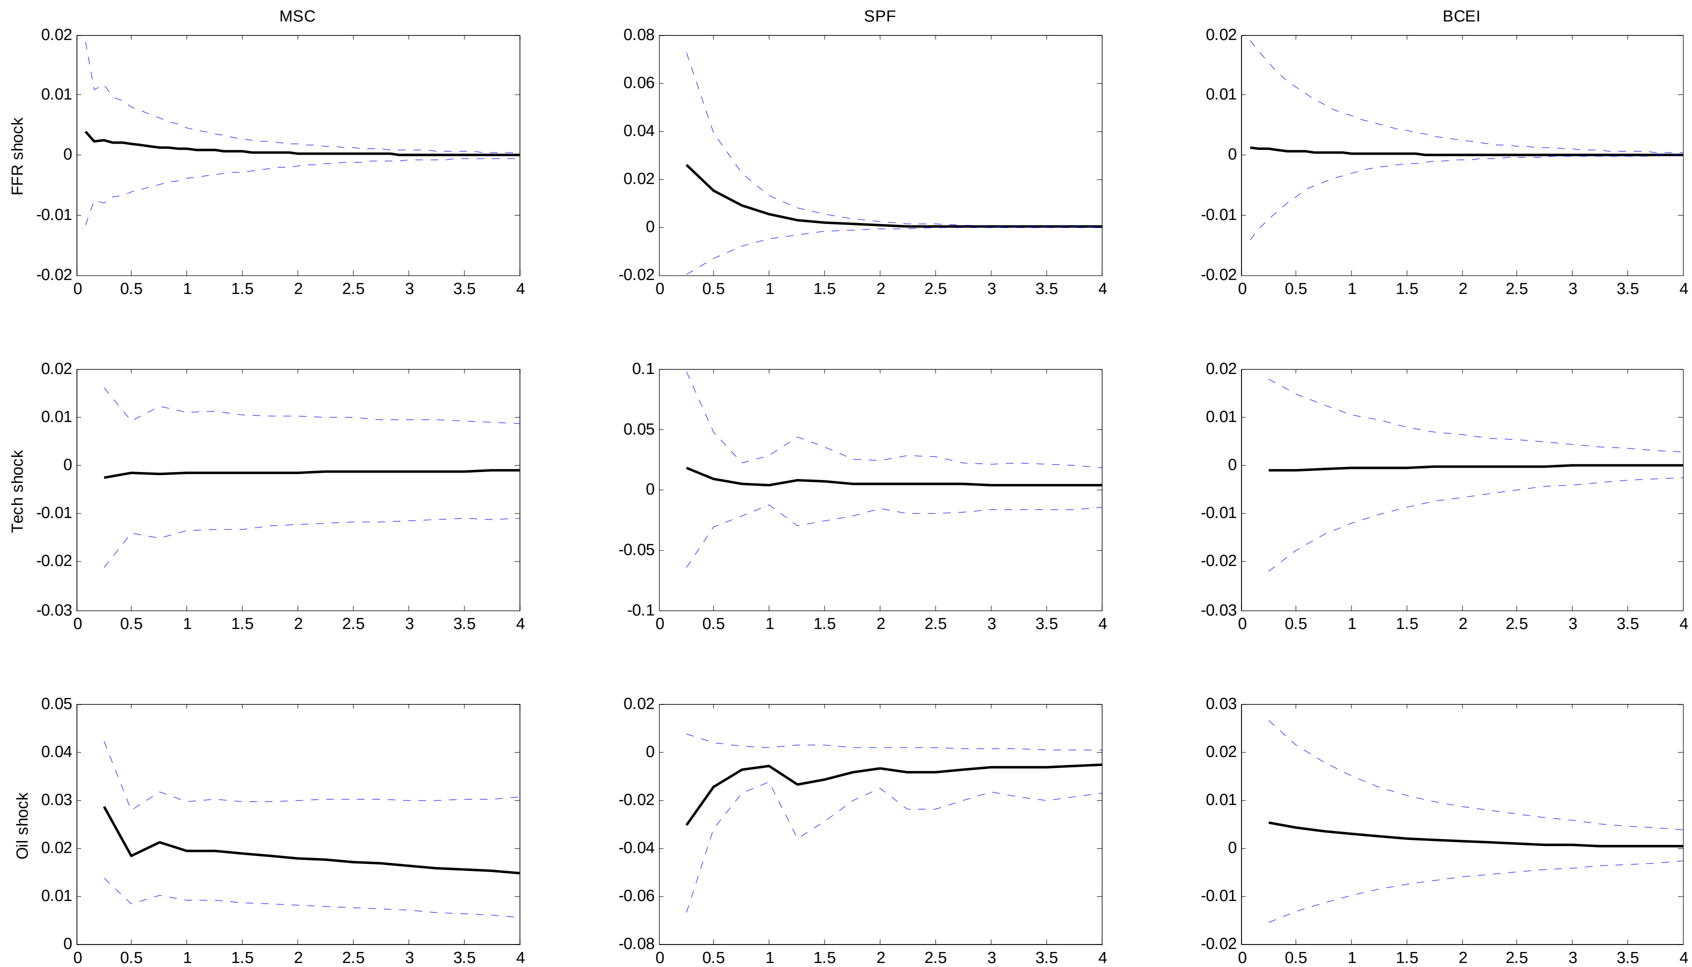

Note: Horizontal axis shows time in years. Technology shocks are identified using Gali (1999) long run restrictions. FFR shocks are identified recursively from a five variable VAR with inflation, unemployment rate, inflation in the commodity price index, index of industrial production and fed funds rate (monthly data). Oils shocks are taken from Hamilton (1996). Responses of MSC and BCEI forecasts to FFR shocks are based on monthly data. Dashed lines are 95% confidence intervals. Responses in all other panels are based on quarterly data. See section 4.3 for details.

Figure 6: Monte Carlo Simulations of Conditional Forecast Error Responses to Innovations

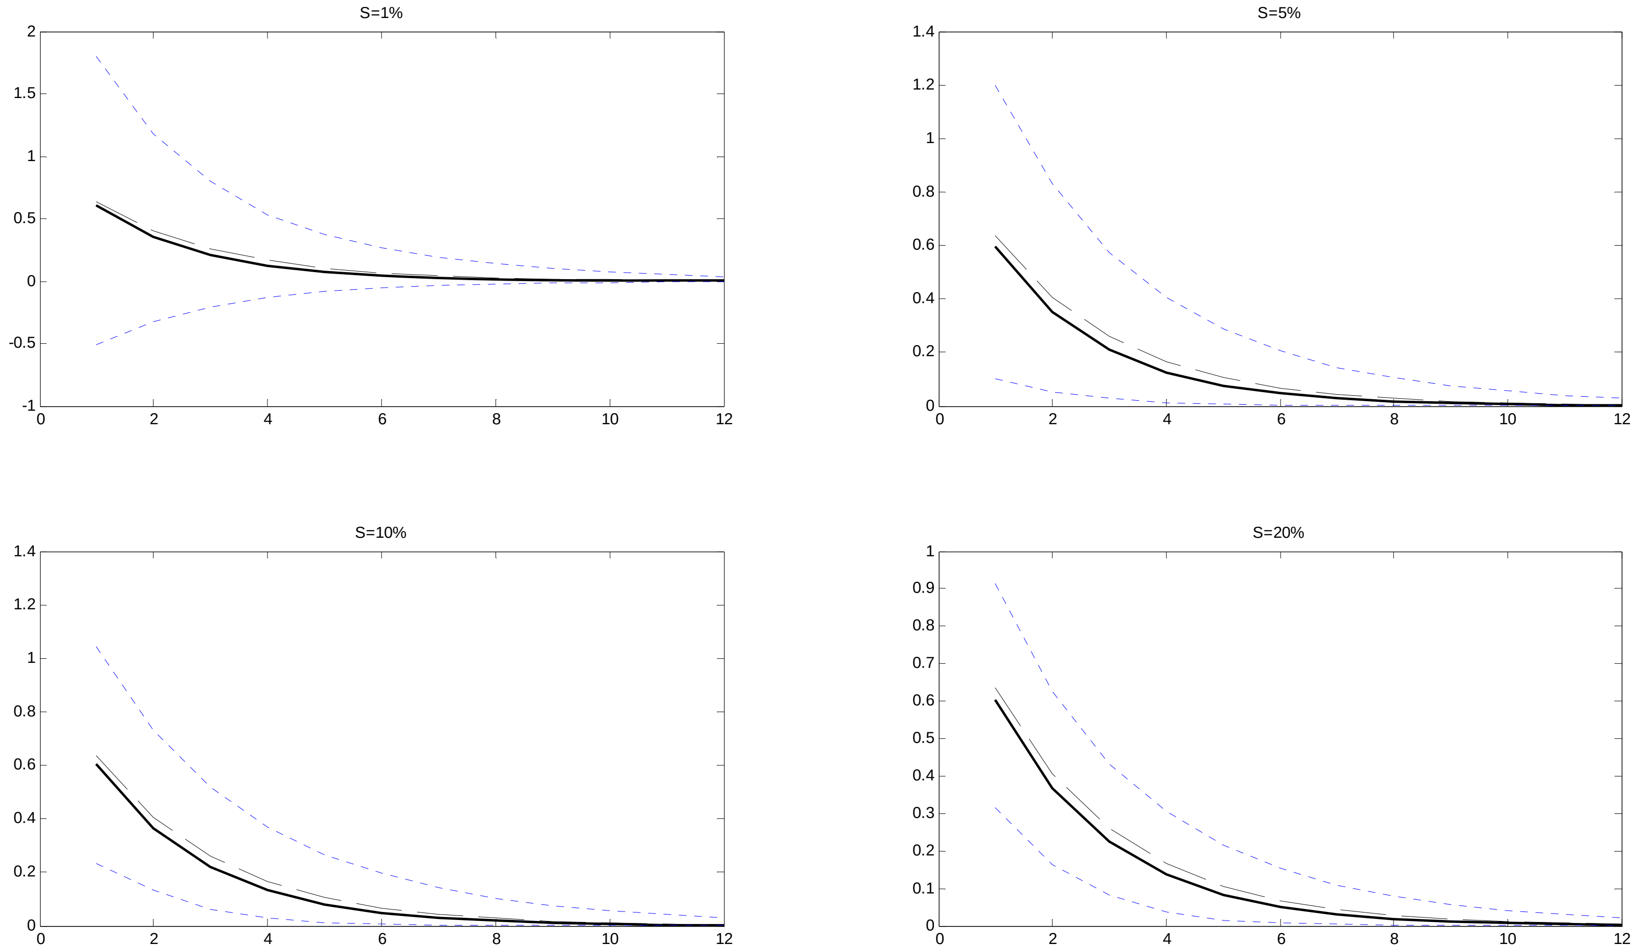

Note: Each figure illustrates the median (black bold line) estimated conditional response of forecast errors to an innovation to the model described in section 4.4 from 1,000 simulations. The true response is the dashed black line, while the blue dotted lines indicate the 95% confidence interval from the distribution of impulse responses. *S* is the fraction of the inflation variance accounted for by the innovation used in the estimation.

Figure 7: Monte Carlo Simulations of Conditional Forecast Dispersion Responses to Innovations

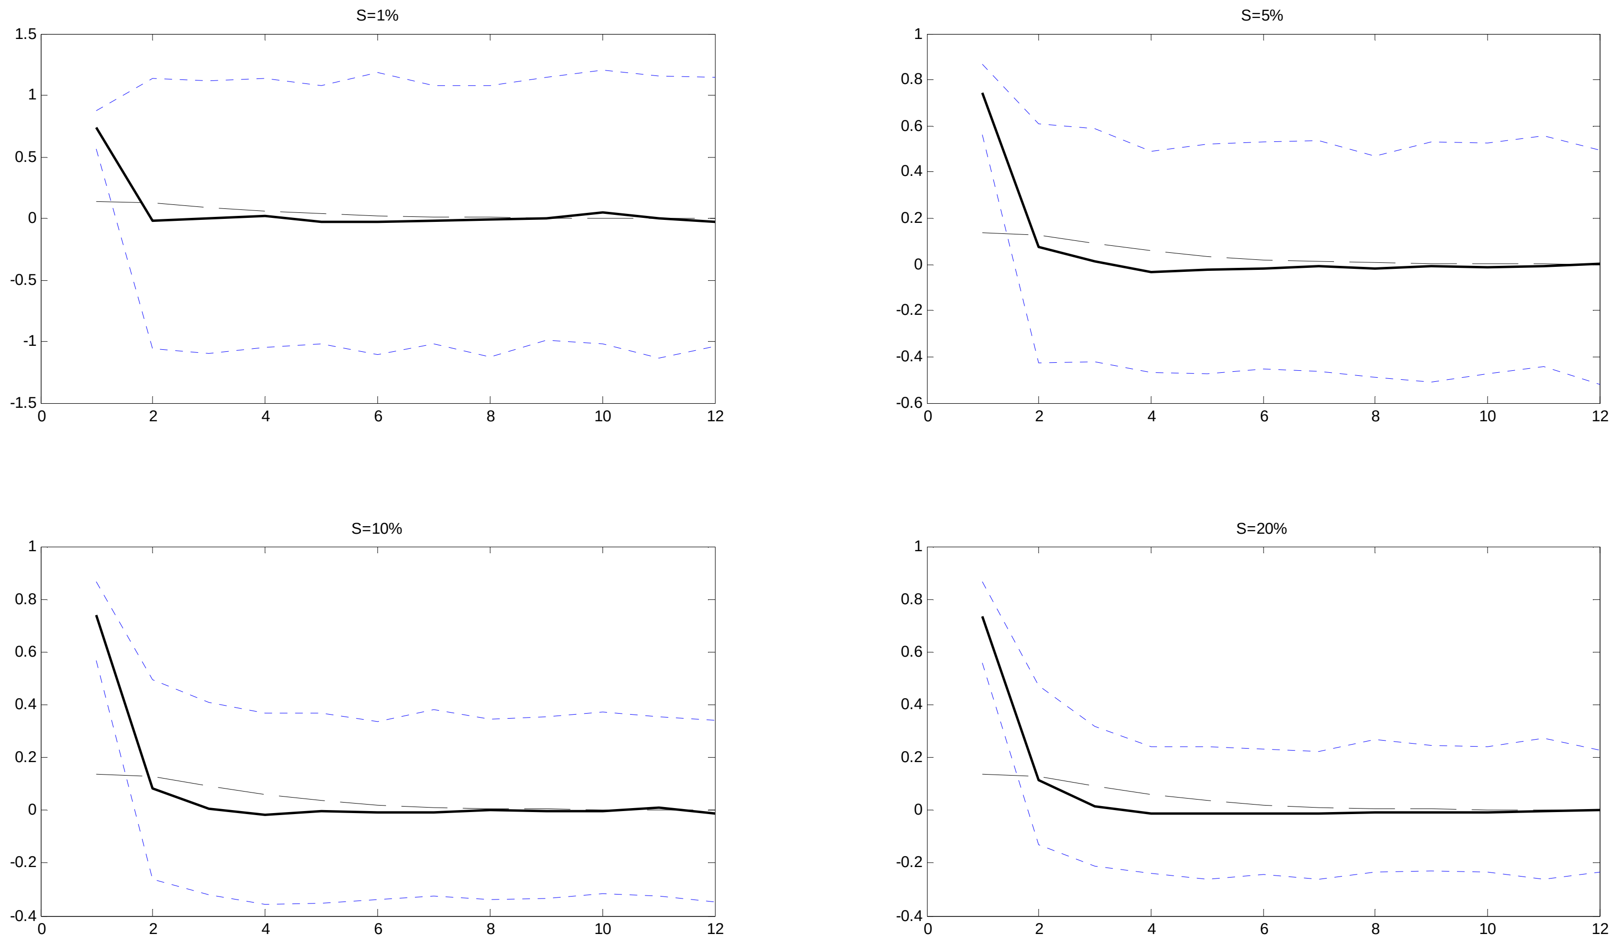

Note: Each figure illustrates the median (black bold line) estimated conditional response of forecast dispersion to the absolute value of an innovation to the model described in section 4.4 from 1,000 simulations. The true response is the dashed black line, while the blue dotted lines indicate the 95% confidence interval from the distribution of impulse responses. *S* is the fraction of the inflation variance accounted for by the innovation used in the estimation.

Figure 8: Response of forecast errors to baseline shocks using a VAR

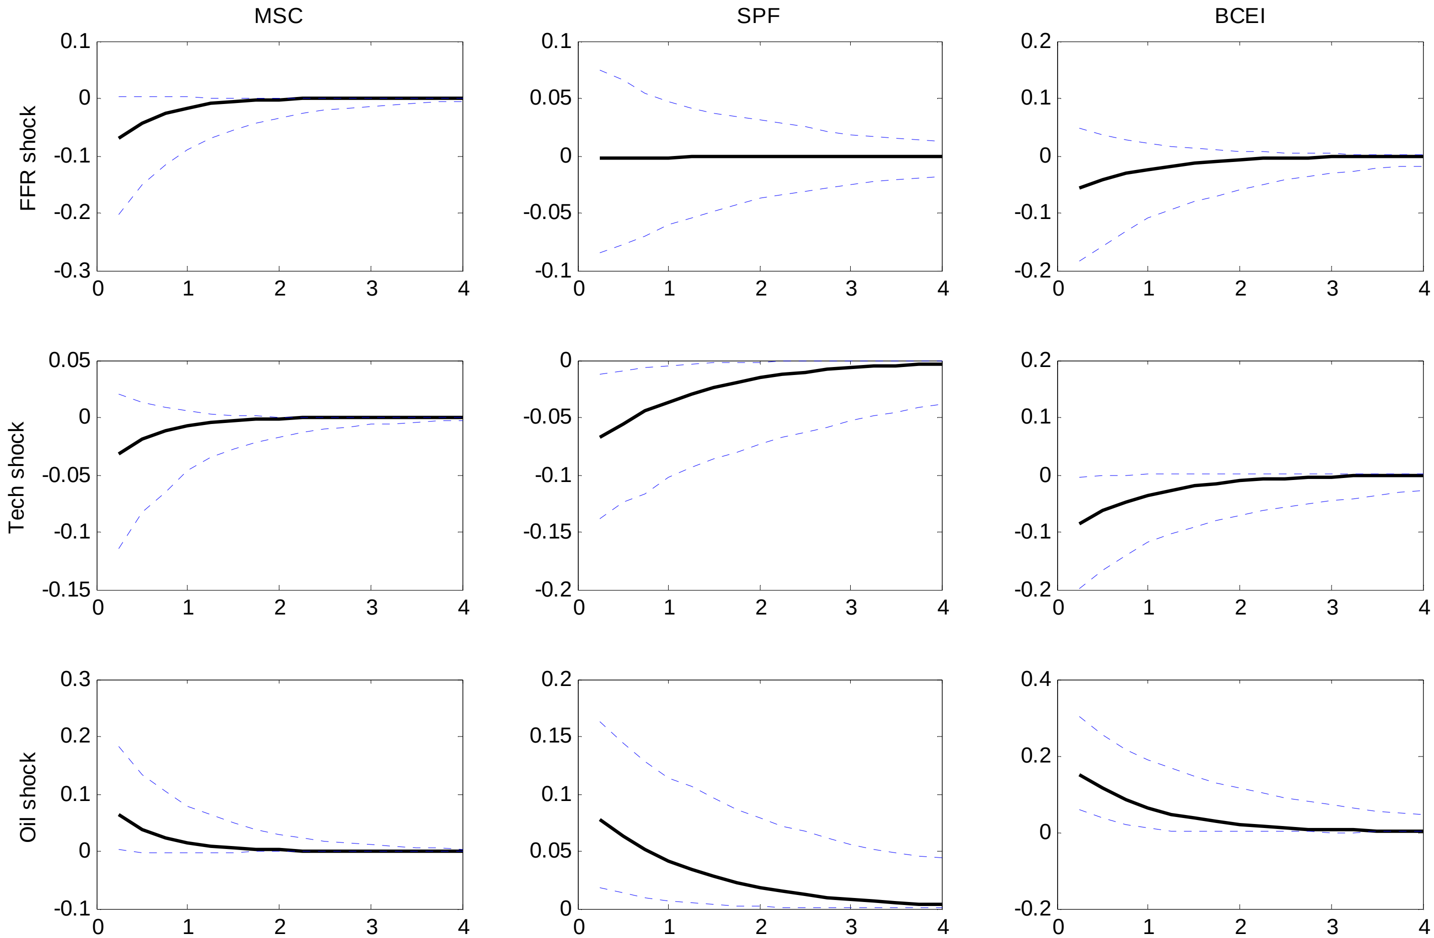

Note: Horizontal axis shows time in years. Technology shocks are identified using Gali (1999) long run restrictions. FFR shocks are identified recursively from a five variable VAR which includes inflation, unemployment rate, inflation in the commodity price index, real GDP and federal funds rate. Oils shocks are taken from Hamilton (1996). Impulse responses are from a 4-variable VAR with (absolute value) of shock measures ordered as in graph followed by the forecast error measure. Impulse responses start one year after each innovation. Dashed lines are 95% confidence intervals. Data is quarterly from 1984Q1 to 2007Q3. See section 5.1 for details.

Figure 9: Response of forecast dispersion to baseline shocks using a VAR

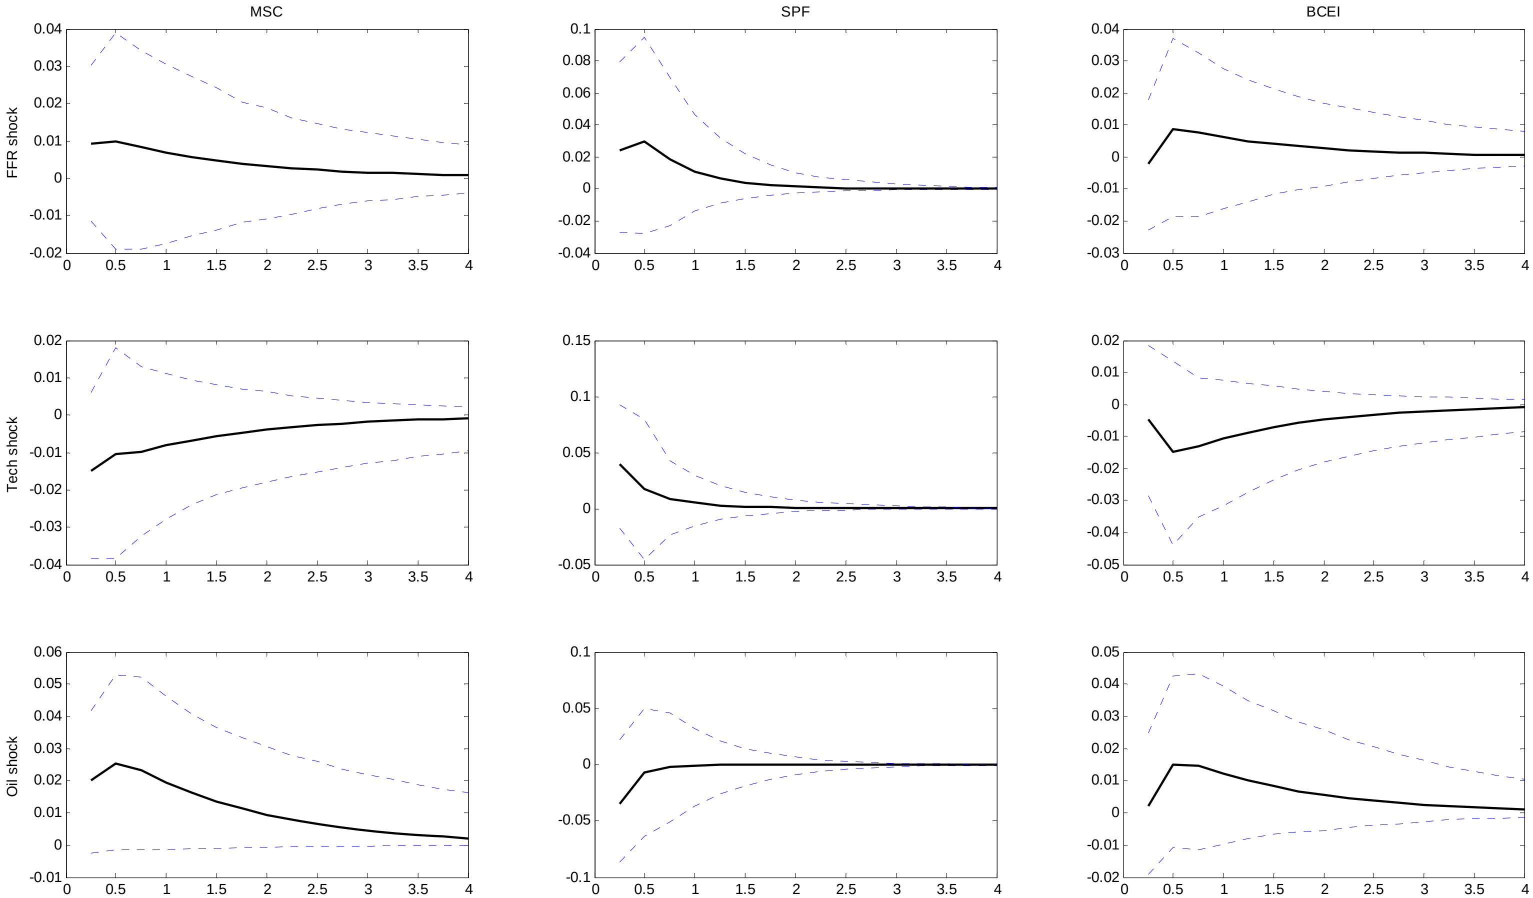

Note: Horizontal axis shows time in years. Technology shocks are identified using Gali (1999) long run restrictions. FFR shocks are identified recursively from a five variable VAR which includes inflation, unemployment rate, inflation in the commodity price index, real GDP and federal funds rate. Oils shocks are taken from Hamilton (1996). Impulse responses are from a 4-variable VAR with (absolute value) of shock measures ordered as in graph followed by the forecast dispersion measure. Dashed lines are 95% confidence intervals. Data is quarterly from 1984Q1 to 2007Q3. See section 5.1 for details.

Figure 10: Response of Forecast Errors to Alternative Shocks

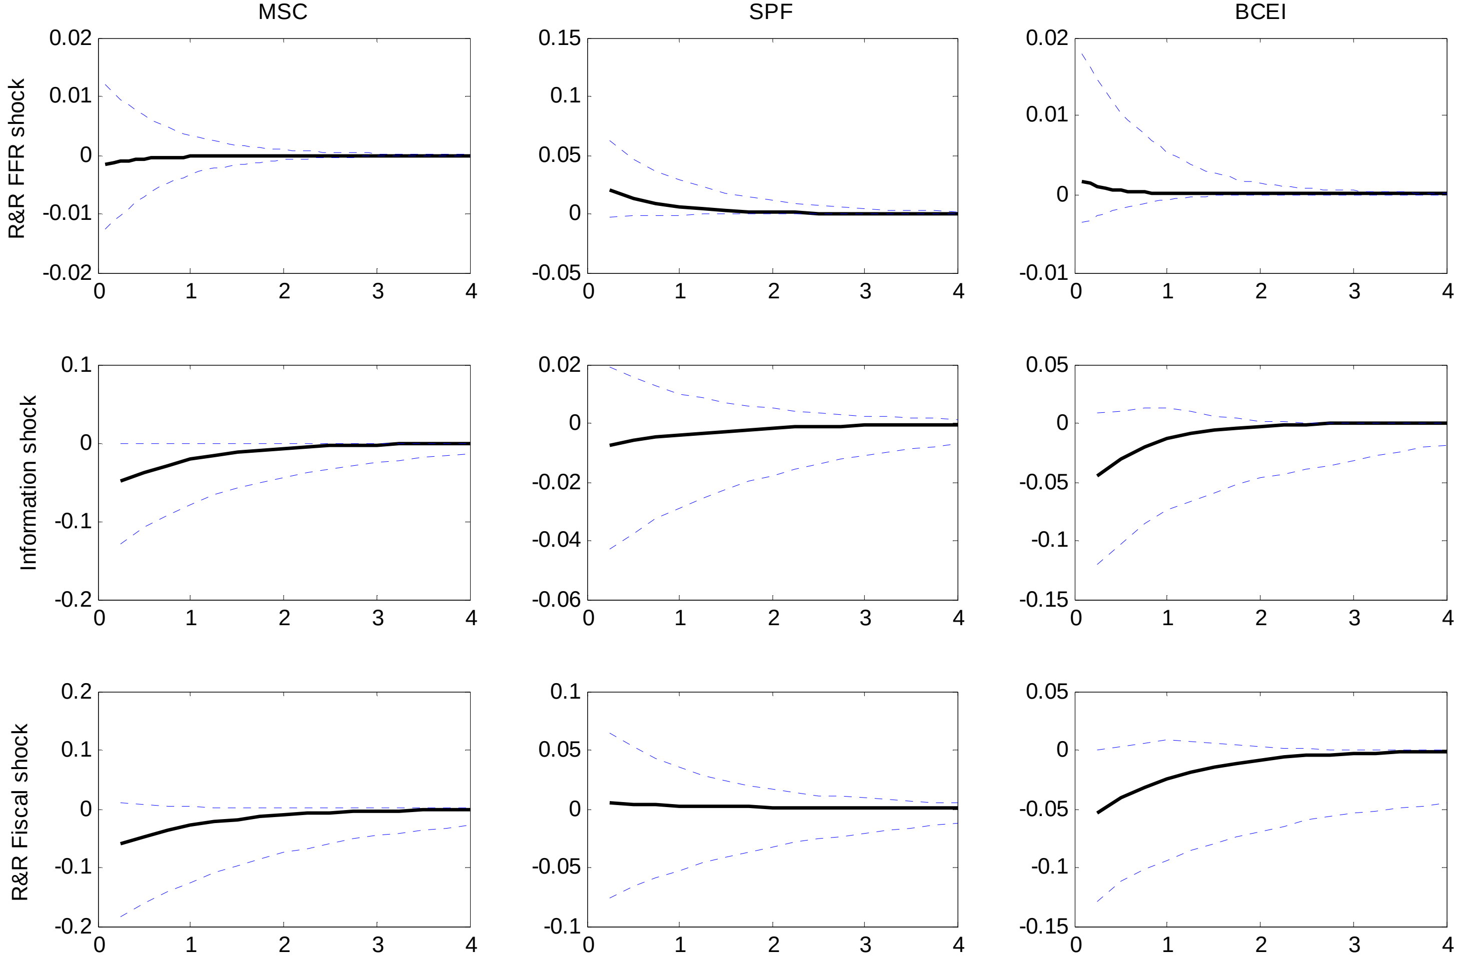

Note: Horizontal axis shows time in years. Information shocks are identified using Beaudry and Portier (2006) long run restrictions. FFR shocks are taken from Romer and Romer (2004). Fiscal shocks are taken from Romer and Romer (2007). Responses of MSC and BCEI forecasts to FFR shocks are based on monthly data. Dashed lines are 95% confidence intervals. Responses in all other panels are based on quarterly data. See section 5.2 for details.

Figure 11: Response of Forecast Dispersion to Alternative Shocks

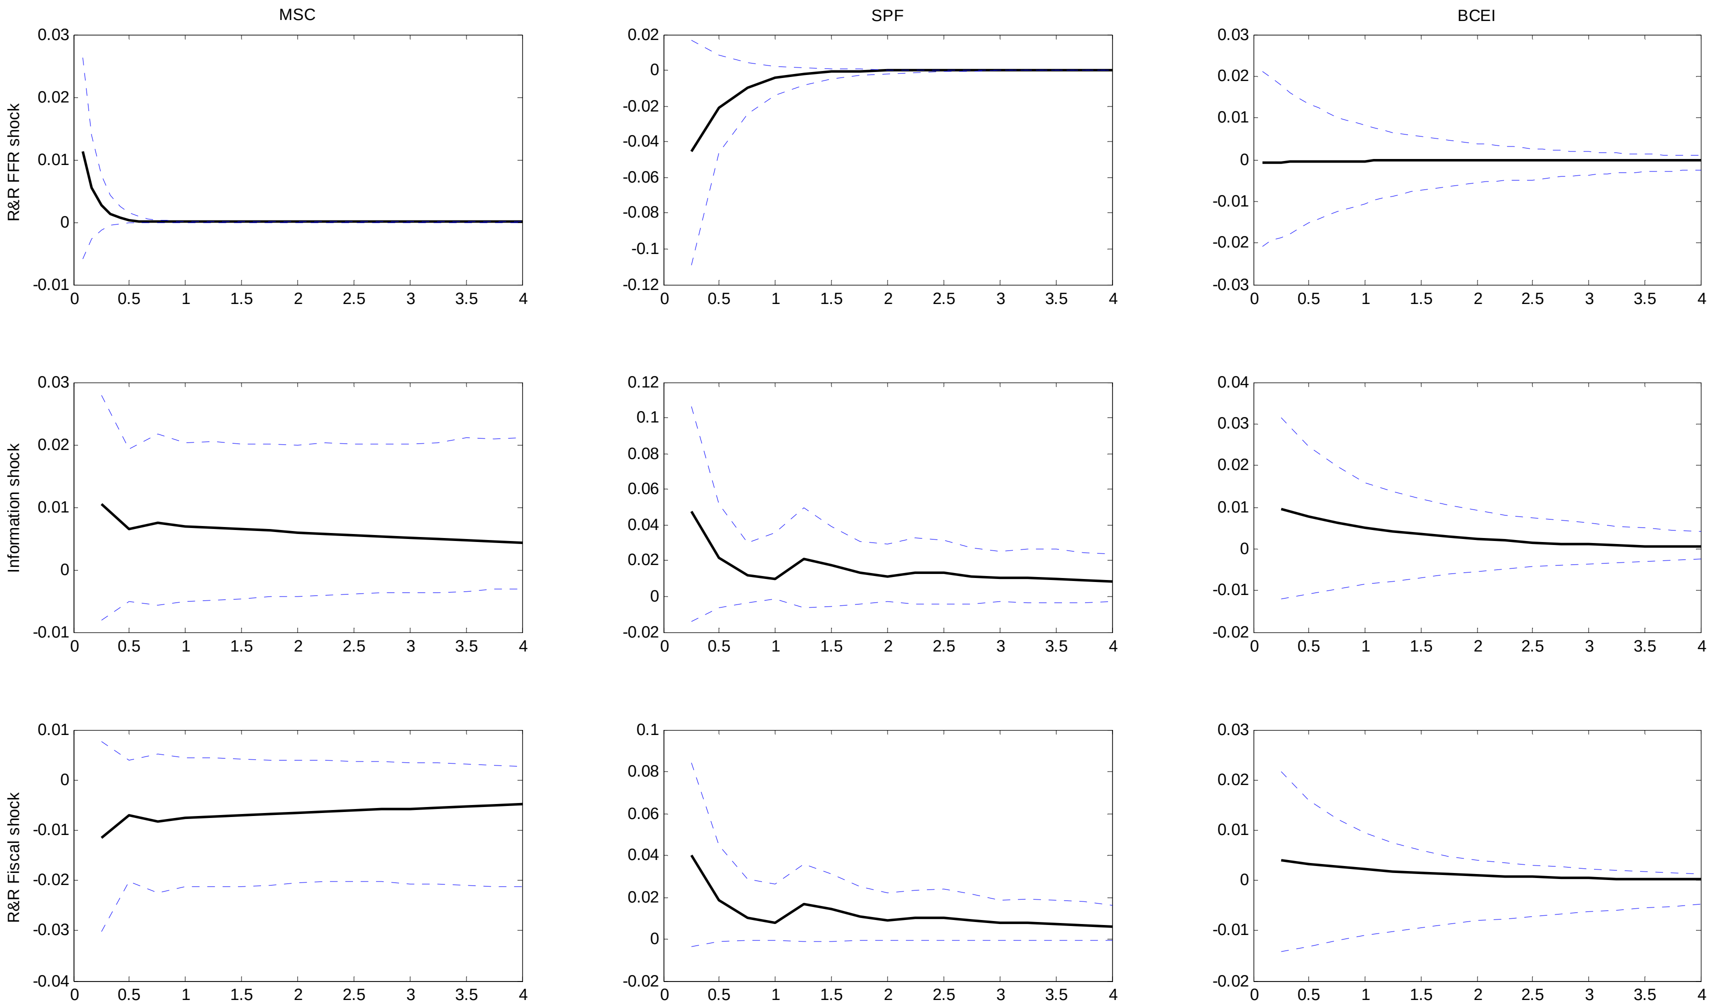

Note: Horizontal axis shows time in years. Information shocks are identified using Beaudry and Portier (2006) long run restrictions. FFR shocks are taken from Romer and Romer (2004). Fiscal shocks are taken from Romer and Romer (2007). Responses of MSC forecasts to FFR shocks are based on monthly data. Dashed lines are 95% confidence intervals. Responses in all other panels are based on quarterly data. See section 5.2 for details.

Figure 12: Response of forecast errors and forecast dispersion to composite shock

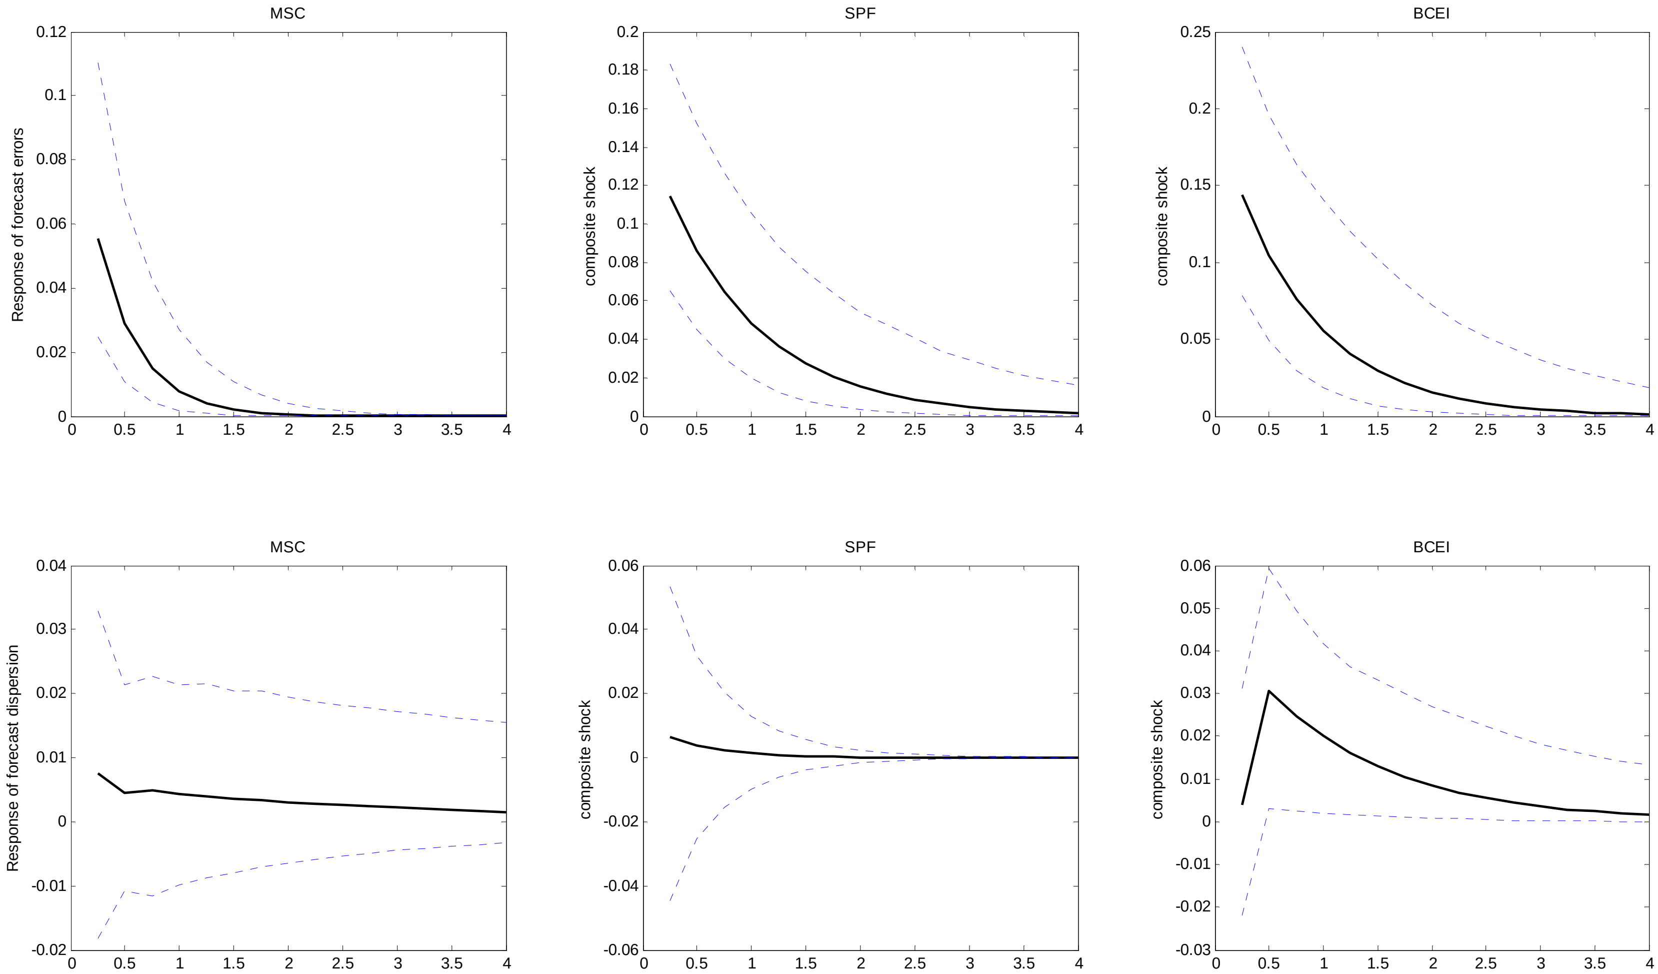

Note: Horizontal axis shows time in years. The composite shock is the predictable component of inflation residuals from AR(4) when regressed on monetary policy, technology, and oil price shocks. Response of forecast errors begins one year after the shock. Dashed lines are 95% confidence intervals. Responses in all panels are based on quarterly data. See section 5.3 for details.

Figure 13: Response of forecast errors and forecast dispersion to unidentified inflation innovations

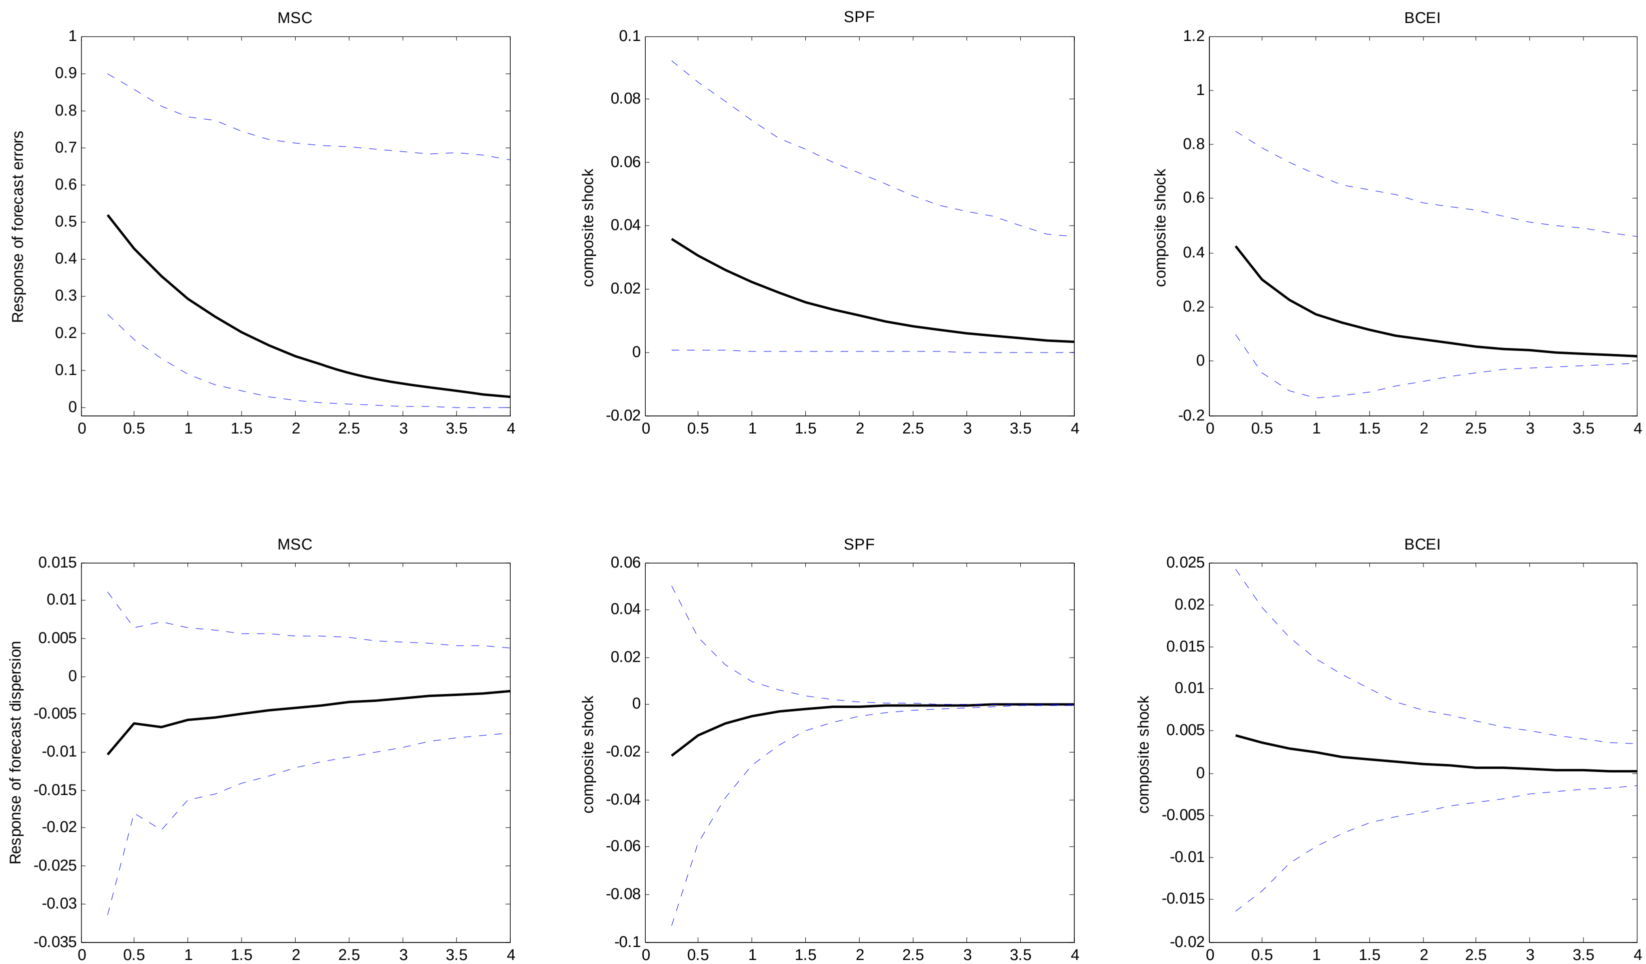

Note: Horizontal axis shows time in years. Unidentified inflation innovations are the unpredictable component of inflation residuals from AR(4) when regressed on monetary policy, technology, and oil price shocks. Response of forecast errors begins one year after the shock. Dashed lines are 95% confidence intervals. Responses in all panels are based on quarterly data. See section 5.4 for details.

Figure 14: Response of UE Forecast Errors and UE Forecast Dispersion to Shocks

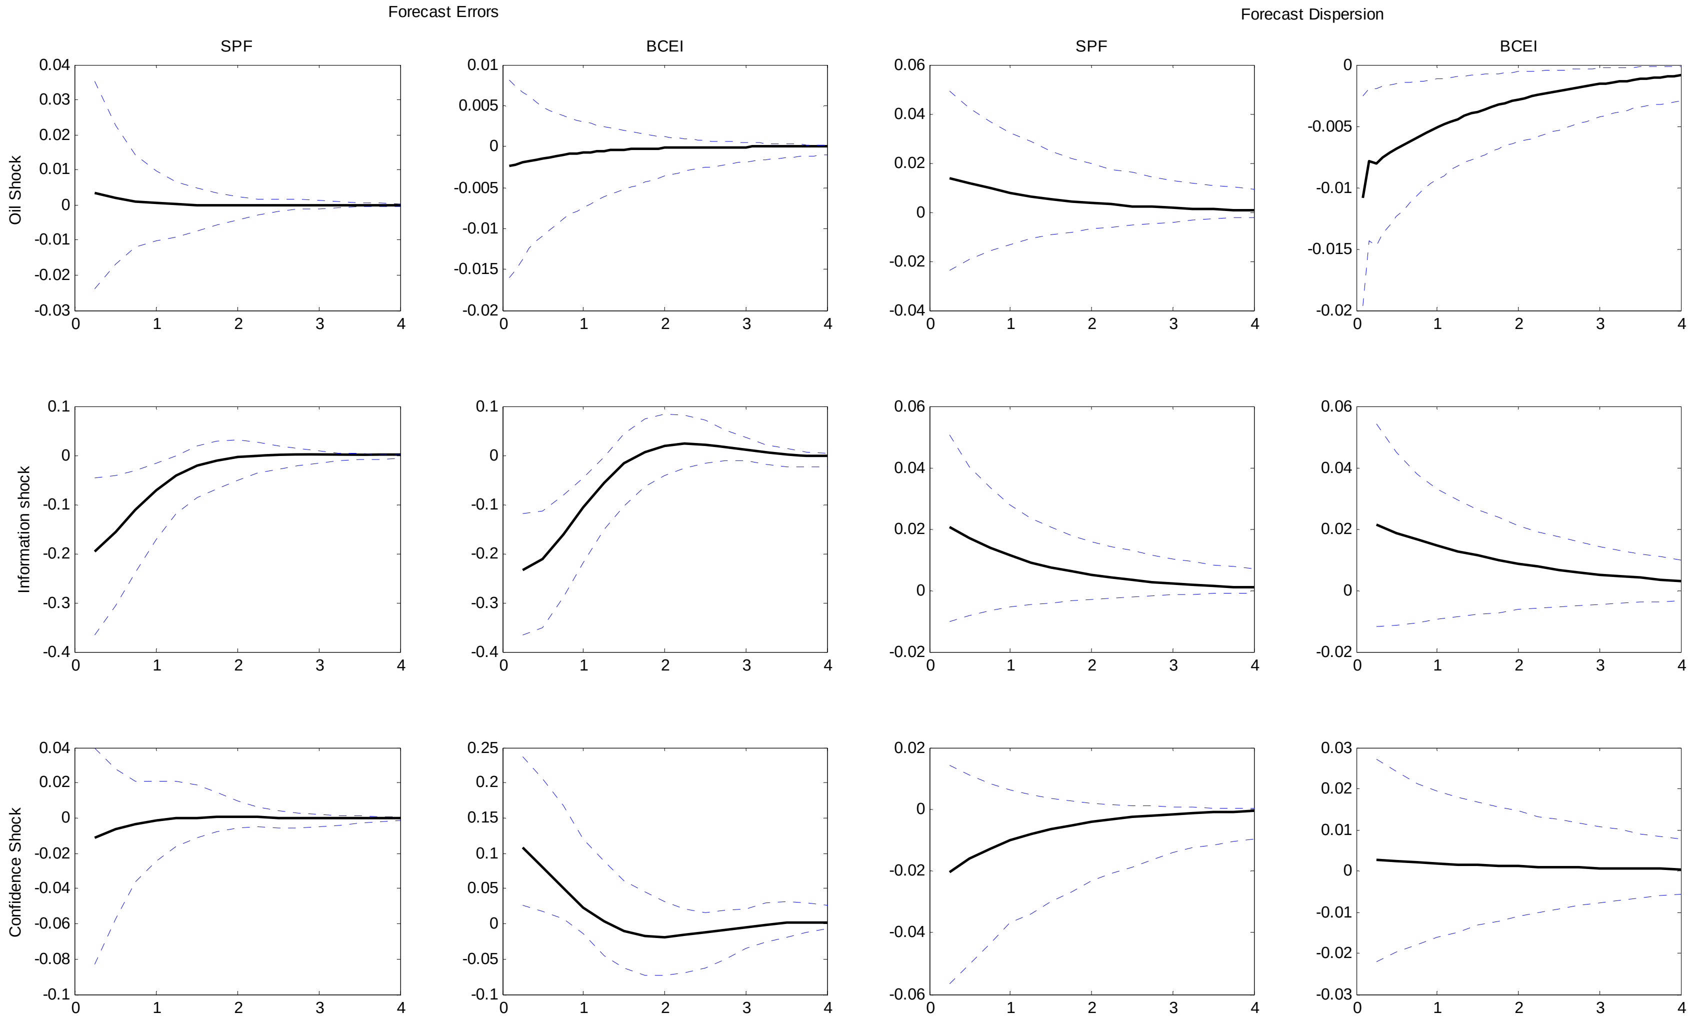

Note: Horizontal axis shows time in years. Oil price shocks are from Hamilton (1996), Information shocks are identified using Beaudry and Portier (2006) long run restrictions and confidence shocks follow Barsky and Sims (2008). Response of forecast errors begins one year after the innovation. See section 5.5 for details.