



Testing the Sticky Information Phillips Curve

Olivier Coibion^{*}
College of William and Mary

College of William and Mary
Department of Economics
Working Paper Number 61

October 2007

^{*} I am grateful to Bob Barsky, Michael Elsby, Yuriy Gorodnichenko, Chris House, Ed Knotek, Oleg Korenok, N. Gregory Mankiw, Peter Morrow, Julio Rotemberg, Matthew Shapiro, Clemens Sialm, Eric Sims for very helpful comments as well as the Robert V. Roosa Dissertation Fellowship for financial support. This paper was previously distributed under the title “Empirical Evidence on the Sticky Information Phillips Curve.”

Testing the Sticky Information Phillips Curve

Abstract

I consider the empirical evidence for the sticky information model of Mankiw and Reis (2002) relative to the basic sticky price model, conditional on historical measures of inflation forecasts. Overall, the evidence is unfavorable to the sticky information model of price-setting: the estimated structural parameters are inconsistent with an underlying sticky information model and the sticky-information Phillips Curve is statistically dominated by the New Keynesian Phillips Curve. I find that the poor performance of the sticky information approach is driven by two key elements. First, predicted inflation in the sticky information model places substantial weight on old forecasts of inflation. Because these consistently underestimate inflation in the 1970s and overestimate inflation since the 1980s, particularly at long forecast horizons, predicted inflation from the sticky information model inherits these patterns. Second, predicted inflation from the sticky information model is excessively smooth.

JEL Codes: E30, E37

Keywords: Sticky Information, Expectations, Inflation

Olivier Coibion
College of William and Mary
Williamsburg, VA 23187-8795
ocoibion@wm.edu

1 Introduction

Empirical research on the response of the economy to monetary policy shocks has identified important stylized facts that can be used to differentiate among competing models.¹ The delayed response of inflation to such shocks has been of primary interest, because the basic sticky price model cannot replicate this key feature of the data.² As a result, much recent research has been devoted to developing models that can match this stylized fact. While some of this work has been done within the context of sticky price models, others have proposed dropping the assumption of sticky prices entirely and focusing instead on informational rigidities.³ A leading example is the sticky information model of Mankiw and Reis (2002), in which firms update their information infrequently according to a time dependent process but are free to change prices at all times.

The gradual diffusion of information across the population, which is the key assumption in the sticky information model, has received some empirical support. For example, Carroll (2003) estimates the rate of diffusion of information from professional forecasters to the general population from an epidemiological model and finds results in line with those assumed by Mankiw and Reis. Dopke et al (2005) provide similar support for the diffusion of information from forecasters to households in European countries. Mankiw and Reis (2003) estimate a sticky information model applied to wage setting and find that the average wage setter updates his information about once per year. Khan and Zhu (2006) directly estimate the structural parameters of the sticky information model applied to price setting and conclude that the evidence is not inconsistent with firms updating their information approximately once a year. Klenow and Willis (2007) find micro-level evidence consistent with firms responding to old information in price-setting decisions.

¹ See Christiano, Eichenbaum and Evans (1999).

² Mankiw (2001) emphasizes this point. Other failures of the basic sticky price model include predicting costless deflations, economic booms under pre-announced credible disinflations (Ball 1994), and failing to reproduce the positive correlation between changes in inflation and the level of output.

³ In the context of sticky price models, Galí and Gertler (1999) add rule-of-thumb firms to a sticky price model and derive a hybrid New Keynesian Phillips Curve. Trabandt (2005) shows that such a model can yield a delayed response of inflation to monetary policy shocks. Christiano et al (2005) allow sticky price firms to index their prices to some measure of inflation in non-reoptimizing periods. Calvo et al (2003) allow sticky price firms to choose a reset price and a rate at which prices will be automatically increased.

The empirical approach used in this paper follows the distinction drawn by Carroll (2003) between the formation of expectations by professional forecasters and the diffusion of those forecasts to the population. Sticky information is interpreted as a gradual (time-dependent) diffusion of forecasts from professional forecasters to firms combined with the additional assumption that price changes are costless. The Sticky Information Phillips Curve (SIPC) then gives the relationship between inflation and the output gap conditional on past forecasts of the current state. One can similarly derive the New Keynesian Phillips Curve (NKPC) conditional on professional forecasts by assuming a time-dependent process for changing prices with the additional assumption of costless acquisition of professional forecasts. The purpose of writing these models conditional on inflation forecasts is to separate two issues: whether forecasts are consistent with rational expectations and whether price-setting decisions should best be modeled as sticky-prices or sticky-information. Previous work on assessing the validity of these models has imposed rational expectations on the forecasts, making the empirical exercise a joint test of price-setting decisions and rational expectations. The approach taken here allows us to separate the two and focus on the question of the validity of each price-setting assumption conditional on observed historical forecasts.

This distinction proves to have important implications for the empirical results. Whereas previous work has found little evidence strongly favoring sticky prices or sticky information, the results of this paper are strongly at odds with the sticky information assumptions. Using historical survey measures of inflation forecasts, the estimated structural parameters of the Sticky Information Phillips Curve point to no statistically significant degree of information rigidity nor is there a discernible link between the nominal side and the real side of the economy. The sticky information Phillips Curve, since the 1970s, can thus be rejected on structural grounds. Second, the SIPC is also strongly rejected statistically in favor of the NKPC, the very model that it was supposed to replace.

I show that this rejection of the Sticky Information Phillips Curve is due to two elements. The first is a real-time forecast error effect. Professional forecasters consistently underestimated inflation in the 1970s but overestimated inflation in the 1980s and 1990s. This feature of forecasts is increasingly true at longer forecast horizons. Because the sticky

information Phillips Curve places an important weight on older forecasts of current inflation, this leads to predicted inflation being too low in the 1970s and too high since the 1980s. The real-time forecast error effect plays an important role in explaining why the estimated degree of information rigidity is close to zero. Importantly, this effect is absent when one uses in-sample forecasts, as implicitly done in Dupor et al (2006), Kiley (2006), Korenok (2004), and Korenok et al (2006). Thus, whereas previous work has demonstrated that relying on historical inflation forecasts helps the New Keynesian Phillips Curve empirically (see Roberts (1995) and (1997)), I show that it impairs the ability of the SIPC to match the data.

A second contribution of the paper is to identify another implication of the SIPC at odds with the data which I refer to as the inflation inertia effect: predicted inflation from the SIPC using the preferred parameter estimates of Mankiw and Reis is excessively persistent and insufficiently volatile. This result, unlike the real-time forecast error effect, is robust to the forecasts used. The basic sticky price model, on the other hand, comes much closer to matching both the persistence and volatility of inflation conditional on inflation forecasts and the output gap. This result is particularly surprising given the fact that the sticky information model was designed explicitly to account for inflation inertia missing from the sticky price model.

The paper also attempts to explain the fact that estimates of the degree of information rigidity from the SIPC are very sensitive to the time period. While the estimates over the whole sample point to no information rigidity at all, the sub-sample estimates using data since 1984 are consistent with firms acquiring new forecasts less than once a year on average, although the SIPC continues to be dominated by the NKPC statistically even in the subsample analysis. I argue that the high estimated levels of information rigidity are likely to be capturing the fact that forecast errors were highly predictable over this time period. Because the structural form of the SIPC is very similar to tests of the rationality of the forecasts, periods of predictable forecast errors can mistakenly lead one to conclude that there is a high level of sticky information when, in fact, there is none. I illustrate this using the sticky price and imperfect information model of Erceg and Levin (2003), which delivers a pattern of predictable forecast errors in subsamples similar to that observed in the data, even though there is no delay in the

diffusion of information from professional forecasters to firms, i.e. no sticky information in the model. I estimated the SIPC in Monte Carlo simulations of this model and closely replicate the empirical findings from the US data over the whole time sample as well as since the mid-1980s. These results indicate that a promising area of future research may be the role of imperfect information in the formation of expectations by professional forecasters rather than the gradual diffusion of information from professionals to firms and households.

The structure of the paper is as follows. Section 2 presents the econometric approach used to estimate the SIPC as well as the non-nested model tests. Section 3 presents and discusses the baseline results. Section 4 considers some robustness checks while section 5 contains a discussion and interpretation of the results. Section 6 concludes.

2 Empirical Approach

The goal of the paper is to evaluate the empirical support for the sticky information Phillips Curve relative to the basic sticky price model. I do this conditional on historical forecasts. Specifically, I first follow Carroll (2003) and assume that each quarter, professional forecasters generate a set of forecasts of macroeconomic variables denoted by $F_t[\bullet]$. There is a continuum of firms, each of which knows that its instantaneously optimal price is given by

$$p_t^{\#}(j) = p_t + \alpha x_t$$

where p_t is the aggregate price level, x_t is the output gap, and α is the degree of real rigidity.

In general, one could assume that both the acquisition of new forecasts and changing prices are costly. Instead, I will focus on the two extreme cases: the basic sticky information and sticky price models. Following Mankiw and Reis (MR henceforth), the sticky information model consists of two assumptions. First, the acquisition of new forecasts by firms follows a Poisson process in which there is a probability $1-\lambda$ that any given firm will acquire a new set of forecasts. Second, price changes are costless. Jointly, these two assumptions yield the SIPC

$$\pi_t = \frac{(1-\lambda)}{\lambda} \alpha x_t + (1-\lambda) \sum_{j=0}^{\infty} \lambda^j F_{t-1-j}(\pi_t + \alpha \Delta x_t)$$

which relates inflation to the output gap and past forecasts of current inflation and changes in the output gap.⁴

Alternatively, one can reverse the assumptions: the acquisition of forecasts is costless and immediate, whereas price changes are costly. Assuming a Poisson process for changing prices, in which $(1-\gamma)$ is the probability of changing prices each quarter, we get the NKPC

$$\pi_t = \frac{(1-\beta\gamma)(1-\gamma)}{\gamma} \alpha x_t + \beta F_t \pi_{t+1}$$

which relates inflation to the current output gap and the current forecast of future inflation.⁵ The key difference lies in the timing of the expectations in each Phillips Curve: the sticky-price model implies that the relationship between the nominal and the real side of the economy is conditional on current expectations of future inflation, whereas the sticky information model implies that past forecasts of the current state are the relevant measure of expectations in the Phillips Curve. This distinction reflects the alternative assumptions about the diffusion of information and the costliness of price changes underlying each model.

The purpose of writing these models conditional on inflation forecasts is to separate two issues: whether forecasts are consistent with rational expectations and whether price-setting decisions should best be modeled as sticky-prices or sticky-information. Previous work on assessing the validity of these models has imposed rational expectations on the forecasts, making the empirical exercise a joint test of price-setting decisions and rational expectations. The approach taken here allows us to separate the two and assess the validity of each price-setting assumption conditional on observed historical forecasts.

To assess the empirical support for the sticky information Phillips Curve, I will use two sets of criteria. The first will be whether estimation of the structural parameters of the SIPC yields values that are consistent with the theory of the model. The second will be to compare its performance statistically to the New Keynesian Phillips Curve. I consider first how to

⁴ See Caballero (1989) and Reis (2006) for microfoundations of the SIPC based on firms facing fixed costs to acquiring information. Note that MR impose the additional assumption that professional forecasters have rational expectations.

⁵ Deriving this is standard, but requires the assumption that forecasters know and impose in generating their forecasts the equation describing the price level. Specifically, it requires $F_t b_{t+1} = [F_t p_{t+1} - \gamma p_t] / (1-\gamma)$ where b_t is the optimal reset price for firms.

adequately estimate the structural parameters of the SIPC, then turn to the issue of assessing its validity relative to the sticky price model.

2.1 Estimating the Sticky Information Phillips Curve

To assess the empirical validity of the SIPC, I first augment the SIPC with an error term ε_t , assumed to be i.i.d.⁶

$$\pi_t = \frac{(1-\lambda)}{\lambda} \alpha x_t + (1-\lambda) \sum_{j=0}^{\infty} \lambda^j F_{t-1-j} (\pi_t + \alpha \Delta x_t) + \varepsilon_t \quad (1)$$

Estimating λ and α using equation (1) presents several difficulties. First, the output gap on the RHS will tend to be correlated with the error term. This endogeneity issue can typically be addressed by instrumental variables. However, the infinite amount of regressors on the RHS must be truncated, adding a source of error that will not be uncorrelated with lagged instruments. Therefore the identification condition that instruments be uncorrelated with the error term will typically fail. Second, other than the output gap, all variables on the RHS are past expectations of current values of aggregate inflation and changes in the output gap. While expectational terms in NKPC estimations are frequently replaced with ex-post values (e.g. Gali and Gertler (1999)), doing so in the SIPC would yield an error process that would be highly correlated with both regressors and instruments. It is thus critical to have actual measures of past forecasts as regressors. I address each of these points in turn.

2.1.1 Endogeneity, Instruments, and Truncation

Consistent estimation of the parameters of the SIPC requires an identification condition. Given that the current output gap is a RHS variable, it will generally not be uncorrelated with the error term. Therefore, estimation of equation (1) by Ordinary Least Squares or Nonlinear Least Squares will be inconsistent. However, under the assumption of i.i.d. error terms, past information embodied in lagged values will be orthogonal to the error term, thereby justifying the estimation of equation (1) by instrumental variables.

Consider first a truncated version of (1)

⁶ The error term can come from measurement error on the LHS or i.i.d. markup shocks.

$$\pi_t = \frac{(1-\lambda)}{\lambda} \alpha x_t + (1-\lambda) \sum_{j=0}^{J-1} \lambda^j F_{t-1-j} (\pi_t + \alpha \Delta x_t) + \varepsilon_t \quad (2)$$

where I temporarily ignore the truncated subset of the SIPC. Under the assumption of i.i.d. error terms, one can use the orthogonality condition $E[\varepsilon_t Z_{t-1}] = 0$, where Z_{t-1} is a set of k variables dated $t-1$ or earlier, to consistently estimate λ and α by nonlinear IV. Efficient estimation of these parameters requires a set of instruments that satisfy the orthogonality condition and are sufficiently correlated with the regressors of (2). Note that all past forecasts on the RHS of (2) are valid instruments, as are lags of the output gap. In the baseline estimation, I will use lags of the output gap and a subset of the past forecasts as instruments.

In practice, the truncation that must be imposed on the SIPC adds an additional source of error into equation (2). Specifically, equation (2) should be written as

$$\pi_t = \frac{(1-\lambda)}{\lambda} \alpha x_t + (1-\lambda) \sum_{j=0}^{J-1} \lambda^j F_{t-1-j} (\pi_t + \alpha \Delta x_t) + \varepsilon_t + v_{t,t-J} \quad (2')$$

where $v_{t,t-J} = (1-\lambda) \sum_{j=J}^{\infty} \lambda^j F_{t-1-j} (\pi_t + \alpha \Delta x_t)$. Because this additional source of error is dated $t-J$ and earlier, the orthogonality condition will generally fail. However, consider the covariance of any variable z with $v_{t,t-J}$:

$$\text{cov}(z, v_{t,t-J}) = (1-\lambda) \sum_{j=J}^{\infty} \lambda^j \left[\text{cov}(z, F_{t-1-j} \pi_t + \alpha F_{t-1-j} \Delta x_t) \right].$$

This covariance will be nonzero unless z is uncorrelated with all forecasts dated $t-1-j \forall j \geq J$ of current inflation and changes in the output gap. However, because each covariance is weighted by $0 < \lambda^j < 1$, it follows that as the truncation point J rises, the covariance of any regressor with $v_{t,t-J}$ falls and will converge to 0 as J goes to infinity as long as the covariance of z with past expectations is not too explosive. Quantitatively, truncating past expectations should thus have little effect on the estimation for a large enough J . Monte Carlo exercises confirm that when the degree of information rigidity is low to moderate, we can consistently estimate α and λ even at low truncation points.⁷ As the true value of λ rises, we require ever higher truncation points to consistently estimate λ and α . When the true level of information rigidity is $\lambda=0.75$,

⁷ These are available from the author upon request.

so that firms update their information once a year on average, consistent estimation requires J close to 12.

2.1.2 Forecast Measures

To separate the issue of price-setting decisions from the rationality of forecasts, I rely on historical measures of forecasts. The first approach is to use median expectations data from the Survey of Professional Forecasters (SPF).⁸ The SPF data provide an ideal source of expectations because they are a direct measure of what economists were forecasting and are available on a quarterly basis.⁹ Specifically, the SPF provides expected future paths for prices and real output over each of the subsequent four quarters from each vintage period. To generate expectations of changes in the output gap, I assume that forecasters knew the actual changes in the CBO measure of potential output and derive expectations of future changes in the output gap as expected changes in output minus actual changes in the CBO measure of potential output. The main limitation is that forecasts are only provided for the next four quarters.

As an alternative, I also generate forecasts for each quarter in a way designed to closely replicate what forecasters would have believed each time period. Specifically, for each quarterly observation (e.g. 1982Q1), I follow Stock and Watson (2003) and run a set of bivariate VARs for both inflation and changes in the output gap with a set of predictive variables using real time data available to agents at that time.¹⁰ These are used to generate forecasts of future values of inflation and changes in the output gap from each set of VAR's of that vintage which are then averaged across (excluding the maximum and minimum forecasts).¹¹ I create lagged forecasts going as far as 12 periods earlier for each quarter from 1971Q2 until 2004Q2. For inflation forecasts, I use real time data of inflation, unemployment, and changes in the output

⁸ SPF data is available at the Philadelphia Federal Reserve Board <http://www.phil.frb.org/econ/spf/index.html>. Mean forecasts were also used and yielded qualitatively similar results.

⁹ Other survey measures are not in an appropriate form for this type of analysis. Either they do not contain forecasts of future quarters one by one or they do not yield precise estimates of future values of inflation and changes in the output gap.

¹⁰ Real time data was taken from Philadelphia Federal Reserve at <http://www.phil.frb.org/econ/forecast/realindex.html>. See Croushore and Stark (2001) for a description.

¹¹ In addition, I impose that the AR forecast be one of the variables to be averaged over.

gap (though the CBO measure used in the output gap is not real time data), as well as the final series for the level of short-term interest rates, the interest rate spread (10 year minus 3 month T-bills), the second difference of oil prices, the first difference of industrial production index, and capacity utilization. Each VAR uses only the previous twenty years of data.¹² For forecasts of changes in the output gap, I replace oil prices with the first difference of M0. The lag length in each VAR is selected using the AIC.

2.2 The New Keynesian Phillips Curve and Non-Nested Model Tests

The second criterion to assess the validity of the SIPC is whether it statistically outperforms alternative models of inflation dynamics. The natural alternative is the New Keynesian Phillips Curve, which the SIPC was designed to replace. I first discuss the estimation procedure for the NKPC, and then turn to the tests used to empirically differentiate between the two models.

2.2.1 New Keynesian Phillips Curve

The New Keynesian Phillips Curve can be expressed as

$$\pi_t = \beta F_t \pi_{t+1} + \kappa x_t + \varepsilon_t \quad (3)$$

where β is the discount factor and κ is a function of both real rigidity and the degree of price stickiness. This relationship implies current inflation is proportional to the current forecast of the present discounted sum of future output gaps. Because of the purely forward-looking nature of inflation, this model has been criticized on the grounds that it over-predicts the speed at which inflation responds to monetary policy shocks. MR motivate the sticky information model as a direct substitute for the NKPC on the grounds that the sticky information model can address the failures of the NKPC. Knowing whether the SIPC outperforms the NKPC empirically is thus a particularly interesting question.

I propose to estimate the parameters of the NKPC in manner consistent with that used for the SIPC. Namely, the measures of inflation expectations from section 2.1.2 can be used as

¹² Some series are not available over the whole sample. Additional forecasting variables are added as soon as twenty years' worth of data becomes available for that series.

RHS variables in estimating equation (3).¹³ Under the assumption that ε_t is uncorrelated with all past information, equation (3) can be estimated by instrumental variables. Instruments used include a constant, three lags of the output gap, and the time $t-1$ forecast of time $t+1$ inflation.¹⁴ Using SPF expectations, the GDP deflator for inflation, and the log-deviation of output from the CBO measure of potential output for the output gap, estimation of (3) by instrumental variables from 1971Q2 to 2004Q2 yields

$$\pi_t = -0.38 + 1.12F_t\pi_{t+1} + 0.04x_t + \varepsilon_t$$

(0.26) (0.07) (0.02)

with Newey-West (1987) HAC standard errors in parentheses. Note that the coefficient on expected inflation is greater than, though not statistically different from, one. The coefficient on the output gap is positive and statistically significant, as implied by the theory and noted in Adam and Padula (2003).¹⁵

2.2.3 Non-Nested Model Tests

Because the SIPC and the NKPC are non-nested, I use two approaches to test the empirical validity of the SIPC relative to the sticky-price alternative. First, I apply the Davidson-McKinnon (DM) J-test. This entails estimating each model augmented with the fitted value from the alternative model and testing the null that the coefficient on the fitted value of the alternative is zero. For example, under the null of the NKPC, we can estimate

$$\pi_t = \beta F_t\pi_{t+1} + \kappa x_t + \delta_{SI}\hat{\pi}_t^{SI} + \varepsilon_t \quad (4)$$

where $\delta_{SI}=0$ under the null of the NKPC and $\hat{\pi}_t^{SI}$ is the fitted value from estimating (2).¹⁶

Similarly, we can test the null of the sticky information model using

¹³ Roberts (1997) and Adam and Padula (2003) provide evidence that using survey measures of expectations of future inflation improves the empirical performance of the New Keynesian Phillips Curve.

¹⁴ While weak instruments are typically an issue in estimates of the NKPC, the use of expectations measures on the RHS mitigates this problem. One can strongly reject the null of weak instruments using the tests of Stock and Yogo (2004).

¹⁵ Equivalent estimates using VAR-based expectations yield $\beta=0.97$ (0.05) and $\kappa=0.02$ (0.02). Newey-West standard errors allow for serial correlation of four quarters. Almost identical results hold if labor's share is used instead of the output gap. Because much work has been done on estimating the NKPC, I will not report subsequent estimates of the NKPC unless these differ from those reported here.

¹⁶ This is also estimated by IV using the same instruments as when estimating the NKPC with the addition of $F_{t-1}\pi_t$.

$$\pi_t = \frac{(1-\lambda)}{\lambda} \alpha x_t + (1-\lambda) \sum_{j=0}^J \lambda^j F_{t-1-j} (\pi_t + \alpha \Delta x_t) + \delta_{SP} \hat{\pi}_t^{SP} + \varepsilon_t \quad (5)$$

where $\hat{\pi}_t^{SP}$ is the fitted value from estimating (3) and $\delta_{SP}=0$ is the null under the sticky information model.¹⁷ Possible outcomes of the test include rejecting both models, rejecting neither, or rejecting one and not the other.¹⁸

As an alternative but closely related approach, I also consider an encompassing model test. Specifically, I estimate the following encompassing model

$$\pi_t = \omega \pi_t^{SP}(\gamma, \kappa) + (1-\omega) \pi_t^{SI}(\lambda, \alpha) + \varepsilon_t \quad (6)$$

where $\pi_t^{SP}(\gamma, \kappa)$ is the NKPC of equation (3) and $\pi_t^{SI}(\lambda, \alpha)$ is the SIPC of equation (2).¹⁹ Hence under this approach, I estimate the parameters of the two models jointly along with the weighting parameter ω . Under the null of the sticky price model, we should have $\omega=1$, while the null of sticky-information is $\omega=0$.²⁰ As with the DM tests, this approach can accept one model and reject the other, reject both, or fail to reject either.

3 Results

The inflation data is measured using the implicit GDP price deflator. The output gap is measured as the annualized log-deviation between real GDP and the CBO measure of potential output. I consider alternative measures of inflation and the output gap as robustness checks subsequently. All estimating equations include a constant.

¹⁷ In this case, instruments are the same as when estimating the SIPC plus one lag of inflation and $F_{t-1}\pi_{t+1}$.

¹⁸ See Davidson and McKinnon (2002). Because these estimates are sometimes sensitive to initial values, I use two sets of initial values ($\delta_t=0$ and $\delta_t=1.0$) and present results from the one that achieves the lowest value of objective function. The initial values used for other parameters are the estimated parameters from each Phillips Curve.

¹⁹ For the encompassing equation, I use all instruments from estimating the hybrid NKPC and SIPC.

²⁰ Because this expression is highly nonlinear in five parameters, I estimate the parameters using a Markov Chain Monte Carlo approach, following Chernozhukov and Hong (2003). I impose that $0 < \beta < 1$, $0 < \lambda < 1$, $0 < \omega < 1$, $\alpha > 0$, $\kappa > 0$. Starting values for the iterations are $\beta=0.99$, $\lambda=0.75$, $\alpha=0.10$, $\kappa=0.01$, and $\omega=0.5$. I use 10,000 burn-in iterations and 100,000 subsequent iterations for the estimation. The standard deviation of shocks is taken from standard errors of parameter estimates from single-equation estimations, and set to 0.1 for ω . The objective function is that of nonlinear IV.

3.1 Baseline Results

Table 1 presents estimates of the SIPC in (2), the DM tests of (4) and (5), and the encompassing model (6) on the full sample from 1971Q2 until 2004Q2 for different truncation points in the SIPC for the two measures of expectations. Looking first at the results based on SPF forecasts, the estimates of informational and real rigidities are both negative and insignificantly different from zero, contradicting the theoretical assumptions of the SIPC that both be positive. In addition, the SIPC is rejected under both non-nested model tests. The NKPC, on the other hand, is not rejected by the encompassing model test and only weakly so using the DM test (at the 10% level). Using the real-time VAR forecasts, the results are broadly similar regardless of the truncation used in the SIPC. Again, the estimates of both informational and nominal rigidities are insignificantly different from zero, though both are now positive. The non-nested model tests all reject the SIPC but fail to reject the NKPC. The evidence is thus unfavorable to the sticky information model along both sets of criteria considered. First, unlike the NKPC, the estimated structural parameters of the SIPC are inconsistent with an underlying sticky information model since we cannot reject that firms update their information every quarter. Secondly, the estimated SIPC is statistically inferior to the NKPC. Thus, by both metrics considered, the sticky information Phillips Curve finds little support in the data.

To see why this may be, it is worthwhile examining the predicted values of the models. Figure 1 plots inflation and predicted inflation from the NKPC with $\beta=0.99$ and $\kappa=0.01$.²¹ Overall, predicted inflation from the NKPC tracks actual inflation closely. It captures the two increases in inflation of the 1970s and early 1980s, but over-predicts inflation throughout the mid to late 1980s. This version of the NKPC accounts for approximately 80% of the variation in inflation. Figure 3.2 plots actual inflation and that predicted by the SIPC using the real-time VAR forecasts with a truncation of three years with the parameter values proposed by MR: $\lambda=0.75$ and $\alpha=0.10$.²² Predicted inflation from the SIPC accounts for a much smaller fraction of the variation in inflation, approximately 55%. In addition, this series differs from the time series of inflation along two dimensions. First, predicted inflation fails to replicate the two inflation spikes of the 1970s and early 1980s and is also unable to reproduce the disinflation of the mid-

²¹ $\kappa=0.01$ is approximately equivalent to firms updating prices once a year on average with $\alpha=0.10$.

²² Constants are set to zero for the NKPC and SIPC.

1980s. I will refer to this as the *real-time forecast error* effect. Second, predicted inflation is much smoother than actual inflation. I will refer to this as the *inflation inertia* effect.

3.2 The Real-Time Forecast Error Effect

The first effect is labeled the real-time forecast error effect because it reflects a feature specific to real-time forecasts of inflation: forecast errors are consistently too low in the 1970s, but too high in the 1980s and 1990s. Figure 2 illustrates this using a moving average (centered 4-quarter) of SPF forecast errors at horizons of one and four quarters. During both inflationary episodes in the 1970s, forecast errors are positive, reflecting the fact that forecasters were caught off-guard by rising inflation rates. On the other hand, since the Volker disinflation, professional forecasters have been consistently overestimating inflation. Both features are increasing in the forecasting horizon. The VAR forecasts based on the real-time data that was available to forecasters each quarter yield a very similar pattern.

This has important implications when estimating the parameters of the SIPC. In particular, a high value of λ in the SIPC places substantial weight on older forecasts of current inflation. This accounts for why predicted inflation, under the parameters of Mankiw and Reis, is lower than actual inflation in the 1970s, but consistently higher than actual inflation in the 1980s and 1990s. Because the estimation seeks to minimize persistent departures between predicted and actual inflation, we get estimated values of λ that are close to zero: low values of λ shift the weight in the SIPC from old forecasts of current inflation to more recent forecasts of current inflation, which exhibit a less pronounced pattern of persistent forecast errors.

3.3 The Inflation Inertia Effect

The inflation inertia effect refers to the excessive persistence and insufficient volatility of predicted inflation from the SIPC. To see this, suppose again we impose the preferred values of MR: $\lambda=0.75$ –firms update their information once a year on average– and $\alpha=0.10$ –a significant amount of real rigidity– on real-time VAR forecasts with a truncation of three years. The standard deviation of predicted inflation is 1.78. Actual inflation over the same time period had a standard deviation of 2.61, which implies that the SIPC under predicts the volatility of

inflation by over 30 percent. In addition, predicted inflation from the SIPC has an AR(1) coefficient of 0.999, whereas actual inflation has persistence of 0.88. For comparison, predicted inflation from the NKPC –assuming $\beta=0.99$ and $\kappa=0.01$ – has a standard deviation of 2.35 and persistence of 0.94.

The inflation inertia effect reflects the fact that the SIPC implies that inflation depends on a weighted average of past expectations of inflation. When there is a lot of information rigidity (λ is high), the SIPC places substantial weight on past expectations. This averaging across past expectations then filters out the volatility in past expectations, leaving only a smooth series in its wake. Note that this is another factor that pushes λ down in the estimation. With a low λ , most of the weight is placed on the most recent expectation and little on past forecasts. This eliminates the filtering process and enables the SIPC to more closely match the volatility and persistence of inflation. I discuss the source of the inflation inertia effect in section 5.2.

4 Robustness

In this section, I investigate several issues that arise in the context of estimating sticky price and sticky information models. The first is the choice of series. I verify that my results are robust to using alternative measures of inflation as well as to using labor’s share instead of the output gap, a point that has received much attention in the sticky price literature. Second, I consider the use of in-sample forecasts, as implicitly done in most other empirical work on the SIPC. Third, I redo the estimation while imposing a coefficient of real rigidity. Fourth, I examine the evidence for sticky-information in the sub-period since the Great Moderation.

4.1 Robustness to Data Series

In the baseline estimation, the choice of the GDP Deflator and the output gap (defined as the deviation of output from the CBO measure of potential) were based on limited availability of SPF forecasts for other series. In this section, I reproduce out-of-sample VAR forecasts for two alternative measures of inflation as well as for the use of labor’s share instead of the output gap. In each case, I generate forecasts from each quarter using the data preceding that date. I

then replicate the estimation procedures outlined in section 2. The results for a truncation of the SIPC of three years are presented in Table 2.

With the Non-Farm Business Deflator as our measure of inflation, estimates of the degree of informational and real rigidities are small, but positive, and insignificantly different from zero, confirming the baseline results of Table 1. The SIPC is again rejected according to both non-nested model tests. However, unlike the baseline results, the NKPC is also rejected by both non-nested model tests, despite the fact that the estimated parameters of the NKPC (not shown) are nearly identical to those found previously. This rejection of the NKPC reflects the fact that the NKPC explains a smaller fraction of the variation in NFB inflation than GDP Deflator inflation, with an R^2 of 0.70, rather than an improved performance of the SIPC. In particular, NFB inflation is more volatile than GDP Deflator inflation, and expectations of future inflation are unable to account for this increased variation in inflation. With CPI inflation, the point estimate of information rigidity, at 0.40, is larger than in previous cases and is significantly different from zero at the 10% level. The estimated coefficient of real rigidity remains insignificantly different from zero. However, the SIPC continues to be strongly rejected in the non-nested model tests. The NKPC is also rejected, reflecting the fact that CPI inflation is even more volatile than NFB inflation, and again this increased volatility is not sufficiently accounted for by expectations of future inflation.

I also consider the use of labor's share as the relevant forcing variable in each Phillips Curve. Gali and Gertler (1999) argue that labor's share is a better measure to use than the output gap since it is more closely tied to marginal costs. The use of labor's share in the estimation of the two Phillips Curves has little effect on the estimation results here. The estimated degrees of informational and real rigidities are insignificantly different from zero. The non-nested model tests continue to strongly reject the SIPC, but fail to reject the NKPC. Thus, the use of labor's share does not qualitatively change any of the results relative to the baseline estimation.

4.2 In-Sample vs Out-of-Sample Forecasts

Previous work on the empirical validity of the SIPC has typically not rejected the SIPC on structural grounds, with most finding estimated levels of information rigidity consistent with firms updating their information between once and twice a year. A key difference between the approach used here and this previous work is the nature of the forecasts used. Rather than relying on real-time forecasts, previous authors have relied on a single VAR estimated over the whole period to generate expectations.²³ Such an approach, by construction, eliminates the real-time forecast error effect since forecast errors in the VAR must be i.i.d. To see that this is important for the estimation, I construct an alternative set of forecasts using a single VAR with inflation and changes in the output gap estimated over the whole sample.²⁴ I then use the VAR coefficients to generate forecasts from each time period. The baseline estimation, using these in-sample forecasts, is redone and the results are presented in Table 3.2. Note that the estimated levels of information rigidity are now positive and statistically significant, implying that firms update their information a little over twice a year on average. However, the estimated degree of real rigidity remains insignificantly different from zero and the non-nested model tests continue to strongly reject the null of the SIPC, but fail to reject the null of the NKPC.²⁵

This alternative set of forecasts illustrates the importance of the real-time forecast error effect. By construction, in-sample VAR forecasts eliminate the real-time forecast error effect. Yet the real-time beliefs of forecasters differed substantially from what they would have forecasted had they had access to information from future values. Since the key idea behind sticky information is that inflation depends largely on agents' beliefs about the current state, the use of historical forecasts is more appropriate given the very different patterns exhibited by in-sample forecasts over the same time period. One should also note that the inflation inertia effect is present regardless of whether in-sample or out-of-sample forecasts are used. With in-

²³ See Dupor et al (2006), Kiley (2006), and Korenok et al (2006). Note that these authors estimate the VAR and the SIPC jointly, thereby imposing the cross-equation restrictions implied by the rational expectations solution. Khan and Zhu (2006) is an exception as they rely on out-of-sample forecasts.

²⁴ The VAR is estimated from 1967:Q1 to 2004:Q2. Lag length is chosen using the AIC.

²⁵ For the NKPC, the results are largely unchanged. The estimated β is 1.01 (0.01) and the estimate of κ is 0.008 (0.003).

sample forecasts, the standard deviation of predicted inflation from the SIPC under the assumed parameters of MR is 25 percent less than that of actual inflation.

4.3 Imposing the Degree of Real Rigidity

In this section, I consider the implications of imposing a degree of real rigidity in the estimation of the SIPC as a way of more precisely estimating the degree of information rigidity.²⁶ In particular, I focus on the case of $\alpha=0.10$, the value assumed by MR. Low values of α imply substantial strategic complementarities in price setting among firms and are necessary for the sticky information model to deliver a delayed response of inflation to monetary policy shocks.²⁷ In addition, because substantial amounts of real rigidity are also necessary for sticky price models to match the persistence in the data, imposing this value does not bias, ex ante, the exercise in favor of either model.²⁸ For the NKPC and SIPC to have identical degrees of freedom, I also restrict the coefficient on the output gap in the NKPC to be $\kappa=0.01$. Note that the latter is equivalent to imposing $\alpha=0.10$ and firms update prices approximately once a year on average. These values are also imposed in each non-nested model test.

The results are also presented in Table 2. Note that the estimated levels of information rigidity are now 0.52 and 0.53 for SPF and real-time VAR (truncation of three years) forecasts respectively and are significantly different from zero at the 1% level. However, the non-nested model tests again reject the null of the SIPC but fail to reject the NKPC. Thus, while the SIPC continues to be fare poorly on statistical grounds, it appears to fare better on structural grounds, i.e. according to the first criterion. The reason why estimates of information rigidity are higher with this imposed value of α is as follows. In the unrestricted case, the estimate of λ must be close to zero to minimize both the real-time forecast error and the inertia effects. But the data imply a small and positive link between inflation and the output gap, as seen in the estimates of the NKPC. Note that the coefficient on the output gap term in the SIPC is $(1-\lambda)\alpha/\lambda$. If the estimate of λ must be close to zero, then α must be small as well to avoid having a large coefficient on the output gap. This is what occurs in the unrestricted estimation. But when α is

²⁶ I am grateful to an anonymous referee for this suggestion.

²⁷ See Coibion (2006).

²⁸ Woodford (2003) argues that plausible values of α are between 0.10 and 0.15.

imposed to be greater than its unrestricted estimated value, this magnifies the coefficient on the output gap. To offset this effect requires higher estimated values of λ .

To illustrate this effect more clearly, I reproduce estimates of the degree of information rigidity for levels of α between 0 and 0.5. Parameter estimates and standard errors are shown in Figure 3. Note that estimates of λ are rising monotonically with α , consistent with the explanation above. However, as the estimated degree of information rigidity rises, the real-time forecast error and inflation inertia effects become increasingly present and the empirical fit of the model declines. This is illustrated by the fact that the R^2 of the SIPC is rapidly declining in α . Interestingly, this is not the case for the NKPC, for which the empirical fit is much more robust to the assumed value of α .²⁹ Figure 4 illustrates this by showing the implied R^2 of predicted inflation from the NKPC with imposed values of κ . Essentially, there is an empirical tradeoff between the two criteria for assessing the SIPC: when we impose values of α that yield levels of information rigidity consistent with a delayed response of inflation to monetary policy shocks, the statistical fit of the SIPC worsens substantially relative to the NKPC, reflecting the real-time forecast error and inflation inertia effects.

4.4 Sub-Sample Estimates

One could argue that applying the SIPC to the 1970s is expecting too much of the model. Since this was a period of volatile output and inflation, in which these economic variables were much in the news, the time-dependent process underlying the sticky information model may be a particularly poor assumption (though the same could potentially be said for the sticky price model). In addition, Khan and Zhu (2006) perform a similar analysis for the SIPC and find plausible and statistically significant values of λ , but their estimates are from 1980Q1 on. To see whether the time sample is important, Table 2 presents results from replicating the baseline estimation since the first quarter of 1984. The post-1984 period is frequently referred

²⁹ This is due to the fact that a change in α has a smaller effect on the coefficient on the output gap in the NKPC than in the SIPC, when one assumes identical degrees of price stickiness and sticky information.

to as the “Great Moderation”, in which the volatility of output and inflation is greatly reduced relative to the previous period. As such, it is a natural break point to impose.³⁰

Note first that the estimated levels of information rigidity differ from those over the whole time period. Point estimates of the degree of information rigidity are all statistically positive and relatively high. With SPF forecasts, λ is estimated to be 0.75, exactly the value assumed by MR. Real-time VAR forecasts point to higher levels of information rigidity, reaching 0.94 at a truncation of 12 quarters. Note that $\lambda=0.94$ implies that firms update their information once every four years on average. The estimated levels α remain insignificantly different from zero in each case. The non-nested model tests again reject the SIPC but fail to reject the NKPC. However, the point estimates of ω imply that a larger weight is now placed on the SIPC than was the case over the whole sample, implying that its empirical fit has improved relative to that of the NKPC over this sub-sample period. Nonetheless, we cannot reject that $\omega=1$ but can strongly reject the null of $\omega=0$.

Overall, the sticky-information model clearly performs better over this sub-sample period along one dimension: estimates of the degree of information rigidity are now significantly different from zero. However, the fact that α remains insignificantly different from zero implies that it is still difficult to find any strong link between the nominal and real side of the economy when conditioning on past forecasts of the current state. In other words, there is still little evidence of a sticky-information *Phillips Curve*. In addition, the sticky information model continues to be strongly rejected against the alternative of the basic sticky price model using non-nested model tests, confirming the notion that statistically, the SIPC is outperformed by the simple sticky price model it was designed to replace.

5 Discussion

In this section, I delve more deeply into two puzzling results presented in the paper. The first is the difference in the estimated levels of information rigidity over the whole sample and since the 1980s. The second is the inflation inertia effect: why predicted inflation from the SIPC

³⁰ See McConnell and Perez-Quiros (2000). The results are qualitatively unchanged for different break points from the early to mid-1980s.

under the parameters of Mankiw and Reis appears to be so much more inertial than actual inflation.

5.1 Sub-Sample Estimates of Information Rigidity and Rationality of Forecasts

The most striking feature of the sub-sample estimates is the high estimated degrees of information rigidity. These stand in sharp contrast to those found over the whole sample which were low and not statistically different from zero. While the estimate of α remains insignificantly different from zero and the non-nested model tests are consistent across periods, the large difference across periods in estimated degrees of information rigidity appears puzzling. In this section, I argue that this sub-sample difference arises because of the near-observational equivalence of the SIPC and tests of the rationality of the forecasts used.

To see what drives the difference in estimated values of information rigidity across the two time periods, I consider a more reduced form of the SIPC

$$\pi_t = c + \theta x_t + (1 - \lambda_1) \sum_{j=0}^{J-1} \lambda_1^j F_{t-1-j} \pi_t + \alpha (1 - \lambda_2) \sum_{j=0}^{J-1} \lambda_2^j F_{t-1-j} \Delta x_t + \varepsilon_t \quad (7)$$

This specification makes the coefficient on the output gap a free parameter and allows for a different distribution of weights for past expectations of inflation and past expectations of changes in the output gap. Under the null of the sticky information model, the two values of λ should of course be the same.

Table 3 presents estimates of equation (7) using VAR forecasts with a truncation of three years over the entire sample and since 1984. Consider first the estimates over the whole time period. The coefficient on the output gap is positive and statistically significant, as was found with the NKPC. The estimated degree of λ_2 is almost one, which implies that there is little predictive power in the weighted sum of past expectations of changes in the output gap. This also renders α unidentified, explaining the absurdly large coefficient and standard errors of this parameter. Note also that the estimated value of λ_1 is 0.26, nearly identical to the baseline estimate of λ Table 1. If we eliminate past expectations of changes in the output gap, we find nearly identical estimates of λ_1 . If we drop the output gap as well, again there is little change in the estimated value of λ_1 . In the period since 1984, we get similar results: there is little

predictive power in past expectations of changes in the output gap, the estimated value of λ_1 matches the estimated degree of information rigidity over this time period in the baseline results, and dropping the output gap and past expectations of changes in the output gap does not change the estimated value of λ_1 . What this indicates is that estimated levels of information rigidity are largely driven by past forecasts of inflation, and almost exclusively so since 1984.

While this result goes against the null of the sticky information model, it is informative about what drives the estimated levels of λ in different specifications. In particular, to understand the empirical source of estimated levels of λ , we need only look at the restricted specification:

$$\pi_t = c + (1 - \lambda)[F_{t-1}\pi_t + \lambda F_{t-2}\pi_t + \lambda^2 F_{t-3}\pi_t + \dots] + \varepsilon_t$$

which can be rewritten as

$$\pi_t - F_{t-1}\pi_t = c - \lambda F_{t-1}\pi_t + (1 - \lambda)\lambda[F_{t-2}\pi_t + \lambda F_{t-3}\pi_t + \dots] + \varepsilon_t \quad (8)$$

The standard interpretation of such an estimating equation is that it tests the null of rational expectations. If forecasts are rational, the forecast error should be unbiased and unpredictable using old information. In equation (8), the null of rational expectations would not be rejected if $\hat{c} = \hat{\lambda} = 0$. If expectations are biased, then we should tend to find $\hat{c} \neq 0$. Alternatively, if the forecast error is predictable using old forecasts, then we may find $\hat{\lambda} \neq 0$. Note that equation (8) is a valid test of the rationality of expectations regardless of what the true Phillips Curve representation may be.

The point of estimating the SIPC and NKPC conditional on observed inflation forecasts was to separate two issues: the generation of aggregate forecasts by professional forecasters and whether price-setting decisions reflect sticky prices or sticky information. But what this last exercise shows is that the structural form of the sticky information Phillips Curve is observationally very similar to a test of the rationality of professional forecasts, despite the fact that the two are independent hypotheses. If the SIPC is the true price-setting model, estimating the structural parameters will correctly identify the structural degrees of informational and real rigidities. However, if the SIPC is not the true model, imposing the structural form of the SIPC on the data can reduce to testing a separate hypothesis: whether

professional forecasts of inflation are rational or not. Deviations from rational expectations by professional forecasters could lead to non-zero estimates of c or λ which could be mistakenly taken as evidence of information rigidity.

There is a long literature on the rationality of survey measures of inflation expectations, the principal result of which has been that inflation forecasts often appear to be rational in the long-run but not in the short-run.³¹ To verify this, I perform the following test of the rationality of the forecasts used in the estimation of the SIPC

$$\pi_t - F_{t-1}\pi_t = c + \psi F_{t-j}\pi_t + \varepsilon_t. \quad (9)$$

for different lagged expectations.³² The resulting values of $\hat{\psi}$, along with 95% confidence intervals, are presented in Figure 4 for SPF and real-time VAR forecasts. The key issue here is whether forecasts errors are predictable using older forecasts. The results confirm previous findings that forecasts appear to be consistent with rationality in the long run (at least according to this test), but not in the short run. Over the whole sample, there is little evidence that forecast errors of inflation are predictable using old forecasts. However, since the early 1980s, forecast errors have been highly predictable, even using very outdated forecasts. The latter is due to the pattern of falling inflation over this time period, which has left professional forecasters consistently overestimating inflation. The result, as could be seen in Figure 2, is a period in which forecast errors are persistent and highly predictable.³³ This pattern of predictable forecast errors could explain the difference in sub-sample estimates. Suppose the SIPC is not the true Phillips Curve. Over the whole sample, unpredictable forecast errors (at least according to specification (9)) would tend to imply estimates of λ not different from zero. However, since the 1980s, the predictability of forecast errors could lead to non-zero estimates of λ even if there was no sticky information in the data.

³¹ See Croushore (1998) and Roberts (1998) for evidence and a survey of this literature respectively.

³² Under the null of rational expectations, Working (1960) shows that the error term in equation (10) is MA(1). Hence, for $j=1$ I estimate equation (10) by IV using a constant and three lags of inflation as instruments. For $j>1$, the RHS variables are orthogonal to the error term and OLS is consistent. Newey-West HAC standard errors are calculated using a truncation equal to $j+1$ to take into account overlapping observations.

³³ Note that predictable forecast errors do not necessarily imply that forecasters are irrational. For example, they could be setting forecasts optimally conditional on limited information about the state of the economy or uncertainty about the goals (or credibility) of monetary policy-makers.

To assess whether this pattern of predictable forecast errors is sufficient to explain the difference between sub-sample estimated values of λ , I consider a Monte Carlo exercise in which firms face no costs to acquiring new forecasts. Specifically, I assume that firms face sticky prices but not sticky information, so that the NKPC holds, conditional on professional forecasts. The main components of the model closely follow Erceg and Levin (2003):

$$\pi_t = \beta \tilde{\mathbb{E}}_t \pi_{t+1} + \kappa x_t + \varepsilon_t^\pi \quad (10)$$

$$x_t = \theta \tilde{\mathbb{E}}_t x_{t+1} + (1 - \theta)x_{t-1} - (i_t - \tilde{\mathbb{E}}_t \pi_{t+1}) \quad (11)$$

$$i_t = \rho i_{t-1} + (1 - \rho)[\phi_\pi(\pi_t - \pi_t^*) + \phi_{\Delta y} \Delta y_t] \quad (12)$$

where (10) is the NKPC conditional on expected inflation, (11) is the Euler equation for the output gap x_t conditional on expectations of the future output gap and the real interest rate, and (12) is a Taylor rule describing interest-rate decisions by the central bank. The Taylor rule implies that the central bank responds to deviations of inflation from a stochastic target rate π_t^* . The target rate is observable by forecasters, but its components are not. The target rate follows

$$\pi_t^* = \pi_t^p + \varepsilon_t^q \quad (13)$$

where ε_t^q is an i.i.d. shock and π_t^p follows an AR(1) process with persistence parameter ρ_p . Forecasters use a Kalman filter to generate optimal forecasts of inflation, the inflation target, and the output gap conditional on the observable variables and the known structure of the model. These expectations are denoted by $\tilde{\mathbb{E}}_t$.

To generate a pattern of forecast errors similar to that observed in US data, the shocks to the permanent component of the inflation target are chosen so that the permanent component of target inflation in each simulation is equal to US HP-filtered inflation over the whole sample. This ensures an upward trend in inflation in the 1970s followed by a gradual disinflation in the 1980s and 1990s. The other shocks in the model are assumed to be normally distributed.³⁴ The model is simulated 10,000 times, with each simulation having a burn-in period of 45 quarters followed by 133 period approximating the period from 1971:Q2 to

³⁴ I set the standard deviation of iid shocks to the inflation target equal to 0.002.

2004:Q2.³⁵ The standard deviation of i.i.d. shocks to the NKPC is set equal to the standard deviation of the observed residuals from the NKPC (0.0027). The natural level of output follows an AR(1) process with persistence of $\rho_n=0.99$ and i.i.d. shocks with standard deviation of 0.007.

This approach, on average, replicates the qualitative patterns of forecast error bias and predictability in each sub-sample as well as over the whole period. Given this generated data, I estimate the SIPC in the same way as done in the empirical analysis with a truncation $J=4$ and $J=12$ over the whole sample and the post-1984 equivalent period using \tilde{E} as the measure of forecasts. Mean estimates of the degree of information rigidity from 10,000 simulations are presented in Table 4. First, the median estimated level of information rigidity over the whole sample is close to zero, as found in the baseline results. This follows from the fact that there is no sticky information in the model and that forecast errors are unpredictable on average over the whole period, as in the data. Median estimates of α are also quite small, as observed in the baseline results. Importantly, the mean estimate of λ in the sub-sample period is high: 0.60 on average using a truncation of one year and 0.82 with a truncation of three years. The median estimated α is also increasing in the estimated level of λ , as observed in the empirical results. This exercise illustrates how the gradual adjustment of forecasts, here due to imperfect information, can mistakenly lead one to conclude that sticky information is an important component of price-setting decisions. Thus, the presence of predictable forecast errors since the 1980s can account for what appears to be high estimated levels of information rigidity over this time period.

5.2 The Inflation Inertia Effect

A striking feature of predicted inflation from the SIPC under the parameters of Mankiw and Reis is that it is much more inertial than either actual inflation or predicted inflation from the sticky price model. This result, which is largely independent of whether in-sample or out-of-sample forecasts are used, is particularly surprising because the sticky-information model was designed specifically to match the inflation inertia observed in response to monetary shocks that sticky price models could not replicate. Mankiw and Reis explicitly argue that this is the key empirical

³⁵ I set $\beta=0.99$, $\kappa=0.01$, $\theta=0.5$, and $\rho=0.90$. Following Erceg and Levin (2003), I set the Kalman gain parameter equal to 0.13 and $\rho_p=0.999$. The standard deviation of temporary shocks to the inflation target is set to 0.002.

fact that is hard to match. Their model, with a sufficient combination of informational and real rigidities, is successful at matching this fact. Yet empirically, with this same combination of informational and real rigidities, it appears to do a worse job of matching the aggregate unconditional time series properties of inflation than the model it was designed to replace.

To see how this dichotomy could exist, consider the pattern of responses to identified monetary policy shocks as typically observed in the empirical literature: interest rates adjust rapidly, the response of output peaks after approximately six months, but the peak response of inflation does not occur until one to two years after the shock.³⁶ Now suppose that there were no other shocks in the economy, or that other shocks yielded a similar pattern of responses, then in the aggregate data inflation would tend to lag other macroeconomic variables, particularly interest rates and output. Table 5 presents cross-correlations of inflation with other macroeconomic variables from 1971:Q2 to 2004:Q2. Note first that the peak correlation (in absolute value) between inflation and output growth is contemporaneous. However, the correlation between current inflation and future values of output growth is greater than the correlation between current inflation and lagged values of inflation. Clearly, inflation does not lag output growth. With respect to labor's share, the peak correlation is also contemporaneous, and correlations with leads and lags are highly symmetric, indicating little evidence of inflation leading or lagging labor's share. With respect to the output gap and unemployment, inflation is most highly correlated with future values of these variables, indicating that it strongly leads. A similar pattern holds with respect to both short and long-term interest rates: inflation leads these variables. Overall, inflation does not appear to be a lagging indicator.

The fact that the lead-lag patterns observed after monetary policy shocks differ substantially from those in the aggregate data need not be surprising. Empirical work typically finds that monetary policy shocks do not explain much of the variation in inflation: Andres et al (2005) find that monetary policy shocks account for approximately ten percent of the variance in inflation, Coibion and Gorodnichenko (2007) find less than ten percent, and Altig et al's (2005) results point to fourteen percent of the variance of inflation being accounted for by

³⁶ See Christiano et al (1999).

monetary policy shocks.³⁷ Secondly, other identified shocks appear to yield more rapid responses of inflation. In particular, recent work on identifying technology shocks through long-run restrictions has uncovered that inflation appears to respond most strongly on impact.³⁸ In addition, variance decompositions indicate that technology shocks account for approximately fifty percent of the variance in inflation. The fact that aggregate inflation does not appear to lag other macroeconomic variables could thus just be a reflection of the small fraction of the inflation variance accounted for by monetary policy shocks and the fact that other more quantitatively important shocks yield a more rapid response of inflation.

This has important implications when estimating models based on aggregate data that inherits these features. The sticky-information model, while able to match the response to monetary shocks, does so by building in inflation inertia. This also produces a delayed response of inflation to other shocks, including technology shocks (see Trabandt (2005)). Because these seem to account for a much greater fraction of the variance of inflation than monetary policy shocks, the SIPC is strongly rejected when tested using aggregate unconditional data. The sticky price model, on the other hand, builds in forward-looking inflation. This can readily replicate the response of inflation to technology shocks, but not monetary policy shocks. However, because the latter account for a smaller fraction of the variance in inflation, the NKPC fares better in the data.

The implication is that the delayed response of inflation to monetary policy shocks is only one piece of the puzzle. The other piece is that inflation, overall, is *not* a lagging indicator. Matching both sets of facts requires a model that can predict a rapid response of inflation to certain shocks, like technology shocks, *and* a gradual response of inflation to policy shocks.³⁹ This issue is not limited to the basic sticky price and sticky information models. For example, Trabandt (2005) shows that a hybrid sticky price model that can replicate the delayed response of inflation to monetary policy shocks also implies a delayed response of inflation to technology shocks. Altig et al (2005) estimate a sticky wage model with indexation by minimizing the distance between the model's predicted responses for macroeconomic variables and those

³⁷ See also Leeper et al (1996) and Christiano et al (1999) for similar results for the price level.

³⁸ See Galí (1999), Christiano et al (2003), Edge et al (2003), and Dupor et al (2007).

³⁹ This point is emphasized in Dupor et al (2007).

observed in the data in response to identified monetary policy and technology shocks. Their estimates imply significant indexation of wages, allowing them to match the delayed response of inflation to monetary policy shocks. However, their model then counterfactually predicts a delayed response of inflation to technology shocks. Edge et al (2003) pursue a similar exercise, but reach the opposite conclusion: they find little evidence for indexation of wages. Their model matches the response of inflation to identified technology shocks, but is then unable to match the response of inflation to monetary policy shocks. This differentiated response of inflation to shocks in the data implies that models with standard mechanisms for inducing inflation inertia are unlikely to be able to replicate both sets of facts jointly.

6 Conclusion

Mankiw and Reis (2002) argue that the sticky information model, in which firms acquire new information infrequently but update prices continuously, should replace the sticky price model, in which firms acquire information continuously but change prices infrequently, as the workhorse model of monetary economics. Their primary motivation is that, unlike the sticky price model, the sticky information model can replicate a key stylized fact: monetary policy shocks have a delayed effect on inflation. I show that empirically the sticky price model strongly dominates the sticky information model conditional on historical inflation forecasts. There are two primary sources for this result. First, real-time inflation forecasts are consistently too low in the 1970s but too high since the 1980s, particularly at long horizons. The sticky information model, based on parameters that enable it to match the delayed response of inflation to monetary policy shocks, places significant weight on older forecasts and therefore under-predicts inflation in the 1970s but over-predicts it since the 1980s. Second, the sticky information model predicts excessive smoothness in inflation and that inflation should lag other macroeconomic variables. I argue that, in the data, inflation is *not* a lagging indicator, which reflects the fact that some non-monetary policy shocks predict rapid adjustment of inflation. The sticky-information model, which builds-in a delayed response of inflation to shocks, is thus unable to match this important element in the data and, as a result, fares poorly empirically.

There is a strong symmetry in the weaknesses of the two models. The literature has long emphasized the inability of the basic sticky price model to predict gradual responses of inflation to monetary policy shocks. However, recent work has also demonstrated that in response to technology shocks, inflation adjusts rapidly. The sticky price model can naturally replicate this feature. The sticky information model, on the other hand, can reproduce the delayed response of inflation to monetary policy shocks but is then unable to match the rapid response of inflation to technology shocks. Because aggregate data indicates that inflation fails to lag other macroeconomic variables, shocks that yield rapid adjustment of inflation explain much of the variance in inflation. It should thus not be surprising, *ex post*, to see the NKPC dominate the SIPC empirically.

The key implication is that a successful model must be able to match both sets of facts: rapid adjustment of inflation to technology (and possibly other) shocks as well as a delayed response of inflation to monetary policy shocks. Reconciling the two facts will likely require elements of both price stickiness and informational frictions. The fact that the New Keynesian Phillips Curve does so well empirically when estimated conditional on historical inflation forecasts points to the need to model both price and informational rigidities.⁴⁰ For example, in the model of Erceg and Levin (2003) used in the Monte Carlo exercise, the presence of uncertainty about the monetary policy target induces rational agents to slowly adjust their forecasts about future inflation after a permanent shock to the inflation target, causing inflation to respond more slowly than in a full-information environment. Yet the fact that technology is observable yields a rapid adjustment of inflation to productivity shocks.

Another approach would be to model professional forecasters as following a sticky-information approach in updating forecasts. This could be imbedded in an otherwise standard sticky price model. Thus we would have the basic NKPC conditional on inflation forecasts, which fares so well in the data, combined with a mechanism to explain the gradual adjustment of forecasts to shocks. This approach would yield a very tractable hybrid sticky-price sticky-information model.⁴¹ Alternatively, one could model a menu-cost environment in which firms

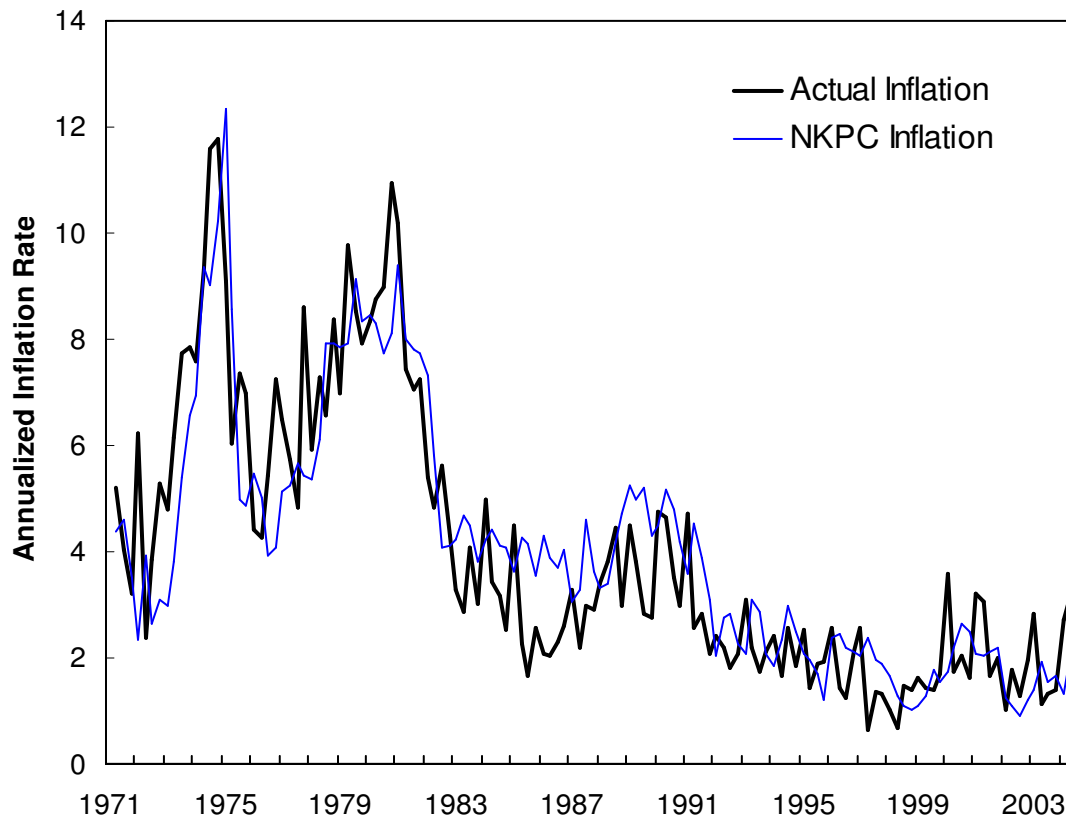
⁴⁰ See Dennis (2007), Knotek (2007) and Gorodnichenko (2007) for recent attempts at modeling price and information rigidities jointly and Dupor et al (2006) for an empirical application.

⁴¹ I am grateful to an anonymous referee for this suggestion.

observe their inputs and productivity perfectly, but face idiosyncratic unobservable shocks to their demand. In such an environment, inflation will respond rapidly to productivity shocks, but slowly to demand shocks since firms will seek to delay paying the menu cost until they are sufficiently confident that a persistent aggregate demand shock has occurred.⁴² Making further progress on modeling inflation dynamics is likely to require more refined approaches to modeling the interaction of price setting decisions with information acquisition and diffusion processes across the population.

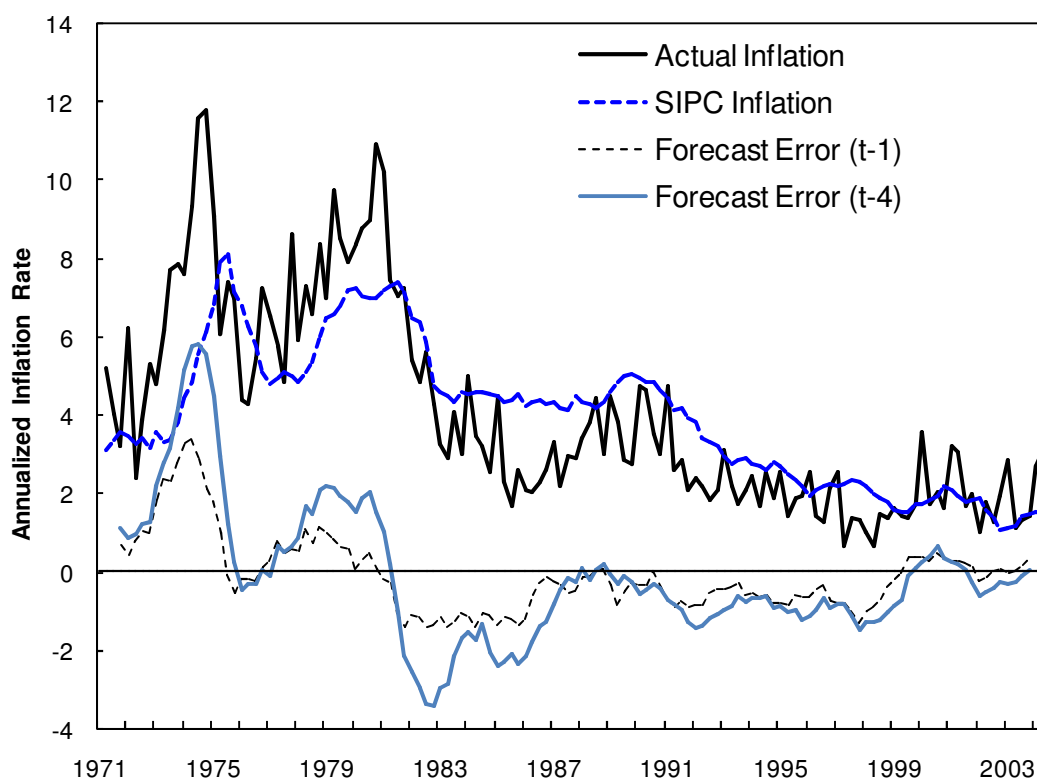
⁴² See Gorodnichenko (2007) for an example of menu costs and imperfect information jointly implying a delayed response of inflation to demand shocks.

Figure 1: Predicted Inflation from New Keynesian Phillips Curve



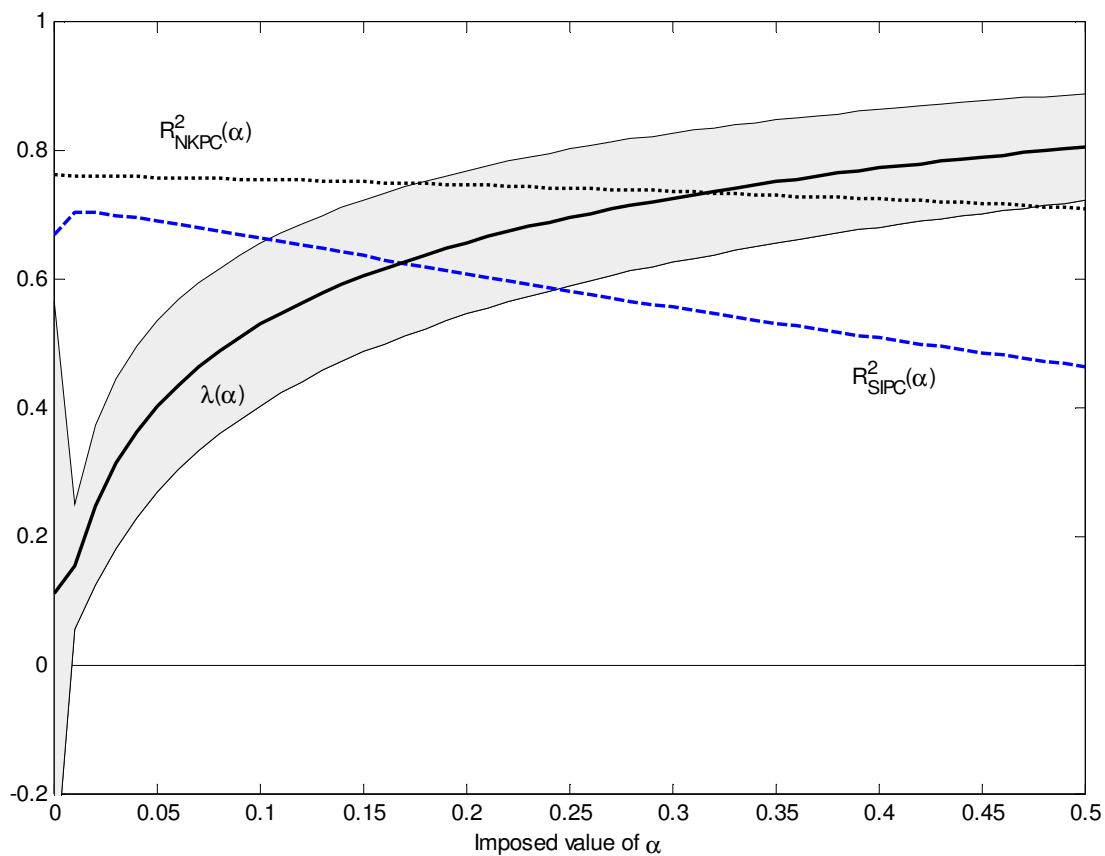
Note: Figure plots inflation and predicted inflation from the NKPC using expectations from VARs on real time data and $c=0$, $\beta=0.99$, and $\phi=0.01$.

Figure 2: Predicted Inflation from the Sticky Information Phillips Curve



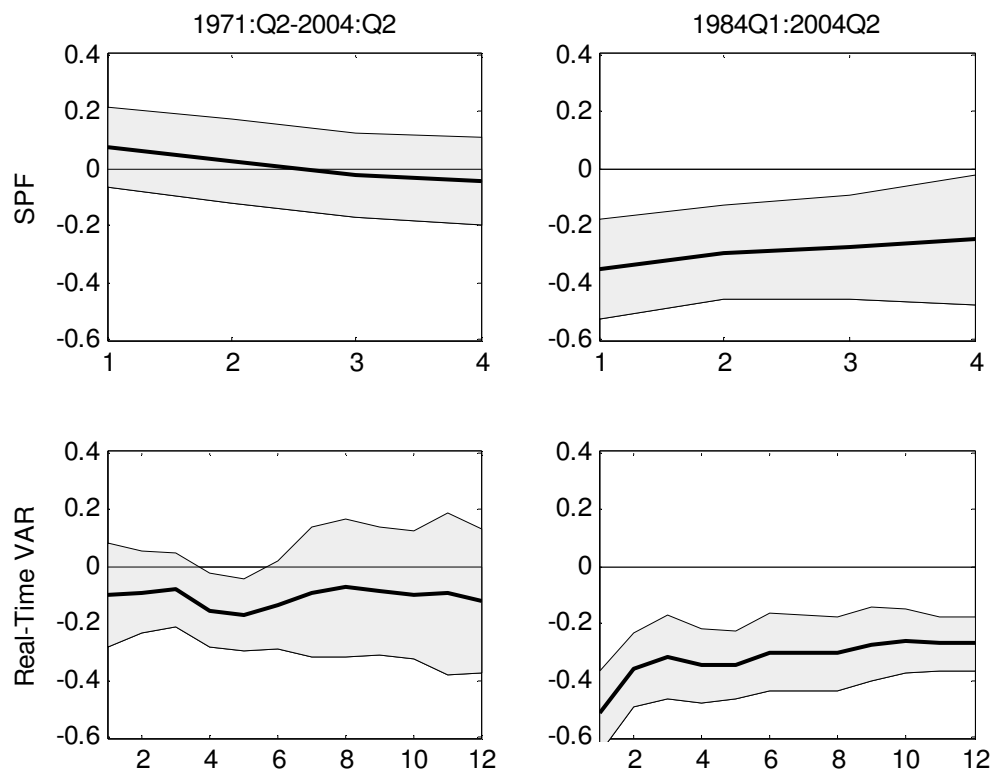
Note: Figure plots inflation and predicted inflation from the SIPC using expectations from VARs on real time data and $c=0$, $\lambda=0.75$, and $\alpha=0.10$ with a truncation of $J=12$. It also includes a 4-quarter centered moving average of SPF forecast errors at the one and four-quarter horizon.

**Figure 3: Estimated Levels of Information Rigidity
Conditional on the Degree of Real Rigidity**



Note: The figure displays estimated levels of information rigidity (λ) for different imposed values of real rigidity (α) when estimating the SIPC. Shaded areas indicate the 95% confidence intervals for estimates. The figure also displays the implied R^2 for the SIPC for each value of λ , as well as the corresponding R^2 from the estimated NKPC.

Figure 4: Tests of Rationality of Real-Time Inflation Forecasts



Note: The figure summarizes estimates of $\hat{\psi}$ from the regression of $\pi_t - F_{t-1}\pi_t = c + \psi F_{t-1} \pi_t + \varepsilon_t$, for SPF and real-time VAR forecasts over different samples. The grey shaded area is the 95% confidence intervals around parameter estimates. Standard errors are Newey-West HAC allowing for serial correlation of order $j+1$. The horizontal axis denotes the timing (lag) of the expectation on the RHS.

Table 1: Baseline Estimates of SIPC and Non-Nested Model Tests

	<i>SPF</i>			<i>Real-Time VAR</i>		
	<i>J=4</i>			<i>J=4</i>		<i>J=12</i>
Panel A: Estimates of Sticky-Information Phillips Curve						
c	0.14	(0.22)	0.07	(0.22)	0.07	(0.23)
λ	-0.27	(0.21)	0.21	(0.23)	0.20	(0.22)
α	-0.01	(0.01)	0.02	(0.02)	0.01	(0.02)
Panel B: Non-Nested Model Tests						
δ_{si}	0.47*	(0.28)	-0.37	(0.27)	-0.37	(0.27)
δ_{sp}	1.28***	(0.20)	1.20***	(0.22)	1.13***	(0.10)
ω	0.94***	(0.07)	0.97***	(0.03)	0.97***	(0.03)

Note: SPF and Real-Time VAR refer to the source of forecasts used in each equation. c , λ , and α are the constant and degrees of information and real rigidity in the SIPC (equation (2) in text) respectively. δ_{si} and δ_{sp} are the coefficients on the fitted values of the SIPC and NKPC in equations (4) and (5) respectively, and ω is the weighting parameter in equation (6). J is the truncation used in estimates of the SIPC. All estimates done by nonlinear IV. See text for instruments. Standard errors in parentheses are Newey-West HAC allowing for four quarters of serial correlation. Statistical difference from zero (and from one for ω) at the 10%, 5% and 1% levels are denoted by a *, **, and *** respectively.

Table 2: Robustness Analysis

	Alternative Data Series			In-Sample Forecasts		Impose $\alpha=0.1$		Post-1984	
	<i>NFB</i>	<i>CPI</i>	<i>LS</i>	<i>J=4</i>	<i>J=12</i>	<i>SPF</i>	<i>VAR(J=12)</i>	<i>SPF</i>	<i>VAR(J=12)</i>
Panel A: Estimates of Sticky-Information Phillips Curve									
c	0.19 (0.26)	0.17 (0.26)	-0.17 (0.20)	0.33*** (0.11)	0.27*** (0.10)	0.62*** (0.22)	0.2 (0.28)	0.41 (0.35)	0.81*** (0.31)
λ	0.10 (0.23)	0.40* (0.20)	0.17 (0.20)	0.37*** (0.12)	0.42*** (0.13)	0.52*** (0.06)	0.53*** (0.06)	0.75*** (0.07)	0.94*** (0.01)
α	0.01 (0.02)	0.06 (0.06)	0.01 (0.01)	0.03 (0.02)	0.03 (0.02)	0.10 n.a.	0.10 n.a.	0.04 (0.05)	0.54 (0.33)
Panel B: Non-Nested Model Tests									
δ_{si}	0.27** (0.11)	0.30** (0.15)	-0.34 (0.25)	0.28 (0.29)	0.27 (0.27)	0.14 (0.24)	-0.14 (0.23)	-0.08 (0.79)	-0.14 (0.66)
δ_{sp}	0.95*** (0.12)	1.00*** (0.13)	1.31*** (0.17)	0.82*** (0.11)	0.98*** (0.03)	1.53*** (0.15)	1.15*** (0.10)	1.03*** (0.34)	0.97*** (0.31)
ω	0.73***,aaa (0.05)	0.72***,aaa (0.05)	0.98*** (0.02)	0.93*** (0.07)	0.92*** (0.08)	0.95*** (0.03)	0.96*** (0.03)	0.79*** (0.15)	0.78*** (0.16)

Note: NFB uses Non-Farm Business Deflator series for inflation. CPI uses the Consumer Price Index for inflation. Labor's Share replaces the output gap as the forcing term in the Phillips Curves. In each case under Alternative Data Series, expectations are generated by VAR's using data from these series and others as described in section 2.1.2 of text. In-Sample Forecasts uses forecasts based on a single VAR estimated over the whole sample. c , λ , and α are the constant and degrees of information and real rigidity in the SIPC (equation (2) in text) respectively. Parameters δ_{si} and δ_{sp} are the coefficients on the fitted values of the SIPC and NKPC in equations (4) and (5) respectively, and ω is the weighting parameter in equation (6). All estimates done by nonlinear IV from 1971Q2 to 2004Q2, except for post-1984 estimates. See text for instruments. Standard errors in parentheses are Newey-West HAC allowing for four quarters of serial correlation. Statistical difference from zero (and from one for ω) at the 10%, 5% and 1% levels are denoted by a *^(a)., **^(a a), and ***^(a a a) respectively.

Table 3: Decomposing Estimates of SIPC

	<i>1971Q2-2004Q2</i>			<i>1984Q1-2004Q2</i>		
	(1)	(2)	(3)	(1)	(2)	(3)
c	-0.48 (0.58)	0.05 (0.23)	-0.17 (0.21)	0.87*** (0.31)	0.77** (0.32)	0.76** (0.31)
β	0.06*** (0.02)	0.06** (0.02)		0.03 (0.04)	0.02 (0.02)	
λ_1	0.26 (0.16)	0.21 (0.22)	0.11 (0.23)	0.94*** (0.03)	0.95*** (0.01)	0.95*** (0.01)
λ_2	0.99*** (0.26)			0.79 (0.69)		
α	-10.30 (252.00)			0.68 (1.77)		

Note: Estimates of equation (7) in the text by nonlinear IV. Standard errors in parentheses are Newey-West HAC allowing for four quarters of serial correlation. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and *** respectively. The truncation used is $J=12$.

Table 4: Estimating the SIPC in the Absence of Sticky Information

	<i>Whole Sample</i>		<i>Sub-Sample</i>	
	<i>J=4</i>	<i>J=12</i>	<i>J=4</i>	<i>J=12</i>
λ	-0.06 (0.29)	-0.08 (0.12)	0.60 (0.35)	0.82 (0.18)
α	0.00 (0.02)	0.00 (0.00)	0.04 (0.03)	0.33 (0.14)

Note: The table presents mean estimates of the degree of information rigidity (λ) in the SIPC when the true model has sticky prices but no sticky information, and forecast errors replicate, on average, those in the data. See section 4.1 for details. J is the truncation applied to the SIPC in the estimation. Results come from 10,000 simulations of 133 periods each. Standard deviation of parameter estimates are in parentheses. The sub-sample period is equivalent to post-1984 estimates in previous sections.

Table 5: Cross-Correlation of Inflation with other Macroeconomic Variables

j	$\rho(\pi_t, X_{t+j})$																
	-8	-7	-6	-5	-4	-3	-2	-1	0	1	2	3	4	5	6	7	8
GDP Growth	0.07	0.09	0.05	0.01	0.01	-0.10	-0.16	-0.18	-0.22	-0.21	-0.19	-0.21	-0.20	-0.13	-0.09	-0.05	-0.08
Labor's Share	0.39	0.40	0.43	0.49	0.53	0.56	0.59	0.62	0.63	0.60	0.56	0.54	0.53	0.52	0.49	0.46	0.46
Unemployment Rate	0.07	0.07	0.07	0.07	0.07	0.10	0.15	0.19	0.26	0.34	0.42	0.50	0.57	0.62	0.67	0.72	0.75
GDP Gap	0.04	0.08	0.10	0.11	0.12	0.08	0.02	-0.04	-0.12	-0.20	-0.27	-0.35	-0.42	-0.48	-0.52	-0.55	-0.59
T-Bill 3-month	0.07	0.13	0.20	0.29	0.37	0.43	0.49	0.55	0.59	0.61	0.63	0.63	0.61	0.60	0.58	0.56	0.56
T-Bond 10-year	-0.01	0.04	0.09	0.14	0.20	0.26	0.32	0.38	0.44	0.50	0.55	0.59	0.60	0.62	0.63	0.63	0.63

Note: The table presents the correlation between inflation at time t and each variable at time $t+j$, where j is given by the numbers at the top of the table. The values in bold indicate the timing of the highest correlation (in absolute value) between inflation and that variable. All correlations are done from 1971:Q2 until 2004Q2.

References

- Altig, David, Lawrence Christiano, Martin Eichenbaum, and Jesper Linde. 2005. "Firm-Specific Capital, Nominal Rigidities, and the Business Cycle" NBER WP 1104.
- Adam, Klaus and Mario Padula. 2003. "Inflation Dynamics and Subjective Expectations in the United States." Manuscript.
- 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.
- Ball, Laurence. 1994. "Credible Disinflation with Staggered Price-Setting," *American Economic Review* 84(1), 282-289.
- Caballero, Ricardo. 1989. "Time Dependent Rules, Aggregate Stickiness, and Information Externalities," Columbia University Discussion Paper Series No. 428.
- Calvo, Guillermo, Oya Celasun, and Michael Kumhof, 2003. "A Theory of Rational Inflationary Inertia" in Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund Phelps, ed. by P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, Princeton University Press.
- Carroll, Christopher D. 2003. "Macroeconomic Expectations of Households and Professional Forecasters," *Quarterly Journal of Economics* 118(1), 269-298.
- Chernozhukov, Victor and Han Hong. 2003. "An MCMC Approach to Classical Estimation," *Journal of Econometrics* 115(2), 293-346.
- Christiano, Lawrence, Martin Eichenbaum, and Charles Evans. 1999. "Monetary Policy Shocks: What Have We Learned And To What End?" *Handbook of Macroeconomics*, J.B. Taylor and Michael Woodford, eds. (Amsterdam, The Netherlands), 65-148.
- Christiano, Lawrence, Martin Eichenbaum, and Robert Vigfusson. 2003. "What Happens after a Technology Shock," NBER Working Paper #9819.
- Christiano, Lawrence, Martin Eichenbaum, and Charles Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy* 113(1), 1-45.
- Coibion, Olivier, 2006. "Inflation Inertia in Sticky-Information Models," *Contributions to Macroeconomics* 6(1), Article 1.
- Coibion, Olivier and Yuriy Gorodnichenko. 2007. "Strategic Interaction among Heterogeneous Price-Setters in an Estimated DSGE Model," Manuscript.
- Croushore, Dean. 1998. "Evaluating Inflation Forecasts," Federal Reserve Bank of Philadelphia Working Paper No. 98-14.
- Croushore, Dean and Tom Stark. 2001. "A Real-Time Data Set for Macroeconomists," *Journal of Econometrics* 105, 111-130.
- Davidson, Russell and James G. MacKinnon. 2002. "Bootstrap J Tests of Non-Nested Linear Regression Models," *Journal of Econometrics* 109, 167-193.
- Dennis, Richard. 2007. "The Frequency of Price Adjustments and New Keynesian Business Cycle Dynamics." Federal Reserve Bank of San Francisco Working Paper 2006-22.
- Dopke, Jorg, Jonas Dovern, Ulrich Fritsche, and Jirka Slacalek. 2005. "European Inflation Expectation Dynamics," Deutsche Bundesbank Discussion Paper, Economic Series, No. 37(2005).
- Dupor, Bill, Tomiyuki Kitamura, and Takayuki Tsuruga. 2006. "Do Sticky Prices Need To Be

- Replaced With Sticky Information?" Manuscript.
- Dupor, Bill, Jing Han, and Yi Chan Tsai. 2007. "What Do Technology Shocks Tell Us about the New Keynesian Paradigm?" Manuscript.
- Edge, Rochelle M, Thomas Laubach, and John C. Williams. 2003. "The Responses of Wages and Prices to Technology Shocks." Finance and Economics Discussion Series 2003-65. Washington: Board of Governors of the Federal Reserve System, 2003.
- Gali, Jordi. 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?," *American Economic Review*, 89(1), 249-271.
- Gali, Jordi and Mark Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis," *Journal of Monetary Economics* 44(?), 195-222.
- Gorodnichenko, Yuriy. 2006. "Endogenous Information, Menu Costs, and Inflation Persistence." Manuscript.
- 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. 2005. "A Quantitative Comparison of Sticky Price and Sticky Information Models of Price Setting." Forthcoming in *Journal of Money, Credit, and Banking*.
- Klenow, Peter J. and Jonathan L. Willis. 2007. "Sticky Information and Sticky Prices," *Journal of Monetary Economics* 54, S79-S99.
- 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. 2004. "Empirical Comparison of Sticky Price and Sticky Information Models." Manuscript.
- Korenok, Oleg, Stanislav Radchenko, and Norman R. Swanson. 2006. "International Evidence on the Efficacy of New Keynesian Models of Inflation Persistence." Manuscript.
- Leeper, Eric M, Christopher A. Sims, and Tao Zha. 1996. "What Does Monetary Policy Do?" *Brookings Paper on Economic Activity* 2.
- Mankiw, N. Gregory. 2001. "The Inexorable and Mysterious Tradeoff Between Inflation and Unemployment," *Economic Journal* 111, C45-C61.
- 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. 2003. "A Model of Monetary Non-neutrality and Structural Slumps," in Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund Phelps, ed. by P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, Princeton University Press.
- McConnell, Margaret M. and Gabriel Perez-Quiros. 2000. "Output Fluctuations in the United States: What Has Changed Since the Early 1980's?" *American Economic Review* 90(5), 1464-1476.
- Newey, Whitney K. and Kenneth D. West. 1987. "A Simple, Positive, Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica* 55(3), 703-708.
- Reis, Ricardo. 2006. "Inattentive Producers," *Review of Economic Studies* 73(3), 793-821.
- Roberts, John M. 1995. "New Keynesian Economics and the Phillips Curve," *Journal of*

- Money, Credit, and Banking* 27, 975-84.
- Roberts, John M. 1997. "Is Inflation Sticky?" *Journal of Monetary Economics* 39, 173-96.
- Roberts, John M. 1998. "Inflation Expectations and the Transmission of Monetary Policy," Finance and Economics Discussion Series 1998-43. Washington: Board of Governors of the Federal Reserve System.
- Stock, James H. and Mark Watson. 2003. "Forecasting Output and Inflation: The Role of Asset Prices," *Journal of Economic Literature* 41(3), 788-829.
- Stock, James H. and Motohiro Yogo. 2004. "Testing for Weak Instruments in Linear IV Regressions." NBER Technical Working Paper No. 284.
- Trabandt, Mathias. 2005. "Sticky Information vs. Sticky Prices: A Horse Race In A DSGE Framework." Manuscript.
- Woodford, Michael. 2003. Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton: Princeton University Press.
- Working, Holbrook. 1960. "Note on the Correlation of First Differences of Averages in a Random Chain," *Econometrica* 28(4), 916-918.