

Sticky Information and Sticky Prices*

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

June 2007

Abstract

In the U.S. and Europe, prices change at least once a year. Yet nominal macro shocks seem to have real effects lasting well beyond a year. “Sticky information” models, as posited by Mankiw and Reis (2002), Sims (2003), and Woodford (2003), can reconcile micro flexibility with macro rigidity. We simulate a sticky information model in which price setters update information on macro shocks less frequently than information on micro shocks. We then examine price changes in the micro data underlying the U.S. CPI. Empirical price changes react to old information, just as sticky information models predict.

JEL: D8, E3, L16

Keywords: sticky information, state dependent pricing

*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, Andre Kurmann, John Leahy, and Chris Sims provided very useful comments. Josephine Smith provided excellent research assistance.

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

[‡]Federal Reserve Bank of Kansas City. Corresponding author: Jonathan Willis, Research Department, Federal Reserve Bank of Kansas City, 925 Grand Blvd., Kansas City, MO 64198; Phone: 816-881-2852; Email: jonathan.willis@kc.frb.org

1 Introduction

Individual consumer and producer prices change every six months to one year.¹ In contrast, many studies find that nominal macro shocks have real effects with a half-life well over a year.² “Sticky information” theories can reconcile macro price rigidity with micro price flexibility.³ These theories, advanced recently by Sims (1998, 2003), Mankiw and Reis (2002, 2006), and Woodford (2003), 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 Sims’ version of sticky information, 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. We seek to answer the question, do price changes reflect dated information on macro states? Answering this question is difficult given the lack of consensus on a measure of monetary policy shocks, especially one that explains inflation movements well. We therefore simulate simple GE models to derive responses of price changes to past inflation movements.

We simulate models featuring exogenous money growth, a cash-in-advance constraint, and monopolistically competitive firms. The firms face idiosyncratic productivity and 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 infrequency 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

¹See Bils and Klenow (2004) and Nakamura and Steinsson (2006) for U.S. evidence, and Dhyne et al. (2005) for studies of Euro Area countries.

²See, for example, Christiano, Eichenbaum and Evans (1999), Romer and Romer (2003), and Bernanke, Boivin and Eliasch (2004).

³Strategic complementarities can also generate a “contract multiplier”, i.e., real effects lasting well beyond price durations. We neglect such real rigidities to focus on sticky information theories.

on older innovations?

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 (1980). This model is closest to Mankiw-Reis in having periodic full updating of macro information. Our benchmark model also shares some of the spirit of Sims, however, in having firms observe their idiosyncratic shocks every period. As expected, the less frequent the updating of macro information in the model, 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 approximate the mean, standard deviation and serial correlation of inflation in the data. We choose the size of menu costs and the size of idiosyncratic productivity shocks to match the frequency and size of micro price changes in the data.

We test whether price changes in the data respond to old inflation innovations, or only those arriving since the firm last changed its price. We find evidence that price changes reflect macro inflation innovations older than they should according to the flexible information model. Our empirical regression results more closely resemble those obtained from our sticky information models than those from our flexible information model.⁴

We also examine whether specific types of price changes reflect macro information or, instead, purely idiosyncratic forces. The BLS labels each price as either a “sale” price or a “regular” price, and also keeps track of when products turn over (“substitutions”). Price changes related to sales and substitutions are often filtered out of price data by macro researchers (e.g., Golosov and Lucas (2007), and Nakamura and Steinsson (2006)) on the

⁴Knotek (2006) also concludes that a model containing both sticky information and sticky prices is consistent with micro and macro evidence.

grounds that they may reflect idiosyncratic considerations rather than macroeconomic information. We find that sales- and substitution-related price changes respond to macro information in much the same way that regular price changes do, which suggests that they should not be dropped from the data in macro studies.

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

2 Model

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 substitution between them. Monopolistically competitive firms produce goods to meet demand at their posted prices. 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 macro variables (exogenous and endogenous) arrives in staggered fashion. By changing the frequency of information arrival we can investigate different degrees of information stickiness. Finally, we assume that firms use a boundedly rational forecast for inflation. This assumption allows us to obtain a solution to the model with a finite state space.

2.1 Households

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

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

where L_t is labor input and C_t is the consumption good. We assume linear utility in order to reduce the number of aggregate states, which allows us to incorporate more heterogeneity while retaining computational feasibility. The consumption good is a Dixit-Stiglitz composite of individual goods with elasticity of substitution 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 are 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 using 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 each period in the form of money. Income consists of wages earned by working for firms at a per-period wage rate, W_t , profits from their ownership of firms, Π_t , real returns from bond holdings, r_t , and lump sum transfers of

money from the central bank, X_{t+1} .⁵ Income earned in period $t - 1$ provides money holdings for consumption in period t :

$$M_t = W_{t-1}L_{t-1} + \Pi_{t-1} + P_{t-1}(1 + r_{t-1})B_{t-1} + X_t. \quad (5)$$

The household budget constraint specifies that money spent on purchases in the current period not exceed money income earned in the previous period. Combining (3) and (5):

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

The solution to the household's optimization decision provides the demand function, real interest rate, and wage rate that firms use in their dynamic programming problem. Since the intertemporal marginal rate of substitution in consumption is 1 in equilibrium, the real interest rate is constant at $r = \frac{1-\beta}{\beta}$. 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 using bonds. We solve for an equilibrium in which households spend all money holdings on consumption in the current period, as bonds are in zero net supply. 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 a constant expected real

⁵We date the money transfer the following period to signify when it affects economic activity.

wage. 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:

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

Rearranging this condition, we can solve for the real wage in the current period as a function of the expected change in the price level:

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

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 . Firms are assumed to meet all demand at a given price, which implies that $Y_i = C_i$.

Contemporaneous real profits for firm i are given by

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

where P_i is the price for good i . 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 a fixed factor of production.

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

real profit function

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

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

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 the same 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

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

2.2.2 Information cost

To explore the implications of sticky information, we assume information regarding macro state variables arrives in a staggered fashion. If new information does not arrive, we assume the firm is not able to determine anything about the current innovation to money growth. This requires that pricing managers not interact with the production managers or accountants within the firm, otherwise they could see production or profits and draw inferences about current money innovations. 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 measurement error to the model. Firms would then solve a signal extraction problem when they do not have updated information. See Zbaracki, Bergen and Levy (2006), however, for a case study suggesting that limited communication between price setters and others within the firm is not an implausible assumption.

Ideally, we would like to specify a model in which firms face a cost of acquiring information about the macro state variables. In Reis (2006), firms decide each period whether to pay for updated information on the aggregate states. As we assume in this model, Reis establishes conditions under which firms find it optimal to update their information at fixed intervals.

Let \bar{A} be the number of periods between observing the aggregate money growth rate, inflation rate, and real money supply. For a given firm in a given period, let A represent the number of periods since aggregate information was last observed, i.e., the age of aggregate information. If a firm has updated information, then $A = 0$. Similarly, let \bar{I} represent the number of periods between observing idiosyncratic information. We set $\bar{I}=0$ so firms always have current information on their idiosyncratic productivity shock. This assumption follows in the spirit of Sims' rational inattention story, wherein firms pay more attention to idiosyncratic than aggregate shocks because the former are much larger.

2.3 Dynamic Optimization

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 timing of information updating impacts the state variables of the firm's optimization problem. The seven state variables are the firm's current nominal price relative to the aggregate price level the last time aggregate information was observed ($p_{i,A}$), the money growth rate when last observed ($g_{M,A}$), the inflation rate when last observed (π_A), the level of real money balances when last observed ($m_A \equiv \frac{M}{P_A}$), the idiosyncratic productivity index (Z_i), the age of aggregate information (A), and the information set Ω used to form future expectations of the endogenous state variables.

Given the state vector, $S = \{p_{i,A}, g_{M,A}, \pi_A, m_A, Z_i, A, \Omega\}$, 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

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

with $S' = \{p_{i,A'}^*, g'_{M,A'}, \pi'_{A'}, m'_{A'}, Z'_i, A', \Omega'\}$. The firm's value function is discounted by β , which reflects 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 the firm does not observe current information, 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 A periods ago, $p_{i,A}$, and integrates over all of the possible sequences of inflation over A periods conditional on information in the state vector. Regardless of the age of the information, the firm will always compute conditional expectations of the future value function. The firm chooses the nominal price relative to the price level A periods ago, $p_{i,A}^*$, that generates the highest expected value.

The value conditional on no price change is expressed as

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

with $S' = \{p_{i,A'}, g'_{M,A'}, \pi'_{A'}, m'_{A'}, Z'_i, A', \Omega'\}$.

For the exogenous state variables, money growth and idiosyncratic productivity, we as-

sume 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 reduce heterogeneity to a manageable scope, as in Dotsey, King and Wolman (1999) (hereafter DKW). An alternative is to assume firms form inflation expectations based on a limited set of information. We choose the latter for two reasons. First, the heterogeneity restrictions required for the DKW model do not match up well with the micro evidence.⁶ Second, due to the heterogeneity introduced by staggered updating of information, assuming bounded rationality helps keep the model tractable.

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

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

where $\nu_{\pi,t+1}$ is the forecast error. Firms will use their inflation forecast along with their forecast of money growth, from (18), to determine a forecast for 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

⁶See 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 VAR system

$$\pi_{t+1}^f = \alpha_0 + \alpha_1 \pi_t + \alpha_2 \ln m_t + \alpha_3 g_{M,t} + \nu_{\pi,t+1} \quad (23)$$

$$\begin{aligned} \ln m_{t+1}^f &= \mu_{g_M} - \alpha_0 - \alpha_1 \pi_t + (1 - \alpha_2) \ln m_t + (\rho_{g_M} - \alpha_3) g_{M,t} \\ &\quad + \nu_{g_M,t+1} - \nu_{\pi,t+1} \end{aligned} \quad (24)$$

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

The equilibrium solution of the model requires the selection of an appropriate inflation forecast rule, $\Theta = \{\alpha_0, \alpha_1, \alpha_2, \alpha_3\}$. Using this forecast rule, the firm solves the optimization problem in (15) by determining a policy function for the updating of prices: $p_{i,A}^* = f(p_{i,A}, g_{M,A}, \pi_A, m_A, Z_i, A, \Omega)$.

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 model simulations.

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

205 increments of 0.13%, or one-third 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).⁷ This conversion provides us with the transition matrix, $\Phi_Z(Z, Z')$, expressing the expected probability of any given realization of $Z_{i,t+1}$ as a function of the current state variables $Z_{i,t}$. The three-variable VAR for inflation, real money balances, and money growth is similarly converted into a first-order Markov chain.⁸ We use the transition matrices to compute the discounted expected value of the future period as well as expected contemporaneous profits if firms have out-of-date information. Another transition matrix, $\Phi_A(A, A')$, provides the probability of moving from information of age A in the current period to information of age A' next period. The parametrization of this matrix will determine the stickiness of macro information.

Table 1 displays the parameter values we use in our model simulations. We calibrate the structural parameters using information from the BLS price data and other sources. We use a bimonthly frequency (six periods per year) in order to match the sampling frequency in the BLS microdata. We set the discount rate, β , to 0.993 ($=0.96^{\frac{1}{6}}$), to arrive at a 4% annual real interest rate. We set the elasticity of substitution between different consumer goods, θ , to 5, corresponding to a 25 percent markup for the firm. This is at the intersection of values used in the IO (3-5) and macro (5-10) literatures. We set returns to scale in production, η , to 0.9. This is a compromise between the more conventional constant returns and labor's share of around 0.7, as we have only labor in the model. We set κ , the marginal disutility of labor divided by the discount rate, to 0.5. The results of interest from the model are not sensitive to changing κ .

We calibrate the remaining parameter values using statistics calculated with the BLS price data, which we discuss in more detail in the next full section. We set the parameters

⁷The discrete grid for idiosyncratic productivity contains 5 points spread equally in terms of the cumulative distribution function of the variable.

⁸The 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.

of the money growth process, $\{\mu_{gM}, \rho_{gM}, \sigma_{gM}\}$, to produce inflation dynamics similar to the data. A random walk for money ($\rho_{gM} = 0$) turns out to be the closest we can come to mimicking the low persistence of inflation in the data. Values of $\mu_{gM} = 0.0038$ and $\sigma_{gM} = 0.013$ allow us to closely match the mean and standard deviation of actual inflation. We base the persistence of the idiosyncratic productivity shock, ρ_Z , on estimates in Klenow and Willis (2006). In that study we looked at the persistence of relative prices within categories of consumption, thereby controlling for different industry price trends (e.g., computers vs. medical care). Translating our monthly serial correlation of 0.68 to our bimonthly frequency here results in $\rho_Z = 0.46$. Midrigan (2006) and Golosov and Lucas (2007) use similar values based on grocery scanner data and BLS data, respectively. We base σ_Z on the absolute size of price changes in the BLS data. Our value of $\sigma_Z = 0.083$ is likewise similar to values used in Klenow and Willis (2006), Midrigan (2006), and Golosov and Lucas (2007). Finally, we set the cost of implementing price changes, ψ , to 1.3% of firm revenue. Combined with the other parameter values, this enables us to match the frequency of price changes observed in the data of 30% per bimonth.

Following Willis (2003), we compute a rational expectations equilibrium of the model using the inflation forecasting rule expressed in (20). For a given specification of the structural parameters along with the inflation forecasting parameters, $\Theta = \{\alpha_0, \alpha_1, \alpha_2, \alpha_3\}$, we solve the model and generate the policy function. We then use the policy function to simulate a panel of 6,000 firms over 500 periods (bimonths).⁹

Simulating data from the model requires an updating process to determine the evolution of the endogenous aggregate-level state variables. The collective actions of firms in the simulation determines the aggregate inflation rate and the level of real money balances.

⁹The size of the panel was chosen as follows. First, the cross-section should have a large number of firms given the 80,000+ price observations per period in the BLS data. 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 yield close to “asymptotic” results. We found that lengthening the sample beyond 500 did not materially affect the simulated moments. Below we do not report standard errors from regressions based on simulated data, as they are quite small given our large panels.

When setting prices in the current period, firms with updated information, $A = 0$, possess the current values 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 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 inflation forecasting rule. An OLS regression of the forecasting rule in (20) is run with simulated values for inflation, real money balances, and money growth. The initial values of the forecast parameters, Θ_0 , are 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 continues until a fixed point is reached. The fixed point represents a bounded rational expectations equilibrium wherein the inflation forecasting rule assumed by firms matches up with the simulated data.

2.3.3 Sticky Information

The setting for \bar{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 consequences of information stickiness, we will consider four cases corresponding to the maximum age of aggregate information ranging from 0 to 3 periods: $\bar{A} \in \{0, 1, 2, 3\}$.

We assume firms always have current information on their idiosyncratic shocks ($\bar{I} = 0$). This assumption allows us to focus on the implications of aggregate information stickiness.

To illustrate the role of sticky information, Figures 1 and 2 display responses of inflation and real output to a 1 percent shock to money growth. Figure 1 shows 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.

Each of the four cases has a different equilibrium inflation process and hence different parameters values in the forecast rule (20). The parameters for each case are displayed in Table 2a. The coefficients vary modestly, but the rule’s explanatory power is enhanced by sticky information as it makes inflation more persistent.¹⁰

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{A} periods old.

As before, we 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 aggregate 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 2b 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 as the forecast rule only contains information lagged one period, and it suggests 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.

¹⁰Following Krusell and Smith (1998), we checked whether additional variables would improve the inflation forecasts. We regressed simulated inflation forecast errors on additional lags of the state variables as well lagged dispersion of prices (i.e., the standard deviation of relative prices in the previous period). We found that these variables did little to explain the forecast errors.

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 80,000 items are surveyed each month, with an item being a specific product (brand and detailed features) sold by a particular outlet. The data is 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.¹¹ 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.

The BLS collects prices monthly for food and energy items in all areas, and for all items in New York, Los Angeles, and Chicago. For other areas, they check prices bimonthly for “core” items (items other than food or energy). Each bimonthly 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 bimonthly dataset. We label half the monthly items odds and half evens, and follow their odd or even prices accordingly. The disadvantage is that we are ignoring half the price quotes for monthly items. Yet in so doing we incorporate the 80,000 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 three cities. Just as important, looking at bimonths rather than months allows us to consider models with greater stickiness of information without adding as many states (e.g., three bimonths as opposed to six months).

To pin down key parameters in our model, we calculate five statistics from the CPI data.

¹¹See Klenow and Kryvtsov (2005) for a more detailed description of the CPI Research Database.

Three are the mean, standard deviation, and serial correlation of the aggregate bimonthly inflation rate. In terms of our model, these help us in 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, let P_{sit} denote the price of item i in sector s in bimonth t , and ω_{sit} the BLS weight on item i within category s in bimonth t . The weights in sector s sum to ω_s^{95} in every bimonth, the BLS consumption expenditure weight of category s in 1995 (which themselves sum to 1). We then define the aggregate inflation rate in bimonth t to be

$$\pi_t = \sum_s \sum_i \omega_{sit} [\ln(P_{sit}) - \ln(P_{sit-1})]. \quad (26)$$

When we calculate model moments for inflation, we use this geometric mean inflation.

We then take the simple average across the 101 bimonths from 1988 through 2004 to arrive at 0.384% per bimonth (2.3% per year) for inflation:

$$\mu_\pi = \sum_{t=1}^{101} \pi_t / 101 = 0.00384. \quad (27)$$

In similar fashion we calculate the standard deviation (0.397%) and serial correlation (0.170) of the inflation rate:

$$\sigma_\pi = \sqrt{\sum_{t=1}^{101} (\pi_t - \mu_\pi)^2 / 100} = 0.00397. \quad (28)$$

$$\rho_\pi = \sqrt{\sum_{t=1}^{100} (\pi_t - \mu_\pi)(\pi_{t-1} - \mu_\pi) / 99} = 0.170. \quad (29)$$

Our fourth moment is the fraction of items changing price from one bimonth to the next. Let $I(\Delta P_{sit} \neq 0)$ be a price-change indicator for item i in sector s in bimonth t , which equals

1 if the item changed price from bimonth $t - 1$ to t , and 0 otherwise. We calculate the mean of this indicator for an item, then take the weighted median value across items to arrive at 0.300 (30.0% per bimonth). Easier to express explicitly is the cousin of this statistic, the weighted *mean* frequency of price changes, which is higher at 38.0%:

$$\overline{I(\Delta P \neq 0)} = \sum_s \sum_i \omega_{si} \frac{\sum_t I(\Delta P_{sit} \neq 0)}{\sum_t 1} = 0.380. \quad (30)$$

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 better approximate a model with heterogeneity. Bills and Klenow (2004) examine this for the Taylor model, and Carvalho (2006) for the Calvo model.

Our final moment is the weighted median absolute size of price changes, which is 0.0853 (8.53%). Again, this value is easier to explicitly define as the weighted *mean*, which is higher at 12.0%:

$$\overline{|\Delta P|} = \sum_s \sum_i \omega_{si} \frac{\sum_t |\Delta P_{sit}|}{\sum_t I(\Delta P_{sit} \neq 0)} = 0.120. \quad (31)$$

As stressed by Klenow and Kryvtsov (2005) and Golosov and Lucas (2007), absolute price changes are much larger than needed to keep up with the trend inflation rate. The trend is about 0.4% per bimonth and the frequency of price changes is around 1/3, so price changes only need average about 1.2% to keep up with trend inflation. Yet the average price change is an order of magnitude larger 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 around 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 much less importance.

In Table 3 we collect these moments. Klenow and Willis (2006) report small bootstrapped

standard errors for similar statistics, which reflect the large number of observations underlying them (about 8 million prices and 3 million price changes). We also give the corresponding moments in our baseline model. We chose the parameter values in our baseline model to try to match these moments. The moments from the baseline model match the empirical moments well, with the exception of the serial correlation of inflation. Even without sticky information and with iid money growth, sticky prices generate more persistent inflation than observed in the data. Making money growth negatively correlated over time actually increases the serial correlation of model inflation, so iid money growth produces the closest serial correlation to the low level in the data.

4 Simulation and Estimation

We now devise a test to empirically discriminate flexible and sticky information. To do so, we first express firm price changes as a function of variables in the information set for the “null” flexible information model.

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

$$\frac{\partial \Pi_{i,t}}{\partial P_{i,t}^*} + \beta E_t \left[(1 - \vartheta) \frac{\partial V(S_{i,t+1})}{\partial P_{i,t}^*} \right] = 0, \quad (32)$$

where ϑ is the probability of a firm changing its price. Here we make the 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). A flat hazard rate 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.¹²

¹²This begs the question of why we did not assume Calvo pricing to begin with. We chose to model state dependent pricing because of the “selection effect” in who chooses to change prices. We would like our test

Iterating forward on the Euler equation and assuming all prices last at most J periods:

$$\sum_{j=0}^{J-1} \varrho^j E_t \left[\frac{\partial \Pi_{i,t+j}}{\partial P_{i,t}^*} \right] = 0, \quad (33)$$

where $\varrho = \beta(1 - \vartheta)$. The derivative of the profit function is expressed as

$$\frac{\partial \Pi_{i,t+j}}{\partial P_{i,t}^*} = (1 - \theta) \Upsilon_{t+j} P_{i,t}^{* - \theta} + \theta \Psi_{i,t+j} P_{i,t}^{* \frac{-\theta}{\eta} - 1}. \quad (34)$$

with Υ_{t+j} representing terms associated with marginal revenue ($\Upsilon_{t+j} \equiv P_{t+j}^{\theta-1} \frac{M_{t+j}}{P_{t+j}}$) and $\Psi_{i,t+j}$ representing terms associated with marginal cost ($\Psi_{i,t+j} \equiv \frac{\kappa Z_{i,t+j}^{\frac{-1}{\eta}}}{\eta P_{t+j} E_t \left[\frac{1}{P_{t+j+1}} \right]} P_{t+j}^{\frac{\theta}{\eta}} \left(\frac{M_{t+j}}{P_{t+j}} \right)^{\frac{1}{\eta}}$).

We can then solve (33) for the optimal price:

$$P_{i,t}^* = \left(\frac{\theta}{\theta - 1} \frac{\sum_{j=0}^{J-1} \varrho^j E_t \Psi_{i,t+j}}{\sum_{j=0}^{J-1} \varrho^j E_t \Upsilon_{t+j}} \right)^{\chi_1} \quad (35)$$

where $\chi_1 \equiv \frac{\eta}{\eta + \theta(1 - \eta)}$.

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

$$\begin{aligned} d \ln P_{i,t}^* &= \chi_1 \left(\sum_{j=0}^{J-1} \rho_j E_t \left[\chi_2 d \ln P_{t+j} + d \ln P_{t+j+1} + \frac{1 - \eta}{\eta} d \ln \left(\frac{M_{t+j}}{P_{t+j}} \right) + \frac{-1}{\eta} d \ln Z_{i,t+j} \right] \right) \\ &+ \chi_1 \sum_{j=0}^{J-1} \zeta_j E_t \left[(\theta - 1) d \ln P_{t+j} - d \ln \left(\frac{M_{t+j}}{P_{t+j}} \right) \right]. \end{aligned} \quad (36)$$

where $\chi_2 \equiv \frac{\theta(1 - \eta)}{\eta}$. For sufficiently low steady state inflation, as in our model, ρ_j can be approximated by $\rho_j = \frac{\beta^j (1 - \vartheta)^j}{\sum_{h=0}^{J-1} \beta^h (1 - \vartheta)^h}$ and ζ_j is approximately zero.¹³

The difficulty in using (36) to test the responsiveness of price changes to new versus old information is that we only observe price changes and inflation in the BLS data. We do not

to reveal the presence of sticky information even if such selection operates. We will check this by running our test on simulated data from state dependent pricing models.

¹³See the derivation in DKW.

observe any disaggregate information nor do we have a good sense of what constitutes an aggregate nominal shock for the economy. We would like to use an estimated process for exogenous monetary and/or technology shocks, and then test to see how long it takes prices to fully respond to those shocks. Such shocks are difficult to consider, however, because there is no consensus on how best to identify them. Moreover, existing identification strategies have had more success replicating empirical output dynamics than inflation dynamics.

As an alternative, we focus on the change in price one would expect based only on current information about inflation. A drawback 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 (the selection effect being an important caveat here).

Given that we observe prices fixed over discrete intervals, we modify (36) to explain the observed size of a price change in period t when the price was last adjusted τ periods ago:

$$\begin{aligned} \Delta \ln P_{i,t} = & \chi_1 \sum_{j=0}^{J-1} \rho_j \left(\chi_2 \left(E_t [\ln P_{t+j}] - E_{t-\tau_{i,t}} [\ln P_{t-\tau_{i,t}+j}] \right) + \right. \\ & \left. \left(E_t [\ln P_{t+j+1}] - E_{t-\tau_{i,t}} [\ln P_{t-\tau_{i,t}+j+1}] \right) \right) + \Xi_{i,t} \end{aligned} \quad (37)$$

where $\Xi_{i,t}$ contains the additional terms in (36) corresponding to real money balances and the idiosyncratic productivity shock.

Since inflation is the only aggregate variable we can use in 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.¹⁴ Instead, we search for an ARMA(p,q) specification that best captures inflation dynamics in the baseline model. We find that an MA(4) specification maximizes

¹⁴Although we could use money stock data to construct series for the real money supply and money growth innovations, we do not do so because in our model "money" is merely a stand-in for a variety of macro shocks that push around the inflation rate.

the adjusted R^2 . This implies that $\ln P$ dynamics are expressed by

$$\ln P_t = \mu + \ln P_{t-1} + \epsilon_t + \sum_{j=1}^4 \delta_j \epsilon_{t-j}. \quad (38)$$

Table 4 presents the moving average coefficients in the data and in the baseline model. Note that we estimate an MA(4) process in the data as well as in the model. An MA(7) actually fits better in the data, but this could reflect precisely the sticky information we want to identify. Rather than incorporate such lagged information into the “flexible information” predicted price change, we maintained the same order MA(4) in the data as in the model.

With this specification for price-level dynamics, we can evaluate (37) as

$$\Delta \ln P_{i,t} = \sum_{s=0}^{\tau_{i,t}-1} \pi_{t-s} + \chi_1 \sum_{j=0}^3 \chi_3 \Delta_{\tau_{i,t}} \epsilon_{t-j} + \Xi_{i,t} \quad (39)$$

where $\Delta_{\tau_{i,t}} \epsilon_t \equiv \epsilon_t - \epsilon_{t-\tau_{i,t}}$ and $\chi_3 = \sum_{k=0}^{3-j} \left((1 + \chi_2) \left(1 - \sum_{l=0}^{k-1} \rho_l \right) - \chi_2 \rho_k \right) \delta_{j+k+1}$. The MA terms affect price changes because they help forecast future inflation. Price setters wish to respond to forecastable movements in the aggregate price level over the life of a price.

To simplify, 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:

$$PPC_{i,t} = \sum_{s=0}^{\tau_{i,t}-1} \pi_{t-s} + \chi_1 \sum_{j=0}^3 \chi_3 \Delta_{\tau_{i,t}} \epsilon_{t-j} \quad (40)$$

Evaluating this expression using the estimated MA(4) for inflation for each case of the model, we run the following regression on the simulated data:

$$\Delta \ln P_{i,t} = \gamma PPC_{i,t} + v_{i,t}. \quad (41)$$

To reiterate, this specification estimates how price changes respond to inflation that has

accumulated since the previous change $\tau_{i,t}$ periods ago. In the baseline model, wherein firms always have current information on the aggregate state variables, we expect an estimate of $\gamma = 1$ if the omitted terms from equation (36) are uncorrelated with inflation information. This test has a close antecedent in Reis (forthcoming), who regresses consumption growth on income innovations and shows that the coefficient falls as information becomes stickier.

The estimates from four model cases are displayed in Table 5a. In the baseline case, all firms have current information on aggregate state variables. In the case with 1 period of information stickiness (labeled Sticky 1), one-half of firms 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 0.55, markedly lower than the unit value in our specification. This discrepancy presumably reflects the various approximations we have made (linearization, flat hazard, discrete grids, MA(4) forecast rule) plus the selection effect and omitted variables. But the γ coefficient is not uniform across the baseline and sticky information models: as information becomes stickier, the γ coefficient steadily falls to 0.22 in the Sticky 3 case. The older the information, the less related price changes are to the price change predicted under flexible information. The final row of Table 5a displays the estimate from the BLS micro data based on over 3 million consumer price changes from 1988 through 2004. The estimate of γ is 0.61, modestly above the baseline, flexible information case.

Table 5b shows similar behavior of the γ coefficient when firms have equally old information. In this alternative model, firms always possess information that is A periods old.

To try to gauge the age of 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 innovations older than those found in equation (36). If firms set their prices based on old information, however, then they should respond to the lagged information. In order to test this hypothesis, we add six lagged

inflation innovation terms (one year of old information) to the estimation equation:

$$\Delta \ln P_{i,t} = \gamma PPC_{i,t} + \sum_{j=1}^6 \lambda_j \Delta \tau_{i,t} \epsilon_{t-3-j} + v_{i,t}. \quad (42)$$

Estimation results for four cases are displayed in Table 6a. In the baseline case, where all firms have current information, the coefficients on lagged innovations are small and often negative, indicating that firms are not putting a lot of weight on old information. However, as the amount of information stickiness is increased in cases Sticky 1 through Sticky 3, we find that the λ coefficients steadily increase. This result is true only for the three information lags, corresponding to the degree of information stickiness in each case. If we simulated older information, however, we would presumably see additional lags attracting higher coefficients.

The final row of Table 6a displays estimates from the BLS data. Here we find some very positive and significant coefficients on old information terms. Five of the six appear economically and statistically significant when compared to the predictions of the baseline vs. sticky information models. These results provide evidence consistent with information being up to a year old.

The results for the alternative model with old information are displayed in Table 6b. The pattern is not as clear: the old information coefficients are not as significant, and do not increase steadily with the degree of information stickiness. The empirical results, therefore, appear more in line with staggered information than with uniformly old information.

4.1 Responses of sales and substitutions to aggregate information

We end by testing whether certain types of price changes exhibit an extreme form of sticky information: namely, that they reflect no macro information at all. We first consider regular price changes vs. sale-related price changes. Golosov and Lucas (2007) and Nakamura and Steinsson (2006) focus on regular price changes by excluding temporary price discounts.

Their rationale is that sales may follow a sticky plan (e.g., 10% off Cheerios the first weekend of every month). In our context, such sales should be purely idiosyncratic and unconnected from aggregate inflation. 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 1 million of the roughly 3 million price changes are sales-related. Given that many sales are temporary, sale-related price changes might, by construction, be negatively correlated with cumulative inflation since the last price change for an item. We therefore add a “down” dummy for regular-to-sale price changes and an “up” dummy for sale-to-regular price changes:

$$\Delta \ln P_{i,t} = \gamma_1 PPC_{i,t} + \gamma_{down} D_{i,t} + \gamma_{up} U_{i,t} + v_{i,t}. \quad (43)$$

Table 7 presents the results. For the full sample the dummies have the expected sign and improve the fit dramatically. Their inclusion more than doubles the coefficient on the predicted price change to about 1.3. When we look at regular price changes alone (those not involving sales prices), the coefficient is approximately equal to 1. For sale-related price changes, the down and up dummies are helpful as expected. But, perhaps surprisingly, the coefficient on macro information is over 1.8. Thus it appears that sales are at least 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.

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 7% of all price changes involve substitutions in the BLS data. Golosov and Lucas

(2007) 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 (1.21 for same product price changes only vs. 1.31 for all price changes). More striking, substitution-related price changes appear less related to recent inflation (0.66 at turnover vs. 1.21 within-product). This finding supports the idea that substitutions reflect some idiosyncratic or longer-range forces, rather than being responses to recent inflation. Still, substitution-related price changes are very related to macro price trends and should probably not be excluded from macro research on price stickiness.

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.

We have argued that sticky information theories have testable implications for micro price changes. Simple GE models demonstrate that the stickier the information, the older the inflation innovations firms respond to when they change prices. Just as these theories predict, price changes in the U.S. CPI microdata reflect information older than predicted by a flexible information model.

In addition, we find that sale-related price changes respond to macro information at least as much as regular price changes do. This suggests that sale prices should not be filtered out of data used for analysis of macroeconomic responsiveness. More muted statements apply to substitution-related price changes, which respond very much to overall inflation, but still half as much as price changes at other times.

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 Lunnermann, 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,” *Journal of Political Economy*, 2007, *115* (2), 171–199.

- 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.
- Knotek, Edward S.**, “A Tale of Two Rigidities: Sticky Prices in a Sticky-Information Environment,” Federal Reserve Bank of Kansas City Research Paper No. 06-15, December 2006.
- Krusell, Per and Anthony Smith**, “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, October 1998, 106 (5), 867–96.
- 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.
- Midrigan, Virgiliu**, “Menu Costs, Multiproduct Firms, and Aggregate Fluctuations,” mimeo, Federal Reserve Bank of Minneapolis, 2006.
- Nakamura, Emi and Jon Steinsson**, “Five Facts About Prices: A Reevaluation of Menu Cost Models,” mimeo, Harvard, 2006.
- Reis, Ricardo**, “Inattentive Producers,” *Review of Economic Studies*, 2006, 73 (3), 793–821.
- , “Inattentive Consumers,” *Journal of Monetary Economics*, forthcoming.
- 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.
- Taylor, John B.**, “Aggregate Dynamics and Staggered Contracts,” *Journal of Political Economy*, February 1980, *88* (1), 1–23.
- 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.
- Zbaracki, Mark J., Mark Bergen, and Daniel Levy**, “The Anatomy of a Price Cut: Discovering Organizational Sources of the Costs of Price Adjustment,” mimeo, Bar-Ilan University, 2006.

Table 1: Parameter Values

| Parameters not based on BLS CPI data | |
|--|--------|
| Discount rate (β) | 0.993 |
| Elasticity of substitution (θ) | 5 |
| Returns to scale (η) | 0.9 |
| Parameters calibrated using BLS CPI data | |
| Average money growth (bimonthly) (μ_{gM}) | 0.0038 |
| Serial correlation of money growth (ρ_{gM}) | 0 |
| Std. dev. of innovation to money growth ($\sigma_{\nu_{gM}}$) | 0.013 |
| Serial correlation of idiosyncratic productivity (ρ_Z) | 0.46 |
| Std. dev. of innovation to idiosyncratic productivity (σ_{ν_Z}) | 0.083 |
| Implementation (menu) cost (ψ) | 0.013 |

Notes: The frequency of the model is bimonthly to match the sampling frequency of the BLS CPI survey. The implementation cost is expressed as a fraction of steady state revenue.

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

| Model | α_1 | α_2 | α_3 | R^2 |
|----------------------------|------------|------------|------------|-------|
| Baseline ($\bar{A} = 0$) | 0.04 | 0.22 | -0.01 | 0.35 |
| Sticky 1 ($\bar{A} = 1$) | 0.08 | 0.22 | 0.02 | 0.79 |
| Sticky 2 ($\bar{A} = 2$) | -0.05 | 0.25 | -0.03 | 0.81 |
| Sticky 3 ($\bar{A} = 3$) | -0.05 | 0.25 | -0.03 | 0.79 |

Notes: This table provides parameters for the following forecast rule used by firms in four separate model specifications: $\pi_{t+1}^f = \alpha_0 + \alpha_1 \pi_t + \alpha_2 \ln m_t + \alpha_3 g_{M,t}$. In each model specification, information on aggregate state variables arrives on a *staggered*, deterministic schedule. \bar{A} indicates the maximum age of information before updating occurs. The R^2 provides the fit of the forecast rule on the simulated data.

Table 2b: Equilibrium forecast rules for model with old information

| Model | α_1 | α_2 | α_3 | R^2 |
|----------------------------|------------|------------|------------|-------|
| Baseline ($\bar{A} = 0$) | 0.04 | 0.22 | -0.01 | 0.35 |
| Old 1 ($\bar{A} = 1$) | -0.05 | 0.25 | -0.03 | 0.94 |
| Old 2 ($\bar{A} = 2$) | -0.30 | 0.27 | -0.17 | 0.72 |
| Old 3 ($\bar{A} = 3$) | -0.11 | 0.21 | -0.07 | 0.40 |

Notes: This table provides parameters for the following forecast rule used by firms in four separate model specifications: $\pi_{t+1}^f = \alpha_0 + \alpha_1\pi_t + \alpha_2 \ln m_t + \alpha_3 g_{M,t}$. Within each model specification, information on aggregate state variables is *identical* for all firms. \bar{A} indicates the age of aggregate information. The R^2 provides the fit of the forecast rule on the simulated data.

Table 3: Moments

| | μ_π | σ_π | ρ_π | $\overline{I(\Delta P \neq 0)}$ | $\overline{ \Delta P }$ |
|----------------|-----------|--------------|------------|---------------------------------|-------------------------|
| BLS CPI Data | 0.0038 | 0.0040 | 0.170 | 0.300 | 0.0853 |
| Baseline Model | 0.0039 | 0.0042 | 0.510 | 0.301 | 0.0865 |

Notes: The data moments are computed using the BLS CPI survey with a bimonthly frequency from 1988 to 2004. μ_π is weighted mean bimonthly inflation for items in the CPI survey, σ_π is the standard deviation of inflation, ρ_π is the serial correlation of inflation, $\overline{I(\Delta P \neq 0)}$ is the median frequency of price adjustment, and $\overline{|\Delta P|}$ is the median absolute size of a price change.

Table 4: Moving-average representation for inflation

| | δ_1 | δ_2 | δ_3 | δ_4 | R^2 |
|----------------|------------|------------|------------|------------|-------|
| BLS CPI Data | 0.17 | 0.07 | -0.06 | -0.03 | 0.038 |
| | (0.10) | (0.10) | (0.10) | (0.10) | |
| Baseline Model | 0.46 | 0.31 | 0.22 | 0.10 | 0.264 |

Notes: This table displays estimates for the following MA(4) inflation specification: $\pi_t = \mu + \epsilon_t + \sum_{j=1}^4 \delta_j \epsilon_{t-j}$. The mean inflation rate, μ , is shown in Table 3. BLS CPI data are measured at a bimonthly frequency from 1988 to 2004.

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

| model | | γ | R^2 |
|----------|----------------------------|----------|-------|
| Model | Baseline ($\bar{A} = 0$) | 0.550 | 0.009 |
| | Sticky 1 ($\bar{A} = 1$) | 0.466 | 0.006 |
| | Sticky 2 ($\bar{A} = 2$) | 0.330 | 0.003 |
| | Sticky 3 ($\bar{A} = 3$) | 0.219 | 0.001 |
| BLS Data | | 0.606 | 0.001 |
| | | (0.016) | |

Notes: This table displays the coefficient from a regression of price changes on the predicted price change due to new information on the aggregate price level since the previous change. For details on this specification, see equations (39) and (41). In each model specification, information on aggregate state variables arrives on a *staggered*, deterministic schedule. \bar{A} indicates the maximum age of information before updating occurs. A panel of 6000 firms and 500 periods is simulated for each model. BLS CPI data are measured at a bimonthly frequency from 1988 to 2004.

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

| | | γ | R^2 |
|----------|----------------------------|----------|-------|
| Model | Baseline ($\bar{A} = 0$) | 0.550 | 0.009 |
| | Old 1 ($\bar{A} = 1$) | 0.367 | 0.003 |
| | Old 2 ($\bar{A} = 2$) | 0.212 | 0.001 |
| | Old 3 ($\bar{A} = 3$) | 0.159 | 0.001 |
| BLS Data | | 0.606 | 0.001 |
| | | (0.016) | |

Notes: This table displays the coefficient from a regression of price changes on the predicted price change due to new information on the aggregate price level since the previous change. For details on this specification, see equations (39) and (41). Within each model specification, information on aggregate state variables is *identical* for all firms. \bar{A} indicates the age of aggregate information. A panel of 6000 firms and 500 periods is simulated for each model. BLS CPI data are measured at a bimonthly frequency from 1988 to 2004.

Table 6a: Response to new and old information in the sticky information model

| | | γ | λ_1 | λ_2 | λ_3 | λ_4 | λ_5 | λ_6 | R^2 |
|----------|----------------------------|----------|-------------|-------------|-------------|-------------|-------------|-------------|-------|
| Model | Baseline ($\bar{A} = 0$) | 0.553 | 0.017 | 0.024 | -0.090 | -0.105 | -0.173 | -0.099 | 0.009 |
| | Sticky 1 ($\bar{A} = 1$) | 0.462 | 0.261 | 0.091 | -0.003 | -0.123 | -0.191 | -0.176 | 0.006 |
| | Sticky 2 ($\bar{A} = 2$) | 0.315 | 0.526 | 0.333 | 0.164 | -0.009 | -0.046 | -0.016 | 0.003 |
| | Sticky 3 ($\bar{A} = 3$) | 0.203 | 0.500 | 0.457 | 0.307 | 0.038 | 0.114 | 0.050 | 0.001 |
| BLS Data | | 0.455 | 0.730 | 0.164 | 1.020 | 0.332 | -0.018 | 0.416 | 0.001 |
| | | (0.018) | (0.037) | (0.041) | (0.044) | (0.044) | (0.042) | (0.038) | |

Notes: This table displays the coefficients from a regression of price changes on the predicted price change due to new information on the aggregate price level since the previous change ($\gamma PPC_{i,t}$) and additional lags of inflation innovations ($\lambda_j \Delta_{\tau_{i,t}} \epsilon_{t-3-j}$). For details on this specification, see equation (42). In each model specification, information on aggregate state variables arrives on a *staggered*, deterministic schedule. \bar{A} indicates the maximum age of information before updating occurs. A panel of 6000 firms and 500 periods is simulated for each model. BLS CPI data are measured at a bimonthly frequency from 1988 to 2004.

Table 6b: Response to new and old information in the old information model

| | | γ | λ_1 | λ_2 | λ_3 | λ_4 | λ_5 | λ_6 | R^2 |
|----------|----------------------------|----------|-------------|-------------|-------------|-------------|-------------|-------------|-------|
| Model | Baseline ($\bar{A} = 0$) | 0.553 | 0.017 | 0.024 | -0.090 | -0.105 | -0.173 | -0.099 | 0.009 |
| | Old 1 ($\bar{A} = 1$) | 0.363 | 0.211 | 0.069 | 0.050 | -0.118 | -0.069 | -0.027 | 0.004 |
| | Old 2 ($\bar{A} = 2$) | 0.202 | 0.355 | 0.071 | 0.145 | -0.072 | 0.082 | 0.060 | 0.001 |
| | Old 3 ($\bar{A} = 3$) | 0.152 | 0.162 | -0.248 | 0.202 | 0.100 | -0.154 | 0.003 | 0.001 |
| BLS Data | | 0.455 | 0.730 | 0.164 | 1.020 | 0.332 | -0.018 | 0.416 | 0.001 |
| | | (0.018) | (0.037) | (0.041) | (0.044) | (0.044) | (0.042) | (0.038) | |

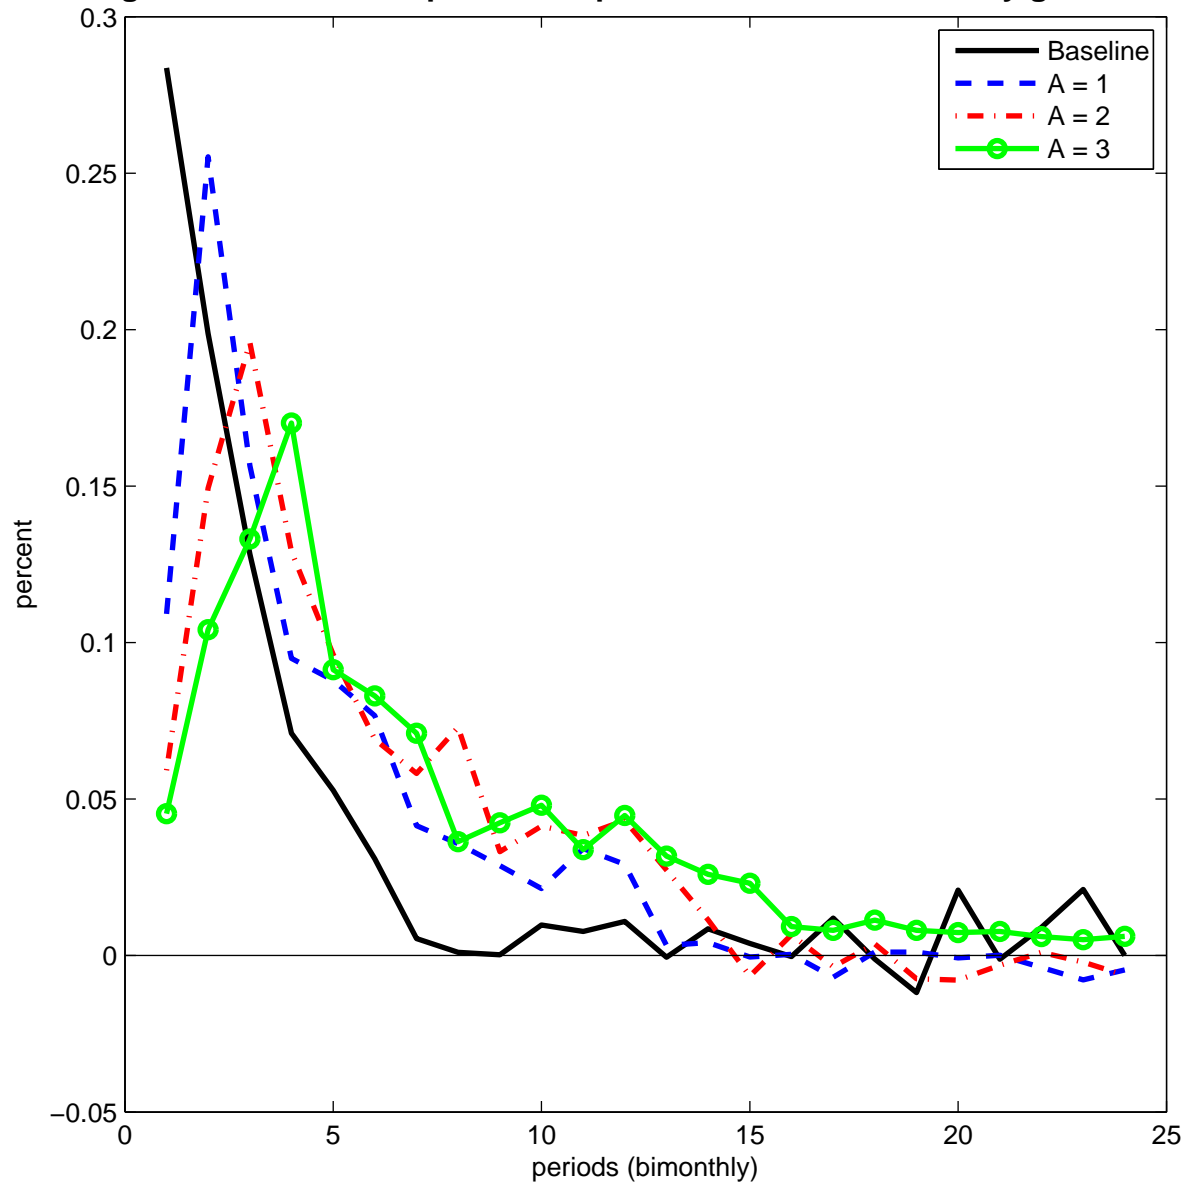
Notes: This table displays the coefficients from a regression of price changes on the predicted price change due to new information on the aggregate price level since the previous change ($\gamma PPC_{i,t}$) and additional lags of inflation innovations ($\lambda_j \Delta_{\tau_{i,t}} \epsilon_{t-3-j}$). For details on this specification, see equation (42). Within each model specification, information on aggregate state variables is *identical* for all firms. \bar{A} indicates the age of aggregate information. A panel of 6000 firms and 500 periods is simulated for each model. BLS CPI data are measured at a bimonthly frequency from 1988 to 2004.

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

| | | γ_1 | γ_{down} | γ_{up} | R^2 |
|----------|----------------------|------------|-----------------|---------------|-------|
| BLS Data | Full Sample | 1.307 | -0.324 | 0.312 | 0.319 |
| | | (0.013) | (0.0004) | (0.0004) | |
| | Regular | 0.996 | NA | NA | 0.003 |
| | | (0.014) | | | |
| | Sales-Related | 1.861 | -0.285 | 0.352 | 0.489 |
| | | (0.030) | (0.0007) | (0.0007) | |
| | Same Product | 1.205 | -0.324 | 0.301 | 0.339 |
| | | (0.014) | (0.0004) | (0.0004) | |
| | Substitution-Related | 0.657 | -0.300 | 0.385 | 0.189 |
| | | (0.056) | (0.003) | (0.002) | |

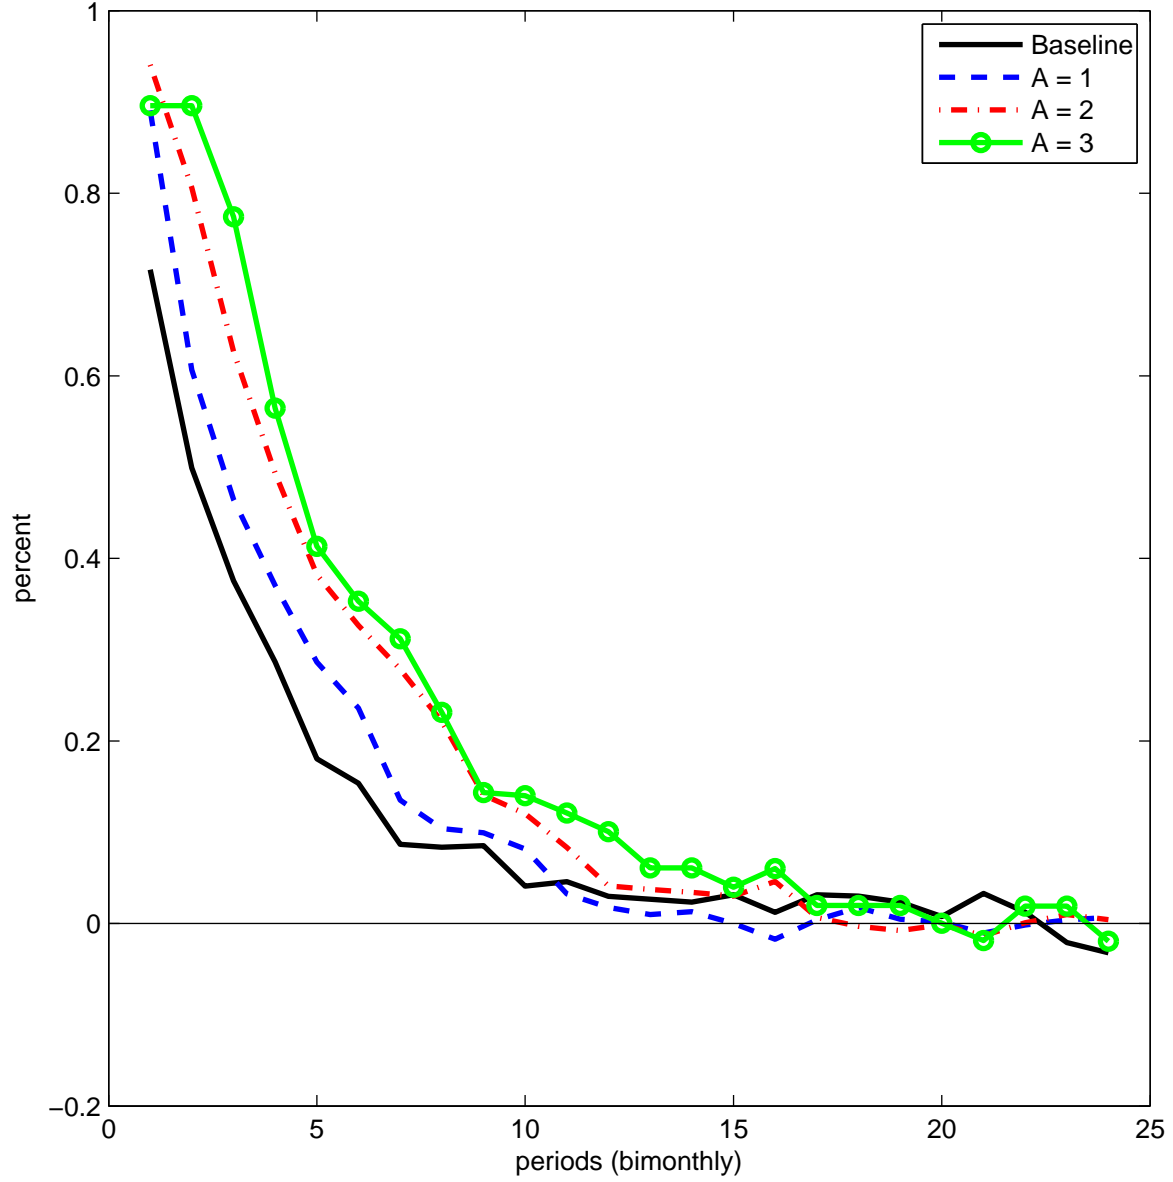
Notes: This table displays the coefficients from a regression of price changes on the predicted price change due to new information on the aggregate price level since the previous change ($\gamma_1 PPC_{i,t}$) along with dummy variables for a regular-to-sale price change ($\gamma_{down} D_{i,t}$) and a sale-to-regular price change ($\gamma_{up} U_{i,t}$). For details on this specification, see equation (43). BLS CPI data are measured at a bimonthly frequency from 1988 to 2004. The various specifications consider all price changes (Full Sample), price changes unrelated to sales (Regular), sales-related price changes (Sales-Related), price changes that do not involve a product substitution (Same Product), and price changes where a product substitution has occurred (Substitution-Related).

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



Notes: In each model specification, information on aggregate state variables arrives on a *staggered*, deterministic schedule. A indicates the maximum age of information before updating occurs. In the baseline model, $A = 0$.

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



Notes: In each model specification, information on aggregate state variables arrives on a *staggered*, deterministic schedule. A indicates the maximum age of information before updating occurs. In the baseline model, $A = 0$.