

Sticky Information and Sticky Prices*

Peter J. Klenow[†] and Jonathan L. Willis[‡]

October 2006

Abstract

In the U.S. and Europe, prices change somewhere between every six months and once a year. Yet nominal macro shocks seem to have real effects lasting well beyond a year. “Sticky information” models, as posited by Sims (2003), Woodford (2003) and Mankiw and Reis (2002), can reconcile micro flexibility with macro rigidity. We simulate a sticky information model in which price setters do not update their information on macro shocks as often as they update their information on micro shocks. Compared to a standard menu cost model, price changes in this model reflect older macro shocks. We then examine price changes in the micro data underlying the U.S. CPI. These price changes do not reflect older information, thereby exhibiting a similar response to that of the standard menu cost model. However, the empirical test hinges on staggered information updating across firms; it cannot distinguish between a full information model and a model where firms have equally old information.

*Prepared for the Swiss National Bank / Journal of Monetary Economics Conference on “Microeconomic Adjustment and Macroeconomic Dynamics,” October 2006. This research was conducted with restricted access to Bureau of Labor Statistics (BLS) data. The views expressed here are those of the authors and do not necessarily reflect the views of the BLS or the Federal Reserve Bank of Kansas City. Ed Knotek and John Leahy provided very useful comments. Josephine Smith provided excellent research assistance.

[†]Department of Economics, Stanford University and NBER, pete@klenow.net

[‡]Federal Reserve Bank of Kansas City, jonathan.willis@kc.frb.org

1 Introduction

Individual consumer and producer prices change every six months to a year. See Bils and Klenow (2004), Klenow and Kryvtsov (2005), and Nakamura and Steinsson (2006) for U.S. evidence, and Dhyne, Alvarez, Bihan, Veronese, Dias, Hoffman, Jonker, Lunnemann, Rumber and Vilmunen (2005) for a survey of studies on Euro Area countries. In contrast, many studies find that nominal macro shocks have real effects with a half-life well over a year. See, for example, Christiano, Eichenbaum and Evans (1999), Romer and Romer (2003), and Bernanke, Boivin and Elias (2004).

“Sticky information” theories can reconcile the macro price rigidity and micro price flexibility. These theories, advanced recently by Sims (1998, 2003), Woodford (2003), and Mankiw and Reis (2002, 2006), feature imperfect information about macro shocks. As a result, many rounds of micro price changes are needed to fully reflect a given macro shock. In versions such as Sims’, the micro flexibility is at the expense of macro flexibility, as firms face convex costs of processing information.

Our aim is to explore whether the tell-tale predictions of sticky information models are borne out in data on micro price changes. Specifically, do price changes reflect dated information on macro shocks and macro states? Given the lack of consensus on a measure of monetary policy shocks, especially one that explains inflation movements well, this is not a straightforward task. We therefore simulate simple general equilibrium models to derive responses of price changes to past inflation movements.

The models we simulate feature exogenous money growth, a cash-in-advance constraint, and monopolistically competitive firms. The firms face idiosyncratic productivity shocks as well as the aggregate money shocks, but do not change prices every period because they face costs of implementing price changes (i.e., menu costs). We model sticky nominal prices alongside sticky information for two reasons. First, 80-90% of prices do not change in the typical month, an important fact for a monetary business cycle model to match. Second, we exploit the lumpiness of price changes to test for sticky information. When a firm changes its price, we ask, does the change reflect only inflation innovations since their last price

change, or does it put weight on older innovations? Related, we can use lumpy price changes to explore whether firms that face bigger idiosyncratic shocks (i.e., exhibit larger and more frequent price changes) update their macro information less frequently. This is precisely the prediction of Sims' model of convex costs of processing information: the more micro shocks firms have to deal with, the less attention they will pay to macro shocks.

As a benchmark, we first consider a model with flexible information (i.e., constant updating on macro states). We then introduce staggered updating of information on macro states a la Taylor. As expected, the less frequent the updating of macro information, the more persistent the real output effects of money shocks. And the stickier the information, the more individual price changes reflect old inflation innovations as opposed to recent ones.

We choose several model parameters to match moments in the CPI Research Database maintained by the U.S. Bureau of Labor Statistics. We choose the mean, standard deviation and serial correlation of money growth in the model to match the mean, standard deviation and serial correlation of inflation in the data. We choose the size of menu costs and the size of idiosyncratic firm productivity shocks to match the frequency and size of micro price changes in the data. Our test is then whether the price changes in the data respond to old inflation innovations, or only those arriving since the firm last changed its price. In the data, we find little evidence that price changes reflect old information. However, the empirical test hinges on staggered information updating across firms; it cannot distinguish between a full information model and a model where firms have equally old information.

We use the test to examine two additional hypothesis. First, items with large and frequent price changes (big idiosyncratic shocks) respond strongly to recent information but also to older information. The first pattern is not so consistent with the Sims' rational inattention story but the second pattern is. Second, temporary price discounts are often filtered out on the grounds that they reflect idiosyncratic considerations rather than macroeconomic information. However, we find that sales-related price changes respond to macro information in much the same way that regular price changes do.

The rest of the paper is organized as follows. In section 2 we lay out the general equilib-

rium models featuring sticky prices (due to menu costs) and exogenously sticky information. In section 3 we describe the CPI micro dataset, and report statistics that we use to set parameter values in our models. In section 4 we compare the price changes produced by the models to those in the CPI microdata. In section 5 we offer tentative conclusions.

2 Model

In order to investigate the role of sticky information in the micro data, we construct a model with several key features. The basic structure of the model follows from Blanchard and Kiyotaki (1987). Households consume a wide variety of goods with a constant elasticity of consumption. Monopolistically competitive firms produce goods to meet demand at their posted prices. In order to generate a motive for holding money, we assume that households must pay for their consumption goods in cash before receiving their income. In order to generate the nominal price rigidities observed in the data, firms face a “menu” cost of implementing a price change. To examine the role of sticky information, we assume that information on the exogenous shocks to the economy arrives in staggered fashion. By changing sequencing of information arrival we can investigate different forms of information stickiness. Finally, we assume that firms use a boundedly rational forecast for inflation. This assumption allows us to obtain a nonlinear solution to the model.

2.1 Households

Households consume a variety of m goods and provide labor for production of the goods. Their choices are made to maximize

$$E_t \left[\sum_{t=0}^{\infty} \beta^t (C_t - \varphi L_t) \right] \tag{1}$$

where the consumption good, C_t , represents an aggregation of individual goods according to the Dixit-Stiglitz aggregator with a constant elasticity of θ :

$$C_t = \left(\sum_{j=1}^m C_{j,t}^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}}. \quad (2)$$

Households make their spending decisions at the beginning of the period before receiving their income, and we assume that their purchases must be paid for out of money holdings, M_t . Money holdings can be used to purchase consumption goods and real bonds, B_t :

$$\sum_{j=1}^m P_{j,t} C_{j,t} + P_t B_t = M_t. \quad (3)$$

Real bonds are priced by the cost of purchasing a unit of the aggregate consumption good, which is given by

$$P_t = \left(\sum_{j=1}^m P_{j,t}^{1-\theta} \right)^{\frac{1}{1-\theta}}. \quad (4)$$

Households receive income at the end of the period in the form of money. Income consists of wages earned by working for firms at a per-period wage rate of W_t , profits from their ownership share of firms, Π_t , returns from bond holdings including a real rate of return, r_t , and lump sum transfers of money from the central bank, X_{t+1} .¹ Income earned in period t provides the money holdings used for consumption in period $t + 1$:

$$M_{t+1} = W_t L_t + \Pi_t + P_t (1 + r_t) B_t + X_{t+1}. \quad (5)$$

The household budget constraint specifies that money spent on purchases in the current period does not exceed the money income earned in the previous period. Combining (3) and

¹We date the money transfer the following period to signify the period in which the transfer affects economic activity.

(5), we get

$$\sum_{j=1}^m P_{j,t} C_{j,t} + P_t B_t = W_{t-1} L_{t-1} + \Pi_{t-1} + P_{t-1} (1 + r_{t-1}) B_{t-1} + X_t. \quad (6)$$

The solution to the household's optimization decision provides the demand function, real interest rate, and wage rate that will be used by firms in their dynamic programming problem. Since the intertemporal marginal rate of substitution in consumption is equal to 1 for households, the real interest rate is constant, $r = \frac{1-\beta}{\beta}$. The first order condition for consumption of the differentiated goods can be transformed into the following demand function for good $C_{i,t}$ relative to good $C_{k,t}$:

$$C_{i,t} = \left(\frac{P_{i,t}}{P_{k,t}} \right)^{-\theta} C_{k,t}. \quad (7)$$

Households are indifferent between consuming today and saving for consumption in the next period. We solve for an equilibrium in which households spend all money holdings on consumption in the current period. Using the cash-in-advance constraint, the demand for a differentiated good can be expressed as a function of real money balances:

$$C_{i,t} = \left(\frac{P_{i,t}}{P_t} \right)^{-\theta} \frac{M_t}{P_t}. \quad (8)$$

Finally, using the households' labor supply decision, we derive an expected real wage that is constant. Since wage income earned today is not spent until the following period, households equate the marginal disutility of labor with the discounted expected marginal utility of consumption produced by marginal income earned from working today:

$$\phi = \beta E_t \left[\frac{W_t}{P_{t+1}} \right] \quad (9)$$

Rearranging this condition, we can solve for the real wage in the current period as a function

of inflation, π :

$$\frac{W_t}{P_t} = \frac{\phi}{\beta} \frac{1}{E_t \left[\frac{1}{1+\pi_{t+1}} \right]} \quad (10)$$

2.2 Firms

In the economy, there are m monopolistically competitive firms. Each firm produces a differentiated good, Y_i , using labor input, L_i . Producers are assumed to meet all demand at a given price, implying that $Y_i = C_i$.

Contemporaneous real profits for firm i are given by

$$\Pi_i = \frac{P_i}{P} Y_i - w L_i, \quad (11)$$

where P_i is the price for good i and w is the real wage. The firm faces the demand function given by (8), the real wage given by (10), and the production function

$$Y_i = Z_i L_i^\eta. \quad (12)$$

Here Z_i is an idiosyncratic productivity shock, and η governs returns to scale of production, allowing for decreasing returns due to (say) a fixed factor of production.

After substituting in the demand, real wage, and production functions, we arrive at the real profit function

$$\Pi_i = \left(\frac{P_i}{P} \right)^{1-\theta} \frac{M}{P} - \frac{\kappa (Z_i)^{\frac{-1}{\eta}}}{E \left[\frac{1}{1+\pi^l} \right]} \left(\frac{P_i}{P} \right)^{\frac{-\theta}{\eta}} \left(\frac{M}{P} \right)^{\frac{1}{\eta}}, \quad (13)$$

where $\kappa = \frac{\phi}{\beta}$.

2.2.1 Price adjustment cost

In order to generate nominal price rigidity, we assume that firms must pay a cost, ψ , in order to implement a price change. This cost is constant for all firms and in all periods and is expressed as a fraction of revenue in the steady-state symmetric equilibrium, where steady-state (ss) revenue for all firms is $R_{ss} \equiv \frac{M}{P}_{ss}$. If firm i chooses to change its price in the current period, then net contemporaneous profits, Π_i^C , will be

$$\Pi_i^C = \left(\frac{P_i}{P}\right)^{1-\theta} \frac{M}{P} - \frac{\kappa (Z_i)^{\frac{-1}{\eta}}}{E \left[\frac{1}{1+\pi'}\right]} \left(\frac{P_i}{P}\right)^{\frac{-\theta}{\eta}} \left(\frac{M}{P}\right)^{\frac{1}{\eta}} - \psi R_{ss}. \quad (14)$$

2.2.2 Information cost

To explore implications of sticky information, we assume that information regarding the exogenous state variables arrives in a staggered fashion. The two exogenous state variables in the model are the idiosyncratic profitability shock, Z , and the growth rate of nominal money supply, g_M . Potentially, idiosyncratic information may arrive at a different rate than aggregate information. Through changes in the timing assumptions, we will explore various forms of sticky information, including some cases in the spirit of rational inattention.

If new information does not arrive in the current period, we assume that the firm is not able to determine anything about the realizations of innovations to the exogenous state. This assumption is similar to claiming that the pricing managers do not interact with the production managers and the accountants of the firm, i.e., they do not see how many goods are produced nor do they observe the profits of the firm. We make this assumption to keep the model tractable and to present the starkest implications of sticky information. The assumption could potentially be relaxed by adding an additional shock to the model such as measurement error. Firms would then solve a signal extraction problem, as in Lucas (1973), when they do not have updated information.

Ideally, we would like to specify a model in which firms face a cost of acquiring information about the exogenous state variables. In such a model, firms would make a state-dependent decision each period regarding whether to pay the costs associated with obtaining updated

information on the idiosyncratic and/or aggregate shock. The assumption of staggered arrival of information could potentially be justified by a model in which the state-dependent decision on information updating results in a constant updating rule.

Given these information updating assumptions, firms will receive new information on a fixed schedule. Let \bar{n}_A be the number of periods between observing the aggregate money growth rate, and let \bar{n}_I be the number of periods between observing the idiosyncratic productivity shock. For a given firm in a given period, let n_A represent the number of periods since aggregate information was last observed, i.e., the age of aggregate information. Similarly, let n_I represent the age of idiosyncratic information. If a firm has updated information on both states, then $n_A = n_I = 0$.

2.3 Dynamic Optimization Problem

Given the presence of an implementation cost of a price change, the firm solves a dynamic optimization problem to maximize profits. In each period the firm decides whether or not to adjust its price. If it decides to adjust, it pays the implementation cost and resets its price. If it does not adjust, its nominal price remains fixed, and its relative price, $p_i = \frac{P_i}{P}$, decreases at the rate of inflation.

The state variables of the firm's optimization problem are impacted by the timing of information updating. Given our assumption that firms are not able to extract any signals about innovations if information is not updated, this implies that firms will not be able to update the endogenous aggregate state variables, inflation and real money balances, unless they received updated information on the money growth rate. The eight state variables are the firm's current nominal price relative to the aggregate price level at the last time that the aggregate information was observed ($p_{i,-n_A}$), the money growth rate when last observed ($g_{M,-n_A}$), the inflation rate when aggregate information was last observed (π_{-n_A}), the level of real money balances when aggregate information was last observed ($m_{-n_A} \equiv \frac{M}{P}_{-n_A}$), the idiosyncratic productivity index when last observed ($Z_{i,-n_I}$), the age of aggregate information (n_A), the age of idiosyncratic information (n_I), and the information set Ω used to form future

expectations of the endogenous state variables.

Given the state vector, $S = \{p_{i,-n_A}, g_{M,-n_A}, \pi_{-n_A}, m_{-n_A}, Z_{i,-n_I}, n_A, n_I, \Omega\}$, the firm maximizes the following value function:

$$V(S) = \max(V^C(S), V^{NC}(S)), \quad (15)$$

where $V^C(S)$ represents the firm's value conditional on changing its price and $V^{NC}(S)$ its value conditional on not changing its price. The value of a price change is expressed as

$$V^C(S) = \max_{p_{i,-n_A}^*} \{E_{-n_A, -n_I} [\Pi_i^C] + \beta E_{S'|S} [V(S')]\}, \quad (16)$$

with $S' = \{p_{i,-n'_A}^*, g'_{M,-n'_A}, \pi'_{-n'_A}, m'_{-n'_A}, Z'_{i,-n'_I}, n'_A, n'_I, \Omega'\}$. The firm's value function is discounted by β , reflecting the household's real interest rate.

In order to solve this optimization problem, the firm must be able to form expectations over the state variables. In periods in which current information is not observed, the firm computes expected profits conditional on the most recent information they have on the state variables. For example, to form an expectation of the current relative price, p_i , the firm takes the current nominal price relative to the price level n_A periods ago, $p_{i,-n_A}$, and integrates over all of the possible sequences of inflation over n_A periods conditional on information in the state vector. Regardless of the age of the information, the firm will always need to take conditional expectations of the future value function. The firm chooses the nominal price relative to the price level n_A periods ago, $p_{i,-n_A}^*$, that generates the highest expected value.

The value conditional on no price change is expressed as

$$V^{NC}(S) = E_{-n_A, -n_I} [\Pi_i] + \beta E_{S'|S} [V(S')], \quad (17)$$

with $S' = \{p_{i,-n'_A}, g'_{M,-n'_A}, \pi'_{-n'_A}, m'_{-n'_A}, Z'_{i,-n'_I}, n'_A, n'_I, \Omega'\}$.

For the exogenous state variables, money growth and idiosyncratic productivity shocks,

we assume autoregressive processes:

$$g_{M,t} = \mu_{g_M} + \rho_{g_M} g_{M,t-1} + \nu_{g_{M,t}}, \quad \nu_{g_M} \sim N(0, \sigma_{\nu_{g_M}}^2) \quad (18)$$

$$\ln Z_{i,t} = \rho_Z \ln Z_{i,t-1} + \nu_{Z_{i,t}}, \quad \nu_Z \sim N(0, \sigma_{\nu_Z}^2). \quad (19)$$

2.3.1 Bounded rationality

In order to compute a fully rational expectation of inflation, a firm needs to know the state variables of all firms in the economy, including the joint distribution of relative prices and idiosyncratic productivity shocks. One way to solve this model would be to introduce restrictions that reduce the heterogeneity to a manageable scope, as in Dotsey, King and Wolman (1999), hereafter DKW. An alternative solution is to assume that firms form inflation expectations based on a limited set of information. We choose the latter solution method for two reasons. First, the heterogeneity restrictions required for the DKW model do not match up well with the micro evidence.² Second, due to the heterogeneity introduced by staggered updating of information, assuming bounded rationality helps keep the model tractable.

We assume that firms use the following linear forecasting rule to form expectations of inflation:

$$\pi_{t+1}^f = \alpha_0 + \alpha_1 \pi_t + \alpha_2 \ln m_t + \alpha_3 g_{M,t} + \nu_{\pi,t}. \quad (20)$$

Firms will use the inflation forecast along with the forecast of money growth, from (18), to come up with a forecast for the log of real money balances, $\ln m_{t+1}^f$:

$$\ln m_{t+1}^f = \ln m_t + g_{M,t+1}^f - \pi_{t+1}^f. \quad (21)$$

The dynamic system used for forming aggregate expectations can be expressed as a

²See Klenow and Kryvtsov (2005) and Willis (2000).

three-variable autoregressive VAR:

$$\begin{bmatrix} \pi_{t+1}^f \\ \ln m_{t+1}^f \\ g_{M,t+1} \end{bmatrix} = A_0 + A_1 \begin{bmatrix} \pi_t \\ \ln m_t \\ g_{M,t} \end{bmatrix} + \xi_{t+1}. \quad (22)$$

With a little manipulation, we can convert (20), (21), and (18) into the following VAR system:

$$\pi_{t+1}^f = a_0 + a_1 \pi_t + a_2 \ln m_t + a_3 g_{M,t} + \nu_{\pi,t+1} \quad (23)$$

$$\begin{aligned} \ln m_{t+1}^f &= \mu_{g_M} - a_0 - a_1 \pi_t + (1 - a_2) \ln m_t + (\rho_{g_M} - a_3) g_{M,t} \\ &\quad + \nu_{g_M,t+1} - \nu_{\pi,t+1} \end{aligned} \quad (24)$$

$$g_{M,t+1} = \mu_{g_M} + \rho_{g_M} g_{M,t} + \nu_{g_M,t+1} \quad (25)$$

The equilibrium solution of the model requires the selection of an appropriate inflation forecast rule, $\Theta = \{\alpha_1, \alpha_2, \alpha_3\}$. Using this forecast rule, the firm will solve the optimization problem in (15) by determining a policy function for the updating of prices: $p_{i,-n_A}^* = f(p_{i,-n_A}, g_{M,-n_A}, \pi_{-n_A}, m_{-n_A}, Z_{i,-n_I}, n_A, n_I, \Omega)$.

The recursive equilibrium of the model consists of the functions V and f along with the inflation forecast rule, Θ , such that (i) V and f solve the firm's optimization problem and (ii) the expected inflation dynamics from the forecast rule matches the actual inflation dynamics resulting from firms' pricing decisions in a simulated economy.

2.3.2 Calibration and Simulation

Due to the presence of a discrete-choice decision in the optimization problem expressed in (15), the model is solved numerically using value function iteration. In this solution, all state variables are placed on discrete grids. The bounds of the relative price state are set wide enough to include all optimal pricing decisions, and prices are placed on the grid in increments of 0.4%, or about half the steady state inflation rate for this economy. The

autoregressive process for idiosyncratic productivity is transformed into a discrete-valued Markov chain following Tauchen (1986).³ This conversion provides us with the transition matrix expressing the expected probability of any given realization of Z_{t+1} as a function of the current state variables Z_t . The three-variable VAR for inflation, real money balances, and money growth is similarly converted into a first-order Markov chain.⁴ These transition matrices are used to compute the discounted expected value of the future period as well as expected contemporaneous profits if firms have out-of-date information.

Regarding information updating, there are two transition matrices for the respective updating of aggregate information and idiosyncratic information. For aggregate information, the transition matrix $\Phi_A(n_A, n'_A)$ provides the probability of moving from information of age n_A in the current period to information of age n'_A next period. A similar transition matrix, $\Phi_I(n_I, n'_I)$, exists for idiosyncratic information. The parametrization of these matrices will determine the stickiness of information.

We calibrate the structural parameters of the model to approximate several features of the BLS microdata that will be described in the next section. Table 1 displays the parameter values. A trimester frequency is used for the model in order to match up with the sampling frequency studied in the data. Therefore, we set the discount rate, β , equal to $0.96^{\frac{1}{3}}$. The elasticity of substitution between different consumer goods, θ , is set at 5, corresponding to a 25 percent markup for the firm. The implementation cost of a price change is set at 1.1 percent of revenue to induce a frequency of adjustment similar to the micro data. The parameters for the money growth process, ρ_{gM} and σ_{gM} , are set to produce inflation dynamics similar to the data. A random walk turns out to be a good approximation. The idiosyncratic productivity shock parameters, ρ_Z and σ_Z , are based on estimates in Klenow and Willis (2006) and the BLS facts reported below. Finally, the parameter κ , which is the marginal disutility of labor divided by the discount rate, is set at 0.5. The results of interest from the model are not sensitive to changes in κ .

³The discrete grid for idiosyncratic productivity contains 5 points spread equally in terms of the cumulative distribution function of the variable.

⁴The discrete grids for inflation, real money balances, and money growth contain 11, 7, and 5 points, respectively, spread equally in terms of the cumulative distribution function of the variables.

Table 1: Parameter Values

β	0.983
θ	5
η	0.9
ψ	0.011
ρ_{gM}	0
σ_{gM}	0.008
ρ_Z	0.25
σ_Z	0.08
κ	0.5

Following Willis (2003), the inflation forecasting rule expressed in (20) is used to compute a rational expectations equilibrium of the model. For a given specification of the structural parameters of the model along with the inflation forecasting parameters, $\Theta = \{\alpha_1, \alpha_2, \alpha_3\}$, the model is solved and the policy function is generated. A panel of 6,000 firms over 51 trimesters is then simulated using the policy functions.⁵

Simulating data from the model requires an updating process to determine the evolution of the endogenous aggregate-level state variables. The aggregate inflation rate and the level of real money balances are determined by the collective actions of firms in the simulation. When setting prices in the current period, firms with updated information, $n_A = 0$, possess the current value of inflation and real money balances. To determine the current-period inflation rate while simulating the model, which in turn determines the level of real money balances using equation (21), we locate the grid point in the discretized inflation state space that most closely matches equation (4), where the inflation rate is combined with P_{-1} to get P .

After simulating the full panel, we evaluate the forecasting rule used to form expectations

⁵The size of the panel was chosen to capture two features. First, the panel should have a large number of firms given the large number of price observations by the BLS. We found increasing the number of firms above 6,000 did not alter the results in any significant fashion. Second, the number of periods should match the length of the BLS sample, namely 51 trimesters. To replicate the panel for model moments, we lengthened the number of periods and then divided them into subsamples.

of inflation. An OLS regression of the linear forecasting rule in (20) is executed on the simulated data using the simulated values for inflation, real money balances, and money growth. The initial assumed values of the forecast parameters, Θ_0 , are then compared to the OLS estimates, Θ_1 . If these values differ, then the forecast parameters are updated based on Θ_1 and a new solution for the model is derived. This updating process continues until a fixed point is reached. This fixed-point solution represents a rational expectations equilibrium where the inflation forecasting rule assumed by firms matches up with the behavior of the simulated data.⁶

2.3.3 Sticky Information

We first consider a model in which aggregate information arrives in a deterministic fashion. The setting for \bar{n}_A provides the interval between updates of information. The updating across firms will be staggered so that a constant fraction of firms receive new information each period. To illustrate the deterministic updating of information, the transition matrix for the case with $\bar{n}_A = 2$ is shown below, where the rows represent the age of information in the current period, $n_A \in \{0, 1, 2\}$ and the columns represent the age of information next period, $n'_A \in \{0, 1, 2\}$:

$$\Phi_A(n_A, n'_A) = \begin{array}{|c|c|c|} \hline 0 & 1 & 0 \\ \hline 0 & 0 & 1 \\ \hline 1 & 0 & 0 \\ \hline \end{array}$$

To illustrate the features of information stickiness, we will consider four cases corresponding to the maximum age of aggregate information ranging from 0 to 3 periods, $\bar{n}_A \in \{0, 1, 2, 3\}$.

In terms of idiosyncratic information, we assume that firms always have current information on their idiosyncratic shocks ($\bar{n}_I = 0$). This assumption allows us to focus on the implications of aggregate information stickiness.

⁶Following Krusell and Smith (1998), we plan to explore whether the inclusion of additional variables into the forecasting rule will lead to a significant improvement in the inflation forecast. Candidate variables include additional lags of the state variables and moments of the price distribution.

To illustrate the role of sticky information, Figures 1 and 2 display the impulse responses of inflation and output to a 1 percent shock to the money growth rate. As shown in Figure 1, an increase in information stickiness leads to a delayed, hump-shaped response of inflation. The delayed inflation response suggests that there will be a stronger output response for sticky information models than for the baseline model. This pattern is clearly observed in Figure 2.

It is important to note that each of the four cases has a slightly different equilibrium inflation forecast rule. The parameters of the inflation forecast rule, equation (20), for each case are displayed in Table 2. The coefficients vary only somewhat, but the explanatory power of this forecast rule when estimated on simulated data increases dramatically when information is sticky. This suggests that this simple forecast rule for inflation is more reasonable for a model with sticky information than for a model with full information.

Table 2: Equilibrium forecast rules for model with sticky information updating

Model	a_1	a_2	a_3	R^2
Baseline ($\bar{n}_A = 0$)	-0.04	0.50	0.08	0.10
Sticky 1 ($\bar{n}_A = 1$)	0.01	0.49	0.16	0.73
Sticky 2 ($\bar{n}_A = 2$)	-0.03	0.47	-0.05	0.76
Sticky 3 ($\bar{n}_A = 3$)	-0.05	0.36	-0.03	0.67

Figure 1: Inflation response to 1 percent shock to the money growth rate





Figure 2: Output response to 1 percent shock to the money growth rate



2.3.4 Old information

As an alternative model of information stickiness, we also consider an economy in which all firms have equally old information. This assumption approximates a model in which information processing costs are such that it takes firms several periods to discern an aggregate shock. In our model, this would be represented as a case where firms **always** have aggregate information that is \bar{n}_A periods old.

As before, we will consider four different information assumptions. In the baseline model, firms always have current information. In the second case, firms always have aggregate information that is 1 period old. This differs from the previous model in that firms are now restricted so that they *never* possess current information, whereas in the sticky information model, half of firms possess current information and half possess information that is 1 period old. We also consider cases in which information is 2 and 3 periods old, respectively. Table 3 displays the equilibrium inflation forecast parameters. All of the coefficients change with the age of information, and the explanatory power of this equation on simulated data is strongest when information is 1 period old. This makes intuitive sense the forecast rule only contains information lagged one period, and it suggests that additional information lags should be added to the cases with older information. However, since each additional lagged variable becomes a state variable for the optimization problem, we cannot maintain tractability of the solution with an expanded forecast rule.

Table 3: Equilibrium forecast rules for model with old information

Model	a_1	a_2	a_3	R^2
Baseline ($n_A = 0$)	-0.04	0.50	0.08	0.10
Old 1 ($n_A = 1$)	-0.04	0.50	0.08	0.91
Old 2 ($n_A = 2$)	-0.34	0.39	-0.24	0.68
Old 3 ($n_A = 3$)	-0.21	0.26	-0.11	0.35

In future research, we also plan to explore another form of sticky information in which there is heterogeneity in the maximum age of information that is related to heterogeneity

in the volatility of idiosyncratic shocks and the size of implementation costs. This distinction is intended to approximate a model of rational inattention where firms that face larger idiosyncratic shocks choose to spend more effort processing idiosyncratic as opposed to aggregate information. In our setup, this can be modelled by modifying \bar{n}_A and \bar{n}_I so that firms that face larger idiosyncratic shocks “choose” to update idiosyncratic information more frequently and aggregate information less frequently. A high value of the implementation cost leads to less frequent updating of both information processes because the marginal value of new information decreases as implementation costs increase. Preliminary investigation of this model shows that the frequency of adjustment is a key element in determining whether rational inattention of this form can deliver results similar to Mackowiak and Wiederholt (2005). If firms update aggregate information frequently but they do not adjust because of implementation costs, then these firms may respond to aggregate shocks with an even longer delay than firms that update aggregate information less frequently but adjust prices more frequently.

3 CPI Data

For producing the Consumer Price Index, the U.S. Bureau of Labor Statistics conducts a monthly Commodities and Services Survey. This Survey covers all types of consumer products and services other than shelter, or around 70% of consumer spending. About 85,000 items are surveyed each month, with an item being a specific product (brand and detailed features) sold by a particular outlet. The data are collected from around 20,000 outlets located mostly in 45 large urban areas.

The CPI Research Database, maintained by the BLS Division of Price and Index Number Research, contains all prices in the Commodities and Services Survey from January 1988 to the present.⁷ We base our statistics on data through December 2004. The BLS tracks individual items for about five years, affording many opportunities to observe price changes.

⁷See Klenow and Kryvtsov (2005) for a more detailed description of the CPI Research Database.

The BLS labels each collected price as either a “regular” price or a “sale” price (i.e., a temporarily low price). Although sales may entail menu costs, Golosov and Lucas (2003) and others have argued that one should focus on regular prices for macro questions. We therefore report results for both regular prices and all posted prices. To construct a continuous series of regular prices, we substitute the most recent regular price whenever the current price is a sale price. To minimize the importance of measurement error, we drop price changes that exceed 10 natural log points in absolute value. These price jumps constitute less than one-tenth of one percent of all price changes.

The BLS collects prices monthly for food and energy items in all areas, and monthly for all items in New York, Los Angeles, and Chicago. For other areas, they check prices bi-monthly for “core” items (items other than food or energy). Each bi-monthly item is either odd (checked in months 1=January, 3=March, 5=May, 7=July, 9=September and 11=November) or even (checked in months 2=February, 4=April, 6=June, 8=August, 10=October, and 12=December). To use all items from all areas, and yet have a single frequency, we construct a “trimester” dataset. The first trimester is months 1-4 (January-April), and contains prices from January for “odds” and February for “evens”. The second trimester (months 5-8, or May-August) includes May prices for odds and June prices for evens. Finally, the third trimester (months 9-12, or September-December) has September prices for odds and October prices for evens. We label half the monthly items odds and half evens, and follow a subset of their prices accordingly. The disadvantage of looking at trimesters is that we are ignoring price quotes in months 3, 7 and 11 for odds and months 4, 8 and 12 for evens. Yet in so doing we incorporate the 85,000 items coming from all areas. If we were to stick with a monthly dataset, in contrast, we would have only around 14,000 items from the top 3 cities. Just as important, looking at trimesters rather than months allows us to consider models with greater stickiness of information without adding as many states (e.g., three trimesters as opposed to three months).

To help pin down key parameters in our model, we calculate five statistics from the CPI data. Three of the moments are the mean, standard deviation, and serial correlation of the

aggregate trimester inflation rate. In terms of our model, these can be thought of as helpful for setting the mean, standard deviation, and serial correlation of money growth. The other two statistics are the median frequency of price changes and the median size of price changes. These two moments guide our choices for the size of menu costs and the size of idiosyncratic productivity shocks.

To define the statistics precisely, let P_{sit} denote the price of item i in sector s in trimester t , and ω_{sit} the BLS weight on item i within category s in trimester t . The weights in sector s sum to ω_s^{95} in every trimester, the BLS consumption expenditure weight of category s in 1995 (which themselves sum to 1). We then define the aggregate inflation rate in trimester t to be

$$\pi_t = \sum_s \sum_i \omega_{sit} [\log(P_{sit}) - \log(P_{sit-1})].$$

We then take the simple average across the 50 trimesters from 1988 through 2004 to arrive at 0.805% per trimester (2.43% per year) for regular price inflation:

$$\mu_\pi = \sum_{t=1}^{50} \pi_t / 50 = 0.00805.$$

In similar fashion we calculate the standard deviation (0.482%) and serial correlation (0.163) of the inflation rate:

$$\sigma_\pi = \sqrt{\sum_{t=1}^{50} (\pi_t - \mu_\pi)^2 / 49} = 0.00482.$$

$$\rho_\pi = \sqrt{\sum_{t=1}^{49} (\pi_t - \mu_\pi)(\pi_{t+1} - \mu_\pi) / 48} = 0.163.$$

Our fourth moment is the frequency of items changing price from one trimester to the next. Let $I(\Delta P_{sit} \neq 0)$ be a price-change indicator for item i in sector s in trimester t . It takes on the value 1 if the item changed price from trimester $t - 1$ to t , and 0 otherwise. We calculate the mean value of this indicator for an item, then take the weighted median

value across items to arrive at 0.357 (35.7% per trimester). Easier to express explicitly is the cousin of this statistic, namely the weighted *mean* frequency of price changes, which is higher at 43.5%:

$$\overline{I(\Delta P \neq 0)} = \sum_s \sum_i \omega_{si} \frac{\sum_t I(\Delta P_{sit} \neq 0)}{\sum_t 1} = 0.435.$$

Here $\omega_{si} = \sum_t \omega_{sit}$. We prefer the median to the mean because, in time-dependent models at least, the median appears to provide a better approximation to a model with heterogeneity. Bils and Klenow (2004) examine this for the Taylor model, and Carvalho (2006) for the Calvo model.

Our fifth and final moment is the median absolute size of price changes, which is 0.0795 (7.95%). Again, it is easier to explicitly define the weighted *mean*, which is higher at 12.1%:

$$\overline{|\Delta P|} = \sum_s \sum_i \omega_{si} \frac{\sum_t |\Delta P_{sit}|}{\sum_t I(\Delta P_{sit} \neq 0)} = 0.121.$$

As stressed by Klenow and Kryvtsov (2005) and Golosov and Lucas (2003), absolute price changes are much larger than needed to keep up with the trend inflation rate. The trend is about 0.8% per trimester and the frequency of price changes is around 1/3, so price changes only need average about 2.4% to keep up with trend inflation. Yet the average price change is five times as large at 12%. These large price changes do not merely reflect different sectoral mean inflation rates, as Klenow and Kryvtsov report large price movements even relative to a sectoral price index defined for 200-300 separate categories of consumption. Given the relative stability of the aggregate inflation rate, idiosyncratic shocks will need to be large to generate such price changes in our model. Such idiosyncratic shocks will dominate individual firm decisions about when and how much to change prices, with aggregate conditions of less importance. As discussed earlier, firms might be rationally inattentive to aggregate state variables, preferring to focus on the first order idiosyncratic shocks.

In Table 4, nearby, we summarize these moments. We also give the corresponding moments in our baseline model. We chose the parameter values in our baseline model to roughly

match these moments.⁸

Table 4: Moments

	μ_π	σ_π	ρ_π	$\overline{I(\Delta P \neq 0)}$	$\overline{ \Delta P }$
BLS CPI Data	0.00805	0.00482	0.136	0.357	0.0795
Baseline Model	0.00810	0.00457	0.142	0.347	0.0830

4 Simulation and Estimation

One way to investigate the plausibility of a sticky information model is devise a test that reveals the extent to which firms respond to current versus lagged information. In terms of the state-dependent model presented here, we can precisely derive the price change of firms as a function of variables in their information set.

Conditional on a firm choosing to adjust its price, the Euler equation for the price decision is expressed as

$$\frac{\partial \Pi_{i,t}}{\partial P_{i,t}^*} + \beta E_t \left[(1 - \alpha) \frac{\partial V(S_{t+1})}{\partial P_{i,t}^*} \right] = 0, \quad (26)$$

where α is the probability of a firm changing its price. Here we make a simplifying assumption that the probability of price adjustment is independent of the time since the previous change. This assumption matches the flat hazard rate found in the micro data by Klenow and Kryvtsov (2005). It also is a reasonable approximation of the hazard function in the model because the volatility of idiosyncratic shocks dominates the small, but increasing, incentive to adjust due to the upward drift in the nominal money supply.

⁸The corresponding moments for posted prices, which include temporary price discounts, are 0.575% for mean inflation, 0.493% for its standard deviation, 0.189 for its serial correlation, 46.7% for the median frequency of price changes, and 8.91% for the median absolute size of price changes.

Iterating forward on the Euler equation and assuming that all prices last at most J periods, we derive

$$\sum_{j=0}^{J-1} \beta^j (1-\alpha)^j E_t \left[\frac{\partial \Pi_{i,t+j}}{\partial P_{i,t}^*} \right] = 0, \quad (27)$$

where the derivative of the profit function is expressed as

$$\frac{\partial \Pi_{i,t}}{\partial P_{i,t}^*} = (1-\theta) (P_{i,t}^*)^{-\theta} P_t^{\theta-1} \frac{M_t}{P_t} + \theta (P_{i,t}^*)^{\frac{-\theta}{\eta}-1} \frac{w_t}{\eta} (Z_{i,t})^{\frac{-1}{\eta}} P_t^{\frac{\theta}{\eta}} \left(\frac{M_t}{P_t} \right)^{\frac{1}{\eta}}. \quad (28)$$

We can then solve (27) for the optimal price:

$$P_{i,t}^* = \left(\frac{\theta \sum_{j=0}^{J-1} \beta^j E_t \left[(1-\alpha)^j \frac{w_{t+j}}{\eta} (Z_{i,t+j})^{\frac{-1}{\eta}} P_{t+j}^{\frac{\theta}{\eta}} \left(\frac{M_{t+j}}{P_{t+j}} \right)^{\frac{1}{\eta}} \right]}{\theta - 1 \sum_{j=0}^{J-1} \beta^j E_t \left[\frac{\omega_{j,t+j}}{\omega_{0,t}} P_{t+j}^{\theta-1} \frac{M_{t+j}}{P_{t+j}} \right]} \right)^{\frac{\eta}{\eta+\theta(1-\eta)}}. \quad (29)$$

Following DKW, we take a total derivative of the optimal pricing equation to show the determinants of an observed price change in the model:

$$\begin{aligned} d \ln P_{i,t}^* &= \chi_1 \left(\sum_{j=0}^{J-1} \rho_j E_t \left[\frac{-1}{\eta} d \ln Z_{i,t+j} + d \ln w_{t+j} + \chi_2 d \ln P_{t+j} + \frac{1-\eta}{\eta} d \ln M_{t+j} \right] \right) \\ &\quad + \chi_1 \sum_{j=0}^{J-1} \delta_j E_t [(\theta-1) d \ln P_{t+j} - d \ln M_{t+j}]. \end{aligned} \quad (30)$$

Here we use $\chi_1 \equiv \frac{\eta}{\eta+\theta(1-\eta)}$ and $\chi_2 \equiv \left(1 + \frac{(\theta-1)(1-\eta)}{\eta} \right)$. For sufficiently small rates of steady state inflation, such as in our model, ρ_j can be approximated by $\rho_j = \frac{\beta^j (1-\alpha)^j}{\sum_{h=0}^{J-1} \beta^h (1-\alpha)^h}$ and δ_j is approximately zero.⁹

The difficulty in using (30) to test the responsiveness of price changes to new versus old information is that, in the BLS data, we only observe price changes and inflation. We do not observe any disaggregate information nor do we have a good sense of what constitutes

⁹See the derivation in DKW.

an aggregate shock for the economy. Ideally, we would like to use an estimated process for exogenous monetary or technology shocks, and then test to see how long it takes prices to fully respond to those shocks. However, for most monetary shocks that have been identified in the literature, the aggregate price does not begin to respond to the shock until around six quarters have passed. Technology shocks are also difficult to consider because there is not a strong consensus on how best to identify them.

Therefore, we propose focusing on the change in price that is directly related to changes in the aggregate price level. The potential problem with this approach is that we will be ignoring all other aggregate variables to which firms may be responding. Ignoring the idiosyncratic information should not be as problematic because we will be using a large panel of observations in which idiosyncratic shocks should wash out.

Given that we observe prices that are fixed over discrete intervals, we modify (30) to explain the observed size of a price change in period t when the price was last adjusted τ periods ago:

$$\Delta \ln P_{i,t} = \chi_1 \chi_2 \sum_{j=0}^{J-1} \rho_j \left(E_t [\ln P_{i,t+j}] - E_{t-\tau_{i,t}} [\ln P_{i,t-\tau_{i,t}+j}] \right) + \Xi_{i,t}, \quad (31)$$

where $\Xi_{i,t}$ contains the additional terms in (30) not related to the aggregate price level.

Since inflation is the only aggregate variable we can use on the actual data, we do not use the firms' forecast rule from the model to evaluate expected changes in the price level in the simulated data. Instead, we search for an ARMA(p,q) specification that best fits simulated inflation dynamics. As a reminder, we consider four cases for the model that differ by the degree of staggering of aggregate information arrival, $\bar{n}_A \in \{0, 1, 2, 3\}$. Across these four cases, we find that an MA(3) specification best fits the inflation dynamics for the BLS data. This implies that $\ln P$ dynamics are best expressed by

$$\ln P_t = \mu + \ln P_{t-1} + \epsilon_t + \delta_1 \epsilon_{t-1} + \delta_2 \epsilon_{t-2} + \delta_3 \epsilon_{t-3}. \quad (32)$$

where the point estimates and standard errors, in parentheses, are $\mu = 0.008$ (0.001), $\delta_1 =$

0.13 (0.14), $\delta_2 = 0.21$ (0.15), $\delta_3 = 0.25$ (0.15), and the adjusted R^2 is 0.057.

With this specification for price-level dynamics, we can evaluate (31) as

$$\begin{aligned}
\Delta \ln P_{i,t} &= \chi_1 \chi_2 \sum_{s=0}^{\tau_{i,t}-1} \pi_{t-s} \\
&+ \chi_1 \chi_2 (1 - \rho_0) (\delta_1 \Delta_{\tau_{i,t}} \epsilon_t + \delta_2 \Delta_{\tau_{i,t}} \epsilon_{t-1} + \delta_3 \Delta_{\tau_{i,t}} \epsilon_{t-2}) \\
&+ \chi_1 \chi_2 (1 - \rho_0 - \rho_1) (\delta_2 \Delta_{\tau_{i,t}} \epsilon_t + \delta_3 \Delta_{\tau_{i,t}} \epsilon_{t-1}) \\
&+ \chi_1 \chi_2 (1 - \rho_0 - \rho_1 - \rho_2) (\delta_3 \Delta_{\tau_{i,t}} \epsilon_t) + \Xi_{i,t},
\end{aligned} \tag{33}$$

where $\Delta_{\tau_{i,t}} \epsilon_t \equiv \epsilon_t - \epsilon_{t-\tau_{i,t}}$.

To estimate this expression on the simulated data, define $PPC_{i,t}$ as the predicted price change due to new information on the aggregate price level since the previous change $\tau_{i,t}$ periods ago:

$$\begin{aligned}
PPC_{i,t} &= \sum_{s=0}^{\tau_{i,t}-1} \pi_{t-s} \\
&+ (1 - \rho_0) (\delta_1 \Delta_{\tau_{i,t}} \epsilon_t + \delta_2 \Delta_{\tau_{i,t}} \epsilon_{t-1} + \delta_3 \Delta_{\tau_{i,t}} \epsilon_{t-2}) \\
&+ (1 - \rho_0 - \rho_1) (\delta_2 \Delta_{\tau_{i,t}} \epsilon_t + \delta_3 \Delta_{\tau_{i,t}} \epsilon_{t-1}) \\
&+ (1 - \rho_0 - \rho_1 - \rho_2) (\delta_3 \Delta_{\tau_{i,t}} \epsilon_t).
\end{aligned} \tag{34}$$

Evaluating this expression using parameters from the model and the estimated MA process for inflation for each respective case of the model, we estimate the following regression on the simulated data:

$$\Delta \ln P_{i,t} = \gamma PPC_{i,t} + v_{i,t}.$$

To reiterate, this specification estimates the responsiveness of price changes to new information on inflation that has arrived since the previous change $\tau_{i,t}$ periods ago. In the baseline model, where firms always have current information on the aggregate state variables, we should expect an estimate of $\gamma = \chi_1 \chi_2 = 0.93$ if the omitted terms from equation (30) are

uncorrelated with inflation information.

The estimates from the four model cases are displayed in Table 5. The point estimates are means across 100 samples, and the “standard errors” are the standard deviations across the 100 samples. In the baseline, all firms have current information on aggregate state variables. In the case with 1 period of information stickiness (labeled Sticky 1), roughly one-half of firms that adjust their price have new information on aggregate state variables and one-half of firms have information that is one period old. In the case with 2 periods of information stickiness (labeled Sticky 2), one-third have new information, and so on. The estimate of γ in the baseline is 0.60, markedly lower than the expected 0.93. This discrepancy presumably reflects the various approximations we have made. The γ coefficient drops only modestly to 0.55 in Sticky 1, but by a larger amount to 0.37 in the Sticky 3 case. Given that the estimate of γ differs sharply in the baseline from its theoretical value, this does not appear to be the best way to gauge whether firms are fully responding to new information on the price level since they last changed their price.

The final row of Table 5 displays the estimate from the BLS micro data, based on over one million consumer price changes in the U.S. from 1988 through 2004. The estimate of γ is 0.76, actually higher than any of our model cases.¹⁰

¹⁰With posted prices rather than regular prices, there are over 1.4 million price changes. But the *ppc* coefficient is very similar at 0.77.

Table 5a: Response of price changes to price-level information in the sticky information model

		γ	R^2
Model	Baseline ($\bar{n}_A = 0$)	0.598 (0.110)	0.015
	Sticky 1 ($\bar{n}_A = 1$)	0.554 (0.092)	0.012
	Sticky 2 ($\bar{n}_A = 2$)	0.459 (0.110)	0.008
	Sticky 3 ($\bar{n}_A = 3$)	0.377 (0.139)	0.006
BLS Data		0.764 (0.010)	0.005

Note: The estimated parameter and R^2 shown for each model represent the average across 100 simulated panels, where each simulation consists of 6000 firms estimated for 51 periods. The standard deviation of the parameter estimate across the 100 panels is shown in parentheses.

The price responsiveness changes markedly when we switch from a model of sticky information to the alternative model where firms have equally old information. As a reminder, in this alternative model, all firms always possess information that is n_A periods old. Estimates in Table 5 show that firms are responding equally to “new” price-level information, even though the age of information changes for each case. This illustrates the endogeneity of inflation. If all firms have the same information, even if it is several periods old, they will still respond to any information updates in a similar way, and thus firms will appear to be responding to “new” information in price-level changes. The only way to determine the age of information would be to estimate the responsiveness to an exogenous shock.

Table 5b: Response of price changes to price-level information in the old information model

		γ	R^2
Model	Baseline ($\bar{n}_A = 0$)	0.598	0.015
		(0.110)	
	Old 1 ($n_A = 1$)	0.609	0.008
		(0.118)	
	Old 2 ($n_A = 2$)	0.531	0.008
		(0.077)	
	Old 3 ($n_A = 3$)	0.566	0.009
		(0.059)	
BLS Data		0.764	0.005
		(0.010)	

Note: The estimated parameter and R^2 shown for each model represent the average across 100 simulated panels, where each simulation consists of 6000 firms estimated for 51 periods. The standard deviation of the parameter estimate across the 100 panels is shown in parentheses.

In search of a sharper test for sticky information, we augment the estimation equation above to include lagged information. If firms all have current information on the aggregate state variables, then their price changes should not respond to information on lagged innovations beyond the innovations found in equation (34). However, if firms set their prices based on old information, then they should respond to the lagged information. In order to test this hypothesis, we add four lagged inflation innovation terms to the estimation equation:

$$\Delta \ln P_{i,t} = \gamma PPC_{i,t} + \lambda_1 \Delta_{\tau_{i,t}} \epsilon_{t-3} + \lambda_2 \Delta_{\tau_{i,t}} \epsilon_{t-4} + \lambda_3 \Delta_{\tau_{i,t}} \epsilon_{t-5} + \lambda_4 \Delta_{\tau_{i,t}} \epsilon_{t-6} + v_{i,t}. \quad (35)$$

The results of this estimation for the four cases of the model are displayed in Table 6a. In the baseline case, where all firms have current information, the coefficients on lagged innovations are either insignificant or negative, indicating that firms are not responding positively to old information. However, as the amount of information stickiness is increased

in cases Sticky 1 through Sticky 3, we find that the estimates of λ increase.

The final row of Table 6a displays estimates from the BLS data. Here we find that the lagged coefficients are mixed, with two insignificant, one significantly *negative*, and one (the oldest) very positive and significant. The posted price results (not shown) are much the same. These results provide modest support for sticky information models.

Table 6a: Response of price changes to new and old price-level information in the sticky information model

		γ	λ_1	λ_2	λ_3	λ_4	R^2
Model	Baseline ($\bar{n}_A = 0$)	0.580	-0.028	-0.022	-0.034	0.005	0.016
		(0.120)	(0.353)	(0.289)	(0.286)	(0.303)	
	Sticky 1 ($\bar{n}_A = 1$)	0.540	0.151	0.071	0.100	0.017	0.013
		(0.092)	(0.313)	(0.299)	(0.293)	(0.309)	
	Sticky 2 ($\bar{n}_A = 2$)	0.439	0.523	0.399	0.317	0.221	0.009
		(0.114)	(0.432)	(0.429)	(0.484)	(0.402)	
	Sticky 3 ($\bar{n}_A = 3$)	0.360	0.851	0.543	0.518	0.390	0.006
		(0.141)	(0.397)	(0.444)	(0.514)	(0.424)	
BLS Data		0.763	-0.008	-0.193	-0.042	0.364	0.005
		(0.011)	(0.033)	(0.034)	(0.035)	(0.034)	

Note: The estimated parameters and R^2 shown for each model represent the average across 100 simulated panels, where each simulation consists of 6000 firms estimated for 51 periods. The standard deviations of the parameter estimates across the 100 panels are shown in parentheses.

The results for the alternative model with old information are displayed in Table 6b. The results show strong responses of prices to old information in the cases with old information. However, there appears to be no clear pattern for the sign and magnitude of the responses. The response captured by λ_1 is strongly negative, where as λ_2 is strongly positive. As discussed above, the endogeneity of inflation appears to limit the effectiveness of this empirical test for the case in which firms have equally old information.

Table 6b: Response of price changes to new and old price-level information in the old information model

		γ	λ_1	λ_2	λ_3	λ_4	R^2
Model	Baseline ($\bar{n}_A = 0$)	0.580	-0.028	-0.022	-0.034	0.005	0.016
		(0.120)	(0.353)	(0.289)	(0.286)	(0.303)	
	Old 1 ($n_A = 1$)	0.479	-2.510	1.032	-1.301	-0.151	0.063
		(0.093)	(0.855)	(0.591)	(0.558)	(0.945)	
Old 2 ($n_A = 2$)		0.521	-0.570	0.121	-0.175	-0.012	0.012
		(0.071)	(0.443)	(0.383)	(0.394)	(0.427)	
Old 3 ($n_A = 3$)		0.521	-1.990	0.662	-0.819	-0.112	0.034
		(0.066)	(0.663)	(0.541)	(0.565)	(0.824)	
BLS Data		0.763	-0.008	-0.193	-0.042	0.364	0.005
		(0.011)	(0.033)	(0.034)	(0.035)	(0.034)	

Note: The estimated parameters and R^2 shown for each model represent the average across 100 simulated panels, where each simulation consists of 6000 firms estimated for 51 periods. The standard deviations of the parameter estimates across the 100 panels are shown in parentheses.

Thus far, we have tested for the general presence of sticky information. We next modify this test to explore whether there is evidence for rational inattention in particular. According to models in which firms face information processing constraints, one would expect that firms facing large idiosyncratic shocks would spend more effort processing idiosyncratic information and less effort processing aggregate information. In the micro data, we can plausibly test this theory by splitting the sample into four groups. Group 1 consists of products in the BLS data that have a below-median frequency of adjustment, $\overline{I(\Delta P \neq 0)}_L$, and below-median average absolute size of price changes, $|\overline{\Delta P}|_L$. Group 2 consists of products in the BLS data that have an above-median frequency of adjustment, $\overline{I(\Delta P \neq 0)}_H$, and an above-median average absolute size of price changes, $|\overline{\Delta P}|_H$. There are over 250,000 price changes in each group. If firms setting prices face convex information processing costs, we should expect to find that the products with more frequent and larger price changes respond less strongly to

new information and more strongly to old information.

The estimation results for these two groups are displayed in Table 6c. The estimate for the responsiveness to new information, γ , is well below one for Group 1, whereas it is way *above one* for Group 2. The coefficients on old information are negative or insignificant for Group 1, compared to mostly very positive for Group 2. To us, these results provide mixed support for rational inattention. Contrary to the theory, the group with less frequent, smaller shocks responds less to the latest information. Yet consistent with the theory, the group with the more frequent, larger shocks responds more positively to older information.¹¹ We plan to investigate this issue further, as it could be that shocks common to the Group 2 goods are driving more of the movements in the aggregate inflation rate, making it seem as though these goods are responding more to aggregate information than they are.

Table 6c: Response of price changes to new and old price-level information in BLS subsamples

		γ	λ_1	λ_2	λ_3	λ_4	R^2
BLS Data	Full Sample	0.763 (0.011)	-0.008 (0.033)	-0.193 (0.034)	-0.042 (0.035)	0.364 (0.034)	0.005
	$\overline{I(\Delta P \neq 0)}_L, \overline{ \Delta P }_L$	0.331 (0.007)	-0.173 (0.022)	-0.001 (0.022)	0.046 (0.022)	-0.104 (0.020)	0.010
	$\overline{I(\Delta P \neq 0)}_H, \overline{ \Delta P }_H$	3.856 (0.054)	1.505 (0.098)	0.248 (0.114)	-0.625 (0.120)	1.810 (0.105)	0.015
	$\overline{I(\Delta P \neq 0)}_H, \overline{ \Delta P }_L$	1.239 (0.022)	-0.048 (0.026)	-0.064 (0.031)	0.029 (0.032)	-0.059 (0.028)	0.013
	$\overline{I(\Delta P \neq 0)}_L, \overline{ \Delta P }_H$	0.674 (0.028)	-0.205 (0.100)	-0.314 (0.102)	0.485 (0.103)	-0.079 (0.102)	0.003

As a final test for sticky information, we break price changes into different types. We first consider regular price changes vs. sale-related price changes. Golosov and Lucas (2003)

¹¹Results with posted prices are broadly similar.

and Nakamura and Steinsson (2006) focus on regular price changes, i.e., those excluding temporary price discounts. Their rationale is that sales follow a sticky plan (e.g., 10% off Cheerios the first weekend of every month), and do not contribute to macro price flexibility. In our context, such sales should be purely idiosyncratic and unconnected from aggregate inflation. In other words, these price changes should not reflect new information on the aggregate price level. To test this hypothesis, we split the sample of price changes into those involving only regular prices (both the old and new prices are “regular” prices according to the BLS) and those involving a sales price (either the old and/or the new price is a “sale” price according to the BLS). In this breakdown, about 25% of the roughly 1.5 million price changes are sales-related. Because sales tend to be temporary, they should generate a pattern of price decreases after longer durations (high $\tau_{i,t}$ values) and price increases after shorter durations (low $\tau_{i,t}$ values), controlling for aggregate inflation (the ppc term). We therefore specify

$$\Delta \ln P_{i,t} = \gamma_1 PPC_{i,t} + \gamma_2 \tau_{i,t} + v_{i,t}.$$

In the baseline model, we are not sure what to expect for γ_2 or how the inclusion of $\tau_{i,t}$ should affect γ_1 . Similarly for regular price changes in the data. But we do expect the sales-related price changes to exhibit $\gamma_1 = 0$ and $\gamma_2 < 0$ in the data.

Table 7 presents results for the baseline model first, then results using the BLS data. Interestingly, in the baseline model we see $\gamma_1 > 1$ and $\gamma_2 < 0$. For regular price changes in the data we see the opposite, namely $\gamma_1 < 1$ and $\gamma_2 > 0$. But the results for sales-related price changes do not entirely conform to expectations either. We do find $\gamma_2 < 0$, as we expected. But we also find $\gamma_1 \gg 0$. Indeed, the coefficient on new information is not far below that for regular price changes. Thus it appears that sales are just as responsive to recent inflation as are regular price changes. Since sales tend to be temporary, the upshot is that their declines are not as deep and they give way to higher regular prices when recent inflation has been high. These results appear to undermine the hypothesis that sales do not reflect recent information on the aggregate price level.

Table 7: Response of sales and substitutions to new price-level information

		γ_1	γ_2	R^2
Model	Baseline	1.664	-0.010	0.024
		(0.196)	(0.0017)	
BLS Data	Full Sample	0.702	0.0022	0.004
		(0.021)	(0.0001)	
BLS Data	Regular	0.633	0.0044	0.009
		(0.020)	(0.0001)	
	Sales-Related	0.521	-0.0410	0.021
		(0.077)	(0.0005)	
BLS Data	Same Product	0.779	0.0024	0.005
		(0.021)	(0.0001)	
	Substitution-Related	0.236	-0.0013	0.000
		(0.125)	(0.0008)	

Finally, we split the sample of price changes into those related to product turnover, or “substitutions” in the BLS vernacular, and those involving precisely the same product. About 6% of all price changes involve substitutions in the BLS data. Golosov and Lucas (2003) and Nakamura and Steinsson (2006) likewise filter out these price changes. The regression results are in the bottom panel of Table 7. The same product regression looks similar to the full sample regression, although the coefficient on new information is higher (0.78 vs. 0.70). More striking, substitution-related price changes appear unrelated to recent inflation, or even the duration of the price. This finding supports the idea that substitutions reflect idiosyncratic or longer-range forces, rather than being responses to recent inflation. Bils and Klenow (2004) and Klenow and Kryvtsov (2005) find that filtering out such substitutions adds a month or two to the typical duration of a price.

5 Conclusion

Researchers are striving to develop micro foundations for apparently long-lasting real effects of nominal shocks. Nominal rigidities may be an important component, but prices do not appear to be sticky for long enough to do the job alone. Hence, Sims, Woodford and Mankiw-Reis have formulated theories in which macro information is stickier than micro prices. In Sims' incarnation the two are tightly related: micro shocks demand micro flexibility, thereby undercutting macro flexibility because of convex costs of processing all types of information.

In this paper we have argued that sticky information theories have testable implications for micro price changes. We use a simple GE models to demonstrate that the stickier the information, the older the inflation innovations firms respond to when they change prices. When we examine price changes in the U.S. CPI-RDB, we find little evidence that price changes reflect old information. However, our empirical test hinges on staggered information updating across firms; it cannot distinguish between a full information model and a model where firms have equally old information.

In addition, we find that items with larger, more frequent price changes respond more to recent information, but also more to older information. This lends only mixed support to Sims' hypothesis that the seeds of sticky macro information are sown by flexible micro information: the more micro shocks firms have to worry about, the less attention they pay to macro shocks. And finally, we find that sales-related price changes respond to macro information in much the same way that regular price changes do. This suggests that temporary price discounts should not be filtered out of data used for analysis of macroeconomic responsiveness.

References

- Bernanke, Ben S., Jean Boivin, and Piotr Elias**, “Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach,” NBER Working Paper 10220, 2004.
- Bils, Mark and Peter J. Klenow**, “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, October 2004, *112*, 947–985.
- Blanchard, Olivier and Nobuhiro Kiyotaki**, “Monopolistic Competition and the Effects of Aggregate Demand,” *American Economic Review*, September 1987, *77*, 647–666.
- Carvalho, Carlos**, “Heterogeneity in Price Stickiness and the New Keynesian Phillips Curve,” mimeo, Princeton University, 2006.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles Evans**, “Monetary Policy Shocks: What Have We Learned and to What End?,” *Handbook of Macroeconomics*, volume 1A, Elsevier: New York, 1999. John B. Taylor and Michael Woodford, eds.
- Dhyne, Emmanuel, Luis J. Alvarez, Herve Le Bihan, Giovanni Veronese, Daniel Dias, Johannes Hoffman, Nicole Jonker, Patrick Lunnemann, Fabio Ruller, and Jouko Vilmunen**, “Price Setting in the Euro Area: Some Stylized Facts from Individual Consumer Price Data,” European Central Bank Working Paper 524, 2005.
- Dotsey, Michael, Robert King, and Alexander Wolman**, “State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output,” *Quarterly Journal of Economics*, May 1999, *114*, 655–690.
- Golosov, Mikhail and Robert E. Lucas Jr.**, “Menu Costs and Phillips Curves,” NBER Working Paper 10187, 2003.
- Klenow, Peter J. and Jonathan L. Willis**, “Real Rigidities and Nominal Price Changes,” Federal Reserve Bank of Kansas City Research Paper No. 06-03, March 2006.

- and **Oleksiy Kryvtsov**, “State Dependent or Time Dependent Pricing: Does It Matter for Recent U.S. Inflation?,” NBER Working Paper 11043, 2005.
- Krusell, Per and Anthony Smith**, “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, October 1998, *106* (5), 867–96.
- Lucas, Robert E. Jr.**, “Some International Evidence on Output-Inflation Trade-Offs,” *American Economic Review*, 1973, *63*, 326–334.
- Mackowiak, Bartosz and Mirko Wiederholt**, “Optimal Sticky Prices Under Rational Inattention,” mimeo, Humboldt University Berlin, 2005.
- Mankiw, Gregory N. and Ricardo Reis**, “Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, November 2002, *117*, 1295–1328.
- and —, “Pervasive Stickiness,” *American Economic Review*, May 2006, *96* (2), 164–69.
- Nakamura, Emi and Jon Steinsson**, “Five Facts About Prices: A Reevaluation of Menu Cost Models,” mimeo, Harvard, 2006.
- Romer, David H. and Christina D. Romer**, “A New Measure of Monetary Shocks: Derivation and Implications,” NBER Working Paper 9866, 2003.
- Sims, Christopher**, “Stickiness,” *Carnegie-Rochester Conference Series on Public Policy*, December 1998, *49*, 317–56.
- , “Implications of Rational Inattention,” *Journal of Monetary Economics*, April 2003, *50* (3), 665–90.
- Tauchen, George**, “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions,” *Economics Letters*, 1986, *20*, 177–181.

Willis, Jonathan, “Estimation of Adjustment Costs in a Model of State-Dependent Pricing,” Federal Reserve Bank of Kansas City Research Paper No. 00-07, December 2000.

—, “General Equilibrium of a Monetary Model with State-Dependent Pricing,” unpublished paper, Federal Reserve Bank of Kansas City, 2003.

Woodford, Michael, *Interest and Prices*, Princeton University Press: Princeton, NJ, 2003.